

Wolfgang E. Nagel
Dietmar H. Kröner
Michael M. Resch *Editors*

High Performance Computing in Science and Engineering '12

H L R I S

 Springer

High Performance Computing in Science and Engineering '12

Wolfgang E. Nagel • Dietmar H. Kröner
Michael M. Resch
Editors

High Performance Computing in Science and Engineering '12

Transactions of the High Performance
Computing Center, Stuttgart (HLRS) 2012

 Springer

Editors

Wolfgang E. Nagel
Zentrum für Informationsdienste
und Hochleistungsrechnen (ZIH)
Technische Universität Dresden
Dresden
Germany

Dietmar H. Kröner
Abteilung für Angewandte Mathematik
Universität Freiburg
Freiburg
Germany

Michael M. Resch
Höchstleistungsrechenzentrum
Stuttgart (HLRS)
Universität Stuttgart
Stuttgart
Germany

Front cover figure: Horizontal slice through a snapshot of simulated turbulence in the atmosphere. Red and green regions are strongly turbulent while blue regions are relatively calm. (Details can be found in 'Numerical Investigation of Stratified Turbulence' by S. Remmler and S. Hickel)

ISBN 978-3-642-33373-6 ISBN 978-3-642-33374-3 (eBook)
DOI 10.1007/978-3-642-33374-3
Springer Heidelberg New York Dordrecht London

Library of Congress Control Number: 2012954941

Mathematics Subject Classification (2010): 65Cxx, 65C99, 68U20

© Springer-Verlag Berlin Heidelberg 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Contents

Part I Physics

Peter Nielaba

Constraints on the Two-Flavour QCD Phase Diagram from Imaginary Chemical Potential	5
O. Philipsen and Ph. de Forcrand	
Parallel Computer Algebra and Feynman Integrals	13
P.A. Baikov, K.G. Chetyrkin, J.H. Kühn, P. Marquard, M. Steinhauser, and T. Ueda	
Electronic and Optical Excitations of Aminopyrimidine Molecules from Many-Body Perturbation Theory	17
A. Rieger, M. Rohrmüller, M. Landmann, S. Sanna, E. Rauls, U. Gerstmann, and W.G. Schmidt	
Spinodal Decomposition Kinetics of Colloid-Polymer Mixtures Including Hydrodynamic Interactions	29
Alexander Winkler, Peter Virnau, and Kurt Binder	
Multi Relaxation Time Lattice Boltzmann Simulations of Multiple Component Fluid Flows in Porous Media	39
Sebastian Schmieschek, Ariel Narváez, and Jens Harting	
Simulation of Pre-planetesimal Collisions with Smoothed Particle Hydrodynamics II	51
R.J. Geretshausen, F. Meru, K. Schaal, R. Speith, and W. Kley	
The Stellar IMF at Very Low Metallicities	69
Gustavo Dopcke, Simon C.O. Glover, Paul C. Clark, and Ralf S. Klessen	

The SuperN-Project: Porting and Optimizing VERTEX-PROMETHEUS on the Cray XE6 at HLRS for Three-Dimensional Simulations of Core-Collapse Supernova Explosions of Massive Stars	81
F. Hanke, A. Marek, B. Müller, and H.-Th. Janka	
Stability of the Strongest Magnetic Fields	95
Konstantinos D. Kokkotas, Burkhard Zink, and Paul Lasky	
Part II Solid State Physics	
Holger Fehske	
Molecular Dynamics Simulations of Laser Ablation in Metals: Parameter Dependence, Extended Models and Double Pulses	105
Johannes Roth, Johannes Karlin, Marc Sartison, Armin Krauß, and Hans-Rainer-Trebin	
Electronic Surface Properties of Transparent Conducting Oxides: An Ab Initio Study	119
B. Höffling and F. Bechstedt	
Surface Magnetism: Relativistic Effects at Semiconductor Interfaces and Solar Cells	129
U. Gerstmann, M. Rohrmüller, N.J. Vollmers, A. Konopka, S. Greulich-Weber, E. Rauls, M. Landmann, S. Sanna, A. Riefer, and W.G. Schmidt	
Ab-Initio Calculations of the Vibrational Properties of Nanostructures	145
Gabriel Bester and Peng Han	
Part III Reacting Flows	
Dietmar Kröner	
Conservative Implementation of LES-CMC for Turbulent Jet Flames	159
P. Siwaborworn and A. Kronenburg	
Numerical Investigation of a Complete Scramjet Demonstrator Model for Experimental Testing Under Flight Conditions	175
Yann Simson, Peter Gerlinger, and Manfred Aigner	
Application of the Unified Turbulent Flame-Speed Closure (UTFC) Combustion Model to Numerical Computation of Turbulent Gas Flames	187
Feichi Zhang, Peter Habisreuther, and Henning Bockhorn	
Lagrangian Approach for the Prediction of Slagging and Fouling in Pulverized Coal Combustion	207
Olaf Lemp, Uwe Schnell, and Günter Scheffknecht	

Part IV Computational Fluid Dynamics

Siegfried Wagner

Discontinuous Galerkin for High Performance Computational Fluid Dynamics 225
 Christoph Altmann, Andrea Beck, Andreas Birkefeld,
 Gregor Gassner, Florian Hindenlang, Claus-Dieter Munz,
 and Marc Staudenmaier

Computational Aeroacoustics with Higher Order Methods..... 239
 E. Rebecca Busch, Michael S. Wurst, Manuel Keßler,
 and Ewald Krämer

Effects of an Oblique Roughness on Hypersonic Boundary-Layer Transition..... 255
 Gordon Groskopf and Markus J. Kloker

Effect of Wall Roughness Seen by Particles in Turbulent Channel and Pipe Flows 277
 Michael Breuer and Michael Alletto

Large-Eddy Simulation of Supersonic Film Cooling at Incident Shock-Wave Interaction 295
 Martin Konopka, Matthias Meinke, and Wolfgang Schröder

Large-Eddy Simulations of Stratified and Non-stratified T-junction Mixing Flows 311
 David Klören and Eckart Laurien

Simulation of Compressible Viscous Flow with an Immersed Boundary Method..... 325
 B. Jastrow and F. Magagnato

Numerical and Experimental Examination of Shock Control Bump Flow Physics 333
 K. Nübler, S.P. Colliss, T. Lutz, H. Babinsky, and E. Krämer

Water Droplet Flow Paths and Droplet Deposition in Low Pressure Steam Turbines 351
 J. Starzmann, M.V. Casey, and J.F. Mayer

MPC and Coarray Fortran: Alternatives to Classic MPI Implementations on the Examples of Scalable Lattice Boltzmann Flow Solvers 367
 Markus Wittmann, Georg Hager, Gerhard Wellein, Thomas Zeiser,
 and Bettina Krammer

Part V Transport and Climate

Christoph Kottmeier

Modelling Near Future Regional Climate Change for Germany and Africa	375
H.-J. Panitz, G. Fossier, R. Sasse, A. Sehlinger, H. Feldmann, and G. Schädler	
Setting Up Regional Climate Simulations for Southeast Asia	391
Patrick Laux, Van Tan Phan, Christof Lorenz, Tran Thuc, Lars Ribbe, and Harald Kunstmann	
The Agulhas System as a Key Region of the Global Oceanic Circulation	407
J.V. Durgadoo and A. Biastoch	
Numerical Investigation of Stratified Turbulence	415
S. Remmler and S. Hickel	

Part VI Miscellaneous Topics

Wolfgang Schröder

Software Framework ug4: Parallel Multigrid on the Hermit Supercomputer	435
Ingo Heppner, Michael Lampe, Arne Nägel, Sebastian Reiter, Martin Rupp, Andreas Vogel, and Gabriel Wittum	
Simulation of Liquid-Liquid Equilibria with Molecular Models Optimized to Vapor-Liquid Equilibria and Model Development for Hydrazine and Two of Its Derivatives	451
Stefan Eckelsbach, Thorsten Windmann, Ekaterina Elts, and Jadran Vrabec	
A Particle-in-Cell Method to Model the Influence of Partial Melt on Mantle Convection	461
Ana-Catalina Plesa, Doris Breuer, and Tilman Spohn	
A Forward Model of Mantle Convection with Evolving Continents and a Model of the Andean Subduction Orogen	473
Uwe Walzer, Roland Hendel, Christoph Köstler, Markus Müller, Jonas Kley, and Lothar Viereck-Götte	
Euler Deconvolution of GOCE Gravity Gradiometry Data	503
M. Roth, N. Sneeuw, and W. Keller	
Parameterization of Threshold Accepting: The Case of Economic Capital Allocation	517
H.-P. Burghof and J. Müller	

Distributed FE Analysis of Multiphase Composites Regarding 3D Elasticity Problems 531
 Kai Schrader and Carsten Könke

Reduction of Numerical Sensitivities in Crash Simulations on HPC-Computers (HPC-10) 547
 Christiana Eck, Oliver Mangold, Raphael Prohl, and Anton Tkachuk

Characterization of Carrier Sense Multiple Access in Vehicular Propagation Channels 561
 J. Mittag and H. Hartenstein

HMDB51: A Large Video Database for Human Motion Recognition 571
 Hilde Kuehne, Hueihan Jhuang, Rainer Stiefelhagen, and Thomas Serre

Efficient Parallelization of a Three-Dimensional High-Order Particle-in-Cell Method Applied to Gyrotron Resonator Simulations 583
 J. Neudorfer, A. Stock, T. Stindl, R. Schneider, S. Roller, C.-D. Munz, and M. Auweter-Kurtz

Part I Physics

Professor Dr. Peter Nielaba

Many important results have been achieved by the computer time granted at the HLRS and the projects are partly embedded in research clusters, e.g. “clusters of excellence” or “Sonderforschungsbereiche”. The contributions in the proceedings present the results of large scale simulations for elementary particle models, for systems on nano- and micrometer length scales, and for astrophysics phenomena, and are summarized and commented below.

O. Philipsen from the University of Frankfurt and Ph. de Forcrand from the ETH Zürich and CERN (“*muQCD*”) have investigated the critical surface bounding the region featuring chiral phase transitions in the quark mass and chemical potential parameter space of quantum chromodynamics (QCD) with three flavors of quarks. Their calculations are valid for small to moderate quark chemical potentials, $\mu \leq T$. The authors study the situation for the limiting case of two light flavours and show first results on the nature of the chiral transition at zero chemical potential from extrapolations using imaginary chemical potential.

For their Monte Carlo simulations the authors used the standard Wilson gauge and Kogut-Susskind fermion actions. Configurations are generated using the Rational Hybrid Monte Carlo (RHMC) algorithm. In order to investigate the critical behavior of the theory, the authors use the Binder cumulant as an observable. For each set of fixed quark mass and chemical potential, the critical coupling β_c has been interpolated from a range of typically three to five simulated β -values by Ferrenberg-Swendsen reweighting. The simulations have been performed on the NEC-SX9 at the HLRS in Stuttgart. An estimate of the Binder cumulant for one set of mass values consisted of at least 200k trajectories, and the estimate of a critical point required at least 500k trajectories.

P. Baikov, K. Chetyrkin, J.H. Kühn, P. Marquard, M. Steinhauser and T. Ueda from the KIT Karlsruhe (“*ParFORM*”) have investigated multi-loop Feynman integrals by a computer algebra program (“*PARFORM*”), using MPI on the XC4000. The computation of Feynman integrals is important for the computation of quantum corrections to physical processes, the code is under constant development, and the authors give an overview on the present status.

A. Riefer, M. Rohrmüller, M. Landmann, S. Sanna, E. Rauls, U. Gerstmann, and W.G. Schmidt from the University of Paderborn (“*MolArchI*”) have investigated the electronic structure and optical response of 2-aminopyrimidine molecules by a combination of density functional theory and many-body perturbation theory. The calculations predict quasiparticle gaps, i.e., differences between the ionization energies and electron affinities, of about several eV for the molecules, and the result indicate a near cancellation of the electronic self-energy and exciton binding energies for the lowest excitations of 2-aminopyrimidines. The authors find a strong influence of local-field effects as well as resonant-nonresonant coupling terms in the electron-hole Hamiltonian on the optical properties.

The authors have used the Vienna Ab-initio Simulation Package (“*VASP*”) implementation of the gradient-corrected density functional theory (DFT-GGA) for their computations of the ground state and GWA calculations. The HSE06 hybrid functional has been used as well. For the electronic self energy calculations applying perturbation theory (“ G_0W_0 ”) and Bethe-Salpeter type calculations the cell size has been varied. The calculations within this project were performed on the CRAY XE6 at the HLRS with good scaling properties.

K. Binder, P. Virnau and A. Winkler from the University of Mainz (“*colloid*”) have investigated the spinodal decomposition of colloid-polymer mixtures between walls including hydrodynamic interactions by the multi particle collision dynamics and domain decomposition methods on Hermit. Polymers are described as soft spheres weakly interacting with each other, while colloid-polymer and colloid-colloid pairs interact with the (repulsive) Weeks- Chandler-Andersen potential, so that a depletion attraction between colloids results, similar to the Asakura-Oosawa model. Large system sizes and long simulation times are required for this study, and interesting results of the effect of confining geometry on the decomposition dynamics have been achieved and will be studied in the future. Important results are the effects of hydrodynamics and of different boundary conditions on the growth exponents. The authors used the standard halo layer domain decomposition technique to treat the embedded particles and the parallelization approach proposed by Sutmann et al. for the solvent particles. This level of parallelization allowed the authors to use 1,024 cores (or more) to study successfully the phase separation kinetics in the Asakura-Oosawa model in huge systems (max. 1,100,000 MD particles and 52,000,000 solvent particles) over multiple scales of MD time. The authors find that for huge systems at the late time stages of phase separation the net cost of explicit solvent particles in the framework of MPCD is only about approx. 10 % in comparison to standard Molecular Dynamics simulations.

S. Schmieschek, A. Narvez, and J. Harting from the University of Stuttgart and the Eindhoven University of Technology (“*icpsusp*”) investigate multiple

component flows in porous media by lattice Boltzmann simulations on the XC2 cluster at the SCK. The authors integrated a “multi relaxation time” collision scheme with their pseudo-potential multiphase lattice Boltzmann implementation “lb3d” for fluids with multiple components, and good scaling is achieved. The steps taken to optimize the performance have allowed to limit the increase in calculation time close to the expected minimum. The authors show that for multi-component systems the additional calculation time spent in the collision step is very small in comparison to the time increase due to the calculation of the interaction forces. The authors give an overview on the present status.

R.J. Geretshauer, F. Meru, K. Schaal, R. Speith, and W. Kley (University of Tübingen and ETH Zürich, “SPH-PPC”) have analyzed pre-planetesimal collisions with their solid body smoothed particle hydrodynamics (SPH) code *parasph*. The main focus of the project is on the investigation of the growth conditions for macroscopic pre-planetesimals. By parameter studies the authors investigated fragmentation criteria in dust collisions depending on aggregate size and aggregate porosity, and they extended their previous study on bouncing criteria of equally sized aggregates depending on their porosity and the presence of compacted shells of various porosities. The authors derive fragmentation criteria for dust cylinders depending on angular velocity as well as porosity and perform corresponding simulations.

The code *parasph*, used by the authors, is based on the “parasph” library, featuring domain decomposition, load balancing, nearest neighbor search, and inter-node communication, extended for the simulation of ductile, brittle, and porous media and by an implementation of a porosity model and by SPH enhancements. The parallel implementation utilizes the Message Passing Interface (MPI) library, and HDF5 was included as a compressed input and output file format with increased accuracy, decreasing the amount of required storage space. The simulations were carried out on the NEC Nehalem cluster of the HLRS with 240,143–476,476 SPH particles depending on the size of the projectile. Thirty-two to eighty cores were used, and simulations roughly took 72–240 h for 1 s of simulated time, depending on the size of the problem and the involved physical process.

S.C.O. Glover, P.C. Clark and R.S. Klessen from the University of Heidelberg (“EDuCool”) investigate the evolution of star-forming clouds for a wide range of metallicities and the effects on the mass function of the fragments that form, on cooling and heating rates, and on the number of Bonnor-Ebert masses of the fragmenting clouds. The authors use a modified version of the Gadget-2 (SPH) code, a typical simulation run is done with 40 million SPH particles and requires 130 kCPU-h on 256 or 512 CPUs. Interesting results have been obtained on the thermodynamical evolution of gas and dust, the fragmentation and the properties of the fragments.

F. Hanke, A. Marek, B. Müller, and H.-Th. Janka from the MPI for Astrophysics in Garching (“SuperN”) investigate core-collapse supernova explosions of massive stars. The authors have developed a fully MPI-OpenMP parallelized version of their VERTEX-PROMETHEUS code in order to perform three-dimensional simulations of stellar core-collapse and explosion. The simulations typically require 10^{20}

floating point operations, and good scaling up to 32,000 cores has been achieved on Hermit.

K.D. Kokkotas, P. Lasky and B. Zink from the University of Tübingen (“*Magnetar*”) have investigated the dynamic stability of strong magnetic fields inside highly magnetized neutron stars by general relativistic magnetohydrodynamics simulations by the Horizon code on GPUs. Parts of the simulations have been done on the Nehalem cluster at the HLRS. The authors have varied the stiffness of the equation of states and the rotation of the neutron star and were able to simulate the time evolution for hundreds of milliseconds or even seconds. Their results are important contributions to an Sonderforschungsbereich on “Gravitational Wave Astronomy”.

Constraints on the Two-Flavour QCD Phase Diagram from Imaginary Chemical Potential

O. Philipsen and Ph. de Forcrand

In a long term project, we calculate the critical surface bounding the region featuring chiral phase transitions in the quark mass and chemical potential parameter space of QCD with three flavours of quarks. Our calculations are valid for small to moderate quark chemical potentials, $\mu \lesssim T$. The presence of tricritical lines at imaginary chemical potential $\mu = i \frac{\pi}{3} T$, with known scaling behaviour in their vicinity, puts constraints on this phase diagram. Here we undertake first steps to study the situation for the limiting case of two light flavours. In this case, the nature of the chiral transition at zero chemical potential is not yet established. We show first results of our project to extract this behaviour from extrapolations using imaginary chemical potential.

1 Introduction

The fundamental theory describing the strong interactions is Quantum Chromodynamics (QCD) with two light quark flavours, the u - and d -quarks, and a heavier s -quark. Since the interaction weakens at asymptotically large energy scales, QCD predicts at least three different forms of nuclear matter: the usual hadronic matter at low temperature and density, a quark gluon plasma at high temperature and low density, and colour superconducting nuclear matter at low temperatures and high density. Direct Monte Carlo simulations of the finite density QCD phase diagram are impossible because of the so-called sign problem, so that indirect methods need

O. Philipsen (✉)

Institut für Theoretische Physik, Goethe-Universität Frankfurt, 60438 Frankfurt am Main, Germany

e-mail: philipsen@th.physik.uni-frankfurt.de

Ph. de Forcrand

Institut für Theoretische Physik, ETH Zürich, CH-8093 Zürich, Switzerland Physics Department, TH-Unit, CERN, CH-1211 Geneva, Switzerland

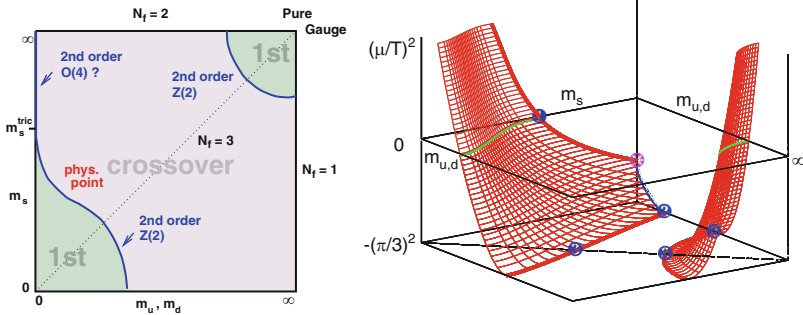


Fig. 1 *Left*: Schematic phase transition behaviour of $N_f = 2 + 1$ QCD for different choices of quark masses ($m_{u,d}, m_s$) at $\mu = 0$. *Right*: The same with chemical potential for quark number as an additional parameter. The critical lines sweep out surfaces as μ is turned on. At imaginary chemical potential $\mu = i\pi/3T$, the critical surfaces terminate in tricritical lines which determine their curvature through critical scaling

to be employed, which work for small enough μ/T only (for an overview and references, see [1]).

At zero chemical potential, the nature of the quark-hadron phase transition depends on the quark masses, as summarised in Fig. 1. In the limits of zero and infinite quark masses, order parameters for the breaking of the global chiral and centre symmetry, respectively, can be defined, and one finds numerically that first order phase transitions take place at some finite temperature T_c . On the other hand, for intermediate quark masses the transition is an analytic crossover. Hence, each corner of first order phase transitions is bounded by a second-order critical line as in Fig. 1. The physical quark masses are light, so our interest is in the lower left boundary, which is called the chiral critical line, as opposed to the deconfinement critical line in the heavy mass region.

In previous work the location of the boundary line has been determined for the case of degenerate quark masses, $N_f = 3$ [2, 3], where it was also shown that it belongs to the 3d Ising, or 3d $Z(2)$, universality class. On the lattice, temperature and lattice spacing are related by $T = 1/(aN_t)$, i.e. larger N_t corresponds to finer lattices for a fixed physical temperature. We have used $N_t = 4$ lattices, corresponding to a lattice spacing $a \sim 0.3$ fm, to map out how this line changes (i) for $N_f = 3$ as a function of chemical potential μ [3, 4] and (ii) for $\mu = 0$ in the case of non-degenerate quark masses $m_{u,d} \neq m_s$ [4]. It was found that the physical point is located close to the boundary line on the crossover side.

When a chemical potential for the baryon density is switched on, the chiral critical point of the $N_f = 3$ theory recedes as in Fig. 1 (right). The critical quark mass marking the boundary line can be expanded as

$$\frac{m_c(\mu)}{m_c(0)} = 1 + c_1 \left(\frac{\mu}{\pi T}\right)^2 + c_2 \left(\frac{\mu}{\pi T}\right)^4 + \dots \quad (1)$$

with $c_1 = -3.3(3)$, $c_2 = -47(20)$ [5, 6]. The same behaviour is found for non-degenerate quark masses. Tuning the s -quark mass to its physical value, we calculated $m_c^{u,d}(\mu)$ with $c_1 = -39(8)$ and $c_2 < 0$ [7]. Similar behaviour is found for heavy quarks. Hence, the critical lines sweep out surfaces, as shown in Fig. 1 (right). Contrary to common expectations based on simplified models, the curvature of both surfaces is such that the region of first order shrinks, i.e. the phase transitions weaken when a real chemical potential is switched on. As a consequence, the physical point remains in the crossover region and there is no chiral critical point in QCD for moderate baryon densities.

In the current and future project, we explore in particular the implications for the behaviour of QCD in the two-flavor chiral limit ($m_u = m_d = 0$, $m_s = \infty$). In that limit, it is widely believed that QCD undergoes a finite-temperature, second-order $O(4)$ chiral transition at $\mu = 0$, which turns first-order at a tricritical point for some real μ . However, other possibilities exist. At $\mu = 0$ in particular, the finite-temperature transition might be first-order. The present numerical evidence is inconclusive: using Wilson fermions, $O(4)$ scaling is preferred [8], while with staggered fermions $O(4)$ scaling has been elusive, and first-order behaviour has also been claimed [9]. Note that behaviour consistent with $O(4)$ has been seen with improved staggered fermions, in an $N_f = 2 + 1$ setup where the strange quark mass is fixed at its physical value [10]. Approaching the chiral limit from the imaginary μ direction, i.e. working in the back plane of Fig. 1 (right), offers a novel, independent method to help settle the issue.

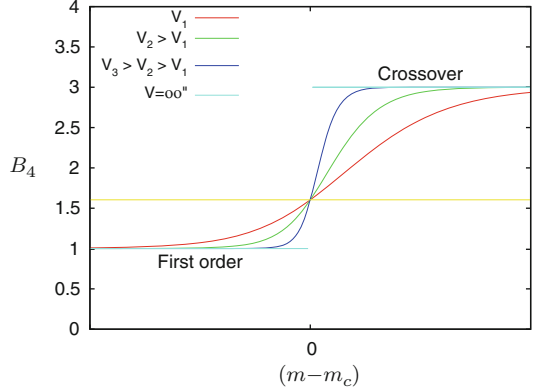
2 The Binder Cumulant and Universality

In order to investigate the critical behaviour of the theory, we use the Binder cumulant [11] as an observable. It is defined as

$$B_4(m, \mu) = \frac{\langle(\delta X)^4\rangle}{\langle(\delta X)^2\rangle^2}, \quad (2)$$

with the fluctuation $\delta X = X - \langle X \rangle$ of the order parameter of interest. Since we investigate the region of chiral phase transitions, we use the chiral condensate, $X = \bar{\psi}\psi$. For the evaluation of the Binder cumulant it is implied that the lattice gauge coupling has been tuned to its pseudo-critical value, $\beta = \beta_c(m, \mu)$, corresponding to the phase boundary between the two phases. In the infinite volume limit the Binder cumulant behaves discontinuously, assuming the values 1 in a first order regime, 3 in a crossover regime and the critical value ≈ 1.604 reflecting the 3d Ising universality class at a chiral critical point. On a finite volume the discontinuities are smeared out and B_4 passes continuously through the critical value. This is sketched in Fig. 2. In the neighbourhood of the chiral critical point at zero chemical potential it can be expanded linearly

Fig. 2 Schematic behaviour of the Binder cumulant as a function of quark mass for $\mu = 0$. First order transitions and crossovers correspond to $B_4 = 1, 3$, respectively, whereas a second order 3d Ising transition is characterised by $B_4 \approx 1.604$. On finite volumes the step function gets smeared out



$$B_4(m, \mu) = A + B(am - am_c) + C(a\mu)^2 \dots, \quad (3)$$

with $A \rightarrow 1.604$ for $V \rightarrow \infty$.

3 Continuation of the Critical Surfaces to Imaginary μ

Let us consider imaginary chemical potential, $\mu = i\mu_i$. The QCD partition function exhibits two important exact symmetries, reflection symmetry in μ and $Z(3)$ -periodicity μ_i , which hold for quarks of any mass [12],

$$Z(\mu) = Z(-\mu), \quad Z\left(\frac{\mu}{T}\right) = Z\left(\frac{\mu}{T} + i\frac{2\pi n}{3}\right), \quad (4)$$

for general complex values of μ . The symmetries imply transitions between adjacent centre sectors of the theory at fixed $\mu_i^c = (2n + 1)\pi T/3, n = 0, \pm 1, \pm 2, \dots$. The $Z(3)$ -sectors are distinguished by the Polyakov loop

$$L(\mathbf{x}) = \frac{1}{3} \text{Tr} \prod_{\tau=1}^{N_\tau} U_0(\mathbf{x}, \tau) = |L| e^{-i\varphi}, \quad (5)$$

whose phase φ cycles through $\langle \varphi \rangle = n(2\pi/3), n = 0, 1, 2, \dots$ as the different sectors are traversed. Moreover, the above also implies reflection symmetry about the $Z(3)$ phase boundaries, $Z(\mu_i^c + \mu_i) = Z(\mu_i^c - \mu_i)$.

Transitions in μ_i between neighbouring sectors are of first order for high T and analytic crossovers for low T [12–14], as shown in Fig. 3 (left). The order parameter is the shifted phase of the Polyakov loop, $\phi = \varphi - \mu_i/T$. Away from $\mu_i = \mu_i^c$, there is a chiral or deconfinement transition line separating high and




Fig. 3 *Left*: phase diagram for imaginary μ . *Vertical lines* are first order transitions between $Z(3)$ -sectors, *arrows* show the phase of the Polyakov loop. The $\mu = 0$ chiral/deconfinement transition continues to imaginary μ , its order depends on N_f and the quark masses. *Right*: Schematic phase diagram of the Roberge-Weiss endpoint

low temperature regions. This line represents the analytic continuation of the chiral or deconfinement transition at real μ . Its nature (first, second order or crossover) depends on the number of quark flavours and masses. We have investigated the nature of the junction of those lines for $N_f = 3$ [15], similar investigations have been carried out for $N_f = 2$ [16]. The combined outcome is shown in Fig. 3 (right), which corresponds to the bottom plane of Fig. 1 (right). The Roberge-Weiss point is of first order for heavy and light quarks, and of second order for intermediate mass quarks. These regimes are separated by tricritical lines. Four points, corresponding to the $N_f = 3$ [15] and $N_f = 2$ boundary points, have been explicitly computed. The most natural generalisation to non-degenerate quark masses is that they are continuously connected by tricritical lines.

Figure 1 (right) represents the connection between the $\mu = 0, i\pi T/3$ phase diagrams. In the vicinity of a tricritical point, scaling laws apply. In our case, the scaling exponents governing the behaviour near the tricritical point are mean-field. This is shown for the example of heavy quarks, where

$$[(\mu/T)^2 + (\pi/3)^2] \propto (m_{u,d} - m_{\text{tric}})^{5/2}. \quad (6)$$

Figure 4 shows the corresponding behaviour of the critical heavy quark mass as a function of chemical potential.

Note that the $N_f = 2$ (i.e. $m_s = \infty$) “backplane” contains two tricritical points on the chiral critical surface: one in the Roberge-Weiss plane, the other on the $m_{u,d} = 0$ vertical axis. The location of the latter is related to the value of the tricritical strange quark mass. They should be again connected by critical lines related to tricritical scaling.




Fig. 4 For heavy quarks, tricritical scaling in the vicinity of the Roberge-Weiss imaginary- μ value extends far into the region of real μ [15]





Fig. 5 Binder cumulants for the lightest fermion mass, $am = 0.0025$, determining the critical point

4 Preliminary $N_f = 2$ Results

Following the above discussion, we joined forces with an Italian collaboration and performed simulations of $N_f = 2$ QCD. We used staggered quarks of masses $am_q = 0.01$ and 0.005 on $N_t = 4$ lattices scanning in $(\mu/T)^2$ to determine the value of imaginary μ corresponding to a second-order transition. Our observable is the Binder cumulant of the quark condensate. For an example at the smallest quark mass, see Fig. 5. Consistent results are obtained from the finite-size scaling of the

Fig. 6 Critical line joining the known tricritical point on the *horizontal axis*, and extrapolated to the unknown tricritical point on the *vertical axis*



plaquette distribution. So far we have obtained four critical points, Fig. 6. It seems impossible to smoothly match two tricritical scaling curves passing through these points. Additional masses are needed to determine the critical curve. Nevertheless, assuming convexity of the critical curve already constrains the $m_{u,d} = 0$ tricritical point to lie at $(\mu/T)^2 \gtrsim -0.3$. The figure illustrates the case where this point lies at $\mu = 0$. It might also lie at $(\mu/T)^2 > 0$, so that the $\mu = 0$ chiral transition would be first-order. Preliminary results have been reported in [17]. Additional small-mass measurements are underway and will settle this issue.

5 Simulation Details

For our Monte Carlo simulations we use the standard Wilson gauge and Kogut-Susskind fermion actions. Configurations are generated using the Rational Hybrid Monte Carlo algorithm [18]. Our numerical procedure to compute the Binder cumulant is as follows. For each set of fixed quark mass and chemical potential, we interpolate the critical coupling β_c from a range of typically three to five simulated β -values by Ferrenberg-Swendsen reweighting [19]. For each simulation point $\sim 50k$ RHMC trajectories have been accumulated, measuring the gauge action, the Polyakov loop and up to four powers of the chiral condensate after each trajectory. Thus, the estimate of B_4 for one set of mass values consists of at least 200k trajectories, and the estimate of a critical point at least 500k trajectories.

The simulations are performed on the NEC-SX9 at the HLRS in Stuttgart. A scan in parameter space involves simulations of many parameter sets. For such a problem, parallelisation is achieved trivially by running one set of couplings per node, each node running in vector mode. This way of parallelising allows to explore large regions of the parameter space at the same time, which is necessary when mapping out a phase diagram. At the same time, there is no overhead for parallelisation and communication, ensuring maximal computing efficiency and one-to-one scaling of

compute power with the number of processors. The vector mode ensures maximal throughput for each individual lattice. Typically we work with several nodes at once using each of their eight cores in parallel and avoiding communication between the nodes.

References

1. O. Philipsen, Eur. Phys. J. ST **152** (2007) 29 [arXiv:0708.1293 [hep-lat]].
2. F. Karsch, E. Laermann and C. Schmidt, Phys. Lett. B **520** (2001) 41 [arXiv:hep-lat/0107020].
3. P. de Forcrand and O. Philipsen, Nucl. Phys. B **673** (2003) 170 [arXiv:hep-lat/0307020].
4. P. de Forcrand and O. Philipsen, JHEP **0701** (2007) 077 [arXiv:hep-lat/0607017].
5. P. de Forcrand, S. Kim and O. Philipsen, PoS **LAT2007** (2007) 178 [arXiv:0711.0262 [hep-lat]].
6. P. de Forcrand and O. Philipsen, JHEP **0811** (2008) 012 [arXiv:0808.1096 [hep-lat]].
7. J. T. Moscicki, M. Wos, M. Lamanna, P. de Forcrand and O. Philipsen, Comput. Phys. Commun. **181** (2010) 1715 [arXiv:0911.5682 [Unknown]].
8. Y. Iwasaki, K. Kanaya, S. Kaya and T. Yoshie, Phys. Rev. Lett. **78** (1997) 179 [hep-lat/9609022].
9. M. D'Elia, A. Di Giacomo and C. Pica, Phys. Rev. D **72**, 114510 (2005) [hep-lat/0503030]; G. Cossu, M. D'Elia, A. Di Giacomo and C. Pica, arXiv:0706.4470 [hep-lat]; C. Bonati, G. Cossu, M. D'Elia, A. Di Giacomo and C. Pica, PoS **LATTICE 2008** (2008) 204 [arXiv:0901.3231 [hep-lat]].
10. S. Ejiri, F. Karsch, E. Laermann, C. Miao, S. Mukherjee, P. Petreczky, C. Schmidt and W. Soeldner *et al.*, Phys. Rev. D **80** (2009) 094505 [arXiv:0909.5122 [hep-lat]].
11. K. Binder, Z. Phys. B **43** (1981) 119.
12. A. Roberge and N. Weiss, Nucl. Phys. B **275**, 734 (1986).
13. P. de Forcrand and O. Philipsen, Nucl. Phys. B **642**, 290 (2002) [arXiv:hep-lat/0205016].
14. M. D'Elia and M. P. Lombardo, Phys. Rev. D **67**, 014505 (2003) [arXiv:hep-lat/0209146].
15. P. de Forcrand and O. Philipsen, Phys. Rev. Lett. **105** (2010) 152001 [arXiv:1004.3144 [hep-lat]].
16. M. D'Elia and F. Sanfilippo, Phys. Rev. D **80** (2009) 111501 [arXiv:0909.0254 [hep-lat]].
17. C. Bonati, P. de Forcrand, M. D'Elia, O. Philipsen and F. Sanfilippo, arXiv:1201.2769 [hep-lat].
18. M. A. Clark and A. D. Kennedy, Phys. Rev. Lett. **98** (2007) 051601 [arXiv:hep-lat/0608015].
19. A. M. Ferrenberg and R. H. Swendsen, Phys. Rev. Lett. **63** (1989) 1195.

Parallel Computer Algebra and Feynman Integrals

P.A. Baikov, K.G. Chetyrkin, J.H. Kühn, P. Marquard, M. Steinhauser,
and T. Ueda

1 Introduction

Relativistic quantum field theories are the mathematical tools which are indispensable in modern particle physics. In particular in combination with perturbation theory, realized with the help of so-called Feynman diagrams, it is possible to make predictions which can be confronted with experiment.

The main motivation of the current project is to study multi-loop Feynman integrals, the mathematical expressions of the Feynman diagrams, which have to be computed when considering quantum corrections to physical processes. From the mathematical point of view a Feynman integral corresponds to a multi-dimensional integral over the time and space components of momenta which are defined in a complex space-time dimension d . The integrand consists of rational functions of scalar products formed by all d -dimensional momenta involved in the problem.

There are several technical challenges one has to master in the course of such calculations. Among the most demanding ones is the processing of huge amount of data which has to be processed in intermediate steps. For this reason some years ago parallel versions of the computer algebra program FORM [1] have been developed which are refined and optimized to date in order to solve even more complex tasks. The main advantage of FORM is the ability to process huge expressions (basically only limited by disk space) in a fast and effective way.

There are two parallel versions of FORM: PARFORM [2] and TFORM [3]. Whereas the parallelization of TFORM is realized with threads which restricts

This is the report for the project PARFORM for the period June 2011–2012.

P.A. Baikov · K.G. Chetyrkin · J.H. Kühn · P. Marquard · M. Steinhauser (✉) · T. Ueda
Institut für Theoretische Teilchenphysik, Karlsruher Institut für Technologie,
76128 Karlsruhe, Germany
e-mail: matthias.steinhauser@kit.edu

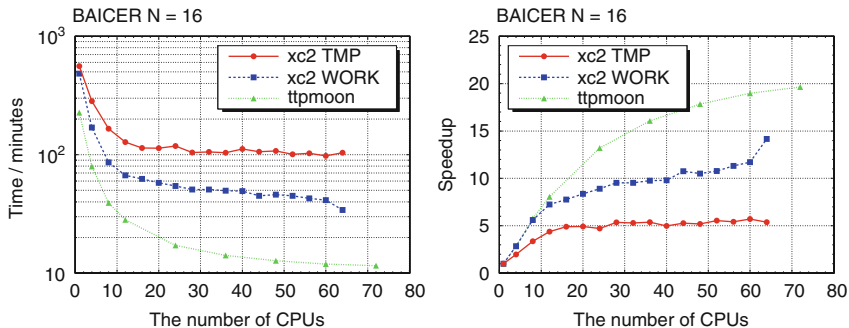


Fig. 1 Run time and speed-up for the benchmark program BAICER on XC4000. For comparison, we also plot the result on ttpmoon, a computer cluster in our institute

TFORM to multi-core computers, PARFORM uses MPI (“Message Passing Interface”) for the communication among the various processors. Thus PARFORM can also be used for computer clusters with a fast interconnection, e.g. infiniband. Due to the structure of the XC4000 cluster which mostly consists of four-core nodes essentially only PARFORM can be used for calculations.

Let us mention that since August 2010 all versions of FORM are open source and can be downloaded from [4].

In the past we have used the cluster XC4000 (and its predecessor XC6000) for various calculations running up to ten jobs in parallel. On April 2011 we were, however, asked to reduce the number of jobs to at most two since the performance of the whole system is significantly reduced if PARFORM starts writing or reading the intermediate results to or from the hard disk. This makes the use of XC4000 quite unattractive. For this reason we have not submitted any production job in the past 12 months but have used XC4000 essentially to further develop PARFORM and to benchmark against other platforms.

2 PARFORM Running on XC4000

In this section we update the performance tests presented in the report from June 2010. Figure 1 compares the run time and speed-up for a typical benchmark job as a function of the number of used CPUs. The red (circles) and blue (squares) lines correspond both to calculations performed on XC4000. Whereas for the red data the local disks have been used the blue data have been obtained by using global disks. The former show a very bad performance (of at most 60 MB/s) and thus reach a speed-up of 5 only for about 15 CPUs. On the other hand the use of the global disks leads to good results (with a theoretical peak performance of 320 and 400 MB/s for reading and writing, respectively) up to about 8–12 CPUs which is what we are

usually using for the production jobs. Of course, the use of the global hard disks explains the performance reduction of the system.

The green curve (triangles) has been obtained by running the benchmark jobs on a cluster with 12-core nodes interconnected via QDR infiniband (with about 1 GB/s read/write performance). As can be seen from Fig. 1, up to about ten cores the green and blue curves are identical. For more CPUs, however, the green curve shows a significant better behaviour reaching a speed-up of 10 for about 16 cores and 15 for 30 cores.

3 FORM 4 and PARFORM

The serial version of FORM is under constant development in order to be able to be applicable to a broad class of problems which recently resulted in version 4 [5]. Of course, modifications of FORM require also the adaption of the parallel versions. In the last year the main emphasis has been in implementing the new features of version 4 into PARFORM.

In FORM version 4 two essentially new features, which has been requested by many users over the recent years, are introduced: the factorization of expressions and the ability to work with rational polynomial functions as coefficients of terms. They are based on newly written routines for (multivariate) polynomial operations (additions, subtractions, greatest common divisors, factorizations, etc.), for which FORM uses a number of well-known algorithms.

All new features added in version 4 are also available in PARFORM. Although FORM and PARFORM share most of the source code, some additional code was needed to synchronize data between the master process and all worker processes. An example is the factorization of $\$$ -variables, which are small expressions stored in memory. When factorizing a $\$$ -variable, each factor is stored in an additionally allocated buffer. Therefore the routine for synchronization of $\$$ -variables between the master and workers was extended such that all factors are also synchronized via MPI if the $\$$ -variable has been factorized.

References

1. FORM version 3.0 is described in: J. A. M. Vermaseren, “New features of FORM”, arXiv:math-ph/0010025; for more developments, see also: M. Tentyukov and J. A. M. Vermaseren, “Extension of the functionality of the symbolic program FORM by external software”, arXiv:cs.sc/0604052; FORM can be obtained from the distribution site at <http://www.nikhef.nl/~form>.
2. M. Tentyukov, D. Fliegner, M. Frank, A. Onischenko, A. Retey, H. M. Staudenmaier and J. A. M. Vermaseren, “ParFORM: Parallel Version of the Symbolic Manipulation Program FORM”, arXiv:cs.sc/0407066; M. Tentyukov, H. M. Staudenmaier and J. A. M. Vermaseren,

- “ParFORM: Recent development”, Nucl. Instrum. Meth. A **559** (2006) 224. H. M. Staudenmaier, M. Steinhauser, M. Tentyukov, J. A. M. Vermaseren, “ParFORM”, Computeralgebra Rundbriefe **39** (2006) 19. See also <http://www-ttp.physik.uni-karlsruhe.de/~parform>.
3. M. Tentyukov and J. A. M. Vermaseren, “The multithreaded version of FORM”, arXiv:hep-ph/0702279.
 4. <http://www.nikhef.nl/~form/formcvs.php>
 5. J. Kuipers, T. Ueda, J. A. M. Vermaseren and J. Vollinga, arXiv:1203.6543 [cs.SC].

Electronic and Optical Excitations of Aminopyrimidine Molecules from Many-Body Perturbation Theory

A. Riefer, M. Rohrmüller, M. Landmann, S. Sanna, E. Rauls, U. Gerstmann, and W.G. Schmidt

Abstract Calculations based on (occupation constrained) density functional theory using local as well as hybrid functionals to describe the electron-electron exchange and correlation are combined with many-body perturbation theory in order to determine the electronic and optical excitation properties of 5-(pentafluorophenyl)pyrimidin-2-amine, 5-(4-methoxy-2,3,5,6-tetrafluorophenyl)pyrimidin-2-amine, and 5-(4-(dimethylamino)-2,3,5,6-tetrafluorophenyl)pyrimidin-2-amine. Large quasiparticle shifts and exciton binding energies of about 4 eV are found. They cancel each other partially and thus allow for a meaningful description of the molecular optical response within the independent-particle approximation. We find a surprisingly strong influence of local-field effects as well as resonant-nonresonant coupling terms in the electron-hole Hamiltonian on the optical properties.

1 Introduction

Organic semiconductors are important materials for various applications due to their low cost fabrication processes and the possibility to fine-tune desired functions by chemical modification of their building blocks. While the last years have seen a tremendous progress in the understanding of the excitation properties of inorganic semiconductors, fueled in part by the availability of advanced computational schemes for electronic structure and optical response calculations such as the GW approximation (GWA) for obtaining accurate electronic quasiparticle energies and

A. Riefer · M. Rohrmüller · M. Landmann · S. Sanna · E. Rauls · U. Gerstmann · W.G. Schmidt (✉)
Lehrstuhl für Theoretische Physik, Universität Paderborn, 33095 Paderborn, Germany
e-mail: W.G.Schmidt@upb.de




Fig. 1 Schematic model of 5-(pentafluorophenyl)pyrimidin-2-amine (FAP), 5-(4-methoxy-2,3,5,6-tetrafluorophenyl)pyrimidin-2-amine (OFAP), and 5-(4-(dimethylamino)-2,3,5,6-tetrafluorophenyl)pyrimidin-2-amine (NFAP) (from left to right). Dark (red), light (yellow), gray, lightgray and small balls indicate O, C, N, F and H atoms, respectively

the Bethe-Salpeter approach (BSE) to calculate electron-hole interaction effects [1–6], far less is known about the electronic and optical properties of organic crystals.

Recently, a novel class of organic electronic material has been synthesized by the self-assembly and silver(I) complex formation of 2-aminopyrimidines [7]. The compounds were structurally as well as optically characterized [8] and it was found that the solid state absorption differs remarkably from the parent compound 2-aminopyrimidine. The optical properties could be tuned by changing the silver counterion or by the reversible solvent extrusion and interchange. Furthermore, the electrical conductivity of the material was proven for a thin crystalline film.

In order to gain a better understanding of the excitation properties of this class of systems, we first study molecular excitations in the respective parent molecules. In detail, we present *first-principles* calculations on the electronic and optical properties of 5-(pentafluorophenyl)pyrimidin-2-amine (FAP), 5-(4-methoxy-2,3,5,6-tetrafluorophenyl)pyrimidin-2-amine (OFAP), and 5-(4-(dimethylamino)-2,3,5,6-tetrafluorophenyl)pyrimidin-2-amine (NFAP) in order to clarify the impact of many-body effects and chemical trends. The three aminopyrimidine molecules (APM) are shown in Fig. 1. They consist of 22 (FAP), 26 (OFAP) and 30 (NFAP) atoms forming a 2-aminopyrimidine ring (atoms 1–8 in Fig. 1) and a (per)fluorinated phenylring (12–17) where the position no. 22 is either a fluorine atom F_1 (FAP), a methoxy group (OFAP) or an amino group (NFAP).

2 Methodology

Ground-state and GWA calculations are performed using the Vienna Ab-initio Simulation Package (VASP) implementation [9] of the gradient-corrected [10] density functional theory (DFT-GGA). In addition, the hybrid functional due to

Heyd, Scuseria, and Ernzerhof (HSE06) [11] was used. The electron-ion interaction is described by the projector augmented-wave (PAW) method [12, 13]. We expand the valence wave functions into plane waves up to an energy cutoff of 400 eV. DFT calculations for single molecules were performed using a $14 \times 15 \times 20 \text{ \AA}^3$ super cell and Γ point sampling for the Brillouin zone (BZ) integration. Test calculations show that the eigenenergies are converged within a few hundredths an eV. For electronic self-energy calculations applying perturbation theory (G_0W_0) and Bethe-Salpeter type calculations (see, e.g., Ref. [5]) as well as for calculations of charged molecules the cell size was varied as described below.

DFT calculations are known to often considerably underestimate electronic excitation energies [4]. Reliable quasiparticle gaps, exciton pair energies and Stokes shifts, however, can be obtained from occupation constraint DFT (or Δ SCF) methods, cf. Refs. [14–16]. Thereby the quasiparticle (QP) gap is obtained directly as difference between the ionization energy and electron affinity

$$E_g^{QP} = E(N + 1, \mathbf{R}) + E(N - 1, \mathbf{R}) - 2E(N, \mathbf{R}), \quad (1)$$

where $E(N, \mathbf{R})$, $E(N + 1, \mathbf{R})$, and $E(N - 1, \mathbf{R})$ represent the energy of a N , $N + 1$, and $N - 1$ electron system, respectively, with the equilibrium geometry \mathbf{R} of the N electron system. The energy of the lowest excitonic excitation corresponding to the situation that one electron occupies the lowest unoccupied molecular orbital (LUMO) leaving a hole behind in the highest occupied molecular orbital (HOMO) is given by

$$E_{ex} = E(e - h, \mathbf{R}) - E(N, \mathbf{R}), \quad (2)$$

where $E(e - h, \mathbf{R})$ is the total energy of the system in presence of the electron-hole-pair with fixed geometry \mathbf{R} . Alternatively, as can be derived from Janak's Theorem (see Ref. [14]), the energy of the lowest optical excitation can be obtained from the difference of the eigenenergies of the half-occupied HOMO $\varepsilon_{H,0.5}$ and LUMO $\varepsilon_{L,0.5}$, respectively

$$E_{ex} \simeq E_{ex}^J = \varepsilon_{L,0.5} - \varepsilon_{H,0.5}. \quad (3)$$

Relaxing the atomic coordinates to the geometry \mathbf{R}^* for fixed occupation numbers yields the lowest emission energy

$$E_{em} = E(e - h, \mathbf{R}^*) - E(N, \mathbf{R}^*), \quad (4)$$

which can be used to calculate the Stokes shift

$$\Delta_S = E_{ex} - E_{em}. \quad (5)$$

From calculations of the ground-state energy for different cell sizes one can conclude an error for the Δ SCF values of 0.1 eV. The QP gaps E_g^{QP} are compared to the gap $E_g^{G_0W_0}$ that has been obtained from the G_0W_0 approximation of the electronic self energy and is obtained by postprocessing the PW91 wave functions and eigenvalues. The implementation details are given in Ref. [17].

For systems where the electronic states have either the occupancy 0 for conduction states, $n = c$, or 1 for valence states, $n = v$ one obtains the dielectric tensor in independent-particle approximation [18–20] (IPA)

$$\varepsilon_{ij}(\omega) = \delta_{ij} + \frac{4\pi e^2}{\Omega} \lim_{q \rightarrow 0} \frac{1}{q^2} \sum_{\mathbf{k}c\nu} 2 \times \frac{1}{\varepsilon_c(\mathbf{k}) - \varepsilon_\nu(\mathbf{k}) - (\hbar\omega + i\eta)} \langle u_{c\mathbf{k}+\mathbf{q}_i} | u_{\nu\mathbf{k}} \rangle \langle u_{c\mathbf{k}+\mathbf{q}_j} | u_{\nu\mathbf{k}} \rangle^*, \quad (6)$$

where the sum $\sum_{\mathbf{k}}$ is to be taken over the first BZ, \mathbf{q}_i is the reciprocal vector in the cartesian direction i , $u_{n\mathbf{k}}$ are the periodic parts of the Bloch wave functions, $\varepsilon_n(\mathbf{k})$ the respective eigenenergies, Ω is the crystal volume, and η is the broadening. In order to allow for comparison with the experimental data we average over the three cartesian directions

$$\varepsilon(\hbar\omega) = \frac{1}{3} \sum_{i=x,y,z} \varepsilon_{ii}(\hbar\omega). \quad (7)$$

The dielectric function within the IPA or by solving the BSE is based on the electronic structure as obtained from either the PW91/HSE06 calculations (partially with scissors shifted eigenvalues) or from the GWA.

Solving the Bethe-Salpeter equation includes the electron-hole attraction and local-field effects in the dielectric function. For practical calculations, the BSE is transformed into a two-particle Schrödinger equation. Neglecting dynamical screening and umklapp processes, the resonant part of the exciton Hamiltonian (Tamm-Dancoff-Approximation, TDA, cf. Ref. [21]) for direct transitions and spin-singlets can be calculated in reciprocal space according to

$$\begin{aligned} \hat{H}_{vc\mathbf{k},v'c'\mathbf{k}'}^{res} &= (\varepsilon_c^{QP}(\mathbf{k}) - \varepsilon_{v'}^{QP}(\mathbf{k})) \delta_{vv'} \delta_{cc'} \delta_{\mathbf{k}\mathbf{k}'} \\ &+ \frac{4\pi}{\Omega} \sum_{\mathbf{G},\mathbf{G}'} \left\{ 2 \frac{\delta_{\mathbf{G}\mathbf{G}'}(1 - \delta_{\mathbf{G}\mathbf{0}})}{|\mathbf{G}|^2} B_{cv}^{\mathbf{k}\mathbf{k}}(\mathbf{G}) B_{c'v'}^{\mathbf{k}'\mathbf{k}'}(\mathbf{G}) \right. \\ &\left. - \frac{\varepsilon^{-1}(\mathbf{k} - \mathbf{k}' + \mathbf{G}, \mathbf{k} - \mathbf{k}' + \mathbf{G}', 0)}{|\mathbf{k} - \mathbf{k}' + \mathbf{G}|^2} B_{cc'}^{\mathbf{k}\mathbf{k}'}(\mathbf{G}) B_{vv'}^{\mathbf{k}\mathbf{k}'}(\mathbf{G}') \right\}, \quad (8) \end{aligned}$$

where the Bloch integral

$$B_{nn'}^{\mathbf{k}\mathbf{k}'}(\mathbf{G}) = \frac{1}{\Omega} \int d\mathbf{r} u_{n\mathbf{k}}^*(\mathbf{r}) e^{i\mathbf{G}\mathbf{r}} u_{n'\mathbf{k}'}(\mathbf{r}) \quad (9)$$

over the periodic parts u of the Bloch wave functions has been introduced. In the actual calculations we replace the inverse dielectric matrix ε^{-1} by a diagonal model dielectric function suggested by Bechstedt et al. [22]. It depends on the static dielectric constant ε_∞ and reduces the computational effort substantially. In case of inorganic semiconductors [23, 24], molecular crystals [16, 25] and even surfaces

[26], the application of this model dielectric function leads to rather accurate results. This is related to the fact that the model dielectric function depends on the local charge density and therefore carries some information about the local screening. For molecule calculations, the correct choice of ε_∞ is difficult. The authors of Ref. [27] defined an effective volume Ω_{eff} where the screening takes place in order to address this problem in their optical response calculation of poly-para-phenylenevinylene. In our work we use for molecular calculations $\varepsilon_\infty = 1$, which marks the lower limit for the screening interaction. If one assumes $\Omega_{eff} = 18^3 \text{ \AA}^3$, the IPA calculations for FAP result in $\varepsilon_\infty = 1.05$, which leads to a blueshift of the excitonic eigenvalues by about 0.3 eV. Calculations for further values of ε_∞ indicate a nearly linear dependence of the exciton binding energies on the screening, as may be expected. The dimension of the exciton Hamiltonian (Eq. 8) is determined by the size of the energy window for conduction and valence states. The spectra are calculated including either all states satisfying $\varepsilon_c(\mathbf{k}) - \varepsilon_v(\mathbf{k}) < 6 \text{ eV}$ (DFT) or the lowest 96 states (GWA). For the actual calculation of the spectra we use the time-evolution algorithm proposed by one of the present authors [26]. In addition to BSE-TDA, also calculations with the full exciton Hamiltonian were performed (BSE). For the comparison with measured optical spectra we use real and imaginary parts of the dielectric function, $\varepsilon'(\hbar\omega)$ and $\varepsilon''(\hbar\omega)$, respectively to obtain the attenuation coefficient α using the approximation

$$\alpha(\hbar\omega) \propto \hbar\omega \sqrt{\left[\sqrt{\varepsilon'(\hbar\omega)^2 + \varepsilon''(\hbar\omega)^2} - \varepsilon'(\hbar\omega) \right]}. \quad (10)$$

The calculated data are compared with optical absorption measurements on powder samples.

The HLRS CRAY XE6 is the main computational resource used for the calculations in this work. As can be seen in Fig. 2, the scaling is nearly linear up to 200 cores, allowing for highly efficient calculations.

3 Results

The structural relaxation of FAP, OFAP and NFAP in gas phase shows that the geometry of the aminopyrimidine and pentafluorophenyl rings does barely change upon attachment of either a fluorine atom (FAP), a methoxy group (OFAP) or an amino group (NFAP). The comparison of our calculated data with x-ray data of two polymorphic crystals of the hydrogen analogue 5-phenyl-pyrimidin-2-ylamine (HAP) and a HAP-hexafluorobenzene co-crystal[7] as well as the recently crystallized NFAP ligand itself shows only small differences in bond length and angles. Only for the hydrogen bonds we observe deviations of up to 0.10–0.16 Å between measured and calculated data. The geometries calculated here also closely agree with Møller Plesset perturbation theory (MP2) results for APM [28]: The bond lengths deviate by less than 0.02 Å and the largest deviation of bond angles amounts to 3°.

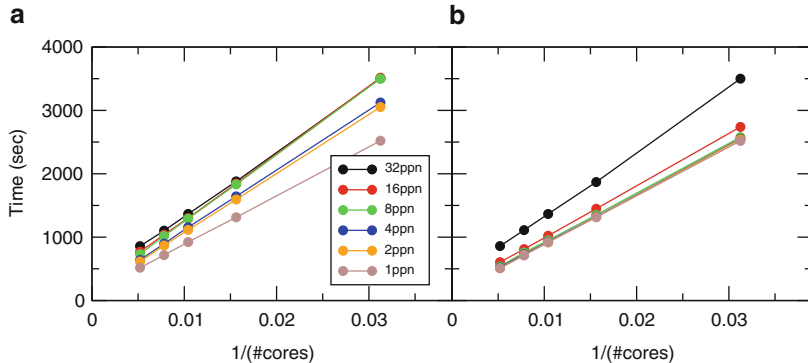


Fig. 2 Wall clock time for the DFT-GGA calculation of 5-(pentafluorophenyl)pyrimidin-2-amine (FAP) in a $20 \times 20 \times 20 \text{ \AA}^3$ cell including 1,152 electronic states on the HLRS CRAY XE6. In (a) the behavior of the wall clock time with respect to the number of cores and tasks per node (*ppn*) is shown. As can be seen, the scaling is nearly linear up to 200 cores. The time required is reduced with increased distribution of the tasks on several nodes. Additionally, (b) shows that the wall clock time can be further reduced if the cores employed for the calculations are equally spaced (spacing $32/\text{ppn}$) on the nodes

Table 1 Molecular excitation energies (in eV, see text)

	FAP	OFAP	NFAP
E_g^{PW91}	3.46	3.35	3.00
E_g^{HSE06}	4.53	4.55	4.21
$E_g^{G_0W_0}$	$\lesssim 7.7$	$\lesssim 7.4$	$\lesssim 7.1$
E_g^{QP}	7.36	7.06	6.47
E_{ex}	3.51	3.46	3.21
E_{ex}^J	3.50	3.46	3.22
E_{em}	2.08	1.97	1.98
Δ_S	1.43	1.49	1.23

Starting from the relaxed structures we calculated the quantities defined in Eqs. (1)–(5). The results for FAP, OFAP, and NFAP are compiled in Table 1. We find that the difference of the HOMO and LUMO eigenenergies, $E_g^{PW91} = \varepsilon_L - \varepsilon_H$, is largest for FAP and decreases by going from OFAP to NFAP (see also Fig. 4 for the electronic levels), i.e., with increasing electron-donating properties. In HSE06 the ordering between FAP and OFAP is reverse compared to the GGA calculation. However, the gaps are very close. The trend observed with GGA holds also for the G_0W_0 gaps $E_g^{G_0W_0}$ and the Δ SCF gaps E_g^{QP} . The calculation of a QP Δ SCF gap requires the determination of the total energies $E(N + 1, \mathbf{R})$ and $E(N - 1, \mathbf{R})$ of charged molecules. Due to the interactions with the periodic images the dependence of the latter and thus the gap E_g^{QP} on the cell size is not negligible. In order to correct the calculated excitation energies, the gaps were determined for a cubic cell with varying size $L = 18, \dots, 30 \text{ \AA}$.

As shown in Fig. 3, the gap values depend linearly on $1/L$. Extrapolation to $L \rightarrow \infty$ leads to the gaps cited in Table 1.




Fig. 3 Dependence of the calculated quasiparticle-gaps E_g^{QP} and G_0W_0 gaps $E_g^{G_0W_0}$ on the cell size L . The filled/stripped symbols for G_0W_0 values denote calculations with a cutoff for the response function of 60/40 eV (see text). The inset shows the respective values for the energy difference between the LUMO+1 and HOMO

A dependence on the unit cell size is also noted for the calculated G_0W_0 gaps, see Fig. 3. To some extent, this is to be expected due to the periodic repetition of the molecules. The restriction of the calculations with respect to further parameters due to numerical limitations, however, is even more important in the present case. The self-energy calculations for cubic cells with the size $L = 18\text{--}20 \text{\AA}$ ($22\text{--}24 \text{\AA}$) were performed with a maximum cutoff for the response function of 60 eV (40 eV), 90 frequency points, and a cutoff of 15–16 eV for the sum over empty states (including up to 1,056 bands).

The dependence of the G_0W_0 on the numerics is obvious from the inset in Fig. 3, where the energy difference between the FAP HOMO and LUMO+1 states is shown, but also from Fig. 4, where the energetic ordering of the electronic states is visualized. Obviously, the order changes upon inclusion of electronic self-energies calculated with the G_0W_0 approximation, but is itself not yet converged, at least for the unoccupied states. Nevertheless, as will be shown below, the reordering due to state-dependent self-energy corrections calculated in G_0W_0 improves the agreement between the measured and calculated optical absorption. The present data suggest that the band gaps calculated within the GWA decrease with increasing cell size for the molecules studied here. The numbers given in Table 1 should thus be considered as approximate upper limits. We find that the values are by about 0.5 eV larger than the respective energy gaps determined from the ΔSCF calculations. The fundamental gaps calculated with the HSE06 scheme, on the other hand, are between the PW91 and the quasiparticle gaps.

Interestingly, the quasiparticle shifts are nearly cancelled by electron-hole attraction effects: The lowest electron-hole excitation energies E_{ex} are remarkably close to the difference of the HOMO and LUMO single particle eigenenergies obtained from DFT. This near cancellation of many body effects due to the electron-electron and the electron-hole interaction suggests that optical excitation spectra calculated




Fig. 4 Energies of molecular orbitals as obtained from DFT (PW91) and G_0W_0 calculations for cubic cells with $L = 22 \text{ \AA}$ (left) and $L = 24 \text{ \AA}$ (right). The influence of the self-energy corrections and cell size on the energy order of the states is indicated by different colors. *Thick bars* refer to the orbitals that correspond to HOMO, LUMO, and LUMO+1 in the PW91 calculations. The fundamental gap is indicated. Note the different energy-region for the empty states

in the independent-particle approximation may be a reasonable description at least for the low-energy excitations.

The calculation of the electron-hole excitation energies is computationally robust: The approaches according to Eqs. (2) and (3) result in energies that agree within 0.01 eV. The lowest-energy excitations calculated for structural relaxation differ appreciable from the respective vertical excitation energies. We calculate Stokes shifts between 1.2 and 1.5 eV for the three molecules. Thereby, the energetic ordering changes between absorption and emission. While NFAP is predicted to have the lowest vertical excitation energy, its deexcitation occurs at slightly larger energies than OFAP.

Our calculated values are in reasonable agreement with the experimental data available: For FAP dissolved in ethanol Stoll et al.[8] measured a Stokes shift of 1.28 eV. Given that the optical response of the molecules will be influenced by the solvent, these data confirm the validity of the present calculations.

From the eigenfunctions and eigenvectors obtained in DFT one can directly calculate the dielectric function in independent-particle approximation. Figure 6 shows the resulting spectra for FAP, OFAP, and NFAP. Obviously, in all three cases




Fig. 5 Orbital character of the states HOMO (a)–(c) and LUMO (d)–(f) for FAP, OFAP and NFAP

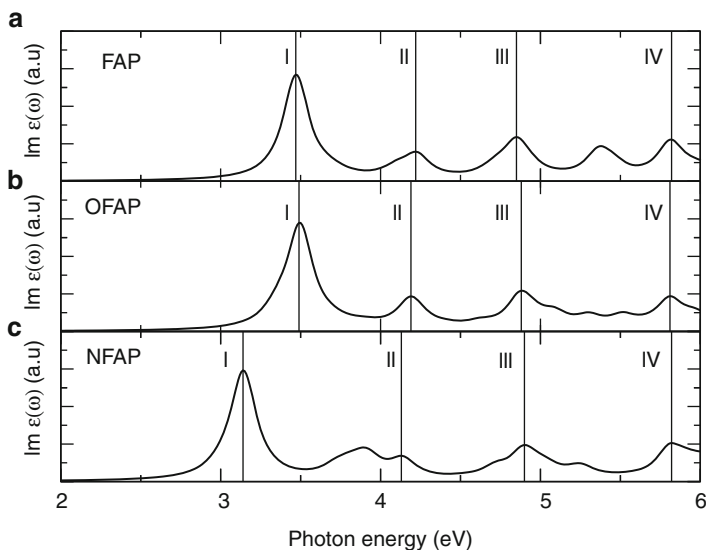


Fig. 6 Imaginary part of the dielectric function calculated in independent-particle approximation for FAP, OFAP and NFAP. A broadening of $\eta = 0.10$ eV has been used

the onset of the optical absorption is larger than $E_g^{DFT} = \varepsilon_L - \varepsilon_H$ due to the small transition probability between HOMO and LUMO. There are more similarities in the spectra. In particular FAP and OFAP agree largely concerning the positions and line shapes of the main peaks I-IV (see Fig. 6).

Since the dielectric function in independent-particle approximation is composed of independent transitions between occupied and empty electronic states, it is straightforward to interpret. In particular we find that transitions between HOMO and LUMO+1 are essentially causing the first absorption peak for all three molecules. The data show furthermore that the optical absorption occurs largely due to states localized at the aminopyrimidine and pentafluorophenyl rings. This explains why the optical response of the three molecules shown in Fig. 6 is rather similar. A notable exception is the first absorption peak of NFAP. In this case the HOMO is strongly influenced by amino-group localized states (cf. Fig. 5).

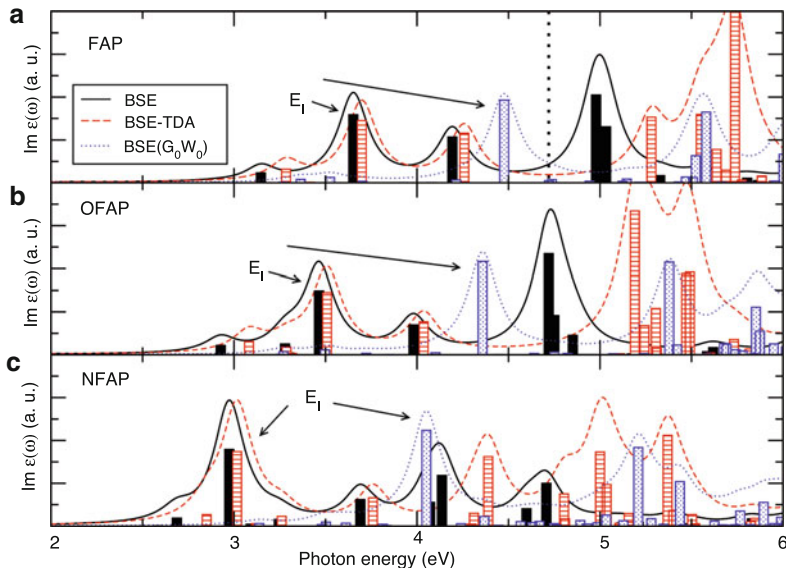


Fig. 7 Imaginary part of the dielectric function calculated by solving the BSE based on G_0W_0 calculations or by applying a respective scissors-shift to reproduce the Δ SCF-gaps or for FAP (a), OFAP (b) and NFAP (c). A broadening of $\eta = 0.10$ eV has been used. The *solid (dashed)/dotted curve* and *bars* gives the spectra and oscillator strengths versus the eigenvalues calculated within BSE (BSE-TDA) on the basis of the scissors shifted PW91/ G_0W_0 electronic structure. The eigenvalues contributing to the first peak are labeled. See Ref. [23] for details. The strongest absorption maximum of FAP dissolved in ethanol[8] is shown by a *dotted line*

Contributions of the attached fluorine atom or the methoxy group are – to a much smaller extent – also present in the first absorption peak of FAP or OFAP (cf. Fig. 5).

In Fig. 7 the molecular dielectric functions calculated by taking many body effects into account are shown. The calculations have been performed using the full excitonic Hamiltonian as well as applying the TDA. The empty electronic levels were either shifted such that the respective molecular Δ SCF gaps are reproduced or the G_0W_0 electronic structure was used as input. The red-shift of the first peak in the NFAP spectra compared to FAP and OFAP as observed in IPA occurs also on the BSE level of theory. It is even enhanced by the smaller respective value of the Δ SCF gap. In general the positions of the first optical absorption maxima calculated within the BSE agree within 1 eV with the IPA calculation, which is indicative for some cancelation of quasi-particle and excitonic shifts as already concluded from the values in Table 1.

Similar to the IPA spectra discussed above we perform a systematic analysis of the states contributing to the respective absorption peaks (for details see Ref. [29]). It turns out that – as already found on the IPA level of theory – HOMO and LUMO+1 are the states that mostly contribute to the first adsorption peak.

Comparing the spectra obtained from the full Hamiltonian and in TDA one finds distinct differences: (i) a redshift of the eigenvalues going from TDA to the full

Hamiltonian and (ii) strong modifications of the line shape for energies above 4.5 eV (FAP, OFAP) or 4.0 eV (NFAP). While the Tamm-Dancoff-Approximation clearly affects the calculated optical absorption, in particular for excitations beyond the lowest absorption peak, we find the influence of the electronic structure that is used as input for the BSE calculations to be even more important. The optical spectrum based on the G_0W_0 electronic structure differs appreciable from the one based on scissors-shifted PW91 eigenvalues. This is due to the state-dependent self-energy corrections leading to an energetic reordering of the eigenvalue spectrum that results in a significant blue-shift of the optical absorption data.

The measured position of the optical absorption peak of FAP dissolved in ethanol[8] in the energy window 2.3–5.7 eV is at 4.72 eV (vertical line in Fig. 7). Clearly, the BSE spectrum based on the G_0W_0 electronic structure agrees best with this value. It yields an optical absorption peak at 4.48 eV. From Table 1 it is clear that the error bar of the calculated excitation energies is of the order of several tenths of an eV. Moreover, our choice for the static dielectric constant used in the molecule calculations is bound to result in excitation energies that approach the real values from below. An additional uncertainty in the experiment-theory comparison is related to the fact that the solvent molecules are not included in the present gas-phase calculations. Therefore the deviation between measured and calculated data of less than 0.3 eV is not surprising.

Comparing the computational results for the electronic states of FAP, OFAP, and NFAP summarized in Fig. 4 and the optical response from Figs. 6 and 7 one finds that the former are far more sensitive to the attachment of functional groups than the latter. Since the optical absorption essentially takes place at the aminopyrimidine and pentafluorophenyl rings, modifications in the molecular wave functions due to methoxy or amino group are only partially reflected in the optical data.

4 Summary

In the present work the electronic structure and optical response of 2-aminopyrimidines is analyzed on the basis of DFT as well as many-body perturbation theory calculations. The calculations predict quasiparticle gaps, i.e., differences between the ionization energies and electron affinities, of about 7 eV for the molecules. The energies of the lowest optical excitations of the respective molecules are considerably lower. In fact, our result indicates a near cancellation of the electronic self-energy and exciton binding energies for the lowest excitations of 2-aminopyrimidines. In addition to electron-hole attraction effects, we observe a very strong influence of local fields, i.e., the unscreened electron-hole exchange on the optical absorption spectra. Moreover, the resonant-nonresonant coupling terms in the excitonic Hamiltonian usually neglected in calculations for inorganic semiconductors are found to noticeably modify peak positions and oscillators strengths in case of the systems studied here.

Acknowledgements We gratefully acknowledge financial support from the DFG as well as supercomputer time provided by the HLRS Stuttgart and the Paderborn PC².

References

1. M. Rohlfling, S.G. Louie, Phys. Rev. Lett. **81**, 2312 (1998)
2. S. Albrecht, L. Reining, R. DelSole, G. Onida, Phys. Rev. Lett. **80**, 4510 (1998)
3. M. Rohlfling, Int. J. Quant. Chem. **80**, 807 (2000)
4. G. Onida, L. Reining, A. Rubio, Rev. Mod. Phys. **74**, 601 (2002)
5. P.H. Hahn, W.G. Schmidt, F. Bechstedt, Phys. Rev. B **72**, 245425 (2005)
6. W.G. Schmidt, phys. stat. sol. (b) **242**, 2751 (2005)
7. I. Stoll, R. Brodbeck, B. Neumann, H.G. Stammler, J. Mattay, Crys. Eng. Comm. **11**(2), 306 (2009)
8. I. Stoll, R. Brockhinke, A. Brockhinke, M. Boettcher, T. Koop, H.G. Stammler, B. Neumann, A. Niemeyer, A. Huetten, J. Mattay, Chem. Materials **22**(16), 4749 (2010)
9. G. Kresse, J. Furthmüller, Comp. Mat. Sci. **6**, 15 (1996)
10. J.P. Perdew, J.A. Chevary, S.H. Vosko, K.A. Jackson, M.R. Pederson, D.J. Singh, C. Fiolhais, Phys. Rev. B **46**, 6671 (1992)
11. J. Heyd, G.E. Scuseria, M. Ernzerhof, J. Chem. Phys. **118**(18), 8207 (2003)
12. P.E. Blöchl, Phys. Rev. B **50**, 17953 (1994)
13. G. Kresse, D. Joubert, Phys. Rev. B **59**, 1758 (1999)
14. R.M. Dreizler, E.K.U. Gross, *Density Functional Theory* (Springer-Verlag, Berlin, 1990)
15. M. Preuss, W.G. Schmidt, K. Seino, J. Furthmüller, F. Bechstedt, J. Comp. Chem. **25**, 112 (2004)
16. A. Hermann, W.G. Schmidt, P. Schwerdtfeger, Phys. Rev. Lett. **100**, 207403 (2008)
17. M. Shishkin, G. Kresse, Phys. Rev. B **74**, 035101 (2006)
18. H. Ehrenreich, M.H. Cohen, Phys. Rev. **115**, 786 (1959)
19. S.L. Adler, Phys. Rev. **126**, 413 (1962)
20. N. Wisner, Phys. Rev. **129**, 62 (1963)
21. V. Olevano, L. Reining, Phys. Rev. Lett. **86**, 5962 (2001)
22. F. Bechstedt, R. Del Sole, G. Cappellini, L. Reining, Solid State Commun. **84**, 765 (1992)
23. A. Riefer, F. Fuchs, C. Rödl, A. Schleife, F. Bechstedt, R. Goldhahn, Phys Rev B **84**(7), 075218 (2011)
24. W.G. Schmidt, M. Albrecht, S. Wippermann, S. Blankenburg, E. Rauls, F. Fuchs, C. Rödl, J. Furthmüller, A. Hermann, Phys. Rev. B **77**, 035106 (2008)
25. P.H. Hahn, W.G. Schmidt, K. Seino, M. Preuss, F. Bechstedt, J. Bernholc, Phys. Rev. Lett. **94**, 037404 (2005)
26. W.G. Schmidt, S. Glutsch, P.H. Hahn, F. Bechstedt, Phys. Rev. B **67**, 085307 (2003)
27. A. Ruini, M. Caldas, G. Bussi, E. Molinari, Phys. Rev. Lett **88**(20) (2002)
28. A.Y. Golovacheva, A.N. Romanov, V.B. Sulimov, J. Phys. Chem. A **109**, 3244 (2005)
29. A. Riefer, E. Rauls, W.G. Schmidt, J. Eberhard, I. Stoll, J. Mattay, Phys Rev B **85**, 165202 (2012)

Spinodal Decomposition Kinetics of Colloid-Polymer Mixtures Including Hydrodynamic Interactions

Alexander Winkler, Peter Virnau, and Kurt Binder

Abstract The phase separation dynamics of a model colloid-polymer mixture is studied by taking explicitly the hydrodynamic interactions caused by the solvent into account. Based on the studies on equilibrium phase behavior we perform a volume quench from the homogeneous region of the phase diagram deep into the region where colloid-rich and polymer-rich phases coexist. We demonstrate that the Multiparticle Collision Dynamics (MPCD) algorithm is well suited to study spinodal decomposition and present first results on the domain growth behavior of colloid-polymer mixtures in quasi two-dimensional confinement. On the one hand side we find that the boundary condition of the solvent with respect to the repulsive walls strongly influences the phase separation dynamics and on the other hand we show that the wetting behavior of the system leads to changes in the demixing pattern morphology over time and hence affects the domain growth laws.

1 Project Description

This project deals with computer simulations of suspensions of colloidal particles in confinement under non-equilibrium conditions. Colloidal suspensions are a prototype system of soft matter, and the softness of these systems make them ideally suited for studying them under far from equilibrium conditions.

The term colloid in general refers to systems that are structured on the nano- to micrometer scale. Food, detergents, paints and many biological substances are colloids. In the context of soft matter physics, however, the term is usually restricted to suspensions of nano- to micrometer sized particles or droplets dispersed in a liquid. Additionally to their obvious use in technological applications, colloids are

A. Winkler (✉) · P. Virnau · K. Binder
Johannes Gutenberg-Universität Mainz, Staudinger Weg, 7 55099 Mainz, Germany
e-mail: winkalex@uni-mainz.de

interesting model systems for research on statistical mechanics. The interactions between colloidal particles can be tuned e.g. by addition of salt to suspensions of charged particles or by grafting polymers onto their surfaces, and a large number of experimental techniques are available to study their properties (for overviews see [1–5]).

In binary fluid mixtures, the growth of the linear dimension $l(t)$ with time t after the system was quenched from the one-phase region into the two-phase region can be described by simple power laws $l(t) \propto t^\alpha$. The exponent α not only depends on the dimensionality of the system, but also whether or not hydrodynamic effects are considered. In two dimensions, droplet coagulation yields $l(t) \propto t^{0.5}$ [6], while the Lifshitz-Slyozov mechanism gives $l(t) \propto t^{1/3}$ [7] like in three dimensions (see [8] for a recent study of a pure two-dimensional system). For compositions, where domain structures are bicontinuous, hydrodynamic mechanisms yield $l(t) \propto t$ in three dimensions and $l(t) \propto t^{2/3}$ in two dimensions if $l \gg L_{in} = \eta^2/(\rho\gamma)$ (here η is the shear viscosity, ρ the density and γ the interfacial tension between A and B-rich domains) [9]. However, experiments and simulations often find slow transients before asymptotic power laws occur [10, 11], and for two-dimensional systems it is still controversial whether a scaling description in terms of universal power laws holds at all. In quasi two-dimensional systems the situation becomes even more complex because the wetting behavior of the system leads to additional dynamics perpendicular to the walls which affect the coarsening pattern morphology. There exist reports where the Lifshitz-Slyozov ($\alpha = 1/3$) power law for droplet patterns was exceeded in quasi two dimensions which was commented as a result of a wetting layer backflow [12, 13].

In this project we shall study the dynamics of a model colloid-polymer mixtures between two walls. Polymers are described as soft spheres weakly interacting with each other, while colloid-polymer and colloid-colloid pairs interact with the (repulsive) Weeks-Chandler-Andersen potential, so that a depletion attraction between colloids results, similar to the Asakura-Oosawa model. This model was developed and tested by Zausch et al. [14]. In this context, the phase behavior of this model was determined in the bulk.

1.1 Description of Methods and Algorithms

In mesoscale simulations a simple continuum description based on the Navier-Stokes equations is sometimes not sufficient. A fully atomistic approach on the other hand may also not be well suited since only the short time behavior is accessible with present-day computers due to the high number of microscopic degrees of freedom.

The Multi Particle Collision Dynamics (MPCD) algorithm [15–17] is easy to implement and a comparatively fast approach which takes advantage of explicit ideal solvent particles for which analytic expressions, e.g. transport coefficients, exist. The basic idea is to extend the ordinary Velocity Verlet based Molecular

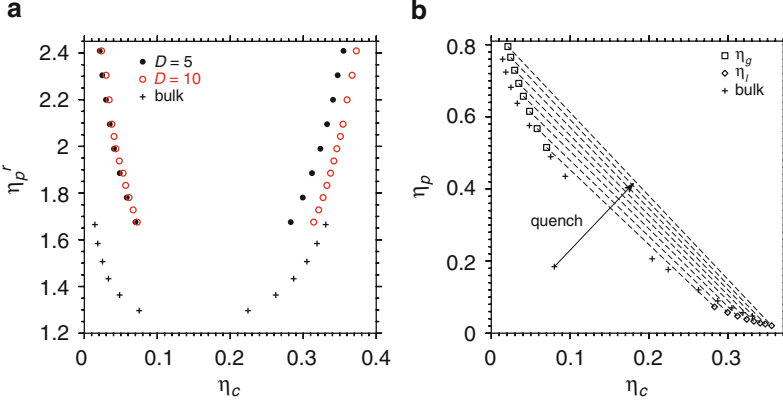


Fig. 1 Phase diagrams of the continuous AO-model. **(a)** Colloid coexistence densities as a function of polymer reservoir packing fraction $\eta_p^r = \frac{\pi}{6} d_p^3 \exp(\beta \mu_p)$ (here, d_p is the polymer diameter, μ_p is the chemical potential of the polymers and β is the inverse temperature) for wall distances $D = 5$ and $D = 10$. **(b)** The same colloid coexistence densities as a function of polymer packing fraction η_p for wall distance $D = 5$. In both plots the crosses represent the bulk phase diagram data for comparison

Dynamics (MD) simulations by a stochastic collision step followed by a streaming step of the solvent. The momentum is locally conserved by the division of the system into cells which contain the ideal, point like solvent particles with mass m and velocity v . The collision step consists of an independent random rotation of the relative velocities $v_i - u$ of the solvent particles in each cell, where u is the average velocity in the cell. Then, a streaming step is performed where the solvent particles propagate the distance $x = v_i \tau$, τ being the discretization time step of the solvent.

There are different techniques to couple the embedded fluid particles (like colloids and polymers) to the solvent. The simplest way is to exchange momentum between solvent particles and fluid particles by inserting also the fluid particles in addition to the solvent particles into the cells during the collision step [18].

1.2 Previous Work

We were able to determine the static properties of the continuous Asakura-Oosawa model confined between two planar walls by means of free energy calculations. We used cluster-move based grand canonical Monte Carlo simulations [14, 19] together with successive umbrella sampling [20] and parallel Wang-Landau sampling [21, 22] to calculate the coexistence densities and interfacial free energies for various parameter sets. The resulting phase diagrams are shown in Fig. 1.

2 Progress Report (Jan.–Apr. 2012)

We were able to implement and test extensively the MPCD algorithm. We used the standard halo layer domain decomposition technique to treat the embedded particles and the parallelization approach proposed by Sutmann et al. [23] for the solvent particles. This level of parallelization allows us to use 1,024 cores (or more) to study successfully the phase separation kinetics in the Asakura-Oosawa model in huge systems (max. 1,100,000 MD particles and 52,000,000 solvent particles) over multiple scales of MD time. We performed a volume quench (see Fig. 1b) from the homogeneous phase into the region where the colloid-rich phase and the polymer-rich phase coexist and recorded the average domain growth as a function of MD time. The investigation of spinodal decomposition kinetics of colloid-polymer mixtures, which are intrinsically asymmetric binary mixtures [14], is a very challenging problem in terms of performance on parallel super computers. The domain growth of colloid and polymer phases leads to an asymmetric distribution of the computational effort, because the polymer-rich phase consists typically out of two to four times more particles than the colloid rich phase. This fact leads to a strong slowdown of the simulation with respect to the Molecular Dynamics force calculation when the system demixes over time. The treatment of solvent particles in the framework of MPCD algorithms hides this imbalance between the various processes, so that we achieve almost the same performance as in simple Molecular Dynamics simulations even though 52,000,000 solvent particles are present (see Fig. 2). The simulations corresponding to the NOHI (no hydrodynamic interactions) performance curve are very closely related to standard MD simulations, since no explicit solvent particles are present and the overall performance is set by the velocity Verlet integration step. We conclude that the problem of spinodal decomposition is very well suited for MPCD algorithms, since the imbalance of computational effort in the velocity Verlet integration of fluid particles gets hidden by the homogeneous computational effort which is needed to handle the solvent. For huge systems at the late time stages of phase separation the net cost of explicit solvent particles in the framework of MPCD is only about approximately 10 % in comparison to standard Molecular Dynamics simulations.

For the first wall distance under consideration, $D = 5$, which is rather small in comparison to the lateral system dimensions of $L_x = L_y = 256$ (unit length scale is a colloid diameter), the system is expected to exhibit spinodal decomposition kinetics with a two-dimensional character. Via the MPCD algorithm it was possible to change the solvent properties like wall boundary conditions and hydrodynamic coupling without influencing the static properties of the model. From the results for wall distance $D = 5$ we are able to conclude that without hydrodynamic interactions (NOHI) the phase separation kinetics is distinctively slower than with included hydrodynamic interactions (MBST Stick, MBST Slip). Moreover, we discover a strong influence of the confinement due to the boundary conditions of solvent particles. When the solvent interacts via perfect slip (specular reflection at the walls) we discover that the domain growth is accelerated in comparison

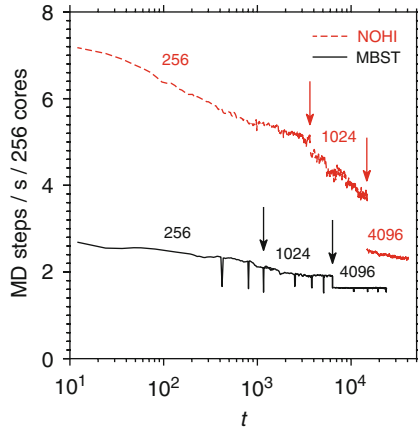


Fig. 2 Performance of the MPCD simulations on the super computer Hermit for system size $256 \times 256 \times 10$ for two different thermostat types as a function of MD time. The MBST (Maxwell Boltzmann Stochastics Thermostat) curve corresponds to fully included hydrodynamics by explicit solvent particles, while in the NOHI (no hydrodynamic interactions) performance curve no explicit solvent particles are present. *Arrows* indicate a change in the number of used cores. The overall performance for 1,024 and 4,096 cores was divided by 4 and 16 respectively. The sharp minima in the MBST data correspond to checkpoints where a full snapshot is written to the hard disk

to stick boundary conditions (bounce back reflection rules). When switching off hydrodynamic interactions the slowest domain growth mechanism is observed (see the snapshot series in Fig. 3). The development of the average domain size l_d over time is shown in Fig. 4. The domain growth law is altered from an $1/3$ power law behavior to a $2/3$ power law when comparing perfect stick boundary conditions and perfect slip boundary conditions [24]. We believe to explain this change in the phase separation kinetics via hydrodynamic screening caused by the walls.

We discover further a non-monotonic behavior of the dynamics perpendicular to the walls which influences the morphology of the observed demixing patterns (see the transition from percolated structures to droplets in the case of slip boundaries in Fig. 3). These dynamics perpendicular to the walls result from the fact that the colloids try to form a wetting layer at the walls while the polymers do not favor the walls. This becomes even more pronounced in the case of wall distance $D = 10$. The morphology of the coarsening pattern plays a crucial rule with respect to the growth law dynamics. Only in the case of percolating structures the hydrodynamic long range interactions are expected to cause a $2/3$ power law for the late times. In contrast, the coarsening via droplets is predicted to happen according to a power law with exponent $1/3$.

For quasi two-dimensional systems we find that the dynamics perpendicular to the walls change the morphology of the phase separation patterns over time. By choosing different ratios between the colloid and polymer concentration (labeled as quench1, quench2, quench3 respectively) we can influence the morphology and




Fig. 3 Snapshot series of the demixing process. Only the polymers are shown (*black dots*). Time increases from top to bottom. The *left series* corresponds to hydrodynamics with perfect slip boundary conditions. The *middle column* corresponds to hydrodynamics with stick boundary conditions and the *right column* corresponds to switched off hydrodynamics. (a) Time $t = 12$. (b) Time $t = 1,200$. (c) Time $t = 12,000$. (d) Time $t = 30,000$

hence, the domain growth laws (Fig. 5). The Euler characteristic (Fig. 5a) represents a measure of morphology, values larger than zero correspond to a droplet structure, values below zero correspond to a structure dominated by holes and values close

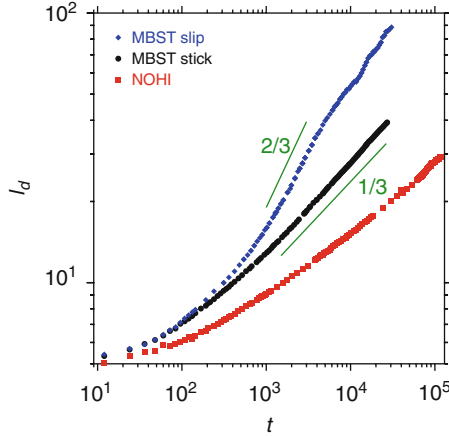


Fig. 4 Growth of the average domain size l_d as a function of MD time (1 MD time unit corresponds to 500 integration steps) for three different solvent properties in a system of size $256 \times 256 \times 5$. *Red squares* represent the data for switched off hydrodynamic interactions (NOHI), while the *blue diamonds* and *black full circles* show the data obtained with the Maxwell Boltzmann Stochastics Thermostat (MBST) with perfect slip and perfect stick boundaries respectively

to zero represent percolation patterns. The data becomes meaningful from time $t = 600$ on, when colloid and polymer curves start to behave in an antisymmetrical manner. The data was obtained from single runs and has a statistical error of typically $\delta\chi \approx 0.005$. Part (b) of the same figure shows the corresponding growth laws of the average domain size the corresponding decomposition pattern of which are shown just beneath in the same figure. Here, the perfect slip boundary condition was used, since only in the case of not screened hydrodynamics we expect to see a difference in the domain growth with respect to the underlying morphology. One can clearly see that being closer to a percolation pattern (quench3) leads to the fastest domain growth while crossing the percolation line slows down the demixing process drastically (which might be evidence for the so-called pinning effect [25–27]). From these first results for $D = 10$ we conclude that the spinodal decomposition of asymmetric binary mixtures (with respect to their bulk phase behavior as well as with respect to the condensation behavior at the walls) in quasi two-dimensional confinement does not show a universal behavior of the domain growth but rather the dynamic changes in morphology influence strongly the growth mechanism of the decomposition patterns. Additionally we see that in contrast to pure two-dimensional systems a screening of the hydrodynamic solvent correlations due to walls is present and that hydrodynamic acceleration of the coarsening mechanism can be only present for length scales comparable to the wall separation D in the case of stick boundary conditions.




Fig. 5 (a) Euler characteristic for colloid and polymer pattern at various concentration ratios using an MBS thermostat with slip boundaries. (b) Average domain size l_d for the three different colloid-polymer concentrations obtained from a single run as a function of time. As a guide for the eye power laws with different slopes are included. At the *bottom* a snapshot series for the various concentration ratios is given. Only the polymers are shown (*black dots*)

3 Outlook

It is planned to complete the studies for wall distance $D = 10$ with the focus on quench 3 and compare similar as for $D = 5$ the influence of hydrodynamics on the phase separation kinetics. Furthermore, the extreme case of an almost pure

two-dimensional system with $D = 1.5$ (where two particles do not fit on top of each other anymore) will be studied to clarify to which extent the dimensionality influences the spinodal decomposition and whether it is a continuous or a more sharp transition of the domain growth behavior when the system gets access to the third dimension. The extreme case of $D = 1.5$ is additionally useful in the sense that the morphology pattern is not influenced dynamically anymore by a competing dynamics perpendicular to the walls. Finally, we want to address the question whether wetting layer backflow (which can only exist in quasi two-dimensional systems) can enhance the phase separation kinetics in the case of droplet patterns.

Acknowledgements We would like to thank the DFG (TR6, A5) for funding of the project. Computing resources used to test the algorithms, obtain first results and the static properties of the continuous AO model at JUROPA supercomputer in Jülich are gratefully acknowledged. Especially acknowledged are G. Gompper, R. Winkler, C. Huang and G. Sutmann for extremely helpful discussions and advices.

References

1. T. Palberg. Crystallization kinetics of repulsive colloidal spheres. *J. Phys.: Condens. Matter*, 11:R323, 1999.
2. P. N. Pusey. Colloidal suspensions. In J. P. Hansen, D. Levesque, and J. Zinn-Justin, editors, *Liquids, freezing, and glass transition*, page 763, Amsterdam, 1991. North Holland.
3. E. Bartsch. Diffusion in concentrated colloidal suspensions and glasses. *Curr. Opin. Colloid Interface Sci*, 3:577, 1998.
4. H. Löwen. Colloidal dispersions in external fields: recent developments. *J. Phys.: Condens. Matter*, 20:404201, 2008.
5. New developments in colloid science. *MRS Bull.* 29 (2004).
6. K. Binder and D. Stauffer. Theory for slowing down of relaxation and spinodal decomposition of binary-mixtures. *Phys. Rev. Lett.*, 33:1006, 1974.
7. I.M. Lifshitz and V.V. Slyozov. The kinetics of precipitation from supersaturated solid solutions. *J. Phys. Chem. Solids*, 19:35, 1961.
8. Daniel Reith, Katarzyna Bucior, Leonid Yelash, Peter Virnau, and Kurt Binder. Spinodal decomposition of polymer solutions: molecular dynamics simulations of the two-dimensional case. *J. Phys.: - Condens. Matter*, 24(11):115102, MAR 21 2012.
9. H. Furukawa. Effect of inertia on droplet growth in a fluid. *Phys. Rev. A*, 31:1103, 1985.
10. S.K. Das, J. Horbach, K. Binder, and S. Puri. Molecular dynamics study of phase separation kinetics in thin films. *Phys. Rev. Lett.*, 96:016107, 2006.
11. Leonid Yelash, Peter Virnau, Wolfgang Paul, Kurt Binder, and Marcus Mueller. Spinodal decomposition of polymer solutions: A parallelized molecular dynamics simulation. *Phys. Rev. E*, 78(3, Part 1):031801, SEP 2008.
12. SK Das, S Puri, J Horbach, and K Binder. Spinodal decomposition in thin films: Molecular-dynamics simulations of a binary Lennard-Jones fluid mixture. *Phys. Rev. E*, 73(3, Part 1):031604, MAR 2006.
13. Michael J. A. Hore and Mohamed Laradji. Dissipative particle dynamics simulation of the interplay between spinodal decomposition and wetting in thin film binary fluids. *J. Chem. Phys.*, 132(2):024908, JAN 14 2010.
14. J. Zausch, P. Virnau, J. Horbach, R. L. C. Vink, and K. Binder. Statics and dynamics of colloid-polymer mixtures near their critical point of phase separation: A computer simulation of a continuous AO model. *J. Chem. Phys.*, 130:064906, 2009.

15. G. Gompper, T. Ihle, K. Kroll, and R.G. Winkler. Multi-particle collision dynamics: A particle-based mesoscale simulation approach to the hydrodynamics of complex fluids. *Adv. Poly. Sci.*, 221:1, 2009.
16. A. Malevanets and R. Kapral. Mesoscopic model for solvent dynamics. *J. Chem. Phys.*, 110:8605, 1999.
17. A. Malevanets and R. Kapral. Solute molecular dynamics in a mesoscale solvent. *J. Chem. Phys.*, 112:7260, 2000.
18. A. Malevanets and J. Yeomans. Dynamics of short polymer chains in solution. *EPL*, 52:231, 2000.
19. R. L. C. Vink and J. Horbach. Grand canonical Monte Carlo simulation of a model colloid-polymer mixture: Coexistence line, critical behavior, and interfacial tension. *J. Chem. Phys.*, 121:3253, 2004.
20. P. Virnau and M. Müller. Calculation of free energy through successive umbrella sampling. *J. Chem. Phys.*, 120:10925, 2004.
21. F. Wang and D. P. Landau. Efficient, multiple-range random walk algorithm to calculate the density of states. *Phys. Rev. Lett.*, 86:2050, 2001.
22. A. Winkler, D. Wilms, P. Virnau, and K. Binder. Capillary condensation in cylindrical pores: Monte Carlo study of the interplay of surface and finite size effects. *J. Chem. Phys.*, 133:164702, 2010.
23. G. Sutmann, R.G. Winkler, and G. Gompper. Simulating hydrodynamics of complex fluids: Multi-particle collision dynamics coupled to molecular dynamics on massively parallel computers. (in preparation).
24. A. Winkler, R.G. Winkler, P. Virnau, G. Gompper, and K. Binder. Multiparticle collision dynamics study of hydrodynamic mechanisms of spinodal decomposition in confined colloid-polymer mixtures. (in preparation).
25. H. Chen and A. Chakrabarti. Hydrodynamic effects on domain growth in off-critical polymer blends. *J. Chem. Phys.*, 108(14):6006–6013, APR 8 1998.
26. B. Crist. On “pinning” domain growth in two-phase polymer liquids. *Macromolecules*, 29(22):7276–7279, OCT 21 1996.
27. J. Lauger, R. Lay, and W. Gronski. The percolation-to-cluster transition during spinodal decomposition of an off-critical polymer mixture - observation by light-scattering and optical microscopy. *J. Chem. Phys.*, 101(8):7181–7184, OCT 15 1994.

Multi Relaxation Time Lattice Boltzmann Simulations of Multiple Component Fluid Flows in Porous Media

Sebastian Schmieschek, Ariel Narváez, and Jens Harting

Abstract The flow of fluid mixtures in complex geometries is of practical interest in many fields, ranging from oil recovery to freeze-dried food products. Due to its inherent locality the lattice Boltzmann method allows for straightforward implementation of complex boundaries and excellently scaling parallel computations. The widely applied Bhatnagar Gross Krook (BGK) scheme, used to model the contribution of particle collisions to the velocity field, does however suffer from limitations in precision that become more prominent with increasing surface to volume ratio. To increase the accuracy of simulations of mixtures in porous media, we integrated a so-called multi relaxation time (MRT) collision scheme with a pseudo-potential method for fluids with multiple components. We describe some optimisation details of the implementation and present test results verifying the physical accuracy as well as benchmarks obtained on the XC2 Opteron cluster at the Scientific Supercomputing Centre Karlsruhe (SSCK).

1 Introduction

Fluid flows involving multiple components and phases are of interest in many different fields of application, e.g. oil recovery [1], soil treatment and decontamination [2], food processing [3] and fuel cell optimisation [4]. This has motivated ongoing theoretical [5] and experimental [6] as well as numerical [7–12] research.

S. Schmieschek · A. Narváez · J. Harting (✉)

Institute for Computational Physics, University of Stuttgart, Pfaffenwaldring 27,
70569 Stuttgart, Germany

Department of Applied Physics, Eindhoven University of Technology, Den Dolech 2,
5600MB Eindhoven, The Netherlands

e-mail: jens@icp.uni-stuttgart.de

An important measure to describe the behaviour of fluid flows in porous media is the permeability. It is relating an externally applied mean pressure gradient to a resulting mean flow rate through the medium.

Numerous numerical models have been utilized to model flow in porous media, exhibiting strengths in different areas [7–9, 13]. While finite element methods (FEM) allow for a very efficient calculation of certain systems, the incorporation of complex boundary conditions may become very tedious. Particulate models such as dissipative particle dynamics (DPD) or molecular dynamics (MD) ease the integration of various fluid components but require time averaging to produce strong statistics. Furthermore, microscopic modeling as in MD does not allow to reach length- and time-scales relevant for the simulation of porous media. While in direct comparison of algorithmic costs to macroscopic (FEM) and other mesoscopic methods (DPD), lattice Boltzmann methods (LBM) seem to be expensive, their inherent thermostatics suppresses noise in the results and thus makes the method very competitive. In addition, the localised nature of the model eases parallelization as well as the integration of complex boundary geometries.

In the decades since its introduction, the LBM has been subject to several extensions, e.g. to account for multiphase flows [16–18], thermal transport phenomena [19, 20] or reactive flows [21]. Furthermore weaknesses of the model were discovered and accounted for [22, 23]. One prominent issue of the widely used Bhatnagar Gross Krook (BGK) collision scheme is the unphysical dependence of the position of simple bounce back boundaries on the relaxation time parameter τ . While in straight channel geometries, the thusly introduced error is easy to account for, in complex boundary geometries a significant relaxation time dependent error is observed [12, 23]. Amongst several other improvements, the introduction of a so-called multi relaxation time (MRT) collision scheme to the LB algorithm allows to mend this shortcoming.

Since many flow situations in porous media involve multiple fluid phases or components, a consistent algorithm which takes phase separation or interfacial tension into account is required. In this contribution we report aspects of the integration of a MRT collision scheme with our pseudo-potential multiphase lattice Boltzmann implementation *lb3d* and present results of ongoing physics- and performance benchmarks towards a more accurate and precise simulation tool for multiphase flows in porous media. Figure 1 shows example visualisations of our test cases: on the left hand side a rendering of the model porous medium, a BCC sphere packing and on the right hand side three snapshots of the domains of two immiscible fluids in a ternary mixture for different concentrations of an amphiphilic species.

2 Simulation Method

The lattice Boltzmann equation

$$\underbrace{f_k(\mathbf{x} + \mathbf{c}_k \Delta t, t + \Delta t) - f_k(\mathbf{x}, t)}_{\text{Advection}} = \underbrace{\Omega(f_k, f_k^{\text{eq}})}_{\text{Collision}}$$




Fig. 1 *Left:* Rendering of the surface area of the model porous medium used in our simulations. A body centred cubic (BCC) sphere packing. The permeability κ of this system is analytically accessible [14, 15]. *Right:* Rendering of the dynamics of the spinodal decomposition of two immiscible components in a ternary mixture. Depicted is the state reached after simulation of 10,000 timesteps with surfactant concentrations of $C_s = \{0, 18.75, 31.25\}\%$

is modeling fluid flows in terms of the single particle distribution function f . It is discretised on a lattice with velocities \mathbf{c}_k , $k = 0..b$. Our model relies on the so-called D3Q19 lattice, comprised by $b = 18$ velocities in direction of the face and edge-centers of a three dimensional cube, as well as an additional zero-velocity. The time development of f is described by an undisturbed linear motion (advection) and an operator Ω representing particle interactions. The time step Δt and the length of the lattice constant Δx are further on chosen to equal one. The expression f_k^{eq} designates the local equilibrium distribution, depending on the conserved properties density ρ and mean velocity \mathbf{u} . Its form can be chosen depending on the problem to simulate. In the scope of this work a third order expansion of the local Maxwell distribution with the lattice speed of sound $c_s = 1/\sqrt{3}$ and lattice weights w_k is used [24]:

$$f_k^{\text{eq}} = w_k \rho \left[1 + \frac{(\mathbf{c}\mathbf{u})}{c_s^2} + \frac{(\mathbf{c}\mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u}\mathbf{u}}{2c_s^2} + \frac{(\mathbf{c}\mathbf{u})^3}{6c_s^6} - \frac{\mathbf{u}\mathbf{u}(\mathbf{c}\mathbf{u})}{2c_s^4} \right]$$

For the BGK approach the collision operator

$$\Omega_{\text{BGK}} = -\frac{1}{\tau} [f_k(\mathbf{x}, t) - f_k^{\text{eq}}(\mathbf{x}, t)],$$

with a single relaxation time τ is used. The relaxation time is determining the kinematic viscosity $\nu = c_s^2 (\tau - \frac{1}{2})$. With the above described choices the equation of state (EOS) $P = \sum_k f_k^{\text{eq}} c_s^2 = \rho c_s^2 = \rho k_B T$ of an ideal gas with density ρ and Temperature T is recovered, where k_B denotes the Boltzmann constant.

While this formulation readily allows to solve the Navier-Stokes equations up to second order accuracy, there exist several limitations to the model. Due to the single relaxation time scale the Prandtl number, i.e. the ratio of momentum to energy transport, is fixed. In the low viscosity regime $\tau \rightarrow 0.5$ strong numerical instabilities are observed. Furthermore, the position of a bounce-back (no-slip) boundary is viscosity dependent.

These issues can be addressed by introducing a different collision operator, allowing to vary the Prandtl number and increasing the overall numerical stability. The collision operator of the generalised (MRT) LBM can be written as [25]

$$\Omega_{\text{MRT}} = -\mathcal{M}^{-1} \hat{\mathcal{S}} \mathcal{M} [|f(\mathbf{x}, t) - |f^{\text{eq}}(\mathbf{x}, t)|].$$

Herein \mathcal{M} is an invertible transformation matrix, relating the stochastic moments of the single particle velocity distribution f to linear combinations of its discrete components f_k . It can be obtained by a Gram Schmidt orthogonalization of a matrix representation of the stochastic moments. The latter can be related to physical properties such as density ρ , velocity \mathbf{u} , momentum \mathbf{p} , energy e , etc. The collision process is performed in the space of moments, where $\hat{\mathcal{S}}$ is a diagonal matrix of the individual relaxation times. Additional to the shear viscosity, the bulk viscosity $\xi = \frac{2}{3}c_s^2 \left(\frac{1}{\tau_{\text{bulk}}} - \frac{1}{2} \right)$ can be varied.

The Shan Chen pseudo-potential method alters the local mean velocity

$$\mathbf{u} = \frac{\sum_{\alpha} (\mathbf{u}_{\alpha} \cdot \rho_{\alpha}) / \tau_{\alpha}}{\sum_{\alpha} (\rho_{\alpha}) / \tau_{\alpha}} + \frac{\mathbf{F}^{\alpha} \cdot \tau_{\alpha}}{m_{\alpha}},$$

common to all species α of density ρ and mass m by introducing a phenomenological force

$$\mathbf{F}^{\alpha\bar{\alpha}} = \frac{dp^{\alpha}}{dt}(\mathbf{x}) = -\psi^{\alpha}(\mathbf{x}) \sum_{\alpha} g_{\alpha\bar{\alpha}} \sum_{k=0}^b \psi^{\bar{\alpha}}(\mathbf{x}') \mathbf{c}_k,$$

which is defined to be proportional to a functional of the fluid densities $\psi^{\alpha} = 1 - e^{-\rho}$ and a coupling parameter $g_{\alpha\bar{\alpha}}$ determining the nature and magnitude of interaction between the components α and $\bar{\alpha}$. This accounts for additional (non-ideal) terms in the equation of state [16, 26]. For attractive self-interactions ($g_{\alpha\alpha} < 0$) of a magnitude larger than a critical value g_c a phase separation in a vapour and a liquid phase can be observed. Repulsive interactions between two components ($g_{\alpha\bar{\alpha}} > 0$) are utilised to model systems of partly miscible or immiscible fluid mixtures. While the input parameters are determined strictly phenomenological, this approach has recently been shown to be equivalent to the explicit adjustment of the free energy of the system [27].

Amphiphiles (surfactants) are comprised of a hydrophilic (water-loving) head group and a hydrophobic (oil-loving) tail. Amphiphilic behaviour is modeled by a dipolar moment \mathbf{d} with orientation ϑ defined for each lattice site. It is relaxing in a BGK-like process, where the equilibrium moment is dependent on the surrounding

fluid densities [28]. The introduction of the dipole vector accounts for three additional Shan Chen type interactions, namely an additional force-term

$$\mathbf{F}^{\alpha s} = -2\psi^\alpha(\mathbf{x}, t)g_{\alpha s} \sum_{k \neq 0} \tilde{\mathbf{d}}(\mathbf{x} + \mathbf{c}_k, t) \cdot \Theta_k \psi^s(\mathbf{x} + \mathbf{c}_k, t)$$

for the regular fluid components α imposed by the surfactant species s . Therein, the tilde denotes post collision values and the second rank tensor $\Theta_k \equiv \mathbb{I} - 3\frac{\mathbf{c}_k \mathbf{c}_k}{c^2}$, with the identity operator \mathbb{I} weights the dipoles force contribution according to the orientation relative to the density gradient. The surfactant species is subject to forcing as well, where the contribution of the regular components α is given by

$$\mathbf{F}^{\alpha s} = 2\psi^s(\mathbf{x}, t)\tilde{\mathbf{d}}(\mathbf{x}, t) \cdot \sum_{\alpha} g_{\alpha s} \sum_{k \neq 0} \Theta_k \psi^\alpha(\mathbf{x} + \mathbf{c}_k, t),$$

and

$$\begin{aligned} \mathbf{F}^{ss} = & -\frac{12}{\|\mathbf{c}_s\|^2} \psi^s(\mathbf{x}, t)g_{ss} \cdot \sum_{k \neq 0} \psi^s(\mathbf{x} + \mathbf{c}_k, t) \cdot \left(\tilde{\mathbf{d}}(\mathbf{x} + \mathbf{c}_k, t) \cdot \Theta_k \cdot \tilde{\mathbf{d}}(\mathbf{x}, t)\mathbf{c}_k \right. \\ & \left. + [\tilde{\mathbf{d}}(\mathbf{x} + \mathbf{c}_k)\tilde{\mathbf{d}}(\mathbf{x}, t) + \tilde{\mathbf{d}}(\mathbf{x}, t)\tilde{\mathbf{d}}(\mathbf{x} + \mathbf{c}_k, t)] \cdot \mathbf{c}_k \right) \end{aligned}$$

is the force due to self-interaction of the amphiphilic species [29].

3 Implementation

The lattice Boltzmann implementation *lb3d* provides functionality to simulate three-dimensional simple, binary immiscible and ternary amphiphilic fluid systems using the Shan Chen pseudo-potential model for non-ideal fluid interactions.

The boundary conditions available include periodic boundaries, body forcing, and bounce-back boundaries as well as Lees-Edwards shearing for simple and binary fluid mixtures. The software is written in Fortran 90 and parallelized using MPI. It supports XDR and parallel HDF5 formats for I/O and provides checkpoint and restart for long-running simulations.

The code has been developed at University College London, University of Stuttgart and Eindhoven University of Technology. It has been ported to many supercomputers worldwide, where it has shown excellent scalability. Recently it has been shown to scale linearly on up to 262,144 cores on the European Blue Gene/P system Jugene [30]. We were able to confirm very good scaling of our ternary systems on the Opteron based XC2 as well (Fig. 2, left).




Fig. 2 *Left:* Speedup measure for a ternary system of 256^3 lattice sites on the XC2 at SSC Karlsruhe. The *solid line* represents ideal scaling relative to 16 core performance, the *dotted line* represents ideal scaling relative to 128 core performance. We observe very good scaling. *Right:* Performance comparison, measured in thousand lattice updates per second (kLUPS) produced on a single processor core based on the Intel Nehalem microarchitecture, running at a clock rate of 2.4 GHz. Results for a variation in component count using a BGK or MRT collision scheme, respectively. The components *a* and *b* denote immiscible Newtonian fluids, whereas component *s* is an amphiphilic fluid species. With increasingly complex fluid interaction calculations, the contribution of a more costly collision scheme is rendered less important

In recent years, *lb3d* has been coupled to a molecular dynamics solver in order to simulate complex fluids containing particulate components, such as blood [31], or nano-particle stabilised emulsions and suspensions [32]. Moreover, the code has been used to study for example the self-assembly of cubic mesophases [33], micro-mixing [34], flow through porous media [10] and fluid surface interactions [35, 36]. A refactored version of the code with limited functionality that focuses mainly on multiphase fluid simulation functionalities, has recently been released under the LGPL [37].

In order to improve the accuracy and numerical stability of *lb3d* with respect to the application to the simulation of multiphase flows in porous media, recently a MRT collision model was integrated with the software. While the MRT collision algorithm is more complex than the BGK collision scheme and can cause significant performance loss when implemented naively, the increase in calculation cost can be dramatically reduced. We take advantage of two properties of the system to minimize the impact of the additional MRT operations on the code performance:

1. The symmetry of the lattice allows to precalculate the sum and difference of discrete velocities which are linear dependent, thus saving at least half of the calculation steps [38].
2. The equilibrium stochastical moments can be expressed as functions of the conserved properties density and velocity, thus saving the transformation of the equilibrium distributions [25].

As such, the performance penalty could be reduced below 17%, which is close to the minimal additional cost reported in [25]. Since in multiphase systems the relative cost of the collision scheme is further reduced, the use of the MRT scheme has even less impact on the performance (Fig. 2, right).




Fig. 3 *Left*: Double logarithmic plot of the dynamics of the average domain size for different surfactant concentrations. After an initial phase, resulting from the chosen homogenous random initialisation of the fluid densities, the average domain size follows a power-law. Increasing surfactant concentration is slowing down the domain growth. With exceeding a concentration of $C_s \approx 35\%$ the critical micelle concentration is reached and a stable microemulsion with a minimal domain size is formed. *Right*: Time exponent of the dynamics of the average domain size as a function of the surfactant concentration, determined by a fitted power-law behaviour depicted on the left-hand side. At low concentrations $C_s \lesssim 10\%$ the impact of added surfactant on the domain size dynamics is close to negligible. As the concentration approaches the saturation of the interface the retardation of the decomposition becomes very sensitive to added surfactant

4 Results

To evaluate the functionality of the newly implemented MRT collision scheme, simulations of the influence of surfactant concentration on the decomposition process of two immiscible fluid components, earlier performed with *lb3d* utilising the BGK method are reproduced [33].

While keeping the total density of all fluid components constant at $\rho_{\text{tot}} = 0.8$ – equivalent to constant pressure – the concentration of the surfactant component is varied between 0 and 37.5% in steps of 6.25%. Simulations are carried out on a domain of 256^3 lattice sites. The coupling parameter between the immiscible components b (blue) and r (red) is set to $g_{br} = 0.08$. The coupling parameters to the surfactant species are set to $g_{\alpha s} = -0.006$ and $g_{ss} = -0.003$, respectively. The fluid components are initialised at homogenous density distribution and a random dipole orientation. The average domain size is extracted by evaluation of the Fourier transformation of the field of the order parameter $\Phi = \rho_b - \rho_r$ which becomes zero at the interface of the immiscible components [33].

As can be observed from Fig. 3, two regimes of the dynamics can be distinguished. At low to intermediate surfactant concentrations ($C_s < 35\%$) spinodal decomposition is taking place, resulting in completely demixed fluids. Here, after an initial phase related to the formation of domains from the homogenous initialised fluids, the average domain size clearly follows a power-law with respect to time. Figure 3 (left) shows a double-logarithmic plot of the dynamics of the average domain sizes for varying surfactant concentration. The extracted numerical values for the power-law relation can be found in Fig. 3 (right). Average domain sizes

exceeding 80 lattice sites have been excluded from the fit, as in this range the periodic boundary conditions begin to influence the results. In the range of higher concentration close to the critical micelle concentration ($C_s = 31.25\%$) the initial phase is extended considerably as compared to lower concentrations. The impact of adding surfactant on the dynamics of the average domain size is almost an order of magnitude large here as compared to the low concentration range around 10%. With exceeding a concentration of $C_s \approx 35\%$ the critical micelle concentration is reached and stable microemulsions with a constant domain size are formed. These results are in agreement with the findings of the reference study and allow confidence in the MRT collision scheme.

To test the behaviour of the MRT scheme in a complex geometry, the flow of two miscible phases through the model BCC sphere packing porous medium, illustrated in Fig. 1 (left), is simulated. The permeability for this model system

$$\kappa = \frac{a^2}{6\pi\chi h}$$

can be derived analytically in terms of the side length of the unit cell a , the dimensionless drag h and a geometry-factor χ giving the ratio of actual to maximal sphere radius. We choose the maximum radius here, so $\chi = 1$. The dimensionless drag can be shown to depend on the system geometry only and can be expressed as an expansion of χ [14, 15].

In our simulation setup, we measure the permeability using Darcy's law

$$\kappa = -\rho\nu \frac{\langle u \rangle}{\langle \nabla P \rangle},$$

with the fluid density ρ , the kinematic viscosity ν , the average velocity $\langle u \rangle$ and the average pressure gradient $\langle \nabla P \rangle$. For the test runs two miscible fluids of equal density $\rho_r = \rho_b = 1.0$ and equal kinematic viscosity, varied by the relaxation parameter τ , are driven through the porous medium by applying a body force of $F = 0.001$ in lattice units in a four lattice sites long layer located in the center of a buffer region of 20 lattice sites width.

Figure 4 shows a comparison of absolute permeability measurements conducted in simulations with varying relaxation parameter τ using the MRT and BGK collision scheme, respectively. In these preliminary simulations the only relaxation parameter modified in the MRT scheme is the one related to the dynamic viscosity, all other parameters are kept to equal 1 or 0 for parameters related to conserved properties (ρ and \mathbf{u}). In this configuration at $\tau = 1$, the MRT scheme is therefore equivalent to the BGK scheme, as can be observed from the identical outcome of the permeability measure at that point. However, already for this parameter set, the dependence of the system's behaviour on the relaxation parameter is less pronounced. It has been shown that this dependence can be eliminated by a suitable choice of parameters [23]. However, previous attempts by other groups were not combined with the Shan Chen multicomponent method and thus did not allow to

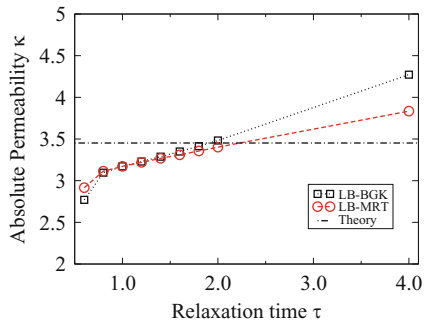


Fig. 4 Preliminary results of measurements of absolute permeability of the model BCC sphere packing porous medium for varying relaxation times τ . In this case we use non-optimised values for the MRT relaxation parameters, all equal to 1. In the case of $\tau = 1.0$, this reduces the MRT to the BGK-scheme as is here illustrated by identical results. By adjusting the relaxation parameter set, the dependence of the permeability measurement on the relaxation parameter can be completely resolved [23]. Since equal fluid properties were chosen, the relative permeability κ_α/κ is 0.5 for all cases

simulate multicomponent flow in porous media. All relative permeabilities κ_α/κ measured were found to be exactly 0.5 which is to be expected, since in the runs presented here all properties of the liquid components were set to be equal.

These promising results allow to proceed with further tests, such as to optimise the MRT parameter set and study the accuracy of relative permeability measurements for fluid mixtures with different properties. Possible future applications of this model include then permeability calculations for ternary amphiphilic mixtures and studies of imbibition and demixing processes in porous media, where our model can take effects of wettability into account as well.

5 Summary

To improve accuracy and efficiency of the application of our lattice Boltzmann implementation for fluid flows with multiple components in porous media we have implemented a multi relaxation time collision scheme. The steps taken to optimize the performance have allowed to limit the increase in calculation time close to the expected minimum. Furthermore we could show that for multi-component systems the additional calculation time spent in the collision step is very small in comparison to the time increase due to the calculation of the interaction forces. The reproduction of prior simulations of the dynamics of an amphiphilic fluid mixture as well as a currently modest but optimisable improvement in accuracy on a broader scale of viscosity values allow confidence in the implementation. Preliminary tests of permeability measurements for multicomponent flows in porous media are promising and allow to proceed further towards productive application of the model to the simulation of ternary amphiphilic mixtures in porous media.

Acknowledgements This work was financed within the DFG priority program “nano- and microfluidics” and by the NWO/STW VIDI grant of J. Harting. Sebastian Schmieschek likes to express his gratitude to the HPC-Europa2 programme for funding during development. We thank the Scientific Supercomputing Center Karlsruhe for providing the computing time and technical support for the presented work.

References

1. L. Paterson. Diffusion-limited aggregation and two-fluid displacements in porous media. *Phys. Rev. Lett.*, 52:1621–1624, Apr 1984.
2. Z. Zhang and J. E. Smith. The velocity of dnapl fingering in water-saturated porous media: laboratory experiments and a mobile-immobile-zone model. *Journal of Contaminant Hydrology*, 49(3–4):335–353, 2001.
3. A. Datta. Porous media approaches to studying simultaneous heat and mass transfer in food processes. i: Problem formulations. *Journal of Food Engineering*, 80(1):80–95, 2007.
4. Z. Wang, C. Wang, and K. Chen. Two-phase flow and transport in the air cathode of proton exchange membrane fuel cells. *Journal of Power Sources*, 94(1):40–50, 2001.
5. R. de Boer. Highlights in the historical development of the porous media theory: Toward a consistent macroscopic theory. *Applied Mechanics Reviews*, 49(4):201–262, 1996.
6. R. Lenormand, E. Touboul, and C. Zarcone. Numerical models and experiments on immiscible displacements in porous media. *Journal of Fluid Mechanics*, 189:165–187, 1988.
7. M. Liu, P. Meakin, and H. Huang. Dissipative particle dynamics simulation of pore-scale multiphase fluid flow. *Water Resour. Res.*, 43(4):W04411, April 2007.
8. X. Yi, J. Ghassemzadeh, K. Shing, and M. Sahimi. Molecular dynamics simulation of gas mixtures in porous media. I. Adsorption. *J. Chem. Phys.*, 108(5):2178–2188, Feb 1998.
9. T. Y. Hou and X.-H. Wu. A multiscale finite element method for elliptic problems in composite materials and porous media. *Journal of Computational Physics*, 134(1):169–189, 1997.
10. J. Harting, M. Venturoli, and P. V. Coveney. Large-scale grid-enabled lattice-Boltzmann simulations of complex fluid flow in porous media and under shear. *Phil. Trans. R. Soc. Lond. A*, 362:1703–1722, 2004.
11. A. Narváez and J. Harting. A D3Q19 lattice-Boltzmann pore-list code with pressure boundary conditions for permeability calculations. *Advances in Applied Mathematics and Mechanics*, 2:685, 2010.
12. A. Narváez, T. Zauner, F. Raischel, R. Hilfer, and J. Harting. Quantitative analysis of numerical estimates for the permeability of porous media from lattice-Boltzmann simulations. *J. Stat. Mech: Theor. Exp.*, 2010:P211026, 2010.
13. J. Harting, T. Zauner, A. Narvaez, and R. Hilfer. Flow in porous media and driven colloidal suspensions. In W. Nagel, D. Kröner, and M. Resch, editors, *High Performance Computing in Science and Engineering '08*. Springer, 2008.
14. H. Hasimoto. On the periodic fundamental solutions of the Stokes equations and their application to viscous flow past a cubic array of spheres. *Journal of Fluid Mechanics*, 5(02):317–328, 1959.
15. A. Sangani and A. Acrivos. Slow flow through a periodic array of spheres. *International Journal of Multiphase Flow*, 8(4):343–360, 1982.
16. X. Shan and H. Chen. Lattice Boltzmann model for simulating flows with multiple phases and components. *Phys. Rev. E*, 47(3):1815, 1993.
17. M. R. Swift, W. R. Osborn, and J. M. Yeomans. Lattice Boltzmann simulation of nonideal fluids. *Phys. Rev. E*, 75(5):830, 1995.
18. A. K. Gunstensen, D. H. Rothman, S. Zaleski, and G. Zanetti. Lattice Boltzmann model of immiscible fluids. *Phys. Rev. A*, 43(8):4320, 1991.

19. B. Dünweg, U. D. Schiller, and A. J. C. Ladd. Statistical mechanics of the fluctuating lattice Boltzmann equation. *Phys. Rev. E*, 76:36704, 2007.
20. R. Adhikari, K. Stratford, M. E. Cates, and A. J. Wagner. Fluctuating lattice Boltzmann. *Europhys. Lett.*, 71:473, 2005.
21. Q. Kang, P. C. Lichtner, and D. Zhang. Lattice Boltzmann pore-scale model for multicomponent reactive transport in porous media. *J. Geophys. Res.*, 111(B5):B05203, May 2006.
22. X. He, Q. Zou, L.-S. Luo, and M. Dembo. Analytic solutions of simple flows and analysis of nonslip boundary conditions for the lattice Boltzmann BGK model. *Journal of Statistical Physics*, 87:115–136, 1997.
23. C. Pan, L.-S. Luo, and C. T. Miller. An evaluation of lattice Boltzmann schemes for porous medium flow simulation. *Computers & Fluids*, 35(8–9):898–909, 2006.
24. S. Succi. *The lattice Boltzmann equation for fluid dynamics and beyond*. Oxford University Press, 2001.
25. D. d’Humières, I. Ginzburg, M. Krafczyk, P. Lallemand, and L. Luo. Multiple-relaxation-time lattice Boltzmann models in three dimensions. *Phil. Trans. R. Soc. A*, 360:437451, 2002.
26. X. Shan and G. Doolen. Multicomponent lattice-Boltzmann model with interparticle interaction. *J. Stat. Phys.*, 81(112):379, 1995.
27. M. Sbragaglia and X. Shan. Consistent pseudopotential interactions in lattice Boltzmann models. *Phys. Rev. E*, 84:036703, Sep 2011.
28. M. Nekovee, P. V. Coveney, H. Chen, and B. M. Boghosian. Lattice-Boltzmann model for interacting amphiphilic fluids. *Phys. Rev. E*, 62(6):8282–8294, Dec 2000.
29. M. Nekovee, P. V. Coveney, H. Chen, and B. M. Boghosian. Lattice-Boltzmann model for interacting amphiphilic fluids. *Phys. Rev. E*, 62:8282, 2000.
30. D. Groen, O. Henrich, F. Janoschek, P. Coveney, and J. Harting. Lattice-boltzmann methods in fluid dynamics: Turbulence and complex colloidal fluids. In W. F. Berndt Mohr, editor, *Jülich Blue Gene/P Extreme Scaling Workshop 2011*. Jülich Supercomputing Centre, 52425 Jülich, Germany, apr 2011. FZJ-JSC-IB-2011-02; <http://www2.fz-juelich.de/jsc/docs/autoren2011/mohr1/>.
31. F. Janoschek, F. Toschi, and J. Harting. Simplified particulate model for coarse-grained hemodynamics simulations. *Phys. Rev. E*, 82:056710, 2010.
32. F. Jansen and J. Harting. From Bijels to Pickering emulsions: a lattice Boltzmann study. *Phys. Rev. E*, 83:046707, 2011.
33. N. González-Segredo, J. Harting, G. Giupponi, and P. V. Coveney. Stress response and structural transitions in sheared gyroidal and lamellar amphiphilic mesophases: lattice-Boltzmann simulations. *Phys. Rev. E*, 73:031503, 2006.
34. A. Sarkar, A. Narváez, and J. Harting. Quantification of the degree of mixing in chaotic micromixers using finite time Lyapunov exponents. *Microfluidics and Nanofluidics*, in press, 2010.
35. C. Kunert and J. Harting. Roughness induced apparent boundary slip in microchannel flows. *Phys. Rev. Lett.*, 99:176001, 2007.
36. S. Schmieschek, A. V. Belyaev, J. Harting, and O. I. Vinogradova. Tensorial slip of superhydrophobic channels. *Phys. Rev. E*, 85:016324, 2012.
37. 'lb3d open source version 7' available online, March 2012. <http://mtp.phys.tue.nl/lb3d>.
38. B. Chun and A. J. C. Ladd. Interpolated boundary condition for lattice Boltzmann simulations of flows in narrow gaps. *Phys. Rev. E*, 75:066705, Jun 2007.

Simulation of Pre-planetesimal Collisions with Smoothed Particle Hydrodynamics II

R.J. Geretshauer, F. Meru, K. Schaal, R. Speith, and W. Kley

Abstract In the frame of planet formation by coagulation the growth step from millimetre-sized highly porous dust aggregates to massive kilometre-sized planetesimals is not well constrained. In this regime of pre-planetesimals, collisional growth is endangered by disruptive collisions, disintegration by rotation as well as mutual rebound and compaction. Since laboratory studies of pre-planetesimal collisions are infeasible beyond centimetre-size, we perform numerical simulations. For this purpose, utilise the parallel smoothed particle hydrodynamics (SPH) code `paraspH`. This program has been developed to simulate macroscopic highly porous dust aggregates consisting of protoplanetary material. We briefly introduce our porosity model and use it to perform simulations on the growth criteria of pre-planetesimals. With the aid of parameter studies we investigate fragmentation criteria in dust collisions depending on aggregate size and aggregate porosity. We extend a previous study on bouncing criteria of equally sized aggregates depending on their porosity and the presence of compacted shells of various porosities. Regarding the rotational stability of highly porous dust aggregates we theoretically derive fragmentation criteria for dust cylinders depending on angular velocity as well as porosity and perform suitable simulations.

R.J. Geretshauer (✉) · K. Schaal · W. Kley
Institut für Astronomie und Astrophysik, Abteilung Computational Physics, Eberhard Karls
Universität Tübingen, Auf der Morgenstelle 10, 72076 Tübingen, Germany
e-mail: ralf.j.geretshauer@uni-tuebingen.de

F. Meru
Institut für Astronomie und Astrophysik, Abteilung Computational Physics, Eberhard Karls
Universität Tübingen, Auf der Morgenstelle 10, 72076 Tübingen, Germany

Institut für Astronomie, ETH Zürich, Wolfgang-Pauli-Strasse 27, 8093 Zürich, Switzerland

R. Speith
Physikalisches Institut, Eberhard Karls Universität Tübingen, Auf der Morgenstelle 14, 72076
Tübingen, Germany

1 Introduction

According to the coagulation scenario planet formation occurs in protoplanetary discs by a sequence of successive mutual collisions from micrometre-sized dust grains to planets. Three regimes can be distinguished: during the first step from grains to millimetre-sized pre-planetesimals growth proceeds in a fractal way [5]. The growth mechanism in the second growth step from millimetre-sized pre-planetesimals to kilometre-sized planetesimals is not well constrained and subject of this work. Once a sufficient population of planetesimals exists, the third growth step to full-size planets proceeds due to gravity-aided accretion [15]. An overview of this process can be found in Ref. [9].

We turn to the second growth step from pre-planetesimals to planetesimals which is currently subject of extensive numerical and experimental effort. Two important obstacles could be identified in the pre-planetesimal regime: fragmentation and bouncing barrier. The most serious barrier is the fragmentation barrier: the relative velocities between pre-planetesimals increase as the size of the pre-planetesimals increases. The consequence is an increased probability for catastrophic disruption. Since experimental data on pre-planetesimal collisions are rather sparse, collision maps often have to contain simplistic assumptions such as a constant velocity threshold for fragmentation over several orders of magnitude in size [2, 16]. Transition velocities between parameter regions of positive and negative growth are often treated as independent from object porosities.

The second obstacle to planetesimal formation is the bouncing barrier. In this scenario [16, 28], millimetre to centimetre-sized pre-planetesimals become more and more compacted during their collisional evolution. The relative velocities for these objects are still rather low. However, due to their compact structure the pre-planetesimals rebound instead of sticking to each other or destroying each other upon collision. Recent studies indicate that if growth is halted at smaller sizes than millimetre size, the bouncing barrier could even be beneficial to planetesimal growth if some objects grow out of this barrier and sweep up all the smaller particles [27].

The need for a detailed numerical investigation of pre-planetesimal collisions arose from the lack of experimental data under realistic protoplanetary disc conditions in this field. For this reason we developed a smoothed particle hydrodynamics (SPH) code to simulate porous pre-planetesimal material [9, 11] and calibrated it with laboratory benchmark experiments [17]. For a detailed classification of the collision outcome we developed the four-population model [10]. Numerical model and classification scheme were utilised to show that the stability of pre-planetesimals is significantly decreased if they possess an inhomogeneous structure [12, 13]. Additionally, preliminary studies showed that the conditions for rebound are drastically changed if pre-planetesimals feature a compacted outer shell [9, 13]. We also begun with a large-scale parameter study of pre-planetesimal collisions in the decimetre regime which already indicated that size ratio and porosity of the collision partners have an influence on the collision outcome, which is not negligible [9, 13].

In this article we continue this work. Initially, in Sect. 2, we briefly review the code and porosity model. Detailed quantitative investigations of size- and porosity-dependence of the fragmentation velocity threshold are presented in Sect. 3.1. The influence of the compactness of hard shells is studied in Sect. 3.2. Finally, in Sect. 3.3, analytical formulae for the stability of rotating dust aggregates are presented and counterchecked with simulations. Further areas of research are described in the outlook Sect. 4.

2 SPH Code and Porosity Model

2.1 Solid Body SPH

For the simulation of pre-planetary collisions we utilise the numerical method *smoothed particle hydrodynamics (SPH)* [14, 21] together with extensions for solid body mechanics [1, 20, 22] and a suitable porosity model [11, 25].

As a method which was originally developed for pure hydrodynamics, SPH relies on the usual equations of continuum mechanics to ensure conservation of mass, momentum, and energy. Under external forces, stresses develop inside a solid body to restore its original shape. Hence, the main properties of solid body mechanics enter in the momentum equation via the stress tensor $\sigma_{\alpha\beta}$, which is defined as

$$\sigma_{\alpha\beta} = -p\delta_{\alpha\beta} + S_{\alpha\beta} , \quad (1)$$

where p and $S_{\alpha\beta}$ denote the hydrostatic pressure and deviatoric stress tensor, respectively. Throughout this article Greek indices denote the spatial components and the Einstein summation notation is applied. The time evolution of the deviatoric stress tensor is computed according to Hooke's law in frame invariant Jaumann rate form [1, 9, 11, 13, 24]. The transition to plastic deformation due to shear is determined by the von Mises criterion

$$S^{\alpha\beta} \rightarrow f S^{\alpha\beta} , \quad (2)$$

where $f = \min [Y^2(\phi)/3J_2, 1]$ and $Y = Y(\phi)$ is the shear strength. The quantity $J_2 = S_{\alpha\beta} S^{\alpha\beta}$ is the second irreducible invariant of the deviatoric stress tensor. Sub-grid porosity is modelled by the filling factor ϕ , which is defined as

$$\phi = \frac{\rho}{\rho_s} = 1 - \Phi , \quad (3)$$

where ρ , ρ_s , and Φ are the density of the porous material, the density of the matrix material, and the porosity, respectively. The pre-planetary material is mono-disperse spherical SiO_2 dust [3, 4] dust for which $\rho_s = 2,000 \text{ kg m}^{-3}$.




Fig. 1 Calibrated strength curves. The compressive strength $\Sigma(\phi)$, tensile strength $T(\phi)$, and shear strength $Y(\phi)$ as a result of the calibration process described in Ref. [11]

The elastic as well as the plastic evolution of the hydrostatic part of the stress tensor are computed according to the filling-factor-dependent bulk modulus $K(\phi)$ as well as the compressive strength $\Sigma(\phi)$ and tensile strength $T(\phi)$ [11–13], respectively. The bulk modulus scales with the filling factor by the following relation

$$K(\phi) = K_0 \left(\frac{\phi}{\phi_{\text{RBD}}} \right)^\gamma \quad (4)$$

where $\gamma = 4$ and K_0 is the bulk modulus of an uncompressed random ballistic deposition (RBD) dust sample with $\phi_{\text{RBD}} = 0.15$ [3]. This value was calibrated to be $K_0 = K(\phi_{\text{RBD}}) = 4.5$ kPa [11].

The strength quantities [11, 17] are illustrated in Fig. 1 and represent transition thresholds from the elastic to the plastic regime. The tensile strength is given by a power-law

$$T(\phi) = -10^{2.8+1.48\phi} \text{ Pa} . \quad (5)$$

The compressive strength takes the form

$$\Sigma(\phi) = p_m \left(\frac{\phi_{\text{max}} - \phi_{\text{min}}}{\phi_{\text{max}} - \phi} - 1 \right)^{\Delta \ln 10} , \quad (6)$$

with $\phi_{\text{max}} = 0.58$ and $\phi_{\text{min}} = 0.12$, which denote maximum and minimum filling factor of the compressive strength relation. The quantity $\Delta \ln 10$ is the power of the expression with $\Delta = 0.58$. The constant $p_m = 260$ Pa is its mean pressure.

The shear strength is given by the geometric mean of Eqs. (5) and (6)

$$Y(\phi) = \sqrt{\Sigma(\phi)|T(\phi)|} . \quad (7)$$

2.2 *The Code `Parasph`*

The `parasph` code was developed by Hipp [19] and extended by Schäfer for the simulation of ductile, brittle, and porous media [23, 24]. The porosity model was improved and calibrated by Geretshauer for pre-planetesimal material [9, 11] together with other extensions. The program is based on the `parasph` library [6, 7]. This is a set of routines developed for a easier and faster handling of parallel particle codes. By means of this library the physical problem and the parallel implementation are clearly separated. `parasph` features domain decomposition, load balancing, nearest neighbour search, and inter-node communication. Moreover, SPH enhancements such as additional artificial stress and XSPH were implemented. The adaptive Runge-Kutta-Cash-Karp integrator has been used for the simulations presented here. For testing purposes an adaptive second order Runge-Kutta, and an Euler integrator was added. HDF5 (Hierarchical Data Format) was included as a compressed input and output file format with increased accuracy, which decreases the amount of required storage space considerably. The parallel implementation utilises the Message Passing Interface (MPI) library. Test simulations yielded a speedup of 120 on 256 single core processors of a Cray T3E and of 60 on 128 single core processors on a Beowulf-Cluster.

3 Results

The simulations in this section are carried out with 240,143–476,476 SPH particles depending on the size of the projectile. The program `parasph` is run on the NEC Nehalem cluster of the HLRS. Depending on the size of the problem 32–80 cores were used. The simulation time is strongly dependent on the collision velocity. In particular for fragmenting collisions the adaptive time step becomes as small as $\sim 10^{-5}$ s throughout the whole simulation. In contrast in bouncing and sticking collisions, the time step is initially low and then increases for the rest of the simulation. To give a rough number, simulations take 72–240 h wall-clock time for 1 s of simulated time depending on size of the problem and involved physical process.

3.1 *Velocity Thresholds*

The main focus of the project lies on the investigation of the growth conditions for macroscopic pre-planetesimals. In this section we carry out simulations to constrain the influence of projectile size and aggregate porosity on these conditions for centimetre to decimetre-sized pre-planetesimals. Although the velocity threshold




Fig. 2 Fragmentation threshold velocity against projectile size. The different colours represent the filling factors $\phi = 0.15$ (blue), $\phi = 0.25$ (red), and $\phi = 0.35$ (black). The arrows denote upper and lower limits. The target is a sphere of radius 10 cm. Due to the higher collision energy the threshold velocity for fragmentation decreases with increasing projectile size. Highly porous aggregates ($\phi = 0.15$) are less stable than aggregates of intermediate porosity ($\phi = 0.35$)

for fragmentation might be significantly lowered if the pre-planetesimals possess an inhomogeneous filling factor distribution [10], we restrain to initially homogeneous aggregates as a simplification and reference case.

Figure 2 illustrates the dependence of the fragmentation threshold velocity on the size of the impacting projectile. In each case the target is a porous dust sphere of radius 10 cm. Also the projectile has spherical shape. Both objects are homogeneous and of identical filling factor. We investigate three of the latter: $\phi = 0.15$, $\phi = 0.25$, and $\phi = 0.35$. This represents a range between high and intermediate porosity. All three curves have a similar shape. Very small projectiles intrude into the target and get swallowed. Only few dust fragments are ejected. With increasing projectile size the impact energy increases and hence the threshold for disruption decreases. However, compared to the immense range of sizes in pre-planetesimal collisions ranging from millimetre to kilometre sizes the difference in threshold fragmentation velocity within the quite narrow investigated size range is remarkable. We conclude that even small differences in size have to be considered in the field of pre-planetesimal growth. Figure 2 already gives rise to another important dependence: the aggregate porosity.

Figure 3 shows a surprising dependence of the threshold velocity on the aggregate porosity. As a standard setup for this study we collide two homogeneous




Fig. 3 Threshold velocity for fragmentation against filling factor. The target and projectile are spheres of radius 10 cm and 6 cm, respectively. The *arrows* denote *upper* and *lower* limits. Starting from low filling factors the threshold velocity increases because the strength of the aggregate increases. At a filling factor of $\phi \sim 0.37$ a sharp drop occurs

spherical aggregates with vanishing impact parameter. The target and projectile radii are 10 cm and 6 cm, respectively. At low filling factors or, equivalently, high porosity, the tensile and shear strength are low (see Eqs. (5), (7), and Fig. 1). As a consequence, such pre-planetesimals are rather fragile. However, the compressive strength is also rather low (Eq. (6) and Fig. 1) and hence energy can easily be dissipated by plastic deformation since the pressure threshold for plastic compression is exceeded in each collision.

As the filling factor increases the compressive strength, shear strength, and tensile strength increase. On the one hand the plastic deformability of the aggregate and hence the ability to dissipate energy hereby decrease. On the other hand, however, the aggregates become less fragile. Overall the aggregates gain a higher stability and the threshold velocity further increases for higher filling factors.

At a filling factor of $\phi \sim 0.35$ a sudden drop in the threshold velocity occurs. This drop is surprising and is caused by a complex interplay between the elastic and plastic properties of the aggregates which makes it difficult to capture this feature quantitatively. This issue is still object of ongoing research but the influence of the following aspects could be identified. (1) With increasing compressive strength $\Sigma(\phi)$ an increasing fraction of the initial kinetic energy can be stored in elastic loading of the aggregates (see also Sect. 3.2). As a consequence, this elastic energy is available for the separation of the aggregates in a bouncing collision. (2) With increasing ϕ also the bulk modulus of the aggregates increases (Eq. (4)) and the

aggregates become stiffer. As a result, the contact area between the aggregates decreases and less energy is necessary to separate the aggregates in a bouncing collision. (3) Because of partial sticking of the aggregates, a clean growth-neutral bouncing is unrealistic. Instead some mass transfer or even the rip out of larger chunks can be expected. (4) Due to the increase in $\Sigma(\phi)$ and $K(\phi)$ density waves of increasing amplitude propagate across the aggregate and lead to a local rarefaction of the material and consequential local reduction of shear and tensile strength (see also a similar phenomenon in Sect. 3.3). These density waves finally rip the aggregate apart even at low collision velocities.

The aspects (1) and (2) decrease the probability for sticking of the projectile to the target and increase the probability of bouncing of the targets. Aspect (4) increases the probability for fragmentation of the aggregate. The collisions at the drop of the threshold velocity in Fig. 3 show a behaviour where the aggregates mainly bounce but partly also fragment or some mass transfer (aspect 3) occurs. Since sticking and consequential growth and bouncing (with some fragmentation or mass transfer), resulting in neutral or negative growth, are two clearly distinct physical processes, a sharp drop in the describing threshold velocity can be expected which at least qualitatively explains the shape of the curve in Fig. 3. We note that a collision model based on experimental data [16] which features only the categories “porous” and “compact” aggregates uses a filling factor of $\phi = 0.40$ as separation criterion which is close to the drop filling factor of $\phi = 0.37$. Other experimental investigations [26] find that the maximum filling factor that can be reached for porous aggregates in protoplanetary discs is roughly $\phi \sim 0.33$. However, this value is obtained for polydisperse SiO_2 dust which is less compressible. This value can be expected to be higher for our monodisperse dust. To summarise, the drop in our threshold velocity curve occurs where also experimental data indicate a significant change in the collision behaviour of dust aggregates. However, the reasons for this drop require a closer investigation.

3.2 *Bouncing, Hard Shells and Porosity*

In our first study we investigate the influence of porosity on the bouncing and sticking behaviour. This is to assess whether the experiments with intermediate porosity [2] and high porosity [18] can be combined into the collision map presented by Güttler et al. [16].

For this, we conduct simulations of collisions between a homogeneous target and projectile. We distinguish three cases where both objects feature a uniform filling factor of $\phi = 0.15$, $\phi = 0.35$, and $\phi = 0.55$. According to the compressive strength relation this is nearly equivalent to maximum, intermediate, and minimum porosity. The aggregates thus feature very low, intermediate, and very high pressure thresholds for plastic deformation. Resulting from this property we expect a strong dependence of the filling factor on the sticking and bouncing behavior. The




Fig. 4 Cross-section through the outcome of collision simulations with homogeneous aggregates for high ($\phi_i = 0.15$, *left*), intermediate ($\phi_i = 0.35$, *middle*) and low porosity ($\phi_i = 0.55$, *right*). The projectile radius is equal to the target radius. The collision velocities are in the first line (*a*, *e*, *i*) 0.1, in the second line (*b*, *f*, *j*) 0.3, in the third line (*c*, *g*, *k*) 0.5, and in the fourth line (*d*, *h*, *l*) 1.0 ms^{-1} . The highly porous dust aggregates stick for each of the simulated collision velocities. The filling factor is increased in the vicinity of the impact site. Additionally the aggregates are deformed to almost kidney shape for 1.0 ms^{-1} (*d*). In contrast, the homogeneous aggregates with $\phi_i = 0.35$ exclusively rebound. The region of increased filling factor in the interior increases with collision velocity. The impact site is flattened. The aggregates with low filling factor also exclusively rebound at the given velocities. However, no significant compaction or deformation is visible except some fragmentation for the highest collision velocity (*l*)

impacting projectile has a radius $r_p = 10 \text{ cm}$ just like the target. The impact velocity v_0 is 0.1, 0.3, 0.5, and 1.0 ms^{-1} .

The theoretical demarcation between sticking and bouncing in Ref. [16] makes the following assumptions: (1) elastic deformation of the aggregates, and (2) the filling factor in the contact region is not changed in the collision. These assumptions are too simplistic.

1. The assumption of elastic deformation of the aggregates is only valid for filling factors close to the maximum filling factor. In particular for highly porous aggregates the deformation is highly plastic and as a consequence the contact area between the aggregates is increased compared to elastic contact (see e.g. Fig. 4d).




Fig. 5 Cross-section through collision results involving aggregates with a hard shell of $0.1r$. The interior of all aggregates is highly porous $\phi_c = 0.15$, which would lead to perfect sticking if the aggregates were homogeneous. However, in the *left* (a–d), *middle* (e–h), and *right column* (i–l) the filling factors of the hard shell are 0.35, 0.45, and 0.55, respectively. Initially, both aggregates are spheres of 10 cm radius. The displayed cross-sections show the situation after the impact with 0.1 (a, e, i), 0.3 (b, f, g), 0.5 (c, g, k), and 1.0 ms^{-1} (d, h, l). Even thin hard shells can lead to bouncing if the filling factor is sufficiently high

- As it can clearly be seen in Fig. 4, the filling factor is highly increased in the contact area. This leads to an increased tensile strength in this region (Eq. (5)). An increased tensile strength, however, also increases the contact energy, which promotes sticking.

We conclude, that because of these effects in particular for highly porous dust aggregates the threshold velocity for sticking is much larger than presented in [16]. Consequently, the parameter space where bouncing occurs is much smaller than assumed in this reference and sticking dominates for low velocities and low filling factors.

We use the same setup as above regarding ratio of target and projectile radius and collision velocity, but we add a hard outer shell (Fig. 5). The thickness of the hard shell is given as a fixed fraction of target and projectile radii, respectively, i.e. 0.1, 0.2, 0.3, 0.4. The core of all aggregates has a filling factor $\phi_c = 0.15$, whereas the filling factor of the hard shell ϕ_h takes the values 0.35, 0.45, and 0.55.

From the comparison with aggregates without hard shell, which exclusively resulted in sticking (cf. Fig. 4, first row), it is evident that hard shells do have an influence on the bouncing behaviour of dust aggregates. This is because for aggregates with $\phi = 0.15$ the compressive strength is very low. As a consequence, nearly the total kinetic energy of the impact is dissipated by plastic deformation and nearly no elastic loading of the colliding objects is possible.

Conversely, for aggregates with hard shells ($\phi_h = 0.35$) the plastic deformation threshold is higher for the shell. During the impact, the shell is elastically loaded and the aggregates rebound. However, in the immediate area around the impact site the deformation threshold for the shell is exceeded and plastic deformation takes place in the hard shell. Therefore, the tensile $T(\phi)$ and shear $Y(\phi)$ strengths are increased in this region and counteract the bouncing. However, hard shells of intermediate porosity are still sufficiently plastically deformable such that bouncing is still a rare event in the investigated parameter space.

The situation changes for hard shell filling factors of $\phi_h = 0.45$ and 0.55. The value of $\phi_h = 0.55$ is close to maximum compaction. Thus, the hard shell rather breaks than being plastically deformed. As a consequence of the lacking plastic deformability the collisions of aggregates with denser hard shells exclusively result in bouncing and partial fragmentation (of the shell).

We conclude that hard shells of intermediate porosity in general do not prevent sticking. Only in a few cases, where the hard shell was thick enough for sufficient elastic loading or at the right velocity for a thin hard shell bouncing occurs. However, if the hard shells become more compact, even thin shells suffice to yield bouncing collisions. For very compact hard shells and higher collision velocities partial fragmentation of mostly shell material can be expected.

3.3 *Stability of Rotating Dust Aggregates*

Particularly in aggregate collisions where the impact parameter is not vanishing, the aggregates or fragments start to rotate [24]. However, high spinning rates may lead to additional disruption of the aggregates since tidal forces may tear them apart. To assess this effect for pre-planetesimal collisions, we analytically derive critical angular velocities for a long rotating dust cylinder depending on its filling factor and radius. Additionally we carry out simulations of rotating cylinders with the same specifications. As a consistency check of the SPH code we compare the stress evolution inside the cylinder with the theoretical reference. We also compare the simulated critical angular velocity with the expected value.

Table 1 Results of the rotating cylinder study. The cylinders have an initial filling factor ϕ_i . The rotation is accelerated by $\dot{\omega}$. The onset of fragmentation starts at a critical time t_c and a critical angular velocity ω_c . Simulated and theoretical value of the latter are compared

ϕ_i	$\dot{\omega}[\text{s}^{-1}]$	$t_c[\text{s}]$	$\omega_c[\text{s}^{-1}]$ simulation	$\omega_c[\text{s}]$ theory	%
0.15	1	9.95	9.64	9.52	101.26
0.25	1	10.2	9.86	14.3	68.95
0.35	0.2	31.25	6.08	17.98	33.82
0.45	0.2	28.4	5.50	23.56	23.34
0.55	0.2	27.05	5.20	34.36	15.13

According to our porosity model presented in Sect. 2 disruption can occur when the stress inside the cylinder represented by the full stress tensor $\sigma_{\alpha\beta}$ exceeds either the tensile strength $T(\phi)$ or the shear strength $Y(\phi)$. For each angular velocity ω the stress state is well defined [8]. From this one can derive two critical angular velocities ω_c^T and ω_c^Y . The first arises from the tensile strength criterion and reads:

$$\omega_c^T = \sqrt{\frac{-12(1-\nu)}{\rho_0 r_0^2 (3-\nu)}} T(\phi), \quad (8)$$

where r_0 and ν are the radius of the cylinder and the Poisson ratio, respectively. Fragmentation due to exceeding the shear strength occurs when the following critical angular velocity is reached

$$\omega_c^Y = \sqrt{\frac{8(1-\nu)}{\rho_0 r_0^2 (3-4\nu)}} Y(\phi). \quad (9)$$

The cylinder fragments at the minimum of one of these velocities. The theoretical values for the fragmentation criterion can be found in Table 1 for the investigated initial filling factors ϕ_i .

To check our theoretical predictions, we carry out simulations of a rotating cylinder consisting of porous SiO_2 dust. The radius of the cylinder is $r_0 = 10$ cm and its height is $h = 60$ cm. It rotates about the z-axis which represents the central axis of the cylinder and is accelerated by an constant angular acceleration $\dot{\omega}$. The exact values are given in Table 1.

Initially the cylinder deforms elastically and the tensions inside the body increase. We measure the principal stresses of $\sigma_{\alpha\beta}$ along a line from (0,0,0) which represents the centre of mass of the cylinder radially outwards along the x-axis. The resulting principal stresses in cylinder coordinates $\sigma_{\phi\phi}$, σ_{rr} and σ_{zz} are displayed in Fig. 6 for $t = 10$ s and $t = 25$ s, shortly before the cylinder fragments. During its whole elastic evolution the stresses match the theoretical predictions extremely well. This confirms the validity of our code in the elastic regime.




Fig. 6 Stresses $\sigma_{\alpha\beta}$ of the stable rotating cylinder with $\phi_i = 0.55$. The principal stresses in cylinder coordinates $\sigma_{\phi\phi}$ (top), σ_{rr} (middle), and σ_{zz} are measured along a line from the centre of the cylinder radially outwards at two different times $t = 10$ s (left) and $t = 25$ s (right). In this elastic regime the simulated stresses (black) almost perfectly follow the curves expected from the theory (colored). The radius of the cylinder is 10 cm, however due to SPH smoothing this boundary is exceeded in the graphs. At the rim the stresses start to deviate from the expected values

With respect to fragmentation, the simulation for the cylinder with $\phi_i = 0.15$ reproduces the critical angular velocity very well (see Table 1). For higher initial filling factors however, the predicted critical angular velocity is much larger than the simulated value. The discrepancy increases with ϕ_i . The reason for this is illustrated in Fig. 7 which displays the evolution of the density inside the cylinder. As the body rotates density waves develop inside. These lead to a local rarefaction of the material. As a consequence the tensile and shear strength is diminished in these




Fig. 7 Density evolution of a rotating cylinder with $\phi_i = 0.55$ after the stable phase. The images show a cut along the x - z plane through the centre of the cylinder which is rotating about the z -axis. The density is colour-coded. While at $t = 25$ s the density distribution is still homogeneous, at $t = 26$ s density fluctuations at the *top* and *bottom* of the cylinder start to develop. Together with the density the tensile and shear strength are locally diminished. These regions are weak spots at which rupture starts ($t = 27$ s and $t = 27.4$ s)

regions (Fig. 8). However, the stresses increase with increasing angular velocity. The rarefied regions serve as weak spots from which rupture can develop due to material failure which arises from a stress exceeding the strength values.

4 Outlook

The presented results encourage further studies for a deeper insight into the outcome of pre-planetary collisions. First of all, the sudden drop in fragmentation threshold velocity at a given porosity requires further investigations. More simulations have to be carried out to assess other dependencies of this behaviour, for example the projectile size. A high resolution in time is desirable at this transition point to understand the complex interplay between elastic and plastic behaviour. With respect to the rotating cylinder the development of the elastic waves should be studied in more detail.

Recently, new measurements for another possible pre-planetary material became available. This demands for an extensive parameter study using this new material to explore possible differences with respect to growth conditions. Off-centre collisions of pre-planetary material are still a desiderate because they represent the more realistic situation in the disc.




Fig. 8 Local shear strength $Y(\phi)$ (red) and von Mises stress $\sqrt{3J_2}$ (black) for $r > 0$ at $z = 0.2$ for a rotating cylinder with $\phi_1 = 0.55$. According to the von Mises yielding criterion the cylinder starts to fragment once the von Mises stress exceeds the shear strength. In the stable regime ($t = 20$ s) $Y(\phi)$ is two magnitudes larger than the stresses. However, density fluctuations lead also to fluctuations in $Y(\phi)$ which are already visible at $t = 26$ s and $t = 27$ s. Eventually, the shear strength locally falls below the von Mises stress at $t = 27.15$ s and $r \sim 1$ cm and rupture sets in

To gain an idea of what realistic processed pre-planetesimals look like, it is desirable to study the collision history of an aggregate. For this, a homogeneous aggregate is subsequently collided with other aggregates of different sizes and filling factor at various collision velocities. The input information for a realistic collision sequence can be derived from global coagulation simulations.

References

1. Benz, W., Asphaug, E.: Impact Simulations with Fracture. I. Method and Tests. *Icarus* **107**(1), 98–116 (1994)
2. Blum, J., Münch, M.: Experimental Investigations on Aggregate-Aggregate Collisions in the Early Solar Nebula. *Icarus* **106**(1), 151–167 (1993)
3. Blum, J., Schräpler, R.: Structure and Mechanical Properties of High-Porosity Macroscopic Agglomerates Formed by Random Ballistic Deposition. *Phys. Rev. Lett.* **93**(11), 115,503 (2004)

4. Blum, J., Schräpler, R., Davidsson, B.J.R., Trigo-Rodríguez, J.M.: The Physics of Protoplanetary Dust Agglomerates. I. Mechanical Properties and Relations to Primitive Bodies in the Solar System. *The Astrophysical Journal* **652**(2), 1768–1781 (2006)
5. Blum, J., Wurm, G.: The Growth Mechanisms of Macroscopic Bodies in Protoplanetary Disks. *Annual Review of Astronomy and Astrophysics* **46**(1), 21–56 (2008)
6. Bubeck, T., Hipp, M., Hüttmann, S., Kunze, S., Ritt, M., Rosenstiel, W., Ruder, H., Speith, R.: Parallel SPH on Cray T3E and NEC SX-4 using DTS. In: E. Krause, W. Jäger (eds.) *High performance computing in science and engineering '98*, pp. 396–410. Springer, Berlin and New York (1999)
7. Bubeck, T., Hipp, M., Hüttmann, S., Kunze, S., Ritt, M., Rosenstiel, W., Ruder, H., Speith, R.: SPH test simulations on a portable parallel environment. In: *Proceedings of the Workshop on Physics and Computer Science*, pp. 139–155 (1999)
8. Chakrabarty, J.: *Theory of plasticity*, 3 edn. Elsevier/Butterworth-Heinemann, Amsterdam and Boston and Oxford (2006)
9. Geretshausen, R.J.: *Simulation of pre-planetary collisions with smoothed particle hydrodynamics*. Ph.D. thesis, Universität Tübingen, Tübingen (2011)
10. Geretshausen, R.J., Meru, F., Speith, R., Kley, W.: The four-population model: a new classification scheme for pre-planetary collisions. *Astronomy & Astrophysics* **531**, A166 (2011)
11. Geretshausen, R.J., Speith, R., Güttler, C., Krause, M., Blum, J.: Numerical simulations of highly porous dust aggregates in the low-velocity collision regime: Implementation and calibration of a smooth particle hydrodynamics code. *Astronomy and Astrophysics* **513**, A58 (2010)
12. Geretshausen, R.J., Speith, R., Kley, W.: Collisions of inhomogeneous pre-planetary bodies. *Astronomy & Astrophysics* **536**, A104 (2011)
13. Geretshausen, R.J., Speith, R., Kley, W.: Simulation of Pre-Planetary Collisions with Smoothed Particle Hydrodynamics. In: W.E. Nagel, D.B. Kröner, M. Resch (eds.) *High performance computing in science and engineering '11*, pp. 29–45. Springer, Berlin and New York (2012)
14. Gingold, R.A., Monaghan, J.J.: Smoothed particle hydrodynamics – Theory and application to non-spherical stars. *Monthly Notices of the Royal Astronomical Society* **181**, 375–389 (1977)
15. Goldreich, P., Lithwick, Y., Sari, R.: Final Stages of Planet Formation. *The Astrophysical Journal* **614**(1), 497 (2004)
16. Güttler, C., Blum, J., Zsom, A., Ormel, C.W., Dullemond, C.P.: The outcome of protoplanetary dust growth: pebbles, boulders, or planetesimals? - I. Mapping the zoo of laboratory collision experiments. *Astronomy and Astrophysics* **513**, A56 (2010)
17. Güttler, C., Krause, M., Geretshausen, R.J., Speith, R., Blum, J.: The Physics of Protoplanetary Dust Agglomerates. IV. Toward a Dynamical Collision Model. *The Astrophysical Journal* **701**, 130–141 (2009)
18. Heißelmann, D., Fraser, H.J., Blum, J.: Experimental Studies on the Aggregation Properties of Ice and Dust in Planet-Forming Regions. *International Astronautical Congress Abstracts* **58**, 1–6 (2007)
19. Hipp, M., Rosenstiel, W.: Parallel hybrid particle simulations using mpi and openmp. In: M. Danelutto, D. Laforenza, M. Vanneschi (eds.) *Euro Par 2004 parallel processing*, pp. 189–197. Springer (2004)
20. Libesky, L.D., Petschek, A.G.: Smooth Particle Hydrodynamics with Strength of Materials. In: H. Trease, M.J. Fritts, W.P. Crowley (eds.) *Advances in the Free-Lagrange method: including contributions on adaptive gridding and the smooth particle hydrodynamics method*, *Lecture notes in physics*, vol. 395. Springer (1991)
21. Lucy, L.B.: A numerical approach to the testing of the fission hypothesis. *The Astronomical Journal* **82**, 1013–1024 (1977)
22. Randles, P.W., Libesky, L.D.: Smoothed Particle Hydrodynamics: Some recent improvements and applications. *Computer Methods in Applied Mechanics and Engineering* **139**(1–4), 375–408 (1996)

23. Schäfer, C.: Application of Smooth Particle Hydrodynamics to selected Aspects of Planet Formation. Ph.D. thesis, Universität Tübingen, Tübingen (2005)
24. Schäfer, C., Speith, R., Kley, W.: Collisions between equal-sized ice grain agglomerates. *Astronomy and Astrophysics* **470**(2), 733–739 (2007)
25. Sirono, S.: Conditions for collisional growth of a grain aggregate. *Icarus* **167**(2), 431–452 (2004)
26. Teiser, J., Engelhardt, I., Wurm, G.: Porosities of Protoplanetary Dust Agglomerates from Collision Experiments. *The Astrophysical Journal* **742**(1), 5 (2011)
27. Windmark, F., Birnstiel, T., Güttler, C., Blum, J., Dullemond, C.P., Henning, T.: Planetesimal formation by sweep-up: how the bouncing barrier can be beneficial to growth. *Astronomy and Astrophysics* **540**, A73 (2012)
28. Zsom, A., Ormel, C.W., Güttler, C., Blum, J., Dullemond, C.P.: The outcome of protoplanetary dust growth: pebbles, boulders, or planetesimals? II. Introducing the bouncing barrier. *Astronomy and Astrophysics* **513**, A57 (2010)

The Stellar IMF at Very Low Metallicities

Gustavo Dopcke, Simon C.O. Glover, Paul C. Clark, and Ralf S. Klessen

Abstract The theory for the formation of the first population of stars (Pop. III) predicts an initial mass function (IMF) dominated by high-mass stars, in contrast to the present-day IMF, which tends to yield mostly stars with masses less than $1 M_{\odot}$. The leading theory for the transition in the characteristic stellar mass predicts that the cause is the extra cooling provided by increasing metallicity. In particular, dust can overtake H_2 as the leading coolant at very high densities. The aim of this work is to determine the influence of dust cooling on the fragmentation of very low metallicity gas. To investigate this, we make use of high-resolution hydrodynamic simulations with sink particles to replace contracting protostars, and analyze the collapse and further fragmentation of star-forming clouds. We follow the thermodynamic response of the gas by solving the full thermal energy equation, and also track the behavior of the dust temperature and the chemical evolution of the gas. We model four clouds with different metallicities (10^{-4} , 10^{-5} , $10^{-6} Z_{\odot}$, and 0), and determine the properties of each cloud at the point at which it undergoes gravitational fragmentation. We find evidence for fragmentation in all four cases, and hence conclude that there is no critical metallicity below which fragmentation is impossible. Nevertheless, there is a clear change in the behavior of the clouds at $Z = 10^{-5} Z_{\odot}$, caused by the fact that at this metallicity, fragmentation takes longer to occur than accretion, leading to a flat mass function at lower metallicities.

G. Dopcke (✉) · S.C.O. Glover · P.C. Clark · R.S. Klessen
Zentrum für Astronomie der Universität Heidelberg, Institut für Theoretische Astrophysik,
Albert-Ueberle-Str. 2, 69120 Heidelberg, Germany
e-mail: gustavo@uni-hd.de

1 How Are the Stars Formed in the Early Universe?

The first stars to form during the earliest stages of the the Universe were thought to give rise to massive stars, the so-called Population III (Pop. III), with numerical simulations predicting masses in the range 20–150 M_{\odot} [e.g. 1, 6, 28, 41]. However, recent results show that lower mass stars can also be formed, albeit with characteristic masses above the solar value [11, 12, 17, 18, 36, 38]. This contrasts with present-day star formation, which yields typically stars with masses less than 1 M_{\odot} [9, 24], and so at some point in the evolution of the Universe there must have been a transition from primordial (Pop. III) star formation to the mode of star formation we see today (Pop. II/I).

Metal line cooling and dust cooling are effective at lower temperatures and larger densities, and so it has been proposed that metal enrichment of the interstellar medium by previous generations of stars causes the transition from Pop. III to Pop. II. This suggests that there might be a critical metallicity Z_{crit} at which the mode of star formation changes.

The main coolants that have been studied in the literature are C II and O I fine structure emission [7, 8, 15, 21, 23, 29, 34, 35], and dust emission [e.g. 26, 27, 30, 32, 33]. Carbon and oxygen are identified as the key species because in the temperature and density conditions that characterize the early phases of Pop. III star formation, the O I and C II fine-structure lines dominate over all other metal line transitions [20]. By equating the C II or O I fine structure cooling rate to the compressional heating rate due to free-fall collapse, one can define critical abundances $[C/H] = -3.5$ and $[O/H] = -3.0$ ¹ for efficient metal line cooling [8].

A more promising way to form low mass Pop. II stars involves dust cooling. Dust cooling models [e.g. 13, 26, 27, 32, 33] predict a much lower critical metallicity ($Z_{\text{crit}} \approx 10^{-4} - 10^{-6} Z_{\odot}$), with most of the uncertainty coming from the nature of the dust in high-redshift galaxies.

At densities $n \gtrsim 10^{11} \text{ cm}^{-3}$ dust cooling becomes efficient [26], since inelastic gas-grain collisions are more frequent [19]. This cooling enhances fragmentation, and since it occurs at high densities, the distances between fragments can be very small [25, 27, 30, 31, 33]. In this regime, interactions between fragments will be common, and analytic models of fragmentation are unable to predict the mass distribution of the fragments. A full 3D numerical treatment, following the fragments, is needed.

In [13], we improved upon these previous treatments by solving the full thermal energy equation, and calculating the dust temperature through the energy equilibrium equation. We assumed that the only significant external heat source is the cosmic microwave background (CMB), and included its effects in the calculation of the dust temperature. We found that model clouds with metallicities as low as

¹ $[X/Y] = \log_{10}(N_X/N_Y)_{\star} - \log_{10}(N_X/N_Y)_{\odot}$, for elements X and Y, where \star denotes the gas in question, and where N_X and N_Y are the mass fractions of the elements X and Y.

$10^{-4} Z_{\odot}$ or $10^{-5} Z_{\odot}$ do indeed show evidence for dust cooling and fragmentation, supporting the predictions of [39, 40] and [10].

In this work, we simulate the evolution of star-forming clouds for a wider range of metallicities (10^{-4} , 10^{-5} , $10^{-6} Z_{\odot}$, and 0), and study the effect that this has on the mass function of the fragments that form. We also investigate how properties such as cooling and heating rates, and number of Bonnor-Ebert masses [5, 14] of the fragmenting clouds vary with metallicity and whether there is any systematic change in behavior with increasing metallicity.

2 Simulations

2.1 Numerical Method

We model the collapse of a low-metallicity gas cloud using a modified version of the Gadget-2 [37] smoothed particle hydrodynamics (SPH) code. To enable us to continue our simulation beyond the formation of the first very high density protostellar core, we use a sink particle approach [3, 22], in the same way as in [13]. Sink particles are created once the SPH particles are bound, collapsing, and within an accretion radius, h_{acc} , which we take to be 1.0 AU. The threshold number density for sink particle creation is $5.0 \times 10^{13} \text{ cm}^{-3}$. At the threshold density, the Jeans length at the minimum temperature reached by the gas is approximately one AU, while at higher densities the gas becomes optically thick and begins to heat up. Further fragmentation on scales smaller than the sink particle scale is therefore unlikely to occur. For further discussion of the details of our sink particle treatment, we refer the reader to [11].

For the metallicities and dust-to-gas ratios considered in this study, the dominant sources of cooling are the standard primordial coolants (H_2 bound-bound emission and collision-induced emission) and energy transfer from the gas to the dust. Collisions between gas particles and dust grains can transfer energy from the gas to the dust (if the gas temperature T is greater than the dust temperature T_{gr}), or from the dust to the gas (if $T_{gr} > T$). Full details of the dust cooling treatment can be found in [13].

3 Code Performance

Each of the simulations need approximately 130 kCPU-h to be concluded. In order to run them in a reasonable time, we need to make use of 256 or 512 CPUs. A minimum of 40 million SPH particles are required for resolving the fragment mass and correct predict the mass distribution of the protostellar cores.

Table 1 Parallel efficiency for different number of SPH particles in simulations run

CPUs	Particles			
	10^5	10^6	10^7	$4 \cdot 10^7$
8	1.00			
16	0.82			
32	0.52	1.00		
64	0.26	0.91	1.00	1.00
128		0.44	0.92	0.95
256			0.50	0.75
512				0.52

We also run simulations to evaluate the parallel efficiency using different number of CPUs and SPH particles. The time is normalized for the minimum number of CPUs, which is related to the number of SPH particles in the simulation. The simulation was run until a predicted result was achieved, without file writing, so that we could measure just the time needed for the CPUs to calculate the evolution of the gas and dust. The result is shown in Table 1.

In addition, to cover the parameter space, we run four simulations, resulting in 0.55 Mio. CPU-h. For resolution studies and post-processing, we had an additional overall cost of approximately 0.15 Mio. CPU-h. In total, this project required approximately 0.7 Mio. CPU-h

4 Computational Challenge

We performed a set of four simulations, with metallicities $Z/Z_{\odot} = 10^{-4}$, 10^{-5} , 10^{-6} , and the metal-free case. Each simulation used 40 million SPH particles. We used these simulations to model the collapse of an initially uniform gas cloud with an initial number density of 10^5 cm^{-3} and an initial temperature of 300 K. The cloud mass was $1,000 M_{\odot}$. We included small amounts of turbulent and rotational energy, with $E_{\text{turb}}/|E_{\text{grav}}| = 0.1$ and $\beta = E_{\text{rot}}/|E_{\text{grav}}| = 0.02$, where E_{grav} is the gravitational potential energy, E_{turb} is the turbulent kinetic energy, and E_{rot} is the rotational energy. The mass resolution is $2.5 \times 10^{-3} M_{\odot}$, which corresponds to 100 times the SPH particle mass [see e.g. 4]. The redshift chosen was $z = 15$, when the cosmic microwave background temperature was 43.6 K. The dust properties were taken from [16], and the dust grain opacities were calculated in the same fashion as in [2]. In the calculations, the opacities vary linearly with Z , which means for instance that for the $Z/Z_{\odot} = 10^{-4}$ calculations, the opacities were 10^{-4} times the original values.

5 Scientific Results

5.1 *Thermodynamical Evolution of Gas and Dust*

Dust cooling is a consequence of inelastic gas-grain collisions, and thus the energy transfer from gas to dust vanishes when they have the same temperature. We therefore expect the cooling to cease when the dust reaches the gas temperature. In order to evaluate the effect of dust on the thermodynamic evolution of the gas and verify this assumption, we plot in Fig. 1, the temperature and density for the various metallicities tested. We compare the evolution of the dust and gas temperatures in the simulations, at the point of time just before the formation of the first sink particle (see Table 2). The dust temperature (shown in blue) varies from the CMB temperature in the low density region to the gas temperature (shown in red) at much higher densities.

The efficiency of the cooling is also expressed in the temperature drop at high densities. The gas temperature decreases to roughly 400 K in the $10^{-5} Z_{\odot}$ simulation, and 200 K in the $Z = 10^{-4} Z_{\odot}$ case. This temperature drop significantly increases the number of Jeans masses present in the collapsing region, making the gas unstable to fragmentation. The dust and the gas temperatures couple for high densities, when the compressional heating starts to dominate again over the dust cooling. The subsequent evolution of the gas is close to adiabatic.

5.2 *Fragmentation*


The transport of angular momentum to smaller scales during the collapse leads to the formation of a dense disk-like structure, supported by rotation. This disk then fragments into multiple objects.

Figure 2 shows the density structure of the gas immediately before the formation of the first protostar. The top-left panel shows a density slice on a scale comparable to the size of the initial gas distribution. The structure is very filamentary and there are two main over-dense clumps in the center. If we zoom in on one of the clumps, we see that its internal structure is also filamentary. Observe that at large scales the gas cloud properties are the same for all metallicities. Differences in the thermodynamic evolution appear only at $n \gtrsim 10^{11} \text{ cm}^{-3}$ (see Fig. 1). As a consequence, we observe variations in the cloud structure only in the high-density regions.

For $Z = 10^{-6} Z_{\odot}$ and 0, the formation of spiral structures is not observed. In these two runs, star formation occurs mainly in the central clump.

At the beginning of the simulation, the cloud had $\sim 3 M_{BE}$. During the collapse, the gas cools and reaches $\sim 6 M_{BE}$ in all cases. Cooling and heating are different

Fig. 1 Dependence of gas and dust temperatures on gas density for metallicities 10^{-4} , 10^{-5} , and 10^{-6} and zero times the solar value, calculated just before the first sink particle was formed (see Table 2). In *red*, we show the gas temperature, and in *blue* the dust temperature. The *dashed lines* are lines of constant Jeans mass



depending on the metallicity, and this difference is seen for distances smaller than ~ 400 AU. The $Z = 10^{-4} Z_{\odot}$ case, for instance, has twice the number of M_{BE} for distances smaller than ~ 10 AU, when compared to the other cases. This will have direct consequences for the fragment mass function as we will see in the following section (Fig. 3).

Table 2 Sink particle properties for the different metallicities at the point where $4.7M_{\odot}$ have been accreted by the sink particles. “ST” (start time) is the time when sink particles start to form. “FT” (formation time), is the time taken to accrete $4.7M_{\odot}$ in the sinks. “SFR” is the mean star formation rate. Mean and median refer to the final mean and median sink mass. Finally, “N” is the number of sink particles formed

Z/Z_{\odot}	ST (10^3 year)	FT (year)	SFR (M_{\odot}/year)	Mean (M_{\odot})	Median (M_{\odot})	N
0	171.6	73	0.064	0.24	0.12	19
10^{-6}	171.2	72	0.065	0.29	0.06	16
10^{-5}	170.8	88	0.053	0.24	0.11	19
10^{-4}	169.2	138	0.034	0.10	0.05	45




Fig. 2 Number density maps for a slice through the high density region for $Z = 10^{-4} Z_{\odot}$ (top), $10^{-5} Z_{\odot}$, $10^{-6} Z_{\odot}$, and 0 (bottom). The image shows a sequence of zooms in on the density structure in the gas immediately before the formation of the first protostar

5.3 Properties of the Fragments

The simulations were stopped at a point when $4.7M_{\odot}$ of gas has been accreted into the sink particles, because the high computational cost made it impractical to continue. Figure 4 shows the mass distribution of sink particles at that time. We typically find sink masses below $1M_{\odot}$, with somewhat smaller values in the


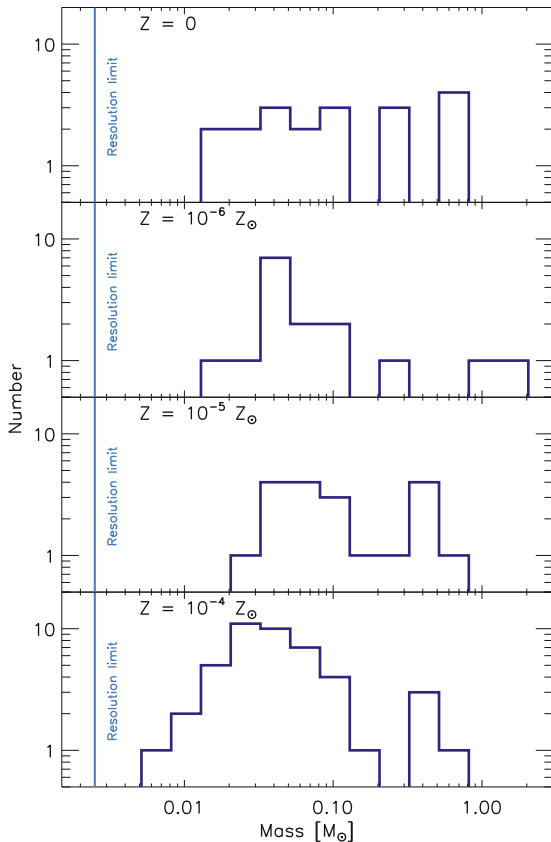


Fig. 3 Number density map showing a slice through the densest clump, and the star formation efficiency (SFE) for $Z = 10^{-4} Z_{\odot}$ (bottom), $10^{-5} Z_{\odot}$, $10^{-6} Z_{\odot}$, and 0 (top). The box is $100 \text{ AU} \times 100 \text{ AU}$ and the percentage indicates the star formation efficiency, i.e. the total mass in the sinks divided by the cloud mass ($1,000 M_{\odot}$)

$10^{-4} Z_{\odot}$ case compared to the other cases. No sharp transition in fragmentation behavior was found, but rather a smooth and complex interaction between kinematic and thermodynamic properties of the cloud.

Table 2 lists the main sink particle properties. It shows that the time taken to form the first sink particle is slightly shorter for higher metallicities. This shorter time is

Fig. 4 Sink particle mass function at the point when $4.7 M_{\odot}$ of gas had been accreted by the sink particles. The mass resolution of the simulations is indicated by the vertical line



a consequence of the more efficient cooling by dust, which decreases the thermal energy that was delaying the gravitational collapse. In Table 2 we also observe that the star formation rate is lower for $Z = 10^{-4} Z_{\odot}$. This is because star formation started at an earlier stage of the collapse, when the mean density of the cloud was lower and there was less dense gas available to form stars.

To better understand whether the resulting stellar cluster was affected by varying the metallicity, we plot the final sink mass distribution in Fig. 4. It shows that for the simulations with $Z \leq 10^{-5} Z_{\odot}$, the resulting sink particle mass function is relatively flat. There are roughly equal numbers of low-mass and high-mass stars, implying that most of the mass is to be found in the high-mass objects. This mass function is consistent with those found in other recent studies of fragmentation in metal-free gas [18, 36]. If the sink particle mass function provides a reliable guide to the form of the final stellar IMF, it suggests that at these metallicities, the IMF will be dominated by high-mass stars.

All of the histograms in Fig. 4 have the lowest sink particle mass well above the resolution limit of $0.0025 M_{\odot}$. Note that in all cases, we are still looking at the very

early stages of star cluster evolution. As a consequence, the sink particle masses in Fig. 4 are not the same as the final protostellar masses – there are many mechanisms that will affect the mass function, such as continuing accretion, mergers between the newly formed protostars, feedback from winds, jets and luminosity accretion, etc.

6 Conclusions

In this paper we have addressed the question of whether dust cooling can lead to the fragmentation of low-metallicity star-forming clouds. For this purpose we performed numerical simulations that follow the thermodynamical and chemical evolution of collapsing clouds. The chemical model included a primordial chemical network together with a description of dust effects, where the dust temperature was calculated by solving self-consistently the thermal energy equilibrium equation.

As a result, we found that dust can cool the gas, for number densities higher than 10^{11} , 10^{12} , and $3 \times 10^{13} \text{ cm}^{-3}$ for $Z = 10^{-4}$, 10^{-5} , and $10^{-6} Z_{\odot}$, respectively. Higher metallicity implies larger dust-to-gas fraction, and consequently stronger cooling. This is reflected in a lower temperature of the dense gas for the higher metallicity simulations, and this colder gas permitted a faster collapse. Therefore, the fragmentation behavior of the gas depends on the metallicity, and higher metallicities lead to a faster collapse.

For example, the characteristic fragment mass was lower for $Z = 10^{-4} Z_{\odot}$, since a lower temperature reduces the Bonnor-Ebert masses at the point where the gas undergo fragmentation. This also implies a lower ratio of fragmentation and accretion time, $t_{\text{frag}}/t_{\text{acc}}$, which will lead to a mass function dominated by low-mass objects. For $Z \leq 10^{-5} Z_{\odot}$, fragmentation and accretion timescales are comparable, and the resulting mass spectrum is rather flat, with roughly equal numbers of stars in each mass bin.

In addition to that, dust cooling appears to be insufficient to change the stellar mass distribution for the $Z = 10^{-5}$ and $10^{-6} Z_{\odot}$ cases, when compared with the metal-free case. This can be seen in the sink particle mass function (Fig. 4), which shows that the $Z \leq 10^{-5} Z_{\odot}$ cases do not appear to be fundamentally different.

Finally, we conclude that the dust is not an efficient coolant at metallicities below or equal to $Z_{\text{crit}} = 10^{-5} Z_{\odot}$, in the sense that it can not change the fragmentation behavior for these metallicities. Our results support the idea that low mass fragments can form in the absence of metals, and clouds with $Z \lesssim Z_{\text{crit}}$ will form a cluster with a flat IMF.

Acknowledgements We would like to acknowledge the support from the High Performance Computing Center Stuttgart (HLRS) of the University of Stuttgart. The present work is supported by the *Baden-Württemberg Stiftung* in the program *Internationale Spitzenforschung* via contract P-LS-SPII/18, the German *Bundesministerium für Bildung und Forschung* via the ASTRONET project STAR FORMAT (grant 05A09VHA), a Frontier grant of Heidelberg University sponsored by the German Excellence Initiative, the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). All computations described

here were performed at the *Leibniz-Rechenzentrum*, National Supercomputer HLRB-II (*Bayerische Akademie der Wissenschaften*), Jülich Supercomputing Centre – Jülich Research on Petaflop Architectures, and on the HPC-GPU Cluster Kolob (University of Heidelberg).

References

1. Abel, T., Bryan, G. L., & Norman, M. L. 2002, *Science*, 295, 93
2. Banerjee, R., Pudritz, R. E., & Anderson, D. W. 2006, *MNRAS*, 373, 1091
3. Bate, M. R., Bonnell, I. A., & Price, N. M. 1995, *MNRAS*, 277, 362
4. Bate, M. R. & Burkert, A. 1997, *MNRAS*, 288, 1060
5. Bonnor, W. B. 1956, *MNRAS*, 116, 351
6. Bromm, V., Coppi, P. S., & Larson, R. B. 2002, *ApJ*, 564, 23
7. Bromm, V., Ferrara, A., Coppi, P. S., & Larson, R. B. 2001, *MNRAS*, 328, 969
8. Bromm, V. & Loeb, A. 2003, *Nature*, 425, 812
9. Chabrier, G. 2003, *PASP*, 115, 763
10. Clark, P. C., Glover, S. C. O., & Klessen, R. S. 2008, *ApJ*, 672, 757
11. Clark, P. C., Glover, S. C. O., Klessen, R. S., & Bromm, V. 2011a, *ApJ*, 727, 110
12. Clark, P. C., Glover, S. C. O., Smith, R. J., et al. 2011b, *Science*, 331, 1040
13. Dopcke, G., Glover, S. C. O., Clark, P. C., & Klessen, R. S. 2011, *ApJ*, 729, L3
14. Ebert, R. 1955, *ZAp*, 37, 217
15. Frebel, A., Johnson, J. L., & Bromm, V. 2007, *MNRAS*, 380, L40
16. Goldsmith, P. F. 2001, *ApJ*, 557, 736
17. Greif, T. H., Bromm, V., Clark, P. C., et al. 2012, *ArXiv e-prints*
18. Greif, T. H., Springel, V., White, S. D. M., et al. 2011, *ApJ*, 737, 75
19. Hollenbach, D. & McKee, C. F. 1979, *ApJS*, 41, 555
20. Hollenbach, D. & McKee, C. F. 1989, *ApJ*, 342, 306
21. Jappsen, A.-K., Klessen, R. S., Glover, S. C. O., & Mac Low, M.-M. 2009a, *ApJ*, 696, 1065
22. Jappsen, A.-K., Klessen, R. S., Larson, R. B., Li, Y., & Mac Low, M.-M. 2005, *A&A*, 435, 611
23. Jappsen, A.-K., Mac Low, M.-M., Glover, S. C. O., Klessen, R. S., & Kitsionas, S. 2009b, *ApJ*, 694, 1161
24. Kroupa, P. 2002, *Science*, 295, 82
25. Omukai, K. 2000, *ApJ*, 534, 809
26. Omukai, K., Hosokawa, T., & Yoshida, N. 2010, *ApJ*, 722, 1793
27. Omukai, K., Tsuribe, T., Schneider, R., & Ferrara, A. 2005, *ApJ*, 626, 627
28. O’Shea, B. W. & Norman, M. L. 2007, *ApJ*, 654, 66
29. Santoro, F. & Shull, J. M. 2006, *ApJ*, 643, 26
30. Schneider, R., Ferrara, A., Natarajan, P., & Omukai, K. 2002, *ApJ*, 571, 30
31. Schneider, R. & Omukai, K. 2010, *MNRAS*, 402, 429
32. Schneider, R., Omukai, K., Bianchi, S., & Valiante, R. 2012, *MNRAS*, 419, 1566
33. Schneider, R., Omukai, K., Inoue, A. K., & Ferrara, A. 2006, *MNRAS*, 369, 1437
34. Smith, B. D. & Sigurdsson, S. 2007, *ApJ*, 661, L5
35. Smith, B. D., Turk, M. J., Sigurdsson, S., O’Shea, B. W., & Norman, M. L. 2009, *ApJ*, 691, 441
36. Smith, R. J., Glover, S. C. O., Clark, P. C., Greif, T., & Klessen, R. S. 2011, *MNRAS*, 414, 3633
37. Springel, V. 2005, *MNRAS*, 364, 1105
38. Stacy, A., Greif, T. H., & Bromm, V. 2012, *MNRAS*, 2508
39. Tsuribe, T. & Omukai, K. 2006, *ApJ*, 642, L61
40. Tsuribe, T. & Omukai, K. 2008, *ApJ*, 676, L45
41. Yoshida, N., Omukai, K., & Hernquist, L. 2008, *Science*, 321, 669

The SuperN-Project: Porting and Optimizing VERTEX-PROMETHEUS on the Cray XE6 at HLRS for Three-Dimensional Simulations of Core-Collapse Supernova Explosions of Massive Stars

F. Hanke, A. Marek, B. Müller, and H.-Th. Janka

Abstract Supernova explosions are among the most powerful cosmic events, whose physical mechanism and consequences are still incompletely understood. We have developed a fully MPI-OpenMP parallelized version of our VERTEX-PROMETHEUS code in order to perform three-dimensional simulations of stellar core-collapse and explosion on Tier-0 systems such as Hermit at HLRS. Tests on up to 64,000 cores have shown excellent scaling behavior. In this report we present the system of equations and the algorithm for its solution that are employed in our code VERTEX-PROMETHEUS. We also discuss the parallelization of VERTEX-PROMETHEUS and present our progress in porting, optimizing, and performing production runs on a large variety of machines, starting from vector machines and reaching to modern systems. In particular the results of our efforts to achieve good parallel scaling on the new Cray XE6 at HLRS Stuttgart are highlighted.

1 Introduction

A star more massive than about eight solar masses ends its life in a catastrophic explosion, a supernova. Its quiescent evolution comes to an end, when the pressure in its inner layers is no longer able to balance the inward pull of gravity. Throughout its life, the star sustained this balance by generating energy through a sequence of nuclear fusion reactions, forming increasingly heavier elements in its core. However, when the core consists mainly of iron-group nuclei, central energy generation ceases. The fusion reactions producing iron-group nuclei relocate to the core's surface, and their "ashes" continuously increase the core's mass. Similar to a white dwarf, such a core is stabilised against gravity by the pressure of its

F. Hanke (✉) · A. Marek · B. Müller · H.-Th. Janka
Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Strasse 1, Postfach 1317, D-85741
Garching bei München, Germany
e-mail: fhanke@mpa-garching.mpg.de

degenerate gas of electrons. However, to remain stable, its mass must stay smaller than (roughly) the Chandrasekhar limit. When the core grows larger than this limit, it collapses to a neutron star, and a huge amount ($\sim 10^{53}$ erg) of gravitational binding energy is set free. Most ($\sim 99\%$) of this energy is radiated away in neutrinos, but a small fraction is transferred to the outer stellar layers and drives the violent mass ejection, which disrupts the star in a supernova.

Despite 40 years of research, the details of how this energy transfer happens and how the explosion is initiated are still not well understood. Observational evidence about the physical processes deep inside the collapsing star is sparse and almost exclusively indirect. The only direct observational access is via measurements of neutrinos or gravitational waves. To obtain insight into the events in the core, one must therefore heavily rely on sophisticated numerical simulations. The enormous amount of computer power required for this purpose has led to the use of several, often questionable, approximations and numerous ambiguous results in the past. Fortunately, however, the development of numerical tools and computational resources has meanwhile advanced to a point, where it is becoming possible to perform multi-dimensional simulations with unprecedented accuracy. Therefore there is hope that the physical processes which are essential for the explosion can finally be unravelled.

An understanding of the explosion mechanism is required to answer many important questions of nuclear, gravitational, and astro-physics like the following:

- How do the explosion energy, the explosion timescale, and the mass of the compact remnant depend on the progenitor's mass? Is the explosion mechanism the same for all progenitors? For which stars are black holes left behind as compact remnants instead of neutron stars?
- What is the role of the – incompletely known – equation of state (EoS) of the proto-neutron star? Do softer or stiffer EoSs favour the explosion of a core collapse supernova?
- How do neutron stars receive their natal kicks? Are they accelerated by asymmetric mass ejection and/or anisotropic neutrino emission?
- What are the generic properties of the neutrino emission and of the gravitational wave signal that are produced during stellar core collapse and explosion? Up to which distances could these signals be measured with operating or planned detectors on earth and in space? And what can one learn about supernova dynamics or nuclear and particle physics from a future measurement of such signals in the case of a Galactic supernova?
- How do supernovae contribute to the enrichment of the intergalactic medium with heavy elements? What kind of nucleosynthesis processes occur during and after the explosion? Can the elemental composition of supernova remnants be explained correctly by the numerical simulations? Does the rapid neutron capture process (r-process), which produces e.g. gold and the actinides, take place in supernovae?

2 Numerical Modeling

2.1 *History and Constraints*

According to theory, a shock wave is launched at the moment of “core bounce” when the neutron star begins to emerge from the collapsing stellar iron core. There is general agreement, supported by all “modern” numerical simulations, that this shock is unable to propagate directly into the stellar mantle and envelope, because it loses too much energy in dissociating iron into free nucleons while it moves through the outer core. The “prompt” shock ultimately stalls. Thus the currently favoured theoretical paradigm exploits the fact that a huge energy reservoir is present in the form of neutrinos, which are abundantly emitted from the hot, nascent neutron star. The absorption of electron neutrinos and anti-neutrinos by free nucleons in the post-shock layer is thought to reenergize the shock, thus triggering the supernova explosion.

Detailed *spherically symmetric* hydrodynamic models, which recently include a very accurate treatment of the time-dependent, multi-flavour, multi-frequency neutrino transport based on a numerical solution of the Boltzmann transport equation [1, 2], reveal that this “delayed, neutrino-driven mechanism” does not work as simply as originally envisioned. Although in principle able to trigger the explosion (e.g., [3–5]), neutrino energy transfer to the post-shock matter turned out to be too weak. For inverting the infall of the stellar core and initiating powerful mass ejection, an increase of the efficiency of neutrino energy deposition is needed.

A number of physical phenomena have been pointed out that can enhance neutrino energy deposition behind the stalled supernova shock. They are all linked to the fact that the real world is multi-dimensional instead of spherically symmetric (or one-dimensional; 1D) as assumed in the works cited above:

- (1) Convective instabilities in the neutrino-heated layer between the neutron star and the supernova shock develop to violent convective overturn [6]. This convective overturn is helpful for the explosion, mainly because (a) neutrino-heated matter rises and increases the pressure behind the shock, thus pushing the shock further out, (b) cool matter is able to penetrate closer to the neutron star where it can absorb neutrino energy more efficiently, and (c) the rise of freshly heated matter reduces energy losses by the reemission of neutrinos. These effects allow multi-dimensional models to explode easier than spherically symmetric ones [7–9].
- (2) Recent work [10–13] has demonstrated that the stalled supernova shock is also subject to a second non-radial low-mode instability, called the standing accretion shock instability or “SASI” for short, which can grow to a dipolar, global deformation of the shock [12, 14, 15].

- (3) Convective energy transport inside the nascent neutron star [16–18] might enhance the energy transport to the neutrinosphere and could thus boost the neutrino luminosities. This would in turn increase the neutrino-heating behind the shock.

This list of multi-dimensional phenomena (limited to non-magnetized supernova cores) awaits more detailed exploration by multi-dimensional simulations. Until recently, such simulations have been performed with only a grossly simplified treatment of the involved microphysics, in particular of the neutrino transport and neutrino-matter interactions. At best, grey (i.e., single energy) flux-limited diffusion schemes were employed. Since, however, the role of the neutrinos is crucial for the problem, and because previous experience shows that the outcome of simulations is indeed very sensitive to the employed transport approximations, studies of the explosion mechanism require the best available description of the neutrino physics. This implies that one has to solve the Boltzmann transport equation for neutrinos.

2.2 *The Mathematical Model*

As core-collapse supernovae involve such a complex interplay of hydrodynamics, self-gravity and neutrino heating and cooling, numerical modellers face a classical “multiphysics” problem. Although the overall problem can still be formulated as a system of non-linear partial differential equations, rather dissimilar methods – sometimes with conflicting requirements on the computer architecture and the parallelization strategy – need to be applied to treat individual subsystems. In the case of our code, the system of equations that needs to be solved consists of the following components:

- The multi-dimensional Euler equations of (relativistic) hydrodynamics, supplemented by advection equations for the electron fraction and the chemical composition of the fluid, and formulated in spherical polar coordinates;
- Equations for the space-time metric (or in the Newtonian case, the Poisson equation) for calculating the gravitational source terms in the Euler equations;
- The Boltzmann transport equation and/or its moment equations which determine the (non-equilibrium) distribution function of the neutrinos;
- The emission, absorption, and scattering rates of neutrinos, which are required for the solution of the neutrino transport equations;
- The equation of state of the stellar fluid, which provides the closure relation between the variables entering the Euler equations, i.e. density, momentum, energy, electron fraction, composition, and pressure.

In what follows we will briefly summarise the neutrino transport algorithms, thus focusing on the major computational kernel of our code. For a more complete description of the entire code we refer the reader to [19, 20], and the references therein.




Fig. 1 Illustration of the phase space coordinates (see the main text)

2.3 “Ray-by-Ray Plus” Method for the Neutrino Transport Problem

The crucial quantity required to determine the source terms for the energy, momentum, and electron fraction of the fluid owing to its interaction with the neutrinos is the neutrino distribution function in phase space, $f(r, \vartheta, \phi, \epsilon, \Theta, \Phi, t)$. Equivalently, the neutrino intensity $I = c/(2\pi\hbar c)^3 \cdot \epsilon^3 f$ may be used. Both are time-dependent functions in a six-dimensional phase space, as they describe, at every point in space (r, ϑ, ϕ) , the distribution of neutrinos propagating with energy ϵ into the direction (Θ, Φ) at time t (Fig. 1).

The evolution of I (or f) in time is governed by the Boltzmann equation, and solving this equation is, in general, a six-dimensional problem (as time is usually not counted as a separate dimension). A solution of this equation by direct discretization (using an S_N scheme) would require computational resources in the PetaFlop range. Although there are attempts by at least one group in the United States to follow such an approach, we feel that, with the currently available computational resources, it is mandatory to reduce the dimensionality of the problem.

Actually this should be possible, since the source terms entering the hydrodynamic equations are *integrals* of I over momentum space (i.e. over $\epsilon, \Theta,$ and Φ), and thus only a fraction of the information contained in I is truly required to compute the neutrino effects on the dynamics of the flow. It therefore makes sense to consider angular moments of I , and to solve evolution equations for these moments, instead of dealing with the Boltzmann equation directly. The 0th to 3rd order moments are defined as

$$\mathbf{J}, \mathbf{H}, \mathbf{K}, \mathbf{L}, \dots (r, \vartheta, \phi, \epsilon, t) = \frac{1}{4\pi} \int I(r, \vartheta, \phi, \epsilon, \Theta, \Phi, t) \mathbf{n}^{0,1,2,3,\dots} d\Omega \quad (1)$$

where $d\Omega = \sin \Theta d\Theta d\Phi$, $\mathbf{n} = (\cos \Theta, \sin \Theta \cos \Phi, \sin \Theta \sin \Phi)$, and exponentiation represents repeated application of the dyadic product. Note that the moments are *tensors* of the required rank.

So far no approximations have been made. In order to reduce the size of the problem even further, one needs to resort to assumptions on its symmetry. At this point, one assumes that I is independent of Θ and Φ , then each of the angular moments of I becomes a *scalar*, which depends on three spatial dimensions, and one dimension in momentum space: $J, H, K, L = J, H, K, L(r, \vartheta, \phi, \epsilon, t)$. Thus the neutrino moment equations at different angular directions (except for some terms which can be accounted for explicitly in an operator split) decouple from each other. Therefore, for each “radial ray”, i.e. for all zones of same angle, the moment equations can be solved independently. Except for some additional terms, this problem is identical to solving $N_\theta \times N_\phi$ times the moment equations for a spherically symmetric star with $N_\theta \times N_\phi$ being the number of grid zones in polar direction. As we will explain later, the great advantage of our “ray-by-ray” neutrino transport is the easy way to obtain perfect scaling behaviour to a large number of cores.

The System of Equations

With the aforementioned assumptions it can be shown [19], that in the Newtonian approximation the following two transport equations need to be solved in order to compute the source terms for the energy and electron fraction of the fluid:

$$\begin{aligned}
 & \left(\frac{1}{c} \frac{\partial}{\partial t} + \beta_r \frac{\partial}{\partial r} + \frac{\beta_\vartheta}{r} \frac{\partial}{\partial \vartheta} + \frac{\beta_\phi}{r \sin \vartheta} \frac{\partial}{\partial \phi} \right) J \\
 & + J \left(\frac{1}{r^2} \frac{\partial(r^2 \beta_r)}{\partial r} + \frac{1}{r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{r \sin \vartheta} \frac{\partial \beta_\phi}{\partial \phi} \right) \\
 & + \frac{1}{r^2} \frac{\partial(r^2 H)}{\partial r} + \frac{\beta_r}{c} \frac{\partial H}{\partial t} - \frac{\partial}{\partial \epsilon} \left\{ \frac{\epsilon}{c} \frac{\partial \beta_r}{\partial t} H \right\} \\
 & - \frac{\partial}{\partial \epsilon} \left\{ \epsilon J \left(\frac{\beta_r}{r} + \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\phi}{\partial \phi} \right) \right\} \\
 & - \frac{\partial}{\partial \epsilon} \left\{ \epsilon K \left(\frac{\partial \beta_r}{\partial r} - \frac{\beta_r}{r} - \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} - \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\phi}{\partial \phi} \right) \right\} \\
 & + J \left(\frac{\beta_r}{r} + \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} + \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\phi}{\partial \phi} \right) \\
 & + K \left(\frac{\partial \beta_r}{\partial r} - \frac{\beta_r}{r} - \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta)}{\partial \vartheta} - \frac{1}{2r \sin \vartheta} \frac{\partial \beta_\phi}{\partial \phi} \right) \\
 & + \frac{2}{c} \frac{\partial \beta_r}{\partial t} H = C^{(0)}, \quad (2)
 \end{aligned}$$

$$\begin{aligned}
 & \left(\frac{1}{c} \frac{\partial}{\partial t} + \beta_r \frac{\partial}{\partial r} + \frac{\beta_{\vartheta}}{r} \frac{\partial}{\partial \vartheta} + \frac{\beta_{\varphi}}{r} \frac{\partial}{\partial \varphi} \right) H \\
 & + H \left(\frac{1}{r^2} \frac{\partial(r^2 \beta_r)}{\partial r} + \frac{1}{r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_{\vartheta})}{\partial \vartheta} + \frac{1}{r \sin \vartheta} \frac{\partial \beta_{\varphi}}{\partial \varphi} \right) \\
 & + \frac{\partial K}{\partial r} + \frac{3K - J}{r} + H \left(\frac{\partial \beta_r}{\partial r} \right) + \frac{\beta_r}{c} \frac{\partial K}{\partial t} - \frac{\partial}{\partial \epsilon} \left\{ \frac{\epsilon}{c} \frac{\partial \beta_r}{\partial t} K \right\} \\
 & - \frac{\partial}{\partial \epsilon} \left\{ \epsilon L \left(\frac{\partial \beta_r}{\partial r} - \frac{\beta_r}{r} - \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_{\vartheta})}{\partial \vartheta} - \frac{1}{2r \sin \vartheta} \frac{\partial \beta_{\varphi}}{\partial \varphi} \right) \right\} \\
 & - \frac{\partial}{\partial \epsilon} \left\{ \epsilon H \left(\frac{\beta_r}{r} + \frac{1}{2r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_{\vartheta})}{\partial \vartheta} + \frac{1}{2r \sin \vartheta} \frac{\partial \beta_{\varphi}}{\partial \varphi} \right) \right\} \\
 & + \frac{1}{c} \frac{\partial \beta_r}{\partial t} (J + K) = C^{(1)}. \quad (3)
 \end{aligned}$$

These are evolution equations for the neutrino energy density, J , and the neutrino flux, H , and follow from the zeroth and first moment equations of the comoving frame (Boltzmann) transport equation in the Newtonian, $\mathcal{O}(v/c)$ approximation. The quantities $C^{(0)}$ and $C^{(1)}$ are source terms that result from the collision term of the Boltzmann equation, while $\beta_r = v_r/c$, $\beta_{\vartheta} = v_{\vartheta}/c$, and $\beta_{\varphi} = v_{\varphi}/c$, where v_r , v_{ϑ} , and v_{φ} are the components of the hydrodynamic velocity, and c is the speed of light. The functional dependencies $\beta_r = \beta_r(r, \vartheta, \varphi, t)$, $J = J(r, \vartheta, \varphi, \epsilon, t)$, etc. are suppressed in the notation. This system includes four unknown moments (J, H, K, L) but only two equations, and thus needs to be supplemented by two more relations. This is done by substituting $K = f_K \cdot J$ and $L = f_L \cdot J$, where f_K and f_L are the variable Eddington factors, which for the moment may be regarded as being known, but in our case are indeed determined from a separate simplified (“model”) Boltzmann equation.

The moment equations (2) and (3) are very similar to the $\mathcal{O}(v/c)$ equations in spherical symmetry which were solved in the 1D simulations of [21] (see Eqs. (7),(8),(30), and (31) of the latter work). This similarity has allowed us to reuse a good fraction of the one-dimensional version of the transport part, for coding the multi-dimensional algorithm. The additional terms necessary for this purpose have been set in boldface above.

Finally, the changes of the energy, e , and electron fraction, Y_e , required for the hydrodynamics are given by the following two equations

$$\frac{de}{dt} = -\frac{4\pi}{\rho} \int_0^\infty d\epsilon \sum_{\nu \in \{\nu_e, \bar{\nu}_e, \dots\}} C_\nu^{(0)}(\epsilon), \quad (4)$$

$$\frac{dY_e}{dt} = -\frac{4\pi m_B}{\rho} \int_0^\infty \frac{d\epsilon}{\epsilon} \left(C_{\nu_e}^{(0)}(\epsilon) - C_{\bar{\nu}_e}^{(0)}(\epsilon) \right) \quad (5)$$

(for the momentum source terms due to neutrinos see [19]). Here m_B is the baryon mass, and the sum in Eq. (4) runs over all neutrino types. The full system consisting of Eqs. (2–5) is stiff, and thus requires an appropriate discretization scheme for its stable solution.

2.3.1 Method of Solution

In order to discretize Eqs. (2–5), the spatial domain $[0, r_{\max}] \times [\vartheta_{\min}, \vartheta_{\max}] \times [\varphi_{\min}, \varphi_{\max}]$ is covered by N_r radial, N_ϑ latitudinal, and N_φ longitudinal zones, where $\vartheta_{\min} = 0$ and $\vartheta_{\max} = \pi$ correspond to the north and south poles, respectively, of the spherical grid and $\varphi_{\min} = 0$ and $\varphi_{\max} = 2\pi$ covers the full sphere. (In general, we allow for grids with different radial resolutions in the neutrino transport and hydrodynamic parts of the code. The number of radial zones for the hydrodynamics will be denoted by N_r^{hyd} .) The number of bins used in energy space is N_ϵ and the number of neutrino types taken into account is N_ν .

The equations are solved in three operator-split steps corresponding to a lateral, an azimuthal and a radial sweep.

In the first two steps, we treat the boldface terms in the respectively first lines of Eqs. (2–3), which describe the lateral and azimuthal advection of the neutrinos with the stellar fluid, and thus couple the angular moments of the neutrino distribution of neighbouring angular zones. For this purpose we consider the equations

$$\frac{1}{c} \frac{\partial \mathcal{E}}{\partial t} + \frac{1}{r \sin \vartheta} \frac{\partial(\sin \vartheta \beta_\vartheta \mathcal{E})}{\partial \vartheta} = 0, \quad (6)$$

$$\frac{1}{c} \frac{\partial \mathcal{E}}{\partial t} + \frac{1}{r \sin \vartheta} \frac{\partial(\beta_\varphi \mathcal{E})}{\partial \varphi} = 0, \quad (7)$$

where \mathcal{E} represents one of the moments J or H . Although it has been suppressed in the above notation, an equation of this form has to be solved for each radius, for each energy bin, and for each type of neutrino. An explicit upwind scheme is used for this purpose.

In the third step, the radial sweep is performed. Several points need to be noted here:

- Terms in boldface not yet taken into account in the lateral sweep, need to be included into the discretization scheme of the radial sweep. This can be done in a straightforward way since these remaining terms do not include derivatives of the transport variables J or H . They only depend on the hydrodynamic velocities v_ϑ and v_φ , which are a *constant* scalar field for the transport problem.
- The right hand sides (source terms) of the equations and the coupling in energy space have to be accounted for. The coupling in energy is non-local, since the source terms of Eqs. (2) and (3) stem from the Boltzmann equation, which is an integro-differential equation and couples all the energy bins.

- The discretization scheme for the radial sweep is *implicit* in time. Explicit schemes would require very small time steps to cope with the stiffness of the source terms in the optically thick regime, and the small CFL time step dictated by neutrino propagation with the speed of light in the optically thin regime. Still, even with an implicit scheme $\gtrsim 10^5$ time steps are required per simulation. This makes the calculations expensive.

Once the equations for the radial sweep have been discretized in radius and energy, the resulting solver is applied ray-by-ray for each pair of angles (ϑ, φ) and for each type of neutrino; i.e. for constant (ϑ, φ) , N_ν two-dimensional problems need to be solved.

The discretization itself is done using a second order accurate scheme with backward differencing in time according to [21]. This leads to a non-linear system of algebraic equations, which is solved by Newton-Raphson iteration with explicit construction and inversion of the corresponding Jacobian matrix with the Block-Thomas algorithm.

3 Porting and Scaling on the Cray XE6 “HERMIT” at HLRS

3.1 Parallelization Strategy

The ray-by-ray approximation readily lends itself to parallelization over the different angular zones. In order to make efficient use of modern supercomputer systems with relatively small shared-memory units (e.g. 16 CPUs per node on Cray XE6), distributed memory parallelism is indispensable. An MPI version of the VERTEX-PROMETHEUS code using domain decomposition was initially developed within a cooperation between MPA and the Teraflop Workbench at the HLRS in 2007/2008. Since then, the parallelization of VERTEX-PROMETHEUS has been further extended to allow good scaling on several thousands of cores as required for future 3D supernova simulations.

The VERTEX-PROMETHEUS code employs a hybrid MPI-OpenMP parallelization scheme, in which the parallelization of the transport module – the main computational kernel and most CPU-intense part of the code – is along radial “rays” for fixed angular bins of the three-dimensional grid. Hence, every “ray” of the transport is treated by one core using as many OpenMP threads as cores available on an individual node. This strategy allows almost perfect scaling behavior, since almost no MPI communication is necessary between individual rays during the transport step.

The MPI-parallelization of the much less expensive hydrodynamical part PROMETHEUS is based on standard domain decomposition methods. Hereby, the reconstruction scheme used to solve the hydrodynamic equations requires so-called




Fig. 2 Schematic sketch of the MPI communication pattern for an angular direction in the hydrodynamics part of our code. The *red rectangle* symbolizes the data available in a specific MPI task, the surrounding area marked by the *dashed line* reflects the zones to be communicated

“ghost-zones”, which have to be available in each MPI task. In our case, four ghost zones are required on each cell interface in angular directions to integrate one time step and these zones have to be MPI communicated to the neighbouring MPI tasks. A sketch of grid zones to be MPI communicated is illustrated in Fig. 2.

3.2 *Porting VERTEX-PROMETHEUS to the Cray XE6 “HERMIT”*

As demonstrated in Fig. 3, we have already obtained excellent scaling behavior with the explained parallelization strategy. For example, we have performed scaling tests on the BlueGene/P system JUGENE at the Forschungszentrum Jülich to demonstrate that our VERTEX-PROMETHEUS code scales perfectly up to 65,000 cores.

Since our VERTEX-PROMETHEUS code runs successfully on several architectures, the code should in principle work out of the box. However, we had to change several smaller statements in order to be able to compile the code. Furthermore, while performing the first scaling test on the Cray of the HLRS we detected that the routine, which calculates the most important neutrino interaction rates, shows poor performance. Initially, we have used the same version of this part of the transport solver, which performs perfectly using the Intel compiler. To obtain better results on the Cray XE6 we have rewritten this routine and we use now a vectorized version with one main loop.

Employing this single optimization, the code scales well on up to 32,000 cores of the Cray XE6 at HLRS as shown in Fig. 3. However, the scaling behavior is still slightly worse than on Intel platforms. We plan to analyse the detailed code performance on the Cray XE6 further to get better results of the scaling tests.

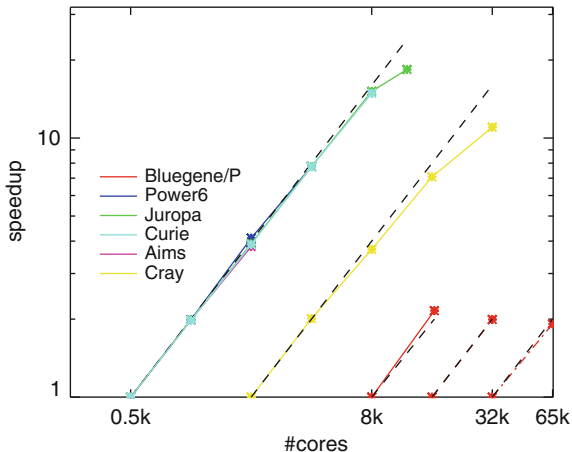


Fig. 3 Strong scaling of VERTEX-PROMETHEUS on different machines and architectures. The *colored lines* show the speedup on the respective machine relative to the run with the smallest number of cores for a given problem size. The symbols mark the number of cores on which the timings were done. The *dashed black lines* indicate a theoretically perfect scaling behavior. Note that several lines lie on top of each other. Also note that due to limited memory on BlueGene/P three different setups were timed, which are shown as three separate speedup curves

Another point concerning the special characteristics of the Cray XE6 is the strong interconnection of the individual nodes. We cannot profit a lot by this feature since our code needs only a low amount of communication (less than 5% of the total computing time).

Furthermore, we want to improve the performance of I/O on the Cray XE6. The I/O is now handled by means of parallel HDF5 to ensure high scalability and to eliminate the excessive memory consumption associated with temporary I/O arrays on the root node. The handling of I/O performs quite well on IBM BlueGene and Intel systems, however we want to optimize I/O on the Cray XE6 further.

4 Conclusion

We have presented our main simulation tool VERTEX-PROMETHEUS. In the past years, we have developed a fully MPI/OpenMP parallelized code version to be able to perform large scale runs on several thousand cores. At the moment our code shows excellent scaling behavior on several platforms. After the new Cray XE6 “HERMIT” had become available at HLRS, we have ported VERTEX-PROMETHEUS to this new system. With minor optimizations (required by the compiler) the code scales now up to 32,000 cores.

Since our code is now ready to run on the new Cray XE6 at HRLS, we are ready to start the first generation of three-dimensional simulations of core-collapse supernova explosions this year. These simulations are extremely expensive (several 10^{20} floating point operations) that we need to strongly rely on Tier-0 systems such as “HERMIT”. Only systems like the new Cray XE6 in Stuttgart give us the possibility to advance our understanding of the details of the explosion mechanism of core-collapse supernovae.

Acknowledgements We thank especially K. Benkert for her extremely valuable and fruitful work on the MPI version of VERTEX. Support by the Deutsche Forschungsgemeinschaft through the SFB/TR27 “Neutrinos and Beyond” and the SFB/TR7 “Gravitational Wave Astronomy”, and by the Cluster of Excellence EXC 153 “Origin and Structure of the Universe” (<http://www.universe-cluster.de>) are acknowledged, as well as computer time grants of the HLRs, NIC Jülich, and Rechenzentrum Garching are acknowledged.

References

1. Rampp, M., Janka, H.T.: Spherically Symmetric Simulation with Boltzmann Neutrino Transport of Core Collapse and Postbounce Evolution of a $15 M_{\odot}$ Star. *Astrophys. J.* **539** (2000) L33–L36
2. Liebendörfer, M., Mezzacappa, A., Thielemann, F., Messer, O.E., Hix, W.R., Bruenn, S.W.: Probing the gravitational well: No supernova explosion in spherical symmetry with general relativistic Boltzmann neutrino transport. *Phys. Rev. D* **63** (2001) 103004–+
3. Bethe, H.A.: Supernova mechanisms. *Reviews of Modern Physics* **62** (1990) 801–866
4. Burrows, A., Goshy, J.: A Theory of Supernova Explosions. *Astrophys. J.* **416** (1993) L75
5. Janka, H.T.: Conditions for shock revival by neutrino heating in core-collapse supernovae. *Astron. Astrophys.* **368** (2001) 527–560
6. Herant, M., Benz, W., Colgate, S.: Postcollapse hydrodynamics of SN 1987A - Two-dimensional simulations of the early evolution. *Astrophys. J.* **395** (1992) 642–653
7. Herant, M., Benz, W., Hix, W.R., Fryer, C.L., Colgate, S.A.: Inside the supernova: A powerful convective engine. *Astrophys. J.* **435** (1994) 339
8. Burrows, A., Hayes, J., Fryxell, B.A.: On the nature of core-collapse supernova explosions. *Astrophys. J.* **450** (1995) 830
9. Janka, H.T., Müller, E.: Neutrino heating, convection, and the mechanism of Type-II supernova explosions. *Astron. Astrophys.* **306** (1996) 167–+
10. Thompson, C.: Accretional Heating of Asymmetric Supernova Cores. *Astrophys. J.* **534** (2000) 915–933
11. Blondin, J.M., Mezzacappa, A., DeMarino, C.: Stability of Standing Accretion Shocks, with an Eye toward Core-Collapse Supernovae. *Astrophys. J.* **584** (2003) 971–980
12. Scheck, L., Plewa, T., Janka, H.T., Kifonidis, K., Müller, E.: Pulsar Recoil by Large-Scale Anisotropies in Supernova Explosions. *Phys. Rev. Letters* **92** (2004) 011103–+
13. Foglizzo, T., Galletti, P., Scheck, L., Janka, H.T.: Instability of a Stalled Accretion Shock: Evidence for the Advective-Acoustic Cycle. *Astrophys. J.* **654** (2007) 1006–1021
14. Scheck, L., Kifonidis, K., Janka, H.T., Müller, E.: Multidimensional supernova simulations with approximative neutrino transport. I. Neutron star kicks and the anisotropy of neutrino-driven explosions in two spatial dimensions. *Astron. Astrophys.* **457** (2006) 963–986
15. Scheck, L., Janka, H.T., Foglizzo, T., Kifonidis, K.: Multidimensional supernova simulations with approximative neutrino transport. II. Convection and the advective-acoustic cycle in the supernova core. *Astron. Astrophys.* **477** (2008) 931–952

16. Keil, W., Janka, H.T., Müller, E.: Ledoux Convection in Protoneutron Stars— A Clue to Supernova Nucleosynthesis? *Astrophys. J.* **473** (1996) L111
17. Burrows, A., Lattimer, J.M.: The birth of neutron stars. *Astrophys. J.* **307** (1986) 178–196
18. Pons, J.A., Reddy, S., Prakash, M., Lattimer, J.M., Miralles, J.A.: Evolution of Proto-Neutron Stars. *Astrophys. J.* **513** (1999) 780–804
19. Buras, R., Rampp, M., Janka, H.T., Kifonidis, K.: Two-dimensional hydrodynamic core-collapse supernova simulations with spectral neutrino transport. I. Numerical method and results for a $15M_{\odot}$ star. *Astron. Astrophys.* **447** (2006) 1049–1092
20. Müller, B., Janka, H., Dimmelmeier, H.: A New Multi-dimensional General Relativistic Neutrino Hydrodynamic Code for Core-collapse Supernovae. I. Method and Code Tests in Spherical Symmetry. *Astrophys. J. Suppl.* **189** (2010) 104–133
21. Rampp, M., Janka, H.T.: Radiation hydrodynamics with neutrinos. Variable Eddington factor method for core-collapse supernova simulations. *Astron. Astrophys.* **396** (2002) 361–392

Stability of the Strongest Magnetic Fields

Konstantinos D. Kokkotas, Burkhard Zink, and Paul Lasky

Abstract Neutron stars with the strongest magnetic fields, known as magnetars, contain fields up to nine orders of magnitude stronger than those produced in the laboratory. While the field exterior to the star is thought to be dominated by a roughly dipolar structure, the interior field is entirely unknown, and is currently a hotly debated topic in astrophysics since it is thought to be connected with huge gamma-ray outburst, the *giant flares*, and possibly also with gravitational radiation.

1 Scientific Background

No known object in the universe carries stronger magnetic fields than a magnetar, a neutron star endowed with a surface magnetic field of approximately 10^{15} G in strength [1]. Neutron stars in general represent the most extreme objects in the Universe, with masses 1.4 times that of the sun compressed in a region less than 30 km in diameter. The state of matter inside these stars is largely unknown, and efforts to predict their internal structure with the use of gravitational-wave detectors are the subject of current research.

Germany is very active in the field of gravitational-wave detection: Our group in Tübingen forms part of one of the prestigious transregional SFB (Sonderforschungsbereiche), which is focused on “Gravitational Wave Astronomy,” and it has recently entered its third period after a very positive review. We are part of these activities in the form of two projects which contribute modeling efforts for neutron stars and magnetars as possible gravitational wave sources.

K.D. Kokkotas (✉) · B. Zink
Eberhard-Karls University of Tübingen, Tübingen, Baden-Württemberg, Germany
e-mail: kostas.kokkotas@uni-tuebingen.de; bzink@tat.uni-tuebingen.de

P. Lasky
University of Melbourne, Parkville, VIC, Australia
e-mail: paul.lasky@unimelb.edu.au

Due to their ultra-strong magnetic fields, magnetars contain a substantial amount of electromagnetic field energy, which they release in both continuous radiation, frequent flares and storms, and sometimes in so-called “giant flares” with energies greater than 10^{44} erg [2]. Moreover, these events are so energetic that they are potential sources of gravitational radiation observable from Advanced LIGO, the upcoming upgrade of the LIGO gravitational wave detectors which have already been in operation for a few years [3]. In addition, the tails after the giant outburst also exhibit a set of normal mode oscillations, some of which are likely core plasma (Alfvén) waves excited by the outburst [4, 5].

Gravitational wave luminosities and post-flare modes are likely connected to the internal structure of the magnetic field inside the magnetar. However, despite their importance, this internal structure is entirely unknown. Simple analytical models with purely dipolar or toroidal field structure are commonly used for modeling purposes, but such models are known to be unstable [6, 9]. It is probable that actual core magnetic fields in neutron stars are much more complex. Since neutron stars are very compact objects, general relativistic magnetohydrodynamics (GRMHD) simulations are needed to investigate these strongest magnetic fields.

2 Report

The main goal of this project is to investigate the dynamical stability of magnetic fields inside highly magnetized neutron stars, and the associated gravitational wave signal possibly connected with giant flare events in magnetars. In our initial study of the subject [9] we have found that purely poloidal magnetic fields in general relativistic neutron star models are dynamically unstable to both low-order varicose (sausage-like) and kink instabilities, and we have followed this evolution over several hundred milliseconds as the instability is driven towards the saturation state. Following this process for over a hundred Alfvén times (and tens of thousand of sound crossing times) is unprecedented in this field, and competitive efforts to do so had to stop evolving after only 30 or 40 ms. This new capability is in part made possible due to the use of GPU computing, which allowed us to employ the massively parallel architectures to evolve for such long evolution times, as well as reduce experimentation turnover times.

The main result obtained in the first paper is that well-known axisymmetric equilibrium fields, which are often used as exact background configurations for two-dimensional simulations or perturbative studies, are in fact always unstable, and therefore these simplified models are generally very approximate. This is consistent with earlier studies on magnetic field stability in main-sequence stars [6]. We do, however, observe a somewhat different saturation state than what would be expected from simple “twisted-torus” configurations which have been discussed in the literature.

Such a twisted-torus state is a mixture of poloidal and toroidal fields, recognizing that both purely toroidal and purely poloidal fields are expected to be dynamically

unstable. In between, Newtonian MHD studies have presumably found stable configurations [7], so it is natural to expect that the result of the reconfiguration of magnetic fields starting from a dynamically unstable state may result in a configuration similar to a twisted torus, in particular since the poloidal field instability itself tends to produce toroidal components. In fact, we do find a saturation state with substantial toroidal components, but only *in parts* of that state, we also observe something akin to a twisted torus. Overall, the final state of our simulations is in fact highly nonaxisymmetric and still quite dynamic, so either the system needs much longer timescales to settle down, or we have found quasi-equilibria of a truly complex structure, which could actually represent a possible more realistic configuration for neutron star interiors in general. If confirmed in the future, this would constitute an intriguing result.

If a giant flare in a magnetar is indeed connected with a large rearrangement in the internal magnetic field, then this process could be an observable source of gravitational radiation. Magnetic fields inside magnetars could exceed 10^{15} or even 10^{16} G in amplitude, and a large-scale dynamical process would not only be a burst source (though possibly a weak one), but also excite various modes inside the star which could oscillate until some damping mechanism kicks in. The prospect of obtaining a direct signal from the magnetic field inside a neutron star is both intriguing from the perspective of understanding the giant flare mechanism better, as well as from the viewpoint of asteroseismology, since oscillation modes such excited could give us insights into the properties of the supernuclear matter in the core.

However, the most important question is: will we be able to actually *observe* a signal from these events, even if a violent rearrangement of a strong magnetic field happens as indicated? This question has been hotly debated in the literature, and in a recent follow-up paper [10] we have given the first answer for a wide range of possible field strengths based on general relativistic magnetohydrodynamics simulations.

In a series of simulations for the poloidal field instability, we have investigated the gravitational wave output over hundreds of milliseconds. The field strengths ranged from about 10^{15} G at the surface up to (unrealistically high) 10^{17} G. This wide range was chosen to measure the scaling behavior of gravitational wave strain and energy with the magnetic field strength, since even if the particulars of our model are not represented in a magnetar giant flare, we expect such scaling laws to be rather robust results. On the other hand, this allowed us to gain an order of magnitude estimate of the kind of signal we could expect from a giant flare under *optimistic* conditions. So far, estimates in the literature spanned a range of many orders of magnitude, so even a rough estimate has not been available until then.

The primary agent for gravitational-wave emission from neutron stars is the fundamental quadrupole mode ($\ell = 2$ *f*-mode) since it couples efficiently into gravitational radiation. If a large dynamical rearrangement of the magnetic field happens inside a star, the resulting fluid motions can have an overlap with this *f*-mode and thereby excite observable oscillations during the ringdown timescale (approximately 100 ms). Such a monotone signal (the frequency hardly shifts due to the presence of the magnetic field, dynamic or not) is an optimal target for

observation even under noisy conditions. However, current interferometric detectors are not very sensitive in the frequency range of the mode (kHz), which somewhat offsets its advantages. Still, it is the most obvious emission mechanism, and has primarily been discussed in the literature.

We have found that the catastrophic rearrangement of magnetic fields following a dynamical instability is in fact a rather weak motor for gravitational radiation, certainly well below some of the more optimistic results in the literature. In fact, assuming typical damping times, an f -mode excited from even a rather strong internal field will not yield an observable signal in present or currently planned detectors according to our simulations. Though disappointing, this is of course a significant insight when discussing potential targets for gravitational wave observations. Some hope comes from the fact that we have not covered the whole space of dynamical instabilities yet, but it is hard to conceive why say a toroidal unstable field of similar strength should produce a signal orders of magnitude stronger.

But there is another result we have obtained in that paper, which is far more promising in terms of observation: The magnetic field dynamics themselves give of course both rise to a burst signal (during the reconfiguration), as well as possibly later local or semi-global modes of oscillation when the field has mostly settled down. The Alfvén waves timescale gives rise to modes in the range of roughly ten to hundreds of Hertz, which is precisely in the range of optimal sensitivity of interferometric detectors, and incidentally also overlaps with the frequency range spanned by the quasi-periodic oscillations (QPOs) seen in the tail of giant flares.

In large part due to the use of GPU computing, we can now afford to simulate neutron stars for hundreds of milliseconds or even seconds. Therefore we were able to investigate gravitational wave output also, and for the first time in the published literature, in this low-frequency range which is so interesting for future observations. We did indeed find a substantial amount of energy in this range, likely a mixture of the (still ongoing) rearrangement due to the dynamical instability, and Alfvén modes travelling along the field lines. While it is hard to disentangle these effects precisely, we do observe a substantial gravitational wave output from these modes. It is still well below what can be observed at present, but future gravitational wave detectors like the proposed Einstein Telescope could be able to observe parts of the signal under benign conditions. The unknown damping timescale of the Alfvén modes, and the question of frequency stability of the modes, will have to be addressed in the future to say more about this, but it is still an exciting result.

In extension of the work above, we have recently finished a more detailed investigation of the nature of the large-scale magnetic field instability under different conditions [11]. While the earlier studies were focussed on a “canonical” model of a neutron star (a nonrotating polytrope with $K = 100$ and $\Gamma = 2$) which is often employed in the literature, we have extended the parameter space to include different equations of state and also rotation. There are a number of reasons for this: first, it is important to estimate how robust results obtained in a special case are under different (but still reasonable) conditions. Connected with that, we wanted to

estimate the parameter dependence of the dynamical instability, saturation state, and of course also the level of gravitational wave emission.

The equation of state of a real neutron star is largely unknown. In fact, finding constraints on the EOS is one of the major goals of modern neutron star research, since this ties directly into nuclear physics and the properties of matter at supernuclear densities (as are present in the core of the neutron star). So far, we employ a popular, but simple and parametrized model, the polytrope, in the hope that the parameter space spanned by the models represents, to some extent, the full variety of behavior of real neutron stars, at least as far as the questions we are trying to address are concerned. In the most recent paper, we have varied the stiffness of the EOS (which also determines the maximum possible mass before collapse to a black hole would occur, and the mass-radius relation) between very low and very high values to observe how the gravitational wave output changes with this important parameter.

The other parameter we have varied is the rotation of the neutron star, and this deserves some explanation. Observed magnetars are slowly rotating (periods in the order of 10 s, which is high compared to the sound crossing timescale of milliseconds), so the initial focus on nonrotating models was a reasonable first step. However, we wanted to understand the behavior of the instability when rotation is introduced on grounds of principle, and we have therefore performed a number of simulations of neutron star models with low to very rapid rotation rates (periods in the order of milliseconds) to observe the dependence of the dynamical instability on the rotation rate. It should be added that strongly magnetized neutron stars are probably not *born* slowly rotating, so understanding the stability of magnetic fields in their interior is not just an academic exercise.

The results of this extended study can be summarized as follows: (i) gravitational waves from f -modes caused by magnetar flares are unlikely to be detected in the current or near-future generation of gravitational waves observatories, (ii) gravitational waves from Alfvén waves propagating inside the neutron star are more likely candidates, although this interpretation relies on the unknown damping time of these modes, (iii) any magnetic field equilibria derived from our simulations are characterized as non-axisymmetric, with approximately 65 % of their magnetic energy in the poloidal field, (iv) rotation acts to separate the timescales of different instabilities in our system, with the varicose mode playing a more major role due to a delayed kink instability and (v) despite the slowing growth rate of the kink mode, it is always present in our simulations, even for models where the rotational period is of the same order as the Alfvén timescale. Details can be found in [11].

We are now at the point where we have explored a good part of the essential parameter space of the magnetic field instability. Future work will include additional field topologies (toroidal fields) and also differential rotation. The latter is particularly interesting since differential rotation can interact with magnetic fields in interesting ways, e.g. via dynamo effects or the magnetorotational instability. We now just begin to really understand how magnetic fields in neutron stars look like, and how they interact with the observed phenomena.

Table 1 Each simulation used 1 GPU. The simulations investigate magnetic field stability in rotating magnetized neutron star models. Each neutron star is a rapid rotator with 400 Hz frequency, and a central magnetic field strength of 4.8×10^{16} G. Some simulations additionally have an imposed toroidal field component

Simulation	Time	SUs	Notes
4.8e16G, 400 Hz	192 h	1536	Rotating magnetized simulation of purely poloidal field.
4.8e16G, 400 Hz, tor 0.1	192 h	1536	Imposed toroidal field.
4.8e16G, 400 Hz, tor 0.5	192 h	1536	Imposed toroidal field.
4.8e16G, 400 Hz, tor 1.0	192 h	1536	Imposed toroidal field.

3 Codes

All studies described above were performed with the Horizon code [8]. The code solves the equations of general relativistic magnetohydrodynamics (GRMHD) on general spacetimes and in three dimensions, and is therefore ideally suited to investigate the complex and often nonsymmetric dynamics of magnetic fields in neutron stars.

Horizon is fully implemented in CUDA, and is therefore entirely ported to make use of NVIDIA GPUs. Since the code operates on a regular uniform mesh, many operations can be parallelized/vectorized in a straightforward manner, and coherency in the memory access patterns (which is important to obtain a high memory bandwidth on GPUs) is possible in many cases. While these statements would equally apply to Newtonian MHD, GRMHD has in addition a higher *arithmetic complexity* than standard MHD codes, so that we, overall, obtain a very high speedup compared to a reference CPU implementation.

All major operations are performed on the GPU, and no memory transfers over the PCI bus are performed to update the simulation as long as only one GPU is employed. The code also has an experimental MPI implementation, but we did not yet employ this part of the code in our studies, though this is planned for the future.

4 Resource Usage

We have employed several machines to perform the simulations leading to the publications mentioned above. Initially, the nehalem cluster at HLRS was the only larger resource available to us, but subsequently one of us (P. Lasky) obtained access to a new, large machine in Australia, which also coincided with his move from Germany to Melbourne. For practical reasons, many of the simulations were performed on this so-called MASSIVE cluster, and only a small number on the nehalem cluster in Stuttgart. We have collected a list of the simulations performed in Table 1.

References

1. R. C. Duncan, C. Thompson, *Formation of very strongly magnetized neutron stars - Implications for gamma-ray bursts*, ApJ 392, L9 (1992).
2. C. Thompson, R. C. Duncan, *The soft gamma repeaters as very strongly magnetized neutron stars - I. Radiative mechanism for outbursts*, MNRAS 275, 255 (1995).
3. The LIGO Scientific Collaboration: B. Abbott et al, *Search for gravitational-wave bursts from soft gamma repeaters*, PRL 101, 211102 (2008).
4. H. Sotani, K. D. Kokkotas, *Alfvén polar oscillations of relativistic stars*, MNRAS 395, 1163 (2009).
5. A. Colaiuda, H. Beyer, K. D. Kokkotas, *On the quasi-periodic oscillations in magnetars*, MNRAS 396, 1441 (2009).
6. J. Braithwaite, *Axisymmetric magnetic fields in stars: relative strengths of poloidal and toroidal components*, MNRAS 397, 763 (2009).
7. V. Duez, J. Braithwaite, S. Mathis, *On the stability of non-force-free magnetic equilibria in stars*, ApJ 724, L34 (2010).
8. B. Zink, *HORIZON: Accelerated general relativistic magnetohydrodynamics*, arXiv:1102.5202 (2011).
9. P. D. Lasky, B. Zink, K. D. Kokkotas, K. Glampedakis, *Hydromagnetic instabilities in neutron stars*, ApJ 735, L20 (2011).
10. B. Zink, P. D. Lasky, K. D. Kokkotas, *Are gravitational waves from giant magnetar flares observable?*, Phys. Rev. D 85, 024030 (2012).
11. P. D. Lasky, B. Zink, K. D. Kokkotas, *Gravitational waves and hydromagnetic instabilities in rotating magnetized neutron stars*, arXiv:1203.3590 (2012).

Part II

Solid State Physics

Prof. Dr. Holger Fehske

The computational treatments in the fields of solid state physics and material science were directed towards the electronic and structural properties of fascinating systems ranging from metals via ceramics, oxides, semiconductors and their surfaces to nanostructures. The six contributions selected reveal that research in this area extraordinarily benefits from the supercomputing facilities of the High Performance Computing Center Stuttgart during this funding period.

A good example is the work by the Stuttgart group (J. Roth et al.) on laser ablation in metals that has been carried out with fully parallelized IMD, a molecular dynamics simulation package. In this problem, laser radiation directly acts on the metal's nearly-free electrons, which after excitation to a non-equilibrium state, quickly equilibrate. For this reason a two-temperature model with separate temperatures for electrons and phonons applies. The resulting set of coupled heat-conduction differential equations for electrons and ions gives the time evolution of the electron and lattice temperature within the system. On an atomistic scale the continuum description has to be replaced by molecular dynamics. Choosing aluminium as a reference system, the authors first determined the electronic material properties (e.g. the heat capacity and the electron-phonon coupling strength) and afterwards analyzed the melting depth and the ablation threshold, also if a two-pulse sequence with long time separation or overlapping pulses were applied.

The *ab initio* study by B. Höffing and F. Bechstedt from the European Theoretical Spectroscopy Facility and the University Jena deals with electronic surface properties of the transparent conducting oxides In_2O_3 , ZnO_2 , and ZnO . Employing the repeated-slab supercell method, the authors have carried out density functional (DFT) calculations which include many-body effects perturbatively in order to determine ionization energies and electron affinities for different surface

Prof. Dr. Holger Fehske

Institut für Physik, Lehrstuhl Komplexe Quantensysteme, Ernst-Moritz-Arndt-Universität
Greifswald, Felix-Hausdorff-Str. 6, 17489 Greifswald, Germany

e-mail: fehske@physik.uni-greifswald.de

orientations and terminations of the oxides. Thereby the quasiparticle equation for the electronic self-energy is solved on top of the self-consistent solution of a generalized Kohn-Sham equation using Hedin's GW method. This allows to analyze the surface stability in dependence on the chemical potential of the oxygen atom respectively molecule. The numerical results indicate a strong dependence of the surface barrier on the orientation and termination of the surface which might be caused by the strong surface dipoles existing in these ionic compounds.

A technically related, but physically different project addresses surface states at silicon surfaces with different hydrogen coverage, which is a central challenge on the way to optimize solar cells. The technologically highly relevant investigations by the Paderborn's group (U. Gerstmann et al.) were based on a gauge-including projector augmented plane wave approach in the framework of DFT with a gradient-corrected PBE functional. The focus is on electron paramagnetic resonance fingerprints of the paramagnetic states created by hydrogen adsorbed on Si(111) and Si(001). At this the numerics provides the microscopic structure and magnetization density of a paramagnetic hydrogen vacancy at different-oriented Si surfaces. Interestingly the calculations could be extended to micro- and nano-crystalline 3C-SiC being basic solar cell materials.

The group from the Stuttgart MPI for Solid State Research (G. Bester and P. Hang) has continued its investigations of the vibrational properties of both passivated and unpassivated nanoclusters. The authors discuss how molecular-type surface-acoustical and surface-optical modes coexist and interact with bulk-type vibrations to the point where structural changes were induced by the surface. Their ab initio DFT calculations give deeper insight into the thermodynamics of III-V and II-VI colloidal nanoclusters with up to thousand atoms. Of particular importance in this respect seems to be the behavior of the low-temperature specific heat. Noteworthy, coherent acoustic modes could be identified in good agreement with experiment.

To re-emphasize, all the projects in the rapidly growing field of computational solid state physics have in common, besides the high scientific quality, the strong need for computers with high performance to achieve their results. Therefore the new leading edge supercomputer technology being available at the HLRS will play an essential role in their physical research.

Molecular Dynamics Simulations of Laser Ablation in Metals: Parameter Dependence, Extended Models and Double Pulses

Johannes Roth, Johannes Karlin, Marc Sartison, Armin Krauß,
and Hans-Rainer-Trebin

1 Introduction

Laser ablation has become a very useful tool in machining today. For example for drilling holes, welding, engraving or coating by deposition of laser-irradiated material. The opposite process, laser removal of material is in general called laser ablation and some aspects of this process shall be discussed here.

If the target of laser ablation are metals, and the pulses applied have durations of a few femtoseconds, then the time evolution of the process can be described in the following way: The laser acts on the free electrons of the metal and excites them. The next step is a thermalization of the electrons which leads to an electron temperature different from the ordinary lattice temperature. Then the electron system and the lattice start to exchange heat which means especially heat conduction into the bulk by the electrons. Up to this point the process is dominated by the quantum nature of the material. The lattice is heated up by the energy obtained from the electrons, it melts, and finally ablation occurs. The latter processes now take place on the scale of several picoseconds and can thus be simulated by classical molecular dynamics simulations.

Since we want to simulate large samples with millions to billions of atoms, we cannot use ab-initio-methods to study the quantum effects. Instead of that we apply a continuum model, the so-called two temperature model (TTM) [1, 3, 7], which consists of two coupled heat balance equations formulated for the electrons and lattice as a function of the temperatures mentioned above. The lattice equation will later be replaced by molecular dynamics (MD) simulations which allows us to obtain atomistic information about the ablation process. The combined model is called TTM+MD.

J. Roth (✉) · J. Karlin · M. Sartison · A. Krauß · H.-R. -Trebin
Institut für Theoretische und Angewandte Physik, Universität Stuttgart, Stuttgart, Germany
e-mail: johannes@itap.physik.uni-stuttgart.de

The behavior of a material in classical MD simulations is given by the structure (initial conditions) and the interaction. The question is now how this behavior is driven by the coupling to the electron system whose quantum nature is described by the continuum model. In the case of metals there are three relevant parameters: electron heat conductivity, electron heat capacity and electron-phonon coupling. With respect to these parameters all metals can be divided into classes. Keeping the interaction and the crystal structure fixed we have varied the parameters within the experimentally observed range. In a first part of this work we will present results on this study.

In Fourier's law for the heat flow generated by a temperature gradient it is assumed that the temperature change is instantaneous. This might not be true in general for femtosecond pulses. Several people have generalized Fourier's law to include a finite relaxation time [4, 5]. If combined with the conservation of energy this leads to a generalization of the electron heat balance equations to a heat wave equation, also called the telegraph equation. For the coupled equations this is called an extended TTM (ETTM) and combined with MD an ETTM+MD model. We show numerically by solution of the ETTM and by simulations with the ETTM+MD model that the simple TTM leads to satisfactory results for example for Al and Cu, but that the ETTM has to be applied in case of Pb for example.

Drilling holes with femtosecond lasers requires thousands of pulses with time intervals of hundreds of microseconds between two consecutive pulses. The complete process is obviously beyond the realm of classical molecular dynamics simulations. But since the sample cools down completely between two pulses we can simply run two simulations at pristine and ablated samples thus skipping the long interval in between. In the work presented here we have reduced the time interval even more such that the sample is not cooled down completely. Yet another aspect is the temporal shape of the pulses which is usually Gaussian. But it is predicted from experimental studies with increasing or decreasing overlapping pulse sequences that non-Gaussian pulses shapes may be more effective in ablation. We have studied two cases of two overlapping Gaussian pulses, one increasing, one decreasing and will report first results.

1.1 *Molecular Dynamics Simulations*

All the simulations have been carried out with IMD, the **ITAP Molecular Dynamics** simulation package [11, 17].¹ IMD contains modules to simulate laser ablation with the two-temperature model [9, 21]. Most of the features are fully parallelized using MPI. With respect to details of the usage and the parameters applied for ablation we refer the interested reader to the IMD homepage.¹

¹Available at <http://www.itap.physik.uni-stuttgart.de/~imd>

2 The Two-Temperature Model and Its Coupling to Molecular Dynamics

In femtosecond laser ablation the laser radiation acts directly on the free electrons of the metal, exciting them to a non-equilibrium state. If the electrons equilibrate fast enough it is possible to define a separate temperature for them. This is the basis of the two-temperature model (TTM) where separate temperatures for the electrons and the lattice or phonons are introduced. The model was introduced by Anisimov et al. [1] based on ideas formulated by Ginzburg, Kaganov and others [3, 7]. Laser ablation is then described by a system of coupled heat conduction equations for the electrons and the ions separately:

$$C_e(T_e) \frac{\partial T_e}{\partial t} = \nabla[K_e \nabla T_e] - \kappa(T_e - T_i) + S(\mathbf{x}, t), \quad (1)$$

$$C_i(T_i) \frac{\partial T_i}{\partial t} = \nabla[K_i \nabla T_i] + \kappa(T_e - T_i). \quad (2)$$

Equations (1) and (2) describe the time evolution of the electronic (T_e) and the ionic (T_i) or lattice temperature within the metal. $C_{e,i}$ are the heat capacities, $K_{e,i}$ the heat conductivities, κ the electron-phonon coupling constant and $S(\mathbf{x}, t)$ the external laser field. With these equations the laser field is coupled physically meaningfully into the system: first the energy is brought into the electronic system via a source term $S(\mathbf{x}, t)$. Then the electronic system transports the heat diffusively into the bulk and at the same time interacts with the ions.

To work on an atomistic scale, Eq. (2), which is a continuum description of the temperature field, has to be replaced by molecular dynamics. This means that instead of (2) the following equations of motion have to be solved together with (1):

$$m_j \frac{d^2 \mathbf{x}_j}{dt^2} = -\nabla_{\mathbf{x}_j} U(\{\mathbf{x}_k\}) - \frac{\kappa}{C_i} \frac{(T_i - T_e)}{T_i} m_j \frac{d \mathbf{x}_j}{dt}. \quad (3)$$

This is the coupled TTM+MD model [6, 14]. The lattice parameters C_i and K_i are no longer present, since they are intrinsic properties of the atoms. The coupling parameter κ has to be translated to atomistic observables which can be done by division by C_i , the atomistic heat capacity which is obtained directly in the simulations.

Thus we are left with the parameters C_e , K_e , and κ which can vary in a broad range. The heat capacity C_e which is a linear function of temperature for metals is replaced by the heat capacity coefficient $\gamma = C_e/T$ which is constant over a broad temperature range.

3 Parameter Studies of Electronic Material Properties

While the heat capacity and heat conductivity of the ions are intrinsic properties of the interaction we note that the electron properties might vary in a broad range. The idea here is to keep the atomistic part, namely the crystal structure and the interaction fixed, vary the electronic parameters in the range observed in experiment, and determine how strongly the melting depth and ablation threshold depend on these parameters. More details may be found in [13].

As atomistic reference system we chose aluminum. First of all, since it is well studied and has previously been used in our simulations [16, 18]. Second, since it turns out that it is a good reference point centering the range of parameter values observed for other metals.

Parameter Range and Sample Properties

The parameters for standard Al are the electron heat capacity coefficient $\gamma = 135 \text{ J}/(\text{m}^3 \text{ K}^2)$, the electron-phonon-coupling constant $\kappa = 2.45 \cdot 10^{17} \text{ J}/(\text{m}^3 \text{ K s})$, and the electron thermal conductivity $K_e = 235 \text{ J}/(\text{s K m})$ (see for example Appendix C of the book of Bäuerle [2]). The laser fluences² for the study of the melting depth were varied between 457.77 and 951.44 J/m^2 in steps of 114.4 J/m^2 . K_e was increased from 67.8 to 406.8 $\text{J}/(\text{s K m})$ in steps of 67.8 $\text{J}/(\text{s K m})$, γ in four steps from 1.19 to 118.97 $\text{J}/(\text{m}^3 \text{ K}^2)$ and κ in four steps from $1.36 \cdot 10^{16}$ to $6.78 \cdot 10^{18} \text{ J}/(\text{m}^3 \text{ K s})$. The sample volume was $184.44 \times 4.84 \times 4.84 \text{ nm}^3$ containing about 260,000 atoms.

Results for the Melting Depth

In all cases we keep two of the three parameters fixed at the aluminum value and vary the third one. If the melting depth is plotted as a function of fluence and heat capacity coefficient γ (Fig. 1, top) we find only a weak dependence on this parameter. The reason is that the amount of energy stored in the electronic system is rather small. If the melting depth is plotted as a function of fluence and the thermal conductivity K_e (Fig. 1, middle) we observe an almost linear increase with respect to this parameter. The reason is that the energy can penetrate deeper into the bulk for higher conductivities. If the melting depth is plotted as a function of fluence and the electron-phonon coupling constant κ (Fig. 1, bottom) we observe a region between 10^{16} and $10^{17} \text{ J}/(\text{m}^3 \text{ K s})$ where no melting is induced. The coupling is too small and the energy stays in the electron system. At high κ and for higher fluences

²The strength of a laser beam is typically given by power per area, called laser fluence.

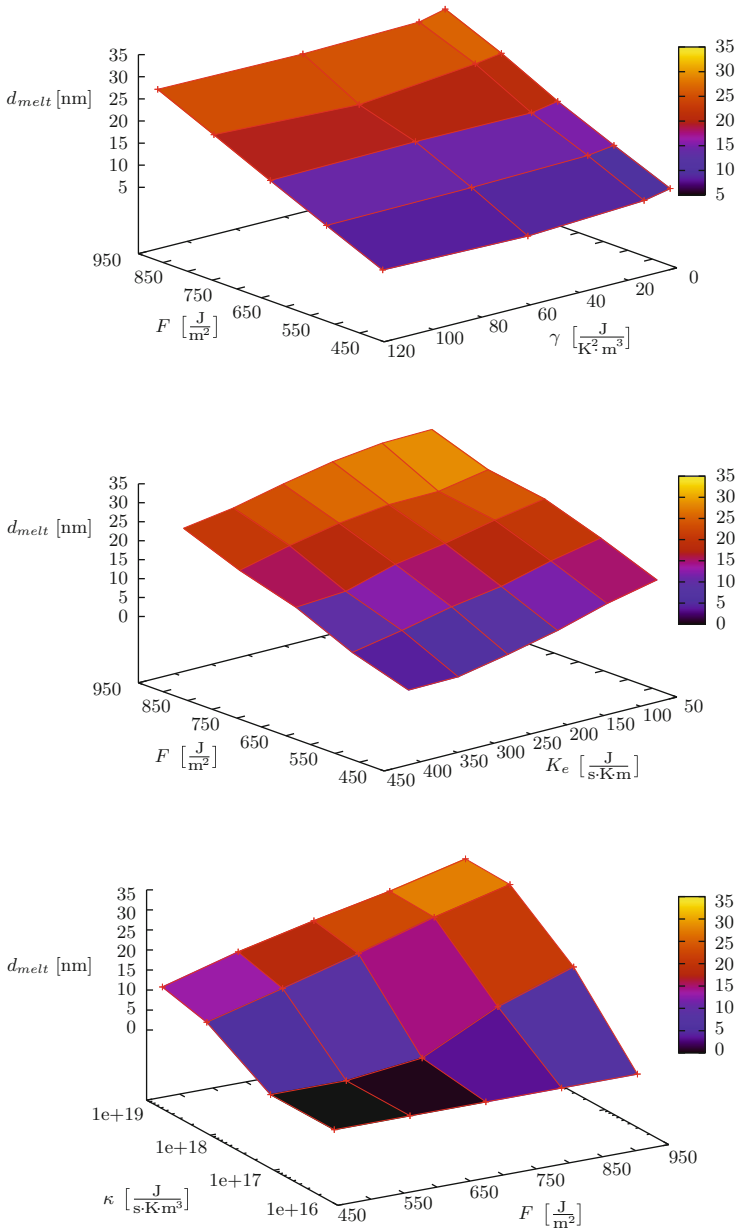


Fig. 1 Melting depth as a function of fluence and parameters. From *top to bottom*: heat capacity coefficient γ , thermal conductivity K_e , electron-phonon-coupling parameter κ

the the melting depth increases again with the parameters. The reason is similar to the previous case: a better coupling leads to a higher energy transfer which finally melts the lattice.

Results for the Ablation Threshold

The ablation threshold can suitably be fitted by simple functions:

1. The dependency of threshold fluence on the heat capacity coefficient can be described by a power law $F(\gamma) = a \cdot \gamma^b + c$ with $a = 0.00706 \text{ m}^4 \text{ K}^4/\text{J} \pm 0.41 \%$, $b = 1.98 \pm 2.5 \%$, $c = 1144 \text{ J/m}^2 \pm 0.044 \%$. Here, γ is limited to the range between 1.19 and $118.97 \text{ J/(m}^3 \text{ K)}$ since outside this region no ablation has been observed. The ablation threshold increases with increased energy storage, but the influence is again rather small.
2. The threshold fluence is a linear function of the thermal conductivity over the whole range of parameters $F(K_e) = a \cdot K_e + b$ with $a = 0.284 \text{ K s/m} \pm 3.6 \%$ and $b = 160 \text{ J/m}^2 \pm 1.7 \%$. For higher conductivities a larger fluence is required since heat is lost by diffusion into the bulk.
3. The threshold fluence with respect to the electron-phonon-coupling has been fitted with a linear function although here the dependency is not so obvious as can be concluded from the error bars: $F(\kappa) = a \cdot \kappa + b$. The values are $a = -1,351 \text{ m K s} \pm 28.9 \%$, $b = 184.711 \text{ J/m}^2 \pm 0.93 \%$. If $\kappa > 6.78 \cdot 10^{18} \text{ J/(m}^3 \text{ K s)}$ the coupling is so strong that no stable system can be simulated. For higher couplings a smaller fluence is sufficient since heat is transfer faster to the lattice and leads to ablation.

Summary of the Parameter Studies

In summary we conclude that the behavior of the ablation threshold is as expected: the electron system stores little energy, if the thermal conductivity is high, the heat flows into the bulk and leads to deep melting and increasing ablation threshold. A high electron-phonon coupling constant on the other hand leads to earlier melting and ablation, thus requiring less fluence. In general the situation is more complicated since the parameters are temperature-dependent. The metals may be partitioned into three classes: Al where the parameters γ and κ are merely rescaled, Ag, Cu, Au with a high temperature region of increased parameters and Ni, Pt, W with decreasing parameters [10]. A study like ours would now require to work with average or effective parameters.

4 An Extended Two-Temperature Model

Derivation of the Model

In Fourier's law of heat flow $\mathbf{q}(\mathbf{x}, t) = -K_e \nabla T(\mathbf{x}, t)$ and in the standard two-temperature model [1,3,7] it is assumed that the temperature change is instantaneous and the relaxation time of the electrons τ_e is zero. If a finite relaxation time τ_e is introduced for the electrons [4,5] we get the modified law for the heat flow

$$\tau_e \partial_t \mathbf{q}(\mathbf{x}, t) = -\mathbf{q}(\mathbf{x}, t) + K_e \nabla T(\mathbf{x}, t) \quad (4)$$

which together with the conservation of energy leads to a heat wave equation:

$$\tau_e C_e(T_e) \frac{\partial^2 T_e}{\partial t^2} + C_e(T_e) \frac{\partial T_e}{\partial t} = \nabla [K_e \nabla T_e] - \kappa (1 + \tau_e \frac{\partial}{\partial t})(T_e - T_i) + (1 + \tau_e \frac{\partial}{\partial t}) S(\mathbf{x}, t) \quad (5)$$

Comparing the new equation to the original equation (1) we observe that in addition to the second time derivative on the left side which has frequently been added phenomenologically, the electron-phonon coupling parameter κ and the laser source term $S(\mathbf{x}, t)$ have to be supplemented by a differential operator.

Results for ETTM

The extended two-temperature model (ETTM) has been solved numerically by Hüttner and Rohr [4] and reproduced by us for 1 ps pulses. For Al ($\tau_e = 0.067$) the differences between the TTM and the ETTM are negligible. They are significant for Cu ($\tau_e = 0.276$) where a difference in peak temperature of 1,000 K and final temperature difference of 500 K is observed. They grow even larger for Pb ($\tau_e = 0.377$) with a peak temperature difference of 6,000 K and a final temperature difference of 1,000 K.

Results for ETTM+MD

We studied the ETTM+MD model for a test sample with 180,000 atoms and a number of different metals. Standard embedded atom interactions were applied [20]. The electron and lattice temperatures were monitored as a function of time at a certain depth within the sample. In all cases we find that the lattice temperature is completely unaffected by a finite relaxation time. The behavior of the electron temperature is presented in Fig. 2 for 0.1 ps pulses. For Al and even for Pb there are only minor differences between the two models although the relaxation time




Fig. 2 Electron temperatures in the TTM+MD and the ETTM+MD model. From left to right: Al, Cu and Pb. The results for TTM+MD are printed in blue and for ETTM+MD in *red*

increases by a factor of 4. For Pb there are deviations between TTM+MD and ETTM+MD, but considerably reduced if compared to the TTM vs. ETTM case, i.e. at the peak temperature from 6,000 to 3,000 K and at the final temperature from 1,700 K to a very low value.

Summary for the ETTM

The results for ETTM and ETTM+MD are not directly comparable due to the different parameters used. In general the standard two-temperature model is certainly sufficient for long pulse durations, i.e. longer than the electron relaxation time τ_e . The extended two-temperature models are required if τ_e is large as in the case of Pb. The heat transport is delayed in ETTM and ETTM+MD which leads to a higher electron temperature since the heat flow to the lattice ions is reduced. The wave nature of the extended equation (5) does not play a role in the regime studied.

5 Properties of Two-Pulse Sequences in Aluminum

A two-pulse sequence with a long time separation and two overlapping pulses formed by two increasing and decreasing Gaussians have been simulated to study the effect of pulse shapes on the effectivity of ablation. More details can be found in [8].

Pulse Shapes

The illumination of the samples with the laser beam was chosen homogeneously due to their small cross sections. Thus the intensity distribution is given by a one dimensional function

$$S_{1D} = (1 - R) \cdot \mu \exp(-\mu x) \sigma_E \cdot g(t) \quad (6)$$

The reflectivity R is set to 0.825, μ is the inverse absorption depth and $g(t)$ the temporal shape of the laser pulse. The energy density per area parameter σ_E for aluminum is 144.18 J/m^2 .

The starting point was a Gaussian pulse with width $\sigma_t = 0.4 \text{ ps}$:

$$g_0(t) = \frac{1}{\sqrt{2\pi}\sigma_t} \exp\left(-\frac{1}{2} \frac{(t - t_0)^2}{\sigma_t^2}\right). \quad (7)$$

For all simulations t_0 was set to 1.018.

The shapes of the double pulses can be described with the following formula

$$g_i(t) = \frac{1}{\sqrt{2\pi}\sigma_t} \left(a_i \exp\left(-\frac{1}{2} \frac{(t - t_0)^2}{\sigma_t^2}\right) + b_i \exp\left(-\frac{1}{2} \frac{(t - c_i \cdot t_0)^2}{\sigma_t^2}\right) \right) \quad (8)$$

with amplitudes a_i and b_i and time interval c_i .

For the separated double pulse we set $a_1 = b_1 = 1$ and $c_1 = 15$. For the increasing pulse we set $a_2 = 3/4$, $b_2 = 1$ and $c_2 = 2$. For the decreasing pulse we set $a_3 = 1$, $b_3 = 3/4$ and $c_3 = 2$.

Results of the Simulation of Two-Pulse Sequences

The sample size for the simulations was $241.92 \times 4.84 \times 4.84 \text{ nm}^3$ containing about 350,000 atoms. The samples were equilibrated at 305 K. Sequences of simulations were carried out at intervals of $\Delta\sigma_E = 8 \text{ J/m}^2$. Thus the data given for the ablation thresholds are only accurate within about $\Delta\sigma_E/2$.

1. **Single pulse:** Ablation sets in at $\sigma_E = 224.28 \text{ J/m}^2$ and $t = 12.6 \text{ ps}$ with an ablation depth of 21.2 nm. The sample melts down to a depth of 35.0 nm.
2. **Two pulses at a fixed time interval:** Ablation is observed at $\sigma_E = 144.18 \text{ J/m}^2$ and $t = 29.8 \text{ ps}$, thus after the second pulse has arrived. The first pulse melts the probe up to a depth of 26.0 nm at the time when the second pulse hits the probe. At first the ablation depth is 10.4 nm, but finally the sample melts down to a depth of 79.6 nm and a layer with 26.0 nm is ablated.

The electron and lattice temperature have nearly reached equilibrium when the second pulse hits the sample, but the sample is still warm.

3. **Increasing pulse:** Ablation sets in at $\sigma_E = 144.18 \text{ J/m}^2$ and $t = 11.9 \text{ ps}$ with an ablation depth of 11.4 nm and a melt depth of 34.5 nm. The sample melts up to a depth of 54.0 nm and the total depth of the ablated layer is 20.0 nm.
4. **Decreasing pulse:** Ablation is observed at $\sigma_E = 137.17 \text{ J/m}^2$ and $t = 12.2 \text{ ps}$ with an ablation depth of 12.2 nm and a melt depth of 34.2 nm. The sample melts up to a depth of 51.0 nm and the total depth of the ablated layer is 20.3 nm.

At first the evolution of the electron temperature is sharp as for a single pulse, but then it continues to rise slowly. This behavior is caused by the inset of the second pulse at that time. The lattice temperature behaves similar to the previous case. Its rise time is longer than for a single pulse. Obviously, the specific pulse shape of the laser beam is smeared out and only its width determines the behavior of the lattice temperature.

The total fluence to achieve ablation is $1,260 \text{ J/m}^2$ for a single pulse, $1,648 \text{ J/m}^2$ for two separate pulses, $1,439 \text{ J/m}^2$ for the increasing and $1,367 \text{ J/m}^2$ for the decreasing pulse.

Summary of the Simulation of Two-Pulse Sequences

Two separate pulses require considerably more fluence than a single pulse to achieve ablation. Furthermore, ablation is observed only *once*. There are several reasons for this behavior: part of the energy of the second pulse is still absorbed by the ablated layer and lost for further heating. The fact that sample is still warm should play a minor role. Another reason for the higher threshold is that the energy of the pulses add up but ablation is initiated by one pulse only. The total number of ablated layers is the similar to the single pulse. The final melting depth is much larger than the depth for a single pulse. Obviously the energy added to the system acts as expected. To supplement this study the distance between the pulses was varied. If the height of the second pulse is between 75 and 87.5 % of the first pulse we find that σ_E has to rise up to about 152.19 J/m^2 to achieve ablation.

The two cases of overlapping pulses are very similar. The fluence for ablation is about 10 % higher than for the single pulse. This is due to the lower maximum and shows that it is better to concentrate the laser energy into a very short pulse. The melting depths are larger than for a single pulse which again reflects the fact that more energy has been added to the system. In contrast to a single sharp pulse ablation splits up into several layers, but in total the depth is similar to the case of a single pulse.

The results should be compared to the work of Rosanti and Urbassek [12]. The pulses they study are half as wide as ours, and they vary the time between the two pulses. Unfortunately they do not report melting depths and ablation thresholds.

Experimental results show that the most effective way to ablate material by multi-pulses is a decreasing sequence of pulses with very short distance between them, almost overlapping. For increasing sequences the small peaks at the beginning do not lead to melting or ablation and thus this pulse shape is not as effective as decreasing pulses. The most effective pulse shape, however, seems to be a sharply rising pulse with a slow decay [19].

Table 1 Pair potentials

Year	Machine	Time (μ s/step/atom)	Cores	Atoms/core
1996	T3E 600	53	8–512	4,630–37,044
1999	T3E 1200	38	128–512	10,117,414
2006	Blue Gene/L	20 ± 1	1–2,048	2,000–128,000
2010	XT5 (Jaguar)	6	131,072	1,048,576
2012	XE6	6.5	512	16,384–1,048,576
2012	XK6 Cuda	4.7	16	16,384–1,048,576
2012	XE6 OMP	6.0	16–512	16,384–1,048,576

Table 2 EAM interactions

Year	Machine	Time (μ s/step/atom)	Cores	Atoms/core
2006	Blue Gene/L	36 ± 5	1–2,048	2,000–128,000
2010	Nehalem	8.3	2,048	29,297
2010	Nehalem	9.7	3,072	19,531
2010	Nehalem	9.9	3,584	16,741

Ranges indicate that the performance is almost the same within this interval.

6 Performance and Benchmarks

General benchmark data for IMD have been given by Stadler et al. [17]. The data demonstrate that IMD scales almost linearly in weak scaling (same number of atoms per processor) and fairly well for strong scaling (total number of atoms constant, thus communication load growing). This behavior is still valid as a systematic study on the Blue Gene/L clearly shows (See the previous HLRS report [9]).

Tables 1 and 2 collect benchmark results for T3E, Blue Gene/L (Jülich), Nehalem, Hermit (XE6/XK6), and Jaguar (XT5). EAM interactions require typically up to twice as much time as pair interactions. Currently we find that IMD is about three times faster on Hermit than on the Blue Gene/L. On Jaguar, the degradation of performance going from 1 to 131,072 cores was only about 11 % for 1,048,576 atoms per core! The total amount of atoms in this case was $137'438'953'472$. IMD is about 8 % faster with the GNU compiler as compared to the Cray compiler. If the extra time required for EAM interactions with respect to pair potentials is taken into account we find that IMD is currently about as fast on the Nehalem as on Hermit.

The numbers for pair potentials may be not as reliable as those for EAM interactions since the former are pure benchmark results from short time runs while the latter are results from true production simulations from multi-hour runs.

In all cases tested up to now, OpenMP lead to a rapid loss of performance, for example by a factor of 10 when going from one to two threads on Hermit (pure OMP). Cray is currently carrying out benchmarks to figure out how to improved the performance with OpenMP. If OpenMP and MPI are used, the performance loss

is down to 44 % for 16,384 atoms and 11 % for 1,048,576 atoms for 8 OpenMP threads.

The GPU version of IMD is still under development. Currently it can be used only for pair potentials and monatomic samples.

A major issue are the development of inhomogeneities during simulation. The performance loss for 60 million atoms was about 75 % after 200,000 time steps. For movies the full sample has to be run until the end. For production runs a simple cure is available: since the simulation has to be restarted anyways, the uninteresting bulk of the sample is cut off after 200,000 steps.

In general load balancing is a very hard task because the communication setup is so central to the code that changing it to improve performance would be almost equivalent to rewriting the code. Furthermore, the performance is strongly dependent on the simulated setup. So there is no general recipe. Since the performance problem is typically one-dimensional it has been discussed to break up the rigid assignment of atoms to equidistant chunks of the simulation box. However, if a reordering of the atoms occurs during simulation this could kill all improvements since it requires communication of huge amounts of atoms and vast amounts of extra memory.

7 Summary

We presented results on ablation simulation with homogeneous laser beams in the two-temperature model combined with molecular dynamics simulations. We studied the influence of electronic parameters, namely heat conductivity, heat capacity and electron-phonon-coupling, the question in which cases the TTM must be extended to finite electron temperature relaxation times. Finally we presented results on pulse sequences and overlapping pulses.

References

1. Anisimov S.I., Kapeliovich B.L., Electron-emission from surface of metals induced by ultrashort laser pulses, Perel'man T.L., *Zh. Eksp. Teor. Fiz.* **66**, 776 (1974) [*Sov. Phys. JETP* **39**, 375–380 (1974)].
2. Bäuerle D., *Laser Processing and Chemistry*, Fourth Edition, Springer Heidelberg 2011.
3. Ginzburg V.L., Shabanskiy V.R., Kineticheskaya temperatura elektronov v metallakh i anomalnaya elektronnaya emissiya, *Dokl. Akad. Nauk SSSR* **100**, 445–448 (1955).
4. Hüttner B., Rohr G., On the theory of ps and sub-ps laser pulse interaction with metals I. Surface temperature, *Appl. Surf. Sci.* **103**, 269–274 (1996).
5. Hüttner B., Thermodynamics and transport properties in the transient regime, *J. Phys.: Cond. Matt.* **11**, 6757–6777 (1999).
6. Ivanov, D.S., Zhigilei, V., Combined atomistic-continuum modeling of short-pulse laser melting and disintegration of metal films, *Phys. Rev. B* **68**, 064114 (2003).

7. Kaganov M.I., Lifshits I.M., Tanatarov L.V., Relaxation between electrons and the crystalline lattice, *Zh. Eksp. Teor. Fiz.* **31**, 232 (1956) [*Sov. Phys.-JETP* **4**, 173 (1957)].
8. Krauß A., Multipulsanregung in Metallen, Bachelor Thesis, Stuttgart (2011).
9. Roth, J., Trichet, C., Trebin, H.-R., Sonntag, S., Laser ablation of metals, in *High Performance Computing in Science and Engineering '10*, eds. W.E.Nagel, D.B. Kröner, M.M. Resch, Springer Heidelberg, 2011, pp. 159–168.
10. Lin Z., Zhigilei L.V., Celli V., *Phys. Rev. B* **77**, 075133 (2008).
11. Roth J., Gähler G., Trebin H.-R., A molecular dynamics run with 5.180.116.000 particles, *Int. J. Mod. Phys. C* **11**, 317–322 (2000).
12. Rosanti Y., Urbassek H.M., Ultrashort-pulse laser irradiation of metal films: the effect of a double-peak laser pulse, *App. Phys. A* **101**, 509–515 (2010).
13. Sartison M., Characterization of Ablation Properties, Bachelor Thesis, Stuttgart (2011).
14. Schäfer C., Urbassek H.M., Zhigilei L.V., Metal ablation by picosecond laser pulses: A hybrid simulation, *Phys. Rev. B* **66**, 115404 (2002).
15. Sonntag, S., Computer Simulations of Laser Ablation from Simple Metals to Complex Metallic Alloys, PhD Thesis, Stuttgart (2011).
16. Sonntag, S., Trichet, C., Roth, J., Trebin, H.-R., Molecular Dynamics Simulations of Cluster Distribution from Femtosecond Laser Ablation in Aluminum, *Appl Phys A* **104**, 559–565 (2011).
17. Stadler, J., Mikulla, R., Trebin, H.-R., IMD: A software package for molecular dynamics studies on parallel computers. *Int. J. Mod. Phys. C* **8**, 1131–1140 (1997).
18. Sonntag S., Roth J., Gähler F., Trebin H.-R., Femtosecond laser ablation of aluminum, *Appl. Sur. Sci.* **255**, 9742–9744 (2009).
19. Siegel J., Deep ablation in dielectrics with temporally shaped femtosecond pulses, Talk at the 11th Conference on Laser Ablation, Cancun, Mexico, (2011).
20. Sheng H., <https://sites.google.com/a/gmu.edu/eam-potential-database/Pb>, <http://cde.gmu.edu/node/39>.
21. Ulrich C., Simulation der Laserablation an Metallen, Diplomarbeit, Stuttgart (2007). <http://elib.uni-stuttgart.de/opus/volltexte/2007/3296/>

Electronic Surface Properties of Transparent Conducting Oxides: An Ab Initio Study

B. Höffling and F. Bechstedt

Abstract We investigate the surface properties of the transparent conducting oxides In_2O_3 , SnO_2 , and ZnO using density functional theory and quasiparticle calculations based on many-body perturbation theory. We employ the repeated-slab supercell method. An energy alignment of valence and conduction states via the electrostatic potential is applied to determine ionization energies and electron affinities for various surface orientations and terminations of the oxides. In addition, surface energies for different orientations of bixbyite In_2O_3 are calculated. We find a strong influence of surface orientation and preparation techniques on these fundamental quantities.

1 Introduction

Transparent conducting oxides (TCOs) like In_2O_3 , SnO_2 , and ZnO are routinely used as transparent electrodes in photovoltaic and optoelectric devices [1] as well as in transparent electronics based on doped oxides [2, 3]. They are transparent in almost the entire range of the solar spectrum and usually exhibit a high electron conductivity [4–6]. They are also used in silicon (Si) photonics and Si-based solar cells [7]. Electronic properties of their surfaces like ionization energy and electron affinity are frequently used to predict natural band discontinuities at the interfaces with other materials such as Si [8, 9]. The existence of surface or interface states within the fundamental gap can lead to electron-hole recombination and limit the efficiency of the device. Consequently, these parameters are of great interest, but due to sample preparation problems, rather poorly known. Modern theoretical approaches can help to address these questions.

B. Höffling (✉)

European Theoretical Spectroscopy Facility (ETSF) and Institut für Festkörpertheorie und -optik, Friedrich-Schiller-Universität Jena, Max-Wien-Platz 1, 07743 Jena, Germany

e-mail: benjamin.hoeffling@uni-jena.de

Density functional theory (DFT) is known to underestimate the fundamental band-gap of semiconductors, therefore many-body effects have to be taken into account correctly to describe the electronic properties of oxides [10–13]. We use modern quasiparticle (QP) calculations based on many-body perturbation theory [12, 14] to predict the electronic bulk properties of the body-centered cubic (*bcc*) bixbyite as well as the rhombohedral (*rh*) geometry of In_2O_3 , the most favored rutile (*rt*) geometry of SnO_2 , and wurtzite (*wz*) structure ZnO . We combine the results with DFT ground-state calculations to obtain surface energies, ionization potentials and electron affinities for various surface orientations and terminations of the TCOs. Due to the large cell size required for surface calculations and the high computational cost of quasiparticle methods, massively parallel machines are required to perform the calculations.

The underlying theoretical and computational methods are described in Sect. 2. In Sect. 3 our results are presented and their reliability discussed in the light of available measured values. Finally, in Sect. 4, we give a brief summary and conclusions.

2 Computational Methods

The ground-state properties of the oxides are determined in the framework of DFT [15], using the local density approximation (LDA) [16] for exchange and correlation (XC). We employ the XC functional of Ceperley and Alder [17]. The calculations for ZnO have been carried out in the generalized gradient approximation (GGA), employing the PW91 functional to model XC [18]. All computations are performed with the Vienna Ab initio Simulation Package (VASP) [19]. The projector-augmented wave (PAW) method [20] is used to describe the electron-ion interaction in the core region. Usually it allows for the accurate treatment of first-row elements such as oxygen and localized semicore states such as $\text{In}4d$, $\text{Zn}3d$, and $\text{Sn}4d$ by modest plane-wave cutoffs. The electronic wave functions are expanded into plane waves up to cutoff energies of 550 (In_2O_3), 450 (SnO_2), and 500 eV (ZnO), respectively [10–14].

Brillouin-zone (BZ) integrations are carried out as summations over special points of the Monkhorst-Pack (MP) type [21]. Monkhorst-Pack meshes of $5 \times 5 \times 5$ (cubic) or $8 \times 8 \times 8$ (rhombohedral) k-points are found to be sufficient for In_2O_3 [10]. For hexagonal ZnO a $12 \times 12 \times 7$ mesh is applied [12]. In the *rt*- SnO_2 case, we use a mesh of $8 \times 8 \times 14$ k-points [13].

All calculations were carried out on the NEC SX-9 system, on which both the scaling behavior and the performance per CPU for our code are most efficient (see Ref. [22] for details on performance and scalability).

By minimizing the total energy obtained within DFT-LDA or DFT-GGA we obtain the cubic (a_0) and non-cubic (a , c) lattice constants for bulk TCOs. For In_2O_3 we find values $a_0 = 10.09 \text{ \AA}$ [10] (experiment: 10.12 \AA [23]) and 5.48 \AA [10] (experiment: 5.49 \AA [23]) for the *bcc* and *rh* structure, respectively. In the case

of rt -SnO₂ we get $a = 4.74 \text{ \AA}$ and $c = 3.20 \text{ \AA}$ [13] (experiment: $a = 4.74 \text{ \AA}$ and $c = 3.19 \text{ \AA}$ [24]), and, finally, for wz -ZnO we observe $a = 3.28 \text{ \AA}$ and $c = 5.28 \text{ \AA}$ (experiment: $a = 3.25$ and $c = 5.2 \text{ \AA}$ [24]). Except for c for rt -SnO₂ the lattice constants differ from the corresponding measured values by less than 1 %.

The structural parameters are then used for the calculation of the QP band structures [14, 25]. The QP equation for the electronic self-energy is solved perturbatively on top of the self-consistent solution of a generalized Kohn-Sham (gKS) equation using Hedin's GW method. In the zeroth approximation we describe the GW self-energy by the spatially non-local XC potential $V_{XC}(\mathbf{x}, \mathbf{x}')$ described by the hybrid functional HSE06 of Heyd et al. [26–29]. The QP shifts for the gKS eigenvalues are computed within the G_0W_0 approach [30]. It has been demonstrated that for the studied compounds this treatment leads to energy gaps in excellent agreement with measured values [10, 12–14, 31].

The surface calculations are carried out using the repeated slab supercell method. The number of layers in the slab is 9 (bcc -In₂O₃(001)), 8 (bcc -In₂O₃(110)), 11 (rh -In₂O₃(001)), 8 (rt -SnO₂(001)), 19 (rt -SnO₂(100)), 6 (wz -ZnO(10 $\bar{1}$ 0)), and 20 (wz -ZnO(0001)), with 12 \AA of vacuum each. Usually orthorhombic slabs are applied resulting typically in $N \times N \times 1$ MP meshes, with $N = 3, 8, 8,$ and 12 for bcc -In₂O₃, rh -In₂O₃, SnO₂(001), and ZnO, respectively. The SnO₂(100) slab is treated using an MP mesh of $8 \times 14 \times 1$. For polar directions, i.e. bcc -In₂O₃(001), SnO₂(001), and ZnO(0001), we encounter the fundamental problem of a net dipole moment within the slab, and the additional difficulty that a slab with two non-equivalent surfaces does not allow the computation of surface energies. To get around these obstacles we employ symmetric slabs by breaking the stoichiometry within the supercell and adding an additional layer of oxygen or, alternatively, metal atoms. Because of different bond lengths in the [0001] and the [000 $\bar{1}$] direction, respectively, the ZnO(0001) slab cannot be symmetrized this way. We introduced a central Zn-layer and constrained the bond lengths to create a symmetric slab, thereby creating an unphysical strain in the center of the slab. Since the lateral cell-size of the (0001) – 1×1 slab is small, we make the slab thick enough to obtain a converged electrostatic potential exhibiting bulklike oscillations and a plateau in the vacuum region. However, due to the additional unknown strain we cannot calculate surface energies for this surface orientation.

To align the QP band levels in the slabs one needs to determine the electrostatic potential $V(\mathbf{x})$ acting on the electrons. It can be derived from the effective single-particle potential occurring in the Kohn-Sham equation [16] or the generalized Kohn-Sham equation [14] and is defined as the local part of the electron-ion interaction described by the pseudopotentials and the Hartree potential of the electrons. This quantity is independent of local (LDA), semi-local(GGA) or non-local (HSE06) description of the exchange-correlation part of the effective single-particle potential and therefore well suited to serve as a universal reference level.

The ionization energy

$$I = E_{\text{vac}} - E_v \quad (1)$$




Fig. 1 Planar average of the electrostatic potential $\tilde{V}(z)$ along the cubic axis $z \parallel [10\bar{1}0]$ for the ZnO($10\bar{1}0$) slab (black) and for the ZnO bulk calculations (red). The QP conduction band minimum E_c and valence band maximum E_v are shown as dotted lines. The ionization energy I and electron affinity A are indicated. The vacuum level is used as energy zero

and the electron affinity

$$A = E_{\text{vac}} - E_c \quad (2)$$

are defined as the energy difference between the vacuum level E_{vac} , i.e. the electrostatic potential as seen by an electron in the vacuum, and the valence band maximum (VBM) E_v and the conduction band minimum (CBM) E_c , respectively. Hence, to obtain QP values for I and A , one has to align the QP bulk band structure with the electrostatic potential as obtained through the DFT-LDA slab calculation. As an example, the alignment for the ZnO($10\bar{1}0$) surface is shown in Fig. 1. The plane averaged electrostatic potential $\tilde{V}(z)$ shows the characteristic atomic oscillations in the area of the slab and reaches a plateau in the vacuum region. By aligning the atomic oscillations in the slab with those obtained via the bulk calculation one derives I and A .

Computing I and A using this method we come up against a theoretical problem in the QP description. The GW approximation sets the vertex function Γ for the calculation of the XC self-energy $\Sigma = GW\Gamma$ to $\Gamma \equiv 1$. It has been shown that the inclusion of vertex corrections changes the position of the SiO₂ VBM by 0.6 eV while the gap is closed by 0.3 eV, so that I and A are reduced by 0.6 eV and 0.3 eV, respectively [32]. Therefore, a variation of the band edges of about 5–10% due to further many-body effects cannot be excluded.

The surface energy E_{SF} is defined as the energy cost for creating the surface. For stoichiometric cells with symmetric slabs with surface area A containing N bulk unit cells it can be easily calculated by

$$E_{\text{SF}} = \frac{E_{\text{slab}} - NE_{\text{bulk}}}{2A}, \quad (3)$$

where E_{slab} and E_{bulk} are the calculated total energies for the slab and the bulk unit cell respectively. For non-stoichiometric slabs we have to generalize the formalism

to include the chemical potential μ_{cat} and μ_{O} of the cation (i.e. In, Sn, or Zn) and oxygen atom, respectively [33]. For a slab with N_{cat} cations and N_{O} oxygen atoms Eq. (3) then turns into

$$E_{\text{SF}} = \frac{1}{2A} [E_{\text{slab}}(N_{\text{cat}}, N_{\text{O}}) - N_{\text{cat}}\mu_{\text{cat}} - N_{\text{O}}\mu_{\text{O}}]. \quad (4)$$

Since the bulk acts as reservoir, the chemical potentials of the atoms are related by

$$n_{\text{cat}}\mu_{\text{cat}} + n_{\text{O}}\mu_{\text{O}} = \mu_{\text{bulk}}, \quad (5)$$

where n_{cat} and n_{O} denote the number of cations and oxygen atoms per formula unit, respectively, $\mu_{\text{bulk}} = E_{\text{bulk}}/n_{\text{unit}}$ is the chemical potential per formula unit, and n_{unit} the number of formula units per bulk unit cell. Substituting Eq. (5) into Eq. (4), we obtain

$$E_{\text{SF}} = \frac{1}{2A} \left[E_{\text{slab}} - \frac{N_{\text{cat}}}{n_{\text{cat}}} \mu_{\text{bulk}} + \left(\frac{n_{\text{O}}}{n_{\text{cat}}} N_{\text{cat}} - N_{\text{O}} \right) \mu_{\text{O}} \right]. \quad (6)$$

This enables us to determine surface stabilities in dependence on the chemical potential of the oxygen atom, generally given in relation to μ_{O_2} , the chemical potential of the free O_2 molecule as obtained via a LDA total energy calculation of a free oxygen molecule.

3 Results and Discussion

We determine the ionization energies and electron affinities of the TCOs for different surface orientations and terminations. The results are listed in Table 1 and displayed in Fig. 2.

There are only a few measurements of the surface properties of In_2O_3 and Sn-doped- In_2O_3 (Indium-Tin Oxide, ITO). The electron affinity seems to vary in the range of $A = 4.1 \dots 5.0$ eV in dependence on the doping concentration (see Ref. [34] and references therein). Together with the measured gap of 3.6 eV, ionization energies of $I = 7.7 \dots 8.6$ eV may be derived. Klein [35] suggested values of $A = 3.5 \pm 0.2$ eV and $I = 7.1 \pm 0.15$ eV for evaporated In_2O_3 films. In a more recent paper [36] the same author gave values of $A = 4.45$ eV and $I = 8.05$ eV for ITO samples. Our theoretical values indicate a strong influence of the surface orientation and termination. The oxygen-terminated (001)-surface differs from the indium-terminated surface by more than 3 eV. To illustrate, the planar average of the respective electrostatic potentials are shown in Fig. 3a. The atomic configuration of the different surface terminations is shown in Fig. 3b, along with isosurfaces of the electrostatic potential, to illustrate the different surface barriers. This effect is due to the strong influence of the direction of the surface dipole in highly polar materials like the TCOs. Dangling bonds located at the oxygen atoms will most

Table 1 Characteristic energies: fundamental gap E_g , electron affinity A and ionization energies I of transparent conducting oxides derived from QP calculations. All values in eV. The surface orientation used for the calculation of I and A is indicated by the Miller indices (hkl) or $(hkil)$. Experimental values are given in parentheses

Crystal	Orientation	E_g	A	I
rh -In ₂ O ₃	(0001)	3.79 (3.02) ^a	6.11	9.41
bcc -In ₂ O ₃	(001) In-terminated	3.15	6.10	9.25
	(001) O-terminated		9.22	12.37
	(110)		5.30	8.45
rt -SnO ₂	(100)	(3.58) ^a 3.64	(3.5–5.0) ^b 4.10	(7.1–8.6) ^b 7.73
	(001) Sn-terminated		3.45	7.08
		(3.6) ^c	(4.44) ^d	(8.04) ^d
wz -ZnO	(0001) Zn-terminated	3.21	5.07	8.24
	(10 $\bar{1}$ 0)		3.65	6.84
		(3.38) ^e	(3.7–4.6, 4.05 4.42, 4.64) ^{d,f,g,h}	(7.1–8.0, 7.45, 7.82, 8.04) ^{d,f,g,h}

^aReferences [37, 38]

^bReferences [35, 36]

^cReference [46]

^dReferences [41, 42]

^eReference [45]

^fReference [43]

^gReference [39]

^hReference [40]




Fig. 2 Valence band (*red*) and conduction band (*green*) edges for (1) rh -In₂O₃(0001), (2) bcc -In₂O₃(001) O-terminated, (3) bcc -In₂O₃(001) In-terminated, (4) bcc -In₂O₃(110), (5) SnO₂(100), (6) SnO₂(001) Sn-terminated, (7) ZnO(0001) Zn-terminated, and (8) ZnO(10 $\bar{1}$ 0). The vacuum level is used as energy zero

likely increase polarity and, hence, the surface dipole of the slab. On the other hand, the In-termination will decrease the dipole, lowering the surface barrier for electrons. All in all, our predictions seem to overestimate the experimental




Fig. 3 (a) Planar average of the electrostatic potential $\tilde{V}(z)$ for the In-terminated (*black*) and O-terminated (*red*) $bcc\text{-In}_2\text{O}_3(001)$ surface. (b) Surface structure of the O-terminated (*above*) and In-terminated (*below*) $bcc\text{-In}_2\text{O}_3(001)$ surface. Isosurfaces of the spatially resolved electrostatic potential are shown for $V(\mathbf{x}) = -3 \text{ eV}$

findings. The discrepancies to the largest experimental values are of the order of 0.5 eV. Apart from uncertainties in the theoretical description, such as the neglect of vertex corrections in the GW approximation, several problems of the real-structure surfaces such as doping influence, surface coverage (and, hence, surface dipole), and sample quality may occur. Also the gap value of 3.6 eV taken from optical measurements deviates by 0.5 eV from the recently predicted one [37], mostly due to the fact that the lowest interband transitions are dipole-forbidden in the bixbyite structure [38].

There are several measurements regarding the surface barrier of $wz\text{-ZnO}$, which vary over a wide range. Jacobi et al. [39] found electron affinities of $A = 3.7, 4.5,$ and 4.6 eV depending on surface orientation and termination. Other electric measurements yield an electron affinity of $A = 4.64 \text{ eV}$ [40]. A value of $A = 4.05 \text{ eV}$ is derived from studies of the semiconductor-electrolyte interface [41] which yields $I = 7.45 \text{ eV}$ taking into account the known gap [42]. Another measurement gave $I = 7.82 \text{ eV}$ [43]. Again, it seems that we overestimate the measured values.

The surface properties of SnO_2 are hardly known. Measurements gave $A = 4.44 \text{ eV}$ [41] which, in combination with the gap of 3.6 eV measured for $rt\text{-SnO}_2$ [42], yields an ionization energy of $I = 8.04 \text{ eV}$. For tetragonal SnO_2 , a variation in the interval $I = 7.9\text{--}8.9 \text{ eV}$, depending on Sb doping, is reported [44]. SnO_2 is therefore the only TCO where our predictions seem to underestimate the experimental value. This might be due to a possible influence of virtual electronic states in the fundamental gap [8].

The large variety of measured values for I and A of the TCOs is probably not only due to sample quality problems, but also to the strong dependence of the surface barrier on surface orientation and termination. To investigate the influence of the surface preparation on the surface energy of the different orientations and terminations we calculate the surface energy of important surfaces for the bcc geometry of In_2O_3 .




Fig. 4 Surface energy E_{SF} per unit area of the non-polar (110) (blue), the In-terminated (001) (black), and the O-terminated (001) (red) surface of $bcc\text{-In}_2\text{O}_3$ in dependence of the chemical potential μ_{O} of the oxygen atom. The energy zero is set to half the chemical potential of a free O_2 molecule

The surface energy is plotted in dependence on the chemical potential of the oxygen atom in Fig. 4. For the polar (001) direction we see a strong dependence of the surface energy on the chemical potential μ_{O} of the O atom. While in the oxygen-rich limit of $\mu_{\text{O}} = 1/2\mu_{\text{O}_2}$ the O-termination is favored over the In-termination by a factor of 2, this changes drastically in oxygen-poor environments. The In-terminated surface becomes equally stable or even energetically favored over the oxidized surface. The non-polar (110) surface with a surface energy $E_{\text{SF}} = 0.11 \text{ eV}/\text{\AA}^2$ is energetically favored over both (001) phases in practically all settings. We therefore conclude, that the ionization energy and electron affinity of $bcc\text{-In}_2\text{O}_3$ strongly depends on the preparation conditions of the sample.

4 Summary and Conclusions

We have investigated the electronic properties of the transparent conducting oxides In_2O_3 , SnO_2 and ZnO by means of quasiparticle calculations based on many-body perturbation theory. The resulting band structures with rather accurate fundamental band gaps were combined with density functional theory calculations of material slabs to investigate electronic surface properties for different surface orientations and terminations. For this purpose the bulk and surface electronic structures have been energetically aligned using the electrostatic potential as reference.

The results were compared with the few experimental data available. Even though the experimental values are at variance, we found nevertheless a slight but systematic overestimation of ionization energy and electron affinity in our predictions. This is most likely due to the omission of vertex corrections in the

quasiparticle calculations. The only exception to this rule was SnO₂, possibly due to the influence of states within the fundamental gap at real surfaces. We found a strong dependence of the surface barrier on the orientation and termination of the surface. This is caused by the strong surface dipoles in highly ionic compounds.

We also analyzed the influence of the chemical potential of oxygen on the surface energy of different surface orientations and terminations of *bcc*-In₂O₃. We found that the surface stability, of different phases and, hence, the surface barrier of the sample, is strongly dependent on the preparation conditions.

Acknowledgements We gratefully acknowledge fruitful discussions with A. Schleife, F. Fuchs, S. Küfner, and J. Furthmüller. The work was financially supported by the EU e-I3 Project ETSF (Grant No. 211956) and the German Federal Government (BMBF Project No. 13N9669). Grants of computer time were given by the Höchstleistungsrechenzentrum Stuttgart.

References

1. D. S. Ginley and C. G. E. Bright, *MRS Bulletin* **25**, 15 (2000).
2. G. Thomas, *Nature* **389**, 907 (1997), ISSN 0028–0836.
3. K. Nomura, H. Ohta, K. Ueda, T. Kamiya, M. Hirano, and H. Hosono, *Science* **300**, 1269 (2003).
4. Z. M. Jarzebski, *Phys. Status Solidi A* **71**, 13 (1982), ISSN 1521–396X.
5. Y. Li, G. S. Tompa, S. Liang, C. Gorla, Y. Lu, and J. Doyle, in: *The 43rd national symposium of the American Vacuum Society*, vol. 15, pp. 1063–1068, AVS, Philadelphia, Pennsylvania (USA) (1997).
6. P. P. Edwards, A. Porch, M. O. Jones, D. V. Morgan, and R. M. Perks, *Dalton Trans.* pp. 2995–3002 (2004), ISSN 1477-9226.
7. R. W. Miles, G. Zoppi, and I. Forbes, *Mater. Today* **10**, 20 (2007), ISSN 1369-7021.
8. B. Höffling, A. Schleife, F. Fuchs, C. Rödl, and F. Bechstedt, *Appl. Phys. Lett.* **97**, 032116 (2010).
9. B. Höffling, A. Schleife, C. Rödl, and F. Bechstedt, *Phys. Rev. B* **85**, 035305 (2012).
10. F. Fuchs and F. Bechstedt, *Phys. Rev. B* **77**, 155107 (2008).
11. A. R. H. Preston, B. J. Ruck, L. F. J. Piper, A. DeMasi, K. E. Smith, A. Schleife, F. Fuchs, F. Bechstedt, J. Chai, and S. M. Durbin, *Phys. Rev. B* **78**, 155114 (2008).
12. A. Schleife, C. Rödl, F. Fuchs, J. Furthmüller, and F. Bechstedt, *Phys. Rev. B* **80**, 035112 (2009).
13. A. Schleife, J. B. Varley, F. Fuchs, C. Rödl, F. Bechstedt, P. Rinke, A. Janotti, and C. G. Van de Walle, *Phys. Rev. B* **83**, 035116 (2011).
14. F. Fuchs, J. Furthmüller, F. Bechstedt, M. Shishkin, and G. Kresse, *Phys. Rev. B* **76**, 115109 (2007).
15. P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).
16. W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
17. D. M. Ceperley and B. J. Alder, *Phys. Rev. Lett.* **45**, 566 (1980).
18. J. P. Perdew, *Electronic Structure of Solids '91*, p. 11, Akademie-Verlag, Berlin (1991).
19. G. Kresse and J. Furthmüller, *Comp. Mater. Sci.* **6**, 15 (1996).
20. G. Kresse and D. Joubert, *Phys. Rev. B* **59**, 1758 (1999).
21. H. J. Monkhorst and J. D. Pack, *Phys. Rev. B* **13**, 5188 (1976).
22. R. Leitsmann and F. Bechstedt, in: W. E. Nagel, D. B. Kröner, and M. M. Resch, eds., *High Performance Computing in Science and Engineering '10*, p. 135, Springer, Heidelberg (2010).

23. ICSD, Inorganics Crystal Structure Database, Fachinformationszentrum Karlsruhe, Germany (2002).
24. J. Haines and J. M. Léger, *Phys. Rev. B* **55**, 11144 (1997).
25. F. Bechstedt, F. Fuchs, and G. Kresse, *Phys. Status Solidi B* **246**, 1877 (2009).
26. J. Heyd, G. E. Scuseria, and M. Ernzerhof, *J. Chem. Phys.* **118**, 8207 (2003).
27. J. Heyd, G. E. Scuseria, and M. Ernzerhof, *J. Chem. Phys.* **124**, 219906 (2006).
28. A. V. Krukau, O. A. Vydrov, A. F. Izmaylov, and G. E. Scuseria, *J. Chem. Phys.* **125**, 224106 (2006).
29. J. Paier, M. Marsman, K. Hummer, G. Kresse, I. C. Gerber, and J. G. Ángyán, *J. Chem. Phys.* **124**, 154709 (2006).
30. M. Shishkin and G. Kresse, *Phys. Rev. B* **74**, 035101 (2006).
31. A. Schleife, C. Rödl, F. Fuchs, J. Furthmüller, F. Bechstedt, P. H. Jefferson, T. D. Veal, C. F. McConville, L. F. J. Piper, A. DeMasi, K. E. Smith, H. Lösch, R. Goldhahn, C. Cobet, J. Z. niga Pérez, and V. M. noz Sanjosé, *J. Korean Phys. Soc.* **53**, 2811 (2008).
32. R. Shaltaf, G.-M. Rignanese, X. Gonze, F. Giustino, and A. Pasquarello, *Phys. Rev. Lett.* **100**, 186401 (2008).
33. Y. Duan, *Phys. Rev. B* **77**, 045332 (2008).
34. A. Klein, C. Körber, A. Wachau, F. Säuberlich, Y. Gassenbauer, R. Schafrank, S. Harvey, and T. Mason, *Thin Solid Films* **518**, 1197 (2009), ISSN 0040-6090.
35. A. Klein, *Appl. Phys. Lett.* **77**, 2009 (2000).
36. Y. Gassenbauer and A. Klein, *The Journal of Physical Chemistry B* **110**, 4793 (2006).
37. P. D. C. King, T. D. Veal, F. Fuchs, C. Y. Wang, D. J. Payne, A. Bourlange, H. Zhang, G. R. Bell, V. Cimalla, O. Ambacher, R. G. Egdell, F. Bechstedt, and C. F. McConville, *Phys. Rev. B* **79**, 205211 (2009).
38. A. Walsh, J. L. F. Da Silva, S.-H. Wei, C. Körber, A. Klein, L. F. J. Piper, A. DeMasi, K. E. Smith, G. Panaccione, P. Torelli, D. J. Payne, A. Bourlange, and R. G. Egdell, *Phys. Rev. Lett.* **100**, 167402 (2008).
39. K. Jacobi, G. Zwicker, and A. Gutmann, *Surf. Sci.* **141**, 109 (1984), ISSN 0039-6028.
40. W. Göpel, L. J. Brillson, and C. F. Brucker, *J. Vac. Sci. Technol.* **17**, 894 (1980).
41. M. A. Butler and D. S. Ginley, *J. Electrochem. Soc.* **125**, 228 (1978).
42. C. Kilic and A. Zunger, *Appl. Phys. Lett.* **81**, 73 (2002).
43. W. Mönch, *Semiconductor Surfaces and Interfaces*, Springer, Berlin (2001).
44. C. Körber, P. Agoston, and A. Klein, *Sensors and Actuators B* **139**, 665 (2009), ISSN 0925-4005.
45. W. Martienssen and H. Warlimont, eds., *Springer Handbook of Condensed Matter and Materials Data*, Springer, Berlin (2005).
46. K. Reimann and M. Steube, *Solid State Communications* **105**, 649 (1998), ISSN 0038-1098.

Surface Magnetism: Relativistic Effects at Semiconductor Interfaces and Solar Cells

U. Gerstmann, M. Rohrmüller, N.J. Vollmers, A. Konopka, S. Greulich-Weber, E. Rauls, M. Landmann, S. Sanna, A. Riefer, and W.G. Schmidt

Abstract Ab initio calculations of the electronic g -tensor of paramagnetic states at surfaces and solar cells are presented, whereby special emphasis is given onto the influence of relativistic effects. After discussing the numerical requirements for such calculations, we show that for silicon surfaces the g -tensor varies critically with the hydrogen coverage, and provides an exceptionally characteristic property. This holds also in the case of powder spectra where only the isotropic part g_{av} is available from experiments. Extending our calculations onto microcrystalline 3C-SiC, our study explains why sol-gel grown *undoped* material can serve as an excellent acceptor material for an effective charge separation in organic solar cells: Due to an auto-doping mechanism by surface-induced states it fits excellently into the energy level scheme of this kind of solar cell and has the potential to replace the usually used rather expensive fullerenes.

1 Introduction

Solar cells provide an increasing market with high potential for a further development. The global market for photovoltaics cells is expected to be doubled during the next decade. However, overcapacities will be an ongoing challenge for the manufacturers and a wedding-out and consolidation process seems to be unavoidable. One way out is the production of highly efficient solar cells by minimal costs. So far, however, such an optimization of the cells is mainly based on try and error. For a further improvement of cell performance a better understanding of the microscopic structures and the basic electronics behind the light-induced separation of charge carriers as well as the efficiency limiting processes is crucial.

U. Gerstmann · M. Rohrmüller · N.J. Vollmers · A. Konopka · S. Greulich-Weber · E. Rauls · M. Landmann · S. Sanna · A. Riefer · W.G. Schmidt (✉)
Lehrstuhl für Theoretische Physik, Universität Paderborn, 33095 Paderborn, Germany
e-mail: W.G.Schmidt@upb.de

Experimentally, electron paramagnetic resonance (EPR) provides a powerful tool to analyse the microscopic structure of paramagnetic systems. Since most of the electrically active centers in solar cells are those with unpaired electrons, these electronic states are paramagnetic. Hence, EPR provides an appropriate possibility to characterize the basic material as well as the solar cells itself. In many cases, however, the wealth of important information available from EPR measurements, cannot be extracted from experimental data alone. For an identification of the microscopic structures, accurate first principles calculations of as many as possible relevant properties are necessary to calculate a fingerprint of the structure that can be compared with the experimental values. From EPR experiments the components of the electronic g -tensor are available also in those cases in which hyperfine (hf) splittings cannot be resolved. However, in contrast to the *ab initio* calculation of hf splittings that already do have an appreciable history, quantitative predictions of electronic g -tensors making use of the machinery of *ab initio* density functional theory (DFT) have become possible only very recently [1]. In semiconductors, this has been already demonstrated successfully for defects in SiC and GaN bulk material [2–4]. For surfaces, however, theoretical data obtained by *first principle* calculation is very rare. In this work we show, that the EPR parameters are mainly influenced by relativistic effects like the spin-orbit coupling (SOC). At surfaces and interfaces these effects become exceptionally anisotropic giving rise to special effects like the so-called Rashba effect [5].

We evaluate our method using hydrogenated silicon surfaces as an example. Such surface states appear in hydrogenated microcrystalline silicon (μc -Si:H), a material that can be used for efficient and low-cost solar cells [6] (see Fig. 1).¹ We show that the *ab initio* calculation of g -tensors can help to elucidate the situation in such microcrystal. We calculate the elements of the electronic g -tensor for some paramagnetic states at silicon surfaces from first principles, using a recently developed gauge-including projector augmented plane wave (GI-PAW) approach [1, 9] in the framework of DFT. According to the in-diffusion of water and atmospheric gases [10] we investigate the EPR fingerprint of those paramagnetic states that are created by hydrogen adsorbed at Si(111) and Si(001) surfaces. Our approach is shown to be able to distinguish between different surface states [11]. For silicon surfaces with different hydrogen coverages, the g -tensor is by far more characteristic than the hf splitting of the Si dangling bonds or that of the adsorbed H atoms. This holds in cases of powder spectra like in the case of amorphous or microcrystalline material for solar cells where only the angular average of the spectra is available experimentally.

A central challenge on the way to optimized solar cells is to make the thickness of the individual layers smaller than the diffusion length of the charge carriers. Recently, 3C-SiC microcrystals grown by a sol-gel based process have been

¹In comparison with cells based on amorphous silicon suffer less from the notorious light-induced degradation, known as the Staebler-Wronski effect [7]. Best cell performance is, however, achieved for material grown close to the transition to amorphous growth [8].

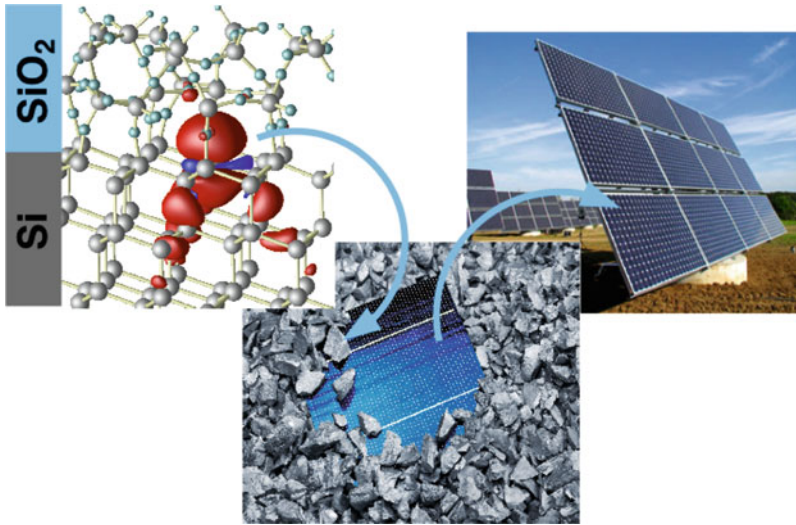


Fig. 1 From microcrystalline material (here: atomic structure of oxidized $\mu\text{c-Si}$) towards single solar cells and photovoltaic plantages

proposed as a promising acceptor material for photovoltaic applications [27]. Such $\mu\text{c-SiC}$ samples have been already characterized by optical spectroscopy and electron paramagnetic resonance (EPR) [32]. In this work, the available experimental data is analyzed with the help of ab initio DFT calculations resulting in electronic band structures and g -tensors. Based on this, a possible scenario for the observed acceptor process is discussed.

2 Methodology

Our first-principles calculations of the EPR parameters are based on density functional theory (DFT) using the generalized gradient approximation for the electron exchange and correlation functional (GGA-PBE) in its spin-polarized form [12]. The hyperfine splittings, i.e. the interaction of the magnetic moments of the electrons with those of the nuclei, are calculated taking into account relativistic effects in scalar-relativistic approximation [13, 14]. Although there exist a non-relativistic derivation for the isotropic contact interaction by Fermi [15], Breit has shown that the origin of the hyperfine splitting can be only described correctly in a relativistic treatment [16]. The static magnetic field caused by the magnetic moment $\boldsymbol{\mu}_I = g_N \mu_N \mathbf{I}$ of a nucleus with gyromagnetic ratio g_N located at the origin is included using the vector potential for this magnetic field

$$\mathbf{A}_I(\mathbf{r}) = \nabla \times \left(\frac{\boldsymbol{\mu}_I}{r} \right) \quad (1)$$

replacing the momentum operator \mathbf{p} by the canonical momentum $\boldsymbol{\pi} = \mathbf{p} + e/c\mathbf{A}_I$ in Dirac's equation

$$(c\boldsymbol{\alpha} \cdot \boldsymbol{\pi} + \beta mc^2 + V - E_{rel}) \Psi = 0. \quad (2)$$

The influence of the resulting magnetic fields $\mathbf{B}_I(\mathbf{r}) = \nabla \times \mathbf{A}_I(\mathbf{r})$ leads to level splittings in the $10^{-12} \dots 10^{-2}$ eV range. The smallness of these splittings allows a simplified computation via perturbation theory. Within first order perturbation theory the expectation value of the hyperfine interaction is given by

$$E_{\text{HF}} = -e \langle \Psi | \boldsymbol{\alpha} \cdot \mathbf{A}_I(\mathbf{r}) | \Psi \rangle, \quad (3)$$

Here, $\boldsymbol{\alpha}$ is a 4×4 matrix constructed from the 2×2 Pauli spin matrices σ_x , σ_y , and σ_z , respectively, whereby $|\Psi\rangle = \begin{pmatrix} \Phi_L \\ \Phi_S \end{pmatrix}$, is given by the Dirac spinor decomposing into the two-component Pauli spinors Φ_L and Φ_S . For light atoms, Φ_L is the large component whereas Φ_S turns out to be small. This leads to

$$E_{\text{HF}} = -e (\langle \Phi_L | \boldsymbol{\sigma} \cdot \mathbf{A}(\mathbf{r}) | \Phi_S \rangle + \langle \Phi_S | \boldsymbol{\sigma} \cdot \mathbf{A}(\mathbf{r}) | \Phi_L \rangle). \quad (4)$$

Thus, E_{HF} is a genuine relativistic term that couples large and small components of Dirac's equation. The small component Φ_S can be expressed in terms of the large component Φ_L as

$$\Phi_S = \frac{c \boldsymbol{\sigma} \cdot \mathbf{p}}{2mc^2 + E - V(\mathbf{r})} \Phi_L = \frac{S(\mathbf{r})}{2mc^2} (\boldsymbol{\sigma} \cdot \mathbf{p}) \Phi_L, \quad (5)$$

whereby $S(r)$ is the inverse relativistic mass correction. This can be used to express Φ_S in terms of Φ_L , leading to expectation values containing the large component exclusively [17]. In the case of orbital quenching we obtain [14]:

$$E_{\text{contact}} = -\frac{8\pi}{3} \mu_B \langle \Phi_L | S(\mathbf{r}) \boldsymbol{\mu}_I \cdot \boldsymbol{\sigma} \delta(\mathbf{r}) | \Phi_L \rangle + \left\langle \Phi_L \left| \frac{1}{r^4} \frac{\partial S}{\partial r} [\boldsymbol{\mu}_I \cdot \boldsymbol{\sigma} r^2 - (\boldsymbol{\mu}_I \cdot \mathbf{r}) (\boldsymbol{\sigma} \cdot \mathbf{r})] \right| \Phi_L \right\rangle \quad (6)$$

$$E_{\text{dip}} = \mu_B \left\langle \Phi_L \left| \frac{1}{r^5} [\boldsymbol{\sigma} \cdot \boldsymbol{\mu}_I r^2 - 3 (\boldsymbol{\sigma} \cdot \mathbf{r}) (\boldsymbol{\mu}_I \cdot \mathbf{r})] \right| \Phi_L \right\rangle \quad (7)$$

The dipolar term E_{dip} is angular dependent and, thus, in the general case gives rise to anisotropic hf tensors. In the non-relativistic case, since $S(\mathbf{r}) \rightarrow 1$, only the first term in (6) contributes to the *isotropic* part, the so-called contact term. By this,

we obtain the results of the classical theory given by Fermi [15], that only the probability amplitude at the nucleus contributes. In the relativistic case, however, this first term does not contribute at all. It is second term in (6) which becomes the relativistic analogue to the contact interaction. For a pure Coulombic potential

$$V(r) = \frac{-Ze^2}{r} \quad (8)$$

the derivative $\partial S(r)/\partial r$ is similar to a broadened δ -function

$$\delta_{\text{Th}}(\mathbf{r}) = \frac{1}{4\pi r^2} \frac{\partial S}{\partial r} = \frac{1}{4\pi r^2} \frac{r_{\text{Th}}/2}{\left[\left(1 + \frac{E}{2mc^2} \right) r + r_{\text{Th}}/2 \right]^2}. \quad (9)$$

In other words, the magnetization density of the electron in the relativistic theory is not evaluated at the origin, where it would be divergent for s electrons, but is averaged over a sphere of radius

$$r_{\text{Th}} = \frac{Ze^2}{mc^2} \quad (10)$$

which is the Thomson radius, about ten times the nuclear radius.

As a result, the divergence of the s electrons presents no problem. Also if we approximate the nuclear potential by that of a charged volume rather than that of a point charge [18], the divergence already disappears. However, it is important to note, that we would obtain divergent contact terms mixing the approximations, e.g. using (scalar)² relativistic orbitals in a non-relativistic formula.

Whereas the hf splittings depend on the magnetization density $m(\mathbf{r}) = n^\uparrow(\mathbf{r}) - n^\downarrow(\mathbf{r})$ exclusively, the main deviation of the g -tensor from its free electron value $g_e \approx 2.002\ 319\ 278$ is given by the spin-orbit coupling of the many-particle system. In physically transparent form it can be written in terms of spin-polarised electron currents $\mathbf{j}_\uparrow^{(1),\mu}$ and $\mathbf{j}_\downarrow^{(1),\mu}$ induced by a unit magnetic field \mathbf{B}^μ applied along the direction μ :

$$\Delta g_{\mu\nu}^{\text{SO}} = \frac{\alpha}{2S} \left[\sum_{\sigma=\pm 1(\uparrow,\downarrow)} \sigma \int \nabla V_{\text{eff}}^\sigma \times \mathbf{j}_\sigma^{(1),\mu}(\mathbf{r}) d^3r \right]_\nu. \quad (11)$$

To obtain this result we start again from Dirac's equation and apply perturbation theory with respect to spin-orbit coupling and to an external magnetic field \mathbf{B} .

² In the scalar relativistic treatment Φ_L is calculated solving Dirac's equation but thereby ignoring spin-orbit interactions. This leaves the electron spin as a "good" quantum number. Already in a scalar relativistic treatment, s -like wave functions diverge at the nuclear site (if the nucleus is taken to be a point charge).

α is the fine structure constant, S the total spin given by the number of unpaired electrons times 1 s. $\nabla V_{\text{eff}}^\sigma$ denotes the gradient of the spin-polarized effective potential. Besides ground state quantities the evaluation of the g -tensor requires the calculation of the spin-currents $\mathbf{j}_\sigma^{(1),\mu}$ in linear magnetic response [9]:

$$\mathbf{j}_\sigma^{(1),\mu}(\mathbf{r}) = 2 \sum_o \text{Re} \left\langle \psi_{(0),o}^\sigma | \mathcal{J}^p \mathcal{G}^\sigma(\varepsilon_o) \mathcal{H}_{(1)} | \psi_{(0),o}^\sigma \right\rangle + \frac{\alpha}{2} n^\sigma(\mathbf{r}) \cdot \mathbf{B}^\mu \times \mathbf{r} \quad (12)$$

$\mathcal{J}^p = \frac{1}{2} \frac{\nabla}{i} |\mathbf{r}\rangle \langle \mathbf{r}| + c.c.$ denotes the current operator for vanishing magnetic field. $\mathcal{H}^{(1)} = \frac{\alpha}{2} \mathbf{L} \cdot \mathbf{B}^\mu$ describes the influence of the *uniform* magnetic field determining the perturbed wavefunctions $|\psi_{(1),o}^\sigma\rangle = \mathcal{G}^\sigma(\varepsilon_o) \mathcal{H}_{(1)} |\psi_{(0),o}^\sigma\rangle$ via Green's function of the unperturbed system

$$\mathcal{G}^\sigma(\varepsilon) = \sum_e \frac{|\psi_{(0),e}^\sigma\rangle \langle \psi_{(0),e}^\sigma|}{\varepsilon - \varepsilon_e^\sigma}, \quad (13)$$

whereby the sum runs over the *empty* orbitals e .

Strictly seen, the formalism so far presented ensures only a faithful description of the nuclear and two-electron spin-orbital coupling. According Ref. [1, 11] higher order contributions can be approximately taken into account via the *spin-other-orbit correction*, given by the screening $\mathbf{B}^{(1),\mu}(\mathbf{r})$ of the external magnetic field \mathbf{B}^μ by the induced currents as experienced by the magnetization density $m(\mathbf{r})$ of the unpaired electrons:

$$\Delta g_{\mu\nu}^{\text{SOO}} = \frac{g_e \alpha}{2S} \int \mathbf{e}_\nu \cdot \mathbf{B}^{(1),\mu}(\mathbf{r}) m(\mathbf{r}) d^3r. \quad (14)$$

In the case of the paramagnetic states at Si surfaces the contribution of the SOO term comes out to be very small (clearly below 10 ppm). In other words, the spin-other-orbit contributions do not contribute considerably to the g -values given in Tables 1–4.

For a modelling of the surfaces we use supercells and periodic boundary conditions. Hence, the explicit treatment of an external magnetic field \mathbf{B} has to be done in a gauge-invariant way in order to retain the translation invariance of the wavefunctions. Here, the gauge-including projector augmented plane wave (GI-PAW) approach satisfies this requirement and allows an ab initio calculation of the all-electron magnetic response using an efficient pseudopotential approach [1, 9]. The GI-PAW approach is implemented in the pwscf-code (QUANTUM-ESPRESSO package) [19] and has been already applied successfully to identify paramagnetic defect structures in SiC and GaN [2–4].

To model the semiconductor surfaces (silicon and SiC) at least eight atomic layers are treated in a supercell, whereby the lowest layer was saturated with H atoms. To ensure a well defined transition to bulk material, the atoms in the lowest three layers were kept on their ideal bulk positions. All other atoms were allowed to relax freely. To minimize the interaction of the periodic images of the surface, 10 Å vacuum is inserted. We use supercells containing up to 175 atoms and

Table 1 Largest hf splittings (MHz) of a H vacancy at Si(111):H surface as calculated by ab initio DFT. θ denotes the angle between the principal axis of the hf tensor and the surface normal. All hf splittings due to H atoms are below 2 MHz

	# nuclei	A_1	A_2	A_3	θ
Si_1	1	-220.0	-220.0	-414.9	0.0°
Si_2	3	1.4	0.0	-6.7	52.1°
Si_3	3	-26.5	-27.4	-42.8	0.4°
Si_4	3	-20.1	-20.4	-26.7	66.3°
Si_5	3	-7.5	-7.7	-10.4	10.0°

Table 2 k -point convergence of the g -values calculated for the H vacancy at Si(111):H surface. Given are the principal values g_i as well as the angular averaged value g_{av} . θ denotes the angle between the principal axis g_3 and the surface normal

k -point mesh	g_{av}	g_1	g_2	g_3	θ
Γ	2.010887	2.01250	2.01250	2.00766	0.0°
$2 \times 2 \times 1$	2.006393	2.00925	2.00925	2.00068	0.0°
$3 \times 3 \times 1$	2.006607	2.00939	2.00939	2.00104	0.0°
$4 \times 4 \times 1$	2.006630	2.00939	2.00939	2.00111	0.0°
$5 \times 5 \times 1$	2.006630	2.00939	2.00939	2.00111	0.0°
$6 \times 6 \times 1$	2.006630	2.00939	2.00939	2.00111	0.0°

Table 3 Comparison of the calculated g -tensors for a single adsorbed H-atom and for a nearly complete H-coverage of the Si(001) surface. For a better comparison with the calculated hf splittings (MHz) the calculated Δg -values are here given in ppm

	A_1	A_2	A_3	θ_{hf}	Δg_{av}	Δg_1	Δg_2	Δg_3	θ
(111): H saturated with H vacancy	-220	-220	-415	0°	4,309	7,070	7,070	-1,213	0.0°
(001): single adsorbed H atom	-189	-189	-373	18°	-2,149	-2,079	-219	-4,139	27.8°
(001): monohydride with H vacancy	-254	-2,554	-450	20°	2,941	5,841	3,431	-449	33.5°

Table 4 Analysis of 3C-SiC microcrystals: Calculated g -tensor values for different surface related defects visualized in Fig. 5. The value experimentally observed for a μc -SiC powder spectrum and the corresponding angular averaged theoretical values g_{av} are also given

	Defect	g_{av}	g_1	g_2	g_3
	<i>Exp.</i>	2.0073			
(a)	Si(001):H	2.00320	2.00187	2.00359	2.00415
(b)	C(001):H	2.00292	2.00245	2.00273	2.00357
(c)	Si(001) + H	2.00268	2.00238	2.00271	2.00294
(d)	C(001) + N_C	2.00271	2.00180	2.00299	2.00333

norm-conserving pseudopotentials with a plane wave energy cutoff of 30 and 50 Ry for silicon and SiC, respectively. The ab initio calculation of the g -tensor is still very time-consuming: Whereas for the geometry optimisation a $2 \times 2 \times 1$ Monkhorst-Pack (MP) [20] k -point set comes out to be sufficient, to obtain converged estimates for the g -tensor in the general case $4 \times 4 \times 1$ samplings come out to be unavoidable. In some cases the number of k -points can be reduced by symmetry, but the number of $(k+q)$ -points has to be multiplied by a factor of 7 in order to obtain the derivatives in

the reciprocal space. A second requirement for the calculation of the spin-currents in linear magnetic response [9] is calculation of the Green's function (see Eq. 14). As a result, the calculation of a g -tensor becomes computationally extremely demanding and takes about an order of magnitude more CPU time than structure optimization. On the other hand, the calculation of the hyperfine splittings can be done on the fly. Given that structure optimization takes 1 day CPU time, and the g -tensor more than 10 days, the hf splittings are already available after 10 CPU minutes.

3 Results

We first discuss the H vacancy at a Si(111):H surface as a reference system. This structure is obtained if removing one H-atom from an otherwise completely hydrogenated Si(111) surface. It provides a simple model for a single paramagnetic dangling bond. Figure 2 shows several views of the microscopic structure and the corresponding magnetisation density $m(\mathbf{r})$. As can be seen from the top view (*lower left corner*), the structure shows perfect C_{3v} symmetry with the symmetry axis along the surface normal. An analysis of $m(\mathbf{r})$ at the nuclei leads to the hyperfine splittings given in Table 1. As intuitively expected, the by far largest hf splitting (-415 and -220 MHz for the magnetic field along and perpendicular to the surface normal, respectively) is due to the unsaturated Si-atom at the surface. The right part of Fig. 2 shows the magnetization density in a plane parallel to the surface normal. It can be considered as a typical 'textbook' fingerprint of a dangling bond. As an additional feature, weaker accumulations of $m(\mathbf{r})$ are found along the Si zig-zag line. As a result, besides that of the dangling bond nucleus itself, our ab initio calculations predict further characteristic hf splitting. With a value of about -43 MHz due to three equivalent nuclei Si_3 in the third layer, the hf splitting could be large enough to be resolved in EPR measurements. In contrast, the hf splittings below 10 MHz (see Table 1), especially that of the H atoms at the Si(111) surface (below 2 MHz) are too small to be resolved. They will contribute to the width of the central line instead. The position of this central line is determined by the g -tensor.

In Table 2, the calculated principal g -values are compiled for different k -point samplings. At least for the $4 \times 4 \times 1$ and larger samplings the values can be considered converged. The vanishing angle θ between g_3 and the surface normal confirms again the perfect alignment of the dangling bond along the surface normal. Perpendicular to the surface normal with $g_1 = g_2 = 2.00939$ comparatively large g -values are predicted. The g -value parallel to the surface normal remains similar to that of the free electron. The reason for this particular anisotropic shape is the perfect alignment of the dangling bond along the surface normal resulting in a strongly anisotropic spin-orbit coupling similar to the Rashba-type [5] characterized by vanishing spin-orbit coupling for the spin along the surface normal.

We also take this textbook dangling bond as an example to evaluate the computational efficiency of our code in massive parallel application. We analyse the hydrogen vacancy at the H-terminated Si(111):H surface in a 119 atom supercell by





Fig. 2 Microscopic structure and magnetisation density of the paramagnetic H vacancy at Si(111):H surface for different side views (*lower left corner*: top view). Plot of the magnetisation density within a plane including the surface normal. The *arrows* describe the direction of the magnetization density (i.e. the principal axis of the hf tensor) at the Si-nuclei

using a $6 \times 6 \times 1$ k -point set. Due to the trigonal C_{3v} symmetry this results in an explicit treatment of 24 k -points. In the standard mode, the QUANTUM-ESPRESSO code uses parallelization with respect to reciprocal lattice vectors. For Fast Fourier Transformations (FFT) and a limited number of processors, already this simple treatment often provides a good scaling with the number of quasi-processors (cores) used. However, especially in cases where the number of lattice vectors does not fit to the number of cores and for more than 32 cores the performance of the calculations becomes more and more saturated (see Fig. 3). We cannot exclude that the observed saturation is also supported by additional communication between the several nodes as unavoidable for $N_{cores} > 32$. However, for band structure calculations, in general for calculation with a large number of k -points³ a second parallelization via pools can be established. Here, the k -points are divided into pools whereby each pool contains a subset of the k -points. By this, for the present architecture we obtain linear scaling for up to 256 processors (16 nodes).

The number of pools is obvious limited by the number of k -points in the calculation and more critical by the memory consumption per core: By doubling the number of pools the required memory per core is increased by a factor of 1.2–1.4. The best scaling, however, is obtained for metallic systems with several thousands of k -points. On the other hand, for a small number of cores per pool the parallelization via k -points can become counter-productive. For a given system size (including cell size, number of electrons/bands, number of k -points) and for a chosen number of cores an optimal number of pools exists, mostly larger than the available memory allows. In other words, the overall limiting factor is given by the memory available

³For spin-polarized calculations, the second spin channel is realized by doubling the k -point set.

Fig. 3 CPU time on the HLRS CRAY XE6 for a self-consistent field calculations (*blue*) and g -tensor calculations (*red*). For exact description of the used parameters see text. The number of pools is marked by the symbols plus, cross, star, square corresponding to 1, 2, 4 and 8 pools, respectively



per core. Here, for future work, especially for larger systems with up to thousand atoms, an architecture with more than 2 GB RAM per core would be desirable.

Coming back to the physical results of our calculations, the situation becomes more difficult in case of the Si(001) surface that has an appreciable history of both experimental and theoretical work (see e.g. Ref. [21] for a review). It shows the famous 2×1 reconstruction into rows of buckling Si dimers. The left part of Fig. 4 shows a side view along these dimer rows. The adsorption of a single H atom breaks the double bond of a dimer. The result is a single dangling bond (left row but one in Fig. 4). If further H atoms are adsorbed at the Si(001) surface, either further dimers are broken or existing dangling bonds are saturated. In the case of complete saturation, each Si atom at the surface bonds one H atom. In Table 3, the EPR parameters for such microstructures are compared. Obviously, the g -tensor varies strongly with the hydrogen coverage. Especially the Δg_3 -values along the principal axis of the g -tensor differ by more than one order of magnitude. In this sense, the g -tensor is by far more characteristic than the hf splittings which vary only within 20%. Since in contrast to the hf splittings the sign of Δg can be determined experimentally, this holds also in case of powder spectra where only the isotropic part g_{av} is available from experiment.

The results and experiences obtained for the H-terminated silicon surfaces gives us confidence that our method will be accurate enough to predict the spectroscopic magnetic properties of dangling bonds in real devices, like solar cells based on microcrystalline SiC:

Microcrystalline silicon carbide ($\mu\text{c-SiC}$) has become an attractive new class of advanced materials for light emitting diodes and heterojunction photovoltaic devices [24]. Here, the microcrystallites are of interest as effective charge carrier collectors in organic solar cells. When a photon is absorbed by an organic photoactive material, an exciton, i.e., a bound state of an electron and a hole, is created. Due to the notably short exciton lifetime of several tens of nanoseconds, the most important design criterion for such solar cells is to make the thickness




Fig. 4 Structure and magnetisation densities for a single adsorbed H-atom (*left*) and for a nearly complete H-coverage of the Si(001) surface (H vacancy at monohydride surface, *right*)

of the individual layers smaller than the diffusion length of the exciton [22, 23]. For this purpose, usually fullerene molecules are dispersed in a polymer matrix. One possibility to further increase the efficiency of this kind of solar cells is to use alternative acceptor materials with a more suitable position of the LUMO (lowest unoccupied molecular orbital) level. By adjusting the latter, the open circuit voltage of the solar cell can be increased at the expense of the energy lost by the electrons connected with charge transfer from donor to acceptor material. Instead of using rather expensive fullerenes, currently wide bandgap micro- and nano-crystals as acceptor materials are under discussion, such as TiO_2 [25], ZnO [26], and last but not least SiC [27].

Usually, the mentioned materials provide unavoidable technical difficulties, such as intrinsic defects, the lack of suitable *p*-doping [28] or omnipresent *n*-doping. As an promising alternative, in Ref. [31] a sol-gel growth process for $\mu\text{c-SiC}$ has been proposed. The resulting material is almost free from usually unavoidable nitrogen donors, allowing arbitrary doping. In Ref. [32] electron paramagnetic resonance (EPR) was used as an analytic tool for the control of the doping success: Doping with N, Al, and P leads to different, characteristic EPR spectra. They are clearly different from those known for usual shallow donors and acceptors in bulk SiC . The nominally undoped and nitrogen-doped samples show an EPR line similar to a center discovered in porous 3C-SiC, assigned to a carbon dangling bond at the 3C-SiC/ SiO_2 interface. The *g*-factor is slightly different, but the line half width is almost the same. At first view, this similarity is very surprising since the microcrystallites are not oxidized: Even if the crystals became as large as $20\ \mu\text{m}$ in diameter, they do not show the typical SiO_2 lines in nuclear magnetic resonance and infrared reflection spectroscopy. Electrical and photoluminescence measurements support the finding that the required acceptor behavior of $\mu\text{c-SiC}$ is caused by surface-related defects in combination with an appropriate position of the Fermi level, which is determined by doping.

Based on the experimental results, the microscopic structure of the responsible defect structure at the clean surface of the micro-crystallites and its influence onto the charge-separation mechanism is discussed with the help of *ab initio* calculations. In order to elucidate the microscopic origin of the observed EPR signals, we calculate the EPR parameters for some possible dangling-bond related structures

and compare them with the experimental values. For a modelling of the surfaces we use the settings already mentioned in connection with the partially hydrogenated silicon surfaces. In close analogy to the (001)-oriented silicon surfaces we first discuss the different defects at the corresponding surface of 3C-SiC (see Fig. 5). In silicon, the clean (001) surface is stabilized by rows of buckled dimers, which are still surviving if the surface is partially hydrogenated. Only, if nearly all silicon atoms are mono-hydrogenated, the dimers at the reconstructed surface lose their buckling. In 3C-SiC, the silicon-terminated Si-surface looks similar at first view. But due to the smaller bond length in SiC the buckling becomes less efficient, resulting in almost all cases in surfaces with very complicated reconstructions [29, 30] showing metallic and diamagnetic properties. Hence, these configurations cannot be responsible for the observed paramagnetic structures.

Only an almost completely hydrogen-passivated configuration with one missing hydrogen (shown in Fig. 5a) leads to a paramagnetic surface state. In the following, we focus onto the carbon-terminated C-surface, where the situation is by far more straightforward. Here, similar to the case of diamond, already the clean surface provides a 2×1 dimer reconstruction without buckling (see also Fig. 5b–d). All these surface related defects provide paramagnetic states and are, thus, possible candidates to explain the $g = 2.0073$ EPR-signal of unknown origin. In Table 4, the calculated elements of the corresponding g -tensor are listed. Since no angular dependent EPR-measurements for the microcrystalline SiC-powder were possible, the only value that has to be compared to the experiment is g_{av} , the average g -value over possible orientations. In all cases, Si-related hyperfine splittings in the range 20–30 MHz can be found, but the calculated average g -values g_{av} are far away from the experimental data. So, the corresponding models have to be discarded as explanation for the surface related EPR-signal.

Nevertheless, it is worth to check the energetic level of the surface related defects. For this, we have to analyse the corresponding electronic band structures. Figure 6 (*middle*) shows the band structure of the clean C-terminated (001) surface, calculated using the gradient-corrected PBE functional. In comparison with the SiC-bulk material (*left part*) additional broad bands appear in the gap. Those in the lower part are occupied covering a 1.2 eV broad region. The unoccupied bands overlap with the conduction bands of the bulk material, so that the position of the LUMO is lowered. In total, the fundamental gap is considerably reduced to 0.2 eV. Having in mind the well-known underestimation of band gaps by the local density approximation (LDA) or the gradient-corrected functionals like PBE (e.g. the fundamental gap comes out about 1.4 eV while the experimental value is 2.4 eV for 3C-SiC), the experimental values should be larger. To determine the exact positions of the band edges further elaborated calculations (e.g. hybrid functionals) would be necessary. Nevertheless, the LDA prediction of the gap considerably reduced from both the valence band as well as from the conduction band can be considered at least qualitatively correct. The result is something like an auto-doping of the 3C-SiC microcrystals. The additional unpaired electron, introduced by a paramagnetic surface structure, leads only to an additional energy level clearly below the highest occupied surface state, but leaves all other features of the band structure




Fig. 5 Magnetization density of different surface-defects at the (001) surface of 3C-SiC. (a) Si-terminated surface, hydrogen passivated but with one missing H-atom Si(001)-H, (b) the corresponding defect at the C-terminated surface C(001)-H, (c) C-terminated surface with one H-atom adsorbed C(001)+H, (d) C-terminated surface with a substitutional nitrogen atom on a C-site C(001)+N_C

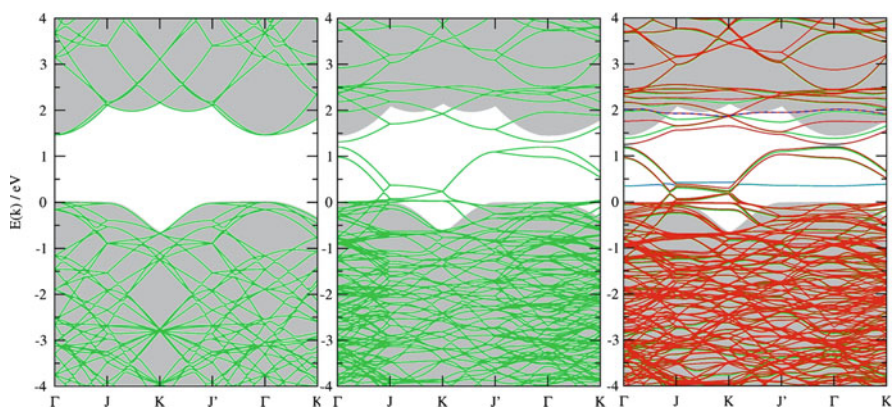


Fig. 6 Calculated band structure of 3C-SiC bulk (*left*) and for a 2×1 dimer-reconstructed SiC (001) C-surface (*middle*). The remarkable reduction of the fundamental gap (*shaded grey*) by the surface leads to an efficient auto-doping. Additional localized defects at the surface (*right*) induce additional defect levels (*blue*) and give rise to a spin-polarized band structure, but do not change the situation with respect to the LUMO considerably

of the clean surface unchanged. In all cases, the gap is significantly reduced, and the pure existence of the surface leads to an efficient auto-doping. In other words, independent from details of the paramagnetic defect structure, already the nominally undoped micro-crystalline sol-gel 3C-SiC behaves as an efficient acceptor for the charge carriers collection.

4 Conclusions

Ab initio calculations of the electronic g -tensor of paramagnetic states at surfaces and solar cells are presented. After discussing the numerical requirements for such calculations, we show that for silicon surfaces the g -tensor varies critically with the hydrogen coverage, and is by far more characteristic than the hf splitting of the Si dangling bonds or the adsorbed H atoms. This holds also in the case of powder spectra where only the isotropic part g_{av} is available from experiments. Extending our calculations onto micro- and nano-crystalline 3C-SiC as a basic material for solar cells, our study shows that sol-gel grown material can serve as an excellent acceptor material for an effective charge separation in organic solar cells. It fits excellently into the energy level scheme of this kind of solar cell and has the potential to replace the usually used, rather expensive fullerenes. It turned out, that already undoped μc -3C-SiC acts as a particular suitable acceptor due to its auto-doping mechanism by a surface-induced band structure.

Acknowledgements We gratefully acknowledge financial support from the DFG as well as supercomputer time provided by the HLRS Stuttgart and the Paderborn PC².

References

1. Ch.J. Pickard and F. Mauri, Phys. Rev. Lett. **88**, 086403 (2002).
2. U. Gerstmann, E. Rauls, S. Greulich-Weber, E.N. Kalabukhova, D.V. Savchenko, A. Pöppel, and F. Mauri, Mat. Sci. Forum **556–557**, 391 (2006).
3. U. Gerstmann, A.P. Seitsonen, F. Mauri, and J. von Bardeleben, Mat. Sci. Forum **615–617**, 357 (2009).
4. U. Gerstmann, A.P. Seitsonen, and F. Mauri, phys. stat. sol. (b) **245**, 924 (2008).
5. E.I. Rashba, Fiz. Tverdogo Tela **2**, 1224 (1960); Sov. Phys. Solid State **2**, 1109 (1960).
6. E. Vallet-Sauvain, U. Kroll, J. Meier, A. Sah, and J. Pohl, J. Appl. Phys. **87**, 3137 (2000).
7. D.L. Staebler, C.R. Wronski, Appl. Phys. Lett. **31**, 292 (1977).
8. S. Klein, F. Finger, R. Carius, T.Dylla, B. Rech, M. Grimm, L. Houben, and M. Stutzmann, Thin Solid Films **430**, 202, (2003).
9. Ch.J. Pickard and F. Mauri, Phys. Rev. B **63**, 245101 (2001).
10. F. Finger, R. Carius, T. Dylla, S. Klein, S. Okur, and M. Günes, IEE Proc.-Circuits Devices Syst. **150**, 309 (2003).
11. U. Gerstmann, M. Rohrmüller, F. Mauri, and W.G. Schmidt, phys. stat. sol. (c) **7**, 157 (2010).
12. J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **78**, 1396 (1997).
13. P.E. Blöchl, Phys. Rev. B **62**, 6158 (2000).
14. U. Gerstmann, phys. stat. solidi (b) **248**, 1319 (2011); p. 305–339 in “Advanced Calculations for Defects in Materials” edited by A. Alkauskas, P. Deák, J. Neugebauer, A. Pasquarello, and C.G. Van de Walle, Wiley-VCH, Weinheim, 2011.
15. E. Fermi, Z. Phys. **60**, 320 (1930).
16. G. Breit, Phys. Rev. **35**, 1447 (1930).
17. S. Blügel, H. Akai, R. Zeller, P.H. Dederichs, Phys. Rev. B. **35**, 3271 (1987).
18. H.A. Bethe, J.W. Negele, Nucl. Phys. A **117**, 575 (1968).
19. P. Giannozzi et al., J. Phys.: Condens. Matter **21**, 395502 (2009); <http://www.quantum-espresso.org>.

20. H.J. Monkhorst and J.D. Pack, *Phys. Rev. B* **13**, 5188 (1976).
21. J. Dabrowski in *Silicon Surfaces and Formation of Interfaces* ed. by J. Dabrowski and H.J. Müssig, p. 82–204 (World Scientific, Singapore 2000).
22. N.S. Saricifti, L. Smilowitz, A.J. Heeger, F. Wudl, *Science* **258**, 1474 (1992).
23. C.J. Brabec, G. Zerza, G. Cerullo, S. De Silvestri, S. Luzatti, J.C. Hummelen, N.S. Saricifti, *Chem. Phys. Lett.* **340**, 232 (2001).
24. G. Yu, J. Gao, J.C. Hummelen, F. Wudl, A.J. Heeger, *Science* **270**, 1789 (1995).
25. Q. Fan, B. McQuillin, D.D.C. Bradley, S. Whitelegg, A.B. Seddon, *Chem. Phys. Lett.* **347**, 325–330 (2001).
26. J.J. Cole, X. Wang, R.J. Knuesel, and H.O. Jacobs, *Nano Lett.* **8**, 1477–1481 (2008).
27. S. Ogawa, *Appl. Phys.* **46**, 518–522 (2007).
28. N. T. Son, Mt. Wagner, C. G. Hemmingsson, L. Storasta, B. Magnusson, W. M. Chen, S. Greulich-Weber, J. -M. Spaeth, and E. Janzen, in *Silicon Carbide, Recent Major Advances*, Eds. W. J. Choyke, H. Matsunami, G. Pensel, Springer Berlin, Heidelberg 2004, 461–492, ISBN 3-540-40458-9.
29. V. Yu. Aristov, *phys. usp.* **44**, 761 (2001).
30. Wenchang Lu, W.G. Schmidt, E.L. Briggs, and J. Bernholc, *Phys. Rev. Lett.* **85**, 4381–4384 (2000).
31. B. Friedel, S. Greulich-Weber, *Mater. Res. Soc. Symp. Proc.* **963**, 0963-Q15-10, Materials Research Society (2007).
32. A. Konopka, S. Greulich-Weber, E. Rauls, W.G. Schmidt, and U. Gerstmann, *Mater. Res. Soc. Symp. Proc.* **1322**, mrss11-1322-b10-02 (2011).

Ab-Initio Calculations of the Vibrational Properties of Nanostructures

Gabriel Bester and Peng Han

1 Introduction

Colloidal semiconductor nanocluster research is a rapidly growing field driven by the attractive idea to tailor material properties by acting on the morphology of the structures. The modification of the optical properties by merely changing the diameter of colloidal quantum dots is one of the figureheads of nanostructure science [1–3]. It is the intense research effort towards the fabrication of nanostructures with favorable properties that has helped to establish most of the knowledge base we rely on today. Till now, the modification of the electronic and optical properties by changing the size of the nanoclusters are well understood theoretically and well controlled experimentally. One open problem of nanostructure science is the effects of temperature on the electronic and optical properties of nanoclusters and hence their vibrational properties. A theory at $T = 0$ K yields very valuable results to unveil certain aspects of the underlying physics, but to make predictions valid in the real world, where the physical properties such as temperature broadening, quantum coherence dephasing, spin-flip transitions and relaxation of charge carriers are key components [4–6], the effects of vibration and temperature on the dynamical processes must be addressed.

The vibrational properties such as the phonon density of states (DOS) and dispersion of bulk semiconductors have been calculated with great accuracy using ab initio density functional perturbation theory (DFPT) since the end of the last century [7]. After the successful applications of DFPT on bulk phonon eigenmodes, ab initio studies on the vibrational properties of semiconductor nanostructures such as fullerenes, nanowires, nanotubes, and nanoclusters with small sizes have been performed [8–10]. However, an accurate density functional theory (DFT) study on the vibrational properties of nanoclusters with the experimentally relevant size of

G.Bester (✉) · P. Han
Max-Planck-Institut für Festkörperforschung, Stuttgart, Germany
e-mail: g.bester@fkf.mpg.de; p.han@fkf.mpg.de

few nm diameter has not been reported until now due to the high computational demand.

With the computational facility available at the Höchst Leistungs Rechenzentrum Stuttgart (HLRS), we have recently calculated confinement and surface effects on the vibrational properties of colloidal semiconductor nanoclusters based on first-principles DFT. We describe how the molecular-type vibrations, such as surface-optical, surface-acoustic, and coherent acoustic modes, coexist and interact with bulk-type vibrations, such as longitudinal and transverse acoustic (LA, TA) and optical (LO, TO) modes. We could link the vibrational properties to structural changes induced by the surface and highlight the qualitative difference between III–Vs and II–VIs semiconductor nanoclusters [11]. We describe the specific heat of nanoclusters at low temperature and link the thermodynamic properties to the low frequency vibrational modes and the surface structure. We suggest that the low temperature specific heat should be a promising avenue to study the surface properties of nanoclusters. Since nanoclusters are believed to have only a certain fraction of their surface atoms directly passivated by ligand atoms, we study the effects of the removal of passivant and the reconstruction of the surface on the vibrational properties. We attribute the strong modification of the vibrational properties to the transformation from sp^3 to sp^2 bonds.

2 Computational Methods

2.1 Research Methodology

A detailed review on density functional theory applied to lattice-dynamical calculations has been given elsewhere [7]. Here, we briefly outline our research methodology. Based on the adiabatic approximation (Born-Oppenheimer approximation), the lattice-dynamical properties of a system are given by

$$\left[-\sum_I \frac{\hbar^2}{2M_I} \frac{\partial^2}{\partial \mathbf{R}_I^2} + E(\{\mathbf{R}\})\right]\Phi(\{\mathbf{R}\}) = \varepsilon\Phi(\{\mathbf{R}\}) \quad (1)$$

where \mathbf{R}_I is the coordinate of atom I , M_I its mass, $\{\mathbf{R}\}$ the nuclei configuration given as a set of atomic positions, \hbar the Planck constant, ε and $\Phi(\{\mathbf{R}\})$ the eigenvalue and eigenvector of the lattice vibrations, respectively. $E(\{\mathbf{R}\})$ is the ground state energy of the system, which is determined by the many body Hamiltonian

$$H_{BO}(\{\mathbf{R}\}) = -\frac{\hbar^2}{2m} \sum_i \frac{\partial^2}{\partial \mathbf{r}_i^2} + \frac{e^2}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - V_R(\mathbf{r}) + V_N(\{\mathbf{R}\}) \quad (2)$$

where m is the mass of the electron, e the electron charge, and \mathbf{r}_i the coordinate of electron i . The electron-nucleus interaction potential $V_R(\mathbf{r})$ is given by

$$V_R(\mathbf{r}) = - \sum_{i,I} \frac{Z_I e^2}{|\mathbf{r}_i - \mathbf{R}_I|} \quad (3)$$

with Z_I represents the charge of the nucleus I . The electrostatic interaction potential $V_N(\{\mathbf{R}\})$ is written as

$$V_N(\{\mathbf{R}\}) = \frac{e^2}{2} \sum_{I \neq J} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|}. \quad (4)$$

Based on the Hellmann-Feynmann theorem, the force acting on the nucleus I is

$$F_I = - \frac{\partial E(\{\mathbf{R}\})}{\partial \mathbf{R}_I} = - \langle \Psi(\mathbf{r}, \{\mathbf{R}\}) | \frac{\partial H_{BO}(\{\mathbf{R}\})}{\partial \mathbf{R}_I} | \Psi(\mathbf{r}, \{\mathbf{R}\}) \rangle \quad (5)$$

and the force constant matrix elements are

$$\frac{\partial^2 E(\{\mathbf{R}\})}{\partial \mathbf{R}_I \partial \mathbf{R}_J} = \int \frac{\partial \rho_R(\mathbf{r})}{\partial \mathbf{R}_J} \frac{\partial V_R(\mathbf{r})}{\partial \mathbf{R}_I} d\mathbf{r} + \int \rho_R(\mathbf{r}) \frac{\partial^2 V_R(\mathbf{r})}{\partial \mathbf{R}_I \partial \mathbf{R}_J} d\mathbf{r} + \frac{\partial^2 V_N(\{\mathbf{R}\})}{\partial \mathbf{R}_I \partial \mathbf{R}_J} \quad (6)$$

where $\Psi(\mathbf{r}, \{\mathbf{R}\})$ is the electronic ground-state wave function and $\rho_R(\mathbf{r})$ the electron charge density for the nuclei configuration $\{\mathbf{R}\}$. The charge density $\rho_R(\mathbf{r})$ is obtained by mapping the problem onto a set of one-particle equations (Kohn-Sham equations):

$$\left[-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial \mathbf{r}^2} + V_R(\mathbf{r}) + e^2 \int \frac{\rho_R(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}' + \frac{\delta E_{xc}}{\delta \rho_R(\mathbf{r})} \right] \psi_n(\mathbf{r}) = \epsilon_n \psi_n(\mathbf{r}) \quad (7)$$

and

$$\rho_R(\mathbf{r}) = 2 \sum_{n=1}^{N/2} |\psi_n(\mathbf{r})|^2 \quad (8)$$

where δE_{xc} is the exchange-correlation energy, ϵ_n and $\psi_n(\mathbf{r})$ are the eigen energy and wave function of the electronic states, respectively.

Based on the harmonic approximation of lattice dynamics, the frequencies ω and the corresponding eigenmodes \mathbf{u}_I are obtained by solving the eigenvalue equation

$$\sum_J \frac{1}{\sqrt{M_I M_J}} \frac{\partial^2 E(\{\mathbf{R}\})}{\partial \mathbf{R}_I \partial \mathbf{R}_J} \mathbf{u}_J = \omega^2 \mathbf{u}_I. \quad (9)$$

To analyze the eigenmodes in terms of core and surface contributions, we calculate the projection coefficients

$$\alpha_{c,s,p}^v = \frac{\sum_I^{(N_c, N_s, N_p)} |\mathbf{X}^v(I)|^2}{\sum_{I=1}^N |\mathbf{X}^v(I)|^2}, \quad (10)$$

where, N_c , N_s , N_p and N are the core, surface, passivant, and total number of atoms, $\mathbf{X}^v(I)$ represents the three components that belong to atom I from the $3N$ -component eigenvector. We define the surface atoms as the atoms belonging to the outermost seven layers of the cluster (around 3 Å thick). From the phonon DOS $D(\omega)$ we obtain the specific heat according to:

$$C_v(T) = N_A k_B \int_0^\infty \left(\frac{\hbar\omega}{k_B T} \right)^2 \frac{e^{\hbar\omega/k_B T}}{(e^{\hbar\omega/k_B T} - 1)^2} D(\omega) d\omega. \quad (11)$$

2.2 Computational Details

The nanoclusters we studied are constructed by cutting a sphere, centered on a cation with T_d point group symmetry, from the zinc blende bulk structure and removing the surface atoms having only one nearest-neighbor bond. The surface dangling bonds are terminated by pseudohydrogen atoms H^* with a fractional charge of 1/2, 3/4, 5/4, and 3/2 for group VI, V, III, and II atoms, respectively. The calculations are performed using the local-density approximation (LDA), Trouiller-Martin norm-conserving pseudopotentials with an energy cutoff of 30 Ry for III–Vs and 40 Ry for II–VIs.

The geometry relaxation is performed using the Broyden-Fletcher-Goldfarb-Shano (BFGS) procedure for the optimization of the ionic positions. The forces are minimized to less than 3×10^{-6} a.u. (5×10^{-4} a.u.) under constrained symmetry for the passivated (unpassivated) nanoclusters. With the optimized geometry, the dynamical matrix elements $\frac{\partial^2 E(\{\mathbf{R}\})}{\partial \mathbf{R}_I \partial \mathbf{R}_J}$ are obtained by solving Eq. (6). In the calculation, the charge density $\rho_R(\mathbf{r})$ are obtained by solving the Kohn-Sham equation self-consistently and the values of $\frac{\partial \rho_R(\mathbf{r})}{\partial \mathbf{R}_J}$ are calculated using a finite difference approach. In principle we need $3N$ atomic displacements to obtain all the elements of the dynamical matrix (N being the number of atoms). In practice we calculate a significantly lower number of displacements ($3N/24$) and use the symmetry elements of the point group to deduce the missing elements. This is a key points to be able to treat these large structures.

All the calculations are performed with the CPMD code package developed at the Max Planck Institute in Stuttgart and at IBM in Zürich [12]. The CPMD code is a high performance parallelized plane wave/pseudopotential implementation of DFT. It offers, at the moment, the best scaling among the DFT codes using a hybrid scheme of MPI and OpenMP.

In this project, all the calculations were carried out on the NEC Nehalem Cluster at HLRS with 2.8 GHz and 12 GB memory per nodes, and infiniband interconnects.

The details of the scaling behavior and the performance per CPU for CPMD code have been given elsewhere [13].

3 Results

3.1 Confinement Effects on Vibrational Properties

We have calculated the vibrational properties for a total of 23 nanoclusters made of InP, InAs, GaP, GaAs, CdS, CdSe, and CdTe. The wave function of the lowest unoccupied molecular orbital (LUMO) state for $\text{Ga}_{531}\text{As}_{532}\text{H}_{412}^*$ nanocluster with isosurface corresponding to 75 % of the maximum value are presented in Fig. 1. In this report, we extract the essence of these calculations and select representative results. The vibrational DOS of InP and CdS nanoclusters along with the bulk phonon DOS are plotted with a broadening of 0.8 cm^{-1} in Fig. 2. Although formally TA, LA, and TO, LO phonon modes cease to exist in a nanocluster, the comparison with the bulk phonon DOS in Fig. 2e, j reveals an obvious bulk parentage. From Fig. 2, we see that: (i) The III–V (InP) nanoclusters show a blueshift of LO-, TO-, and LA-derived cluster modes with decreasing size while this blueshift cannot be found in II–VIs (CdS). (ii) The surface modes tend to completely fill the acoustic-optic phonon gap in II–VIs but not for III–Vs. (iii) The “broadening” of the bulk-like optical phonon branches induced by the confinement is larger for II–VIs than for III–Vs.

These three effects can be understood from the geometry of the relaxed nanoclusters. We plot the nearest-neighbor distances of relaxed III–Vs and II–VIs nanoclusters as a function of the distance of the respective bond to the cluster center in Fig. 3. From this figure, we see that the bond length at the dot center is reduced through the presence of the surface in all cases. The bond length distribution of III–Vs and II–VIs exhibits qualitative differences. For III–Vs, the surface shells show a successive reduction of bond length, going outward, while II–VIs show a large bond length distribution. The overall reduction of bond length in III–Vs along with the positive Grüneisen parameters explains the blueshift of the LO-, TO-, and LA-derived cluster modes. The lack of shift in the TA modes stems from the small negative Grüneisen parameter for this branch. Moreover, we attribute the broadening of optical branches and the filling of the phonon-gap in II–VIs to the results of the large bond length distribution.

After discussing the confinement effects on the high frequency vibrational modes, we now focus on the low frequency modes. In Fig. 4, we plot the size-dependent lowest frequencies $f_{\min} = v/4R$ calculated from the longitudinal and transverse speed of sound v and cluster radius R as solid and dashed curves. The circles and the crosses are the lowest core and surface acoustic modes obtained from the DFT calculations. We see that the lowest core modes follow closely the analytic $1/R$ dependence while the surface acoustic modes are strongly affected

Fig. 1 The wave function of the lowest unoccupied molecular orbit (LUMO) state for a $\text{Ga}_{531}\text{As}_{532}\text{H}_{412}^*$ nanocluster with isosurface corresponding to 75 % of the maximum value. The colors *blue* and *red* give the phase of the wave functions


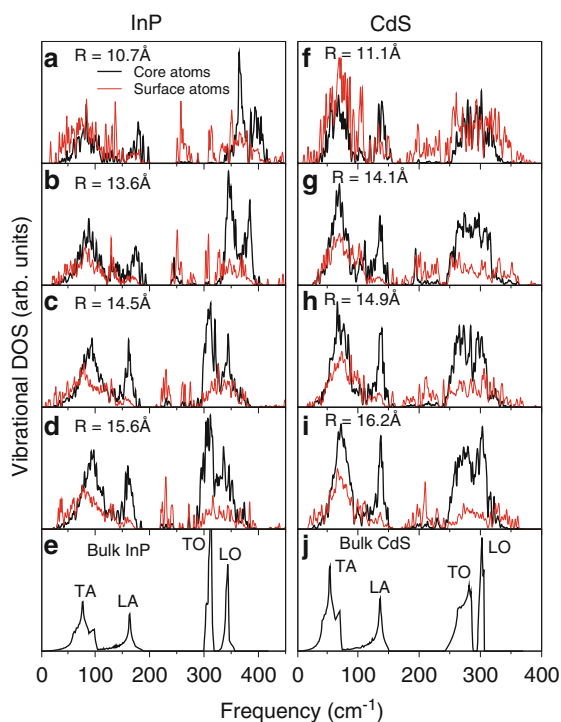


Fig. 2 Vibrational density of states (DOS) contributed by core atoms (*black*) and surface atoms (*red*) for (a)–(d) InP clusters, (e) bulk InP, (f)–(i) CdS clusters, and (j) bulk CdS



by the morphology of the surface and are not monotonous with cluster size. Another important type of vibrational modes are the so-called *coherent acoustic modes*, in which all the atoms vibrate in phase. The coherent phonon modes have been observed with Raman spectroscopy, far-infrared absorption spectroscopy, and resonant high-resolution photoluminescence spectroscopy, and are now the center of

Fig. 3 Bond-length distribution as a function of their distance to the dot center for (a) InP, (b) CdS, (c) III-Vs, and (d) II-VIs. LDA bulk bond lengths are given as *dashed lines*

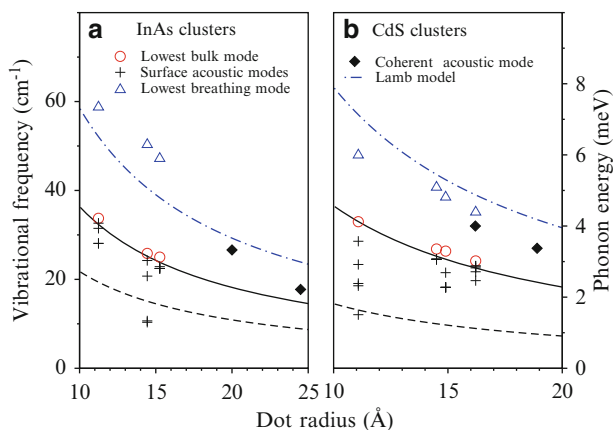
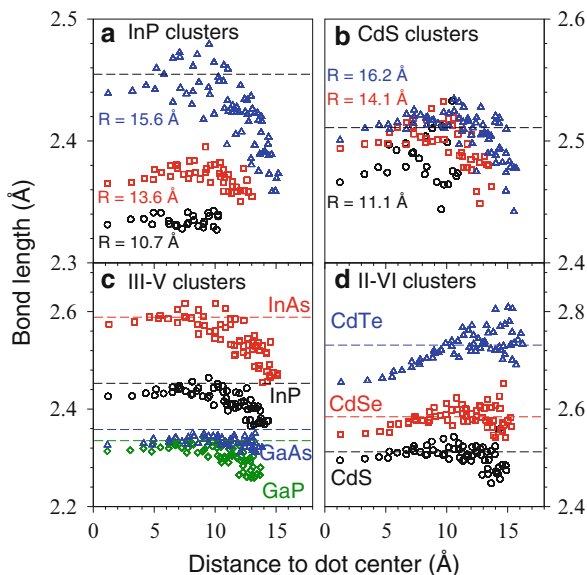


Fig. 4 Size-dependent low frequency vibrational modes for (a) InAs and (b) CdS. Lowest modes with bulk character (*circles*), surface acoustic modes (*crosses*), lowest breathing modes (*triangles*), and experimental results (Oron et al. [14] and Saviot et al. [18]) for the coherent acoustic modes (*diamonds*). Lowest spheroidal mode according to the Lamb model (*dashed-dotted line*), according to the confined bulk model using the sound velocity of the TA- (*solid lines*) and LA-branch (*dashed lines*)

attention when the manipulation of spins and the spin dephasing is investigated [14–16]. We plot our results as triangles along with the experimental data as diamonds and the results from the Lamb model as dashed-dotted line. Our results are in good agreement with the experimental results (although our clusters sizes are still

somewhat smaller than experiment in the case of InAs) and with the simple Lamb model.

3.2 *Passivated versus Unpassivated Nanoclusters*

In the real world applications, the nanoclusters are believed to have only a certain fraction of their surface atoms directly passivated by ligand atoms [17]. Thus, we studied an extreme situation where the nanoclusters are unpassivated [19]. In Fig. 5, we compare the vibrational DOS of a fully passivated with geometry optimization (a), an unpassivated nanocluster without (b) and with (c) geometry relaxation. By removing the passivants while freezing the atomic positions, we effectively create dangling bonds in the sp^3 hybrid orbitals. In this frozen geometry, the orbitals cannot redistribute effectively and the bonds close to the surface are weakened. This leads to the red shift of the modes with surface character. Once the geometry is optimized, the dangling bonds tend to transform from sp^3 to sp^2 hybrid orbitals and strengthen. This leads to a blue shift of vibrational modes with a magnitude even greater than the fully passivated cluster and in general a “broadened” vibrational DOS. To understand this effect, we plot in Fig. 6 the relaxed unpassivated cluster (a) and the nearest neighbor distances of a relaxed passivated and a relaxed unpassivated InP nanocluster as a function of the distance of the respective bond to the cluster center (b). From Fig. 6b, we see that the bond length reduction close to the surface is more significant in the case of the unpassivated cluster and this results in the blue shift. Moreover, we can also find that the variation in bond lengths is significantly larger in the unpassivated cluster, which leads to the broadening of the vibrational DOS.

3.3 *Thermodynamic Properties*

Following the discussion on the vibrational properties, we now turn to the thermodynamic properties of nanoclusters. The vibrational specific heat $C_v(T)$ is calculated using Eq. (11) with the vibrational DOS from DFT computations. Different aspects of the results are summarized in Fig. 7.

In Fig. 7a, we plot $C_v(T)/T$ for CdSeH* nanoclusters with different size as a function of T^2 . With this choice of axis, the Debye T^3 regime would appear linear. We identify two distinct regions, *region 1* below 24K^2 and *region 2* above this temperature. In region 1, $C_v(T)/T$ shows a strongly non-linear dependence and it converts to a nearly linear behavior in region 2. Most surprising is that we find the smaller nanocluster has a larger specific heat than the larger ones. To understand these behaviors, we plot in Fig. 7b the low frequency vibrational spectrum along with the percentage of surface character. The temperature region 1 roughly corresponds to the frequency region below 30cm^{-1} in (b), and the small

Fig. 5 Vibrational DOS of InP nanoclusters with (a) passivant, (b) unpassivated without relaxation and (c) unpassivated with relaxation

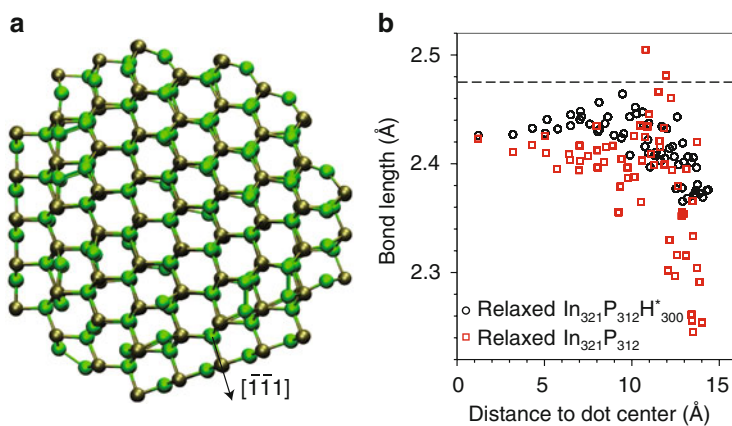
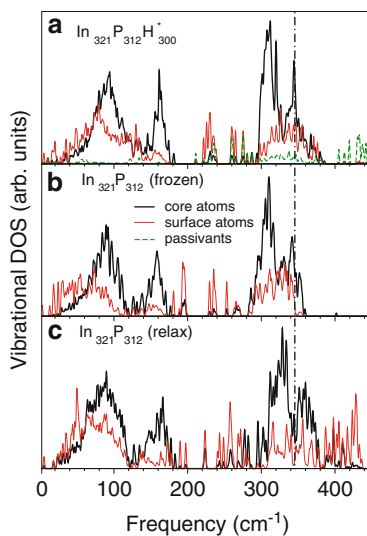


Fig. 6 (a) Relaxed geometry of un-passivated $\text{In}_{321}\text{P}_{312}$ nanoclusters, (b) Bond length distribution as a function of their distance to the dot center for $\text{In}_{321}\text{P}_{312}\text{H}_{300}^*$. The LDA bulk bond length is given as *dashed lines*

nanoclusters have more vibrational modes (higher vibrational DOS) than the larger ones in this frequency region and these modes have surface character. We can see that the surface modes move up in frequency with increasing cluster size and attribute this effect to the atomic configuration on a curved surface. As depicted in the inset of Fig. 7a, the stronger curvature of smaller dots leads to an “open” surface that allows for softer surface modes. From Fig. 7b, we see that the vibrational modes with core characters contribute to the specific heat when the temperature arrives into region 2. These vibrational modes show a red shift with decreasing cluster size due to the negative Grüneisen parameter along with the contraction of the surface.




Fig. 7 Specific heat C_v divided by the temperature T as a function of T^2 for different CdSe nanoclusters. In (b) the low frequency eigenmodes are drawn as *vertical lines* for three different CdSe nanoclusters. The *solid circle* gives the percentage of surface character of each mode. In (c) we estimated the slopes in the (nearly) linear regime of panel (a) for CdSe nanoclusters, changing the passivant mass. In (d) we report the corresponding onset temperature of C_v .

These bulk-like contributions are reflected in the nearly linear behavior in Fig. 7a. We estimated the slopes and the onset temperature of the curves and plotted them in Fig. 7c, d with the passivant mass as hydrogen, carbon, oxygen and phosphorous. The slope reduces with increasing size in all cases and it increases with increasing mass of passivant. This behavior can be understood from the reduced frequency with increasing passivant masses. From Fig. 7d, we find that the smaller nanoclusters have an earlier onset than the larger ones. This is again related to the surface mode softening we reported for decreasing cluster size.

4 Summary

In summary, we have performed ab initio DFT calculations to study the confinement and the surface effects on the vibrational and thermodynamic properties of III–V and II–VI colloidal nanoclusters with up to thousand atoms. We can identify the

following confinement and surface effects. (i) The LA, TO and LO-derived cluster modes of III–V clusters significantly blue shift with decreasing cluster size. For II–VI clusters this shift is absent but the broadening of bulk derived modes is significant and the gap between optical and acoustic phonons is filled by surface modes. (ii) We can clearly ascribe these observations to the large relaxation of the clusters dominated by an inward relaxation of the surface penetrating deep inside the cluster in case of the III–Vs and a large distribution of bond length at the surface of II–VIs. These strong confinement effects tend to disappear for clusters with more than 1,000 atoms. (iii) We find surface optical modes in the phonon gap and surface acoustic modes as the lowest frequency modes. The *coherent acoustic phonons* are identified and found to be in good agreement with results from the Lamb model and experiment. (iv) In the unpassivated nanoclusters, the unpaired electrons in the sp^3 hybrid orbitals reduce the bond strength and result in red shift with frozen geometry. Once the geometry is optimized, the dangling bonds tend to transform from sp^3 to sp^2 hybrid orbitals and leads to a blue shift of vibrational frequencies. (v) The low temperature specific heat reflects the surface acoustic vibrational modes and we suggest to study the low temperature specific heat to access the surface properties.

References

1. Rogach, A. L., Eychmüller, A., Hickey, S. G., Kershaw, S. V.: Infrared-emitting colloidal nanocrystals: Synthesis, assembly, spectroscopy, and applications. *Small* **3**, 536–557 (2007).
2. Gaponik, N., Hickey, S. G., Dorfs, D., Rogach, A. L., Eychmüller, A.: Progress in the light emission of colloidal semiconductor nanocrystals. *Small* **6**, 1364–1378 (2010).
3. Talapin, D. V., Lee, J.-S., Kovalenko, M. V. Shevchenko, E. V.: Prospects of colloidal nanocrystals for electronic and optoelectronic applications. *Chem. Rev.* **110**, 389–458 (2010).
4. Fischer, J. Loss, D: Hybridization and spin decoherence in heavy-hole quantum dots. *Phys. Rev. Lett.* **105**, 266603(4pp) (2010).
5. Kilina, S. V., Kilin, D. S., Prezhdo, O. V.: Breaking the phonon bottleneck in PbSe and CdSe quantum dots: Time-domain Density Functional Theory of charge carrier relaxation. *ACS Nano* **3**, 93–99 (2009).
6. An, J.M., Califano, M., Franceschetti, A., Zunger, A.: Excited-state relaxation in PbSe quantum dots. *J. Chem. Phys.* **128**, 164720 (7pp) (2008).
7. Baroni, S., de Gironcoli, S., Corso A. D.: Phonons and related crystal properties from density-functional perturbation theory. *Rev. Mod. Phys.* **73**, 515–562 (2001).
8. Peelaers, H., Partoens, B., Peeters, F. M.: Phonon Band Structure of Si Nanowires: A Stability Analysis *Nano Lett.* **9**, 107 (5pp) (2009).
9. Caudal, M., Saitta, A. M., Lazzeri, M., Mauri, F.: Kohn anomalies and nonadiabaticity in doped carbon nanotubes. *Phys. Rev. B* **75**, 115423 (11pp) (2007).
10. Chelikowsky, J. R., Zayak, A. T., Chan, T., Tiago, M. L., Zhou, Y., Saas, Y.: Algorithms for the electronic and vibrational properties of nanocrystals. *J.Phys. Condens. Matter* **21**, 064207 (2009).
11. Han, P., Bester, G.: Confinement effects on the vibrational properties of III-V and II-VI nanoclusters. *Phys. Rev. B* **85**, 041306(R) (4pp) (2012).
12. The CPMD consortium page, coordinated by M. Parrinello and W. Andreoni, Copyright IBM Corp 1990–2008, Copyright MPI für Festkörperforschung Stuttgart 1997–2001 <http://www.cpmd.org>.

13. Bester, G., Han, P. in: Nagel, W. E., Kröner, D. B., and Resch, M. M., eds. *High Performance Computing in Science and Engineering* '11, p.119, Springer-Verlag Berlin Heidelberg 2012.
14. Oron, D., Aharoni, A., de MelloDonega, C., van Rijssel, J., Meijerink, A., Banin, U.: Universal Role of Discrete Acoustic Phonons in the Low-Temperature Optical Emission of Colloidal Quantum Dots. *Phys. Rev. Lett.* **102**, 177402 (4pp) (2009).
15. Krauss, T. D., Wise, F. W.: Coherent Acoustic Phonons in a Semiconductor Quantum Dot. *Phys. Rev. Lett.* **79**, 5102 (4pp) (1997).
16. Chilla, G., Kipp, T., Menke, T., Heitmann, D., Nikolic, M., Fromsdorf, A., Kornowski, A., Forster, S., Weller, H.: Direct Observation of Confined Acoustic Phonons in the Photoluminescence Spectra of a Single CdSe-CdS-ZnS Core-Shell-Shell Nanocrystal. *Phys. Rev. Lett.* **100**, 057403 (4pp) (2008).
17. Tang, J., Kemp, K. W., Hoogland, S., Jeong, K. S., Liu, H., Levina, L., Furukawa, M., Wang, X., Debnath, R., Cha, D., Chou, K. W., Fischer, A., Amassian, A., Asbury, J. B., Sargent, E. H.: Colloidal-quantum-dot photovoltaics using atomic-ligand passivation. *Nature materials.* **10**, 765 (7pp) (2011).
18. Saviot, L., Champagnon, B., Duval, E., Ekimov, A. I.: Size-selective resonant Raman scattering in CdS doped glasses. *Phys. Rev. B* **57**, 341 (6pp) (1998).
19. Han, P., Bester, G.: Insights about the surface of colloidal nanoclusters from their vibrational and thermodynamic properties. *J. Phys. Chem. C* **116**, 10790 (6pp) (2012).

Part III

Reacting Flows

Prof. Dr. Dietmar Kröner

Two contributions in this section about Reacting Flows impressively show how the results obtained by high performance computing can be used for the design of realistic combustors in order to increase the efficiency or to avoid damages. Furthermore in two other contributions parameter studies for flames are performed and the results are compared with available experimental data.

In the contribution about “Conservative Implementation of LES-CMC for Turbulent Jet Flames” by P. Siwaborworn and A. Kronenburg, the authors carry out parametric studies for turbulent jet flames, in particular for a so-called Sandia D flame with available experimental results. The method used for this purpose is a combination of large eddy simulation (LES) and of conditional moment closure (CMC) to account for turbulence due to chemistry. Two studies are carried out. One deals with the combustion model and investigates the impact of a so-called FDF-weighting function in the governing equations which is believed to be the main reason for inaccurate prediction. The other deals with the impact of the CMC grid resolution. The study demonstrates that inclusion of the FDF-weighting function leads to more accurate solution. The simulation have been performed on the NEC Nehalem Cluster with up to 720 cores. The parallel efficiency falls below 40 % if more than 180 cores are used.

In “Numerical Investigation of a Complete Scramjet Demonstrator Model for Experimental Testing under Flight Conditions” by Y. Simsont, P. Gerlinger, M. Aigner a scramjet model has been investigated numerically. The simulation corresponds to a realistic scramjet model which has been tested for wind tunnel experiments at Mach 8. The results are used for improving the design, such that the inflow conditions and the distribution of the temperature are optimized. For instance the simulations of the flow through the original geometry do not indicate selfignition. Therefore additional wedges have been attached and it could be shown

that locally the temperature increases such that self-ignition occurs. The underlying mathematical model consists of the compressible Navier-Stokes equations and transport equations for the species with reactive source terms. The results have been obtained with the software package TASC3D on the NEC SX-9 mit 16 CPUs. The performance of the code is carefully studied. It turns out that it was not so strong as results on the older SX-8 machine. This project was supported by the DFG.

The paper “A Unified TFC (Turbulent Flame-Speed Closure) Combustion Model for Numerical Computation of Turbulent Gas Flames” by F. Zhang, P. Habisreuther, M. Hettel and H. Bockhorn concerns the modeling and simulation in turbulent compressible flows, in particular the modeling of premixed, non-premixed and partially premixed combustion. The numerical results are obtained by using the OpenFOAM software, which is based on centered Finite-Volume schemes. Two different turbulent models, RANS and LES are compared. For the combustion mechanisms for the species “state-of-the-art” methods are used, e.g. GRI 3.0 or Maas/Pope. The details of the chemical reactions are computed using CHEMKIN II and the results are integrated into the simulations via tables. The results of all simulations show reasonably good agreement with the experiments. The computational effort for RANS is much lower than that for LES. For the parallelization an efficiency of 55 % has been obtained. The main result has been summarized by the authors as follows: “Nevertheless, there are only very few alternative models until now which could be used to simulate such different flames using one single reaction model. There are models indeed, which could provide better results for some of the demonstrated flames, but these will completely fail elsewhere.” This project was supported by the DFG.

In “Lagrangian Approach for the Prediction of Slagging and Fouling in Pulverized Coal Combustion” by O. Lemp, U. Schnell and G. Scheffknecht the authors present a modeling approach for the prediction of slagging and fouling in industrial coal-fired boilers. Both slagging and fouling may cause damages of the furnace and lead to an immense decline of power plant efficiency. The simulations done so far (in absence of measurements) are thought to contribute to design, investigation of damage processes as well as optimization of boiler performance. Simulation have been performed using the CFD code AIOLOS based on Finite-Volume methods. Chemically reacting flow is described by convection- diffusion-reaction equations. Turbulence is computed via the $k - \varepsilon$ model and for chemistry-turbulence interactions the Eddy-Dissipation concept is employed. Since realistic (industrial) scenarios are considered, high-performance computing must be used in order to get effective hints for a better design, in particular for an optimization of the boiler performance. The computation is based on six million computational cells and the calculation were executed on the NEC-SX8 with 8 CPUs and on the NEC-SX9 with 16 CPUs. The efficiency is lower than for previous computations.

Conservative Implementation of LES-CMC for Turbulent Jet Flames

P. Siwaborworn and A. Kronenburg

Abstract The objective of the present work is to validate a large eddy simulation (LES) approach that has been coupled with a conditional moment closure (CMC) method for the computation of turbulent diffusion flames. Contrary to earlier work, we use a conservative implementation of CMC that ensures mass conservation of the fluxes across the computational cell faces. This is equivalent to a weighting of the fluxes by their probabilities at the cell faces, and it is thought that this weighting leads to a more dynamic response of the conditionally averaged moments to temporal changes induced by the large scale turbulent motion. The first application to the Sandia Flame Series D-F allows for the validation of the method, but further studies with different flame geometries and more pronounced large scale instationary effects will be needed for the demonstration of the benefits of conservative CMC when compared to the conventional (non-conservative) implementation.

1 Introduction

The Large eddy simulation (LES) approach is considered to be the most promising approach for the computation of turbulent flows in applications of engineering interest. LES solves large scales of turbulent flows up to grid sizes using spatial filtering and models subgrid scales using Smagorinsky model. CMC is applied for a turbulent combustion modelling using mixture fraction as a conditional variable. A non-conservative LES-CMC has provided predictions of major and minor species for different flames in a last decade. However, inaccurate predictions occur in CMC cells which have large temporal variations of the mixture fraction field. A lack of FDF-weighting (filtered density function) ratios in a convective term of the non-conservative CMC is believed to be the main reason for inaccurate predictions.

P. Siwaborworn (✉) · A. Kronenburg
Institut für Technische Verbrennung, University of Stuttgart, Stuttgart, Germany
e-mail: p.siwaborworn@itv.uni-stuttgart.de; kronenburg@itv.uni-stuttgart.de

In contrast to non-conservative LES-CMC, the present conservative formulation is inherently mass conserving. It considers FDF-weighting ratios in convective term so that improved predictions of local conditional scalars can be obtained.

In this work, investigations of turbulent jet flames (Sandia Flame D, E and F) are performed by the conservative LES-CMC approach. Flame D is used as the first test case to validate the numerical results by comparison with well-established experimental data. Subsequently, Flames E and F are investigated for extinction and reignition phenomena. Computational results of the present work from using HLRS resources are given first. Subsequently, the computational resources which have been used to simulate Sandia Flames series will be addressed.

2 Results of the Present Work

In this section, the computational setups are described, followed by the summaries of parametric CMC studies in Sect. 2.2. Section 2.3 presents the simulation results for Sandia Flame D, which is carried out as the first test case in order to validate the LES-CMC simulation models as a reference. Predictions of Flame D are also chosen as representatives in this paper, since the simulation results of Flames E and F follow the same tendency as predictions of in Flame D. This section closes with the discussion and conclusions about the performance of various CMC model parameters in Sect. 2.4.

2.1 Computational Setup

Computational grid was generated with dimensions of $80D$ in z -direction and $8D$ in x - and y -directions at the flame base increasing to $60D$ at the outlet of the domain. Fine grid simulations of Sandia Flame series, having $112 \times 112 \times 320$ cell grids for LES and $8 \times 8 \times 80$ cell grids for CMC (reference case), have been carried out with HLRS resources. The regions above the jet and pilot are captured by 28 and 40 LES cells (for a dimension), respectively. Due to grid independence studies by Navarro et al. [1], these computational cells satisfy a condition that the largest fraction of energy spectrum is resolved after the initial break-up of the jet. The CMC grid has 100 nodes in mixture fraction space which has refinement at $\eta = 0$ and 1.

2.2 Parametric Studies

The results of the LES-CMC modelling are presented in three main parts. Parametric studies of flow and mixing field, parametric studies of combustion model and parametric study of CMC grid resolution are investigated in order to study the effects

of each parameter on simulation results. Details of these studies for each parameter are given in the following sections.

2.2.1 Parametric Studies of Flow and Mixing Field

The parametric studies of flow and mixing field are the inflow velocity variances in turbulent inflow generator, the Schmidt number, Sc , the turbulent Schmidt number, Sc_t , and constant value, C_ξ , which is used in modelling of the variance of mixture fraction ($\widetilde{\xi_{sgs}^{\prime\prime 2}} = C_\xi \Delta^2 \left(\frac{\partial \tilde{\xi}}{\partial x_j} \frac{\partial \tilde{\xi}}{\partial x_j} \right)$). Besides C_ξ , all of them are fluid properties which actually are not allowed to be changed. However, the velocity variance levels are adjusted to yield suitable values for the inflow generator which creates the oscillation of the velocity field in this work. The Schmidt number and the turbulent Schmidt number, which are the ratios of momentum transfer rate to mass transfer rate for resolved and unresolved scales, are varied to test a sensitivity of the flow and mixing field. Sandia Flame D is used as a reference case and thus simulation results from these optimal values are reported in Sect. 2.3.

2.2.2 Parametric Studies of Combustion Model

The evaluation of the LES-CMC combustion model is in the focus of the present research project. The parametric studies concerning the combustion model carried out here comprise the evaluation of the CMC formulation, the approximation of the CMC convective fluxes and the model for the conditionally filtered turbulent diffusivity.

- CMC Formulations

The difference of the two CMC formulations (the non-conservative form in Eq. 1 and the conservative form in Eq. 2) is the inclusion of the FDF information into the transport equation, in particular into the convective term (the second term on the LHS) of the conservative CMC formulation.

$$\frac{\partial Q_\alpha}{\partial t} + \tilde{u}_{j,\eta} \frac{\partial Q_\alpha}{\partial x_j} - \frac{1}{\gamma} \frac{\partial}{\partial x_j} \left(\gamma D_\eta \frac{\partial Q_\alpha}{\partial x_j} \right) = \tilde{w}_{\alpha,\eta} + \tilde{N}_\eta \frac{\partial^2 Q_\alpha}{\partial \eta^2}, \quad (1)$$

$$\gamma \frac{\partial Q_\alpha}{\partial t} + \frac{\partial}{\partial x_j} (\gamma \tilde{u}_{j,\eta} Q_\alpha - \gamma D_\eta \frac{\partial Q_\alpha}{\partial x_j}) = \gamma \tilde{w}_{\alpha,\eta} + \gamma \tilde{N}_\eta \frac{\partial^2 Q_\alpha}{\partial \eta^2} + Q_\alpha \frac{\partial}{\partial x_j} (\gamma \tilde{u}_{j,\eta}), \quad (2)$$

where γ denotes $\bar{\rho} \tilde{P}(\eta)$ and $\tilde{P}(\eta)$ is the Favre filtered probability density function (FDF). $Q_\alpha = Y_\alpha | \eta$ is the conditionally filtered mass fraction. \tilde{u}_η is the conditionally filtered velocity, \tilde{w}_η is the conditionally filtered reaction source term and \tilde{N}_η is the conditionally filtered scalar dissipation rate. Term $D_\eta \frac{\partial Q_\alpha}{\partial x_j}$ is

the modelled -by using of a gradient diffusion approximation- of the subgrid-scale conditional scalar flux. Based on finite volume method, both conditional species transport equations can be applied to each control volume (CV) of the computational domain. It is believed that including the FDF information in convection will make the CMC conservative form more precise than the traditional one [2]. Therefore, the study of two different formulations of the combustion model is performed in this work to reveal the results of the assumption.

- Flux Approximations

Since each CMC cell comprises of a number of LES cells, two methods can be applied to approximate the convective flux between CMC cells. The first method calculates convective flux over the CMC cell face from the convective fluxes of the LES cells adjacent to the CMC cell face. The second method computes convective flux from the values at the CMC cell centre, which takes into account the values of all LES cells within the CMC cell. The convective flux from the first method is shown as the summation of the small arrows in Fig. 1, while the convective flux from the second method is shown as the big arrow in the same figure.

- Conditionally Filtered Turbulent Diffusivity Models


There are three methods to model the conditionally filtered turbulent diffusivity, D_η . However, all of them are based on the Smagorinsky model for the subgrid-scale kinematic viscosity, ν_t and the relation of unconditionally filtered turbulent diffusivity, $D_t = \nu_t / Sc_t$. In the first method, D_η is calculated based entirely on CMC cell (named $D_{\eta,1}$). Therefore, ν_t in this method is calculated based on CMC grid resolution. In the second method, $D_{\eta,2}$ is calculated based on D_t from every LES cell which locate inside that CMC cell. The ensemble averaging over a CMC cell can be computed by weighting with FDF. In the third method $D_{\eta,3}$, the ratio of the size of CMC cell to LES cell is included into the second method in order to adjust the length scale during modelling D_η value. This additional value is hoped to predict more accurately since the D_η model should be based on the filter width of CMC instead of on the filter width of LES.

2.2.3 Parametric Study of the CMC Grid Resolution

Another main parametric study for all Sandia Flame series is the CMC grid resolution. For a simple and stable flame, this parameter might not reveal any effect. However, it is believed that a high number of CMC cells may capture the extinction and reignition phenomena due to the turbulence-chemistry interactions in Flames E and F. Thus, three CMC grid resolutions are performed in this study topic. These are $4 \times 4 \times 80$, $8 \times 8 \times 80$ and $16 \times 16 \times 80$ CMC cells for the same LES resolution.

To summarize, all parametric studies are shown in Table 1 and they were varied for Flame D. The values resulting in the best agreement between simulation and experiments of parametric studies of flow and mixing field of Flame D were chosen for further simulations of Flames E and F.

Fig. 1 A schematic of the two approximations of the CMC convective flux



2.3 Results of Sandia Flame D

Results of Sandia Flame D are composed of three principal parts. Firstly, results of the parametric studies of flow and mixing field are reported. Subsequently, results of the parametric studies of combustion model are shown and discussed. Finally, results of the parametric study of CMC grid resolution are given and discussed.

2.3.1 Parametric Studies of the Flow and Mixing Field

Using a reference case of parametric studies in combustion model (CMC-1, flux-1 and $D_{\eta,2}$) and CMC grid resolution of $8 \times 8 \times 80$, the best results of parametric studies in flow and mixing fields are variance-2, Sc_2 , $Sc_{t,1}$ and $C_{\xi,1}$. The meaning of each parameter can be found in Table 1. Overviews of best results in flow and mixing field which use these optimum values can be observed in Figs. 2 and 3.

Figure 2 shows a snapshot of the instantaneous temperature field along a 2D plane through the burner centerline for the whole computational domain (left) and focused on the upstream region (right). The black lines identify the isoline of stoichiometric mixture fraction. It can be observed from a temperature profile (Fig. 2 (right)) that there is a high level of turbulence which comes from the digital turbulent inflow generator at the inlet. Moreover, the local extinction, which would be characterized by discontinuous red color of the temperature along the isoline of stoichiometric mixture fraction, hardly occurs in Flame D, in accordance with the experimental findings.

Radial profiles of mean and RMS axial velocity and mixture fraction at three downstream locations are shown in Fig. 3. Both mean and RMS of axial velocity and mixture fraction agree properly with the experiments [3,4], since the effects of initial inflow from inflow generator are previously checked and adjusted. Moreover, it can be seen from Fig. 3 that the jet spreading is captured well. Small overpredictions of the mean mixture fraction in the range of $1.2 < r/D < 2$ at position $z/D = 3$ may come from the influences of lateral boundary conditions. Small overpredictions

Table 1 Summary of parameters studies

Quantity	Name	Values or methods
Variances of inflow generator	variance-1	$u'u', v'v'$ and $w'w'$ [3]
	variance-2	$\frac{2}{3}u'u', \frac{2}{3}v'v'$ and $\frac{2}{3}w'w'$
Schmidt number	Sc_1	0.4
	Sc_2	0.7
	Sc_3	1.0
Turbulent Schmidt number	$Sc_{t,1}$	0.4
	$Sc_{t,2}$	0.7
Variance of mixture fraction	$C_{\xi,1}$	0.2
	$C_{\xi,2}$	0.3
CMC formulation	CMC-1	Conservative CMC
	CMC-2	Non-conservative CMC
Convective flux	flux-1	Computing fluxes based on LES cells at CMC faces
	flux-2	Computing fluxes based on CMC cell centers
Conditionally filtered turbulent diffusivity	$D_{\eta,1}$	Modelling D_{η} based on CMC cells
	$D_{\eta,2}$	Modelling D_{η} based on LES cells
	$D_{\eta,3}$	Modelling D_{η} with adjusting the length scale
Number of CMC cells	$4 \times 4 \times 80$	4 CMC cells in X- and Y-directions with 80 CMC cells in Z-direction
	$8 \times 8 \times 80$	8 CMC cells in X- and Y-directions with 80 CMC cells in Z-direction
	$16 \times 16 \times 80$	16 CMC cells in X- and Y-directions with 80 CMC cells in Z-direction

of the mean axial velocity and mixture fraction around $1 < r/D < 2$ at position $z/D = 15$ may require more simulation time for more precision. However, the current predictions provide a good basis for the parametric studies of combustion model.




Fig. 2 Snapshots of the temperature field in total computational domain (*left*) and in the upstream region (*right*) for Sandia Flame D. The iso-contour of stoichiometric mixture fraction is presented by the *black lines*

2.3.2 Parametric Studies of the Combustion Model

As shown in Table 2, five case studies are performed to show the effects of each case in CMC model. All cases are based on the optimal conditions from parametric studies of flow and mixing field (variance-2, Sc_2 , $Sc_{t,1}$ and $C_{\xi,1}$) and use $8 \times 8 \times 80$ for the number of CMC cells. A reference case (case-1) includes the models CMC-1, flux-1 and $D_{\eta,2}$, while other cases have at least one varied parameter compared with the reference case.

Preliminary Studies

Preliminary studies of parametric studies of combustion model are required to choose the cases which may predict the simulation results different from the reference case (case-1 from Table 2). These case studies will be further examined for the simulation results in next sections. In the first step, a representative direction has to be defined. Having the highest convective value among three directions, the convective flux in z -direction is chosen as a representative. Note that a consideration of FDF profile is required, since it is applied to transfer the values from a mixture fraction space to a physical space. The low value of FDF means only a small influence of any property can appear in the physical space.

Subsequently, the convective fluxes in z -direction are shown in radial and axial distributions. The instantaneous predictions of CH_4 fluxes in Fig. 4 show that the radial distribution exhibits larger differences of the convective fluxes between each case than the axial distribution. Thus, the radial distribution is applied for investigations of convective flux comparisons for the other species. It should be reminded that a consideration of FDF is necessary as previously discussed.

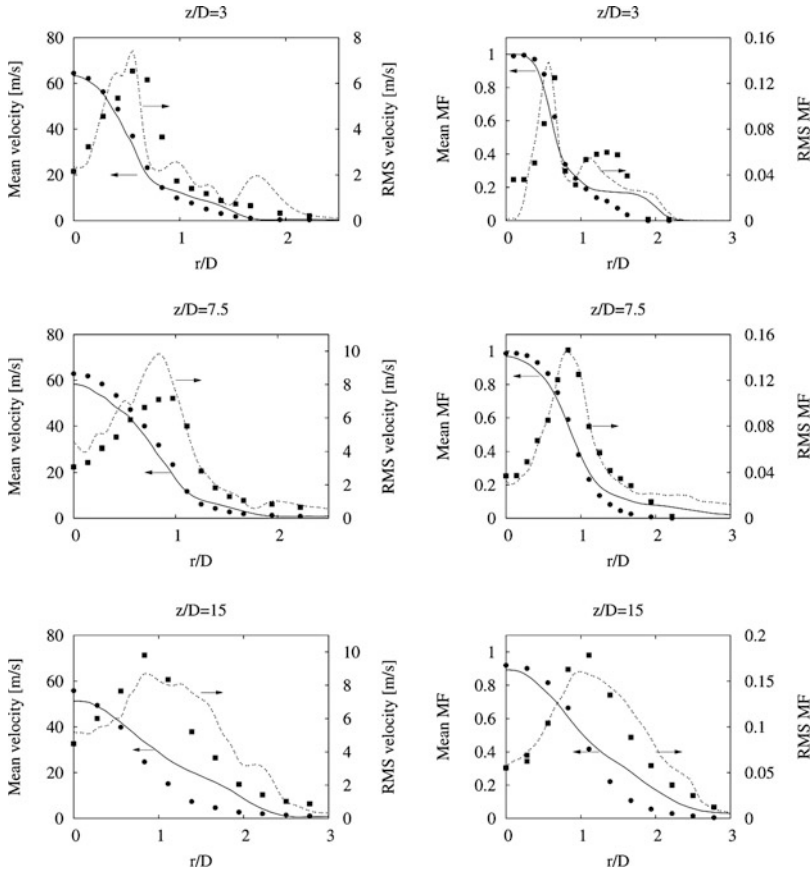


Fig. 3 Radial profiles of mean and RMS axial velocity and mixture fraction at three downstream locations for Flame D. *Symbols* denote experimental values [3,4], while the *solid* and *dashed lines* present the mean and RMS values of LES-CMC (reference case of Table 2)

Table 2 Summary of different parameters in combustion model study. The meaning of each numerical method can be found in Table 1

Name	Combustion model	Flow and mixing field	CMC grid resolution
case-1 (reference case)	CMC-1, flux-1, $D_{\eta,2}$		
case-2	CMC-2, flux-1, $D_{\eta,2}$	variance-2	
case-3	CMC-1, flux-2, $D_{\eta,2}$	$Sc_2, Sc_{t,1}$	
case-4	CMC-1, flux-2, $D_{\eta,1}$	$C_{\xi,1}$	
case-5	CMC-1, flux-1, $D_{\eta,3}$		$8 \times 8 \times 80$

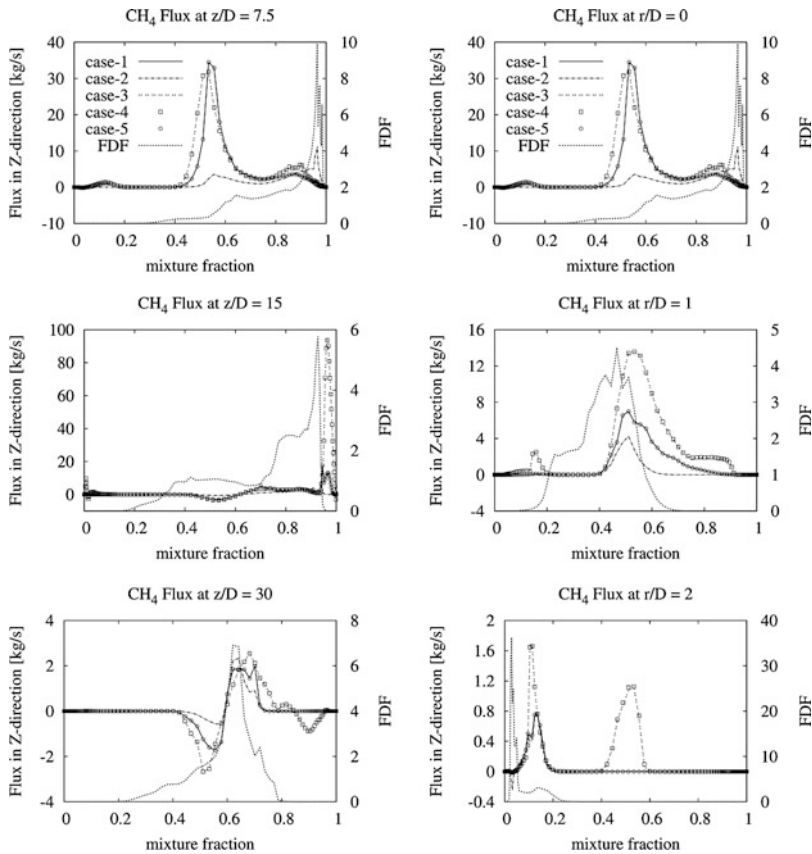


Fig. 4 Axial (*left*) distribution along the centerline and radial (*right*) distribution at $z/D = 7.5$ of convective fluxes in z-direction of CH_4 for a time step (Sandia Flame D)

The investigations of fluxes in other species show that fluxes in different species follow the same tendency of CH_4 fluxes in Fig. 4. It can be observed that case-2, case-3 and case-4 produce different convective fluxes compared with the reference case (case-1). However, case-3 and case-4 produce similar flux, which means their conditional scalar predictions should not be different. Therefore, case-3 is chosen for the further parametric studies. Fluxes of case-1 and case-5 are similar and also case-3 and case-4 are similar because of low effects of different D_η models on the convective term. Since case-1, case-2 and case-3 produce different convective fluxes, the conditional scalar predictions from these cases should differ from each other.

In next step, the statistical results of three cases (case-1, case-2 and case-3) are sampled over 30,000 time steps for the statistics to investigate the different effects of combustion model parameters and to validate with the experiment.

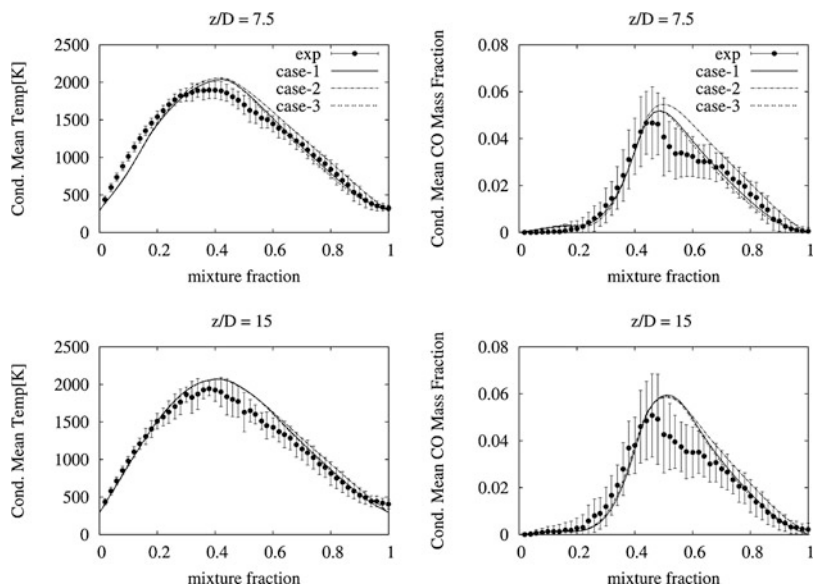


Fig. 5 Conditional profiles of cross-sectionally averaged temperature and CO at two different downstream positions in mixture fraction space for Flame D. Symbols are experimental data [4], while the *solid*, *dashed* and *dotted lines* present the results of LES-CMC in different cases in combustion model (Table 2)

Conditionally Filtered Reactive Scalars

As in CMC methodology, species mass fraction are analyzed in mixture fraction space, the efficiency of the combustion model can be decoupled from flowfield predictions. Therefore, the performance of each case of combustion model can be directly considered from conditional profiles. Initial values for the conditional reactive species are obtained from the SLFM solution.

A good agreement of case studies with experiments is shown in Figs. 5 and 6. The conditional mean temperature and the conditional mean mass fraction of CO, CH₄ and H₂ are given in both figures as a representative of intermediate products, fuel and radicals. Note that the error bars indicate the conditional RMS and they are only plotted to illustrate the turbulent level of each scalar. The reason of different predictions between case-1 and case-2 on the rich side ($\eta > 0.35$) belongs to two different sets of convective fluxes which are calculated from two CMC formulations. Because of the lack of FDF-weighting function in convective term, the low convective fluxes on the rich side are generated in the upstream positions of case-2. These can be observed in Fig. 7. Therefore, case-2 (non-conservative CMC) usually overpredicts on the rich side of mixture fraction in temperature, radical and intermediate product, while the underpredictions occur in the fuel.

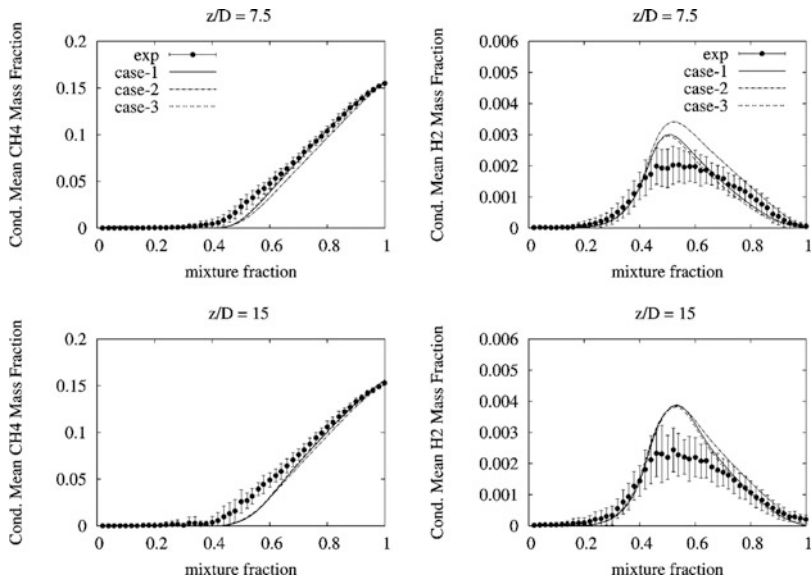


Fig. 6 Conditional profiles of cross-sectionally averaged CH₄ and H₂ at two different downstream positions in mixture fraction space for Flame D. *Symbols* are experimental data [4], while the *solid*, *dashed* and *dotted lines* present the results of LES-CMC in different cases in combustion model (Table 2)

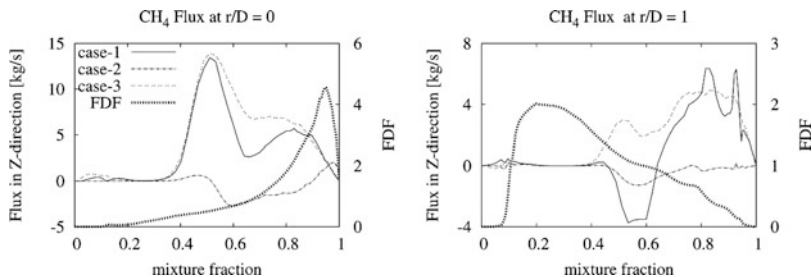


Fig. 7 Radial distribution of mean convective fluxes in z-direction of CH₄ at $z/D = 7.5$ (Sandia Flame D). The *solid*, *dashed* and *dotted lines* present the results of LES-CMC in different cases and FDF (reference case) in combustion model (Table 2)

It can be observed from the Figs. 5 and 6 that case-1 (conservative CMC) shows more accurate results than case-2. Cross-sectional simulation results from case-3 hardly differ from case-1, even though the convective fluxes from both cases are different (Fig. 7). The reason can be explained by using FDF value for each CMC cell in the same cross section to calculate the cross-sectional averages. If the conditional predictions between two cases of any CMC cell have a difference where the low FDF is calculated in mixture fraction space, the conditional predictions in cross-sectional averages will be similar.

Table 3 Summary of different cases in CMC grid resolution study

Name	CMC grid resolution	Flow and mixing fields	Combustion model
res-1	$4 \times 4 \times 80$	variance-2	
res-2 (reference case)	$8 \times 8 \times 80$	$Sc_2, Sc_{t,1}$	CMC-1, flux-1 $D_{\eta,2}$
res-3	$16 \times 16 \times 80$	$C_{\xi,1}$	

2.3.3 Parametric Study of CMC Grid Resolution

As described in Table 1, three cases of CMC grid resolution are varied, while the same conditions of flow and mixing field (variance-2, Sc_2 , $Sc_{t,1}$ and $C_{\xi,1}$) and CMC combustion model (CMC-1, flux-1 and $D_{\eta,2}$) are set up. The variations of the CMC cells in each x - and y - direction are 4 cells for case-1, 8 cells for case-2 (reference case) and 16 cells for case-3 with the same 80 CMC cells in z - direction, as summarized in Table 3.

Since the CMC resolution varies in the radial distribution for three case studies, the radial distribution of conditional value should show more prominent features. Therefore, the radial distribution of mean scalar are investigated at position $z/D = 3, 7.5$ and 15. The mean temperature and CO predictions in radial distributions are shown in Fig. 8.

It can be seen from position $z/D = 3$ that res-2 and res-3 perform better than res-1 since there is an underprediction of the temperature for res-1 in this position. Predictions of res-3 can capture the highest value of CO at position $z/D = 7.5$. Moreover, predictions from res-3 ($16 \times 16 \times 80$ for CMC cells) match better with the experiments than the others at position $z/D = 15$ which show a great advantage of small CMC cells in this resolution. A reason may relate to the size of CMC cell which the big size of CMC cell may predict inaccurately in which a high level of mixture fraction gradient occurs. However, an increasing CMC resolution from res-2 to res-3 requires more computational time than 60%. Considering the computational time and results from all CMC resolutions, the appropriate resolution is res-2 ($8 \times 8 \times 80$ for CMC cells) for Flame D.

2.4 Summary

In this section, parametric studies of LES-CMC are carried out for the Sandia Flame D. The parameter studies comprised investigations of flow and mixing field, variants of the CMC combustion model parameters and CMC grid resolution.

Flame D is used to investigate the influences of various flow and mixing field parameters on the simulation results. These parameters, which are $Sc = 0.7$, $Sc_t = 0.4$ and $C_{\xi} = 0.2$, are optimal values and thus, they are used for further

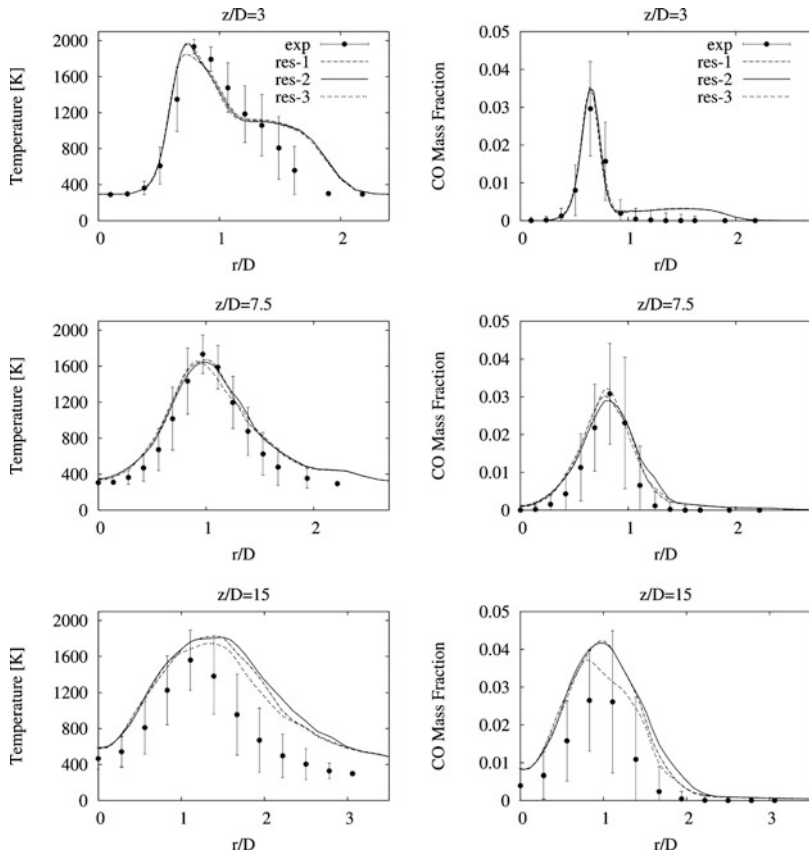


Fig. 8 Radial profiles of mean temperature and CO for Flame D. *Symbols* are experimental data [4], while the *solid, dashed and dotted lines* present the results of LES-CMC from the different CMC grid resolutions (Table 3)

studies of Flame D, as well as for Flames E and F. Values of velocity variance of inflow generator, however, depend on the physical inflow of each flame. It should be noted that an adjustment of the velocity variances is carried out to reduce the high level of turbulence which may come from the implementation of inflow generator. Suitable inflow variances for Flame D are found to be $\frac{2}{3}$ of the measured variances at $z/D = 0.14$ of these flames, respectively.

Parameter studies of different variants of the CMC combustion model are carried out to find the most suitable model. Initial studies of the CMC fluxes have shown that the effect of turbulent diffusivity modelling is negligible. However, a comparison of CMC models, which varies the CMC formulation (conservative vs. non-conservative), reveals considerable differences. Moreover, some slight differences between two methods of CMC convective flux approximation (cell face vs. cell centre based) are detected. Therefore, three dominant cases which differ in both

Table 4 Test cases for the LES-CMC model

	Test cases		
	LES cells	CMC cells	Flux implementation
Bluff-body flame, HM1	$218 \times 218 \times 320$	$16 \times 16 \times 80$ $24 \times 24 \times 80$	Various
4 lifted flames	$96 \times 96 \times 480$	$16 \times 16 \times 80$ $16 \times 16 \times 160$	Various

numerical aspects are investigated for further studies. Conditional mean scalars, which are averaged in the same cross section, show that the conservative CMC formulation with computing convective fluxes based on LES cells located at the CMC cell faces (case-1) is similar to the one with computing convective fluxes based on CMC cell centers (case-3). This can be explained by the low FDF values where the differences of predictions occur in a CMC cell. Consequently, the conditionally averaged predictions with FDF weighting create the similar results over a cross section. Generally, conditional predictions reveal that case-1 can capture better mean measurements than case-2 (the non-conservative formulation using the same flux approximation). This is because the variation of FDF-weighted convective fluxes in different directions of case-1 allows the predictions to be more accurate.

Three different CMC grid resolutions ($4 \times 4 \times 80$, $8 \times 8 \times 80$ and $16 \times 16 \times 80$) are examined in order to find the appropriate number of CMC cells for Sandia Flame D. Basically, the best predictions are found in CMC grid resolution of $16 \times 16 \times 80$. However, the reasonable resolution for Flame D is $8 \times 8 \times 80$ CMC cell due to the computational cost with efficient performance.

3 Future Study Cases

In order to show advantages of conservative CMC, more complicated flames are required for the simulations. Therefore, future test cases will be the Sydney bluff-body flame, HM1, with $218 \times 218 \times 300$ cells for LES for two CMC mesh sensitivity studies. Moreover, four lifted flames investigated at Berkeley [5, 6] and Calgary [7] will be examined (two Berkeley flames and two Calgary flames). The summation of all test cases can be found in Table 4.

4 The Usage of Computational Resources

The fine grid simulation results, which are presented here, have been performed on NEC Nehalem Cluster with 80 processors due to scalability tests for LES-CMC. The using of parallel program MPI and vectorization compiler let the code runs

faster. An approximate wall time is around 48 h per summited job. The analysis of all parametric studies from Table 1, of Flames D, E and F corresponds to 450,000 CPU hours. The amount of CPU use for Sandia Flame series is appropriate to the requirements of computational resources, 450,000 CPU hours for all test cases in future.

References

1. S. Navarro-Martinez, A. Kronenburg and F. Di Mare. Conditional moment closure for large eddy simulations. *Flow, Turb. Combust.*, 75:245–274, 2005.
2. S. Navarro-Martinez and A. Kronenburg. LES-CMC simulations of a turbulent bluff-body flame. *Proc. Combust. Inst.*, 31:1721–1728, 2007.
3. Ch. Schneider, A. Dreizler, J. Janicka and E.P. Hassel. Flow field measurements of stable and locally extinguishing hydrocarbon-fueled jet flames. *Combust. Flame*, 135:185–190, 2003.
4. R.S. Barlow and J. Frank. Piloted CH_4 /Air Flames C,D,E, and F - Release 2.1. *CA*, 94551–0969, 2007.
5. R. Cabra, T. Myrvold, J.-Y. Chen, R. W. Dibble, A. N. Karpets, and R. S. Barlow Simultaneous laser Raman-Rayleigh-LIF measurements and numerical modeling results of a lifted turbulent H_2/H_2 jet flame in a vitiated co-flow. *Proc. Combust. Inst.*, 29:1881–1888, 2002.
6. R. Cabra, J.-Y. Chen, R. W. Dibble, A. N. Karpets, and R. S. Barlow. Lifted methane-air jet flames in vitiated co-flow. *Combust. Flame*, 143:491–506, 2005.
7. T. Leung and I. Wierzba The effect of co-flow stream velocity on turbulent non-premixed jet flame stability. *Proc. Comb. Inst.*, 32(2):1671–1678, 2009.

Numerical Investigation of a Complete Scramjet Demonstrator Model for Experimental Testing Under Flight Conditions

Yann Simsont, Peter Gerlinger, and Manfred Aigner

Abstract In the present paper a complete scramjet demonstrator model for experimental testing at Mach 8 is investigated numerically using the scientific code TASCUM3D (Turbulent All Speed Combustion Multigrid Solver) on the HPC vector system NEC SX-9, installed at the High Performance Computing Center Stuttgart (HLRS). First the three-dimensional intake of the model is simulated. Then the results are used as inlet conditions for the simulation of the combustor, where hydrogen is injected by a lobed strut injector located in the middle of the diverging chamber. The test conditions, corresponding to a flight speed of Mach 8 at an altitude of 30 km, and the results of the simulations are discussed in detail, occurring difficulties are highlighted. In order to ensure self-ignition and prevent the flow from blockage, potential design changes are described and investigated in order to proof their functionality. Finally the performance of the reactive and non-reactive simulations on the NEC SX-9 is analyzed.

1 Introduction

For future hypersonic transportation the use of air breathing engines (i.e. ramjets for flight Mach numbers 2–7 and scramjets for flight Mach numbers 5–15) is of great interest. Contrary to rocket driven systems no oxygen is transported. Therefore air breathing engines provide the opportunity to increase the payload to total mass ratio, and thus reduce the cost per payload unit. The main objective of the Research Training Group GRK 1095/2 (University of Stuttgart, Technical University of Munich, RWTH Aachen University and DLR Cologne) is the design and development of a complete scramjet demonstrator model. Experimental testing of the model is

Y. Simsont (✉) · P. Gerlinger · M. Aigner
Institut für Verbrennungstechnik der Luft- und Raumfahrt, Universität Stuttgart, Pfaffenwaldring
38-40, 70569 Stuttgart, Germany
e-mail: yann.simsont@dlr.de



Fig. 1 Scramjet demonstrator model in the hypersonic test facility IT-302

funded by the DFG grant GA 1332/1-1 and conducted in two hypersonic wind tunnels (IT-302 and AT-303) at the Khristianovich Institute of Theoretical and Applied Mechanics (ITAM), Russian Academy of Sciences, Siberian Branch in Novosibirsk, Russia. A first testing period has been accomplished in October and November 2011, a second testing period has been performed in March and April 2012. Figure 1 shows the complete scramjet demonstrator model mounted in the IT-302 wind tunnel. In the present paper the final numerical simulations before experimental testing are described. Those simulations are directed at distinguishing unfavorable flow conditions, which might occur during testing (e.g. blockage of the intake, deficient self-ignition in the combustor, thermal blockage, etc.), as well as suitable correctives to avoid those phenomena and ensure successful experiments. In respect of the advanced stage of the project – the scramjet demonstrator model had already been assembled – potential design modifications derived from the numerical results have to be easy to implement.

2 Geometry and Test Design

The investigated scramjet demonstrator model consists of an intake, an isolator, a combustor with central strut injection and a nozzle (Fig. 2) and has a total length of 1.046 m. The intake combines a single outer compression ramp (15.5° angle) and side wall compression (3.5° angle) and has been developed based on a 3D mixed compression intake tested at ITAM [1] and various numerical simulations [2]. Gaseous hydrogen is injected at $x = 623$ mm downstream of the ramp leading edge in axial flow direction at the trailing edge of a lobed strut injector through seven horizontal and six vertical ports shown in Fig. 3. The strut is mounted centrally in the model from one sidewall to the other and corresponds to previously investigated lobed strut injectors [3, 4]. The lobed structure creates streamwise vortices and therefore enhances the mixing of fuel and air. The combustion chamber has a rectangular cross section, the top and bottom walls diverge – each with an angle of 2° – beginning at $x = 580$ mm (at half length of the strut).




Fig. 2 Drawing of longitudinal section of the scramjet demonstrator model with instrumentation




Fig. 3 Sketch of central lobed strut injector (top) and areas of hydrogen injection at the trailing edge of the injector highlighted in blue (bottom)

Table 1 Inflow conditions for experimental testing at Mach 8

Air conditions	A	B
Mach number Ma_∞ (—)	8	8
Velocity u_∞ ($\frac{m}{s}$)	2,472	2,201
Total pressure p_0 (bar)	110	110
Static pressure p_∞ (Pa)	1,127	1,127
Total temperature T_0 (K)	3,280	2,600
Static temperature T_∞ (K)	238	188
Wall temperature T_{wall} (K)	293	293

The intake ramp is equipped with 15 thermocouples and two static pressure transducers along the center line. Over 30 pressure transducers are integrated in the top and bottom wall of the combustor and a Pitot rake is mounted at the end of the combustor to measure the Mach number at 5 points of the cross section. The test conditions, corresponding to a flight speed of Mach 8 at an altitude of 30 km, are summarized in column A of Table 1. Column B lists the conditions for a lower total temperature T_0 also tested experimentally. Since the measurement duration in the wind tunnels is relatively short (about 100 ms) the wall temperature of the model is assumed to be constant. To benefit from the symmetry of the investigated scramjet model and economize computation time, only one half of the model is simulated using approximately 24 million structured grid cells (14.4 million for the intake and 9.6 million for the combustor).

3 Governing Equations and Numerical Scheme

The investigations presented in this paper are performed using the scientific code TASC3D. The code has been used successfully in the last two decades simulating reacting and non-reacting flows. It describes reacting flows by solving the full compressible Navier-Stokes, species and turbulence transport equations. Additionally an assumed PDF (probability density function) approach is used to take turbulence chemistry interaction into account. The set of averaged equations in three-dimensional conservative form is given by

$$\frac{\partial \mathbf{Q}}{\partial t} + \frac{\partial(\mathbf{F} - \mathbf{F}_v)}{\partial x} + \frac{\partial(\mathbf{G} - \mathbf{G}_v)}{\partial y} + \frac{\partial(\mathbf{H} - \mathbf{H}_v)}{\partial z} = \mathbf{S}, \quad (1)$$

where

$$\mathbf{Q} = [\bar{\rho}, \bar{\rho}\tilde{u}, \bar{\rho}\tilde{v}, \bar{\rho}\tilde{w}, \bar{\rho}\tilde{E}, \bar{\rho}q, \bar{\rho}\omega, \bar{\rho}\sigma_T, \bar{\rho}\sigma_Y, \bar{\rho}\tilde{Y}_i]^T, \quad i = 1, 2, \dots, N_k - 1. \quad (2)$$

The variables in the conservative variable vector \mathbf{Q} are the density $\bar{\rho}$ (averaged), the velocity components (Favre averaged) \tilde{u} , \tilde{v} and \tilde{w} , the total specific energy \tilde{E} , the turbulence variables $q = \sqrt{k}$ and $\omega = \epsilon/k$ (where k is the kinetic energy and ϵ the dissipation rate of k), the variance of the temperature σ_T and the variance of the sum of the species mass fractions σ_Y and finally the species mass fractions \tilde{Y}_i ($i = 1, 2, \dots, N_k - 1$). Thereby N_k describes the total number of species that are used for the description of the gas composition. The vectors \mathbf{F} , \mathbf{G} and \mathbf{H} specify the inviscid fluxes in x -, y - and z -direction, \mathbf{F}_v , \mathbf{G}_v and \mathbf{H}_v the viscous fluxes, respectively. The source vector \mathbf{S} in Eq.(1) includes terms from turbulence and chemistry and is given by

$$\mathbf{S} = [0, 0, 0, 0, 0, \bar{S}_q, \bar{S}_\omega, \bar{S}_{\sigma_T}, \bar{S}_{\sigma_Y}, \bar{S}_{Y_i}]^T, \quad i = 1, 2, \dots, N_k - 1, \quad (3)$$

where \bar{S}_q and \bar{S}_ω are the averaged source terms of the turbulence variables, \bar{S}_{σ_T} and \bar{S}_{σ_Y} the source terms of the variance variables (σ_T and σ_Y) and \bar{S}_{Y_i} the source terms of the species mass fractions. For turbulence closure a two-equation low-Reynolds-number q - ω turbulence model is applied [5]. The momentary chemical production rate of species i in Eq.(3) is given by

$$S_{Y_i} = M_i \sum_{r=1}^{N_r} \left[(v''_{i,r} - v'_{i,r}) \left(k_{fr} \prod_{l=1}^{N_k} c_l^{v'_{l,r}} - k_{br} \prod_{l=1}^{N_k} c_l^{v''_{l,r}} \right) \right], \quad (4)$$

where k_{fr} and k_{br} are the forward and backward rate constants of reaction r (defined by the Arrhenius function), the molecular weight of a species M_i , the species concentration $c_i = \rho\tilde{Y}_i/M_i$ and the stoichiometric coefficients $v'_{i,r}$ and $v''_{i,r}$ of species i in reaction r . The averaged chemical production rate for a species i due

to the use of an assumed PDF approach is described in detail in Refs. [6, 7]. In the present paper the reactive simulations have been performed using a modified Jachimowski hydrogen/air reaction mechanism with 9 species and 19 steps [8, 9]. The unsteady set of differential equations in Eq. (1) is solved using an implicit lower-upper symmetric Gauss-Seidel (LU-SGS) [8, 10–12] finite-volume algorithm, where the finite-rate chemistry is treated fully coupled with the fluid motion. More details concerning TASC3D may be found in Refs. [7, 8, 12–14].

4 Investigation of the 3D Intake

In this section the three-dimensional intake and the subsequent isolator are investigated using the experimental test conditions listed in Table 1 in column A. The main focus is the generation of convenient flow conditions to enable self-ignition in the combustor while any blockage must be prevented. Figure 4 shows the Mach number distribution in the central plane of the model. The compression of the incoming air is effected by two strong shock waves: the first is the leading ramp shock which shortly misses the cowls lip, the second shock wave originates at the cowl and impinges on the thick boundary layer on the bottom wall of the model close to the central strut injectors leading edge. The interaction of the cowl shock and the boundary layer induces a separation zone with low Mach numbers shown in detail in the lower part of Fig. 4. Although the extent of the separation in the spanwise cross section is relatively small, the flow beneath the central strut injector is perturbed. In order to avoid the impingement of the cowl shock on the boundary layer a passive suction slot is added to the intake model at $x = 525$ mm.

To determine the size of the opening two different designs are investigated numerically. Suction I is 15 mm long and has an angle of 45° at the leading edge and 20° at the trailing edge, suction II is 25 mm long, has the same angle of 45° at the leading edge, though an angle of 37° at the trailing edge. The Mach number distribution in the central plane upstream of the strut is given for both slot configurations in Fig. 5. As intended suction I as well as suction II capture the cowl shock. Consequently the separation zone is successfully eliminated by the suction. Table 2 lists the air mass flows for the intake and the two suction slot configurations. The loss of compressed air exiting through the suction slot is minor for both slot designs with only 1.6 % (suction I) and 4.6 % (suction II) of the total captured mass flow. Because of the enhanced length of the slot, suction II is considered to be more tolerant towards any unsteadiness or modification of the simulated test conditions and is therefore recommended for the experimental testing. The conditions at the interface between isolator and combustor (cross section at half length of the central strut injector, where $x = 580$ mm) shown in Fig. 6 represent a strongly three-dimensional and non-uniform flow field due to the 3D intake. The average levels of pressure ($p_{av} = 0.522$ bar) and temperature ($T_{av} = 1,090$ K)

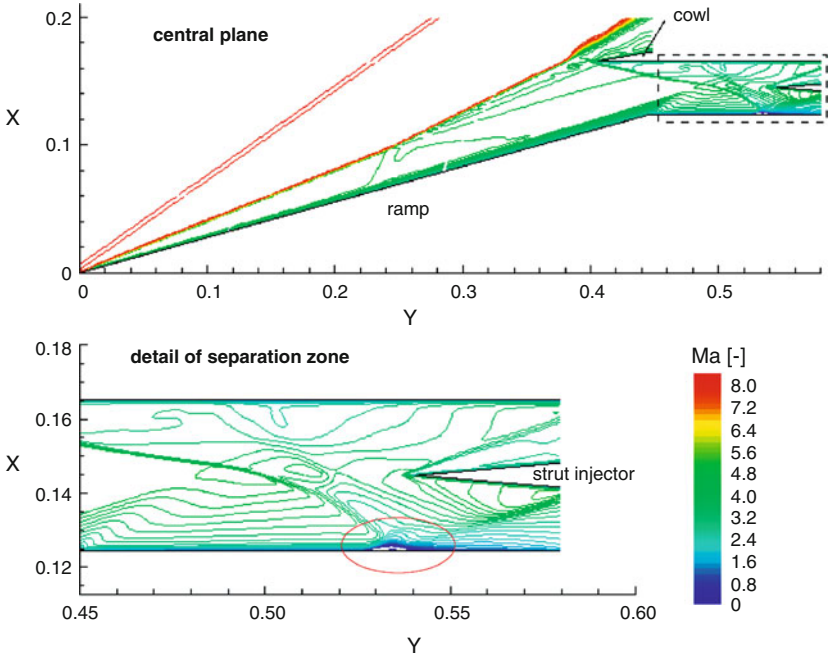


Fig. 4 Mach number distribution in the central plane of the intake (*top*) and detail of the impinging cowl shock at the *bottom* wall causing a separation zone (*bottom*)

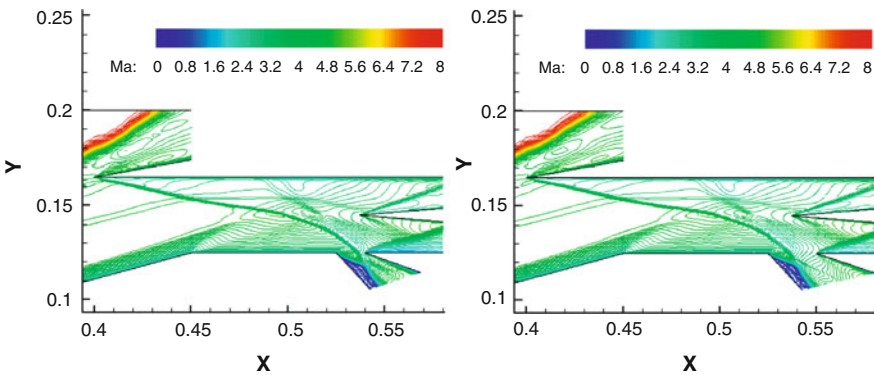
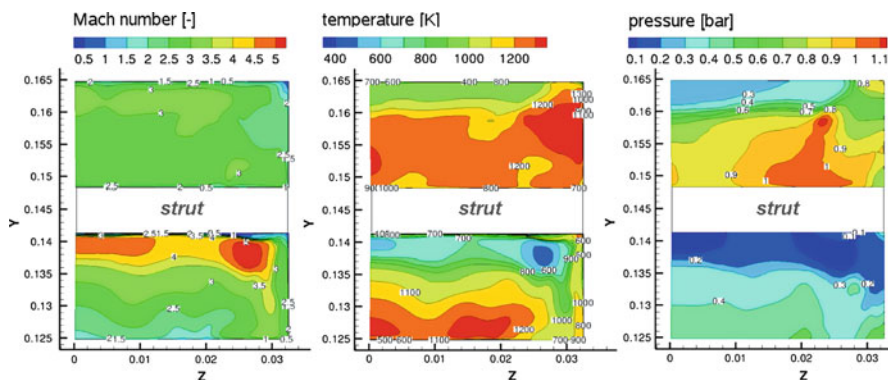


Fig. 5 Mach number distribution in the central plane for suction I (*left*) and suction II (*right*)

are relatively low, while the Mach number is high ($Ma_{av} = 3.02$). With respect to the aspired self-ignition the conditions are particularly poor in the region directly beneath the strut where the temperature is about 800 K and less. These unfavorable conditions are intensified downstream of the interface, where the divergent geometry of the combustor further accelerates the flow and decreases the level of temperature.

Table 2 Calculated air mass flows for intake simulations with passive suction

	Total ($\frac{g}{s}$)	Relative (%)
Mass flow intake \dot{m}_{inflow}	636	100.0
Mass flow suction I $\dot{m}_{suction I}$	10	1.6
Mass flow suction II $\dot{m}_{suction II}$	29	4.6

**Fig. 6** Mach number (*left*), temperature (*middle*) and pressure (*right*) distribution at the cross section at $x = 580$ mm (interface between intake and combustor simulation), respectively**Table 3** Inflow conditions for the hydrogen injection into the combustor (horizontal and vertical ports)

	Hydrogen
Mach number Ma_{H_2} (—)	2.15
Equivalence ratio ϕ	0.6
Mass flow \dot{m}_{H_2} ($\frac{g}{s}$)	10.8
Total pressure $p_{H_2,0}$ (bar)	7.0
Static pressure p_{H_2} (bar)	0.7
Total temperature $T_{H_2,0}$ (K)	293.0
Static temperature T_{H_2} (K)	145.0
Wall temperature strut T_{wall} (K)	293.0

5 Investigation of the Combustor

The conditions for the hydrogen injection are given in Table 3, the conditions for the incoming air flow at the interface are adopted from the intake simulation. In consequence of the suboptimal conditions at the interface discussed at the end of Sect. 4 and the further decrease in temperature due to the diverging combustion chamber, no self-ignition is achieved during the simulation of the combustor. In order to enforce ignition with a design change still feasible at this stage of the project, ignition wedges are integrated in the model. Placed at the top and bottom walls of the combustor close to the trailing edge of the strut injector, they are




Fig. 7 Sketch of the ignition wedges and their position in the model

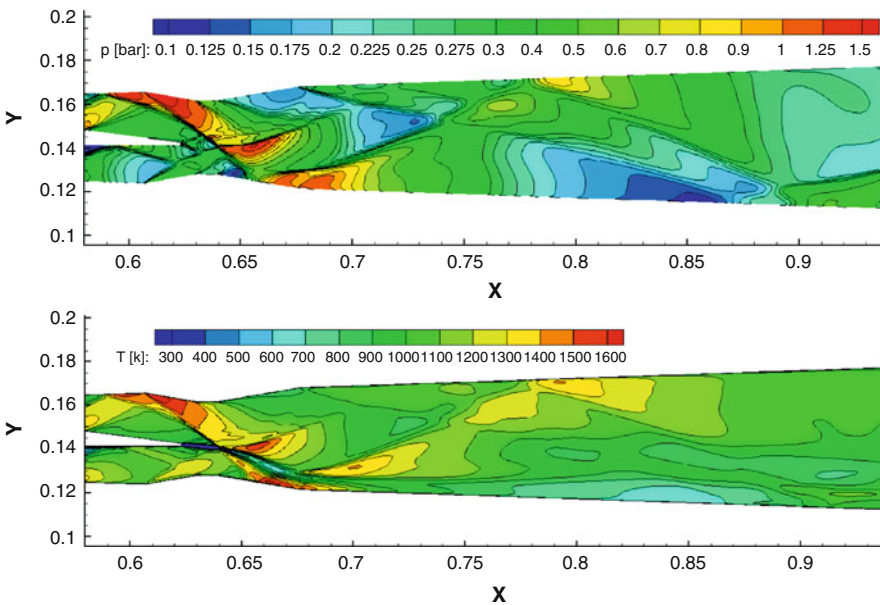


Fig. 8 Pressure (*top*) and temperature (*bottom*) distribution in the central plane of the combustor with ignition wedges (non-reactive simulation), respectively

designed to initiate a shock system that locally increases the temperature close to the injection and, thus, enables ignition. Figure 7 sketches the geometry and the position of the ignition wedges.

The evolving shock system due to the ignition wedges is observed in the pressure and temperature distributions in the central plane in Fig. 8 (non-reactive simulation). Enforced by a reflecting shock wave in the upper region of the incoming flow, the shock wave originating from the front ramp of the top wedge is more dominant than the one generated by the bottom wedge. The top wedge shock crosses the cold injected hydrogen about 16 mm downstream of the injector's trailing edge,

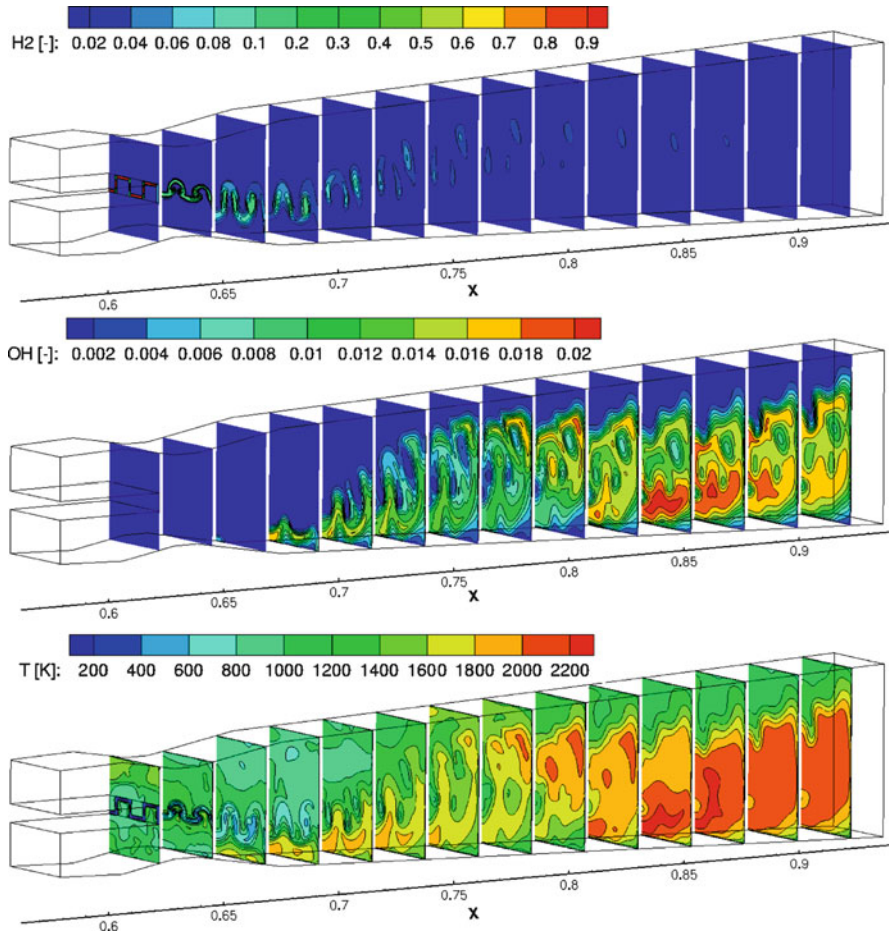


Fig. 9 H₂ (*top*), OH (*middle*) and temperature (*bottom*) distributions at cross sections of the combustor (reactive simulation), respectively

thereby deflecting the fuel towards the lower half of the combustion chamber. Further downstream the shock wave is reflected by the bottom and top walls. The temperature distribution shows three local maxima above 1,500 K: at the front ramp of the top wedge, at about 30 mm downstream of the injector and nearby the rear ramp of the bottom wedge. The upper part of Fig. 9 shows the H₂ distribution for the reacting flow at several cross sections of the model. The deflection of hydrogen by the shock waves is observed as well as the mixing of fuel and air supported by the strong vortices which are induced by the lobed strut. At the hot spot close to the rear ramp of the bottom wedge, hydrogen and air are sufficiently mixed in order to ignite. Accordingly we observe a significant increase in the temperature and OH mass fraction distributions close to the bottom wall (middle and bottom part of

Table 4 Computational performance of the simulations for the ITAM hypersonic wind tunnel experiments

	Intake	Combustor non-reactive	Combustor reactive
Number of volumes (Mio)	14.4	9.6	9.6
Number of species	2	3	9
Number of iterations	140,000	130,000	80,000 (on basis of non-reactive solution)
Vector op. ratio (%)	98.16	98.68	99.12
Average vector length	105.3	191.8	212.5
MFLOPS	4,329	4,912	7,204
Quota peak perf. (%)	4.2	4.8	7.0
Wall-clock time/iter. (s)	1.008	0.744	1.538
Total CPU time (h)	627	430	547

Fig. 9). Further downstream the flame spreads over the whole cross section and a non-uniform, detached flame develops. Therefore the simulation gives evidence of the successful use of the ignition wedges, while no thermal blockage is observed. Due to the short experimental testing times the high temperatures close to the bottom wall are not considered to be critical for the model. Numerical simulations with a reduced total temperature (inflow conditions listed in column B in Table 1) indicate very low temperatures and high Mach numbers at the interface between intake and combustor ($T_{av} = 852$ K, $Ma_{av} = 3.14$). In spite of these unfavorable conditions, the ignition wedges again increase the temperature locally and enforce ignition.

6 Performance Analysis

The numerical investigations of the intake and the combustor have been performed on the NEC SX-9 at the High Performance Computing Center Stuttgart using 1 node and 16 CPUs. Table 4 gives an overview of the performance for the non-reactive and reactive simulations. The average vector length is longer for simulations with a high number of species. Accordingly the best vector performance is achieved by the reactive simulation of the combustor using nine species (vector operation ratio of 99.12%). The low quotas of peak performance on the NEC SX-9 (4.2–7.0%), particularly in comparison to the previous vector-processor based HPC, the NEC SX-8, have already been discussed in Ref. [15].

Table 5 lists the five most time consuming subroutines (using about 73.6% of the computational time) for the reactive simulation of the combustor and their performance data. All subroutines show high vector operation ratios (98.69–99.87%), but the performance varies between 3,413 and 11,247 MFLOPS (3.3–11.0% peak performance) due to significant differences with regard to the bank conflicts. On the right hand side (RHS) of the set of equations subroutine PROP calculates the

Table 5 Performance data for the most time consuming subroutines for the reactive simulation of the combustor

Subroutine	Time (%)	MFLOPS	Vec. op. ratio (%)	Av. vec. length	Bank confl.	Quota peak perf. (%)
PROP	20.8	8,035	98.69	255.9	1.74	7.8
LINE3D	17.5	3,413	99.25	223.3	534.54	3.3
REACTION	16.8	11,247	99.41	240.0	22.93	11.0
LFSWEEP	9.3	3,627	99.87	219.9	647.65	3.5
UFSWEEP	9.2	3,635	99.87	219.7	647.90	3.5

gas properties and subroutine REACTION the chemical source terms. As these are local phenomena, both subroutines only depend on the data of each volume. On the contrary the implicit left hand side (LHS) of the set of equations is solved using an implicit lower-upper symmetric Gauss-Seidel solver (LU-SGS) [8, 10–12]. In order to resolve the data dependency from neighboring cells of the structured i, j, k -ordered grid the LHS subroutines LINE3D, UFSWEEP and LFSWEEP are vectorized along hyperplanes defined by $i + j + k = const.$ using indirect addressing. This probably causes memory latencies and therefore more bank conflicts per iteration (0.5345–0.6479 for the subroutines LINE3D, UFSWEEP and LFSWEEP in comparison to only 0.0017–0.0229 for the subroutines PROP and REACTION).

7 Conclusion

A complete scramjet demonstrator model designated for wind tunnel testing at Mach 8 has been investigated numerically. A separation zone evolving from the interaction of the cowl shock and the boundary layer at the bottom wall of the model has been successfully eliminated by passive suction. The mass flow exiting the model through the suction slot has been evaluated. The first simulation of the combustor showed that the conditions were too cold for self-ignition. Therefore ignition wedges have been integrated in the model. It has been shown for two different test conditions, that the evolving shock system increases the temperature locally and enables self-ignition and a stable combustion.

The performance of the simulations of the intake and the combustor on the NEC SX-9 system has been analyzed. A good vector performance has been shown, especially for the reactive simulation of the combustor. Still, the high number of bank conflicts resulting from the indirect addressing of the LHS has to be further investigated and decreased. When the experimental data is fully analyzed, further simulations reproducing the exact testing conditions have to be performed in order to compare the experimental and the numerical results quantitatively. Furthermore it is intended to port the simulations of the scramjet demonstrator model to the massive parallel scalar Cray XE6 system using more CPUs.

Acknowledgements This work has been supported by the German Research Foundation (DFG) within the GRK 1095/2. The experimental testing has been funded by the DFG grant GA 1332/1-1. The authors would like to thank Dr. Uwe Gaisbauer, Institute of Aerodynamics and Gas Dynamics, University of Stuttgart, who is in charge of the experiments of the project, for the collaboration and support. The simulations have been performed on the national super computer NEC SX-9 at the High Performance Computing Center Stuttgart (HLRS) under the grant number scrcomb.

References

1. Goonko, Y. P., Kharitonov, A. M., Mazhul, I. I., Zvegintsev, V. I., Nalivaichenko, D. G., Chirkashenko, V. F.: Investigations of a Scramjet Model at Hypersonic Velocities and High Reynolds Numbers, AIAA-paper 2002-5273, 2002.
2. Reinartz, B., Behr, M.: Computational Design Study of a 3D Hypersonic Intake for Scramjet Design Testing, High Performance Computing in Science and Engineering 2010, Springer, pp. 429-442, 2010.
3. Gerlinger, P., Stoll, P., Kindler, M., Schneider, F. and Aigner, M.: Numerical Investigation of Mixing and Combustion Enhancement in Supersonic Combustors by Strut Induced Streamwise Vorticity, Aerospace Science and Technology, **12**, pp. 159-168, 2008.
4. Kindler, M., Blacha, T., Lempke, M., Gerlinger, P. and Aigner, M.: Numerical Investigations of Model Scramjet Combustors, High Performance Computing in Science and Engineering 2008, Springer, pp. 153-166, 2008.
5. Coakley, T. J., Huang, P. G.: Turbulence Modeling for High Speed Flows, AIAA-paper 92-0436, 1992.
6. Gerlinger, P.: Numerische Verbrennungssimulation, Springer, ISBN 3-540-23337-7, Berlin-Heidelberg, 2005.
7. Gerlinger, P.: Investigations of an Assumed PDF Approach for Finite-Rate-Chemistry, Combustion Science and Technology, **175**, pp. 841-872, 2003.
8. Gerlinger, P., Möbus, H., Brüggemann, D.: An Implicit Multigrid Method for Turbulent Combustion, Journal of Computational Physics, **167**, pp. 247-276, 2001.
9. Jachimowski, C. J.: An Analytical Study of the Hydrogen-Air Reactions Mechanism With Application to Scramjet Combustion, NASA TP-2791, 1988.
10. Shuen, J. S.: Upwind Differencing and LU Factorization for Chemical Non-Equilibrium Navier-Stokes Equations, Journal of Computational Physics, **99**, pp. 233-250, 1992.
11. Jameson, A., Yoon, S.: Lower-Upper Implicit Scheme with Multiple Grids for the Euler Equations, AIAA Journal, **25**, pp. 929-937, 1987.
12. Gerlinger, P., Brüggemann, D.: An Implicit Multigrid Scheme for the Compressible Navier-Stokes Equations with Low-Reynolds-Number Turbulence Closure, Journal of Fluids Engineering, **120**, pp. 257-262, 1998.
13. Stoll, P., Gerlinger, P., Brüggemann, D.: Domain Decomposition for an Implicit LU-SGS Scheme Using Overlapping Grids, AIAA-paper 97-1896, 1997.
14. Stoll, P., Gerlinger, P., Brüggemann, D.: Implicit Preconditioning Method for Turbulent Reacting Flows, Proceedings of the 4th ECCOMAS Conference, **1**, pp. 205-212, John Wiley & Sons, 1998.
15. Kindler, M., Gerlinger, P. and Aigner, M.: Delayed Detached Eddy Simulations of Compressible Turbulent Mixing Layer and Detailed Performance Analysis of Scientific In-House Code TASC3D, High Performance Computing in Science and Engineering 2011, Springer, pp. 259-272, 2011.

Application of the Unified Turbulent Flame-Speed Closure (UTFC) Combustion Model to Numerical Computation of Turbulent Gas Flames

Feichi Zhang, Peter Habisreuther, and Henning Bockhorn

Abstract The current work presents the numerical computation of turbulent reactive flow by means of three different classes of flame: a premixed, a non-premixed and a partially premixed flame. The aim thereby is to validate the unified turbulent flame-speed closure (UTFC) combustion model developed at our institute. It is based on the presumption that the entire turbulent flame can be viewed as a collection of laminar premixed reaction zones (flamelets) with different mixing ratios. The mixing process is controlled by the mixture fraction ξ and the subsequent chemical reaction by the progress variable θ . The turbulent flame speed S_t is used to describe the flame/turbulence interaction as well as the finite rate reaction. Complex chemistry is included and the pressure dependency (elevated pressure) of the combustion process is included in the model as well. The applicability of the model is explored by means of RANS (Reynolds averaged Navier-Stokes approach) and LES (large eddy simulation) methodologies at a wide range of Damköhler number Da . The results of all simulations show reasonable good agreement with the experiments.

1 Introduction

CFD (computational fluid dynamics) simulation of combustion systems is an important and reliable tool in designing and optimizing combustion devices like combustion engines or gas turbines. For such industrial flows with high Reynolds number Re , the resolution of all turbulent scales and the computation of all reacting species concentrations is not possible due to computational costs. For this reason, modeling is needed to simplify the underlying physics, namely: turbulence

F. Zhang (✉) · P. Habisreuther · H. Bockhorn
Division of Combustion Technology, Engler-Bunte-Institute, Karlsruhe Institute of Technology,
Engler-Bunte-Ring 1, 76131 Karlsruhe, Germany
e-mail: feichi.zhang@kit.edu

and combustion. Since combustion occurs in complex 3D flows characterized by high turbulence intensities and thermal loads in industrial applications, the turbulent length scales are in the order of the typical laminar flame thickness δ_l , suggesting that the flamelet assumptions [1] are in most cases violated. Proper and comprehensive modeling therefore becomes important and is the objective of the work. There are numerous modeling concepts for turbulent combustion which usually underly different physical restrictions, for example, some are only applicable for premixed or non-premixed flames or for fast chemistry.

For turbulence modeling, the classical RANS method is fast, but provides only information about the time mean variables of the flow and lack transient characteristics. On the other hand, the LES offers the possibility to resolve the unsteady flow structures down to a cut-off level and therefore is a good compromise between computational costs and additional accuracy. The main difficulty for LES modeling of turbulent combustion is that the theoretical flame thickness of 0.1–1 mm is generally smaller than the LES mesh size. This phenomena is avoided in the current work by filtering the flame front. In doing so, the unresolved scalar transport increases the flame thickness, both through the turbulent and numerical diffusion so that the filtered or thickened flame can be resolved on the relatively coarse mesh [2]. As a drawback, the response of the flame on turbulence fluctuations becomes less sensitive.

2 Combustion Modelling

The current work proposes a unified turbulent flame-speed closure (UTFC) combustion model, which was developed by the authors, for computation of turbulent reactive flows. It is an extension of the reaction model by Schmid [3] for modeling of non-premixed and partially premixed combustion [4]. The basic idea of the concept is to assume that the mixing of fuel and oxidizer takes place before the chemical reaction, so that the entire turbulent flame can be considered as a collection of distinct reaction zones with different, but individually fixed stoichiometries. The mixing process is controlled by the mixture fraction ξ and the subsequent chemical reaction by the progress variable θ .

To take the effect of turbulence on mixing into account, two transport equations for $\tilde{\xi}$ and its variance $\tilde{\xi}''^2$ are solved in addition to the basic conservation equations and the transport equation for the reaction progress $\tilde{\theta}$ ($\tilde{\cdot}$ denotes Favre-averaged values). Pre-computed laminar flame structures are then averaged by a probability density function (PDF) $P(\xi)$ determined by the mean value of $\tilde{\xi}$ and its variance $\tilde{\xi}''^2$ with a fixed principal shape [1]. All averaged (i.e. over the PDF integrated) quantities are then gathered in a look-up table with three input parameters: the mixture fraction $\tilde{\xi}$, the variance of the mixture fraction $\tilde{\xi}''^2$ and the progress variable $\tilde{\theta}$.

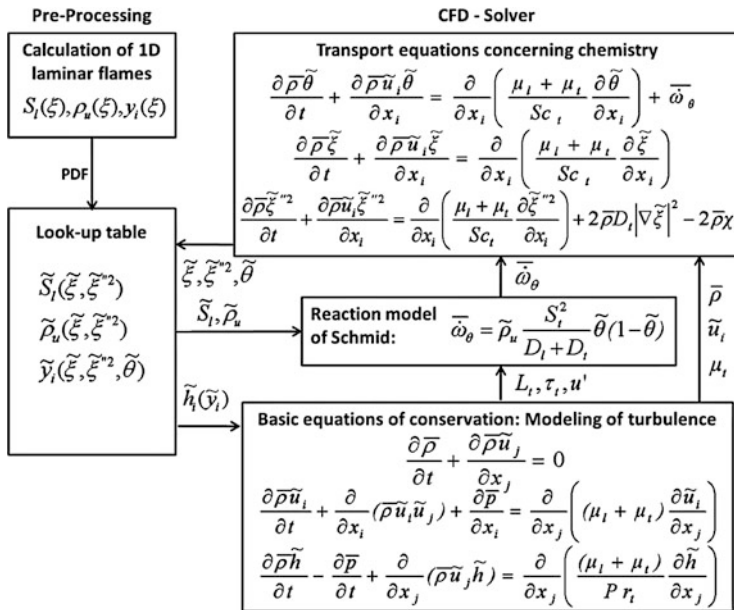


Fig. 1 Connection of the look-up chemistry table and the CFD solver in the UTFC approach


Figure 1 illustrates the interaction between the look-up chemistry table and the CFD solver. The turbulent flow field and the pre-computed flame structures are linked through exchange of $\tilde{\xi}$, $\tilde{\xi}''^2$ and $\tilde{\theta}$ (by multidimensional interpolation). The reaction rate $\bar{\omega}_\theta$ is computed by means of the turbulent flame speed S_l from the turbulence parameters (turbulence intensity u' , turbulent time scale τ_t and turbulent length scale L_t) and the thermophysical characteristics (laminar flame speed \tilde{S}_l and density of unburned mixture $\tilde{\rho}_u$) as a function of the mixing ratio. The species concentrations $\tilde{\gamma}_i$ are needed to evaluate the enthalpy and the temperature, respectively. A more detailed description of the reaction model can be found in [4].

3 Simulation Aspects

3.1 Case Description

To validate the proposed combustion model, it has been implemented into the open source CFD code OpenFOAM [5] which employs the finite volume method of a cell-centered storage arrangement. Subsequently, LES and RANS simulations of a premixed, a partially premixed, and a non-premixed flame have been carried out in sequence. All cases are operated at normal pressure $p_0 = 1$ bar.

Fig. 2 Sketch of the matrix burner [6]




- **Premixed combustion (Matrix Flame) [6]:** A natural gas flame with an air to fuel equivalence ratio of $\lambda = 1.75$ (lean premixed) and a thermal load of $P_{th} = 275.6 \text{ kW}$ with the unburned mixture preheated to $T_0 = 400^\circ \text{C}$ (see Fig. 2). The Re based on the nozzle diameter ($D = 0.15 \text{ m}$) and bulk velocity ($U_{bulk} = 20.7 \text{ m}$) is $Re \approx 48,000$.
- **Partially premixed combustion (Cabra Flame) [7]:** A lifted jet flame with a mixture of 33% natural gas and 67% air by volume (fuel-rich premixed), issued from a central nozzle ($u_{bulk} = 100 \text{ m/s}$, $D = 4.57 \text{ mm}$, $T_0 = 320 \text{ K}$, $Re = 28,000$) into a large coflow of hot combustion products from many lean premixed hydrogen/air flames ($u_{co} = 5.4 \text{ m/s}$, $D_{co} = 210 \text{ mm}$, $T_{co} = 1,350 \text{ K}$), as shown in Fig. 3.
- **Non-premixed combustion (Sandia H3 Flame) [8]:** Fuel jet of 50% H_2 and 50% N_2 in volume which was injected into an air coflow at 0.2 m/s ($T_0 = 300 \text{ K}$). The corresponding Re based on the nozzle diameter ($D = 8 \text{ mm}$) was 10,000.

3.2 Numerical Setups

A fully implicit compressible formulation was used together with the pressure-implicit split-operator (PISO) for the pressure correction. The applied discretizations

Fig. 3 Schematic of the lifted coflow burner [7]



for the momentum equations were based on a second order central difference scheme for LES and a bounded linear scheme for RANS. To show the universality of the approach, only the standard turbulence models, i.e. the Smagorinsky sub grid scale (sgs) model [9] for LES and the standard $k - \epsilon$ model [10] for RANS, were applied for the turbulence modeling. There, the turbulent Schmidt and Prandtl number were assumed as constant $Sc_t = Pr_t = 0.7$.

For all LES simulations, a turbulence inflow-generator [11] was used to provide transient correlated velocity fluctuations at the inlet boundary for each time step. Moreover, the inflow generator has been applied in combination with a velocity profile for the inlet boundary condition. The near-wall region was modelled by wall functions. A non-reflecting boundary condition (NRBC) proposed by Poinso and Lele [12] has been applied for the compressible LES/RANS to avoid nonphysical reflections of traveling pressure waves at the entrainment boundaries. x and r indicate the streamwise and radial axis of the cylindrical coordinate system.

Since the round jet configuration is representative in our cases, the computational domains are chosen to consist of a final part of the nozzle to allow development of the turbulence in the burner and the velocity changes directly at the nozzle exit, and a cylindrical domain ($L \times D$) downstream of the burner where mixing and combustion take place. Due to the simple geometries of the computational domain, in order to achieve better accuracy, the computational grids are built up in a block-structured way employing hexahedral shaped elements. Table 1 shows some statistics of the used computational grids. In each case, the LES and RANS simulations have been computed on the same grid for comparison reasons, although it is not necessary, because only the time mean flow field is solved in RANS where a much coarser grid can be used.

Table 1 Mesh parameters and simulation setups used for the different flames

Case	$L \times D$	N_{cell} (million)	Δ_{min} (mm)	N_{inlet}	Δt (μ s)	t_{total} (s)	$N_{timestep}$
Matrix Flame	2.0 m \times 1.6 m	6.3	0.5	3,000	5	0.6	120.000
Cabra Flame	1.0 m \times 0.2 m	4.3	0.2	900	0.5	0.1	200.000
H3 Flame	0.6 m \times 0.4 m	2.5	0.2	1,200	2.5	0.3	120.000

All simulations are initialized by a uniformly distributed flow. The first stage is to get a fully developed flow field without reaction. In the second step, the non-reactive flows have been ignited and develop further to a turbulent flame. Statistics of the flow parameters are collected after certain volumetric flow-through times in this stage. The time step Δt and the total running time t_{total} for the LES simulation are chosen to keep the CFL no. smaller than 0.4 and to get a well averaged flow field (see Table 1).

3.3 Computational Effort

Since the computational meshes used are fine (see Table 1) and also additional transport equations have to be solved to include chemical reactions, the in-house Linux cluster (with a Gigabit Ethernet Interconnect) at the Division of Combustion Technology is not able to speed up efficiently by increasing used number of processors. For this reason, all LES/RANS simulations have been conducted on the *HP XC4000 cluster of the Steinbuch Centre for Computing (SCC) at the KIT* [13] whose large computational resources enable such a comprehensive comparison of different models on such fine grids.

In our cases, the number of processors used are typically chosen to run a whole simulation as one job using the maximal running time of 4,320 min. For LES simulations, where at least 100.000 time steps are necessary to get a well averaged flow field, 128–196 processors depending on the mesh size have been used to achieve a performance of approximately 3 s per time step in LES. The computational effort for RANS is much lower than that for LES, even though the same mesh is used for both methods, because the LES has to be run for a long time in order to gain statistical convergence. A performance study on parallelization has been done for the finest grid with 6.3 mio. cells. By use of one single node (4 processors), the solver needs 76 s for one time step. Running the same case on 49 nodes (196 processors), a speed-up factor of 25 ($\cong 3$ s) was achieved, which, however, corresponds to an efficiency factor of only 55 %.

3.4 Chemistry Tabulation

For combustion modeling using the UTFIC concept (see Sect. 2), the CHEMKIN II program package [14] has been used for all cases to calculate the internal flame

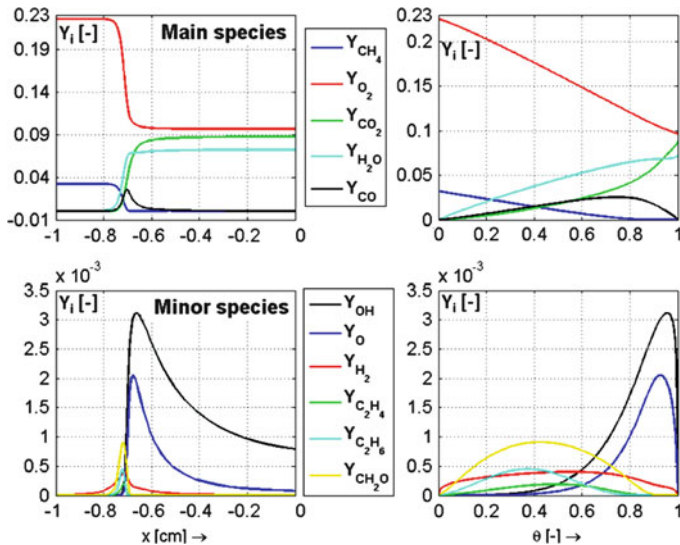


Fig. 4 Species profiles versus the flame coordinate (*left*) and the progress variable (*right*) at $\lambda = 1.75$ and $T_0 = 400^\circ\text{C}$ (Matrix Flame). The resulting laminar flame speed S_l was thereby 0.805 m/s in this case

structure applying a detailed reaction mechanism: the GRI 3.0 (Gas Research Institute) mechanism containing 325 reactions and 53 species [15] for methane/air combustion and the reaction mechanism of Maas and Pope [16] containing 9 species and 27 elementary reactions for combustion of hydrogen with air. In Fig. 4, mass fraction profiles in the flame coordinate $y_i(x)$ and in the θ -space $y_i(\theta)$ used for the **Matrix Flame** are illustrated: at the top for the main species ($O[y_i] = 10^{-2} - 10^{-1}$) and at the bottom for some intermediate species ($O[y_i] = 10^{-3} - 10^{-4}$).

Figure 5 gives an overview of the tabulated laminar flame speed $S_l(\tilde{\xi}, \tilde{\xi}''^2)$ for the **Cabra Flame**. The influence of the hot coflow on the flame speed was incorporated by construction of the look-up table, where each of the flame computations was conducted with a preheating temperature attained due to the mixing of the fuel jet with the hot coflow gas. Since S_l is strongly dependent on the temperature, it increases dramatically by the mixing process and comes up to a maximal value of 3.1 m/s close to the stoichiometric mixture fraction of $\xi_{st} = 0.177$.

Figure 6 shows the integrated table for water vapour $\tilde{y}_{H_2O}(\tilde{\xi}, \tilde{\xi}''^2, \tilde{\theta})$ used for modeling of the **Sandia H3 Flame**, the surfaces from the lowest horizontal position to the highest location indicate values of θ from 0 to 1 or from unburnt to completely burnt state with an increment of $\Delta\theta = 0.2$. Note that $\tilde{\xi}$ and $\tilde{\xi}''^2$ are allocated not equidistantly as indicated by the surface patches.

Fig. 5 Chemistry look-up table for the laminar flame speed used for the Cabra Flame





Fig. 6 Look-up table $\tilde{y}_{H_2O}(\xi, \xi^{1/2}, \theta)$ used for combustion modeling of the H3 Flame




4 Simulation Results

4.1 Premixed Combustion Modeling

In Fig. 7, slices (passing through the centre symmetry axis) of the computed time mean temperature T , streamwise velocity u , and mass fraction of methane y_{CH_4} from LES and RANS simulations are compared with the measured data. A conical shaped flame front stabilizing at the burner rim can be identified. It exhibits a length of about 0.3 m. The RANS predicts a longer flame than the LES and the experiment, which can be attributed to the $k - \epsilon$ model, where the turbulence quantities are underestimated leading to a smaller burning speed and a weakened reaction rate $\bar{\omega}_\theta$, respectively. In contrast, the length and the thickness of the flame provided by LES has a good agreement with the experiment.

At the bottom left of Fig. 8, instantaneous contours of the streamwise velocity \tilde{u} at the inlet BC provided by the inflow generator is illustrated, which forms intense turbulence upstream of the burner exit. The influence of the inflow turbulence can

Fig. 7 Slices of computed and measured time mean variables: LES/Exp./RANS



also be identified in contour-plots of \tilde{u} (top left of Fig. 8) and the vorticity $\omega = |\text{rot} \times \tilde{\mathbf{u}}|$ (bottom right of Fig. 8). The black curves in u - and T -contours indicate the flame surface defined by $\tilde{\theta} = 0.5$. The flame front is strongly corrugated and torn by the intense turbulence from the inflow. However, it is attenuated by passing through the flame surface due to the strong acceleration caused by the gas expansion and the increased fluid viscosity. The hot products mix thereafter with incoming cold air which again enhances turbulence, as indicated by the vorticity field (the white curve is the isoline by $\tilde{T} = 1,700$ K). The temperature increases rapidly in the reaction zone (indicated by the isoline of $\tilde{\theta} = 0.5$) and decreases then by mixing with the ambient air, which is shown at top right of Fig. 8 in the contour plot of instantaneous temperature.

Figure 9 shows a detailed comparison of the computed radial profiles of the time mean values with the experiment for six axial positions (from left to right: the temperature T , the streamwise velocity u , the radial velocity v as well as the Reynolds stresses $R_{xx} = \overline{u'u'}$ and $R_{yy} = \overline{v'v'}$). Note that only the resolved Reynolds stresses are used here, because the sgs part is smaller in magnitude compared to the resolved




Fig. 8 Instantaneous flow and temperature fields derived from the LES computation

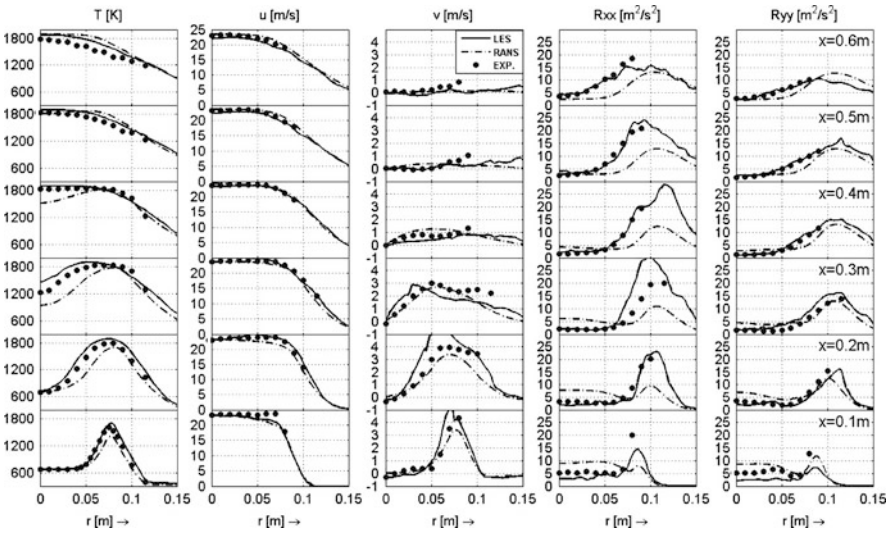


Fig. 9 Comparison of computed and measured radial profiles at different streamwise positions




Fig. 10 Effect of the inflow turbulence on the flame length

one due to usage of the fine grid (the sgs turbulent kinetic energy $k_{sgs} \propto \Delta^2$). The overall agreement for the velocities and the temperature field is very good. Even the Reynolds stresses are reproduced very well. This indicates that the turbulence generator applied for computing the velocity components at the inlet works well. However, the flame length from LES is slightly shorter than the measured one (see also Fig. 7) which may be attributed to the overpredicted R_{xx} , for example, at the position $x = 0.3$ m. To the contrary, the Reynolds stress R_{xx} in the streamwise direction from RANS is largely underestimated since the $k - \epsilon$ model assumes isotropy of turbulence whereby the normal stresses are almost equal. This leads consequently to an over-predicted flame length. Despite of these differences, the results are satisfactory as the flame is preheated and operated at very high turbulence level, so that the interaction of the flame with the turbulence is difficult to capture ($Da < 1$).

It is important to note that the precise setting of the turbulence conditions at the inlet boundary plays a decisive role for the current case. In order to demonstrate this, another simulation using the same numerical setups is applied to the matrix flame, however, applying a laminar inflow condition. As shown in Fig. 10, LES with a steady inflow condition (right) exhibits a rather uniform and laminar flame front due to the missing turbulence and predicts a significantly longer flame than LES using a turbulent inlet BC (left). In this case, the enhanced turbulence leads to a higher turbulent flame speed and a strongly increased reaction rate $\bar{\omega}_\theta$, so that the flame is shorter.

Figure 11 demonstrates the influence of the mesh sizes on the resolved flame brushes computed by LES where the black curves illustrate isolines of $\tilde{\theta} = 0.1 - 0.9$ with the increment of $\Delta\tilde{\theta} = 0.1$. It is obvious that the mesh resolution imposes a grave impact on the turbulence/flame interaction which can be identified by the corrugated and torn flame surfaces. As indicated in Fig. 11, the flame front provided by LES on the coarse mesh with 0.66 million elements is less sensitive to turbulent fluctuations than the one that is result from the finer meshes.




Fig. 11 Effect of mesh size on the resolved flame front from LES of the matrix flame

Also, LES on a finer mesh resolves a somewhat thinner and longer flame due to the attenuated turbulent and numerical diffusion ($\nu_t |D_t \propto \Delta^2$ according to the Smagorinsky model). However, it is noteworthy that the time mean properties, for example, the length and thickness of the flame, should be much more independent on this effect since these are governed by the mean flow rate (of fuel) and resolvable even using a coarse mesh.

4.2 Modeling of the Lifted Partially Premixed Flame

Figure 12 shows instantaneous contours of the temperature and the OH -mass fraction from the LES. A lifted flame with an inner and an outer reaction zone can be identified clearly, which are sustained by a premixed and a non-premixed combustion. The reaction zones due to these two types of flame cause an inner strong and an outer weak formation of OH or heating of the flow. The diluted mixture is pre-burned by a higher reaction rate from the premixed combustion, the remaining fuel mixes thereafter with oxygen from the vitiated coflow and results in a weak diffusive flame. Note that this diffusion flame is not located exactly at the stoichiometric surface because it is sustained by mixing of the flue gas from the previously occurred premixed combustion and incoming oxygen from coflow.

With respect to the underlying physics, the laminar flame speed increases by mixing with the hot coflow (see Fig. 5). Despite of that fact, the flame cannot stabilize or the reaction can not take place directly after the burner mouth due to the very small turbulent time scale or large strain generated at the shear layer, i.e. $\bar{\omega}_\theta \rightarrow 0$ by τ_t or $Da \rightarrow 0$. Further downstream, the fuel jet mixes with the coflow and becomes heated as well due to the high temperature of the coflow (1,350K), so that the flame speed increases strongly. At the same time, the flow is slowed down by the turbulent diffusion process. As a result, there is a counteraction of




Fig. 12 Instantaneous contours of temperature and OH -concentration from the LES computation: *white curve* indicates flame surface defined by $\theta = 0.5$; *black curve* denotes isoline of the stoichiometric mixture fraction $\xi = 0.177$; *horizontal lines* mark positions of the measured radial profiles at $x/d = 15/30/40/50/70$

increasing burning speed and stalled flow motion. As soon as the local flame speed and the flow velocity are in balance on a certain downstream location, the flame is stabilized there. All these phenomena of a lifted partially premixed flame are accounted for by the UTF-C combustion model, so that the stabilization mechanism can be captured properly.

In Fig. 13, the mean and root mean square (rms) values derived from the LES show quantitatively good agreement with the measurement. The small differences between the LES and the measured data may be caused by the used turbulence conditions at the inlet boundary, which was unknown from the experiment. Despite this uncertainty at the inlet BC, the diffusion flux from the RANS computation is underestimated which leads to higher ξ at $x/d = 15$ and $x/d = 30$. This may be caused by the standard $k - \epsilon$ turbulence model due to its isotropic turbulence assumption [10]. Consequently, the lift-off height is predicted too high in the RANS computation, which can also be identified by the lower temperature level at $x/d = 40$ and $x/d = 50$ (dashed line).

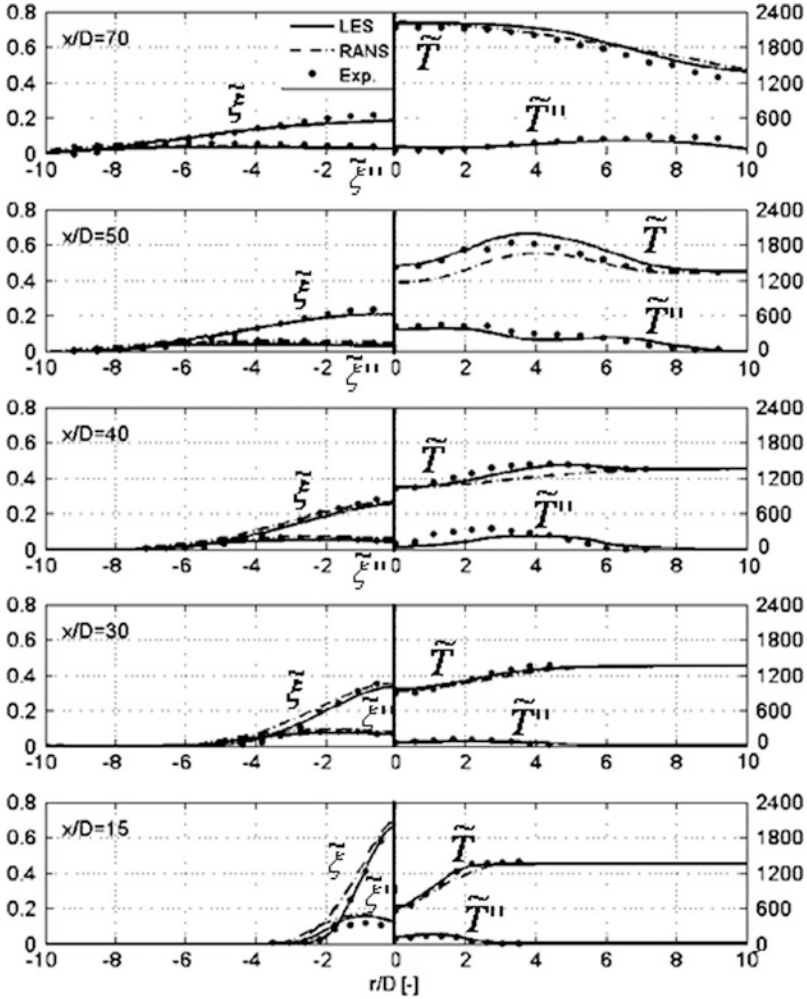


Fig. 13 Comparison of measured and computed time mean and rms profiles of mixture fraction (*left*) and temperature (*right*) at different streamwise positions as indicated in Fig. 12

Figure 14 shows a comparison of the centerline distributions of the LES/RANS simulations with the experiment. The agreement between the predicted and measured data is very good. $\tilde{\xi}$ from LES is slightly smaller than the measured data leading to a higher \tilde{T} , which may be attributed to the enhanced numerical diffusion due to the used relatively coarse grid ($\nu_t \propto \Delta^2$). The RANS simulation overpredicts the lift-off position, so that an increment of temperature occurs further downstream. Despite of these differences, the UFC is an able concept to be used for modeling partially premixed flames at extinction limit and provides even reasonable good solutions.




Fig. 14 Comparison of centerline profiles from measurement and CFD simulations




Fig. 15 Instantaneous contours of the temperature \tilde{T} , the turbulent flame speed S_t and the sub grid scale (sgs) turbulence intensity u' from the LES simulation. The flame surface is indicated by isolines of $\tilde{\xi} = \xi_{st} = 0.3$ in each subplot as well. The horizontal line marks the position of the measured mean flame length at $L_f = 34D$

4.3 Non-premixed Combustion Modeling

In Fig. 15, snapshots from the LES computation give an insight into the resolved reactive flow field. The computed flame exhibits a length of $L_{LES} = 33D$ which is equal to the predicted flame length from RANS, but slightly below the measured one. This may be attributed to the enhanced numerical and turbulent diffusion caused by the relatively coarse grid (of 2.5 mio. cells). The interaction of the flame with the turbulent flow can be identified by the corrugated and disrupted flame surface (denoted by isolines of $\tilde{\xi} = \xi_{st}$). As illustrated in Fig. 15 on the right (u' -contour), there are many vortices generated along the shear layer and the flame




Fig. 16 Comparison of the measured and computed time mean profiles of (from left to right) mixture fraction, temperature, streamwise velocity and turbulent kinetic energy

surface which cause large turbulence intensities and, thus, flame speed (see S_f -contour) in these regions too. The temperature and the flame speed are strongly correlated to the mixing field due to the pre-defined look-up tables. The maximum value of S_f is located somewhat at the fuel rich region $\xi > 0.3$, because H_2 has a much higher diffusion coefficient than air. In addition, regions of large flame speed and temperature are located far downstream due to the large strain and the high variance ξ'^2 caused at the root of the flame or the shear layer.

In Fig. 16, the calculated mean profiles show good agreement with the experimental data [8]. The turbulent kinetic energy (in the last column) is over-predicted close to the burner ($x/D = 5$) in both RANS and LES simulations which may be caused by the used inlet BC, where the turbulence properties were not exactly known from the experiment. Also, the wall thickness of the nozzle was assumed as infinitely thin which causes a higher strain between the jet and the coflow. The turbulent diffusion and dissipation in LES are somewhat overestimated due to the relatively coarse grid used, leading to profiles of ξ and u (in the first and third column) which lie somewhat below the measured data at $x/D = 20$. This explains

the slightly higher temperature (in the second column of Fig. 16) in this region. Compared to the computations presented in the early work [4] using a coarser mesh and otherwise the same conditions, the RANS simulations performed on different grids (1 vs. 2.5 million cells) exhibited results of very similar quality. In contrast, the flame length was considerably underestimated by LES on the coarse mesh in [4] due to the enhanced numerical and turbulent diffusion. This indicates that the RANS method is not as sensitive as LES to mesh resolutions and the $k - \epsilon$ model is well suited for free-shear flows without large adverse pressure gradients [10].

5 Conclusions

The scope of the work is to present numerical studies of turbulent combustion by means of three different flames: a premixed, a partially premixed and a non-premixed flame. Aim of the work is to validate the newly developed unified turbulent flame-speed closure (UTFC) combustion model. The turbulent reactive flows have been computed on relatively fine grids using compressible LES and RANS methods. In the premixed jet flame case, the LES showed very good agreement with the measurements of the mean flow variables and the Reynolds stresses. The inflow turbulence condition was found to play a crucial role for the modelled reaction rate and the resulting flame length, which is well reproduced by use of an inflow generator. The influence of mesh resolutions on the interaction between turbulence and flame front has been pointed out. The second case was represented by a lifted partially premixed flame, the stabilization mechanism of the lifted flame was captured properly by use of the UTFC model. The results from LES/RANS simulations showed a good agreement with the measured streamwise and radial profiles for the mean and rms values of $\tilde{\xi}$ and \tilde{T} . For the non-premixed H3 Flame, the predicted flame length from RANS and LES were very close to the measured one. The computed mean radial profiles are very promising in comparison to the experiment.

In fact, it is doubtful whether the differences between the simulation results and the measured data are solely attributable to the drawbacks involved in the UTFC approach. For example, accurate resolution of the mixing field plays a decisive role for computation of diffusion flames. However, this issue is mainly covered by the turbulence modeling. In addition, the turbulence parameters in S_t are derived from the turbulence model applied as well which influence the reaction rate to an extent. Also, the mesh resolution remains still a quality-determining factor for LES. Nevertheless, there are only very few alternative models until now which could be used to simulate such different flames using one single reaction model. There are models indeed, which could provide better results for some of the demonstrated flames, but these will completely fail elsewhere. However, it is expected that a refined mesh and more accurate specifications of the inlet BCs or, more importantly, use of more sophisticated turbulence or mixing model will improve the results.

In summary, it can be stated that the UTF model is aimed to keep track of the characteristic rate determining time scale of the combustion process to describe the resulting heat release. Therefore, the existence of such a time scale in most of the combustion systems makes this approach suitable for modeling of turbulent premixed, non-premixed and partially premixed combustion. Because of the simplicity, the efficiency and the robustness of the proposed model and according to the different cases already carried out in a variety of systems, it is conceivable that the model may be useful for practical applications. For this purpose, the influence of the chemical kinetics at different operational conditions, like preheating, elevated pressure, and use of different fuels, can be simply implied by modifying the burning speed S_l .

Acknowledgements The authors gratefully acknowledge the financial support by the German Research Council (DFG) through the Research Unit DFG-BO693 "Combustion Noise". Major part of the computation time was kindly provided by the Steinbuch Centre for Computing (SCC) of the Karlsruhe Institute of Technology.

References

1. N. Peters. Turbulent Combustion, Cambridge University Press, Cambridge (2000).
2. F. Zhang, P. Habisreuther, M. Hettel, and H. Bockhorn. Modelling of a Premixed Swirl-stabilized Flame Using a Turbulent Flame Speed Closure Model in LES. *Flow, Turbulence and Combustion* 82, 537–551, 2009.
3. H. Schmid, P. Habisreuther, and W. Leuckel. A model for calculating heat release in premixed turbulent flames. *Combust. and Flame* 113, 79–91, 1998.
4. F. Zhang, P. Habisreuther, M. Hettel and H. Bockhorn. Application of a Unified TFC Model to Numerical Simulation of a Turbulent Non-Premixed Flame. Proceedings of the 8th International ERCOFTAC Symposium on Engineering Turbulence Modelling and Measurements - ETMM8, vol.2, 681–686, 2010.
5. OpenCFD Ltd (2011): OpenFOAM User Guide, Version 2.0.1.
6. M. Zajadatz, M. Hettel and W. Leuckel. Burning Velocity of High-Turbulence Natural Gas Flames for Gas Turbine Application. Proceedings of the International Gas Research Conference, 8.-11.11.98, San Diego, CA, 793–803, 1998.
7. R. Cabra, J. Chen, R. Dibble, A. Karpetis, and R. Barlow. Lifted methane air jet flames in a vitiated coflow. *Combustion and Flame* 143, 491–506, 2005.
8. R. Barlow, (ed.). Proceedings of the TNF Workshops, Sandia National Laboratories, Livermore, CA, available at www.ca.sandia.gov/TNF, 1996–2004.
9. J. Smagorinsky. General circulation experiments with the primitive equations, I, The basic experiment, *Mon. Weather Rev.* 91, 99–164, 1963.
10. D. Wilcox. Turbulence Modelling for CFD, DCW Industries, California, USA (1994).
11. M. Klein, A. Sadiki and J. Janicka. A digital filter based generation of inflow data for spatially developing direct numerical or large eddy simulations. *Journal of Computational Physics* 286, 652–665, 2003.
12. T. Poinso and S. Lele. Boundary conditions for direct simulation of compressible viscous flows, *J. Comput. Phys.* 101, 104–129, 1992.
13. Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany, Rechenzentrum. HP XC4000 User Guide, 2009.

14. J. Kee, F. Rupley and J. Miller. Chemkin-II: A Fortran chemical kinetics package for the analysis of gas-phase chemical kinetics, Report No. SAND89-8009B, Sandia National Laboratories, 1989.
15. G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner, Jr., V. V. Lissianski, and Z. Qin. http://www.me.berkeley.edu/gri_mech/
16. U. Maas, and S. Pope. Simplifying chemical kinetics: intrinsic low-dimensional manifolds in composition space. *Combustion and Flame*, 88, 239–264, 1992.

Lagrangian Approach for the Prediction of Slagging and Fouling in Pulverized Coal Combustion

Olaf Lemp, Uwe Schnell, and Günter Scheffknecht

Abstract The deposition of ash particles in pulverized coal combustion provokes several problems for the operation of utility boilers. In order to avoid such problems, power plant operators have great interest in predicting the slagging and fouling tendency of the used fuel.

For this purpose, an industrially highly relevant tool for the prediction of slagging and fouling which is applicable on high performance computing platforms such as vector machines or massively parallel systems has been developed. The model has been implemented into the CFD code **AIOLOS** and couples several relevant processes that are crucial for the build-up of depositions in power plants. It accounts for the flight of the ash particles through the furnace, the corresponding interaction with the flue gas and considers several deposition mechanisms on walls and tube bundles. In case of a predicted contact between a particle and a surface, the deposition rate is calculated based on the stickiness of the particle and the surface which is correlated with the melting behaviour. The model also takes into account the change of the heat transfer resistance of the already deposited particles and consequently the influence on the flue gas temperature.

The model has exemplary been applied to a utility boiler with a thermal input of 730 MW (360 MW_{el}) in order to demonstrate the capability of this engineering tool.

1 Introduction

Most of the German power plants are fired either with imported bituminous coals or with local lignite coals. Bituminous coals are usually purchased on the world market depending on economic parameters. This can result in the combustion of a coal

O. Lemp (✉) · U. Schnell · G. Scheffknecht
Institute of Combustion and Power Plant Technology – IFK, University of Stuttgart,
Pfaffenwaldring 23, 70569 Stuttgart, Germany
e-mail: olaf.lemp@ifk.uni-stuttgart.de

which has not been considered during the design process of the boiler. However, depending on the composition of the mineral matter of the coal, the slagging and fouling behaviour in the course of the combustion process may vary significantly.

The impact of slagging and fouling on the functioning of a power plant is multiple. Due to the isolating effect of the deposition on the heating tubes the heat transfer from the flue gas to the water–steam cycle can be significantly deteriorated leading to an increase of the temperature of the flue gas at the exit of the furnace. Furthermore, the heat absorption in the convective heat exchangers may be reduced leading to a possible elevated heat transfer in bundles further down stream and consequently to higher steam temperatures that may exceed the design parameters of the pressure parts. Moreover, ash shedding can clog the ash hopper or damage the furnace, as also do the corrosive components of the ashes that can attack the tube material. In addition the constriction of the flue gas path due to depositions in the tube bundles lead to an increase of the flow resistance and to a decrease of the maximum flue gas load. These specified issues cause a significant deterioration of the power plant efficiency or even force the operators to unscheduled shutdowns.

For those reasons power plant operators have a big interest in the prediction of slagging and fouling. In the past mainly slagging indices were used to predict the deposition behaviour of coals. Most of these indices were developed based on a specific boiler geometry and coal and therefore, a transfer to different systems has to be carried out carefully. But as computational power rose in the past decade CFD comes more and more into play as a relevant engineering tool to predict depositions during the design phase and the operation of boilers, or even for the analysis of damages.

2 Description of the Modelling Framework

2.1 Physical and Chemical Phenomena Causing Deposition

Slagging defines the deposition in the furnace where radiative heat transfer is dominant and fouling takes place in the cooler convective heat transfer sections of the boiler. Both phenomena are caused by the inorganic components of the coal which are either bound to the organic fuel matrix or are trapped in it (*included* minerals), or are present as extraneous mineral particles (*excluded* minerals). The respective amount and chemical composition of them varies depending on the origin and geologic age of the coal. Two main processes are relevant for the slagging and fouling: On the one hand the condensation of gaseous inorganic components and on the other hand, the deposition of fly ash particles liberated from the coal matrix during pyrolysis and char combustion [1, 2]. The condensation of species from the gaseous phase is usually the initial step of deposit build-up, leading to an initial layer which is characterized by its poor heat conductivity and low viscosity. As a consequence, this melt layer can “catch” bigger particles. Due to space limitations,

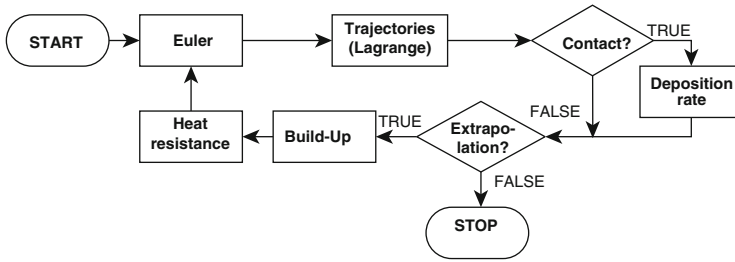


Fig. 1 Flowchart framework

only the modelling of the processes of the solid minerals will be described in the following. The computational framework of the deposition of gaseous minerals is presented in detail in [3].

2.2 Modelling Approach

The model consists of the simulation of the turbulent flow, the temperature and the concentration field of the main flue gas species with an Eulerian approach whose results are set as boundary conditions for the calculation of the particle trajectories using a Lagrangian approach. In case of a predicted collision with a wall or tube, the sticking probability is estimated and a mass deposition rate can be calculated. With a time-dependent extrapolation the thickness of the deposition can be predicted which can be used for subsequent calculations with an adjusted heat transfer resistance. Figure 1 shows a very simplified flowchart of the program with the governing parts of the modelling approach.

2.3 Simulation of Pulverized Coal Combustion

The CFD code AIOLOS, developed by IFK, has been tailored for the simulation of turbulent reacting flows which are typical in pulverized coal combustion. Special emphasis has been put on the interaction between heat transfer, chemical reactions and turbulent flow. AIOLOS solves the transport equations for mass, species, momentum and enthalpy which are obtained from the respective conservation laws. Due to their similar appearance they can be described with a general transport equation which – for the direction j – may be expressed as:

$$\underbrace{\frac{\partial(\rho\Phi)}{\partial t}}_{\text{transient}} + \underbrace{\frac{\partial(\rho u_j \Phi)}{\partial x_j}}_{\text{convective}} = \underbrace{\frac{\partial}{\partial x_j} \left(\Gamma_\Phi \frac{\partial \Phi}{\partial x_j} \right)}_{\text{diffusive}} + \underbrace{S_\Phi}_{\text{sink/source}} \quad (1)$$

with ρ , t , u , x , Γ_ϕ , and S_ϕ denoting density, time, velocity, coordinate, diffusion coefficient, and source term, respectively. The system of equations is a set of strongly coupled non-linear differential equations which demands to be solved numerically. Discretization in AIOLOS is carried out with the Finite Volume method.

To cover all the relevant phenomena of pulverized coal combustion several submodels have been developed. Turbulence phenomena are described with the k - ε model. Pressure-velocity coupling is modelled by the SIMPLE method in combination with the interpolation scheme from Date for pressure correction [4]. Radiation is described with the Discrete Ordinate method [5]. Turbulence-chemistry interaction is modelled with the Eddy Dissipation Concept (EDC) proposed by Magnussen [6], and a global reaction scheme taking into account pyrolysis, char burnout and volatile combustion is applied. The modular structure of the code allows the continuous expansion, e.g. the simulation of new developments in combustion technology (e.g. oxy-coal combustion [7]) or the coupling of the furnace simulation with a highly detailed model of the water-steam cycle [8]. In addition, the code has been optimized to achieve high performance on parallel and vector computers. Further information concerning the AIOLOS code is given elsewhere [5, 9, 10].

2.4 Modelling of the Two-Phase Flow

For the prediction of slagging and fouling, special emphasis has been put on the description of the two-phase flow. The interaction between dispersed coal particles in a continuous carrier gas flow can be described in two ways. The so-called Eulerian approach approximates the two-phase flow as quasi one-phase (continuum approximation). This is a common procedure for the simulation of pulverized coal combustion due to the low concentration of fuel particles in the gas phase. The balance is carried on in a fixed coordinate system, and the slip between gas and dispersed particle phase is neglected (see also [11]). In AIOLOS, the Eulerian approach is always used for the basic simulation, providing the velocity, temperature and main gas species concentrations fields which are the boundary conditions for the prediction of the particle trajectories.

Lagrangian Approach for the Calculation of the Particle Trajectories

The space-resolved calculation of the deposition process is modelled with a Lagrangian approach. As the particles pass through a discrete number of turbulent eddies the respective interaction of the turbulence and the particle has to be taken into account. Therefore, the calculation of the particle trajectories is based on a stochastic separated flow (SSF) model which uses a Lagrangian method of the equation of motion for each particle:

$$\frac{d(m_P \cdot u_P)}{dt} = F_D + F_g \quad (2)$$

with the particle mass m_P , the particle velocity u_P , the drag force F_D and gravitational force F_g . The drag force is defined as

$$F_D = C_D \frac{\mu}{d_P} \frac{Re_P}{2} A_P (u_G - u_P) \quad \text{with} \quad C_D = \frac{24}{Re_P} (1 + 0.15 Re_P^{0.687}) \quad (3)$$

where μ is the molecular viscosity, $u_{G,P}$ the instantaneous velocity of the gas phase or the particle, C_D the drag coefficient – considering the ratio to the Stokes drag – and the Reynolds number Re_P of the particle which is defined as

$$Re_P = \frac{|u_P - u_G| d_P \rho_G}{\mu} \quad (4)$$

Hence the particle velocity can be calculated by solving Eq. (2) analytically:

$$u_P = u_G + (u_P^0 - u_G) \exp\left(-\frac{\Delta t}{\tau_P}\right) + \tau_P g \left(1 - \frac{\rho_G}{\rho_P}\right) \left[1 - \exp\left(-\frac{\Delta t}{\tau_P}\right)\right] \quad (5)$$

with the relaxation time of the particle τ_P which is defined as:

$$\tau_P = \frac{4}{3} \frac{\rho_P d_P^2}{\mu C_D Re_P} \quad (6)$$

and the interaction time Δt which is defined as the minimum of either the transit time t_T of a particle across a turbulent eddy or the eddy–life time t_E [12] which are given as

$$\min\left(t_E = \frac{L_e}{\sqrt{2k/3}}, t_T = -\tau \ln\left(1 - \frac{L_e}{\tau|u_G - u_P|}\right)\right) \quad (7)$$

These times are calculated assuming that the characteristic size of an eddy is the dissipation length scale L_e [13]

$$L_e = C_\mu^{3/4} \frac{k^{3/2}}{\varepsilon} \quad (8)$$

Finally, the new particle position after the interaction time Δt can be calculated:

$$x = x^0 + \Delta t \frac{u_P^0 + u_P}{2} \quad (9)$$

with x^0 and u^0 denoting the position of, and the velocity at the previously considered time step, respectively.

Velocity fluctuations. In the SSF–model velocity fluctuations caused by the gas phase turbulence are assumed to be isotropic with a Gaussian probability density distribution having a standard deviation of $\sqrt{2k/3}$. The local distribution is

randomly sampled when a particle enters an eddy to obtain the instantaneous velocity. Therefore, the mean gas velocity \bar{u}_G is superimposed by a fluctuating component u'_G :

$$u_G = \bar{u}_G + u'_G \quad (10)$$

It should also be mentioned that each Lagrangian particle in the SSF-model represents a fraction of the entire fuel mass stream. Therefore, a high number of particle trajectories has to be calculated in order to achieve a statistically valid result.

Energy Balance

An important aspect for predicting the stickiness of the particle is its temperature. In this framework the temperature is calculated from an energy balance considering the convective heat transfer \dot{Q}_c , the radiative heat transfer between particle and gas \dot{Q}_r and the particle combustion heat release \dot{Q}_{com} [14]:

$$\frac{d}{dt}(m_P \cdot h_P) = \dot{Q}_c + \dot{Q}_r + \dot{Q}_{\text{com}} \quad (11)$$

2.5 Description of Tube Bundles in AIOLOS

AIOLOS treats tube bundles as porous blocks. No special delimitation of the mesh is therefore considered. The corresponding finite volume cells are treated as internal computational cells which are expanded with additional sink and source terms for enthalpy, radiation intensity and momentum. In case of a contact of the flue gas with them, heat is transferred to the tubes and the flue gas experiences a pressure drop. The amount of the source and sink term depends on the packing of the tube bundle [15].

2.6 Calculation of the Deposition Rate

In case of the collision of a particle with a surface the net deposition rate has to be calculated. For this purpose, a two-step model proposed by Pykkönen [16] is used. The model takes into account that in the Lagrangian framework, the particle cloud represents a fraction of the entire fuel stream. After each contact with a surface a fraction of this cloud can be deposited. Therefore, each Lagrangian particle is aggregated with a variable $C_{P,t}$ considering the concentration of its not-yet deposited mass. This value is set at the beginning of the simulation to 1 (=100%) and decreases accordingly to the occurrence of deposition processes:

$$C_{P,t} = C_{P,t-1} \cdot (1 - f_{\text{imp}} \cdot f_{\text{Stck}}) \quad (12)$$

f_{Imp} denotes the impact propensity and f_{Stck} the sticking propensity which are calculated each time-step when a particle has been detected in a next-to-wall cell. This is interpreted by the model as a collision between particle and surface.

Calculation of the Impact Propensity

For the calculation of the impact propensity f_{Imp} , three types of contacts between a particle and a surface are distinguished:

- Contact particle vs. furnace wall
- Contact particle vs. tube in the first row of a tube bundle
- Contact particle vs. tube in a tube bundle

In each case the driving force is different. In the first and the second case inertia is the main driving force. In the third case the main force is due to eddy impaction when a particle has not enough momentum to follow the streamlines (turbophoresis). The so-called thermophoresis which describes the impaction tendency due to a temperature gradient in the boundary layer is neglected. In the first case the impaction propensity is assumed to be equal to 1. In the second case the impaction propensity depends on the cross-sectional area of the tubes A_{Tube} , the cross-sectional of the furnace A_{Cross} and the parameter η_I which considers the impaction of particles on cross flowed cylinders [11]. The propensity is given as

$$f_{\text{Imp}} = \eta_I \left(\frac{A_{\text{Tube}}}{A_{\text{Cross}}} \right) \quad (13)$$

In the third case the packing of the tube bundles and boundary layer mechanisms are decisive. The model calculates the concentration of the corresponding resting time t_{TB} in the next-to-wall cell and calculates a so-called deposition velocity u_D caused by the turbulent impaction. The impaction propensity is calculated as follows:

$$f_{\text{Imp}} = \frac{u_D \cdot t_{TB} \cdot A_{TB}}{V_{TB}} \quad (14)$$

The impaction efficiency is also a function of the tube bundle surface A_{TB} per volume tube bank V_{TB} . A more detailed description of the models and their application can be found in [11, 17].

Calculation of the Sticking Propensity

The calculation of the stickiness of the particle and the surface is based on Walsh's proposal [18] that considers the stickiness of the particle itself and the stickiness of the surface to predict the sticking propensity:

$$f_{\text{Stck}} = \underbrace{P_P(T_P)}_{\text{sticky particles}} + \underbrace{[1 - P_P(T_P)] \cdot P_S(T_S)}_{\text{sticky surface}} \quad (15)$$

$P_{P,S}(T_{P,S})$ denotes the sticking propensity of the particle and surface as a function of the corresponding temperature. Walsh proposed to predict the sticking propensity calculating the actual viscosity of the particle and the surface and setting it into a ratio with a so-called reference viscosity which is defined as the viscosity below each particle would stick completely. The calculation of the viscosity was mostly carried out with empirical correlations based on measurements of glass-rich melts. As the measurement of the reference viscosity is quite time-consuming, and the approaches for the calculation of the viscosity have been observed by the authors as only valid for specific ash compositions [19], alternative models had to be developed. Therefore, an approach which was derived for the deposition of ashes from bio-fuels has been transferred to coal combustion. Frandsen [20] proposed to estimate the sticking propensity based on the melting behaviour of the fuel. The model assumes that if the melted fraction is

- Less than 10 % the particle or surface is non-sticky,
- Between 10 and 70 % a linear approach is assumed,
- Higher than 70 % the particle or surface is fully sticky.

The melt fraction of the coal as a function of the temperature can be calculated with the software FactSage [21]. The calculation relies on the mineral content of the fuel assuming a thermodynamical equilibrium. The basis of this software is a huge database with thermochemical data from measurements of mostly binary and ternary systems which are applied for the multi-component system that represents ash particles [22].

Calculation of the Net Mass Deposition Rate

Finally, the net mass fraction of deposition $\dot{m}_{\text{dep,net}}$ per square meter and second is calculated by

$$\dot{m}_{\text{dep,net}} = \dot{m}_{\text{Ash}} \cdot f_{\text{Stick}} \cdot f_{\text{Imp}} - k_e m_{\text{Ash}} \quad (16)$$

with \dot{m}_{Ash} denoting the arriving mass flow, the specific mass of the already deposited ash m_{Ash} and an erosion coefficient k_e which takes into account several shedding mechanisms (e.g. soot blowing). This coefficient has not yet been extensively investigated and scarce information is found in literature, but it is in the focus of the actual experimental work at IFK.

2.7 Slag Conductivity

Special emphasis has also been put on the description of the thermal conductivity of already deposited particles. Two different models based on the recommendations of Zbogor [23] have been implemented. The conductivity of sintered layers is calculated with the model of Hadley, the one of porous deposits with the model of

Table 1 Boundary conditions of the evaluation case

Inlet flows (kg/s)						
Fuel	OFA	Primary	Secondary	Sidewall	Hopper	
24.46	69	57.89	132	30	1.12	

Yagii/Kunii. Porosity is in both cases the determining resistance for the conductivity of heat between the deposition surface and the heating surface. The porosity ε is calculated in the following way:

$$\varepsilon = 1 - [(1 - \varepsilon_0) + X_{\text{Melt}}(1 - \varepsilon_0)] \quad (17)$$

where ε_0 denotes the initial porosity which depends on the characteristics of the deposit (sulfatic or silicatic) and X_{Melt} which is the calculated melt fraction [24].

3 Simulation Results

3.1 Simulation of a Utility Boiler

In this section results of the simulation of a power plant with a thermal input of 750 MW are presented. The firing system of the utility boiler consists of 12 air-staged swirl burners on three levels. The burners are arranged asymmetrically on each wall of the furnace leading to a rotational flow of the flame. Each burner level is additionally provided with two air nozzles to avoid oxygen lean areas next to the walls that can provoke corrosion of the walls. The over fire air (OFA) is installed at an elevation of 28.4 m. The overall dimensions of the boiler firing imported bituminous coal are $12 \times 12 \times 68$ m. The operating conditions are summarized in Table 1.

For the simulation with AIOLOS a three-dimensional CFD model consisting of 18 domains is used. Each burner is individually discretized, as also the four burner near regions of each wall, one main region and the convective area. The tube bundles are modelled with porous blocks. The CFD model with a total cell number of around 5.8 mio is outlined in Fig. 2 (left). Several validation simulations with AIOLOS have already been made for the considered utility boiler (see [5, 25]). Exemplary, an isosurface at a temperature of $1,400^\circ\text{C}$ is outlined in Fig. 2 (right) for displaying the typical rotational shape of the flame.

In the presented case a North American coal (Pittsburgh #8 coal) is fired. The fuel properties are summarized in Table 2. For the calculation of the melting behaviour the “FactSage” software is used as aforementioned. Figure 3 shows the respective results of the calculation. The green curve represents the amount of melted phases relative to the sum of the solid and the slag phase, the black line describes the sticking propensity. In this case, at a temperature of $1,050^\circ\text{C}$ 10% of melt fraction is reached and the particle and the surface start to become sticky. Between

Fig. 2 CFD model of the utility boiler (*left*; dimensions in m) and isosurface of temperature 1,400 °C (*right*)




Table 2 Properties of Pittsburgh #8 coal

	Proximate analysis				Ultimate analysis					
	C _{fix} (wt.-%)	Volatiles (wt.-%)	Moisture (wt.-%)	Ash (wt.-%)	C (wt.-%)	H (wt.-%)	N (wt.-%)	S (wt.-%)	O (wt.-%)	
ar ^a	53.29	31.68	9.27	12.14	67.20	4.28	1.35	0.88	4.88	
daf ^b	67.80	40.31	–	–	85.11	5.45	1.72	1.12	6.21	
Ash oxide analysis										
	SiO ₂	Al ₂ O ₃	CaO	Fe ₂ O ₃	MgO	Na ₂ O	K ₂ O	TiO ₂	SO ₃	P ₂ O ₅
	63.90	24.12	0.44	1.95	0.32	0.29	2.27	0.41	3.63	0.16

^aAs received

^bDry, ash-free basis

1,050 and 1,430 °C (70 %) the following linear correlation is assumed for the stickiness:

$$f_{\text{Stick}} = -0.0026 \cdot T_{P/S} - 2.7632 \quad (18)$$

Above this temperature full stickiness is assumed for the surface and the particle which implies that the sticking propensity equals 1. This calculation is carried out before the CFD simulation.

The presented results correspond to a theoretical study of two cases. In both cases the operating conditions are equal, only the initial condition of the deposition




Fig. 3 Melt fraction and sticking propensity of Pittsburgh #8 coal




Fig. 4 Deposition rate on rear wall, comparison of clean boiler (left) with slagged boiler (right)

thickness on the surfaces is varied. In the first case an almost clean furnace was assumed and in the second case, a typical condition at the end of the operating time was assumed. Figure 4 shows the predicted deposition rate at the rear wall of the furnace. The rates are normalized by the maximum calculated deposition rate in the second case. It is quite obvious that in the second case the deposition rate is higher than in the first case. Due to the deteriorated heat transfer the flue gas temperature rises, and the resulting higher surface temperatures lead to an increased sticking propensity. A similar result can also be observed in the tube bundles. The results for the first two bundles are shown in Fig. 5 (in both cases a cut through the centre of the bundle is shown). This calculation provides valuable information about the positions with high deposition rates which can be used for the optimization of




Fig. 5 Deposition rate in tube bundles, comparison of clean boiler (*left*) with slagged boiler (*right*)

soot blower operation. The investigated utility boiler has no remarkable slagging and fouling problems, only around the burner quarl high deposition rates are reported from the operators. This can be also observed in the presented results, especially on the second burner level.

3.2 Performance

The presented case consists of almost 6 million computational cells and a total of 18 domains. The calculations were performed at the platforms provided by the High Performance Computing Center Stuttgart (HLRS). The calculations were executed on two vector platforms. On the NEC-SX8 8 CPUs, and on the NEC-SX9 16 CPUs were utilized. An average vector length of 222 (SX9: 196) has been achieved. The computational performance is about 1.6 GFLOPS (SX9: 0.9 GFLOPS) and the memory requirement is 11 GB. Each simulation requires about 120,000 iterations until convergence which results in a total elapsed time for a full simulation of about 120 h (SX9: 64 h). The lower performance compared to previous simulations of pulverized coal combustion with AIOLOS [7, 8] is attributed to the multi-domain grid and the high number of overlapping cells in this specific case. Nevertheless, investigations to increase the performance of this particular case have been initiated showing promising results. The calculation of the particle trajectories is carried out on a massively parallel computer. 500,000 particle trajectories and the consideration of 3,000 time steps take a computing time of about 2 h.

The actual development work focuses specially in the optimization of the code on the new Cray XE6 “Hermit” platform at HLRS.

4 Conclusions and Outlook

The modelling framework of a highly useful industrial application for the prediction of slagging and fouling in coal-fired utility boilers has been presented. The model couples a huge number of physical and chemical processes that determine the deposition of gaseous and solid components of the fuel on heat transfer walls and tube bundles. Due to the lack of experimental data a full validation has not been carried out yet. However, the first simulation results provided plausible and promising results. The software tool can be applied during the design process, for the retrospective investigation of damage events as a consequence of slagging and fouling phenomena, and for the optimization of the boiler performance (e.g. soot blower operation). It is applicable for a wide band of coals and can consider the influence of different operating conditions on the deposition tendency. Due to the multitude of processes the use of high-performance computing (HPC) is a must. Also the influence of the flue gas on the water-steam cycle is currently in the focus of the development. Therefore, the coupling of the tool to IFK's water-steam code [8] will be carried on.

Acknowledgements This work and the specific COORETEC project is funded and supported by Bundesministerium für Wirtschaft und Technologie (BMWi) under grant number 0327744A, EnBW Kraftwerke AG, E.ON Energie AG, RWE Power AG, Vattenfall Europe Generation AG, Steag GmbH, ALSTOM Boiler Deutschland GmbH, Hitachi Power Europe GmbH, AZT Risk & Technology GmbH, Salzgitter Mannesmann Stainless Tubes GmbH, Sandvik Materials Technology Deutschland GmbH, ThyssenKrupp VDM GmbH, V&M Deutschland GmbH. The other project partners: MPA Universität Stuttgart, IEF2 Forschungszentrum Jülich, IWBT TU-Braunschweig and EST TU-Darmstadt are gratefully acknowledged for helpful discussions and suggestions. Computational resources have been provided by the High Performance Computing Center Stuttgart (HLRS).

References

1. G. Couch, Understanding slagging and fouling in pf combustion, IEA Coal Research, London, 1994.
2. E. Raask, Mineral Impurities in Coal Combustion, Hemisphere Publishing Corporation, Washington, 1995.
3. S. Akbar, Numerical Simulation of Deposit Formation in Coal-Fired Utility Boilers with Biomass Co-Combustion, Ph.D. thesis, Universität Stuttgart (2011).
4. A. W. Date, Complete pressure correction algorithm for the solution of compressible Navier-Stokes equations on a nonstaggered grid, Numerical Heat Transfer, Part B: Fundamentals 29 (4) (1996) 441–458.
5. J. Ströhle, Spectral Modelling of Radiative Heat Transfer in Industrial Furnaces, Shaker Verlag, 2003.
6. B. Magnussen, The eddy dissipation concept, in: 11th IEA Task Leaders Meeting "Energy Conservation in Combustion", 1989.
7. M. Müller, U. Schnell, G. Scheffknecht, Oxy-coal Combustion Modeling at Semi-industrial Scale, in: W. E. Nagel, D. B. Kröner, M. M. Resch (Eds.), High Performance Computing in Science and Engineering '11, Springer, 2012, pp. 245–258.

8. A. Matschke, U. Schnell, G. Scheffknecht, Simulation of Triflux Heat Exchangers in Utility Boilers, in: W. E. Nagel, D. B. Kröner, M. M. Resch (Eds.), *High Performance Computing in Science and Engineering '10*, Springer, 2011, pp. 217–228.
9. U. Schnell, Numerical Modelling of Solid Fuel Combustion Processes using advanced CFD-based Simulation Tools, *Progress in Computational Fluid Dynamics* 1 (4) (2001) 208–218.
10. D. Förtsch, A Kinetic Model of Pulverised Coal Combustion for Computational Fluid Dynamics, Ph.D. thesis, Universität Stuttgart, Stuttgart (2003).
11. S. Richter, Numerische Simulation der Flugaschedeosition in kohlestaubgefeuerten Dampferzeugern, Ph.D. thesis, Universität Stuttgart, Stuttgart (2002).
12. C.T.Crowe, T. Troutt, J.N.Chung, Numerical Models for Two-Phase Turbulent Flows, *Annu. Rev. Fluid. Mech* 28 (1996) 11–43.
13. J. Shuen, L. Chen, G. Faeth, Evaluation of a Stochastic Model of Particle Dispersion in a Turbulent Round Jet, *AIChE Journal* 29 (1983) 167–170.
14. H. Knaus, Simulation turbulenter reagierender Zweiphasenströmungen in industriellen Feuerungen mit komplexen Geometrien, Ph.D. thesis, Universität Stuttgart, Stuttgart (2001).
15. C. Sauer, Detaillierte gekoppelte Simulation von Kraftwerksfeuerung und -dampferzeuger, Ph.D. thesis, Universität Stuttgart, Stuttgart (2007).
16. J. Pyykönen, R. Paavilainen, J. Jokiniemi, Modelling the impact of the Burner Arrangement on Pulverised Coal-Fired Boiler Slagging, in: *The Engineering Foundation Conference on 'Impact of Mineral Impurities in Solid Fuel Combustion'*, Kona, Hawaii, 1997.
17. S. K. Kaer, Numerical investigation of ash deposition in straw-fired boilers, Ph.D. thesis, Aalborg University, Aalborg, Denmark (2001).
18. P. Walsh, A. Sarofim, J.M.Beér, Fouling of Convection Heat Exchangers by Lignitic Coal Ash, *Energy and Fuels* 6 (1993) 709–715.
19. O. Lemp, G. Scheffknecht, C. Wieland, H. Spliethoff, Evaluation of deposition models for CFD simulations of coal-fired utility boilers, in: *1st ERCOFTAC Conference on Simulation of Multiphase Flows in Gasification and Combustion*, Dresden, Germany, September 18th-20th 2011.
20. F. Frandsen, L. Hansen, H. Sorensen, K. Hjuler, Characterisation of ashes from biofuels, *Tech. rep.* (1998).
21. GTT-Technologies, *Factsage 6.2* (2012).
22. E. Yazhenskikh, K. Hack, M. Müller, Critical thermodynamic evaluation of oxide systems relevant to fuel ashes and slags, part 5: Potassium oxide-alumina-silica, *Calphad* 35 (2011) 6–19.
23. A. Zbogar, Heat transfer in ash deposits: A modelling tool-box, *Progress in Energy and Combustion Science* 31 (2005) 371–421.
24. G. Richards, J. Harb, P. Slater, Simulation of ash deposit growth in a pulverized coal-fired pilot scale reactor, *Energy and Fuels* 7 (1993) 774–781.
25. T. Sabel, Betriebsoptimierung von Kohlekraftwerksfeuerungen durch Experiment und Simulation, Ph.D. thesis, Universität Stuttgart (2004).

Part IV

Computational Fluid Dynamics

Prof. Dr.-Ing. Siegfried Wagner

The following paragraph represents the selection of papers submitted to HLRS for the Review Workshop 2012 that revealed a very high scientific standard and demonstrated at the same time the unalterable usage of high performance computers (HPCs) for the solution of the problem.

It is very important that users together with members of the HLRS further develop highly sophisticated numerical methods and advanced algorithms. Remarkable progress in this respect is demonstrated by the present contributions over the last year. One part of the users ran their jobs on the NEC SX-9 and on the NEC Nehalem Cluster whereas others already switched to the new platform of HLRS, i.e. the Cray XE6 “Hermit”. Thus, the number of cores that were employed by several users at the “Hermit” has increased remarkably compared to the situation on the NEC SX-9. However, the maximum number of cores used is far below the 113,664 cores that are offered by the “Hermit”. Besides changing from the vector computer NEC SX-9 to the massively parallel platform Cray XE6, the users moved more and more to highly sophisticated methods, e.g. Discontinuous Galerkin (DG) method. The DG method is especially suited to massively parallel platforms and offers higher order methods with an acceptable programming effort.

As an example, Christoph Altmann et al. applied a DG method and an unstructured grid for their numerical simulations using up to 4,096 cores. They plan to extend their code to be able to run reliable simulations on $O(10,000)$ and even on $O(100,000)$ processors. The extension of the codes to be able to use a higher number of cores than so far is and should be a central concern of all users of massively parallel high performance computers.

Rebecca Busch et al. used also a DG method up to fourth order. In addition, they performed the evolution of the surface and volume integrals in small groups which they called patches. The data in the patches are stored in a memory efficient way,

i.e. the data could very quickly be transferred from the memory to the CPU. This procedure enabled them a reduction of computation time of up to 84.4 %.

Groskopf and Kloker studied the effects of an oblique roughness on the hypersonic boundary layer by Direct Numerical Simulation (DNS) using up to tenth order explicit finite differences (EFDs). They ported their NS3D code from the NEC SX-9 to the Cray XE6 “Hermit” and implemented explicit finite differences to speed up the computation of derivatives for a large number of domains. That way they could significantly reduce the amount of necessary MPI communications and increase the performance compared to the more accurate and more expensive compact finite differences they have formerly used. They demonstrated a remarkable speed-up from 14 to 4,096 MPI processes and gained even a superlinear speed-up due to cache effects when using EFDs.

Breuer and Alletto investigated the effect of wall roughness seen by particles in turbulent channel and pipe flows. They could show that LES with a combination of a Lagrangian treatment of the disperse phase could considerably improve the particle statistics in turbulent channel and pipe flow by incorporating a recently published wall roughness model for the solid phase.

Martin Konopka et al. used LES to study supersonic film cooling at incident shock-wave interaction. They reached a performance of 285 GFlop/s on 16 CPUs of the vector computer NEC SX-9. Their code is completely vectorized with a vector operation ratio higher than 99 %.

Klören and Laurien investigated stratified and non-stratified mixing flows and also used Large Eddy Simulation (LES) for their studies in order to reduce the computational effort compared to DNS but to increase the accuracy compared to RANS. An interesting result has been that the speed-up with four cores per node is remarkably better than with eight cores per node.

Jastrow and Magagnato simulated compressible viscous flow and solved the compressible Navier-Stokes equations in the outer flow field using an approximate Riemann solver while using the simplified boundary-layer equations near the wall. The specialty of the code of Jastrow and Magagnato is the Cartesian-grid immersed boundary method (IBM) that offers an interesting approach to realize automatic Cartesian mesh generation. They used 1,344 Opteron cores of the Cray XE6 and consumed about 24 h computational time for one unsteady calculation in three dimensions using about ten million points.

K. Nübler et al. examine the shock control bump flow physics by numerical simulation on the NEC Nehalem Cluster using the well-known RANS code FLOWer and by wind tunnel measurements. In order to identify a good bump design they performed automated shape optimizations and found out that around 80 consecutive computations are necessary to define the optimum shape.

Starzmann et al. solved also the Reynolds-averaged Navier-Stokes (RANS) equations for their numerical simulations and simulated two-phase flows by the ANSYS Code CFX on the NEC Nehalem Cluster. They needed a simulation time of only 2 weeks on the high performance computing cluster for the treatment of their project whereas on a simple PC cluster the pure numerical calculations would have run between 2 and 3 months.

The project of Markus Wittmann et al. aimed at improving the scalable domain decomposition and partitioning scheme for the Lattice Boltzmann flow solver ILBDC. The authors did not use any computers of the HLRS but of other national and international platforms. An important aspect of their contribution is the development and evaluation of GPGPU-enabled codes. Larger scaling and benchmarking tests were carried out on NERSC's Dirac system and on the Japanese Tsubame 2 platform. In addition, some specific results deal with potential alternatives to MPI, especially MPC and Co-Array FORTRAN. Although they did not observe substantial improvement compared to plain MPI usage, the result, although it is negative, is important information for HPC users in general.

Discontinuous Galerkin for High Performance Computational Fluid Dynamics

Christoph Altmann, Andrea Beck, Andreas Birkefeld, Gregor Gassner, Florian Hindenlang, Claus-Dieter Munz, and Marc Staudenmaier

Abstract In this report we present selected simulations performed on the HLRS clusters. Our simulation framework is based on the discontinuous Galerkin method and consists of four different codes, each of which is developed with a distinct focus. All of those codes are written with a special emphasis on (MPI) based high performance computing. Results of compressible flow simulations such as flow past a sphere, compressible jet flow and isotropic homogeneous turbulence as well as an application of our aeroacoustic framework are reported. All simulations are typically performed on hundreds and thousands of CPU cores.

1 Introduction

The central goal of our research is the development of high order discretization schemes for a wide range of continuum mechanic problems with a special emphasis on fluid dynamics. Therein, the main research focus lies on the class of Discontinuous Galerkin (DG) schemes. The inhouse simulation framework consists of four different discontinuous Galerkin based codes with different features, such as structured/unstructured grids, non-conforming grids (h -adaptation), non-conforming approximations spaces (p -adaptation), high order grids (curved) for approximation of complex geometries, modal and nodal hybrid finite elements and spectral elements with either Legendre-Gauss or Legendre-Gauss-Lobatto nodes. The time discretization is an important aspect in our research and plays a major

C. Altmann (✉) · A. Beck · A. Birkefeld · G. Gassner · F. Hindenlang · C.-D. Munz · M. Staudenmaier
Institute of Aerodynamics and Gasdynamics, Universität Stuttgart, Pfaffenwaldring 21, 70569, Stuttgart, Germany
e-mail: altmann@iag.uni-stuttgart.de; beck@iag.uni-stuttgart.de; birkefeld@iag.uni-stuttgart.de; gassner@iag.uni-stuttgart.de; hindenlang@iag.uni-stuttgart.de; munz@iag.uni-stuttgart.de; staudenmaier@iag.uni-stuttgart.de

role in the computing performance of the resulting method. The simulation framework includes standard explicit integrators such as Runge-Kutta, a time accurate local time stepping scheme developed inhouse and an implicit time discretization based on implicit Runge-Kutta methods. The general layout of the framework outsources all aspects of a specific physical problem to be solved (e.g. fluid dynamics) in an encapsulated module separated from the main code by clearly defined interfaces. Thus by exchanging this physical problem definition module, the framework is able to solve various partial differential equations such as the compressible Navier-Stokes equations (fluid dynamics), linearized Euler equations (aeroacoustics), Maxwells equations (electrodynamics) and Magnetohydrodynamik equations (Plasma simulation). One of the major foci in the group is the simulation of unsteady compressible turbulence in the context of Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS). Due to the occurrence of multiple spatial and temporal scales in such problems and the resulting high demand in resolution for both, space and time, a high performance computing framework is mandatory.

2 Description of Methods and Algorithms

Discontinuous Galerkin (DG) schemes may be considered a combination of finite volume (FV) and finite element (FE) schemes. While the approximate solution is a continuous polynomial in every grid cell, discontinuities at the grid cell interfaces are allowed which enables the resolution of strong gradients. The jumps on the cell interfaces are resolved by Riemann solver techniques, already well-known from the finite volume community. Due to their interior grid cell resolution with high order polynomials, the DG schemes can use coarser grids. The main advantage of DG schemes compared to other high order schemes (Finite Differences, Reconstructed FV) is that the high order accuracy is preserved even on distorted and irregular grids.

2.1 High Order Discontinuous Galerkin Solver HALO

A nodal discontinuous Galerkin scheme on a modal basis is implemented in the code HALO (Highly Adaptive Local Operator). The code runs on unstructured meshes composed of hexahedra, prisms, pyramids and tetrahedra. To maintain the high order accuracy at curved wall boundaries, a high order representation of the element boundaries is required. Several techniques for the construction of curved element boundaries are used, see [2,9]. The code is designed for the computation of unsteady flow problems and fully parallelized with MPI [2]. The scheme is explicit and therefore each grid cell only needs direct neighbor information. This property allows a very efficient parallelization. The computation domain is decomposed

by either ParMetis or recently also by the use of space filling curves. A major disadvantage of an explicit DG scheme can be the global time step restriction for guaranteed temporal stability. This restriction depends on the grid cell size, on the degree of the polynomial approximation and on wave speeds for advection terms and on diffusion coefficients for diffusion terms. In HALO, this drawback is overcome by a special time discretization, the so called time-consistent local time stepping [3, 4]. The stability criterion is applied only locally to each grid cell, thus each cell runs with its optimal time step. Hence, the computational effort is concentrated on the grid cells with small time steps. On meshes with strongly varying grid cells as well as flow velocities, the number of operations is greatly reduced compared to an explicit global time stepping approach.

2.2 High Order DGSEM Solver STRUKTI

A very efficient variant of a discontinuous Galerkin formulation is the discontinuous Galerkin spectral element method (DGSEM). This special variation of the DG-method is based on a nodal tensor-product basis with collocated integration and interpolation points on hexahedral elements, allowing for very efficient dimension-by-dimension element-wise operations.

An easy-to-use structured code (STRUKTI) was set up to test the performance of this method, especially for large scale calculations.

2.3 High Order DGSEM Solver FLEXI

To enable the efficient simulation of complex geometries, a second DGSEM based solver was developed. Sharing the same numerical discretization as STRUKTI, FLEXI is tailored to handle unstructured and even non-conforming hexahedra meshes. A common base tool for grid pre-processing shared by the hybrid unstructured solver HALO and by FLEXI was developed. This program allows us to process grid files from different commercial grid generators and translate them into readable HALO/FLEXI meshes. Furthermore, a module for curved grid generation and for non-conforming grid connection is included in this tool. As FLEXI shares the same efficient discretization as STRUKTI, the performance of both codes is comparable as in a high order method, the effort of managing the grid is negligible. The difference lies in the parallelization of both codes. FLEXI uses domain partition based on space filling curves, whereas STRUKTI is optimized for structured meshes. Benchmarking of STRUKTI and FLEXI and improvements of FLEXI's parallelization is an ongoing important task in the group.

2.4 High Order Discontinuous Galerkin Acoustic Solver *NoisSol*

For the simulation of flow induced acoustic phenomena in complex domains a high order discontinuous Galerkin based solver is very well suited. It combines the use of unstructured grids, a low sensitivity to grid quality and low dispersion and dissipation errors.

NoisSol is a solver for the linearized acoustic equations (Linearized Euler Equations and Acoustic Perturbation Equations [17]). It applies a discontinuous Galerkin scheme on triangular or tetrahedral grid cells. The time discretization employs either an ADER scheme (Arbitrary High Order Scheme using Derivatives) or a Taylor-DG scheme [6]. These schemes offer an arbitrary high order of convergence in space and time. NoisSol includes curved elements and allows MPI based parallel computations to reduce the overall wall-clock computation time. A further feature is the coupling to the finite-difference solver PIANO [18] for a domain decomposition between near field and far field.

3 Implicit Large Eddy Simulation of the Taylor Green Vortex

The Taylor-Green Vortex is one of the classical canonical cases for the numerical investigations into turbulence dynamics. Its simple initialization and boundary conditions yet complex non-linear behavior make it an ideal candidate for the evaluation of numerical schemes and the design and development of Large Eddy Simulation (LES) subgrid scale models (SGS). An open question regarding these coarse structure simulations with model closures for the unresolved quantities is whether high order schemes retain their superior accuracy per degree of freedom (which is well established for well-resolved problems) in underresolved simulations. In addition to these issues regarding numerical accuracy, the contest between low and high order schemes in terms of scalability and computational efficiency is still ongoing.

To find answers to some aspects of these problems we have conducted a series of numerical studies of the Taylor-Green Vortex flow at Reynolds number $Re = 1,600$ on both the HLRS Nehalem and Hermit clusters with our DGSEM solver “Strukti”. We used the Nehalem system to conduct fully resolved benchmark simulations of the flow (Direct Numerical Simulations) with at most 134 million spatial degrees of freedom. These simulations served as a reference solution for further studies.

To compare high and low order schemes in a LES-type setting, we now selected a fixed resolution of 256 degrees of freedom per spatial direction, which is theoretically sufficient for a good resolution of the large and medium scales, but not the dissipation-dominated small scales. Our code framework allows an arbitrary choice of the number of elements (E) as well as degree of polynomial approximation (N) within each cell, so 256 degrees of freedom can be achieved by a number of different combinations of E and N . Table 1 lists selected computations with their choices for E and N as well as number of cores and total walltimes.

Table 1 Selected Taylor-Green vortex ($Re = 1,600$) computations

No. of elements E	N	DOF per dir	Stabilization	No. of cores	CPU hours
128	1	256	–	128 (Cray XE6)	3,500
64	3	256	–	256 (Cray XE6)	3,840
32	7	256	–	256 (Cray XE6)	3,900
25	9	250	–	125 (Cray XE6)	3,960
21	11	252	–	343 (Cray XE6)	7,000
16	15	256	Weak	256 (Cray XE6)	9,700
64	7	512	–	512 (NEC Nehalem)	80,400
24	15	384	–	512 (NEC Nehalem)	29,600




Fig. 1 Kinetic energy dissipation rate and zoom in on maximum region

From Table 1, it might seem that low order approximations (low N) clearly outperform their high order counterparts in terms of computational speed. However, especially for simulations of turbulent flows, the quality of the approximation is crucial. Figure 1 compares the results for the computations listed in Table 1 for the dissipation rate, which is the essential key quantity in the Taylor-Green vortex simulation. As obvious from these plots, the high order formulations clearly distance lower order ones in terms of accuracy, in particular in the fully turbulent – i.e. numerically most demanding – flow regime. Thus, although low N computations are at a first glance faster for the same nominal resolution, their deficit in accuracy makes the use of high order schemes very attractive.

This claim is strongly supported by Fig. 2, which compares the computing times for a $Re = 800$ case of a second and 16th order scheme with the exact (DNS) solution. For the same number of degrees of freedom (64^3), the low order scheme is significantly faster than the high order one, but while the high order results is very close to the reference solution, the $N = 1$ result is essentially too far off to capture the nature of the solution. Increasing the resolution of the low order formulation shows a convergence towards the DNS results, but cannot compete




Fig. 2 Comparison of computing time for $\vartheta(2)$ and $\vartheta(16)$ for the Taylor-Green Vortex ($Re = 800$)




Fig. 3 Visualization of vortex detection criterion $\lambda_2 = -1.5$ for $N = 1$, $N = 3$, and $N = 15$ case (left to right, 256^3 DOF in each case)

with the efficiency of the high order formulation. All computations were done on the Nehalem cluster with 64 nodes.

A visual impression of our findings is presented in Fig. 3, where we compare the vortical structure of the solution for increasing order (and fixed total number of degrees of freedom): The significant improvement in solution quality and capturing of the important physical structures for higher order shown in Fig. 1 is immediately obvious. The full investigation with all results can be found in [27].

4 Implicit Large Eddy Simulation of a Compressible Roundjet

We consider the compressible turbulent flow of a roundjet. The considered subsonic test case is proposed by Bogey et al. [24] with Mach number $Ma = 0.9$ and Reynoldsnumber $Re = 65,000$. The setup of the example is plotted in Fig. 4. The




Fig. 4 Setup of the roundjet example. A hexahedra based mesh is used with domain $(30 \times 20 \times 20) r_0$ and analytical initial condition




Fig. 5 Instantaneous vorticity distribution in $z = 0$ plane

mesh consists of 69,445 hexahedra grid cells for the domain $(30 \times 20 \times 20) r_0$, where r_0 is the radius of the roundjet at the inflow.

For the computation with FLEXI, polynomials with degree 5 are chosen, resulting in about 1.5×10^7 degrees of freedom per conservative variable. The initialization of the flow field is analytic with a 5% random disturbance, complemented by a corresponding random forcing at the inflow during the simulation to drive the transition to turbulence. As the spatial resolution is not sufficient to resolve all scales, instabilities can occur due to aliasing of the high order methodology. Similar to the work of Bogey et al., stabilization via filtering is used, where the highest polynomial mode is filtered by a factor of about 95%, resulting in an implicit Large Eddy Simulation (iLES) type approach for this computation.

Figure 5 shows the instantaneous distribution of the z -vorticity component and Fig. 6 compares the characteristic mean centerline velocity with the iLES results of Bogey et al.

The history of the simulation is listed in Table 2 and the overall computational effort for simulating $t = 842$ s sums up to 13,346 core-h. Looking at the core-h

Fig. 6 Comparison of mean centerline velocity with the iLES results of Bogey et al. [24]




Table 2 History of the iLES jet simulation (FLEXI $N = 5, 15 \times 10^6$ DOF) on the HLRS CRAY XE6 system

#cores	DOF per core	Start – end sim.time (s)	Wallclock time (h)	Core-h	Core-h per (s) sim.time
128	118,068	0–62	8.00	1,024	16.36
256	59,034	62–191	7.94	2,033	15.76
512	29,517	191–322	3.98	2,038	15.56
736	20,534	322–842	11.21	8,251	15.87

required for 1 s simulation time, the strong scaling from 128 to 736 cores is nearly perfect. For the simulation on 128 cores, the high load per core seems to slightly slow down the computation, due to caching effects.

5 Direct Numerical Simulation of a Weak Turbulent Flow Past a Sphere

We consider a weakly turbulent flow past a sphere at Mach number $Ma_\infty = 0.3$ and a Reynolds number with respect to the sphere diameter D of $Re_D = 1,000$. This computation was performed with the unstructured DGSEM code FLEXI. The domain extends $25D$ downstream and $4.5D$ upstream and circumferentially. The unstructured mesh shown in Fig. 7 consists of only 21,128 hexahedra, where hexahedra lying on the sphere surface are curved. The computation was performed with polynomial degree 4, yielding 2.64 million degrees of freedom per conservative




Fig. 7 Unstructured mesh of the sphere, 3D views, front view (*left*), slice and back view




Fig. 8 Isosurfaces of $\lambda_2 = -0.001$ of the sphere flow at $Re = 1,000$

variable. The computation was done on the CRAY-XE6 cluster on 4,096 cores, the computational effort for the characteristic time unit $T^* = D/u_\infty$ was 100 core-h and the simulation was run 300 time units, thus resulting in 30,000 CPU-h or about 7 h wallclock time. The mean timestep for this computation was $\Delta t = 1.53 \cdot 10^{-4} T^*$, resulting in a total of about 2×10^6 time steps for this computation.

The mean drag is $C_d = 0.48$ and the Strouhal number $St = 0.32$ compares well to values reported by Tomboulides and Orzag [23] and the references therein. At this Reynolds number, small scales appear in the wake of the sphere due to the Kelvin-Helmholtz-like instabilities of the shear layer. A visualization of the λ_2 criterion in Fig. 8 shows that the behavior of the flow is well captured. As shown in Fig. 9, a single layer of curved hexahedra is sufficient to resolve the boundary layer.

The history of the simulation is listed in Table 3 and the overall computational effort for simulating $t = 286$ s sums up to $\approx 25,000$ core-h. Looking at the core-h needed for 1 s simulation time, the strong scaling from 352 to 2,048 cores is at 77 % and for 4,096 cores is at 61 %. The simulation with 4,096 cores runs with only 5 elements per core.




Fig. 9 Velocity magnitude contours (levels $[0, 1]u_\infty$) in the x-y plane

Table 3 History of the sphere simulation (FLEXI $N = 4$, 2.64×10^6 DOF) on the HLRS CRAY XE6 system

#cores	DOF per core	Start – end		Core-h	Core-h per (s)
		sim.time (s)	Wallclock time (h)		sim.time
352	7,503	62–64	0.32	113	56.46
448	5,895	64–72	0.97	435	54.34
1,024	2,579	72–81	0.50	515	57.17
2,048	1,290	81–108	0.97	1,988	73.62
4,096	645	110–286	4.00	16,395	92.94

6 Slat Noise

One topic that is of great interest for computational aeroacoustic applications in aerospace sciences is the noise generation of an airfoil in high-lift configuration, that is, with deployed slat and flap. This application is also a demanding test case for acoustic simulation programs, since it combines a very inhomogeneous flow, a complex geometry and many different noise generation mechanisms.

In this application a three part airfoil is examined, which was described by Lockard and Choudhari in 2009 [26]. The calculation presented here is based on a RANS computation for an unswept wing with an angle of attack of 4° , a Mach number of 0.17 and a Reynolds number of $1.7e6$. Based on this flow field, sound sources in 2D were calculated by Roland Ewert of the IAS at the German Aerospace Center (DLR) applying their Fast Random Particle Mesh (FRPM) method [19]. The source calculation was limited to a rectangular region around the slat trailing edge (see Fig. 10). The mean flow values have also been taken from the RANS calculation. The acoustic simulations was performed with the Acoustic Perturbation Equations (APE), type 4 [17]. The space and time order of the scheme were set to 4, the time step was $3.45e-5$. A circular domain around the origin with a radius of




Fig. 10 Setup NASA 30P30N (arrow points to center of microphone circle)




Fig. 11 NASA 30P30N, pressure field, $t = 2.0$

3.5 was used. It consisted of 74,700 triangular elements with 747,000 degrees of freedom.

Figure 11 shows the pressure field after 2 time units. It includes all the expected phenomena. For a quantitative evaluation of the results frequency spectra were calculated. The spectrum for a microphone point located at a circle with $r = 1.5$ at $\varphi = -70^\circ$ is shown in Fig. 12. It proves a good qualitative agreement between the frequency spectra of the presented calculation and the reference solution by Dierke et al. [25].

The computation has been performed on the HLRS Nehalem cluster on 320 cores. For the simulation of 9 time units 5.75 h wallclock time or 1,840 core-h were necessary, which results in $27.3e-5$ core-h per degree of freedom and simulated time unit.




Fig. 12 NASA 30P30N, sound pressure level, NoisSol and Dierke [25]

7 Summary and Outlook

In this project, we have successfully applied our high performance discontinuous Galerkin based framework to several test problems to enhance the physical and numerical modeling capabilities. Furthermore, the high performance computing capabilities have been investigated and compared for the different available codes. Compressible benchmark flows such as the turbulent flow past a sphere as well as large eddy simulations of isotropic homogeneous turbulence and a jet flow have been computed and evaluated with typical runs on $>1,000$ processors. Our typical reliable ‘production’ runs use $O(1,000)$ processors. In the future, we plan to extend and improve our framework to support reliable simulation runs on $O(10,000)$ and even on $O(100,000)$ processors to fully unleash the available (and projected) processing power of the HLRS CRAY-XE6 cluster.

Acknowledgements The research presented in this paper was supported in parts by Deutsche Forschungsgemeinschaft (DFG), amongst others within the Schwerpunktprogramm 1276: Met-Strom and the Graduiertenkolleg 1095: Aerothermodynamische Auslegung eines Scramjet-Antriebssystems für zukünftige Raumtransportsysteme and the research projects IDIHOM within the European Research Framework Programme.

References

1. J. S. Hesthaven and T. Warburton, *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications*, 1st Edition, Springer, 2007.
2. F. Lörcher, *Predictor Corrector DG*, PhD thesis, University of Stuttgart, (2008).

3. F. Lörcher and G. Gassner and C.-D. Munz, *A Discontinuous Galerkin Scheme based on a Space-Time Expansion I. Inviscid compressible flow in one space dimension*, J. Sci. Comp., Vol. 32, pp. 175–199, (2007).
4. G. Gassner and F. Lörcher and C.-D. Munz, *A Discontinuous Galerkin Scheme based on a Space-Time Expansion II. Viscous Flow Equations in Multi Dimensions*, J. Sci. Comp., Vol. 34, pp. 260–286, (2007).
5. G. J. Gassner, Frieder Lörcher, Claus-Dieter Munz and Jan S. Hesthaven: "Polymorphic nodal elements and their application in discontinuous Galerkin methods". Journal of Computational Physics, Volume 228, Issue 5, 20 March 2009
6. F. Lörcher, G. Gassner and C.-D. Munz: "Arbitrary High Order Accurate Time Integration Schemes for Linear Problems". European Conference on Computational Fluid Dynamics, ECCOMAS CFD, 2006
7. G. Gassner and F. Lörcher and C.-D. Munz, *An explicit discontinuous Galerkin scheme with local time-stepping for general unsteady diffusion equations*, J. Comput. Phys., Vol. 227, pp. 5649–5670, (2008).
8. G. Gassner and M. Dumbser and F. Hindenlang and C.-D. Munz, *Explicit one-step time discretizations for discontinuous Galerkin and finite volume schemes based on local predictors*, J. Comput. Phys., In Press, Corrected Proof, (2010).
9. F. Hindenlang and G. Gassner and T. Bolemann and C.-D. Munz, *Unstructured high order grids and their application in Discontinuous Galerkin methods*, conference proceedings, V European Conference on Computational Fluid Dynamics ECCOMAS CFD 2010, Lisbon, Portugal, (2010).
10. S. Hickel, *Implicit turbulence modeling for large-eddy simulation*, PhD thesis, TU Dresden, (2005).
11. M.E. Brachet, *Direct simulation of three-dimensional turbulence in the Taylor–Green vortex*, Fluid Dynamics Research, Vol. 8, pp. 1–8, (1991).
12. P.-O. Persson and J. Peraire, *Sub-Cell Shock Capturing for Discontinuous Galerkin Methods*, Proc. of the 44th AIAA Aerospace Sciences Meeting and Exhibit, (2006).
13. G. N. Coleman, J. Kim, R. D. Moser, *A numerical study of turbulent supersonic isothermal-wall channel flow*, J. Fluid Mech, Vol. 305, pp. 159–183, (1995).
14. E. Lenormand, P. Sagaut, L. Ta Phuoc, *Large eddy simulation of subsonic and supersonic channel flow at moderate Reynolds number*, Int. J. Numer. Meth. Fluids, Vol. 32, pp. 369–406, (2000).
15. A. Birkefeld and C.-D. Munz, *A Hybrid Method for CAA*, Proc. of the 36. Jahrestagung für Akustik der Deutschen Gesellschaft für Akustik, DAGA, Berlin, (2010).
16. A. Birkefeld and A. Beck and M. Dumbser and C.-D. Munz and D. König and W. Schröder, *Advances in the Computational Aeroacoustics with the Discontinuous Galerkin Solver NoisSol*, Proc. of the 16th AIAA/CEAS Aeroacoustics Conference (31st AIAA Aeroacoustics Conference), Stockholm, 2010
17. R. Ewert and W. Schröder, *Acoustic Perturbation Equations Based on Flow Decomposition via Source Filtering*, Journal of Computational Physics, Vol. 188, pp. 365–398, (2003).
18. Jan W. Delfs and Marcus Bauer and Roland Ewert and Herwig A. and Grogger and Markus Lummer and Thomas G. W. Lauke, *Numerical Simulation of Aerodynamic Noise with DLR's aeroacoustic code PIANO*, Deutsches Zentrum für Luft- und Raumfahrt e.V., Institut für Aerodynamik und Strömungstechnik
19. R. Ewert, *Broadband slat noise prediction based on CAA and stochastic sound sources from a fast random particle mesh RPM) method*, Computers & Fluids, Vol. 37, pp. 369387, 2008
20. S. Orszag, *Numerical simulation of the Taylor-Green vortex*, Computing Methods in Applied Sciences and Engineering Part 2, Lecture Notes in Computer Science, Vol. 11, pp. 50–64, (1974).
21. D. Fauconnier, *Development of a Dynamic Finite Difference Method for Large-Eddy Simulation*, PhD thesis, University of Gent, Belgium, (2009).
22. A. Beck, G. Gassner, I. Horenko, R. Klein, C.-D. Munz, *Technischer Report im Rahmen des Schwerpunktprogramms 1276 (MetStroem)*, Lugano, Berlin, Stuttgart, (2011).

23. A.G. Tomboulides and S.A. Orszag, *Numerical investigation of transitional and weak turbulent flow past a sphere*, J. Fluid Mech., Vol. 416, pp. 45–73, (2000).
24. Bogey, C., Bailly, C. and Juv, D., *Noise investigation of a high subsonic, moderate Reynolds number jet using a compressible LES*, Theoretical and Computational Fluid Dynamics, 16(4), 273–29, 2003.
25. J. Dierke, C. Appel, J. Siebert, M. Bauer, M. Siefert and R. Ewert, *3D Computation of Broadband Slat Noise from Swept and Unswept High-Lift Wing Sections*, Proceedings of the 16th AIAA/CEAS Aeroacoustics Conference, Portland, AIAA paper 2011–2905 (2011).
26. D. P. Lockard and M. M. Choudhari, *Noise Radiation from a Leading-Edge Slat*, AIAA paper 2009–3101 (2009).
27. G.J. Gassner and A.D.Beck, *On the accuracy of high-order discretizations for underresolved turbulence simulations*, Theoretical and Computational Fluid Dynamics, DOI 10.1007/s00162-011-0253-7

Computational Aeroacoustics with Higher Order Methods

E. Rebecca Busch, Michael S. Wurst, Manuel Keßler, and Ewald Krämer

Abstract The Lighthill acoustic analogy in combination with two different higher-order CFD solvers is used to investigate the sound generation of two test cases. The flow around a cylinder at $Re = 150$ is analysed with a Discontinuous Galerkin method and a counter-rotating open rotor (CROR) with a WENO scheme. The simulation of the cylinder is able to predict both aerodynamic and acoustic behaviour correctly, the vortex street behind the cylinder is responsible for the noise radiation similar to of an acoustic dipole. The analysis of the CROR focuses on the effect of using a higher-order method with a detailed comparison with a standard second order method. While global aerodynamic forces show only small differences, the better transport of vortices, especially of the blade tip vortex, is a benefit for the prediction of interaction noise of the two rotors. This paper includes different investigations on the new HLRS Cray Hermit cluster. The DG code was optimized for single-core usage while still maintaining its good parallel performance. The effect of node-pinning is studied with the CROR configuration which improved the computational time slightly.

1 Introduction

For recent research in aerospace engineering, computational fluid dynamics (CFD) in combination with computational aeroacoustics (CAA) has developed to a useful tool to examine flow fields and acoustic emission caused by flow phenomena, in addition to experiments. With the continuous growth of supercomputing power it is now not only possible to simulate setups with several million cells in acceptable

E.R. Busch (✉) · M.S. Wurst · M. Keßler · E. Krämer
Institut für Aerodynamik und Gasdynamik, Universität Stuttgart, Pfaffenwaldring 21, D-70569,
Stuttgart, Germany
e-mail: busch@iag.uni-stuttgart.de

time, also the use of higher order methods is a possibility. At the Institute of Aerodynamics and Gas Dynamics (IAG) of the University of Stuttgart the development of higher order codes used for CFD simulation of helicopters has been one of the main research fields. Current CFD simulations often include acoustic evaluation since noise emission of aircrafts plays a more and more important role due to its high annoyance to the community. Aeroacoustics are a part of the current research at IAG. In this paper two test cases were simulated with different numerical methods and then examined acoustically. A cylinder at a Reynolds number $Re = 150$ has been simulated with the higher order Discontinuous Galerkin (DG) Code SUNWinT developed by IAG. The SUNWinT Code is still being enhanced and will be used for CFD simulation of helicopters and other applications. Subsequently an acoustic analysis has been carried out with the IAG-tool ACCO, which uses the Ffowcs Williams-Hawkings (FW-H) equation for acoustic modelling. As a practical application, a counter-rotating open rotor at take-off conditions has been simulated with a state of the art finite volume code, namely standard FLOWer, developed by DLR, and WENO FLOWer, enhanced by IAG. The acoustic evaluation has been carried out with ACCO as well. Both FLOWer and ACCO are also used for CFD and CAA simulation of helicopters at IAG. After preceding performance studies all CFD simulations have been carried out on the new Hermit cluster at the High Performance Computing Center in Stuttgart (HLRS).

2 Numerical Methods

2.1 *Discontinuous Galerkin Code SUNWinT*

2.1.1 Basic Features

The Discontinuous Galerkin (DG) method as it is used in fluid mechanics combines ideas from finite volume (FV) discretisation techniques as well as from finite element (FE) discretisation techniques. A typical FE feature of the DG method is that the representation of the solution in a cell is given by a polynomial approximation, whose accuracy can easily be improved by increasing its polynomial order. However, the solution between cells is discontinuous, thus the solution of this Riemann problem requires approximate Riemann solvers known from FV methods. The method was originally developed for the solution of hyperbolic conservation laws as the Euler equation [13] which contains only first-order derivatives. It was extended for the Navier-Stokes equations by Bassi and Rebay [12] and is now capable to solve equations containing second-order derivatives.

2.1.2 Numerical Formulation

The Navier-Stokes equations are given by

$$\frac{\partial U}{\partial t} + \nabla \cdot F_i(U) = \nabla \cdot F_v(U, \nabla U) . \quad (1)$$

Herein U is the vector of conservative variables, F_i and F_v are the convective and the diffusive fluxes. In order to eliminate the second order derivatives of U , the equation is transformed in a first-order system by introducing the gradient of U as an additional solution variable. The new system reads as

$$\Theta - \nabla U = 0 , \quad (2)$$

$$\frac{\partial U}{\partial t} + \nabla \cdot F_i(U) = \nabla \cdot F_v(U, \nabla U) . \quad (3)$$

Discretising both equations with the DG approach results in

$$\int_{\Omega} V_h \Theta_h d\Omega + \int_{\Omega} \nabla V_h \cdot U_h d\Omega - \int_{\partial\Omega} V_h U_h \cdot \mathbf{n} d\sigma = 0 , \quad (4)$$

$$\int_{\Omega} V_h \frac{\partial U_h}{\partial t} d\Omega - \int_{\Omega} \nabla V_h \cdot F_i d\Omega + \int_{\partial\Omega} V_h F_i \cdot \mathbf{n} d\sigma = - \int_{\Omega} \nabla V_h \cdot F_v d\Omega + \int_{\partial\Omega} V_h F_v \cdot \mathbf{n} d\sigma \quad (5)$$

where U_h and V_h are L2-integrable, piecewise polynomial functions. We choose a hierarchical and orthogonal basis according to Sherwin and Karniadakis [19]. Both explicit as well as implicit time integration schemes are available within our code [18].

2.2 The FLOWer Code

FLOWer is a finite-volume code originally developed by DLR and enhanced by IAG. It solves the three-dimensional Reynolds averaged Navier-Stokes equations in a block-structured domain for numerous time steps [1, 2]. For spatial discretisation a second order cell-centered finite volume formulation is used. A hybrid multi-stage Runge-Kutta scheme developed by Jameson is applied [3, 4] for time integration. Additionally, a WENO scheme has been recently implemented into FLOWer by IAG. Movement and complex geometry structures such as rotors or nacelles are embedded with the chimera method. A component grid is integrated into a background mesh where a hole is cut to make room for the component grid cells. In the overlapping zones between the grids the data exchange is carried out. For moving structures like rotors this process is done anew for every time step. With this method complex structures can be meshed with structured grids and a local refinement is possible while keeping the same background mesh.

2.2.1 The Finite Volume Method

Generally, a finite volume scheme is applied since it has the advantage of adaptability to arbitrary meshes without transforming into a computational domain when a good quality grid is used. The flow field is split up into hexahedral cells (control volumes) where the Navier-Stokes equations are applied in integral form on each cell taking care of conservativity [5]. For each cell a discrete flux balance is calculated and the change of flow quantities in time in particular points can be defined. The scheme is given by the equation

$$\frac{d}{dt} u_i = -\frac{1}{C_i} \int_{\partial C_i} \mathbf{f}(u) \mathbf{n} dS$$

where u_i is the vector of conservative variables in this case, \mathbf{n} the normal vector and $\mathbf{f}(u)$ represent the numerical fluxes for a particular control volume arrangement with the volume C_i surrounding the grid node with the indices i, j, k and the cell face ∂C_i . The fluxes are approximated using either a central or an upwind discretisation operator.

2.2.2 The Jameson Scheme

The Jameson scheme utilises a second order approximation and stabilises it with high-order Runge–Kutta time stepping and adding an artificial dissipative term at the end of each time step [3, 4]. It is the standard method used in FLOWER. The conservative variables are simply averaged at the cell faces by

$$u_{i+\frac{1}{2}} = \frac{1}{2}(u_i + u_{i+1})$$

2.2.3 The WENO Scheme

WENO (“weighted essential non-oscillatory”) is a further development of the ENO (“essentially non-oscillatory”) scheme introduced by Harten et al. [6]. The basic idea is reconstructing or combining lower order fluxes for a higher order approximation. In the ENO scheme the least oscillatory stencil is then picked for reconstruction whereas for WENO all stencil candidates are weighed by their smoothness and subsequently used for reconstruction [7]. The recombination of stencils is done by

$$u_{i+\frac{1}{2}} = \sum_{r=0}^{k-1} \omega_r u_{i+\frac{1}{2}}^r$$

where ω_r are the weights so that $\sum_{r=0}^{k-1} \omega_r = 1$ and k the order of stencils. WENO shows a significant performance increase to ENO as the selection constraints

are omitted. Three stencils are implemented at an order of 3, leading to a fifth order reconstruction. However this is only valid on Cartesian grids with equal cell distances in all directions.

2.3 Acoustic Post-processing

Acoustic analysis is done with the IAG tool ACCO which uses the FW-H equation:

$$\begin{aligned} \frac{\partial^2 \rho'}{\partial t^2} - c^2 \nabla^2 \rho' = & \frac{\partial^2}{\partial x_i \partial x_j} [T_{ij} H(f_S)] \\ & - \frac{\partial}{\partial x_i} [(p' n_i + \rho u_i (u_n - v_n)) \delta(f_S)] \\ & + \frac{\partial}{\partial t} [(\rho_0 v_n + \rho (u_n - v_n)) \delta(f_S)] \end{aligned}$$


where ρ' denotes the density fluctuation, c the speed of sound, p' the pressure fluctuation, n_i the normal vector, u_n the surface component of the velocity, v_n its normal component and $\delta(f)$ the Dirac delta function [8, 9]. The wave equation on the left hand side describes the propagation of sound in space while the right hand side specifies the monopole and dipole source terms on surfaces and quadrupole sources terms in volumes. The FW-H equation allows the integration surfaces to surround the noise generating flow structures and geometries. This saves the integration of volume sources inside in addition to the surface terms on the geometric surfaces. However, this requires the choice of the surrounding surfaces far enough from geometric surfaces, so that all source terms are included, but close enough so that numerical dissipation has not significantly damped pressure fluctuations. Hence the use of high order methods with a good vortex transport is worthwhile here. For acoustic analysis surfaces extracted from the flow field are exported for every time step and then analysed by ACCO. ACCO returns total sound pressure levels as well as pressure fluctuation over time. Additionally, a Fast-Fourier transformation can be carried out.

3 Computational Aspects

3.1 Enhancement and Performance Study of the SUNWinT Code

Our code SUNWinT showed an excellent parallel speedup before, e.g. a parallel efficiency up to 85 % using 2,048 cores was reached on both the HLRS NEC Nehalem Cluster as well as the HLRB II SGI Altix system [16, 17]. Sustaining this parallel

Fig. 1 Decomposition of a mesh in patches



speedup on new and upcoming HPC systems is the aim of current developments in our DG code. In a first step the single core performance of the code was improved in a major code redesign. The evolution of the surface and volume integrals in Eq. (5) is done now in small groups called patches. These patches consist of around 15 cells. A typical decomposition of the mesh in patches can be seen in Fig. 1.

The idea behind these patches is that the data in the patches is stored in a memory efficient way and so it could be transferred quickly from the memory to the CPU. Basically, there are two different kind of patches in a mesh, inner patches and border patches. Inner patches are all patches fully inside the computational domain, which do not border to a physical boundary or to a domain handled on another processor. These patches can be handled in the exactly same way. Border patches need a special treatment to ensure that they fulfill their boundary condition or to handle the communication between different processors via MPI. Additionally, the evaluation of the surface and volume integrals inside each patch was simplified in such a way, that all time-independent parts of these integrals are precalculated and efficiently combined with the time-dependent parts. A comparison of the single core performance of the new evaluation is shown in Table 1 for a 2D case and in Table 2 for a 3D case. It can be seen that the computational time reduces significantly both for the Euler equation as well as the Navier–Stokes equations. The saving in computational time is larger doing higher-order calculations, e.g. for fourth order calculations the new evaluation is five times faster. Comparing 2D and 3D cases, no differences are present, the reduction in computational time is about the same. In the future, it is planned to use the patches for a hybrid parallelisation of our code. The idea is, that each node will handle a large amount of patches and will assign a patch to a free core for evaluation if the core is available.

The determination of the parallel performance on the HLRS Hermit cluster is done with two different cases, the 3D simulation from the next chapter and another testcase. Both meshes are decomposed with METIS [15] from 16 zones up to 16,384 zones. The first mesh consists of 336,000 cells, leading to only 82 cells per core, which are calculated on average on the finest decomposition level. The second contains 930,000 cells leading to an extremely low cell count of 57 cells on the finest level. Despite the small number of cells per core, the scaling in Fig. 2a is

Table 1 Comparison of computational time for old and new evaluation for a 2D case with different order and physical modelling on a single processor for 1,000 time steps

Phys. model		First order	Second order	Third order	Fourth order
Euler	Old	10.18	76.1	371.84	1,267.46
	New	5.63	25.16	76.98	201.07
	Reduction(%)	44.7	66.9	79.3	84.1
Navier–Stokes	Old	28.02	235.6	1,042.98	2,725.31
	New	12.23	67.81	227.06	571.53
	Reduction(%)	52.7	72.9	79.5	79.6

Table 2 Comparison of computational time for old and new evaluation for a 3D case with different order and physical modelling on a single processor for 200 time steps

Phys. model		First order	Second order	Third order	Fourth order
Euler	Old	15.26	139.95	1,017.45	5,254.97
	New	7.66	44.05	187.97	1,218.78
	Reduction(%)	49.8	68.5	81.5	76.8
Navier–Stokes	Old	43.53	422.62	2,727.41	13,355.66
	New	24.07	147.43	652.71	2,079.4
	Reduction(%)	44.7	65.1	76.1	84.4 %

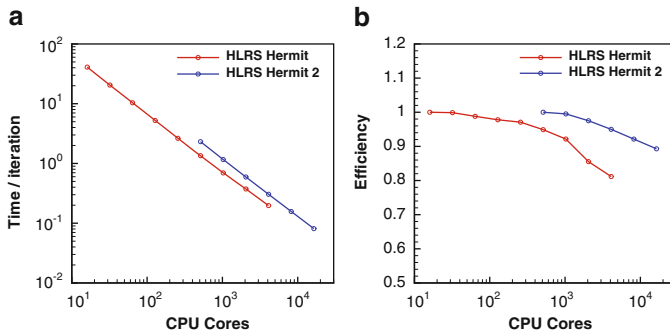


Fig. 2 (a) Strong scaling and (b) efficiency on the HRLS Hermit

very good for both cases, the time step per iteration decreases almost linearly up to 4,096 cores with a parallel efficiency of 81 % (cf. Fig. 2b) for the first case and up to 16,384 cores with a parallel efficiency of 89 % for the second case.

3.2 Performance Study with a Counter-Rotating Open Rotor Setup with FLOWer

In order to determine the fastest way of computation a number of performance tests were run on the Hermit (Cray XE6) cluster at HRLS. The original setup

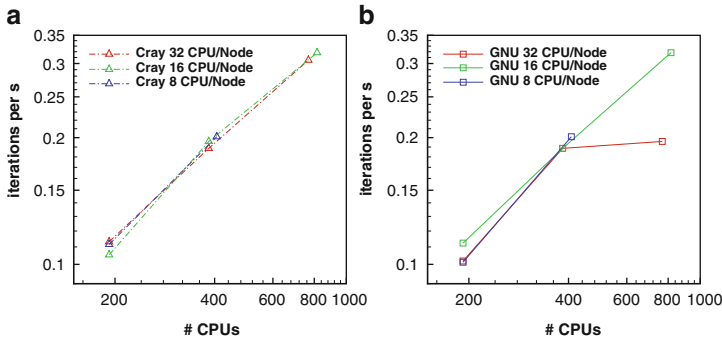


Fig. 3 Performance of FLOWER code with node pinning. (a) Cray compiler and (b) GNU compiler

consists of roughly 50 million cells, where about 25 million cells can be assigned to the background mesh, which contains the cylindrical geometry for a simplified modelling of the CROR's nacelle. The 9×7 rotor setup is modelled with one component grid per blade containing 1.5 million cells for each mesh. The total setup consists of 816 structured blocks with a maximum size of 75,000 cells per block. Figure 3a and b show the iterations per seconds carried out for different number of CPUs used per node. For the Cray compiler all three curves show a good parallelisation while with the GNU compiler the efficiency significantly decreases when all 32 CPUs of a node are used for a large number of CPUs. This may be due to a worse communication between the CPUs of the GNU compiled code, compared to the Cray compiler. For both compilers node pinning does not show a clear improvement except for the use of 816 CPUs on 51 nodes for the Gnu compiler. Only for small numbers of CPUs like 192 a slight improvement with using every second CPU can be seen. Using every fourth CPU has even smaller effects on CPU time.

4 Results

4.1 Laminar Flow Around a Cylinder

In order to study the usage of ACCO with our DG code SUNWinT the flow around a cylinder is calculated. The flow is investigated at $Re = 150$ based on the diameter D of the cylinder, and the freestream Mach number is $Ma = 0.1$. Both 2D as well as 3D simulations were performed. The computational domain is $60D$ large and has a depth of two diameters in the 3D case. The mesh is of O-grid type and consists of $120 \times 70 (\times 40)$ cells. We use fourth order accurate P3 elements for the spatial discretisation and the temporal discretisation is chosen equally with an explicit, fourth order, one-step Runge–Kutta type scheme. The results are compared with

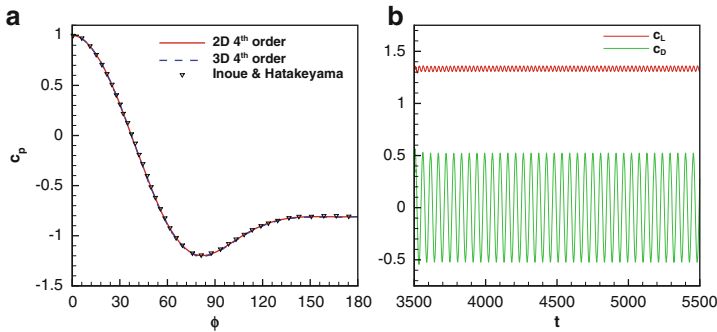


Fig. 4 (a) Time-averaged pressure coefficient c_p for the 2D and 3D simulation, (b) time-dependent lift and drag coefficient for the 2D simulation

Table 3 Global forces C_D , $C_{L,a}$ and the Strouhal number St

	C_D	$C_{L,a}$	St
2D	1.33	0.52	0.187
3D	1.33	0.52	0.184
Inoue and Hatakeyama	1.32	0.52	0.183

the results of Inoue and Hatakeyama [14] who used a sixth order finite difference scheme for this problem.

The aerodynamic results of the 2D and 3D simulations are in very good agreement with the simulation of Inoue and Hatakeyama as the time-averaged pressure coefficients in Fig. 4a indicate. As a consequence, global forces are also quite similar (cf. Table 3). This kind of flow typically develops a vortex street behind the cylinder which can be seen in the time-dependent distribution of the lift and the drag coefficient (cf. Fig. 4b). The Strouhal number of this vortex shedding is in the 2D case 0.187 and in the 3D case 0.184 (Inoue and Hatakeyama: $St=0.183$). Comparing the amplitude of both forces in Fig. 4b it is shown that the amplitude of the lift coefficient is much larger than for the drag coefficient.

For the acoustic analysis the results of the 2D simulations are used. The hull surfaces for the acoustic solver ACCO are placed in a distance of three cylinder diameters away from the center of the cylinder. The observers for the acoustic analysis are placed on a circle $125D$ away from the cylinder. Figure 5a shows the sound pressure level of an observer position located 90° from the freestream direction for different frequencies. The acoustic behaviour is dominated by the aerodynamical behaviour, more precisely, the vortices induced by the cylinder leading to the oscillating forces. Peaks in the SPL are present at the Strouhal frequency and at higher harmonical frequencies. This dipole nature of the generated sound is seen on a directivity pattern (Fig. 5b) for the fundamental frequency and is in good agreement with the findings of Inoue and Hatakeyama who also detected the lift dipole as the major acoustic source.

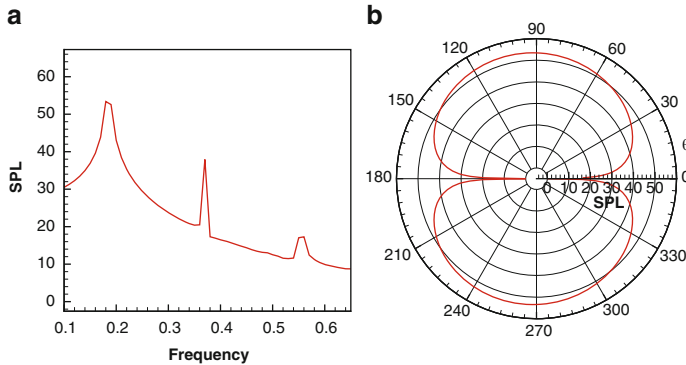


Fig. 5 (a) SPL distribution vs. frequency for an observer at 90° from the freestream direction, (b) directivity pattern for the fundamental frequency

4.2 A Counter-Rotating Open Rotor in Take-Off Conditions

A 9×7 CROR configuration was simulated with standard FLOWer settings and WENO FLOWer at take-off conditions [10]. The environmental conditions were set to ICAO standard atmosphere conditions at sea-level. The simulation was carried out for eight rotor revolutions for both cases on the Hermit cluster at HLRS. The total computation time for standard FLOWer settings was approximately 240 h on 384 CPUs for standard FLOWer settings and 270 h for WENO FLOWer. The WENO scheme was only carried out in a user-set number of blocks in the background grid in proximity to the rotors.

4.2.1 Aerodynamic Results

For comparison of the aerodynamic results integrated values such as thrust coefficient, drive torque coefficient and propulsion efficiency are examined. Additionally the vortex conservation is shown as it is especially important for the rotor-rotor interaction which has a considerable contribution to CROR noise.

Figure 6a and b show the thrust coefficient C_T and the torque coefficient C_M plotted over one rotor revolution. Both front and aft rotor show higher values for C_T and C_M accordingly for WENO. This is due to the higher order of the WENO FLOWer code which is similar to effects seen in previous grid studies, where higher grid resolution showed slightly higher thrust and drive torque levels. Hence the grid resolution used in this calculation still shows effects of grid dependency. Additionally the standard FLOWer code shows higher fluctuations of C_T and C_M especially for the aft rotor, indicating that it cannot represent the highly instationary incoming flow as good as the WENO FLOWer. The fluctuations of C_T and C_M for the front rotor are in roughly the same range for the standard and the WENO version.

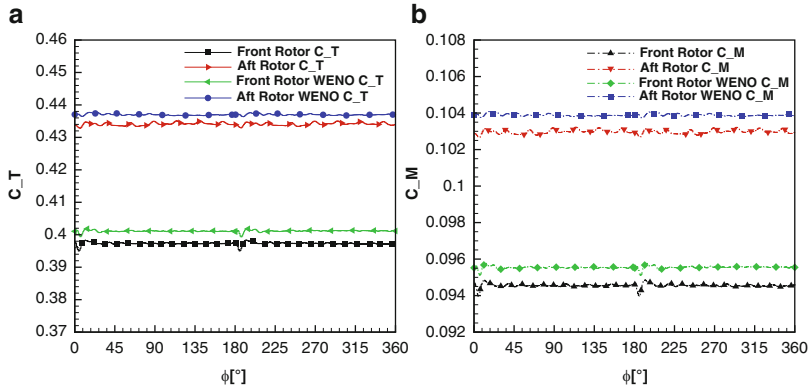


Fig. 6 (a) C_T and (b) C_M over one rotor revolution




Fig. 7 Vortex visualisation of the flow field calculated with (a) standard FLOWer and (b) WENO FLOWer

The characteristic numbers such as C_T and C_M differ by less than 1% for front and aft rotor as well as the CRORs totals. For the front rotor the same propulsion efficiency is reached with standard and WENO FLOWer whereas for the aft rotor and the total CROR the propulsion efficiency differs by 0.1%. However, both FLOWer versions show excellent agreement, but the WENO version is generally more capable of unsteady flow phenomena.

Comparing Fig. 7a and b the vortex transport with the WENO version is better than with the standard version. Especially the vortices that have already passed the aft rotor plane are conserved better with WENO. Wakes and vortices are also more clearly defined with WENO.

4.2.2 Acoustic Results

The acoustic emission of the CROR was evaluated over one rotor revolution with a resolution of 360 time steps for both standard and WENO FLOWer. The IAG tool ACCO was used to obtain sound pressure levels at microphone observers in the x - y plane, where the x -axis is the rotational axis pointing in flow direction and the y -axis is perpendicular so that x and y span a plane perpendicular to the rotors. The microphones were put in a distance of $\theta = 5^\circ$ to each other with $\theta = 0^\circ$ on the negative x -axis and $\theta = 180^\circ$ on the positive x -axis. $\theta = 90^\circ$ is located at a radius of $r = 5$ m to the rotational axis with equal distance to front and aft rotor. Comparing the noise directivities generated from standard FLOWer and WENO FLOWer CFD simulations, it is visible that both show the same characteristics, that is, a main peak at $\theta \approx 95^\circ$ and minor peaks at lower and higher θ , see Fig. 8a. These minor peaks are significantly higher for the WENO case whereas the main peak reaches the same levels for both simulations. The main peak shows a stronger decay for angles below and above $\theta \approx 95^\circ$ for the standard FLOWer simulation.

Figure 8b, c show the tonal noise of front ($BPF1$) and aft rotor blade ($BPF2$) passing frequencies and one higher harmonic ($2BPF1$, $2BPF2$) which are the main contributions of noise in radial direction of the rotors. For the BPFs both code versions show a very good agreement while for the higher harmonic a discrepancy especially on the front rotor can be seen. The nine blade front rotor has a higher BPF and first harmonic than the seven blade aft rotor. The ability of resolving frequencies in the acoustic analysis is strongly dependent on the temporal and spatial resolution of the CFD simulation, hence WENO FLOWer can still resolve the first harmonic of the front rotor opposing to the standard FLOWer. For the aft rotor this effect cannot be seen yet as standard FLOWer is able to resolve the first harmonic as it lies at a lower frequency.

Figure 8d–f show the rotor-rotor interaction noise which occurs at frequencies that are a linear combination of BPFs such as $BPF1 + BPF2$, $2BPF1 + BPF2$ and $BPF1 + 2BPF2$. The lowest interaction frequency $BPF1 + BPF2$ is mainly responsible for the minor peaks on the total noise, cf. Fig. 8a. Here the discrepancies between standard and WENO FLOWer become exceptionally clear. The transport of the front rotor wake and blade tip vortices is worse using standard FLOWer than using WENO FLOWer for CFD calculations. For the second interaction frequency (Fig. 8e) this effect can still be examined while for the third interaction frequency (Fig. 8f) it is not visible at all. This leads to the conclusion that also the resolution of WENO FLOWer is not sufficient to obtain noise directivities at these high frequencies.

For the BPFs, their first harmonics and lower interaction frequencies good results can be achieved with both code versions. These are the relevant frequencies as they are the main contributors to total noise. Their level is roughly 10 dB higher than second or higher harmonics of BPFs and higher interaction frequencies. Both codes show good agreement for the low frequencies and as expected the results differ at higher frequencies. The main noise characteristics can be examined with the data obtained by both CFD simulations, but standard FLOWer reaches the

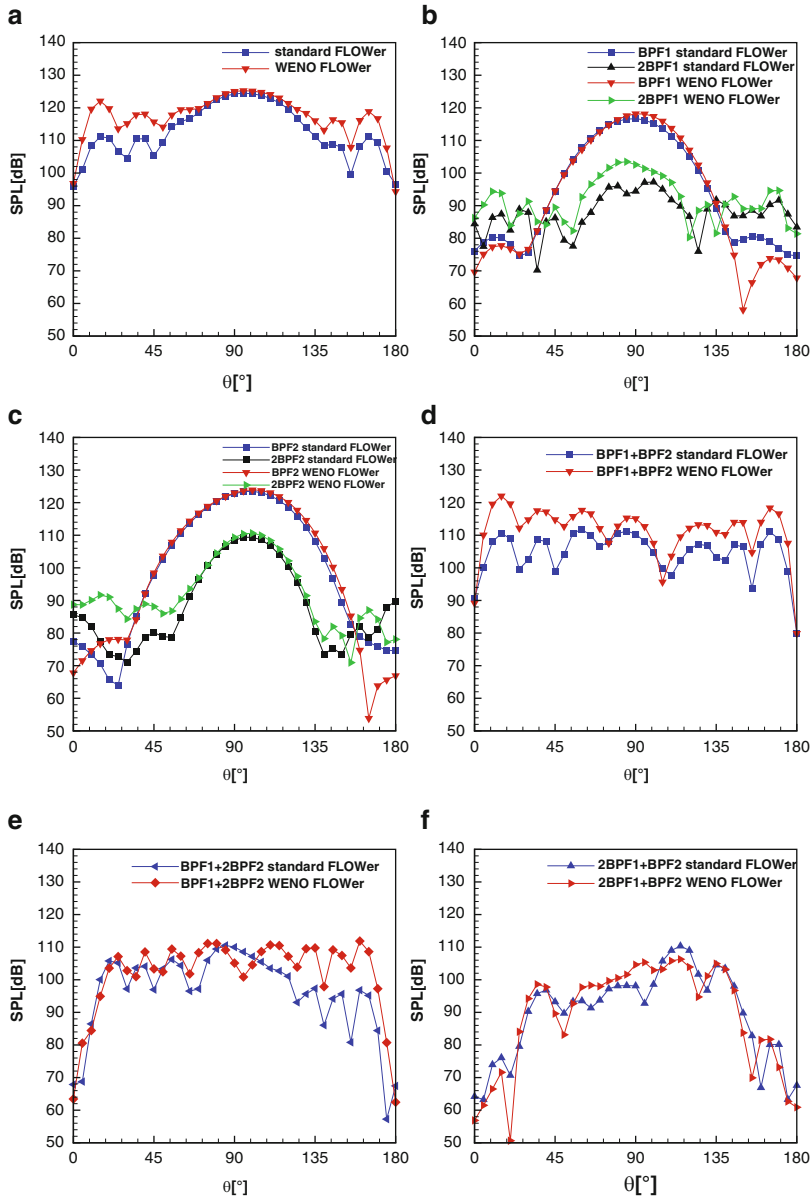


Fig. 8 Noise directivity in x - y -plane. (a) Total sound pressure level, (b) blade passing frequency of front rotor, (c) blade passing frequency of aft rotor, (d) first interaction ($BPF1 + BPF2$) frequency, (e) second interaction frequency ($BPF1 + 2BPF2$), and (f) third interaction frequency ($2BPF1 + BPF2$)

limits of resolution with first harmonics of BPFs and shows weakness for the lower interaction frequencies. Yet the directivities for lower interaction frequencies gained with standard FLOWer still show the same characteristics as with WENO FLOWer. A good qualitative agreement with the results of measurements by Woodward could be achieved for both standard and WENO FLOWer [11].

5 Conclusion

Two different test cases were simulated in this paper with two different CFD codes and analysed acoustically. Both codes are capable of discretising the Navier–Stokes equations with a higher order method. The simulation of the flow around a cylinder at $Re = 150$ with the DG code SUNWinT was a first step to analyse the potential of an acoustic analysis with the DG method. It is in good agreement with the results of Inoue and Hatakeyama [14]. The aerodynamic forces and the distribution of the time-averaged pressure coefficient do not show any differences. The behaviour of the cylinder as an acoustic lift dipole was also clearly verified by the peaks of the SPL in the frequency analysis and in a directivity pattern of the fundamental frequency. In a next step, the method will be used for the acoustic analysis of turbulent flows. Concerning the performance of SUNWinT different improvements have been made which reduced the computational time while maintaining the good parallel performance.

A 9×7 CROR configuration was simulated with standard and WENO FLOWer at take-off conditions. Aerodynamic and acoustic results showed a good agreement with both code versions. Only a slight discrepancy can be detected for the integrated aerodynamic values. Comparison of the flow fields obtained with the different versions illustrates that the vortex transport is better with WENO FLOWer. However with good spatial and temporal resolution an acceptable transport can also be achieved with standard FLOWer. The acoustic evaluation showed that for the single rotor noise the two codes versions hardly differ. The interaction noise contribution can be examined much better with WENO FLOWer. This was expected as WENO FLOWer showed a better vortex transport which causes the interaction of the front rotor wake with the aft rotor leading to the interaction noise. Generally the use of a higher order method is advisable for acoustic analysis as the quality of the acoustic results is directly dependent on the CFD results with the FW-H method used in this paper. With WENO FLOWer the computational effort is acceptable compared to the beneficial improvement of results.

Acknowledgements We greatly acknowledge the provision of supercomputing time and technical support by the High Performance Computing Center Stuttgart (HLRS) for our project HELISIM.

References

1. Kroll, N., Eisfeld, B., Bleecke, H. M.: FLOWer, In: Notes on Numerical Fluid Mechanics Vieweg Braunschweig, Germany, Vol. 71, p. 58–68, 1999
2. Schwarz, Th. O.: Ein blockstrukturiertes Verfahren zur Simulation der Umströmung komplexer Konfigurationen, Institut für Aerodynamik und Strömungstechnik, Universität Braunschweig, PhD thesis, 2005
3. Jameson, A., Schmidt, W., Turkel, W.: Numerical Solution of the Euler Equations by Finite Volume Methods using Runge-Kutta Time-Stepping Schemes, AIAA 14th Fluid and Plasma Dynamic Conference, Palo Alto, California, USA, 1981
4. Jameson, A.: Time Dependent Calculations using Multigrid, with Applications to Unsteady Flows past airfoils and wings, AIAA 10th Computational Fluid Dynamics Conference, Honolulu, Hawaii, USA, 1991
5. Eymard, R., Gallouet, T., Herbin, R. (2006) The Finite Volume Method
6. Harten, A., Enquist S., Osher, S., Chakravarthy, S. (1987) Uniformly high order non-oscillatory schemes, *Journal of Computational Physics*, Vol. 71, p. 231–202
7. Shu, C.-W. (2003) High-order Finite Difference and Finite Volume WENO Schemes and Discontinuous Galerkin Methods for CFD, *International Journal of Computational Fluid Dynamics*, Vol. 17, No. 2, p. 107–118
8. Keßler M., Wagner S. (2004) Source-Time Dominant Aeroacoustics, *Computers & Fluids*, Vol. 33, p. 791–800
9. Ffowcs Williams, J. E., Hawkings, D. L. (1969) Sound Generation by Turbulence and Surfaces in Arbitrary Motion, *Philosophical Transactions of the Royal Society*, Vol. A264, No. 1151, 321–342
10. Busch, R., Keßler M., Krämer E., (2012) Numerical Investigation of Counter-Rotating Open Rotor Noise Emission in Different Flight Conditions, ASME Turbo Expo 2012, Copenhagen, Denmark [accepted]
11. Woodward, R.P. (1992) Noise of two high-speed model counter-rotation propellers at takeoff/approach conditions, *Journal of Aircraft*, Vol. 29, No. 4, p. 679–685
12. Bassi F., Rebay S., (1997) A high-order accurate discontinuous finite element method for the numerical solution of the compressible navierstokes equations, *Journal of Computational Physics*, Vol. 131, p. 267–279.
13. Bassi F., Rebay S., (1997) High-order accurate discontinuous finite element solution of the 2d euler equations, *Journal of Computational Physics*, Vol. 138, p. 251–285.
14. Inoue O., Hatakeyama N., (2002) Sound generation by a two-dimensional circular cylinder in a uniform flow, *J. Fluid. Mech.*, Vol. 471, p. 285–314.
15. Karypis G. and Kumar V., A fast and high quality multilevel scheme for partitioning irregular graphs, *Siam Journal on Scientific Computing*, Vol. 20, No. 1, p. 359–392.
16. Lübon C, Keßler M., Wagner S., (2009) A parallel CFD solver using the discontinuous galerkin approach, *High Performance Computing in Science and Engineering, Garching/Munich 2007*, Springer Berlin Heidelberg, p. 291–302.
17. Lübon C, Keßler M., Wagner S., (2010) A parallel cfd solver using the discontinuous galerkin approach, *High Performance Computing in Science and Engineering, Garching/Munich 2009*, Springer Berlin Heidelberg, 197–205.
18. Landmann B., Keßler M., Wagner S., Krmer E., (2008) A parallel, high-order discontinuous galerkin code for laminar and turbulent flows, *Computers & Fluids*, Vol. 37, No. 4, p. 427–438
19. Sherwin S.J., Karniadakis G.E., (2004) Spectral/hp element methods for computational fluid dynamics

Effects of an Oblique Roughness on Hypersonic Boundary-Layer Transition

Gordon Groskopf and Markus J. Kloker

Abstract The compressible bi-global linear stability theory (B-LST), based on two-dimensional eigenfunctions in flow crossplanes, as well as direct numerical simulation (DNS) are applied to investigate the stability properties of three-dimensional laminar hypersonic boundary-layer flows with discrete surface roughness. The obliquely-placed fence-like roughness element has a height of half the unperturbed thermal boundary-layer thickness at its position on the flat plate. The roughness setup is derived from a Space Shuttle flight experiment. A cold-flow non-reacting gas case at wind-tunnel conditions with Mach 4.8 is considered. The laminar steady base flow is extracted from a (D)NS base-flow solution assuming perfect-gas behavior. Local and integral growth of instabilities in the wake of the discrete roughness element are investigated revealing a considerable persistent amplification in the flow. The cold-flow results are compared to a recent instability analysis of a hot-flow, reacting gas case with identical Mach number and roughness geometry. A full DNS of the cold-flow case validates the B-LST results: The growth of the excited disturbance modes shows good agreement. Performance data for the applied codes running on the NEC SX-9 and CRAY XE6 are given.

1 Introduction

Boundary-layer transition at hypersonic flow speeds is of particular importance to the design of respective vehicles. On one hand the vehicle structure encounters a massive instantaneous as well as integral heat load. The amount of heat the vehicle has to withstand is significantly influenced by surface roughness, e.g., the gap filler protruding from the thermal protection system of the Space Shuttle (Fig. 1). In case of sustained hypersonic flight, aerodynamic drag is another crucial issue, especially

G. Groskopf (✉) · M.J. Kloker
Institut für Aerodynamik und Gasdynamik, Universität Stuttgart, D-70550, Stuttgart, Germany
e-mail: groskopf@iag.uni-stuttgart.de; kloker@iag.uni-stuttgart.de

Fig. 1 Ceramic-tile gap filler protruding from the thermal protection system of the Space Shuttle (Source: NASA (www.nasa.gov))



with regard to fuel consumption. Both heat load and drag are to be minimized by appropriate measures. On the other hand there are cases where the flow is desired to be turbulent to alleviate the effects of flow separation, or to enhance mixing. Therefore it is essential to understand how discrete surface roughness influences and alters the laminar-turbulent transition mechanisms.

Progress has been achieved recently in this field. The experimental approach (see, e.g., [6]) peaked in the STS-119 flight experiment. A numerical ansatz has been applied by, e.g., [3, 4, 15, 16]. Theoretical analyses of the disturbance growth in steady base flows have been done by, e.g., [7–11, 14].

Applying the B-LST in crossplanes in the wake of the roughness element, [7] and [9] show that symmetric as well as oblique, with respect to the oncoming flow, roughness geometries with half the height of the undisturbed boundary layer enhance persistently the growth of first-mode-type instability modes due to the main trailing vortices and ensuing velocity streaks in the wake of the roughness. Within the parameter range investigated the contribution of the horseshoe vortices is of minor importance. This is in accordance with “Condition I” of [3]. The horseshoe vortices attain primary importance if the height of the roughness element reaches or exceeds the undisturbed boundary-layer thickness.

While in [3] the unsteady flow behind a three-dimensional, cylindrical roughness element with a $O(\delta)$ -height has been simulated by means of a Navier–Stokes solver, [15] investigates the evolution of well-defined disturbances in a flat-plate boundary-layer flow altered by a two-dimensional roughness element by means of unsteady DNS, supported by results from primary linear stability theory. It is shown that a roughness with a size of up to 70 % of the undisturbed boundary-layer thickness alters the stability properties of the flow only locally. Such a discrete 2-d element represents a disturbance amplifier with a limited bandwidth: A selected frequency range gains in amplitude locally by the roughness.

In this paper we investigate the evolution of disturbances in a flow at wind-tunnel conditions altered by a three-dimensional, obliquely-placed roughness. The shape of the roughness element follows the geometry used in [10]. It is derived from the STS-119 flight experiment, which investigated the influence of an isolated roughness, mounted on the Space Shuttle’s belly, on laminar-turbulent transition during re-entry. The results of the stability analysis are compared to the hot-flow case from

[10] and [11] exhibiting the same roughness geometry with respect to the thermal boundary-layer thickness δ_T at the same position on the flat plate in terms of Re_{δ_T} and Mach number.

2 Governing Equations, Numerics and Flow Setup

2.1 Governing Equations

The three-dimensional unsteady compressible Navier–Stokes equations are used in a non-dimensional form. Length scales are normalized by a reference length L^* (* marks dimensional quantities). Velocities u , v , and w in streamwise (x), wall-normal (y), and spanwise (z) direction are normalized by the freestream velocity u_∞^* . Furthermore, the respective freestream values are used as reference for density ρ , Temperature T , thermal conductivity θ , and viscosity μ . Non-dimensional pressure p , and time t are based on the reference values $\rho_\infty^* u_\infty^{*2}$, and L^*/u_∞^* .

2.1.1 Direct Numerical Simulation

The equations are applied in conservative formulation (ρ , ρu , ρv , ρw , E), with E being the total energy per unit volume. The fluid is assumed to follow the calorically perfect-gas assumption. Thus, a variation of heat capacities with temperature as well as chemical reactions are excluded. The adiabatic exponent and Prandtl number are fixed to $\kappa = 1.4$ and $Pr = 0.71$, respectively, for air. The viscosity is computed from Sutherland’s law. For more details see [1] and [2].

2.1.2 Bi-global Linear Stability Theory

The B-LST equations are based on the Navier–Stokes equations formulated in primitive variables (ρ , u , v , w , T). All flow quantities are split into a steady base flow and an unsteady perturbation. The assumptions of the (primary) linear stability theory hold complemented by the following specifics for the bi-global approach in crossplanes: A non-zero wall-normal base-flow velocity is allowed as long as its spanwise mean is zero, and the complex amplitude of the perturbation’s modal ansatz is two-dimensional. Analyzing flow crossplanes, the amplitude distribution extends in wall-normal and spanwise direction:

$$\phi(x, y, z, t) = \hat{\phi}(y, z) e^{i(\alpha x - \omega t)}. \quad (1)$$

In case of the temporal approach, used here for the computation of the eigenmodes, α is real and ω complex. In the more natural spatial approach α is complex and ω real. The imaginary part of the complex $\alpha = \alpha_r + i\alpha_i$, however, can be deduced with sufficient accuracy from the temporal theory using Gaster’s relation (see, e.g., [5]):

Table 1 Flow parameters

Case	Cold flow	Hot flow (from [10])
Re_L	10^5	10^5
Ma_∞	4.8	4.84
Pr_∞	0.71	0.67
κ_∞	1.4	1.36
L^* (m)	$8.5 \cdot 10^{-3}$	1.06
u_∞^* (m/s)	716.3	6,668
T_∞^* (K)	55.4	4,006
p_∞^* (Pa)	1,000	2,432
R_{x_r}	1,225	3,226
$Re_{\delta_{T_u}}$	$2 \cdot 10^4$	$2 \cdot 10^4$
Re_k	10^4	10^4
Re_{kk}	434	6,000
$\delta_{T_u}^*$ (m)	$1.7 \cdot 10^{-3}$	$2.2 \cdot 10^{-1}$
$k^*/\delta_{T_u}^*$	0.5	0.5
δ_μ^* (m)	$1.5 \cdot 10^{-3}$	$1.9 \cdot 10^{-1}$
T_w^*	Adiabatic	Radiative equilibrium ($T_{wrad}/T_{wad} \lesssim 0.1$)

$$\alpha_i = -\frac{\omega_i}{c_{gr}}, \quad (2)$$

where c_{gr} is the group velocity of the perturbation. This relation has been found to agree excellently with the spatial solution of the eigenvalue problem for the present flow cases (see [9]). Further details can be found in [7, 9, 10].

2.2 Numerics

2.2.1 Direct Numerical Simulation

For a detailed description of the basic algorithm applied to solve the unsteady Navier–Stokes equations see [1, 2]. Sixth-order compact finite differences (FDs) are applied for the spatial discretization in streamwise, wall-normal and spanwise direction. The time integration is based on the classical fourth-order four-step Runge–Kutta scheme. The equations are solved on a structured curvilinear grid. The implemented algorithm runs on various computer architectures using a combination of distributed- and shared-memory parallelization.

Various boundary conditions are implemented and can be prescribed according to the investigated flow. In spanwise direction periodicity is assumed with a domain width of $\lambda_z = 3.2$ (for L^* see Table 1). For steady supersonic (base) flow all flow quantities are fixed at the inflow plane. An adiabatic no-slip condition is prescribed at the wall.

In case of an unsteady DNS the velocity and temperature disturbances are fixed to zero at the wall; the underlying base flow is computed applying the adiabatic wall

condition. Defined disturbances can be excited via (synthetic) blowing and suction through slits or holes at the wall. For further details see [11].

2.2.2 Bi-global Linear Stability Theory

For the solution of the linear stability equations, base flow as well as perturbations are discretized on a structured curvilinear grid in the crossplane. The flow-quantity derivatives are modeled using compact FDs of up to tenth order in spanwise direction. A spectral Chebyshev collocation method is applied in wall-normal direction. Clustering grid points in regions of high shear away from the wall is possible using a grid transformation.

The temporal approach of the stability theory poses a linear eigenvalue problem (EVP). This is solved applying the implicitly restarted Arnoldi method [13]. The solution according to the spatial approach is obtained either by an iteration of the generalized approach, α complex and ω complex, as the temporal growth rate ω_i converges to zero or by Gaster's relation, described above. The excellent agreement of these methods for the cases investigated so far does not justify the additional computational cost of the former method for the minor gain in accuracy. The algorithm for eigenvalue tracking is based on a best-match approach for the eigenvectors at two consecutive tracking steps. The tracking process may become difficult due to eigenmodes crossing in the complex eigenvalue plane.

The perturbations are assumed to be periodic in spanwise direction. At the wall vanishing velocity and temperature perturbations are prescribed. In the freestream all perturbations are assumed to decay exponentially.

2.3 Flow Setup

2.3.1 Roughness Geometry

The investigated roughness geometry follows the setup of [10], see Fig. 2. The height of the roughness k equals about 50 % of δ_T of the unperturbed flow at the streamwise position of the element. The fence-like element ($e = 8k$ wide and $b = 2k$ thick) is placed obliquely, skewed by $\psi = 45^\circ$, with respect to the oncoming flow. The spanwise spacing between the roughness elements and, thus, the periodicity length is chosen to be four times the spanwise width of the element, $s_z = 4e$. The edges of the element are rounded due to numerical restrictions. The (minor) differences between sharp- and soft-edge elements have been discussed in [10].

2.3.2 Flow Conditions

The flow parameters for the cold-flow (wind-tunnel condition) case, and, for comparison, the hot-flow parameters from [10] are listed in Table 1. Both cases exhibit a Mach number of about 4.8. The roughness element has been placed in both




Fig. 2 Sketch of roughness setup. *Left*: top view with element dimensions b and e , spanwise spacing s_z , and rotation angle ψ . *Right*: front view looking downstream (Source: [10])

cases at about $Re_{\delta_T} = 20,000$ based on the undisturbed flat-plate flow to have this physically important parameter identical at first, see the further discussion below. This results in $R_{x_r} = 1,225$ and $Re_{kk} = 434$ for the cold-flow case, and $R_{x_r} = 3,226$ and $Re_{kk} = 6,000$ for the hot-flow case. The values of Re_{kk} are based on flow quantities of the undisturbed flat-plate flow at the roughness position and height:

$$Re_{kk} = \frac{\rho(x_r, k) u(x_r, k) k}{\mu(x_r, k)}. \quad (3)$$

Re_{kk} is much higher, at identical Re_{δ_T} , for the hot flow because, i.e., $\rho(x_r, k) / \rho_e(x_r)$ and $u(x_r, k) / u_e(x_r)$ are 3.2 and 1.4 times larger, and $\mu(x_r, k) / \mu_e(x_r)$ is 3.1 times lower compared to the cold flow. Thus stronger instability has to be expected if this parameter plays a crucial role.

The Reynolds number

$$Re_k = \frac{\rho_e(x_r) u_e(x_r) k}{\mu_e(x_r)} = 10^4 = Re_{\delta_T} \frac{k}{\delta_T} \quad (4)$$

holds for both flows. Note that according to experimental studies [12] the scaling of 3-d-roughness-induced transition follows a $Re_k = \text{constant}$ criterion like for 2-d roughness in incompressible flow. In this respect the cases compared here might be thought of being equally unstable. However, the hot-flow case implies a strongly cooled wall, $T_{w_{rad}} / T_{w_{ad}} \lesssim 0.1$ at $Ma = 4.8$, translating into a strongly (second-mode) unstable flow. In fact it turns out that, for the R_x value at the roughness considered, the real, uncontrolled flow would have already transitioned to turbulence due to second-mode instability. If, however, an ultrasonically absorptive coating would be used that weakens or suppresses second-mode instability, the roughness would lead to transition anyway because it induces vorticity (first) modes. Thus the laminar hot-flow case may not be realistic in any case but serves as a yet meaningful case for comparison.




Fig. 3 Top view of vortex structures. Visualization by means of the Q -criterion ($Q = 0.07$). HV, MV, and IV denote the horseshoe, main, and inner vortex at leading (L) and trailing (T) edge of the roughness element. Recirculation zones are shown by *black* isosurfaces of $u < 0$. *Gray bars* indicate position and extent of the roughness element. *Top*: color shading indicates streamwise vorticity ω_x ; *red*: clockwise rotation sense as seen in downstream direction; *green*: counter-clockwise rotation. *Bottom*: color shading indicates height y/k above the wall

3 Base Flow

The steady base flow for the B-LST analysis is extracted from the (D)NS base flow. In spite of the unsteady formulation of the Navier-Stokes equations the flow converges to a steady state. Convective exponential growth of numerical background noise (round-off error, 10^{-13}) can be seen downstream the roughness in a temporal Fourier analysis, but the amplitudes of the analyzed frequencies do not exceed a value of 10^{-9} at the end of the integration domain, and can therefore be neglected. For identical setup another base-flow solution had been computed with a different code, see [10]. This code has however not been designed for DNS and thus has a lower accuracy order. Differences in the resulting B-LST growth rates are discussed in Sect. 4.2.

A symmetric roughness configuration excites three pairs of equally strong counter-rotating vortices, with the main vortices being stronger than the horseshoe vortices (see [9]). The naming convention for the oblique setup follows [10]. Here the vortices induced at the leading edge (L) of the roughness element are significantly stronger than the trailing-edge (T) vortices. Thus, the leading-edge main vortex (LMV) becomes the dominant flow structure in the wake of the roughness. This is revealed by a vortex visualization using the Q -criterion, see Fig. 3.

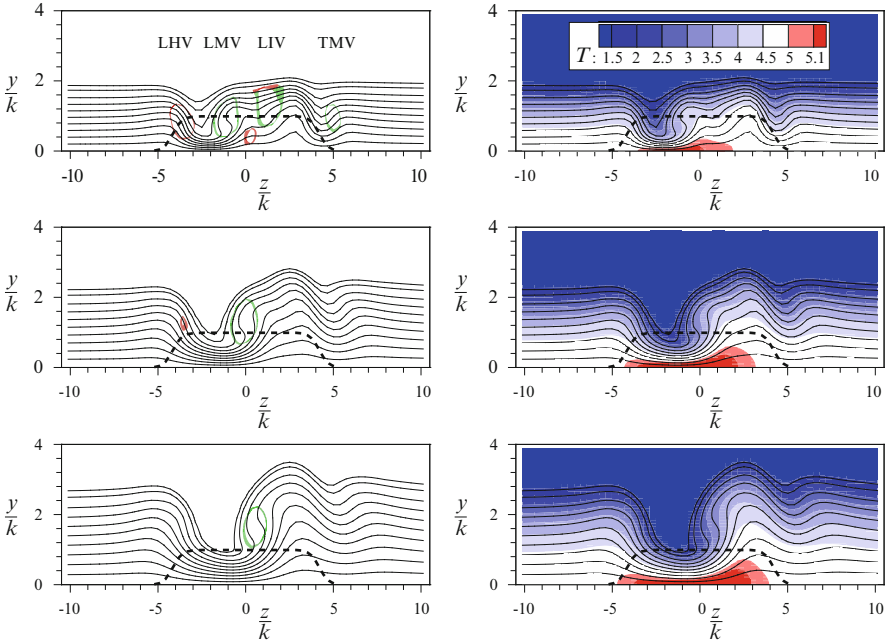


Fig. 4 Streamwise velocity (u) contours. *Solid lines* are isolines of u , starting with $u = 0.1$ near the wall, ending with $u = 0.95$. *Left*: color shading indicates vortices and rotation sense shown in Fig. 3 at $Q = 0.01$. *Right*: contours of temperature. *First row*: streamwise position $(x - x_r)/k = 30$, *second row*: $(x - x_r)/k = 100$, *third row*: $(x - x_r)/k = 200$. *Dashed lines* show a projection of the roughness geometry

The horseshoe vortex (LHV) is less pronounced than the LMV. Whereas the pair of inner vortices (IV) vanishes shortly behind a symmetric roughness [9], in the oblique setup the leading-edge inner vortex (LIV) is part of the formative vortex structure in the wake of the roughness element. The vicinity of LMV and LIV leads to interference between these co-rotating vortices, which in turn inhibits a soon crossflow-vortex-like turnover. The continuous lift-up of the vortices downstream of the roughness element can be observed in Fig. 3 (bottom). The vortices' positions above the wall can also be deduced from Fig. 4 (left).

The influence of the vortices on skin friction and wall temperature is shown in the flow's footprints in Fig. 5. The high-speed and the low-speed streaks can clearly be identified in terms of the $c_{f,n}$ values. The dominant leading-edge high-speed streak induces a pronounced streak of high wall temperature. The roughness element itself, however, is only slightly heated compared to the flat-plate temperature in front of the roughness and the main temperature rise downstream. The extent of the high-temperature streak in wall-normal direction can also be observed in Fig. 4 (right).

The flow around the roughness element can be visualized by means of streamlines, see Fig. 6. Near-wall streamlines in front of the roughness (Fig. 6a and b)




Fig. 5 *Top*: contours of skin friction coefficient $c_{f,n}$ normalized by smooth-wall value. *Bottom*: wall temperature T_w . *Black solid lines* indicate position and extent of the roughness element

are driven around the roughness element in spanwise direction. Coming around the leading edge a small fraction is lifted up by the recirculation zone in the element’s wake, gets accelerated by fluid flowing near the boundary-layer edge, and farther downstream joins the LIV. Figure 6c, d show a similar scenario of the development of LIV and LMV, though the streamlines no longer bend around the roughness element but flow across. The LMV mainly gathers streamlines that come from the height $y/k \approx 1$ in front of the roughness, whereas the LIV streamlines originate from lower heights. Figure 6e, f show the deformation of streamline layers being located above the element height upstream the roughness. The streamlines are deflected to the shape of the velocity streaks, roughly following the upper u -isolines of the crossplane at $(x - x_r) / k = 90$.

4 Stability Analysis

With the B-LST several unstable eigenmodes have been identified at $(x - x_r) / k = 30$, and afterwards been tracked in terms of the streamwise direction and varying frequency to span a stability diagram for each eigenmode. To avoid missing an eigenmode that is not yet amplified in the near wake of the roughness but becomes unstable farther downstream another search has been conducted at $(x - x_r) / k = 250$.




Fig. 6 Streamlines of the cold flow initiated in crossplane $(x - x_r)/k = -10$ at heights (a) $y/k = 0.1$, (b) $y/k = 0.5$, (c) $y/k = 0.75$, (d) $y/k = 1.0$, (e) $y/k = 1.5$ and (f) $y/k = 2.0$. Flow direction is from upper right to lower left corner. *Color shading* indicates the absolute value of the velocity vector $|\mathbf{v}| = \sqrt{u^2 + v^2 + w^2}$. *Black isosurfaces* show regions of separated flow ($u < 0$) enclosing the roughness element. The crossplane in the lower left corner displays u -isolines at $(x - x_r)/k = 90$

The spatial growth rates

$$\alpha_i = -\frac{d}{dx} \ln \frac{A(x)}{A_0} \quad (5)$$

are gained from the computed temporal ones

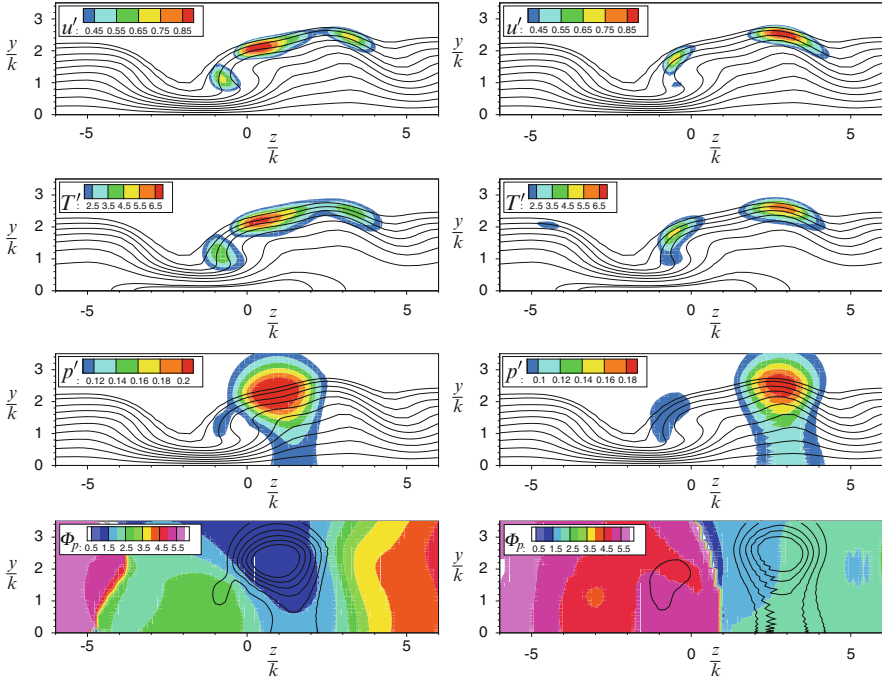


Fig. 7 Normalized modulus of the perturbation amplitude for streamwise velocity $u' = |\hat{u}(y, z)|$ (first row), temperature T' (second row) and pressure p' (third row), plus phase relation of pressure amplitude $\Phi_p = \arg[\hat{p}(y, z)]$ (fourth row) for the two most unstable eigenmodes at $(x - x_r) / k = 90$. *Left*: mode C1 at $\omega_r \delta_u \approx 0.80$. *Right*: mode C2 at $\omega_r \delta_u \approx 0.78$. *Solid lines* are base-flow u -isolines for first and third row, T -isolines for the second row, and pressure perturbation amplitude for the fourth row

$$\omega_i = \frac{1}{t} \ln \frac{A(t)}{A_0} \tag{6}$$

by application of Gaster’s relation, see Eq. (2). Based on the stability diagram the integral growth, quantified by the N -factor development, is computed for fixed frequencies from spatial growth rates extracted as a function of streamwise coordinate. The obtained data for the cold flow are compared to the hot-flow stability data from [11].

4.1 Unstable Eigenmodes

The eigenfunctions of the two most unstable eigenmodes found within the investigated streamwise range of $30 < (x - x_r) / k < 250$ are shown in Fig. 7. Amplitude

distributions of streamwise velocity, temperature and pressure are plotted. Additionally the phase relation of the pressure amplitude is displayed. The maxima of the perturbation amplitudes coincide with regions of large gradients within the base flow, which is shown in the plots for comparison. The shape of the temperature eigenfunctions is similar to the corresponding u -velocity amplitude distribution though the temperature perturbations reach considerably larger amplitudes. The pressure eigenfunctions feature non-negligible amplitudes at the wall.

The modes C1 and C2 exhibit a similarity to the even and odd eigenmodes found in the wake flow of symmetric roughness elements, see [7] and [9]. For the even mode the phase of the local eigenfunction maxima is symmetrical with respect to the symmetry line; the odd mode exhibits an anti-symmetric phase relation. The amplitude distributions of modes C1 and C2 are strongly distorted compared to the modes of the symmetrical flow, but the corresponding phase relations of the eigenfunctions' maxima reveal their even and odd nature, see Fig. 7 (fourth row) for the phase relation of the pressure perturbation amplitude.

4.2 Local and Integral Growth of Unstable Eigenmodes

Tracking the unstable eigenmodes along varying frequency at a fixed streamwise location yields the results shown in Fig. 8 (left). In the near wake at $(x - x_r)/k = 30$ the mode C1 is dominant. For comparison the hot-flow mode results from [11] are also shown: The hot-flow LIV mode has a similar band of amplified frequencies but an amplification rate that is roughly three times higher. The hot-flow LMV mode shows a considerably broader amplified-frequency band. Its maximum growth rate is only slightly lower than the LIV mode's and is located at a four times higher frequency. Note that the hot-flow modes' eigenfunctions significantly differ from the cold-flow modes due to major differences comparing the base flows. For details see [10] and [11]. Going downstream the amplified frequency bands of the modes widen, the corresponding maxima shifted to higher frequencies (not shown). The maximum growth rate of mode C1 persists compared to the near-wake value, whereas in the hot flow the amplification rate of the LMV mode increases considerably.

N -factors, $N = -\int \alpha_i dx$, (Fig. 8, right) are gained from a streamwise tracking of the eigenmodes. For clarity and with the exception of mode C1 only the integrally most amplified frequency of each mode is shown. The hot-flow results are again taken from [11]. Additionally, eigenmodes of the smooth flat-plate flow are shown for comparison. The cold-flow smooth-plate N -factors have been integrated starting shortly behind the edge of the flat plate at about $(x - x_r)/k = -125$ whereas the curves of the roughness eigenmodes inherently start downstream of the roughness where they first have been detected. For the hot-flow condition the corresponding smooth-plate N -factor has been integrated starting downstream of the roughness.




Fig. 8 *Left*: spatial growth rates $\alpha_i \delta_u$ as a function of frequency $\omega_r \delta_u$ (normalized by the undisturbed boundary-layer thickness δ_u at $x = x_r$) at streamwise position $(x - x_r) / k \approx 30$. Cold C1 mode (blue solid line with square symbols), cold C2 mode (blue dashed line, squares). Hot LMV mode (red solid line, circles), hot LIV mode (red dashed line, circles) from [11]. *Right*: N -factors over $(x - x_r) / k$. Cold C1 mode at $\omega_r \delta_u = 1.82$ (blue solid line, diamonds); cold C1 mode at $\omega_r \delta_u = 0.91$ (blue solid line, squares); cold C2 mode at $\omega_r \delta_u = 0.57$ (blue dashed line, squares). For comparison: cold smooth-plate second mode (propagation angle $\varphi = 0^\circ$) at $\omega_r \delta_u = 1.8$ (blue dash-dot line, squares); cold smooth-plate first mode ($\varphi \approx 65^\circ$) at $\omega_r \delta_u = 0.57$ (blue dash-dot line, filled squares). Hot LMV mode at $\omega_r \delta_u = 3.3$ (red solid line, circles); hot LIV mode at $\omega_r \delta_u = 0.82$ (red dashed line, circles); hot smooth-plate second mode at $\omega_r \delta_u = 2.1$ (red dash-dot line, circles) again from [11]. Thin black dash-dot-dot lines indicate $N = 5$ and $N = 9$

In the cold flow the eigenmode C1 gains the highest integral amplification to the end of the investigated range: $N = 6.5$ for $\omega_r \delta_u = 0.91$. Employing a base-flow solution based on a solver with less accuracy order, only $N = 5$ is reached (see [10]), despite the visual differences in the base flow are small. However, $\omega_r \delta_u = 1.82$ exhibits higher local growth for $(x - x_r) / k > 115$ matching the growth rate of mode C2 at the domain’s end; compare the slopes of the corresponding N -factor curves. The still increasing local growth of C2 at that location implies that this mode may take over in integral growth farther downstream as well. Nonetheless mode C1 at $\omega_r \delta_u = 0.91$ reaches an N -factor of about 5, which has been found sufficient for transition to turbulence in wind-tunnel experiments under noisy conditions, at about 180 roughness heights downstream the roughness’ center. The frequency $\omega_r \delta_u = 1.82$ follows at $(x - x_r) / k \approx 230$. The integrally most amplified cold-flow smooth-plate eigenmodes gain N -factors of about 3.6, and 2.9 for the first, and second mode, respectively, at the end of the investigated domain at $(x - x_r) / k = 250$.

The hot-flow modes are much stronger amplified. Within $50 < (x - x_r) / k < 150$ the hot-flow LMV mode grows more than twice as strong as the most amplified smooth-plate eigenmode which is a second mode. The strong wall-cooling ($T_{wrad} / T_{wad} \lesssim 0.1$) does not seem to significantly weaken the shear-layer modes found.

4.3 Superposition of Finite-Amplitude Eigenmodes and Base Flow

The shown eigenmode C1 has been assigned a finite amplitude and superposed on the respective base flow to get an impression of the three-dimensional disturbed flow. Note that the spanwise mean and the shape of the disturbances keep (unnaturally) since the nonlinear interaction of the finite disturbance and the base flow cannot be taken into account within the linearized approach.

The visualization is based on one single crossplane of the base flow and the parameter of this crossplane's eigenmode. The complex amplitude of the perturbation's modal ansatz (see Eq. (1)) defines the oscillation's physical amplitude as well as the relative phase relation of the eigenfunction in the crossplane. The streamwise wavenumber α_r and frequency ω_r determine the spatial wavelength and temporal development, respectively.

The base-flow crossplane at the respective streamwise location has been extruded in x -direction to four times the wavelength of the shown mode. At the beginning of the interval shown the eigenmode has been added to the base flow assuming an initial amplitude of 5% for the streamwise velocity disturbance. In streamwise direction the perturbation is allowed to grow with the spatial amplification rate α_i of the mode at the reference crossplane. The results are shown in Fig. 9.

The even disturbance mode C1 eventually leads to almost symmetrical horseshoe-type secondary vortices traveling on top of the inducing base-flow vortices (Fig. 9, left). The pressure footprint of the perturbed flow (Fig. 9, right) is broader than the respective generating horseshoe vortex. The u -isolines in the crossplanes shown also indicate the modulation of the flow by the developing finite-amplitude disturbances.

5 Direct Numerical Simulation with Controlled Unsteady-Disturbance Input

For the cold-flow scenario an unsteady DNS is conducted. The disturbances are introduced via blowing and suction through a hole at the wall by prescribing $(\rho v)'$ with a smooth function in time and space. For details see [11]. The hole has a diameter of $3.6k$ and is located upstream of the roughness at $(x - x_r)/k = -117$ at the element's streamwise center line ($z = 0$). The frequencies $\omega_r \delta_u = 0.91$ and $\omega_r \delta_u = 1.82$ are excited continuously at a level of $(\rho v)'_{max} = 10^{-7}$ each ($\theta_h = 0$), corresponding to a bi-harmonic point source. The DNS results are analyzed applying a temporal Fourier analysis. The amplitude distributions for the two excited frequencies are shown in Fig. 10 for one streamwise location. For comparison the corresponding B-LST eigenfunctions are also plotted. DNS and B-LST amplitude distributions agree well for $\omega_r \delta_u = 0.91$. The number of maxima is reproduced correctly by the LST results. However, their location slightly differs. For $\omega_r \delta_u = 1.82$ the agreement is good as well.

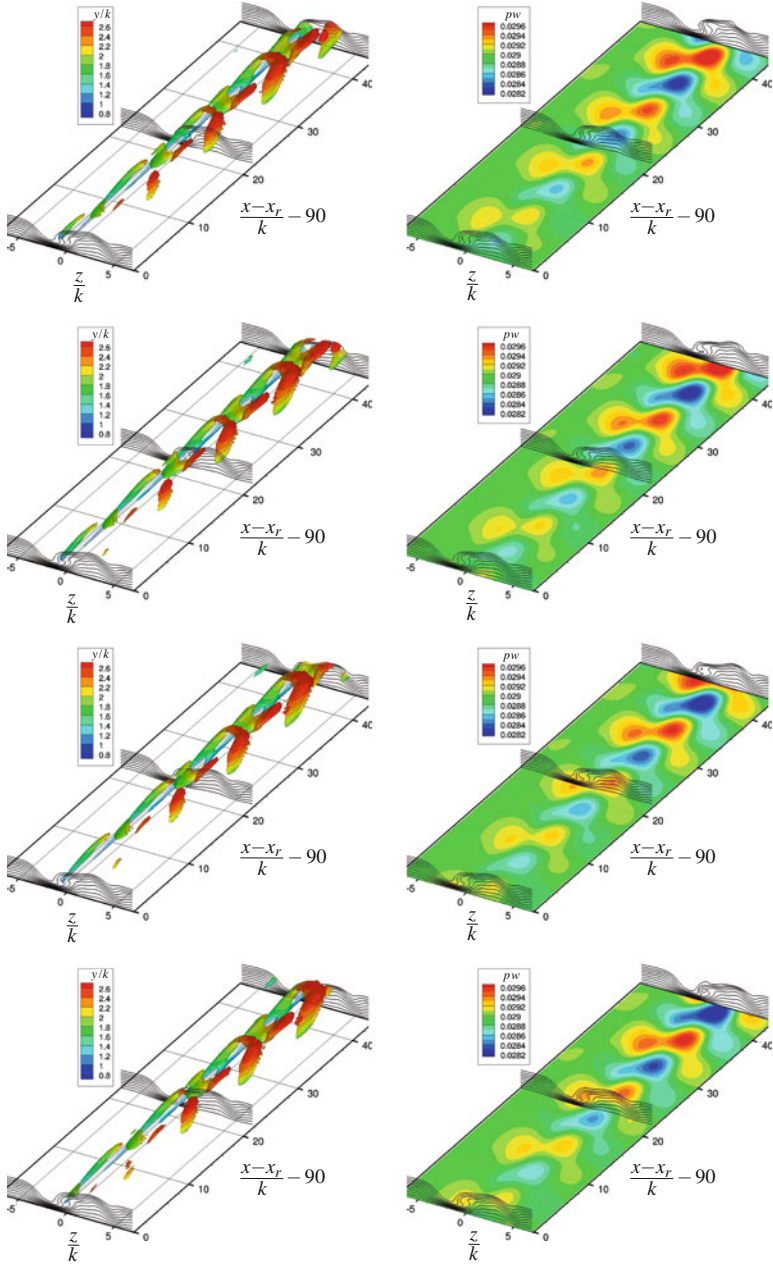


Fig. 9 Timesteps of animated eigenmode C1 at $(x - x_r)/k = 90$. *Left:* λ_2 -criterion ($\lambda_2 = -0.19$), *color shading* indicates the height above the wall. *Right:* pressure at the wall. *First row:* $t/T = 0$, *second row* $t/T = 0.2$, *third row* $t/T = 0.4$ and *fourth row* $t/T = 0.6$. Flow direction is from lower left to upper right corner. The *gray shading* indicates the unperturbed vortices of the steady base flow. Crossplanes display *u*-isolines

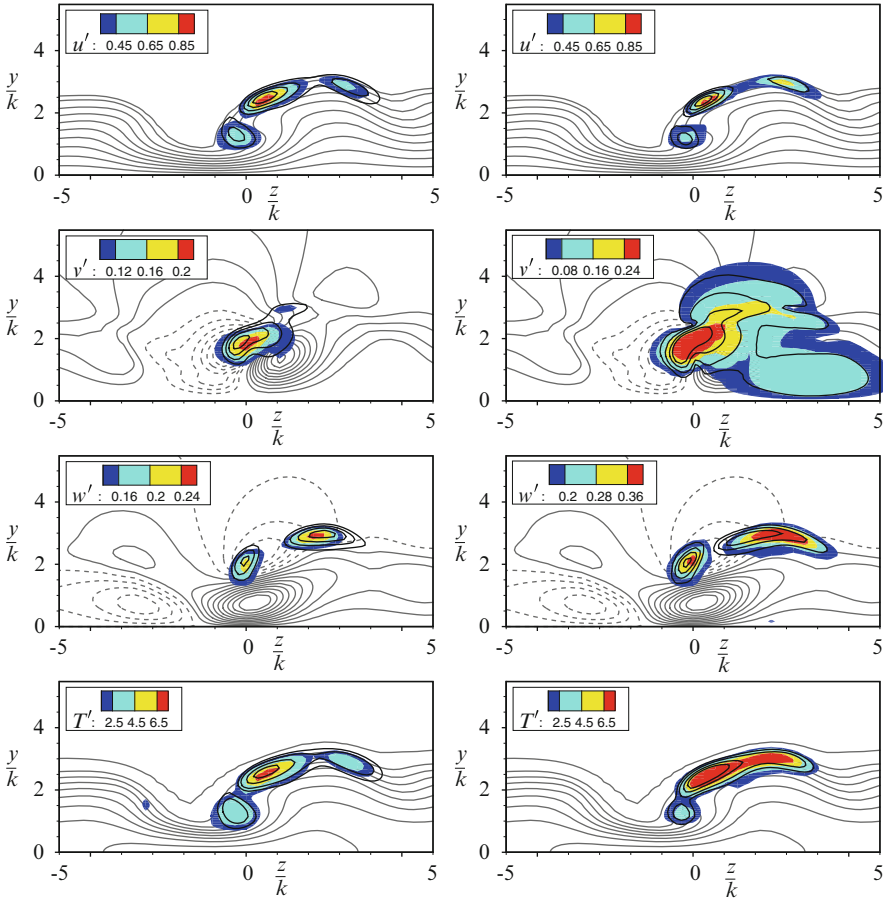


Fig. 10 Comparison of normalized disturbance amplitude distribution from DNS (contours) and LST (*thick black solid lines*) at $(x - x_r)/k = 150$ for frequencies $\omega_r \delta_u = 0.91$ (*left*) and 1.82 (*right*). *Thin solid lines* are corresponding base-flow isolines, *dashed* for negative values. *First row*: streamwise-velocity disturbance u' . *Second row*: wall-normal-velocity disturbance v' . *Third row*: spanwise-velocity disturbance w' . *Fourth row*: temperature disturbance T'

The u' -amplitude evolution with roughness is displayed in Fig. 11, compared to the smooth-plate case. Up to a short distance in front of the recirculation zone upstream of the roughness, $(x - x_r)/k \approx -15$, the disturbance evolution is identical. This zone acts like an amplifier for the disturbances (see the inset of Fig. 11). Marxen et al. [15] have found an analogous behavior for a two-dimensional roughness element under identical flow conditions. Across the roughness itself the disturbances are damped. Whereas the flow behind the two-dimensional element resumes soon the smooth-plate flow, here the flow is persistently deformed by the three-dimensional roughness element. The vortices and subsequent velocity streaks enhance the disturbance amplification significantly. The $(\omega_r \delta_u = 0.91)$ -




Fig. 11 Comparison of u -velocity-disturbance amplitude growth as function of $(x - x_r) / k$ from DNS (lines) and LST (symbols) for frequencies $\omega_r \delta_u = 0.91$ (blue long-dashed line and squares) and 1.82 (blue solid line and diamonds). Blue dash-dot line shows amplitude of mean-flow deformation ($\omega_r \delta_u = 0$), thin blue solid lines represent further nonlinear products of excited frequencies. Green lines indicate disturbance evolution in the smooth-plate boundary layer. The inset displays the region in the vicinity of the roughness for the two excited frequencies in detail, the vertical dot lines indicating the streamwise extent of the oblique element

disturbance grows by a factor of 6,240, corresponding to an N -factor of 8.7, between $-15 < (x - x_r) / k < 385$. The corresponding smooth-plate disturbance reaches an amplitude ratio of only 3.4 ($N = 1.2$) within the same streamwise range. The ($\omega_r \delta_u = 1.82$)-disturbance even reaches an N -factor of 10. The equivalent smooth-plate disturbance experiences a local amplification only. After decaying along the plate up to $(x - x_r) / k = 120$ it finally reaches its maximum amplitude at $(x - x_r) / k = 240$. At the end of the computational domain the disturbances are still linear: The amplitude of the mean flow deformation reaches only $4 \cdot 10^{-5}$ at $(x - x_r) / k = 385$.

To compare with B-LST the amplitudes have been matched at $(x - x_r) / k = 60$ for the frequency $\omega_r \delta_u = 0.91$, and at $(x - x_r) / k = 150$ for $\omega_r \delta_u = 1.82$. For the frequency $\omega_r \delta_u = 0.91$ the agreement is excellent, for $\omega_r \delta_u = 1.82$ the growth rates obtained by the B-LST analysis are too low.

6 Computational Aspects

6.1 Bi-global Linear Stability Theory

The B-LST analyses have been conducted on the NEC SX-9 machine at the High Performance Computing Center Stuttgart (HLRS) of the University of Stuttgart

using a single CPU. A parallel code version of the solver *BIGSTAB* has not been implemented satisfactorily so far. Solving a typical temporal eigenvalue problem in a crossplane of $N_y \times N_z = 61 \times 110$ takes about 1950s for 50 eigenvalues in the vicinity of a specified location in the complex plane. Time consumption of the iterative solver implemented in the *ARPACK* library is dependent on the number of computed eigenvalues, but mainly on the clustering of eigenvalues around the given location of interest. The average vector length is 248 at a vector operation ratio of about 99.6%; 58.5 GFLOPS are achieved. The maximum memory needed is roughly 60 GB. A further reduction of the memory requirements as well as the time consumption may be achieved by taking advantage of the sparse-matrix structure of the stability problem.

The eigenvalue tracking procedure compares the eigenvectors of subsequently solved eigenvalue problems. The additional cost of this comparison is negligible in relation to the EVP solution. Thus, the computational time for an eigenvalue tracking roughly equals the number of tracking steps multiplied by the time needed to solve one EVP. However, the tracking process requires slightly more memory for the eigenvector comparison routine.

Porting of the B-LST code to the new CRAY XE6 is pending, and inherently connected to its parallelization applying the *PARPACK* library as well as taking advantage of the sparse matrices.

6.2 Direct Numerical Simulation

6.2.1 General Performance Analysis for *NS3D* on CRAY XE6

The direct numerical simulations have been conducted on the CRAY XE6 at HLRS applying our *NS3D* code. The computational grid is decomposed into an arbitrary number of domains along streamwise and wall-normal direction. Each domain represents one MPI process. Shared-memory parallelization is applied in spanwise direction.

In the line of porting the *NS3D* code from the afore primarily used NEC SX-9 to the current CRAY XE6 system, explicit finite differences (EFDs) have been implemented in order to speed up the computation of derivatives for a large number of domains. This approach reduces the amount of necessary MPI communication significantly. The gain in performance (for a 2-d flow case) compared to the more accurate but costlier compact finite differences (CFDs) is shown in Fig. 12 for up to 4,096 MPI processes with domain decomposition along the streamwise direction. For the strong-scaling test using the CRAY XE6 installation step 0, 1.28 million grid points ($N_x \times N_y = 25,600 \times 50$) have been distributed to 16, 32, 64, 128, 256 and 512 MPI processes with domain sizes from $N_x \times N_y = 1,600 \times 50$ to $N_x \times N_y = 50 \times 50$ per MPI process. The superlinear speed-up using EFDs is due to cache effects. The CFDs' speed-up suffers from the huge amount of communication required for large numbers of MPI processes.

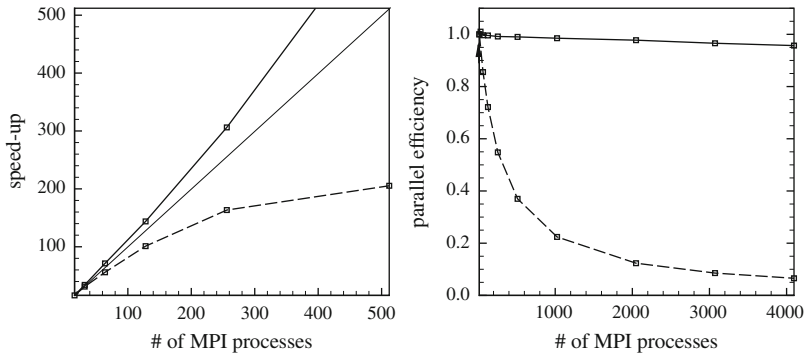


Fig. 12 Speed-up for strong scaling (*left*) and parallel efficiency for weak scaling (*right*) of the *NS3D* code (2-d flow case) on CRAY XE6 at HLRS as a function of the number of MPI processes. *Thick solid line*: eighth-order explicit finite differences (EFD-O8). *Dashed line*: sixth-order compact finite differences (CFD-O6). *Thin solid line* (*left*) represents ideal speed up

For the weak-scaling test using the CRAY XE6 installation step 1, the domain per MPI process has been fixed to $N_x \times N_y = 50 \times 50$ and the number of processes increased. The parallel efficiency of EFDs based computations persists at a high level, above 95 %, whereas the efficiency of CFDs based computations drops to about 6.6 % for 4,096 MPI processes, again suffering from the amount of inter-domain communication.

S. Andersson from CRAY Inc. extended the weak-scaling test for the EFDs to a larger number of cores, and a 3-d symmetrical flow case. This time each MPI process ($N_x \times N_y \times N_z = 20 \times 100 \times 33$) hosted 8 OpenMP threads. Andersson showed that for 14,000 MPI processes in streamwise direction (112,000 cores) *NS3D* yields a parallel efficiency of 94 % with respect to the performance of 100 MPI processes (800 cores).

6.2.2 Performance Data for Current Work Results

The current-work computations have been done applying the compact finite differences version of the *NS3D* code for the asymmetrical flow case. For a typical base flow computation the computational grid of $1,540 \times 150 \times 96$ points in streamwise (x), wall-normal (y) and spanwise (z) direction has been decomposed into 154 domains of size $N_x \times N_y \times N_z = 70 \times 30 \times 96$. On each node one MPI process has been started representing one sub domain. Thirty-two OpenMP threads per node accounted for spanwise parallelization resulting in a total of 4,928 CPUs. This configuration achieved $45.1 \mu s$ per grid point and timestep. For simulations with unsteady disturbance input the size of the computational domain has been reduced to *no. of domains* $\times N_x \times N_y \times N_z = 105 \times 70 \times 30 \times 96$ which yields a total of 3,360 CPUs. This reduction increased the code’s performance to $35.9 \mu s$ per grid point and timestep. If the configuration is chosen according to the one applied by

Stefan Andersson – four MPI processes reside on one node, each using 8 OpenMP threads – the specific time per grid point and timestep can be reduced to $26.6 \mu\text{s}$.

Similar asymmetrical cases (*no. of domains* $\times N_x \times N_y \times N_z = 8 \times 110 \times 150 \times 64$) executed on the NEC SX-9 achieved between 0.95 and $1.5 \mu\text{s}$ per grid point and timestep in single-node mode (8 CPUs with 11.2 GFLOPS each) and multi-node mode (2 nodes, each of the 8 MPI processes hosting 4 Microtasking threads, 7.3 GFLOPS per CPU), respectively. A single-node symmetrical case (*no. of domains* $\times N_x \times N_y \times N_z = 8 \times 110 \times 300 \times 33$) achieved $0.66 \mu\text{s}$ per grid point and timestep with 16.1 GFLOPS per CPU.

7 Conclusions

A laminar hypersonic flat-plate boundary-layer flow altered by a discrete oblique surface roughness element with a height of half the thermal boundary-layer thickness has been analyzed regarding its stability properties. Cold wind-tunnel conditions have been considered and compared to recent hot-flow results from [10] and [11]. In the cold flow with an adiabatic wall the roughness' leading-edge main vortex is the dominant flow structure in the wake of the roughness element. However, a soon crossflow-vortex-like turnover is inhibited by the adjacent inner vortex. The near-wake of the oblique roughness element exhibits stronger gradients than for a symmetric roughness configuration, cf. [7, 8]. In the hot flow with a strongly cooled wall ($T_{w,rad}/T_{w,ad} \lesssim 0.1$), these two vortices gain a larger spanwise distance resulting in a separate development. Due to the cool wall the hot-flow case exhibits larger gradients near the wall.

The instability modes found in the cold-flow wake show a lower amplification rate compared to the unstable hot-flow modes from [11]. This is in accordance to the cold flow's much lower Re_{kk} value of 434, being roughly 1/14th of the hot-flow value ($Re_{kk} = 6,000$) at equal $Re_k = 10^4$. Hence Re_k is not a similarity parameter. The integrally most amplified cold-flow mode gains $N = 6.5$ at about 250 roughness heights downstream the roughness' center compared to the N -factor of 9 at about 110 roughness heights for the most unstable hot-flow mode. The high-order accurate base-flow solver leads to larger N -factors: At the farthest downstream position considered $N = 6.5$ instead of 5 with a lower-order solver. Furthermore, the band of unstable frequencies is narrower for the cold-flow mode, its growth rate maximum shifted to lower non-dimensional values. We note that the hot-flow case has been set up for comparison at identical values of Re_{δ_T} and k/δ_T , but is fundamentally more unstable due to the strong wall cooling. The cold wall does not seem to stabilize the roughness induced instability though its character is of vorticity (first-mode) type rather than acoustic (second-mode) type.

The separated boundary layer in front of the roughness and the coinciding local adverse pressure gradient act as a local disturbance amplifier. Along the roughness, disturbances are strongly damped due to flow acceleration. Both facts are in accordance to [15]. In contrast to the flow behind the two-dimensional element

of [15] the stability properties of the flow behind the three-dimensional roughness do not resume the smooth-plate values. The induced longitudinal vortices and the subsequently generated persistent velocity streaks imprint three-dimensionality on the flow.

Acknowledgements The provided supercomputing time and technical support of the High Performance Computing Center Stuttgart (HLRS) of the University of Stuttgart within the project *LAMTUR* is gratefully acknowledged. The authors thank Michael Keller (IAG, Universität Stuttgart) and Stefan Andersson (CRAY Inc.) for providing the *NS3D* performance analysis data.

References

1. Babucke A (2009) Direct Numerical Simulation of Noise-Generation Mechanisms in the Mixing Layer of a Jet. Dissertation, Institut für Aerodynamik und Gasdynamik, Universität Stuttgart.
2. Babucke A, Linn J, Kloker M J, Rist U (2006) Direct numerical simulation of shear flow phenomena on parallel vector computers. In: High performance computing on vector systems: Proceedings of the High Performance Computing Center Stuttgart 2005, pp. 229–247. Springer Verlag Berlin.
3. Bartkowicz M D, Subbareddy P K, Candler G V (2010) Numerical Simulations of Roughness Induced Instability in the Purdue Mach 6 Wind Tunnel. Fourtyth AIAA Fluid Dynamics Conference & Exhibit, June 28 – July 1, Chicago, Illinois. AIAA Paper 2010–4723.
4. Birrer M, Stemmer C, Groskopf G, Kloker M J (2008) Chemical nonequilibrium effects in the wake of a boundary layer-sized object in hypersonic flows. In: Proceedings of the 2008 Summer Program, CTR, Stanford, Calif., July 6–August 1.
5. Bonfigli G, Kloker M J (2007) Secondary instability of crossflow vortices: validation of the stability theory by direct numerical simulation. *J Fluid Mech* 583: 229–272.
6. Casper K M, Wheaton B M, Johnson H B, Schneider S P (2008) Effect of Freestream Noise on Roughness-Induced Transition at Mach 6. Thirty-eighth AIAA Fluid Dynamics Conference & Exhibit, 23–26 June, Seattle, Wash. AIAA Paper 2008–4291.
7. Groskopf G, Kloker M J, Marxen O (2008) Bi-global secondary stability theory for high-speed boundary-layer flows. In: Proceedings of the 2008 Summer Program, CTR, Stanford, Calif., July 6–August 1.
8. Groskopf G, Kloker M J, Marxen O (2010) Bi-global crossplane stability analysis of high-speed boundary-layer flows with discrete roughness. In: Proceedings of the Seventh IUTAM Symposium on Laminar-Turbulent Transition. Stockholm, Sweden. 2009.
9. Groskopf G, Kloker M J, Stephani K A, Marxen O, Iaccarino G (2010) Hypersonic flows with discrete oblique surface roughness and their stability properties. In: Proceedings of the 2010 Summer Program, CTR, Stanford, Calif., June 27–July 23.
10. Groskopf G, Kloker M J, Stephani K A (2011) Temperature / rarefaction effects in hypersonic boundary-layer flow with an oblique roughness element. Forty-first AIAA Fluid Dynamics Conference and Exhibit, 27–30 June, Honolulu, Hawaii, USA. AIAA Paper 2011–3251.
11. Groskopf G, Kloker M J (2012) Stability analysis of three-dimensional hypersonic boundary-layer flows with discrete surface roughness. RTO-MP-AVT-200, Paper 30.
12. Horvath T J, Berry S A, Merski N R (2004) Hypersonic Boundary/Shear Layer Transition for Blunt to Slender Configurations – A NASA Langley Experimental Perspective. RTO-RSM-AVT-111, Paper 22.
13. Lehoucq R B, Sorensen D C, Yang C (1998) ARPACK User's Guide. SIAM, Philadelphia, Penn.

14. Li F, Choudhari M M, Chang C-L, Edwards J R (2010) Computational Analysis for Roughness-Based Transition Control. In: Proceedings of the Seventh IUTAM Symposium on Laminar-Turbulent Transition. Stockholm, Sweden. 2009.
15. Marxen O, Iaccarino G, Shaqfeh E S G (2010) Disturbance evolution in a Mach 4.8 boundary layer with two-dimensional roughness induced separation and shock. *J Fluid Mech* 648: 435–469.
16. Stephani K, Goldstein D B, Varghese P L (2010) Effects of Rarefaction on Hypersonic Boundary Layer Flow over Discrete Surface Roughness. Forty-eighth AIAA Aerospace Science Meeting Including the New Horizons Forum and Aerospace Exposition., 4–7 January, Orlando, Florida. AIAA Paper 2010–456.

Effect of Wall Roughness Seen by Particles in Turbulent Channel and Pipe Flows

Michael Breuer and Michael Alletto

Abstract In the present contribution it is shown in the context of large-eddy simulation (LES) and a Lagrangian treatment of the disperse phase that it is possible to considerably improve the particle statistics in turbulent channel and pipe flows by adopting a recently published wall roughness model for the solid phase. First, the model presented by Breuer et al. (Int J Multiphase Flow 43:157–175, 2012) is evaluated by means of the experiments conducted for a turbulent channel flow by Kussin (Experimentelle Studien zur Partikelbewegung und Turbulenzmodifikation in einem horizontalen Kanal bei unterschiedlichen Wandrauigkeiten. Ph.D. thesis, Martin–Luther–Universität Halle–Wittenberg, Germany, 2004) and Kussin and Sommerfeld (Exp Fluids 33:143–159, 2002). As a second test case the experiments of Borèe and Caraman (Phys Fluids 17:055108–1–055108–9, 2005) carried out for a turbulent pipe flow were used. For both setups involving rough walls good agreement between experiment and simulation is achieved by considering the effect of the wall roughness on the particle motion. Especially the latter configuration is the precondition for an improved simulation of the complex particle–laden turbulent flow in a combustion chamber reported in the last issue (Breuer and Alletto, High Performance Computing in Science and Engineering '11. Springer, Berlin/Heidelberg, 2012) and in Alletto and Breuer (Int J Multiphase Flow 45: 70–90, 2012) and Breuer and Alletto (Int J Heat Fluid Flow 35:2–12, 2012).

M. Breuer (✉) · M. Alletto

Professur für Strömungsmechanik, Helmut–Schmidt–Universität Hamburg, Holstenhofweg 85,
Postfach 70 08 22, D-22043, Hamburg, Germany

e-mail: breuer@hsu-hh.de; alletto@hsu-hh.de

1 Introduction

Owing to complex phenomena in technically relevant configurations involving particle-laden turbulent flows, e.g., massive recirculation regions in combustion chambers with bluff-body stabilized flames [5] or, e.g., the variety of secondary flows (e.g., precessing vortex cores, recirculation regions) in a cyclone separator [27], it is indispensable to have (i) adequate numerical tools able to reliably predict such complicated flows and (ii) to generate appropriate inflow conditions which closely resemble the real conditions at the domain boundary. Condition (i) can be achieved by means of LES as already shown by a variety of investigations [6, 7]. The effort of the research study carried out recently was to fulfill condition (ii) for particle-laden inflow conditions at high mass loadings.

Former investigations of Breuer and Alletto [8–10] and Alletto and Breuer [1] have shown that the generation of inflow conditions by means of a particle-laden turbulent pipe flow assuming specular reflection of the disperse phase at the walls, led, especially for high mass loadings, to discrepancies between the measurements closely downstream of the entrance of the combustion chamber studied. The mean particle velocities were found to be faster and the particle velocity fluctuations lower than observed in the experiments of Borée et al. [5]. As shown by previous investigations [3, 22, 32, 34, 38] modeling the effect of the wall roughness can correct the aforementioned discrepancies between simulation and experiment. Furthermore, the experiments of Borée et al. [5] showed a considerably flatter mean fluid velocity profile for the high mass loading $\eta = 110\%$ than the unladen flow. Vreman [38] reported in his DNS of a four-way coupled turbulent pipe flow that this effect can be reproduced by the interaction between particles hitting a rough wall and the continuous phase. Hence, the effort of our recent research was to develop a model which mimics the rebound behavior of solid particles at rough walls. In contrast to previous models found in the literature [18, 33, 37, 38] the sandgrain roughness model recently published in Breuer et al. [11] has the big advantage to establish an explicit relation between the carpet of densely packed spheres modeling the rough asperities and commonly used roughness parameters such as the mean roughness height R_z or the root-mean-squared roughness R_q . Kussin and Sommerfeld [21, 22, 35] reported the only investigations known to the authors which explicitly investigated experimentally the roughness effects on turbulent particle-laden flows with a clear specification of the mean roughness R_z of the bounding walls. In order to validate the new model a variety of tests were performed in a turbulent channel flow (see [11]). The results showed good agreement between the simulated particle statistics and the experiments [21, 22, 35].

In the following we briefly summarize the results obtained for the channel flow. Furthermore, the results achieved by adopting the new model to a turbulent particle-laden pipe flow are shown and compared with the experiments of Borée and Caraman [4]. This study are preliminary investigations in order to generate better inflow conditions for the time-consuming simulations in a model combustion chamber already described in the previous report by Breuer and Alletto [10] and additionally in [1, 9].

2 Governing Equations

In this work the multiphase flow is described using an Euler–Lagrange approach in which the two different phases are solved in two different frames of reference.

The continuous phase is solved in an Eulerian frame of reference. The conservation equations of the filtered quantities used in LES [7] can be extended to take into account the feedback effect of the particles on the fluid (two–way coupling). For that purpose the particle–source–in–cell method described in Crowe et al. [15] is used.

The dispersed phase is solved in a Lagrangian frame of reference. The equation of motion is given by Newton’s second law, where the fluid forces are derived from the displacement of a small rigid sphere in a non–uniform flow [25]. For particles with a density much higher than the carrier fluid, i.e. $\rho_p/\rho_f \gg 1$, only the drag, lift, gravity and buoyancy forces have to be considered, leading to:

$$\frac{d\mathbf{u}_p}{dt} = \frac{\mathbf{u}_f - \mathbf{u}_p}{\tau_p/\alpha} + \mathbf{g} \left(1 - \frac{\rho_f}{\rho_p} \right) + \frac{\mathbf{F}_L^{McL}}{m_p} + \frac{\mathbf{F}_L^{Mag}}{m_p} \quad (1)$$

$$\frac{d\boldsymbol{\omega}_p}{dt} = -\frac{10}{3\tau_p} \boldsymbol{\omega}_p. \quad (2)$$

\mathbf{u}_p , \mathbf{u}_f , τ_p , \mathbf{g} , m_p and $\boldsymbol{\omega}_p$ are the particle velocity, the fluid velocity at the particle position, the particle relaxation time, the gravitational acceleration, the mass of the particle and the particle angular velocity. The lift force on a particle due to the velocity shear (Saffman force) \mathbf{F}_L^{McL} is calculated as follows [26]:

$$\mathbf{F}_L^{McL} = \frac{9}{4\pi} \mu_f d_p^2 (\mathbf{u}_f - \mathbf{u}_p) \text{sign}(G) \left[\frac{|G|}{v_f} \right]^{1/2} J^u, J^u = f \left(\text{sign}(G) \frac{[|G|v_f]^{1/2}}{(\mathbf{u}_p - \mathbf{u}_f)} \right) \quad (3)$$

G is the fluid velocity gradient tensor. The lift force on a particle due to rotation (Magnus force) \mathbf{F}_L^{Mag} is calculated as follows [16]:

$$\mathbf{F}_L^{Mag} = \frac{\rho_f \pi}{2} d_p^2 C_{LR} |\mathbf{u}_f - \mathbf{u}_p| \frac{\boldsymbol{\Omega}_{rel} \times (\mathbf{u}_f - \mathbf{u}_p)}{|\boldsymbol{\Omega}_{rel}|} \quad (4)$$

The relative rotation of the particles $\boldsymbol{\Omega}_{rel}$ is given by:

$$\boldsymbol{\Omega}_{rel} = \frac{1}{2} \nabla \times \mathbf{u}_f - \boldsymbol{\omega}_p. \quad (5)$$

The lift coefficient C_{LR} is based on an experimentally determined correlation for $Re_p < 140$ [28]:

$$C_{LR} = 0.45 + \left(\frac{Re_r}{Re_p} - 0.45 \right) \exp(-0.05684 \cdot Re_r^{0.4} Re_p^{0.3}) \quad (6)$$

The Reynolds number of particle rotation Re_r and the particle Reynolds number Re_p are defined as:

$$Re_r = \frac{\rho_f d_p^2 |\boldsymbol{\Omega}_{rel}|}{\mu_f}, \quad Re_p = \frac{\rho_f |\mathbf{u}_f - \mathbf{u}_p| d_p}{\mu_f}. \quad (7)$$

Here d_p denotes the particle diameter. The drag force on the particle is based on Stokes flow around a sphere, where the corresponding drag coefficient is given by $C_D = 24\alpha/Re_p$ with the correction factor $\alpha = 1 + 0.15 Re_p^{0.687}$ to extent the validity towards higher Re_p .

For tiny particles with a relaxation time of the same order as the smallest fluid time scales the unresolved scales in LES become important for the particle motion. To consider the effect of the subgrid scales, a simple stochastic model by Pozorski and Apte [29] is applied. It requires the estimation of the subgrid-scale kinetic energy carried out with the help of the scale similarity approach of Bardina et al. [2]. To account for the rotation of the spherical particles around three Cartesian axes, Newton's second law for the angular momentum (2) is considered. The torque acting on a rotating spherical particle is determined based on the formulation of Rubiñow and Keller [30].

3 Numerical Methods

The continuous phase is solved in an Eulerian frame of reference using the computer code *LESOC* (= **L**arge-**E**ddy **S**imulation **O**n **C**urvilinear **C**oordinates [6, 7, 12]) to integrate the governing equations in space and time. It is based on a 3-D finite-volume method for arbitrary non-orthogonal and block-structured grids. The entire discretization is second-order accurate in space and time. For modeling the non-resolvable subgrid scales the well-known Smagorinsky model [31] with Van Driest damping near solid walls is applied. Alternatively, a dynamic model can be used.

The ordinary differential equation (1) for the particle motion is integrated by a fourth-order Runge-Kutta scheme. To avoid time-consuming search algorithms, the second integration of Eq. (1) to determine the particle position is done in the computational space. Here an explicit relation between the position of the particle and the cell index containing the particle exists [12, 14], which is required to calculate the fluid forces on the particle. Thus a highly efficient particle tracking scheme results allowing to predict the path of millions of particles.

The fluid velocity \mathbf{u}_f at the particle position is calculated with a Taylor series expansion around the cell center next to the particle [24]. This interpolation was shown to have a weaker filtering effect on the fluid velocity than a trilinear interpolation leading to better results for particles with small relaxation times τ_p . The set of three linear ordinary differential equations (2) for the particle angular velocity are solved analytically.

The collisions between particles for a four-way coupled simulation are determined deterministically by a recently developed collision method described in detail in Breuer and Alletto [9]. Based on the technique of uncoupling the calculation of particle trajectories is split into two stages. In the first stage particles are moved based on the equation of motion without inter-particle interactions. In the second stage the occurrence of collisions during the first stage is examined for all particles. If a collision is found, the velocities of the collision pair are replaced by the post-collision velocities without changing their position. The post-collision velocities are calculated by a hard sphere collision model involving friction between the colliding spheres (see e.g., [16,36]) where changes of the velocity and angular velocity of the two involved particles are modeled by the normal coefficient of restitution $e_{n,p}$, the tangential restitution coefficient $e_{t,p}$, the static $\mu_{st,p}$ and the dynamic coefficient of friction $\mu_{dy,p}$.

If the particle passes the center of the cell adjacent to the wall, a recently developed model (Breuer et al. [11]) to mimic the rebound behavior of the particle at rough walls is applied. The model takes into account the momentum loss of the particle during the wall impact by the wall-normal restitution coefficient $e_{n,w}$, the tangential restitution coefficient $e_{t,w}$, the static coefficient of friction $\mu_{st,w}$ and the dynamic coefficient of friction $\mu_{dy,w}$. Furthermore, the model considers the shadow effect leading to a redistribution of the streamwise momentum toward the wall-normal momentum. Based on geometric considerations relying on generally used roughness parameters such as R_z and R_q the local inclination of the wall is determined.

4 HPC Strategies

The major time-consuming computations carried out on the NEC SX-9 were spent to further improve the statistics of the flow in a model combustion chamber published in [1,9,10]. Concerning the HPC aspects the optimization of the solver for the continuous phase in *LESOCC* was already reported in several previous reports, see, e.g., [13,23]. Thus, here we restrict the discussion to the solver for the dispersed phase.

In order to track a huge number of particles, it is important to work with efficient algorithms which are applicable on high-performance computers. The present scheme is highly efficient due to the following reasons:

- No CPU time-consuming search algorithm is needed in the present c-space scheme.
- The particle properties are stored in linear arrays which allows vectorization of all loops in the particle routines over the total number of particles on the processor. If the number of particles is reasonably large (e.g., >256) the loops are efficiently carried out on the vector unit.

- Even if particles are leaving the present domain or deposit, the linear arrays are kept filled by reordering the particles after each time step. This guarantees optimal performance.
- The multi-block exchange between blocks for the particle data completely relies on the same arrangement as used for the continuous phase. The data transfer itself is based on MPI.
- Parallelization of the particle routines (and also of the flow solver) is achieved by domain decomposition, i.e., each processor deals with the particles of its own block. A minor disadvantage of this procedure is the fact that no load balancing of the particle tracking is possible, since the distribution of particles is not known in advance. However, since the tracking is so efficient that the predominant part of the CPU time is still spent for the continuous phase, this imbalance observed for the particle routines does not preponderate for the overall load balancing of the entire code.
- The collision detection procedure is carried out over a small amount of particles contained in a virtual cell which breaks down the computational cost from $O(N_p^2)$ to $O(N_p)$.
- As described in detail in [10] several conditions are introduced to further reduce the number of potential colliding particles step by step. At the end only a very few potential colliding particles remain, which have to be evaluated completely.
- Vectorization of the most time-consuming loop of the collision check routine is achieved by splitting up the loop. Additionally, it is ensured that the loop is reasonably large to be efficiently carried out on vector units.

Regarding more details about the specific topic of particle collisions, we refer to the previous report [10] and [9].

5 Description of the Test Cases

5.1 Horizontal Channel Flow

Kussin and Sommerfeld [22], Kussin [21], and Sommerfeld and Kussin [35] carried out detailed measurements based on phase-Doppler anemometry in a particle-laden horizontal fully developed channel flow. Three different Reynolds numbers were studied experimentally. In the present study a value of $Re = \delta_{Ch} U_B / \nu_f = 21,292$ was chosen ($Re_\tau = \delta_{Ch} u_\tau / \nu_f = 946$). δ_{Ch} denotes the channel half width, U_B the bulk velocity and u_τ the friction velocity. The gravitational acceleration g was made dimensionless by U_B and δ_{Ch} ($g^* = g \delta_{Ch} / U_B^2 = -4.42 \times 10^{-4}$) and pointed towards the bottom wall, i.e., in negative y -direction. The particles were spherical glass beads ($\rho_p = 2,500 \text{ kg/m}^3$) with mean diameters in the range of $60 \mu\text{m}$ to 1 mm subdivided into several classes of nominal size. In the present simulations mono-disperse particles of a nominal diameter of $d_p = 60, 100$ and $195 \mu\text{m}$ were considered characterizing a small size class and two medium size classes

($St = \rho_p d_p^2 U_B / (18 \mu_f \delta_{Ch} \alpha) = 18.8, 41.4$ and 100.4). Exchangeable stainless steel plates allowed to investigate the effect of different degrees of the wall roughness. From all cases studied in [21, 22] the cases R0, R1 and R2 corresponding to a mean roughness of $R_z = 2.32, 4.26$ and $6.83 \mu\text{m}$ were investigated here for the particle size of $d_p = 195 \mu\text{m}$ allowing a detailed evaluation of the behavior of the wall roughness model at a mass loading of $\eta = 30\%$.

Furthermore, the influence of inertia on the particle dynamics was analyzed by examining the behavior of particles with diameters of $d_p = 60, 100$ and $195 \mu\text{m}$ in the R2 channel configuration. The mass loading was set to $\eta = 10\%$ for the particles with a diameter of $d_p = 100$ and $195 \mu\text{m}$. For the $60 \mu\text{m}$ particles a mass loading of $\eta = 12\%$ was considered, since no measurements were available for $\eta = 10\%$.

The parameters modeling the velocity change during a flat wall impact were set to constant values for the cases chosen to evaluate the model: $e_{n,w} = 0.9$, $e_{t,w} = 0.3$, $\mu_{dy,w} = 0.4$, $\mu_{st,w} = 0.5$ and $C_{\text{Surface}} = 3$. Note that for the channel flow only frictionless collisions between the particles were considered with $e_{n,p} = 0.97$ and the Magnus lift force was neglected since it was not implemented at this stage of the development of the code.

The computational domain was $2\pi\delta_{Ch} \times 2\delta_{Ch} \times \pi\delta_{Ch}$ in streamwise, wall–normal and spanwise direction, respectively. The grid employed had $128 \times 150 \times 150$ cells and the dimensionless time step was $\Delta t^* = 0.004$. In streamwise and spanwise direction an equidistant grid was used. In wall–normal direction the grid was stretched geometrically with a stretching factor $r = 1.06$ and the first cell center was located at $\Delta y^+ = 0.8$ for the unladen flow allowing to apply Stokes no–slip boundary conditions. Furthermore, periodic boundary conditions were applied in streamwise and spanwise direction.

5.2 Vertical Pipe Flow

Borèe and Caraman [4] obtained two–component phase–Doppler anemometry measurements of a dilute poly–disperse two–phase flow at the exit of a long aluminum pipe. Unfortunately, no roughness specification was given for the pipe. The radius was $R_{\text{pipe}} = 10\text{mm}$, the bulk velocity was $U_{\text{pipe}} = 3.4\text{m/s}$, the centerline velocity was $U_c = 4\text{m/s}$ and the Reynolds number based on the bulk velocity and the pipe radius was $Re_{\text{pipe}} = 2,253$. The gravity pointed in flow direction. The initial diameter distribution of the spherical glass beads ranged from $d_p = 37$ to $116 \mu\text{m}$ (see Table 1) and had a density of $\rho_p = 2,470\text{kg/m}^3$. The values in Table 1 are taken from the discrete number distribution found in [4]. The bidispersed particle number distribution showed two peaks at $d_p \approx 60 \mu\text{m}$ and $d_p \approx 90 \mu\text{m}$. Borèe and Caraman [4] grouped particles with a diameter $55 \mu\text{m} \leq d_p \leq 65 \mu\text{m}$ to the $60 \mu\text{m}$ size class and particles with a diameter of $85 \mu\text{m} \leq d_p \leq 95 \mu\text{m}$ to the $90 \mu\text{m}$ size class. The particle statistics in the experiment were evaluated for this two size classes. Two different mass loadings of $\eta = 11$ and 110% were considered in the experiment.

Table 1 Initial distribution of the particle sizes by Borée and Caraman [4]

i	1	2	3	4	5	6	7	8	9	10
$d_{p,i}$ (μm)	36.9	42.7	55.8	61.1	67.0	74.1	87.8	97.1	107	116
Number distribution	0.019	0.044	0.120	0.107	0.077	0.093	0.219	0.186	0.102	0.033

In this case the parameters modeling the velocity and angular velocity changes during the wall impact were set to the following values: $e_{n,w} = 0.831$, $e_{t,w} = 0.31$, $\mu_{dy,w} = 0.125$ and $\mu_{st,w} = 0.6$. The parameters chosen for the hard sphere collision model were set to $e_{n,p} = 0.97$, $e_{t,p} = 0.44$, $\mu_{dy,p} = 0.092$ and $\mu_{st,p} = 0.94$. These values were taken from Foerster et al. [17] for the pairings glass–glass (particle–particle collisions) and aluminum–glass (particle–wall collisions). To evaluate the roughness model two equivalent sandgrain roughnesses $k_s = C_{\text{Surface}} R_z = 0.0$ and $15.0 \mu\text{m}$ and additionally a specular wall collision (i.e., only the wall–normal velocity after the wall impact was changed with $e_{n,w} = 1$ and without friction during the collision) were chosen. Note that the chosen roughness height lies in the range given in the literature for aluminum pipes ($k_s = 1.5 - 60 \mu\text{m}$ [19]). To evaluate the simulated particle statistics the particles with the size classes $i = 3 - 5$ (see Table 1) were compared with the measurements of the $60 \mu\text{m}$ size class found in [4] and particles with the classes $i = 7 - 8$ were compared with the $90 \mu\text{m}$ size class of the measurements [4].

The computational domain had a dimensionless radius $R/R_{\text{pipe}} = 1$ and an extension of $L/R_{\text{pipe}} = 24$ in streamwise direction for the low mass loading case and $L/R_{\text{pipe}} = 12R$ for the high mass loading. The longer computational domain for the low mass loading was chosen to ensure that the two–point correlations reached values close to zero at the half of the pipe length. The computational domain was discretized by an O–type grid with 2×10^6 cells for the low mass loading and 1×10^6 for the high mass loading. The first cell center was placed at $y^+ = 0.14$ for the unladen flow and thus no–slip boundary conditions were applied. The flow is again periodic in streamwise direction.

6 Results

6.1 Horizontal Channel Flow

In the following the behavior of the wall roughness model with increasing asperity heights is presented. For a detailed explanation of the governing mechanism in a particle–laden flow confined by rough walls we refer to Breuer et al. [11]. Here only a brief summary of the observations is provided.

The results obtained for R0, R1 and R2 are compared with the experiments [21]. Furthermore, the special case of a smooth wall is included. Figure 1a shows the mean streamwise particle velocity which is in good agreement with the

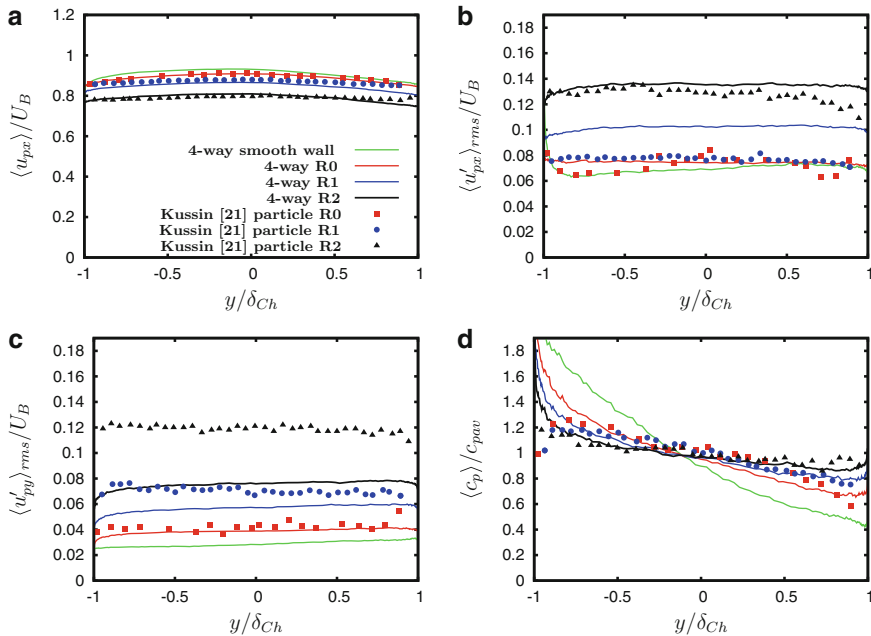


Fig. 1 Variation of the wall roughness: statistical results of the plane channel flow for the particles $d_p = 195 \mu\text{m}$, $\eta = 30\%$, $e_{n,w} = 0.9$, $\mu_{dy,w} = 0.4$: (a) mean velocity, (b) streamwise r.m.s. fluctuations, (c) wall–normal r.m.s. fluctuations, (d) normalized particle number concentration

measurements [21]. It is clearly evident that the mean particle velocity decreases with increasing roughness.

Figure 1b shows the streamwise particle velocity fluctuations. The r.m.s. fluctuations increase with increasing wall roughness. The simulated particle wall–normal velocity fluctuations illustrated in Fig. 1c show an almost linear increase with R_z . It is especially remarkable how the normalized particle number concentration (Fig. 1d) is influenced by the roughness: While a smooth channel wall leads to an accumulation of the particles at the bottom wall, the computed concentration for the roughness R2 delivers an almost uniform profile. Furthermore, there are consistent changes in the concentration profiles if the presented wall model is applied to the R0 roughness rather than assuming a smooth wall. Good agreement with the experiments is observed except in the near–wall regions where the computed profiles show a strong peak at the bottom wall. A possible explanation of this observation is that especially for particles reaching the wall with a small incident angle, the present wall model could lead to a non–zero probability of the rebounded particles to remain grazing, i.e., to rebound with a very small or zero wall–normal velocity (see Konan et al. [20]). Such particles may experience a multiple rebound with the rough wall [20] presently not included in the wall model.

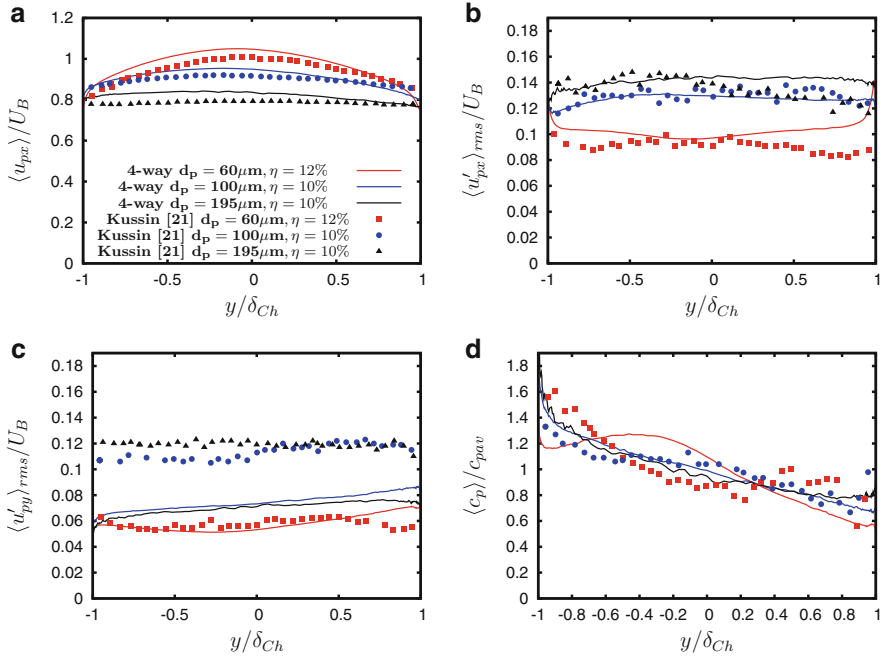


Fig. 2 Variation of d_p : statistical results of the plane channel flow for the particles at $\eta \approx 10\%$, $e_{n,w} = 0.9$, $\mu_{dy,w} = 0.4$, roughness R2: (a) mean velocity, (b) streamwise r.m.s. fluctuations, (c) wall-normal r.m.s. fluctuations, (d) normalized particle number concentration

Figure 2 depicts the particle statistics obtained for the R2 channel configuration varying the particle diameter d_p . The mass loading was set to $\eta = 10\%$ ($d_p = 100$ and $195 \mu\text{m}$) or $\eta = 12\%$ ($d_p = 60 \mu\text{m}$). The influence of inertia is clearly visible considering the mean particle velocity in Fig. 2a: Even though the $60 \mu\text{m}$ particles are much stronger influenced by the roughness structures in the near-wall region, they adjust much quicker to the fluid flow leading to a higher mean velocity compared with the $195 \mu\text{m}$ particles (see also [36]). Good agreement with the experiments of Kussin [21] is found for all cases. Figure 2b shows the particle streamwise velocity fluctuations which also reasonably agree with the experiments [21]. The wall-normal velocity fluctuations (Fig. 2c) computed for the small $60 \mu\text{m}$ particles are in close agreement with the experiments, whereas for the larger particles the fluctuations are underpredicted. Similar to the wall-normal velocity fluctuations no specific trend can be observed for the normalized particle concentration (Fig. 2d).

6.2 Vertical Pipe Flow

In the following results of a turbulent particle-laden pipe flow adopting the model presented in Breuer et al. [11] are compared with the experimental data of Borè and Caraman [4]. Additionally, the DNS data obtained by Vreman [38] in the same pipe are included for comparison. For the low mass loading Vreman [38] presented only simulations with mono-disperse 60 μm particles and thus only the unladen flow results were taken for comparison. For the high mass loading $\eta = 110\%$ Vreman [38] carried out mono-disperse simulations with 90 μm particles. The results achieved are not expected to differ substantially from the poly-disperse scenario (87% of the particle mass belongs to the 90 μm size class [4]) and hence can be compared with the present simulation. Note that the data of Vreman [38] are obtained with a simpler model for the influence of rough walls on the particle rebound. Unfortunately, Borè and Caraman [4] did not provide fluid measurements for the high mass loading and hence these could not be included.

Figure 3 shows the mean streamwise particle ($d_p = 60$ and 90 μm) and fluid velocity for the two mass loadings investigated for different wall roughnesses. The following observations can be made: Adopting the sandgrain wall roughness model leads to (i) a reduction of the mean particle velocity in the pipe center and a slight increase in the near-wall region, (ii) the 90 μm particles (see Fig. 3b, d) are stronger influenced by the wall roughness than the 60 μm particles (see Fig. 3a, c), (iii) deviations from the experiments can be seen for all particle statistics in the near-wall region (see Fig. 3a–d), (iv) the mean fluid flow is only affected in the high mass loading case showing similar to the DNS of Vreman [38] a reduction of the mean fluid velocity in the pipe center. The wall roughness leads to an increase of the particle concentration (not shown for the sake of brevity) in the pipe center and hence due to the four-way coupling to a force on the fluid acting against the flow direction. (Note that the particles in the pipe center are slower than the fluid.) (v) By adopting the sandgrain roughness model good accordance with the experiments [4] can be found except in the near-wall region.

Figure 4 shows the streamwise particle ($d_p = 60$ and 90 μm) and fluid velocity fluctuations for the two mass loadings investigated for different wall roughnesses. It is obvious that for all particle statistics an improvement (Fig. 4a–d) can be seen if the roughness model [11] is adopted. For the low mass loading (Fig. 4a, b) the model leads to an increase of the r.m.s. fluctuations in the pipe center, whereas they remain nearly unchanged close to the wall. Concerning the high mass loading (Fig. 4c, d) the model yields an increase of the r.m.s. fluctuations for the 60 μm particles, whereas for the r.m.s. fluctuations of the 90 μm particles only small changes throughout the pipe can be reported. Remarkable is the strong reduction of the fluid velocity fluctuations from the low mass loading (Fig. 4e) to the high mass loading (Fig. 4f). Note that in his DNS Vreman [38] made the same observation.

Figure 5 shows the wall-normal particle ($d_p = 60$ and 90 μm) and fluid velocity fluctuations for the two mass loadings investigated for different wall roughnesses. As for the other statistics shown, the sandgrain roughness model leads to a

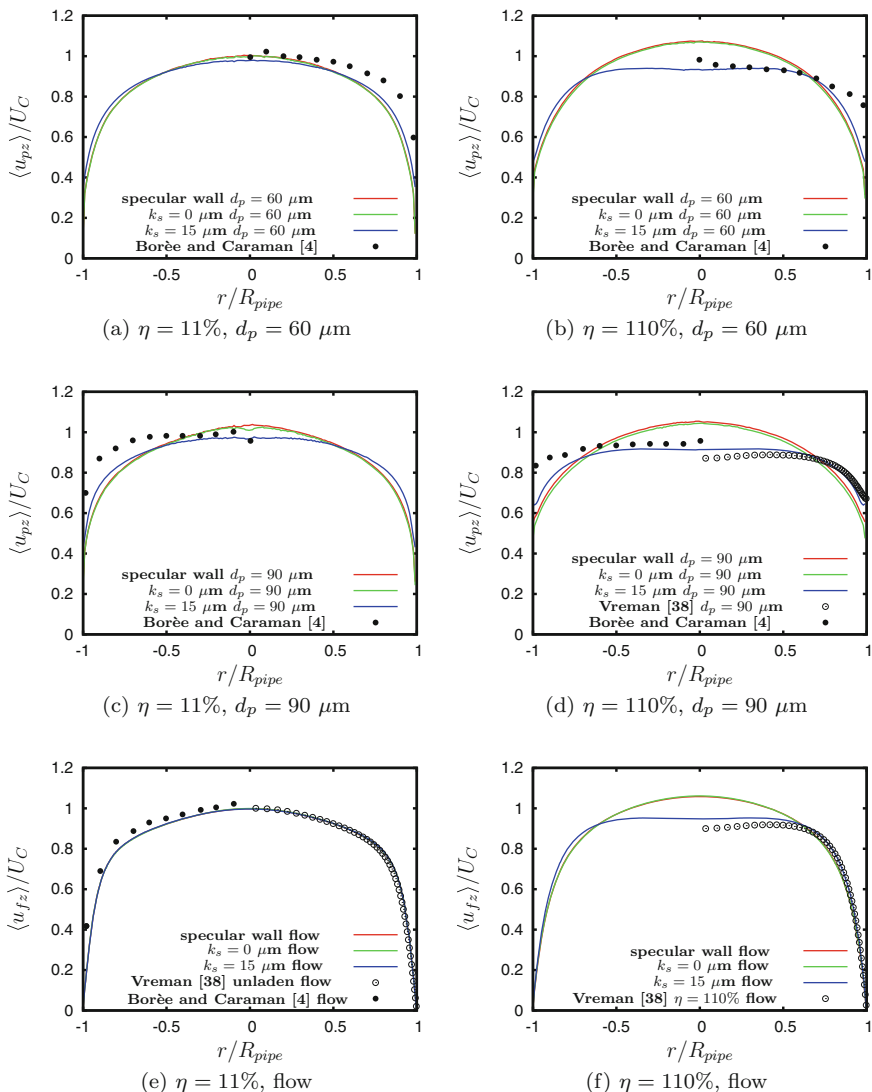


Fig. 3 Variation of the wall roughness: statistical results of the poly-disperse pipe flow; (a)–(d) mean particle velocity for two particle diameters and mass loadings; (e)–(f) mean fluid velocity

considerably better agreement with the experiments than the specular wall or the smooth wall boundary conditions. Remarkable is the strong enhancement of the wall-normal particle velocity fluctuations for both mass loadings and particle classes evaluated (Fig. 5a–d). This indicates that by means of the wall roughness the particles achieve a steeper trajectory in a similar manner as the particles rebounding at the rough walls in the channel flow (see Fig. 1c and also [11]). Analogous to

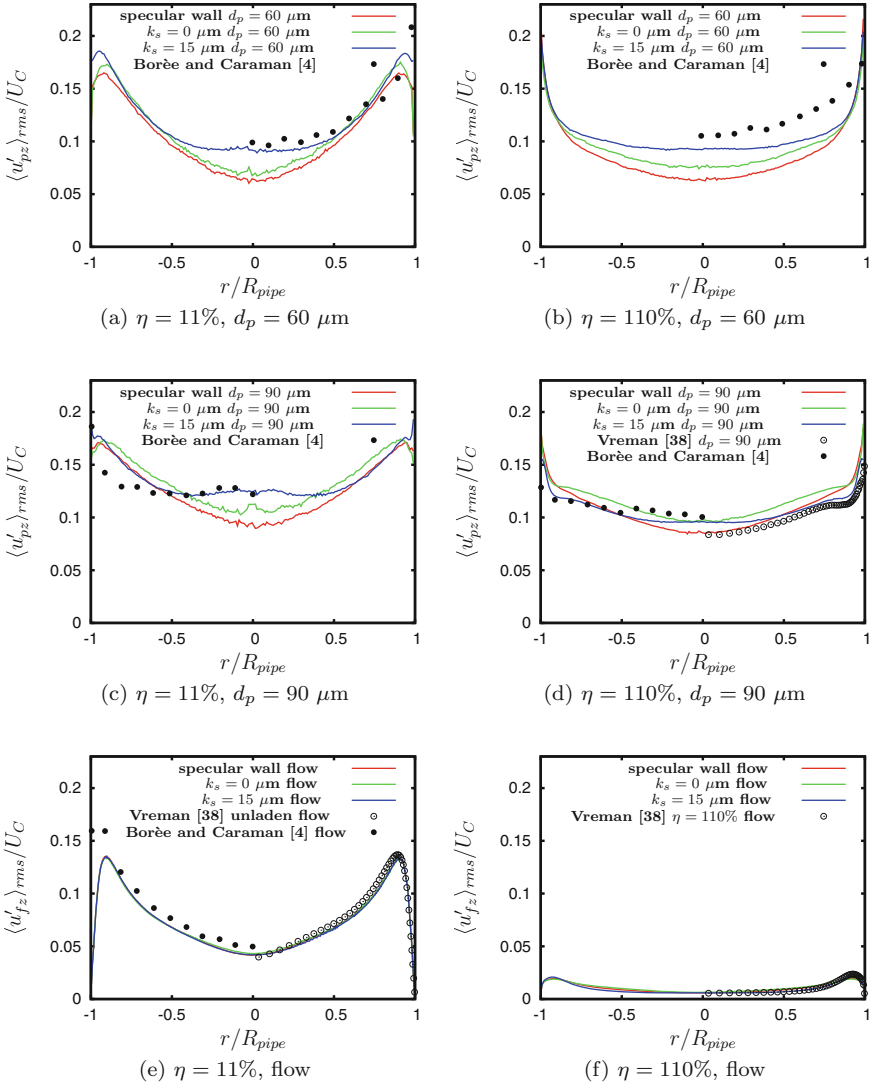


Fig. 4 Variation of the wall roughness: statistical results of the poly-disperse pipe flow; (a)–(d) particle streamwise r.m.s. fluctuations for two particle diameters and mass loadings; (e)–(f) streamwise fluid r.m.s. fluctuations

the streamwise fluid velocity fluctuations also in the wall-normal direction a strong reduction of the fluid velocity fluctuations from the low mass loading (Fig. 5e) to the high mass loading (Fig. 5f) can be observed. Note that the same observation was again made by Vreman [38].

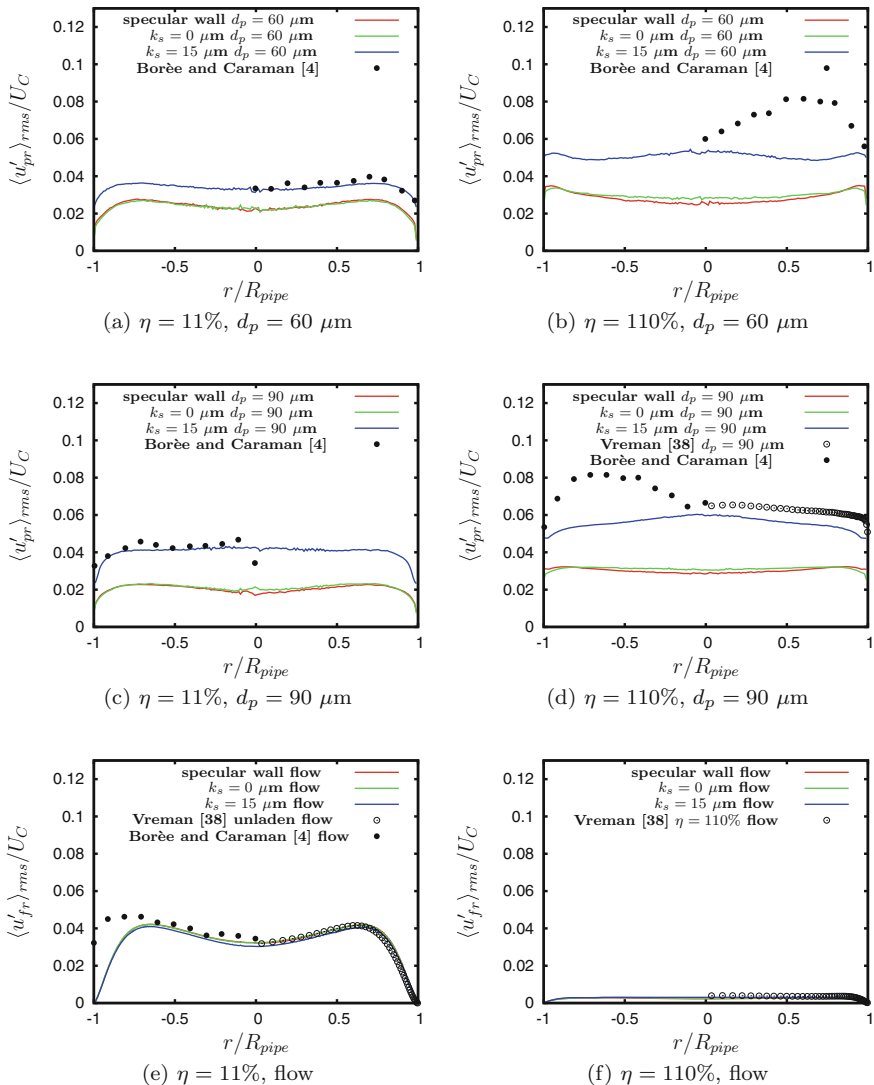


Fig. 5 Variation of the wall roughness: statistical results of the poly-disperse pipe flow; (a)–(d) particle wall-normal r.m.s. fluctuations for two particle diameters and mass loadings; (e)–(f) fluid wall-normal r.m.s. fluctuations

7 Conclusions

In view of the results obtained in Sect. 6 some analogies between the behavior of solid particles bounded by rough walls in a turbulent channel and pipe flow can be found:

- For both configurations the wall roughness leads to a reduction of the mean particle velocity with increasing roughness height which implicates an additional pressure drop due to the four-way coupling.
- An enhancement of the particle wall-normal fluctuations could be found indicating steeper particle trajectories and an additional momentum transfer between the streamwise and the wall-normal direction.
- Both effects were more pronounced for bigger particles than for smaller particles since big particles keep memory of the wall impact for a longer time than small particles.
- An improvement of the particle statistics in both configurations can be achieved by considering the rebound behavior of solid particles at rough walls applying the recently developed sandgrain roughness model.

Regarding the particle concentration differences can be found between the horizontal channel and the vertical pipe. In a channel flow the wall roughness leads to a homogenization of the particle concentration. (Note that gravity pointed towards the bottom wall.) In the pipe flow the wall roughness induces an accumulation of the particles in the pipe center (not shown for the sake of brevity) leading for the high mass loading to a reduction of the mean fluid velocity in this region by means of the four-way coupling.

In conclusion, with the new findings it is possible to substantially improve the inflow conditions of the combustion chamber flow described in the former report of Breuer and Alletto [10] and in [1, 9].

Additionally, for the chamber flow a new grid with a finer resolution is presently used to better resolve the complicated flow in the two recirculation regions extending behind the chamber entrance. Both measures are expected to considerably improve the results of the combustion chamber flow.

Acknowledgements The time-consuming computations were carried out on the national supercomputer NEC SX-9 at the High Performance Computing Center Stuttgart (grant no.: PARTICLE/pfs 12855), which is gratefully acknowledged.

References

1. M. Alletto and M. Breuer. One-way, two-way and four-way coupled LES predictions of a particle-laden turbulent flow at high mass loading downstream of a confined bluff body. *Int. J. Multiphase Flow*, 45:70–90, 2012.
2. J. Bardina, J. H. Ferziger, and W. C. Reynolds. Improved subgrid-scale models for large-eddy simulations. *AIAA Paper*, 80–1357, 1980.
3. M. Benson, T. Tanaka, and J. K. Eaton. Effects of wall roughness on particle velocities in a turbulent channel flow. *Trans. ASME J. Fluids Eng.*, 127:250–256, 2005.
4. J. Borée and N. Caraman. Dilute bidispersed tube flow: Role of interclass collisions at increased loadings. *Phys. Fluids*, 17:055108–1–055108–9, 2005.
5. J. Borée, T. Ishima, and I. Flour. The effects of mass loading and inter-particle collisions on the development of polydispersed two-phase flow downstream of a confined bluff body. *J. Fluid Mech.*, 443:129–165, 2001.

6. M. Breuer. Large-eddy simulation of the sub-critical flow past a circular cylinder: Numerical and modeling aspects. *Int. J. Numer. Meth. Fluids*, 28(9):1281–1302, 1998.
7. M. Breuer. *Direkte Numerische Simulation und Large-Eddy Simulation turbulenter Strömungen auf Hochleistungsrechnern*. Habilitationsschrift, Universität Erlangen-Nürnberg, Berichte aus der Strömungstechnik. Shaker Verlag, Aachen, 2002. ISBN 3-8265-9958-6.
8. M. Breuer and M. Alletto. Two-phase flow predictions of the turbulent flow in a combustion chamber including particle-particle interactions. In *J. Physics: Conference Series*, volume 318, page 052002. IOP Publishing, 2011. 13th European Turbulence Conference (ETC13), Warsaw, Poland, Sept., 12–15, 2011.
9. M. Breuer and M. Alletto. Efficient simulation of particle-laden turbulent flows with high mass loadings using LES. *Int. J. Heat Fluid Flow*, 35:2–12, 2012.
10. M. Breuer and M. Alletto. Numerical simulation of particle-laden turbulent flows using LES. In W. E. Nagel, D. B. Kröner, and M. M. Resch, editors, *High Performance Computing in Science and Engineering '11*, pages 337–352. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-23868-7.
11. M. Breuer, M. Alletto, and F. Langfeldt. Sandgrain roughness model for rough walls within Eulerian-Lagrangian predictions of turbulent flows. *Int. J. Multiphase Flow*, 43:157–175, 2012.
12. M. Breuer, H. T. Baytekin, and E. A. Matida. Prediction of aerosol deposition in 90 degrees bends using LES and an efficient Lagrangian tracking method. *J. Aerosol Science*, 37(11): 1407–1428, 2006.
13. M. Breuer, P. Lammers, T. Zeiser, G. Hager, and G. Wellein. Direct numerical simulation of turbulent flow over dimples - code optimization for nec sx-8 plus flow results. In W.E. Nagel, D. Kröner, and M. Resch, editors, *High Performance Computing in Science and Engineering '07, 10th Results and Review Workshop on High Performance Computing in Science and Engineering, Oct. 04–05, 2007*, pages 303–318. Springer, Berlin, 2008.
14. M. Breuer, E. A. Matida, and A. Delgado. Prediction of aerosol drug deposition using an Eulerian-Lagrangian method based on LES. In *Int. Conf. on Multiphase Flow, July 9–13, 2007*. Leipzig, Germany, 2007.
15. C. T. Crowe, M. P. Sharma, and D. E. Stock. The Particle-Source-In-Cell (PSI-CELL) model for gas-droplet flows. *Trans. ASME J. Fluids Eng.*, 99:325–322, 1977.
16. C. T. Crowe, M. Sommerfeld, and Y. Tsuji. *Multiphase Flows with Droplets and Particles*. CRC Press, 1998.
17. S. F. Foerster, M. Y. Louge, H. Chang, and K. Allia. Measurements of the collision properties of small spheres. *Phys. Fluids*, 6(3):1108–1115, 1994.
18. Th. Frank, K. P. Schade, and D. Petrak. Numerical simulation and experimental investigation of a gas-solid two-phase flow in a horizontal channel. *Int. J. Multiphase Flow*, 19:187–198, 1993.
19. I. E. Idelchik. *Handbook of Hydraulic Resistance*. Springer, second edition, 1986. ISBN 3-540-15962-2.
20. N. A. Konan, O. Kannengieser, and O. Simonin. Stochastic modeling of the multiple rebound effects for particle-rough wall collisions. *Int. J. Multiphase Flow*, 35:933–945, 2009.
21. J. Kussin. *Experimentelle Studien zur Partikelbewegung und Turbulenzmodifikation in einem horizontalen Kanal bei unterschiedlichen Wandrauigkeiten*. PhD thesis, Martin-Luther-Universität Halle-Wittenberg, Germany, 2004.
22. J. Kussin and M. Sommerfeld. Experimental studies on particle behaviour and turbulence modification in horizontal channel flow with different wall roughness. *Exp. Fluids*, 33: 143–159, 2002.
23. P. Lammers, G. Wellein, T. Zeiser, G. Hager, and M. Breuer. Have the vectors the continuing ability to parry the attack of the killer micros? In M. Resch, Th. Bönisch, K. Benkert, T. Furui, Y. Seo, and W. Bez, editors, *High Performance Computing on Vector Systems, Proc. of the High Performance Computing Center Stuttgart, March 17–18, 2005*, pages 25–37. Springer, Berlin, 2006. ISBN 3-540-29124-5.

24. C. Marchioli, V. Armenio, and A. Soldati. Simple and accurate scheme for fluid velocity interpolation for Eulerian-Lagrangian computation of dispersed flow in 3D curvilinear grids. *Computers & Fluids*, 36:1187–1198, 2007.
25. M. R. Maxey and J. J. Riley. Equation of motion for a small rigid sphere in a non-uniform flow. *Phys. Fluids*, 26:883–889, 1983.
26. J. B. McLaughlin. Inertial migration of a small sphere in linear shear flows. *J. Fluid Mech.*, 224:261–274, 1991.
27. S. Obermair, C. Gutsch, J. Woisetschlager, and Staudinger G. Flow pattern and agglomeration in the dust outlet of a gas cyclone investigated by phase Doppler anemometry. *Powder Technology*, 156:34–42, 2005.
28. B. Oesterlé and T. Bui Dinh. Experiments on the lift of a spinning sphere in a range of intermediate Reynolds numbers. *Experiments Fluids*, 19:16–22, 1998.
29. J. Pozorski and S. V. Apte. Filtered particle tracking in isotropic turbulence and stochastic modeling of subgrid-scale dispersion. *Int. J. Multiphase Flow*, 35:118–128, 2009.
30. S. I. Rubinow and J. B. Keller. The transverse force on a spinning sphere moving in a viscous fluid. *J. Fluid Mech.*, 11:447–459, 1961.
31. J. Smagorinsky. General circulation experiments with the primitive equations, I, The basic experiment. *Monthly Weather Review*, 91:99–165, 1963.
32. M. Sommerfeld. Analysis of collision effects for turbulent gas-particle flow in a horizontal channel: Part I. Particle transport. *Int. J. Multiphase Flow*, 29:675–699, 2003.
33. M. Sommerfeld and N. Huber. Experimental analysis and modelling of particle-wall collisions. *Int. J. Multiphase Flow*, 25:1457–1489, 1999.
34. M. Sommerfeld and J. Kussin. Analysis of collision effects for turbulent gas-particle flow in a horizontal channel: Part II. Integral properties and validation. *Int. J. Multiphase Flow*, 29:701–718, 2003.
35. M. Sommerfeld and J. Kussin. Wall roughness effect on pneumatic conveying of spherical particles in a narrow horizontal channel. *Powder Technology*, 142:180–192, 2004.
36. M. Sommerfeld, B. von Wachem, and R. Oliemans. Best practice guidelines for computational fluid dynamics of dispersed multiphase flows. In *SIAMUF, Swedish Industrial Association for Multiphase Flows, ERCOFTAC*, 2008. ISBN 978–91–633–3564–8.
37. K. D. Squires and O. Simonin. LES-DPS of the effect of wall roughness on dispersed-phase transport in particle-laden turbulent channel flow. *Int. J. Heat Fluid Flow*, 27:619–626, 2006.
38. A. W. Vreman. Turbulence characteristics of particle-laden pipe flow. *J. Fluid Mech.*, 584:235–279, 2007.

Large-Eddy Simulation of Supersonic Film Cooling at Incident Shock-Wave Interaction

Martin Konopka, Matthias Meinke, and Wolfgang Schröder


Abstract The impact of shock waves on supersonic cooling films is studied using large-eddy simulations (LES). A laminar cooling film is injected through a slot at a Mach number $Ma_i = 1.8$ into a fully turbulent boundary layer at a freestream Mach number $Ma_\infty = 2.44$. The cooling film is disturbed by oblique shock waves at deflection angles of 5° and 8° at two downstream positions of the slot. At shock impingement close to the slot, i.e., within the potential-core region, at a flow deflection of 5° , a cooling effectiveness decrease of 33 % occurs downstream of the separation bubble compared to a configuration without shock impingement. If the same shock impinges further downstream upon the boundary-layer region, the decrease is only 17 %. The stronger 8° shock wave at the further downstream impingement position leads to a maximum decrease of 33 %. The current report presents a concise version of Konopka et al. (4th European conference for aerospace sciences, 2011).

1 Introduction

In supersonic combustion ramjets (Scramjets) shock waves are present in the isolator and combustion chamber. These shock waves interact with cooling films if supersonic film cooling is used to protect the engine's interior surfaces from the intense aerodynamic heating and hot combustion products. This film cooling problem, i.e., the interaction of shock waves with a supersonic cooling film injected through a slot, was investigated experimentally for turbulent flows [2–4] to assess the impact on heat transfer and cooling effectiveness. More recent experimental studies performed by Kanda and Ono [5] and Kanda et al. [6] on film cooling with shock wave interaction found that shock waves have only little effects on

M. Konopka (✉) · M. Meinke · W. Schröder
Institute of Aerodynamics, RWTH Aachen University, Wüllnerstr. 5a, 52062 Aachen, Germany
e-mail: m.konopka@aia.rwth-aachen.de

Fig. 1 Main flow features of shock-wave impingement upon the potential-core region (case II)



cooling effectiveness. Kanda and Ono and Kanda et al. showed that the cooling effectiveness is mainly reduced by the increased wall-recovery temperature which is caused by the reduced local Mach number downstream of the shock wave. However, Peng and Jiang [7] found in their computational study using the Reynolds averaged Navier Stokes (RANS) equations with the Menter SST [8] turbulence model that the mole fraction of the injected gas decreased at the impingement position of the shock waves. This indicates that increased mixing due to the excited turbulence levels in the cooling flow occurred additionally to the effect of the reduction of the local Mach number. Therefore, the following study focuses on whether the reduction of cooling effectiveness by shock waves is only a function of the local Mach number or if increased turbulence plays a role. Five cooling configurations at the same injection condition are investigated. A zero pressure gradient configuration (case I) is compared to experiments of Juhany et al. [9]. Two configurations where a shock wave generated by a flow deflection angle of 5° at varying shock-impingement positions are analyzed. The first impingement position is located within the potential-core region [10, 11], where the injected cooling flow is laminar (case II). The second position is located further downstream in the boundary-layer region [9, 10] (case III), where the cooling flow mixes with the freestream and boundary-layer-like velocity profiles occur. Additionally, a stronger shock wave generated by an 8° flow deflection impinges at the same downstream position within the boundary-layer region upon the cooling flow (case IV).

The principal flow features of shock-cooling-film interactions within the potential-core region (case I) are sketched in Fig. 1. The potential-core region, which originates at the slot, is encompassed by the laminar slot boundary layer and the mixing layer where the freestream mixes with the cooling flow. On top of the mixing layer there is the shear layer which emanates from the lip and is fed by the turbulent boundary layer. The incident oblique shock wave crosses these layers and causes the laminar separation bubble. Slightly upstream of the separation bubble compression waves decelerate the flow, at the crest of the bubble an expansion fan emerges and then the flow reattaches creating compression waves which unify and form a shock wave. At the end of the separation bubble the laminar slot boundary layer undergoes transition. The report is organized as follows. First, the numerical method will be presented, subsequently, details of the boundary conditions and

the computational mesh will be given. Next, the flow configuration is explained and the results section follows. In the results section, the zero pressure gradient film cooling configuration is validated by experiments from Juhany et al. [9], then, shock-cooling-film interactions are investigated in terms of cooling effectiveness, instantaneous flow properties, and turbulence statistics. Subsequently, some details on the consumed computational resources are given. Finally, some conclusions are drawn.

2 Numerical Method

In the past RANS equations using among other approaches $k - \varepsilon$ turbulence models have been used [12–14] to model film cooling problems with varying success, i.e., depending on the variant of the model, there was quite a discrepancy in the wall temperature distributions. This is caused in part by the modeling of the mixing layer between the cooling flow and the freestream, where no satisfying model to account for density gradients exist [15]. Therefore, in this study high fidelity turbulence modeling is applied, i.e., large-eddy simulations (LES) are performed.

The Navier-Stokes equations are discretized at second-order accuracy using a modified mixed-centered upwind advective upstream splitting method (AUSM) [16] for the Euler terms. The discretization of the non-Euler terms is done using a centered approximation at second-order accuracy. The temporal integration is done by a second-order five-stage low-storage Runge-Kutta method. The non-resolved subgrid scales are implicitly modeled using the MILES ansatz [17]. The viscosity is evaluated by a power law $\mu/\mu_0 = (T/T_0)^{0.72}$ where T_0 denotes the stagnation temperature. A detailed summary of the flow solver used in this study is given by Meinke et al. [18]. The accuracy of its solutions in fully turbulent flows is discussed in [19–21]. The solution algorithm has also shown convincing results in supersonic flows involving shock-boundary-layer interactions [22].

3 Boundary Conditions and Computational Mesh

The prescription of realistic inflow variables for compressible turbulent boundary layers is a challenge in LES since at every time step a different instantaneous inflow distribution is required. This problem can be solved by computing the compressible turbulent boundary layer from the leading edge of a flat plate. To avoid this computationally costly approach an independent boundary layer simulation is performed using the rescaling method proposed by El-Askary et al. [22], which is based on Lund et al.'s [23] approach considering compressibility. In the boundary-layer domain, which is depicted in Fig. 2, the inflow distribution is generated by rescaling the flow variables obtained from a plane within the domain such that a constant boundary layer thickness at the inflow is achieved. A slice of the flow field is then extracted at every time step and injected into the main film cooling domain.




Fig. 2 Sketch of the boundary conditions and the turbulent inflow data generation method

At the lower wall of the film-cooling simulation domain adiabatic no-slip boundary conditions are imposed, at the exit all variables are extrapolated. To avoid any spurious oscillations a sponge layer is used at the exit and at the top of the boundary-layer domain where the flow variables are driven to the desired target variables [22]. At the slot a laminar supersonic cooling flow is prescribed. In the spanwise direction, periodic boundary conditions are used. The desired shock-wave strength and angle is generated at the upper boundary by setting flow conditions which satisfy the Rankine-Hugoniot relations.

The body-fitted computational mesh consists of 17.1 million grid points with an equidistant spacing in the streamwise and spanwise directions. The resolution at the wall in inner coordinates is $\Delta x^+ = 20$, $\Delta y^+ = 0.5$, $\Delta z^+ = 10$ in the streamwise, wall-normal, and spanwise directions. A grid study for this cooling configuration was performed by Konopka et al. [24] where the current resolution was found to be adequate. Due to the spanwise domain size of $z/S = 2.2$ the computations do not resolve some large-scale vortices. A similar cooling configuration with a lower injection Mach number of $Ma_i = 1.2$ was investigated by Konopka et al. [25] and it was found that the spanwise domain size of the present study leads to an overestimation of the cooling effectiveness by about 9 % at a zero pressure gradient. The overestimation is similar at shock interaction. Therefore, reasonable conclusions can be drawn from the current study at the present spanwise domain extent.

4 Flow Configuration

The freestream and injection flow properties match those used in the experiment by Juhany et al. [11]. The freestream Mach number is set to $Ma_\infty = 2.44$ and the freestream Reynolds number $Re_\infty = u_\infty S / \nu_\infty$ based on the slot height S , the freestream velocity u_∞ , and the freestream kinematic viscosity ν_∞ is $Re_\infty = 13,500$. The freestream boundary-layer thickness at the tip of the lip is $\delta/S = 2.2$. The blowing rate M , the injection Mach number Ma_i , and the total

Table 1 Flow configuration

Case	Ma_i	$T_{ii}/T_{i\infty}$	$M = \rho_i u_i / \rho_\infty u_\infty$	β	Ma_3	x_{imp}/S	L_{in}	$\Delta t \cdot u_\infty / 0.03 L_{in}$
I	1.8	0.76	0.74	0°	2.44	–	–	–
II	1.8	0.76	0.74	5°	2.04	17	7.2	1.04
III	1.8	0.76	0.74	5°	2.04	60	3.2	2.58
IV	1.8	0.76	0.74	8°	1.82	60	5.8	1.29

temperature ratio $T_{ii}/T_{i\infty}$ of the cooling flow are listed in Table 1 and are kept constant at all considered cases. Case I is a zero pressure gradient configuration and is validated in Sect. 5.2.3 by the experiments performed by Juhany et al. [11]. The shock waves caused by a $\beta = 5^\circ$ flow deflection angle impinge upon the potential-core region in an inviscid flow at a downstream distance $x/S = 17$ (case II) and in the boundary-layer region at $x/S = 60$ (case III). Additionally a stronger shock wave is considered impinging at the same downstream distance $x/S = 60$ upon the cooling flow (case IV). At the cooling slot, the boundary-layer thickness of the upper and lower laminar supersonic boundary layers are assumed $\delta/S = 0.07$. The pressure is set to the freestream pressure p_∞ .

The duration Δt of the simulations which is used for the averaging process is given in Table 1. The time interval Δt is normalized by the timescale $L_{in}/(0.03u_\infty)$ where L_{in} is the interaction length, i.e. the distance between the point of separation and the location where the shock impinges upon the wall assuming inviscid conditions. This time scale is associated with the low frequency motion of the shock [26]. The current time averaging window covers at least one complete cycle of the shock.

5 Results

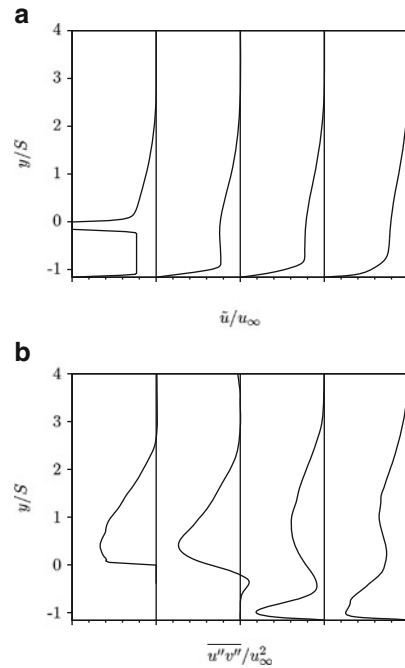
The results section is divided into two major parts, first, the zero pressure gradient film cooling configuration is validated and the length of the potential-core region is determined. Then, shock-cooling-film interactions are investigated in terms of instantaneous and mean flow properties, cooling effectiveness distributions, and turbulence statistics.

5.1 Zero-Pressure Gradient Configuration (Case I)

The zero-pressure gradient configuration (case I) is validated in Sect. 5.2.3 by comparing numerically obtained cooling effectiveness distributions with the experiments by Juhany et al. [11].

To determine the length of the potential-core region for the choice of the shock-impingement positions at cases II–IV, streamwise velocity profiles of case I are

Fig. 3 Streamwise velocity profiles and Reynolds shear stress component $\overline{u''v''}/u_\infty^2$ for case I; grid spacing is $\tilde{u}/u_\infty = 1$ and $\overline{u''v''}/u_\infty^2 = 0.0035$.
(a) $x/S = 0, 15, 30, 40$.
(b) $x/S = 0, 15, 30, 40$

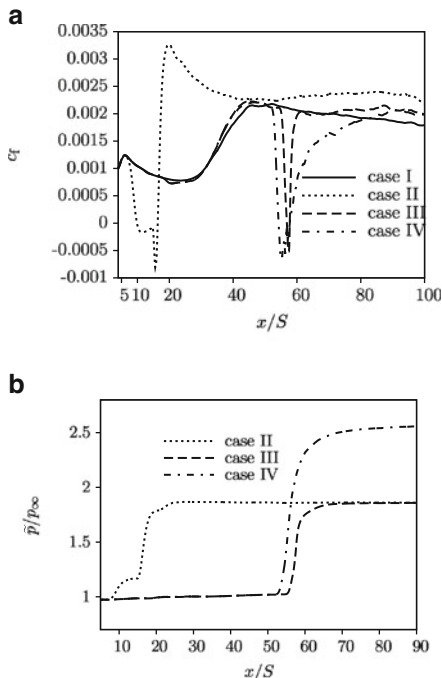


shown in Fig. 3a. At $x/S = 0$, the potential-core region is visible by the constant velocity in the slot flow at $-1.16 \leq y/S \leq -0.16$. At increasing downstream distance, this region of constant streamwise velocity shrinks since the laminar slot boundary layer at the wall and the mixing layer, which emanates from the lip, grow and eventually merge. This becomes apparent at the downstream distance $x/S = 40$, where the velocity profile already resembles that of an undisturbed fully turbulent boundary layer. The Reynolds shear stress component $\overline{u''v''}/u_\infty^2$, which is composed of the fluctuating velocity components in the streamwise and spanwise directions from the Favre-averaged mean, is shown in Fig. 3b. The potential-core region is evident by the zero shear stress at $x/S = 0$, $-1.16 \leq y/S \leq -0.16$. At $x/S = 30$, the laminar slot boundary layer has undergone transition which is indicated by the negative peak shear stress close to the wall. However, a small area with minimal shear stress at $x/S = 30$, $y/S = -0.4$ is still visible, indicating a potential flow. At $x/S = 40$, the slot boundary layer and the mixing layer have merged, marking the end of the potential-core region and the beginning of the boundary-layer region.

5.2 Analysis of Supersonic Film Cooling with Shock Waves

In the following, cases II–IV are considered, i.e., shock-cooling interactions are investigated to show the impact on cooling effectiveness and turbulence statistics.

Fig. 4 Skin friction coefficient and wall pressure distribution plotted vs. the streamwise distance from the slot. **(a)** Skin friction coefficient. **(b)** Wall pressure coefficient.



5.2.1 Flow Characteristics of Supersonic Film Cooling with Shock Waves

Figure 4a shows the skin-friction coefficient distribution vs. streamwise distance of all considered cases to clearly identify regions with separated flow. At the zero-pressure gradient cooling configuration, the skin friction rises quickly at the end of the potential-core region. At the shock-impingement position within the potential-core region at case II, a separation bubble at a length $L_{sep}/S = 6.5$ exists. Downstream of the laminar separation bubble the skin-friction coefficient shows a pronounced peak, which is a clear sign of the transition of the laminar slot boundary layer. At the shock impingement upon the boundary-layer region at case III, a much smaller separation bubble is found with a length of $L_{sep}/S = 0.9$, since the wall-bounded flow is turbulent at this downstream position. The greater shock strength at case IV due to the higher deflection angle $\beta = 8^\circ$ leads to a slightly larger separation $L_{sep}/S = 3.6$, but the minimum skin-friction coefficient has the same level as at case III.

The wall pressure distributions for cases II–IV are juxtaposed in Fig. 4b. The large laminar separation bubble at case II is visible by the plateau in the pressure distribution at $x/S = 15$. Such a plateau is not observed at the shock impingement in the boundary-layer region at both shock strengths.




Fig. 5 Turbulent structures visualized by the Q criterion with mapped-on Mach number contours; shock- and expansion-fan contours at $\partial u_i / \partial x_i \cdot S / u_\infty = -0.2$, case II




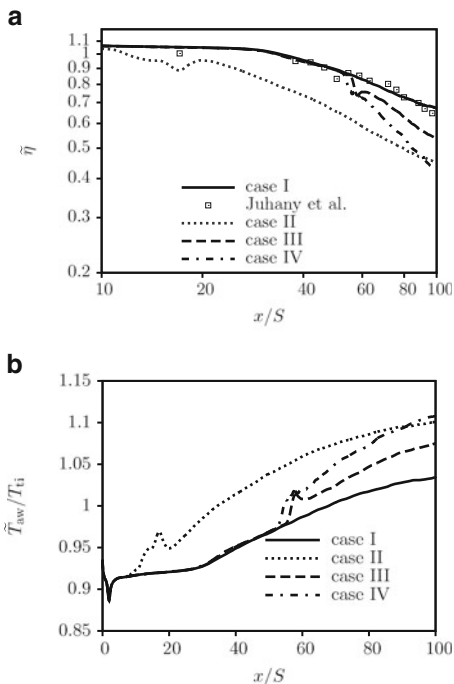
Fig. 6 Instantaneous numerical Schlieren image at the centerline of the computational domain at case II; (a) shear layer, (b) mixing layer, (c) potential-core region, (d) laminar slot boundary layer, (e) laminar separation bubble, (f) incident shock wave, (g) reflected shock wave

5.2.2 Instantaneous Flow Field

Figure 5 shows the vortical structures visualized by the Q criterion [27] with mapped-on Mach number contours for the shock impinging upon the potential-core region (case II). Shock waves are indicated by gray contours of the velocity divergence $\partial u_i / \partial x_i \cdot S / u_\infty = -0.2$. The incoming turbulent boundary layer above the slot is disturbed by an expansion fan which is followed by a shock wave. In the injected cooling film no vortices are present. The upper border of this region is defined by the mixing and shear layers which emanate from the lip. The incident oblique shock wave penetrates through the shear and mixing layers. Downstream of the reflected shock vortices are detected in the cooling flow terminating the potential-core region.

These flow features at the shock-impingement position within the potential-core region (case II) are clearly visible in the instantaneous numerical Schlieren image shown in Fig. 6. The potential-core region (c), which is encompassed by the laminar slot boundary layer (d) and the mixing layer (b), extends even downstream of the shock impingement position at this time level to $x/S = 14.5$. Further downstream at $x/S = 17$, disturbances appear in the potential-core flow which is near the time-averaged reattachment point $x_R/S = 16.3$ of the separation bubble (e).

Fig. 7 Cooling effectiveness and wall temperature vs. streamwise distance for all cases. **(a)** Cooling effectiveness. **(b)** Wall temperature



5.2.3 Cooling Effectiveness

The temporal and spanwise averaged cooling effectiveness definition reads

$$\tilde{\eta} = \frac{\tilde{T}_{aw} - T_{r\infty}}{T_{i\infty} - T_{r\infty}}, \tag{1}$$

where the quantity \tilde{T}_{aw} denotes the Favre-averaged adiabatic wall temperature, $T_{r\infty}$ the recovery temperature of the freestream, and $T_{i\infty}$ the recovery temperature of the cooling flow. The cooling effectiveness distribution in Fig. 7a for case I shows a good agreement with the experiments of Juhany et al. [11]. An in-depth analysis of this cooling configuration can be found in [24]. Konopka et al. [25] investigated a similar cooling configuration with a lower injection Mach number of $Ma_i = 1.2$. Note that Konopka et al. found that the numerical cooling effectiveness is 9 % lower when using a spanwise domain size of $z/D = 2.2$ compared to $z/D = 4.4$ since some large-scale turbulence is not captured. However, the conclusions drawn in this study are still reasonable since the magnitude of the overestimation is the same for cases I–IV. In Fig. 7a it is evidenced that at the shock impingement upon the potential-core region (case II), the cooling effectiveness is reduced beginning at the separation point of the laminar slot boundary layer at $x_S/S = 9.8$. Downstream of the transition of the laminar slot boundary layer downstream of the separation

bubble, the cooling effectiveness decreases further. The turbulent slot boundary layer immediately mixes with the mixing layer emanating from the lip which extends at this downstream position towards the wall. Therefore, it is clear that at case II, besides the reduction of the recovery temperature of the wall by the reduction of the local Mach number, downstream of the shock wave the transition of the laminar slot boundary layer plays an important role. At the more downstream shock impingement position at $x_{\text{imp}}/S = 60$ within the boundary-layer region (case III), the streamwise slope of the cooling effectiveness has increased compared to the zero pressure gradient configuration, indicating increased heat and momentum transfer due to the shock-wave impingement. At case IV, where the flow-deflection angle is increased by 3° compared to case III, the streamwise slope of the cooling effectiveness decline is even steeper.

The cooling effectiveness values which are slightly above unity at $x/S = 10$ are due to the expansion fan and shock wave emanating from the lower tip of the lip which impinge upon the laminar slot boundary layer. This is evidenced by the temperature dip at $x/S = 2$ in the wall-temperature distribution in Fig. 7b.

Besides the averaged wall temperatures the frequent changes of the wall temperature have to be considered in the design process of the wall material of a Scramjet engine, as they might lead to thermal fatigue or local hot spots in the engine. Therefore, the cooling effectiveness fluctuations


$$\sqrt{\eta'^2} = \frac{\sqrt{T''_{\text{aw}}}}{T_{\text{r}\infty} - T_{\text{r}\infty}} \quad (2)$$

are evaluated for the shock-cooling film interactions in Fig. 8a. It is evident that the highest cooling effectiveness fluctuations occur at $x/S = 17.82$ at case II, which is downstream of the averaged reattachment point of the separation bubble. At the more downstream shock-impingement position at cases III and IV, the level of cooling effectiveness fluctuations is 16 % lower than at case II and hardly affected by the shock strength, since there is barely any difference between cases III and IV. To investigate the cause for the high cooling effectiveness fluctuation level at the shock impingement upon the potential-core region, conditional averages of the flow field are considered. In Fig. 8 the dashed line corresponds to the averaging threshold, i.e. all snapshots of the flow field with a local cooling effectiveness $\eta \geq 1.1$ are considered. The averaged total temperature contours shown in Fig. 9 consist of 51 snapshots of the flow field. Figure 9c shows the conditionally averaged total temperature contours collected at the time when the condition $\eta \geq 1.1$ holds. It is visible that a region of cold fluid exists at the position of high cooling effectiveness fluctuations. The origin of the cold fluid is analyzed by considering the conditionally averaged flow field at $\Delta t/(S/u_\infty) = -2.5$ shown in Fig. 9b before the time when $\eta \geq 1.1$ is satisfied. It is evidenced that the region of cold fluid is now located upstream of the point of maximum cooling effectiveness fluctuations. Considering the conditionally averaged flow field at $\Delta t/(S/u_\infty) = -5.5$ in Fig. 9a it is visible

Fig. 8 Cooling-effectiveness fluctuation distributions vs. streamwise distance and instantaneous cooling effectiveness vs. time.

(a) Cooling effectiveness fluctuations.

(b) Instantaneous cooling effectiveness at $x/S = 17.82, z/S = 1.1$ for case II; *dashed line* corresponds to a cooling effectiveness $\eta = 1.1$



that the region of cold fluid is located upstream compared to the time interval $\Delta t/(S/u_\infty) = -2.5$ and is now part of the separation bubble. Note that the separation bubble has shrunk comparing the time intervals $\Delta t/(S/u_\infty) = -5.5$ and $\Delta t/(S/u_\infty) = -2.5$ showing that the region of cold fluid has been shed off the separation bubble. Hence it can be concluded that the shedding of cold fluid off the separation bubble is responsible for very high cooling effectiveness values.

5.2.4 Mean Flow Field

To analyze the impact of the shock waves on the mixing in the flow field near the wall, the dimensionless fluid temperature

$$\Theta = \frac{\widetilde{T}_t - T_{t\infty}}{T_{ti} - T_{t\infty}} \tag{3}$$

is evaluated. Its definition is similar to the cooling effectiveness $\widetilde{\eta}$ except that total temperatures are used as reference temperatures. The quantity Θ reaches a value of 1 in the cooling flow and 0 in the freestream. Figure 10a shows dimensionless fluid temperatures in the slot vicinity for the zero-pressure gradient configuration (case I) and the shock-cooling interaction within the potential-core region (case II). The spreading of the mixing layer at increasing downstream distance is evident in the




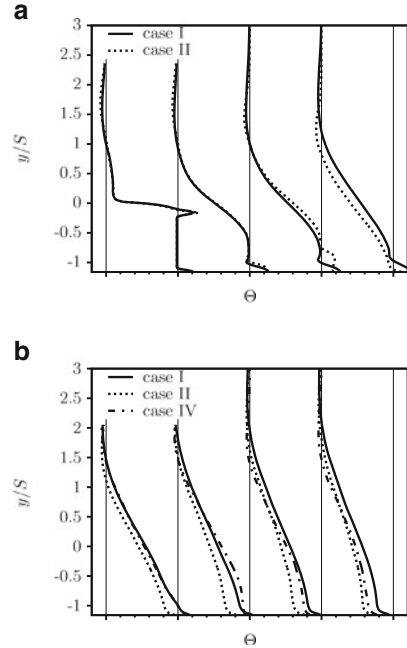
Fig. 9 Conditionally averaged total temperature contours in the x, y -plane at $z/S = 1.1$ for case II, the *arrow* indicates the point of maximum cooling-effectiveness fluctuation and the circle marks regions of low energetic fluid. (a) $\Delta t/(S/u_\infty) = -5.5$. (b) $\Delta t/(S/u_\infty) = -2.5$. (c) $\Delta t/(S/u_\infty) = 0.0$

profiles, e.g., at $x/S = 10$. It is the region where Θ varies from 1 to 0 at increasing wall-normal distances. Downstream of the shock-wave interaction at case II at $x/S = 25$, the dimensionless fluid temperature is lower than at case I indicating increased mixing downstream of the shock compared to the zero-pressure gradient configuration case I. Further downstream, at the stronger shock-cooling interaction in the boundary-layer region at case IV, the dimensionless fluid temperature rises first at $x/S = 50$ due to the displacement of the separation bubble compared to the case I profile. Further downstream, the dimensionless fluid temperature at case IV is quickly reduced throughout the entire boundary-layer profile and at $x/S = 80$ it almost matches the case II profile. At this downstream position, the cooling effectiveness of these two cases is also alike (Fig. 7a).

5.2.5 Turbulence Statistics

Shock waves impinging upon turbulent boundary layers are known [28] to lead to increased turbulence levels downstream of the impingement point. This increased turbulence levels lead to increased mixing and heat and momentum transport of the cooling flow and the freestream.

Fig. 10 Dimensionless fluid temperature profiles Θ at several downstream positions, grid spacing is $\Theta = 1$.
(a) $x/S = 0, 10, 15, 25$.
(b) $x/S = 40, 60, 70, 80$



This is evidenced by the Reynolds shear stress profiles in Fig. 11. At case I the upper boundary of the potential-core region is located where $\overline{u''v''}/u_\infty^2$ begins to deviate from zero at increasing wall-normal distances (Fig. 11a). At increasing streamwise distance from the slot, the upper boundary moves towards the wall. At shock-impingement, the profile at $x/S = 15$ at case II shows a negative peak at $y/S = -0.9$ nearly at the end of the separation bubble. Further downstream of the shock-impingement position, absolute $\overline{u''v''}/u_\infty^2$ levels have risen at $x/S = 25$, $-1.16 \leq y/S \leq 1$. Therefore, it is evident that the shock-wave impinging upon the potential-core region has led to the transition of the laminar slot boundary layer and has reduced the length of the potential core region. At shock impingement within the boundary-layer region negative peaks of $\overline{u''v''}/u_\infty^2$ are detected at $x/S = 60$ close to the wall, which move off the wall further downstream. These higher $\overline{u''v''}/u_\infty^2$ values of the cases III and IV than of the cases I and II in the boundary-layer region at $x/S = 80$ could explain the steeper streamwise slope in cooling effectiveness observed in Fig. 7a.

5.3 Computational Resources

The simulations for this study were performed on the vector computer NEC SX-9 at the HLR Stuttgart. The grids are divided into blocks which reside on a single CPU. Data is exchanged using the message passing interface (MPI). The computational

Fig. 11 Reynolds shear stress component at several downstream positions, grid spacing is $\overline{u''v''}/u_\infty^2 = 0.0035$ **(a)** $x/S = 0, 10, 15, 25$. **(b)** $x/S = 40, 60, 70, 80$

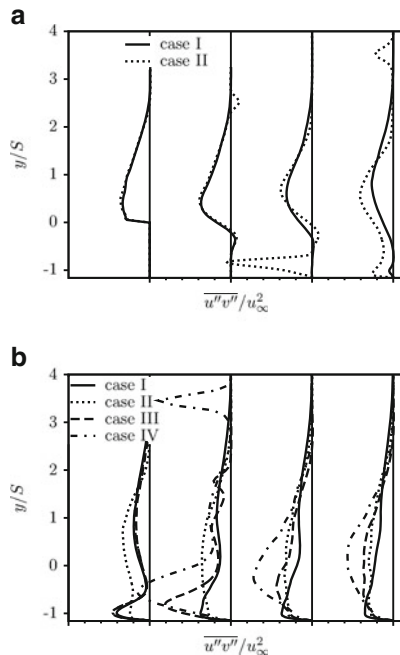


Table 2 Performance on NEC SX-9

	Case I	Case IIlarge	Case multirow
Number of CPUs	7	16	30
Number of Nodes	1	1	2
Grid points/CPU	$2.44 \cdot 10^6$	$1.51 \cdot 10^6$	$0.5 \cdot 10^6$
Total grid points	$17.1 \cdot 10^6$	$24.1 \cdot 10^6$	$15 \cdot 10^6$
Avg. User Time (s)	41,552	41,556	41,761
Avg. Vector Time (s)	36,827	38,502	38,844
Vector Operations Ratio (%)	99.6	99.7	99.54
Avg. Vector Length	240.585	251.439	235.25
Memory/CPU (MB)	2,422	1,998	1,426
Total Memory (GB)	16.95	31.98	42.8
Avg. MFLOPS/CPU	13,997	17,873	8,587
Max. MFLOPS/CPU	19,252	19,590	15,854
Total GFLOPS	97.8	285.881	257.5

details are given in Table 2, where “Case I” denotes the case I of the current slot cooling computation. It is evident that the flow solver is fully vectorized with a vector operation ratio higher than 99 %. Furthermore, a “Case IIlarge” is given, which is similar to “Case I” but with a larger streamwise domain extent. The grid of this case is 1.4-fold larger than that of case I and the computation reaches an overall

performance of 285 GFLOPS. A different computation for the analysis of a cooling configuration using several rows of discrete holes is denoted “Case Multirow” in Table 2. This computation was run on two SX-9 nodes since it has a complex grid with small grid cells and large block surfaces. Therefore, the overall performance at “Case Multirow” is lower than at “Case Ilarge” at 62 % of the grid points of “Case Ilarge”.

6 Conclusion

Large-eddy simulations of shock-cooling-film interactions have been performed. The shock-waves impinge upon the potential-core (case II) and boundary-layer regions (cases III and IV). At the shock-wave impingement position within the potential-core region, the transition of the laminar slot boundary layer occurred downstream of the separation bubble. The increased turbulence levels in the shear- and mixing layer located between the cooling flow and the freestream led to a decrease of cooling effectiveness compared to a zero-pressure gradient configuration (case I) of 33 %. Large cooling effectiveness fluctuations were detected at shock-impingement downstream of the separation bubble since it sheds off large patches of cold fluid. At the shock impingement upon the boundary-layer region at the same shock strength as at the shock impingement upon the potential-core region, a less drastic decrease of cooling effectiveness (17 %) was observed. However, the streamwise cooling effectiveness gradient showed a higher magnitude (case III) compared to the zero-pressure gradient configuration (case I). At increasing shock strength at the further downstream impingement position (case IV), the cooling effectiveness decreased even more rapidly, i.e., the streamwise cooling effectiveness gradient showed an even higher magnitude.

Acknowledgements The support of this research by the German Research Association (DFG) in the framework of the Research Training Group “Aero-Thermodynamic Design of a Scramjet Propulsion System for Future Space Transportation Systems” 1095/2 and the High Performance Computing Center Stuttgart (HLRS) is gratefully acknowledged.

References

1. Konopka, M., Meinke, M., Schröder, W.: Large-Eddy Simulation of High Mach Number Film Cooling with Shock-Wave Interaction. In: 4th European Conference for Aerospace Sciences. (2011)
2. Alzner, E., Zakkay, V.: Turbulent Boundary-Layer Shock Interaction with and without Injection. *AIAA J.* **9** (1971) 1769 – 1776
3. Olsen, G., Nowak, R., Holden, M., Baker, N.R.: Experimental Results for Film Cooling in 2-d Supersonic Flow Including Coolant Delivery Pressure, Geometry, and Incident Shock Effects. In: 28th Aerospace Sciences Meeting, Reno, NV (1990) AIAA Paper 90-0605.

4. Holden, M., Nowak, R., Olsen, G., Rodriguez, K.: Experimental Studies of Shock Wave/Wall Jet Interaction in Hypersonic Flow. In: 28th Aerospace Sciences Meeting, Reno, NV (1990) AIAA Paper 90-0607.
5. Kanda, T., Ono, F., Saito, T., Takahashi, M., Wakamatsu, Y.: Experimental Studies of Supersonic Film Cooling with Shock Wave Interaction. *AIAA J.* **34** (1996) 265–271
6. Kanda, T., Ono, F.: Experimental Studies of Supersonic Film Cooling with Shock Wave Interaction (II). *J. Thermophys. Heat Transfer* **11** (1997) 590–592
7. Peng, W., Jiang, P.X.: Influence of Shock Waves on Supersonic Film Cooling. *J. Spacecr. Rockets* **46** (2009) 67–73
8. Menter, F.M.: Zonal Two Equation k-omega Turbulence Models for Aerodynamic Flows. In: 24th Fluid Dynamics Conference, AIAA (1993) AIAA Paper 93-2906.
9. Juhany, K., Hunt, M.: Flowfield Measurements in Supersonic Film Cooling Including the Effect of Shock-Wave Interaction. *AIAA J.* **32** (1994) 578–585
10. Seban, R.A., Back, L.H.: Velocity and Temperature Profiles in Turbulent Boundary Layers with Tangential Injection. *J. of Heat Transfer* **84** (1962) 45–54
11. Juhany, K., Hunt, M., Sivo, J.: Influence of Injectant Mach Number and Temperature on Supersonic Film Cooling. *J. Thermophys. Heat Transfer* **8** (1994) 59–67
12. Bowersox, R., Schetz, J.: Compressible Turbulence Measurements in a High-Speed High-Reynolds-number Mixing Layer. *AIAA J.* **32** (1994) 758–764
13. Aupoix, B., Mignosi, A., Viala, S., Bouvier, F., Gaillard, R.: Experimental and Numerical Study of Supersonic Film Cooling. *AIAA J.* **36** (1998) 915–923
14. Sarkar, S.: Numerical Simulation of Supersonic Slot Injection Into a Turbulent Supersonic Stream. *Int. J. Turbo Jet. Eng.* **17** (2000) 227–240
15. Aupoix, B.: Modelling of compressibility effects in mixing layers. *J. Turbulence* **5** (2004) N 7
16. Liou, M., Steffen, C.: A new flux splitting scheme. *J. Comput. Phys.* **107** (1993) 23–39
17. Boris, J., Grinstein, F., Orana, E., Kolbea, R.: New Insights into Large Eddy Simulation. *Fluid Dyn. Res.* **10** (1992) 199–228
18. Meinke, M., Schröder, W., Krause, E., Rister, T.: A comparison of second and sixth-order methods for large-eddy simulations. *Comp. & Fluids* **31** (2002) 695–718
19. Alkishriwi, N., Meinke, M., Schröder, W.: A large-eddy simulation method for low mach number flows using preconditioning and multigrid. *Comp. & Fluids* **35** (2006) 1126–1136
20. Renze, P., Schröder, W., Meinke, M.: Large-Eddy Simulation of Film Cooling Flows at Density Gradients. *Int. J. Heat Fluid Flow* **29** (2008) 18–34
21. Rütten, F., Schröder, W., Meinke, M.: LES of low frequency oscillations of the dean vortices in turbulent pipe bend flows. *Phys. Fluids* **17** (2005) 035107
22. El-Askary, W.A., Schröder, W., Meinke, M.: LES of compressible wall-bounded flows. In: 16th AIAA Computational Fluid Dynamics Conference, Orlando, Florida (2003) AIAA-2003-3554.
23. Lund, T., Xiaohua, W., Squires, K.D.: Generation of Turbulent Inflow Data for Spatially-Developing Boundary Layer Simulations. *J. Comput. Phys.* **140** (1998) 233–258
24. Konopka, M., Meinke, M., Schröder, W.: Large-Eddy Simulation of Supersonic Film Cooling at Laminar and Turbulent Injection. In: 17th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, San Francisco, CA (2011) AIAA 2011-2250.
25. Konopka, M., Meinke, M., Schröder, W.: Large-Eddy Simulation of Shock-Cooling-Film Interaction. *AIAA J.* (accepted for publication) (2012)
26. Piponniau, S., Dussauge, J., Debieve, J.F., Dupont, P.: A Simple Model for Low-Frequency Unsteadiness in Shock-Induced Separation. *J. Fluid Mech.* **629** (2009) 87–108
27. Dubief, Y., Delcayre, F.: On coherent-vortex identification in turbulence. *J. Turbulence* **1** (2000) N11
28. Priebe, S., Wu, M., Martin, M.P.: Direct numerical simulation of a reflected-shock-wave/turbulent-boundary-layer interaction. *AIAA J.* **47** (2009) 1173–1184

Large-Eddy Simulations of Stratified and Non-stratified T-junction Mixing Flows

David Klören and Eckart Laurien

Abstract Thermal mixing of coolants with large temperature differences in cooling circuits of power plants may lead to high cycle thermal fatigue in the pipe material. In these mixing regions cracks are often observed in the vicinity of weld seams. In this study the influence of a weld seam in straight pipe flows, isothermal T-junction mixing flows and stratified T-junction mixing flows are investigated with Large-Eddy Simulations (LES). The results are compared to experimental data. Furthermore, T-junction flows with large temperature differences, which are based on an experimental setup at the Institute of Nuclear Energy and Energy Systems (IKE), are numerically investigated and characterized.

1 Introduction

The safety analysis of a nuclear power plant involves many components which are subjected to thermal loading. On the one hand, system transients and steady state stratification effects lead to Low Cycle Fatigue (LCF). On the other hand, the temperature fluctuations caused by flow instabilities of non-isothermal mixing flows may lead to the premature loss of the mechanical integrity of the surrounding material. Such High Cycle Fatigue (HCF) due to thermal striping has been observed, for instance, in a mixing tee pipe configuration of a residual heat removal system [2]. In contrast to LCF, Fluid-Structure-Interaction (FSI) associated with high cycle fatigue is not sufficiently understood. Furthermore, the plant instrumentation is not designed to monitor thermal transients or the flow conditions related to HCF. The assessment of safety and reliability of these components are therefore extremely conservative.

D. Klören (✉) · E. Laurien
Institute of Nuclear Technology and Energy Systems (IKE), University Stuttgart, Stuttgart,
Germany
e-mail: david.kloeren@ike.uni-stuttgart.de; eckart.laurien@ike.uni-stuttgart.de

A strategy based on numerical simulations is adopted in order to provide the characteristic thermal loads which act on the pipe walls. Hence, the models applied in Computational Fluid Dynamics (CFD) have to be validated for the phenomena such as thermal mixing, buoyancy effects and FSI.

The studies performed on the T-junction problem (e.g. Westin et al. [15]) were focused mainly on the mixing zones and the flow downstream of a T-junction. Most experiments are designed for low inlet temperature differences in order to allow optical access for non-intrusive measurements of the velocity and temperature field. For this reason, a new hot T-junction setup that allows optical access and temperature measurements in the near-wall region of the fluid is currently tested at the IKE [8]. Thermally coupled Large-Eddy Simulations of this setup were performed [7] to investigate the fluid-structure interaction.

Usually, straight smooth pipes upstream and downstream of the T-junction were considered. However, cracks were typically observed in the vicinity of weld seams [1]. In the framework of the CFD validation of mixing flows in a T-junction the influence of the weld seam has not been addressed previously.

Firstly, this numerical study investigates the influence of a weld seam on the flow field of fully developed turbulent pipe flow, isothermal T-junction mixing flow and stratified T-junction mixing flow. The results are compared against measurements. The experimental setup is designed for the development and application of optical measurement methods under cold conditions (20 °C) which will be employed in the hot T-junction experiment with temperatures up to 280 °C and a pressure of 75 bar.

Secondly, a first experimental study of the thermal conditions in the hot T-junction flow configuration indicates a stable stratified flow with upstream flow of cold water in the hot line and vice versa. In order to characterize the flow conditions further numerical simulations are performed with temperature differences up to 150 °C.

2 Setup

The geometric proportions are based on a T-junction experiment designed by the Institute of Nuclear Technology and Energy Systems (IKE) and the Material Testing Institute (MPA) at the University Stuttgart for the research in thermal fatigue and CFD model development [8]. The setup consists of a larger pipe with an inner diameter (ID) of $d_1 = 71.8$ mm. The sharp edged, 90° T-junction is connected to the main pipe and a branch pipe with an ID of $d_2 = 38.9$ mm. Three different basic configurations are considered: a single pipe flow (A), a T-junction mixing flow with constant fluid properties (B) and a T-junction mixing flow affected by buoyancy forces and stratification (C). Although each configuration was investigated with and without weld seam model numerically as well as experimentally, only the cases including the weld seam model are considered for this discussion.




Fig. 1 Geometry of straight pipe (left) and T-junction (right) with weld seam. The numerical T-junction model is shown as a cut (in the symmetry plane) for better visualization. The position of the weld seam model is the dent in the surface downstream of the T-junction

Table 1 Simulation test cases for the weld seam investigations. Simulation without weld seam (Case 7) described by Kloeren [7]

Case no.	Setup	Weld seam	Volumetric flow rate	
			Q_1 [l/s]	Q_2 [l/s]
1	A	Yes	0.8	–
2	A	Yes	0.4	–
3	A	Yes	0.4	0.1
4	B	Yes	0.5	0.16
5	B	Yes	0.7	0.51
6	C	Yes	0.4	0.1
7	C	No	0.4	0.1

For the straight pipe case A (Fig. 1, left) the weld seam model is positioned so that the influence on a fully developed turbulent pipe flow is investigated. For the T-junction cases (B, C), it is placed $4d_1$ downstream of the origin of the T-junction coordinate system, which is defined as the intersection point of both the main and branch pipe centerlines (Fig. 1, right). The working fluid is water at ambient temperature of 20°C and ambient pressure level for both the pipe flow and the isothermal T-junction setup (A, B).

In case of the buoyancy-affected T-junction flow (C), a temperature difference of $\Delta T = 100^\circ\text{C}$ at a pressure $p = 75\text{ bar}$ ($T_1 = 120^\circ\text{C}$, $T_2 = 20^\circ\text{C}$) is chosen for the numerical investigations. This is similar to the T-junction simulation by Kloeren [7] which is based on the hot experiment at the IKE. The experimental studies at the cold test-rig employ sugar water in order to realize a comparable density difference. Various volumetric flow rates Q_1 and Q_2 as well as flow rate ratios Q_1/Q_2 are employed. The numerical setups considered for the weld-seam discussion are summarized in Table 1.

An additional case with a temperature difference of 150°C (Case-No. 8) and without a weld seam is considered for the characterization of the hot T-junction flow. The same mass flow rates as case no. 7 (Table 1) are applied. The domain for the inlet flows, however, is extended in upstream direction.

3 Physical Model

Both the turbulent thermal mixing in a T-junction and the flow field affected by a weld seam involves highly anisotropic flow conditions. Furthermore, the unsteady flow field needs to be resolved in order to provide characteristic temperature fluctuations for fatigue analysis. Consequently, the main effort of CFD validation for the T-junction problem has been directed to LES methods or related hybrid turbulence models, such as Detached-Eddy Simulation or Scale-Adaptive Simulation [5].

The LES approach solves the energy-containing large-scale turbulence without turbulence modeling. Hence, the LES is always a three-dimensional and transient computation on a high quality mesh. The separation of turbulent scales is mathematically achieved by a filter operation [3], by which an additional term for the small-scale turbulence is introduced. These stress components $\tau_{ij,SGS}$ (1) are modeled by the so-called Subgrid-Scale (SGS) models. It is assumed that the turbulence becomes increasingly isotropic towards the dissipation scales and thus simple algebraic SGS models based on the mixing-length hypothesis are widely utilized.

For this work, the dynamic Smagorinsky SGS model by Germano [4] and Lilly [10] is chosen. With Reynolds numbers (Re) below 4×10^4 for all test cases the LES of wall bounded internal flows can be performed in reasonable computational time. The SGS shear stress $\tau_{ij,SGS}$ is modeled as follows:

$$\tau_{ij,SGS} = -2\mu_t(\bar{S}_{ij} - \frac{1}{3}\bar{S}_{kk}\delta_{ij}) \quad (1)$$

with the filtered strain rate tensor of the resolved scales

$$\bar{S}_{ij} = \frac{1}{2} \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) \quad (2)$$

In the Smagorinsky-Lilly model the subgrid-scale turbulent viscosity μ_t (3) depends on the strain rate tensor (2), the cell volume V and the model constant C_S . The linear function of the length scale containing the Kármán constant $\kappa = 0.41$ and the wall distance d ensures the correct wall behavior in the near-wall region. The filter length Δ is defined as $\Delta = V^{1/3}$.

$$\mu_t = \rho(\min(\kappa d, C_S \Delta))^2 \sqrt{2\bar{S}_{ij}\bar{S}_{ij}} \quad (3)$$

The local model coefficient C_S is dynamically calculated, assuming scale similarity. C_S is derived from the difference of the resolved solution for both the grid filter width Δ and the test filter width 2Δ . For numerical stability reasons the C_S value is limited in the range from 0 to 0.23.

The subgrid-scale turbulent heat flux $\tau_{\theta,i}$ is modelled as follows:

$$\tau_{\theta,i} = -\frac{\mu_t c_p}{Pr_{t,SGS}} \frac{\partial \bar{T}}{\partial x_i} \quad (4)$$

The subgrid-scale turbulent Prandtl number $Pr_{t,SGS}$ is dynamically calculated according to Lilly [10]. Scale similarity is assumed again and $Pr_{t,SGS}$ can be explicitly determined based on the resolved solution of the grid filter and the test filter.

In case the viscous sub-layer is not sufficiently resolved a wall function with a smooth interpolation between the sub-layer and the logarithmic wall region is applied. The blending function used for the buffer layer is similar to the formulation utilized by Kader [6].

4 Numerical Methods

In the commercial finite volume solver ANSYS FLUENT the SIMPLE algorithm is chosen for the velocity-pressure coupling. The second order central differencing scheme is applied for the convection terms for both the momentum and the energy equation. For the pressure interpolation the PRESTO scheme [12] is applied.

4.1 Boundary Conditions

The weld-seam experiments are conducted in a cold state in pipes made of PVC. For this reason, the simulations apply smooth adiabatic wall boundary conditions in all cases. The velocity inlet conditions are taken from RANS simulations of fully developed turbulent pipe flow. The outlet is set to pressure outlet conditions.

Although LES usually requires unsteady turbulent fluctuations at the inlet, sometimes the inlet turbulence might be insignificant compared to the turbulence produced by the flow instabilities, e.g. in the mixing zone of a T-junction. In the straight pipe flow cases, no such possibly dominating flow instabilities occur upstream of the weld seam model. For this reason, artificial inflow turbulence is modeled by the vortex method [11] for all the considered cases. RANS of fully developed pipe flows are performed with the Reynolds Stress Model (RSM). The random planar inlet fluctuation are emulated based on the profiles of the mean velocity, turbulent kinetic energy, turbulent dissipation rate and the three Reynolds normal stresses. Based on the inlet fluctuation the inlet flow has $3d_1$ to recover the proper turbulent field.


4.2 Time Discretization

Time steps from 1 to 2 ms are chosen with a Courant number less than unity in most of the domain. Courant numbers up to 3 can occur locally. However, due to the second order backward Euler time integration scheme this does not pose a threat

Table 2 Dimensionless size of the wall-adjacent cell volume according to best practice guidelines (BPG) listed in [3]

–	Δx^+	Δy^+	Δz^+
Sagaut [14]	10	2	5
Fröhlich [3]	50	2	15
Piomelli [13]	100	2	40
This study	< 50	< 3	< 15

Fig. 2 Numerical mesh close to the weld seam modeled as a half-circle (width = 8.3 mm, height = 4 mm)



to the numerical stability. Before recording the data for turbulent statistics the flow was given time to develop so that the flow has passed through the full geometry at least once, based on the bulk velocity. At least 8–10 s are calculated for the turbulent statistics.

4.3 Numerical Grid

The geometry and the local coordinate system are shown in Fig. 1. The computational domain extends $10d_1$ downstream of the inlet boundary for the single pipe case (A). The domain for the T-junction cases (B, C) extends from $3d_1$ upstream to $10d_1$ downstream of junction. For all setups the weld seam geometry is modeled with a half-circle profile (width = 8.3 mm, height = 4 mm) and placed $4d_1$ downstream of either the inlet boundary (A) or the intersection (B, C). For this problem the numerical grid contains ca. five million cells.

The numerical mesh is designed according to the best practice guidelines by Fröhlich [3] for a wall-resolved LES. The degree of restrictiveness for the wall adjacent control volume lies between the guidelines suggested by Sagaut [14] and Piomelli [13] as summarized by Fröhlich and shown Table 2. The dimensionless length (stream-wise), height (wall-normal) and width (transversal) of the first cell from the wall surface are denoted as Δx^+ , Δy^+ and Δz^+ respectively. The superscript “+” indicates a non-dimensional length scaled by the friction velocity u_τ and the kinematic viscosity ν . Occasionally, Δy^+ exceeds slightly the recommended value, so that it can still be considered a wall resolved LES. The mesh in the weld seam region is shown in Fig. 2.

Fig. 3 Pipe Flow: Mean value of normalized u_{mean} at $0.3/d_1$ downstream of the weld seam





Fig. 4 Pipe Flow: RMS value of normalized u_{RMS} at $0.3d_1$ downstream of the weld seam



5 Results


The arithmetic mean values and root mean square (RMS) values of the fluctuations are compared. For an entity A with a set of N data points they are defined as:

$$A_{mean} = \frac{\sum A_i}{N}, \quad A_{RMS} = \left(\frac{\sum (A_i - A_{mean})^2}{N} \right)^{1/2} \quad (5)$$

5.1 Pipe Flow

The simulation results are compared to PIV data in Figs. 3 and 4. Initially, the inlet mass flow rates for both the simulations and the experiments were set identically, based on the hydrostatic pressure in the water supply tank. Flow rate measurements based on PIV, however, proved to be more reliable in this case and the experimental mass flow rates are adjusted accordingly.

Fig. 5 T-junction: Mean value of normalized u_{mean} at $0.3/d_1$ downstream of the weld seam for different flow rate ratios Q_1/Q_2



The vertical profiles of the mean (Fig. 3) and RMS values (Fig. 4) of the velocity are compared at the axial position $x/d_1 = 0.3$ downstream of the weld seam. The velocities are scaled by the bulk velocity u_{bulk} . Both LES and the experiments show similarity of the profiles for varying volumetric flow rates Q . The recirculation bubble is indicated by the negative mean velocities close to the wall.

The RMS-values have a minimum in the middle of the pipe and increase towards the walls up to 35 % of u_{bulk} . The reference bulk velocity is calculated by the local one-dimensional vertical velocity profile in order to be consistent with the normalization methods of the experimental data.

While the core flow profiles overlap very well, the recirculation of flow is overestimated in the simulation. Furthermore, the near-wall magnitude of the velocity RMS for flows with similar flow rate is larger for the experimental data. This corresponds to the slightly steeper velocity gradients in the mean velocity profile.

5.2 Non-stratified T-junction Flow

All cases with a weld seam installed display a flattened core mean velocity profile (Fig. 5). The flat, slightly M-shaped core profile is captured well by the simulation. The M-shape resembles a wake flow since the intersecting branch pipe inflow acts as an obstacle for the main water stream. For low flow rates the profile tends to lean towards the bottom of the pipe. The blue and the red simulations display a strong similarity of their mean velocity profiles despite the difference between the absolute values of the inlet flow rates as well as the flow rate ratios ($Q_1/Q_2 = 4/1$ and $Q_1/Q_2 = 2/1$).

For the LES results small recirculation bubbles can be identified on both sides of the wall but not for the experimental results. The peaks of the RMS values (Fig. 6) of the lower flow rates (Case 3, red line) are underestimated compared to

Fig. 6 T-junction: RMS value of normalized u_{mean} at $0.3d_1$ downstream of the weld seam for different flow rate ratios Q_1/Q_2





Fig. 7 Stratified T-junction flow with weld seam (experiment) – Mean value (left) and RMS value (right) near-wall equivalent temperature field are shown. The blue bar indicates the weld seam. The main flow direction is from left to right

the experiment (black dots). The underestimation can be up to 25 %. For larger flow rates (green line, blue “x”) the agreement is very good. The flat profile of the RMS values of the core flow, however, is well reproduced as well the location of the RMS peaks.

5.3 Stratified T-junction Flow

For the stratified case first measurements of the new so-called Near-Wall-LED-Induced-Fluorescence (NW-LED-IF) method [9] are shown. This technique is based on the use of dissolved fluorescent dyes with temperature-dependent properties, which are excited by a green LED light source. It allows the measurement of unsteady temperature or density fields within 1 mm distance from the wall shown in Fig. 7. Sugar was dissolved in the water of the branch line so that the density fractions of sugar-water mixing and hot water mixing ($T_1 = 120^\circ\text{C}$, $T_2 = 20^\circ\text{C}$) are similar.




Fig. 8 Stratified T-junction cases with weld seam (LES) – Mean value (*left*) and RMS value (*right*) of the near-wall temperature field are shown. The *white vertical bar* in the middle of each field indicates the weld seam. The whole pipe is shown in the vertical direction. The main flow direction is from *left* to *right*

So far, only a qualitative comparison can be considered because of the difference of the absolute viscosity and viscosity ratio as well as the different volume flow rates for the weld seam case. However, a wavy character of the stratified flow field, as described by Kloeren [7] and indicated in Fig. 8 is also observed for the experimental case.

Furthermore, a similar flow situation is found close to the weld seam (Figs. 3–6), in which the stratification is shifted downwards (upstream) and upwards (downstream) due to the presence of the weld seam. This leads to increased mixing and fluctuations in the upstream region in the vicinity of the weld. Compared to the simulation, the experiment shows a rather stretched region of increased RMS values, parallel to the weld line.

The weld-seam geometry in combination with the recirculation bubbles up- and downstream of the weld seam contract the effective inner diameter and accelerates the core flow. Due to the wave-like character of the present stratification with the related circumferential velocity components the reduced inner diameter causes an increased rotation in order to maintain the angular momentum of the core flow.

For the same reason a counter rotational component in the recirculation bubbles is induced to balance the global angular momentum. Figure 9 shows the instantaneous tangential velocity in the cross-section (ca. $0.1d_1$ downstream of the weld seam). The counter-current rotation between the core flow and near-wall recirculation area is indicated by the arrows. The tangential velocities can almost reach the bulk velocities. The stably stratified flow experiences a sudden shift by these counter current angular flows. The increased shear flow and the non-stable temperature gradients lead to increased near-wall thermal mixing.




Fig. 9 Instantaneous tangential velocity field in the cross section just shortly downstream of the weld seam




Fig. 10 Mean temperature distribution for temperature differences of 100 °C (Case 7, top) and 150 °C (Case 8, bottom). The arrows in the xz-plane (left) and the yz-plane (right) indicate the inflow of the hot and the cold flow, respectively. Red indicates the maximum temperature and blue is 20 °C

5.4 Onset of Upstream Flow for Large Temperature Differences

First experimental and numerical investigations of the hot T-junction setup indicate a upstream flow of cold fluid into the main pipe. Knowledge about the onset and development of the upstream flow is important in order to conduct measurements of inlet flow conditions.

The results of the LES of Cases 7 and 8 are shown in Fig. 10 and shows that a temperature difference of 150 °C induces an upstream flow in the main pipe of more than $3d_1$ (Fig. 10, left). An elongated recirculation area is formed which extends from the branch pipe to the tip of the cold flow upstream. A similar situation can be identified in the branch pipe. For 150 °C hot fluid can be expected around $3d_2$ upstream and has to be considered for optical measurements in the branch pipe. Additionally, LES of T-junction flow with high temperature differences of 260 °C have been performed as well. However, a domain with an inlet pipe of $25d_1$

upstream of the junction proved to insufficient to ensure a negligible upstream effect on the flow field near the inlet boundary condition. Although simulation results can not be considered valid under those conditions the extend of the upstream flow does not contradict experimental observations.

6 Computational effort

The scalability for the thermal T-junction mixing flow case with weld seam is shown in Fig. 11. Since only multiples of nodes have to be used (one to eight cores per node) the scaling performance is analyzed against the number of nodes instead of number of processors. The speed-up (Fig. 11, left) indicates that using half of the available cores scales better than using all of the eight cores, so that the average time per iteration (Fig. 11, right) is slightly smaller for 15 nodes with 60 instead of 120 cores. However, the savings in average wall clock time per iterations between 40 cores on 5 nodes and 60 cores on 15 nodes are minimal. For this test case the scaling performance for 8 cores per node behaves well up to 5–10 nodes. The meshes are partitioned automatically and the load is equally distributed on the available cores. The data files are written serially. To improve the simulation time the residuals are printed out only once per time step.

For the T-junction simulations considered in the weld seam investigation usually 40 cores were used for a single simulation (8 CPU's per node). For this mesh of ca. five million cells a 24 h run results in a simulation time advancement of ca. 4 s. In average, each computation require ca. 1×10^4 CPU hours. The straight pipe flows grids consists of ca. three million cells and 2–5 nodes were generally used.

Additionally, RANS simulations with the Reynolds Stress Models are calculated in order to provide the inflow profiles for all flow rates of both the main and the branch pipe as well as for all different temperatures for the main inlet flow. These steady state simulations are very small compared to the LES and require ca. 100 CPU hours each.

In order to study the upstream flow conditions for main inlet temperatures of $T_1 \geq 170^\circ\text{C}$ the high temperature T-junction case features a larger mesh with ca. 12 million nodes. These simulations are performed initially with adiabatic wall conditions.

Furthermore, a LES with a refined mesh was performed, which was not discussed above, based on the thermally coupled T-junction flow simulation described in [7]. This mesh ensures a y^+ value of less than unity and contains ca. 19 million nodes. The total physical time for the fine grid solution is 30 s. In average 12 nodes (96 CPU's) were used for this large simulation. A detailed investigation of the grid-refinement study will be conducted, especially in respect to the heat transfer. The computational effort for this particular simulation is ca. 2.6×10^4 CPU hours. The initial values are interpolated from an instantaneous solution on a coarser grid. This initial simulation was run for more than 150 physical seconds to provide a near statistically steady condition for the evaluation of unsteady thermal interaction.




Fig. 11 Scalability for the weld seam simulations of ca. five million cells. Each node features 8 cores of which either 4 (red) or 8 cores (blue) are utilized. The speed-up (left) is based on the solution on 2 nodes. The average wall clock time per iteration is shown for 1,500 iterations

Particularly, the steel pipe wall requires a very long time to establish a proper temperature field determined by the complex flow field, which cannot be achieved with steady state simulations.

7 Conclusion

Large Eddy Simulations of straight pipe flow, isothermal T-junction mixing flow and stratified T-junction flows influenced by a weld seam are performed. The LES data are validated against experimental data of a straight pipe flow and an iso-thermal, not stratified T-junction mixing case.

A new experimental setup for the investigation of the weld seam on the flow field is introduced. PIV measurements and so called Near-Wall-LED-Induced-Fluorescence technique designed for highly pressurized and pipe flow experiments with large temperature differences are presented.

The overall agreement between LES, using a high quality mesh, and the experiment for the isothermal cases is good. The similarity of the profiles in regard to the volumetric flow rates has been shown both in the simulation and the experiments. However, the LES with the dynamic Smagorinsky model slightly underestimates the peaks in the RMS, induced by the weld seam.

The influence on the near-wall scalar showed that LES is capable to reproduce the complex flow condition of wavy stable stratification disturbed by weld seam. It includes flow separation and the sudden rotational shift of stratification layer. This leads to narrow regions with enhanced near-wall scalar fluctuations close to the weld seam, which is consistent with characteristic thermal fatigue pattern in pipe configurations.

This effort is part of an ongoing project in which the experimental conditions and measurement technique will be refined and, eventually, the methods will be employed in the hot experiment in order to provide data for CFD validation.

Acknowledgements This study was conducted in the framework of the project “Thermal Fatigue - Basics of system-, outflow-, and material characteristics of piping under thermal fatigue” funded by the German Federal Ministry of Education and Research (BMBF, PTKA Project No. 02NUK009B). The numerical simulations were performed at the High Performance Computing Center (HLRS) in Stuttgart, Germany.

References

1. S. Chapuliot. Thermal fatigue in mixing areas: Overview of the industrial problem, 2009.
2. S. Chapuliot, C. Gourdin, T. Payen, J.P. Magnaud, and A. Monavon. Hydro-thermal-mechanical analysis of thermal fatigue in a mixing tee. *Nuclear Engineering and Design*, 235(5):575–596, 2005.
3. J. Fröhlich. *Large-Eddy Simulation turbulenter Strömungen*. B.G. Teubner Verlag, Wiesbaden, 2006.
4. M. Germano, U. Piomelli, P. Moin, and W.H. Cabot. A dynamic subgrid-scale viscosity model. *Physics of Fluids A*, 3(7):1760–1765, 1991.
5. W. Haase, M. Braza, A. Revell, F.R. Menter, and Y. Ergorov. *DESider - A European Effort on Hybrid RANS-LES Modelling*. Springer Verlag, Berlin, 2009.
6. B.A. Kader. Temperature and concentration profiles in fully turbulent boundary layers. *Int. J. Heat Mass Transfer*, 24(9):1541–1544, 1981.
7. D. Kloeren and E. Laurien. Large-eddy simulation of thermal mixing in a t-junction with conjugate heat transfer. In *Proceedings of The 14th International Topical Meeting on Nuclear Reactor Thermalhydraulics*, Toronto, Ontario, Canada, 25–30 September 2011.
8. M. Kuschewski, R. Kulenovic, and E. Laurien. Experimental setup for the investigation of fluid-structure interactions in a t-junction. In *Proceedings of The 14th International Topical Meeting on Nuclear Reactor Thermalhydraulics*, Toronto, Ontario, Canada, 25–30 September 2011.
9. M. Kuschewski, R. Kulenovic, and E. Laurien. Novel application of LED-induced fluorescence. In *Proceedings of the Annual Meeting of the German Association for Laser Anemometry*, Ilmenau, Germany, 25–30 September 2011.
10. D. K. Lilly. A proposed modification of the Germano subgrid-scale closure model. *Physics of Fluids A*, 4(3):633–635, 1992.
11. F. Mathey, D. Cokljat, J. P. Bertoglio, and E. Sergent. Specification of LES inlet boundary condition using vortex method. In Y. Hanjalić, I.K. and Nagano and M. Tummers, editors, *4th International Symposium on Turbulence, Heat and Mass Transfer*. Begell House, Inc., Antalya, Turkey, 2003.
12. S.V. Patankar. *Numerical Heat Transfer and Fluid Flows*. Hemisphere, Washington, DC, 2080.
13. U. Piomelli. Large-eddy simulations: Theory and applications. In M. Hallböck, D. Henningson, A. Johansson, and P. Alfredson, editors, *Turbulence and Transition Modelling*, pages 269–331. Kluwer Academic, 1996.
14. P. Sagaut. *Large-Eddy Simulation of Incompressible Flows*. Springer, Berlin, 2. edition edition, 2001.
15. J. Westin, F. Alavyoon, L. Andersson, P. Veber, M. Henriksson, and C. Andersson. Experiments and unsteady CFD-calculations of thermal mixing in a t-junction. In *Workshop Proceedings of Benchmarking of CFD Codes for Application to Nuclear Reactor Safety (CFD4NRS) 2006*, Munich, Germany, 5–7 September 2007. NEA.

Simulation of Compressible Viscous Flow with an Immersed Boundary Method

B. Jastrow and F. Magagnato

Abstract An immersed boundary method combined with a wall-layer approach has been implemented into an established flow solver. In the outer flow field, the compressible Navier-Stokes equations are solved using an approximate Riemann Solver whereas simplified boundary-layer equations are solved near the wall. Turbulence is accounted for by the one-equation model of Spalart-Allmaras in the outer flow region and by a mixing length eddy viscosity model with near wall damping in the wall layer. Computations performed for various test cases show good agreement with reference data found in literature.

1 Introduction

Block-structured, body-fitted grid generation for simulating flow in complex geometry can be very tedious. The Cartesian-grid immersed boundary method (IBM) offers an interesting approach since automatic mesh generation can be realized easily. Furthermore, due to smoothness and orthogonality, Cartesian grids offer high accuracy and efficiency. In the IBM a complex geometry is immersed into a regular Cartesian grid. The effect of the body on the flow is mimicked by the imposition of proper boundary conditions that act as forcing conditions [1, 2]. Several applications of the IBM use a linear interpolation for setting the boundary conditions as proposed in [1]. Their validity is therefore only given for grids resolved down to $\Delta y^+ < 1$. For flow of high Reynolds number these restrictions cannot be held in accordance with acceptable computing time. With the use of a wall-layer model one can overcome the need for high near wall resolution. Wall models based on incompressible turbulent boundary layer equations were proposed

B. Jastrow (✉) · F. Magagnato
Department of Fluid Machinery, University of Karlsruhe, Kaiserstr. 12, 76131 Karlsruhe,
Germany
e-mail: jastrow@kit.edu; franco.magagnato@kit.edu

and tested by Balaras et al. [3], Cabot and Moin [4], Wang and Moin [5] in a body-fitted context. The applicability of such a wall-layer model in the framework of the IBM has been studied by Tessicini et al. [6]. Bond et al. [7] developed a compressible wall model, named the diffusion model. The present report studies the compressible wall-layer model within the framework of the IBM. The consideration of the energy equation allows a temperature distribution in the wall layer which in turn influences the thermodynamic quantities. It is shown, that the implementation in the in-house flow solver SPARC (Structured Parallel Research Code) [8], in the following referred to as SPARC-IBM, provides results for basic test cases in good agreement with literature.

2 Mesh Generation

The mesh is automatically generated via a ray tracing technique. Based on a geometry described by a closed surface triangulation every cell center of a uniform Cartesian grid is marked as internal or external [9]. Those internal cells having at least one external neighbor cell are marked as wall-layer cells. From the latter the normal to the closest wall is computed and stored together with the forcing point and the appropriate interpolation neighbors (Fig. 1).

3 Numerical Method


The Reynolds-averaged Navier-Stokes (RANS) equations for compressible flow together with the Spalart-Allmaras turbulence model are used for the calculations presented in this report. The discretization of the convective terms is obtained by an approximate Riemann solver (HLLC) [10] whereas the viscous terms are discretized with a central difference scheme. The solution vector is updated via a Runge-Kutta scheme.

Inside the wall-layer a simplified set of equations is solved on an embedded grid with a minimum resolution of 40 cells. It is assumed that convection is negligible compared to diffusion, the normal pressure gradient is zero and the streamwise gradients are of orders lower than the gradients normal to the wall. Furthermore, the normal velocity is insignificant compared to the tangential velocity. This leads to the following equations for the x-momentum and the energy:

$$\frac{\partial(\rho u)}{\partial t} = \frac{\partial[(\mu + \mu_t) \frac{\partial u}{\partial y} - y \frac{\partial p}{\partial x}]}{\partial y} \quad (1)$$

$$\frac{\partial(\rho e_t)}{\partial t} = \frac{\partial[u(\mu + \mu_t) \frac{\partial u}{\partial y} + (\kappa + \kappa_t) \frac{\partial T}{\partial y}]}{\partial y}. \quad (2)$$

Fig. 1 Mesh in the near wall region: o : internal cell, Δ : wall-layer cell, x : external cell, f : forcing point, i : interpolation neighbors



The properties μ_t and κ_t are the turbulent viscosity and the turbulent conductivity respectively. In a first stage of implementation, μ_t is obtained by a mixing length eddy viscosity model with near wall damping

$$\frac{\mu_t}{\mu} = \kappa y^+ (1 - e^{-y^+/A})^2 \quad (3)$$

where $\kappa = 0.4$, $A = 19$ and y^+ defines the dimensionless distance to the wall [5]. In the case of laminar flow, the turbulent viscosity is set to zero. At the wall, a no-slip boundary condition is applied. At the forcing point, that defines the outer edge of the wall-layer, the boundary condition is obtained through an interpolation from the flow values of the outer flow field. The equations are further simplified by neglecting the time derivative terms and the streamwise pressure gradient. The wall shear stress needed for the calculation of y^+ is evaluated from the velocity gradient at the wall. Two tridiagonal systems are created for the velocity and the temperature respectively and solved in a segregated way with the other dependent variables held constant. Subsequently the turbulent viscosity is updated. Since the momentum equation and the energy equation are coupled via the viscosity and the velocity, an outer loop runs until the full solution converges. Finally the flow values for the position of the center of the wall-layer cell are extracted and provided to the outer flow field as a boundary conditions. Assuming steady-state in the wall-layer corresponds to an instantaneous response of the wall-layer to the outer flow field which creates some error for unsteady flow calculations.

4 Results

4.1 Flow Past a Flat Plate

The first test case is the flow past a flat plate. The wall-layer model is capable of capturing both, the linear and the turbulent near-wall behavior. The computation

Fig. 2 Comparison of the velocity profile in the boundary layer at $Re_x = 10,000$





Fig. 3 Comparison of the velocity profile in the boundary layer at $Re_x = 10,000,000$



was carried out on a structured multi-block mesh with refinement. The boundary layer was resolved with 32 cells making a total of 54,000 cells in the 2D plane for the laminar case and resolved with 27 cells making a total of 123,000 cells for the turbulent case respectively. The Mach number was chosen to $Ma_\infty = 0.3$. The laminar velocity profile was extracted at a Reynolds number of $Re_x = 10,000$. In Fig. 2 the profile is compared with the analytical results of the Blasius-equation [11] where $f' = \frac{U}{U_\infty}$ is the dimensionless velocity and $\eta = y \sqrt{\frac{U_\infty}{\nu x}}$ is the dimensionless wall distance. The numerical results show very good agreement with the analytic solution of Blasius.

Figure 3 shows the obtained velocity profile for $Re_x = 10,000,000$ together with the experimental data of [12]. The agreement for the turbulent calculation is quite satisfactory as well.




Fig. 4 Streamlines for $Re = 20$

Table 1 Separation length L , Separation angle θ_S and drag coefficient c_D for $Re = 20$

	L	θ_S	c_D
Sucker and Brauer [13]	0.83	43.3	2.02
Dennis and Chang [14]	0.94	43.7	2.05
SPARC-IBM	0.93	41.8	2.09




Fig. 5 Instantaneous contours of the vorticity for $Re = 200$

4.2 Laminar Flow Past a Circular Cylinder

To test the IBM together with curved surfaces, the flow past a circular cylinder was computed. Two flow regimes are shown, a steady flow at $Re = 20$ and an unsteady flow at $Re = 200$. The mesh consists of 105,000 cells in the 2D plane with the near-wall and the wake region refined. The free stream Mach number in both cases was set to $Ma_\infty = 0.3$. The streamlines for the flow at $Re = 20$ are shown in Fig. 4 whereas Table 1 lists the separation length and angle and the drag coefficient in comparison with reference values from literature [13, 14]. The agreement is quite satisfactory, only the separation angle is predicted too low, which might be due to the lack of considering the pressure gradient in the wall-layer.

Figure 5 shows the instantaneous contours of the vorticity for an unsteady laminar flow around the cylinder at $Re = 200$. A comparison of the Strouhal number and the mean drag coefficient with results from literature [15, 16] is provided in Table 2.

Table 2 Strouhal number St and mean drag coefficient c_D for $Re = 200$

	St	c_D
Linnick and Fasel [15]	0.197	1.34
Liu et al. [16]	0.192	1.31
SPARC-IBM	0.193	1.38

5 Computational Efficiency

In the computation of these results we have been using up to 1,344 AMD-Interlagos cores of the CRAY XE6 (Hermit) at HLRS in Stuttgart. The in-house developed code Sparc is parallized with the MPI-2 software. The computational time for one unsteady calculation in three dimesions using about 10 mio points was about 24 h. Since we were using more than 10,000 blocks of the finite volume scheme we could efficiently distribute the blocks on these 1,344 cores with the domain decomposition technique. The load balancing was at about 99 communication is done with the CRAY GEMINI node-node interconnection the parallel efficiency was close to 95 %. This very good parallel effieience could only be obtained because we have reduced the amount of output to a minimum. It is desired that the amount of output data should be increased so we need to improve the read/write performance of the code in the near future. From our recent investigations we know that a higher resolution of the computational mesh is required. We think that using about 80 million points in the next phase will be adequate for a well resolved unsteady calculation.

6 Conclusion and Ongoing Work

The immersed boundary method (IBM) with a wall-layer approach has been implemented into an established block-structured code for solving compressible flow. The computation of basic test cases for laminar flow showed very good agreement with results found in literature. Viscosity dominated turbulent flow like the flow past a flat plate was also captured well. To overcome the lack of consistency, future work comprises the testing of the simplified Spalart-Allmaras turbulence model for the wall-layer. It is believed that in contrast to using the algebraic model, the full coupling of the outer flow field and the wall-layer provides better results for turbulent flows. Furthermore, the streamwise pressure gradient in the wall-layer and a flow adapted mesh refinement is about to be implemented. Since the advantages of the IBM are only given for the simulation of complex geometries where the time for mesh generation is significant, the next step is to validate the implementation for 3D cases. An example for a complex geometry is given in Fig. 6.




Fig. 6 Preliminary solution of the flow field around a complex geometry

References

1. Fadlun E A, Verzicco R, Orlandi P, Mohd-Yusof, J. Combined immersed-boundary finite difference methods for three-dimensional complex flow simulations. *J. Comp. Phys.*, 161, 35–60, 2000.
2. Mohd-Yusof, J. Combined immersed-boundary/B-spline methods for Simulations of flow in complex geometries. *Annual Research Briefs, Center of Turbulence Research*, 317–328, 1997.
3. Balaras, E et al. Two-layer approximate boundary conditions for large-eddy simulations. *AIAA Journal*, 34, pp. 1111–1119, 1996.
4. Cabot, W, Moin, P. Approximate wall boundary conditions in the large-eddy simulation of high Reynolds number flow. *Flow, Turbulence and Combustion*, 63, pp. 269–291, 1999.
5. Wang, M, Moin, P. Dynamic wall modeling for LES of complex turbulent flows. *Physics of Fluids*, 14, pp. 2043–2051, 2002.
6. Tessicini, F, Iaccarino, G, Fatica, M, Wang, M, Verzicco, R. Wall modeling for large-eddy simulation using an immersed boundary method. *Annual Research Briefs, Center of Turbulence Research*, pp. 181–187, 2002.
7. Bond, R B, Blottner, F G. Derivation, implementation, and initial testing of a compressible wall-layer model. *Int. J. Numer. Meth. Fluids*, 66, pp. 1183–1206, 2011.
8. Magagnato, F. KAPPA Karlsruhe Parallel Program for Aerodynamics, *TASK Quarterly Vol. 2*, pp. 215–270, 1998.
9. Verzicco, R, Iaccarino, G. *Immersed Boundary Technique for Large-Eddy-Simulation. Lecture series on Large Eddy Simulation and related techniques: Theory and Applications*, 2006.
10. Toro, E F, Spruce, M, Speares, M. Restoration of the contact surface in the HLL-Riemann solver, *Shock Waves*, 4, pp. 25–34, 1994.

11. Blasius, H. Grenzschichten in Flüssigkeiten mit kleiner Reibung. *Z. Math. Physik*, 56, pp. 1–37, 1908.
12. Wieghardt, K, Tillman, W. On the Turbulent Friction Layer for Rising Pressure. NACA TM-1314, 1951.
13. Sucker, D, Brauer, H. Fluidodynamik bei quer angestromten Zylindern. *Waerme- und Stoffuebertragung*, 8, pp. 149–158, 1975.
14. Dennis, S C R, Chang, G Z. Numerical solutions for steady flow past a circular cylinder at reynolds numbers up to 100. *Journal of Fluid Mechanics*, 42, pp. 471–489, 1970.
15. Linnick, M N, Fasel, H F. A high-order immersed boundary method for unsteady incompressible flow calculations. *J. Comp. Phys.*, 204, pp. 157–192, 2005.
16. Liu, C, Zheng, X, Sung, C H. Preconditioned Multigrid Methods for Unsteady Incompressible Flows. *J. Comp. Phys.*, 139, pp. 35–57, 1998.

Numerical and Experimental Examination of Shock Control Bump Flow Physics

K. Nübler, S.P. Colliss, T. Lutz, H. Babinsky, and E. Krämer

Abstract A method allowing a detailed investigation of the flow physics of shock control bumps (SCBs) on an unswept airfoil has been developed by comparison of the results of experiments and computations. A simple wind tunnel set-up is proposed which is shown to generate representative baseline conditions, allowing fine details of the flow to be measured using an array of techniques. Computational data for the same bump configuration is then validated against the experimental results, allowing a more intimate analysis of the flow physics as well as relating wind tunnel results to the performance of the SCB on an unswept wing.

1 Introduction

Shock control bumps (SCBs) are a promising method to reduce wave drag on transonic wings by inducing a bifurcated shock structure and thus decelerating the flow more isentropically [1]. However, the interaction between the modified surface contour, the wing boundary layer and the system of shocks has been shown to lead to a complex flow field [4, 12–14]. In order to find improved SCB geometries, specifically those which are more robust towards changing free stream conditions, the flow physics need to be better understood [3]. To pursue this target, a joint study is in progress analysing the bump flow both in a supersonic wind tunnel and using RANS-based CFD.

CFD can simulate the whole wing with bumps [2], but may struggle in resolving the small-scale flow physics and the separation prediction requires validation.

K. Nübler (✉) · T. Lutz · E. Krämer
IAG Universität Stuttgart, Stuttgart, Germany
e-mail: nuebler@iag.uni-stuttgart.de

S.P. Colliss · H. Babinsky
Aerodynamics Laboratory, University of Cambridge, Cambridge, UK

Fig. 1 Region of interest for SCB experiments: (a) modelled by wind tunnel experiments; (b) on a typical airfoil





Fig. 2 Comparison of typical pressure gradients on airfoils and in wind tunnels in the absence of a bump



However, once a reliable CFD setup is found, a quick analysis with variable free-stream conditions is possible, allowing shape optimisations or other studies with a high number of required cycles.

For experimental investigations, a blow-down type supersonic tunnel is preferred as runs are relatively inexpensive and repeatable [4]. However, in this method only the flow local to the bump is being considered and the object boundary layer is that grown naturally through the wind tunnel. Although this is a proven technique for ensuring that the Reynolds number based on displacement thickness is sufficiently high, certain other considerations bring the validity of such experiments into question. These are mainly due to the changing curvature of the nozzle in the wind tunnel, which contrasts with the continuous convex surface of the airfoil, shown in Fig. 1.

This is expected to not only have an impact on the development of the boundary layer, but also modify the surface pressure distribution that the SCB is subjected to. A comparison of typical pressure distributions for each situation (Fig. 2) indicates that the post-shock adverse pressure gradient is higher for the airfoil than the wind tunnel. It has been suggested that the flow here is of high importance for the bump performance [15], and thus the question of validity will be addressed in the present work.

2 Methods

2.1 Numerical Setup

Solver

Numerical analyses are performed using FLOWer [7], a well-established structured and density-based RANS solver developed by the German Aerospace Centre (DLR). The two-equation Shear Stress Transport $k-\omega$ [6] turbulence model was used. Except for the finest grid in the grid convergence study, no unsteady computations were necessary within this work.

Governing Equations

FLOWer solves the three-dimensional Reynolds Averages Navier-Stokes equations (RANS) in integral form

$$\frac{\partial}{\partial t} \int_V \mathbf{W} dv + \int_{\partial V} \mathbf{F} \cdot \mathbf{n} ds = 0 \tag{1}$$

with the vector of the conservative variables

$$\mathbf{W} = [\rho, \rho u, \rho v, \rho w, \rho E]^T \tag{2}$$

on block-structured finite volume meshes. The conservative variables are given in a Cartesian coordinate system with ρ, u, v, w, E denoting the density, the Cartesian velocity components of the velocity vector \mathbf{v} and the specific total energy, respectively. V represents the control volume and ∂V its closed outer surface. The flux tensor \mathbf{F} is divided in a convective inviscid part \mathbf{F}^c and a viscous part \mathbf{F}^v such that

$$\mathbf{F} = \mathbf{F}^c - \mathbf{F}^v \tag{3}$$

with

$$\mathbf{F}^c = \begin{bmatrix} \rho u & \rho v & \rho w \\ \rho u^2 + p & \rho uv & \rho uw \\ \rho uv & \rho v^2 + p & \rho vw \\ \rho uw & \rho vw & \rho w^2 + p \\ \rho uE + up & \rho vE + vp & \rho wE + wp \end{bmatrix} \quad \text{and} \quad \mathbf{F}^v = \begin{bmatrix} 0 & 0 & 0 \\ \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \\ \psi_x & \psi_y & \psi_z \end{bmatrix} \tag{4}$$

The ψ_i are abbreviations of the type

$$\psi_i = \left(u\sigma_{ix} + v\sigma_{iy} + w\sigma_{iz} + K \frac{\partial T}{\partial x_i} \right) \quad \text{for } i = x, y, z \tag{5}$$

The pressure p is calculated by the equation of state of the perfect gas

$$p = (\gamma - 1)\rho \left(E - \frac{u^2 + v^2 + w^2}{2} \right) \quad (6)$$

with the specific heats ratio γ . The temperature T is defined by

$$T = \frac{p}{\rho R}. \quad (7)$$

The Reynolds stress tensor σ_{ij} in (4) and (5), that represents correlations between fluctuating velocities and which depends on the fluid viscosity, is given through

$$\sigma_{ij} = \mu(\mathbf{v}_{i,j} + \mathbf{v}_{j,i} - \frac{2}{3}\delta_{ij}\mathbf{v}_{k,k}) \quad (8)$$

where $(\cdot)_{i,j}$ denotes derivation of the i -th component with respect to x_j . Likewise the heat conductivity K in (5) depends on the viscosity μ , as is shown by the relation

$$K = \frac{\gamma}{\gamma - 1} \frac{\mu}{\text{Pr}} \quad (9)$$

with Pr being the Prandtl number. For laminar flow μ in (8) and (9) is set to $\mu = \mu_L$, which is according to the Sutherland law

$$\mu_L = \mu_0 \left(\frac{T}{T_\infty} \right)^{3/2} \frac{T_\infty + S}{T + S}. \quad (10)$$

where $\mu_0 = 1.716 \times 10^{-5}$ kg/(ms) and $S = 110.4$ K.

Meshing and Acceleration

The airfoil chosen for this investigation is ‘pathfinder’, a low camber transonic airfoil designed for research in the field of natural laminar flow [8]. The structured meshing is script-based ensuring consistently high quality meshes for each SCB configuration, thus reducing the grid influence on the solution to a minimum. All calculations performed within this work (except for the grid convergence study) were carried out on a C-grid with 5.8 million cells in total. This consists of 480 cells around the airfoil, of which 130 are in the wake, and 100 cells in the wall-normal direction. The far-field extends to 50 chord lengths from the airfoil, and the grid is refined in the boundary layer so that the nearest cell to the wall corresponds to $Y^+ \approx 1$. Different widths of computational domain, from 0.1 to 0.3 c , were examined for comparisons between experimental and computational results. The number of spanwise cells was 120, which was found to be suitable for the highest

span. Periodic boundary conditions are employed so that the set-up represents an unswept wing of infinite aspect ratio. A single SCB is applied at centre-span, such that the effective bump spacing is set by the width of the computational domain.

For calculations aiming for a target lift value, an angle of attack-update is performed every 1,000 iterations. Figure 3 shows such a case with a target lift value of $C_L = 0.405$. The peaks in residual at 2,000 and 5,000 iterations are caused by switching to a finer mesh; this three-stage multi-grid approach was adopted to accelerate convergence. Convergence is then continued up to a residual of 10^{-6} , or up to 9,000 iterations if this target is missed.

Grid Convergence Study

A grid convergence study was conducted to find a good compromise between computational cost and simulation accuracy. The cell distribution as described in the section above represents level 1 as shown in Table 1. The number of cells in each spatial direction was doubled or halved, respectively, to get the next level. The Y^+ value was kept constant for all meshes at about 1 for $Re = 20 \times 10^6$. The differing number of cells required in the wall normal direction was achieved through varying growth rates.

A contour bump was applied to a pathfinder airfoil section of width 0.3 c; the exact details of placement are described below (see also [2]). This case is considered critical because the flow is very close to separation after the normal shock on the bump crest. Moreover it was considered a good test case to qualitatively assess the prediction of flow features on the bump. For level 0, URANS with 100 inner iterations was required for convergence reasons resulting from the weak numerical damping. A quasi-steady solution with good convergence and coefficients comparable with the steady case was established.

The aerodynamic coefficients in Table 2 show that a convergence in drag prediction is achieved for level 1 and only a small discrepancy remains with respect to lift, whilst for levels 2 and 3 similar results are not achieved. The friction coefficient distribution shown in Fig. 4 provide further evidence that level 1 provides sufficient accuracy, and a similar conclusion could be drawn from the pressure distributions (not shown). The level 1 mesh is therefore chosen for further SCB investigations as a good compromise between accuracy and computational cost.

2.2 Wind Tunnel Setup

Experiments were performed in the No. 1 supersonic wind tunnel at Cambridge University Engineering Department, a blow-down tunnel with a rectangular working section 114 mm wide by 178 mm high. The stagnation pressure is controlled by a manually operated valve with accuracy typically better than 0.5 %, and run times

Fig. 3 Typical solution convergence history for presented setup with target lift on




Table 1 Grid size on levels for convergence study

Level	Approx. number of cells (million)
0	46.0
1	5.78
2	0.75
3	0.10

Table 2 Computed aerodynamic coefficients relative to level 0 (percentage deviations)

Level	C_L	C_D	ΔC_L	ΔC_D	$\Delta L/D$
0	0.424	0.0103	–	–	–
1	0.418	0.0103	–1.45	–0.42	–1.03
2	0.409	0.0109	–3.62	5.83	–8.93
3	0.364	0.0164	–14.3	58.58	–45.96

of up to 30 s are achievable. The nozzle geometry is fixed and set up to operate nominally at $M = 1.3$. The tunnel is equipped with an ejector system to enable a degree of control over the incoming boundary layer properties via either localised or distributed suction. A steady normal shock is positioned by a shock holding plate (following [4]) with accuracy of typically ± 1 mm. An additional post-shock adverse pressure gradient is imposed using a subsonic diffuser, the angle of which was chosen to be 3° based on an (inviscid) estimate of the additional pressure gradient that this would impose although this could be varied if it proved to be necessary. The wind tunnel geometry is shown in Fig. 5.

The flow is studied using a range of experimental techniques. Two-component laser Doppler anemometry (LDA) is used to measure both streamwise and wall-normal velocity components to within typically $\pm 0.5\%$. The flow is seeded with olive oil introduced into the flow in the settling chamber via a centreline seeding rake. Measurements were made on one side of the tunnel and reflected about the

Fig. 4 C_f over airfoil chord for different refinement levels





Fig. 5 Experimental set up: (a) photograph of wind tunnel, the *dashed circle* shows the region of optical access; (b) plan view of tunnel floor showing suction slot arrangement



centreline – the validity of doing this was checked periodically by performing symmetry checks. The surface pressure field is quantified to within $\pm 5\%$ using pressure sensitive paint (PSP), which is calibrated in situ against readings taken from a number of static pressure tappings using a 16 channel NetScanner 9116 series pressure transducer. Total pressure profiles were measured using a 15 hold pitot rake, with an accuracy of approximately 2%. This could be mounted at various spanwise locations in the tunnel. Additionally, total pressure profiles close to the bump were calculated from LDA and PSP data using the assumptions that the pressure and stagnation temperature are constant across the boundary layer. This method is accurate to within $\pm 5\%$.

Table 3 Free stream conditions for pathfinder airfoil meeting the wind tunnel flow properties

M	α	Transition position	Turb. model
0.76	1.8°	10 % chord length	SST k- ω

2.3 Adaptation of Methods

In order to produce relevant validation data, both the experimental and computational set-ups were adapted to produce baseline (no control) conditions which agreed. Several parameters were varied in the computations: free-stream Mach number and angle of attack controlled the inviscid velocity in the region of interest, whilst the transition location determined the incoming boundary layer properties. To reduce computational cost, the coupled Euler-boundary layer code MSES [10] was used, with the best candidate solutions being fine-tuned using FLOWer as detailed in Sect. 2.1. The final values of the parameters are given in Table 3.

Figure 6 shows a comparison of the incoming boundary layer profiles from experiment and CFD using the conditions outlined above, and the corresponding integral parameters are given in Table 4. These show good agreement, indicating that no modification to the wind tunnel boundary layer was required to achieve good agreement.

The surface pressure distributions for experiment and CFD are presented in Fig. 7. Again, good agreement is observed, confirming the choice of diffuser angle in the wind tunnel set-up. It is noteworthy that these results have been achieved for a combination of airfoil parameters which are representative of conditions on a real wing.

3 SCB Application: Validation and Results

3.1 SCB Geometry

A contour bump has been analysed using the methods described in the previous chapter. The longitudinal and lateral height profiles are given by elastic deflection formulae and shown in Fig. 8. Further details of the bump shape and discretisation may be found in [2]. The dimensions in Fig. 8 are scaled with the incoming boundary layer thickness δ_0 which was measured to 6 mm according to Table 4. In absolute dimensions, the bump is 150 mm long and 50 mm wide, with a crest height of 6 mm. The CFD airfoil chord length was set to 1 m.

Two shock positions were examined: ‘on-design’, where the shock is 75 mm from the nose, and ‘off-design’ where the shock is 20 mm further downstream. In the

Fig. 6 Boundary layer profile upstream of shock: CFD clean airfoil data and wind tunnel measurement with Sun and Childs [9] turbulent boundary profile fitted





Table 4 Boundary layer parameters upstream of the shock in the experiment and in the computation


	p_0 (kPa)	M_{shock}	δ_0 (mm)	δ_0^* (mm)
Exp.	170	1.30	6.0	0.67
CFD	173	1.30	7.0	0.67
	θ_0 (mm)	H_0	$Re_{\delta_0^*}$	
Exp.	0.53	1.28	27,650	
CFD	0.51	1.31	25,670	

Fig. 7 Pressure distribution across the shock with diffuser and in CFD



experiments this was achieved by moving the shock, whilst in CFD the bump was moved forwards on the airfoil. The presence of the SCB caused the shock to shift 5 mm upstream on the airfoil and therefore the bump had to be moved 25 mm to match the relative shock position in the wind tunnel. A first analysis of the flow features can be found in [15].

Fig. 8 Contour bump geometry in longitudinal (a) and lateral (b) plane



3.2 Validation

Figure 9 shows a comparison of computed surface streamlines and experimental surface flow visualisation results for both shock positions.¹ On design, a good agreement is observed between the experiment and CFD with only minor discrepancies, both with respect to the flow topology and the region of low shear stress over the SCB tail. The off-design separation topology, which may be classified as an ‘owl-face of the first kind’ [16], is correctly predicted in position by CFD, although is larger. This separation has been found to result from either positioning the bump too far forward on the wing (as was the case here) or from having too tall a bump.

Further favourable comparison may be found in the surface pressure fields presented in Fig. 10. The general shape of the field is seen to be similar in both cases, with respect to the shock footprint as well as the development of a low pressure region just ahead of the main shock on the bump crest. The centreline distributions confirm the similarity in shape, although it is seen that the initial compression is lower for the airfoil than the experiment. This is possibly caused by the underlying curvature of the airfoil which means that whilst the uncontrolled shock strength (at the rear leg of the SCB shock system) is matched, ahead of the bump the flow has not yet reached $M_\infty = 1.3$. Additionally the airfoil curvature is thought to reduce the effective ramp angle, further weakening the initial shock.

It was found that the width of the computational domain (which sets the effective bump spacing) had a noticeable impact on the level of agreement attained between experimental and computational results. The computational results presented in Fig. 9 were obtained for a domain width of $0.1c$, which is similar to the wind tunnel width. For a domain width of $0.3c$, the region of low shear visible on the tail in Fig. 9a was found to be a separation at the trailing edge. This suggests that there is an appreciable interaction between bumps if the spacing (centreline to centreline)

¹The perspective view in (b) is used as the oil accumulated at the foci in the wind tunnel experiment was smeared by tunnel shutdown, obscuring the topology. This image was taken from a high-definition video recorded during the tunnel run.

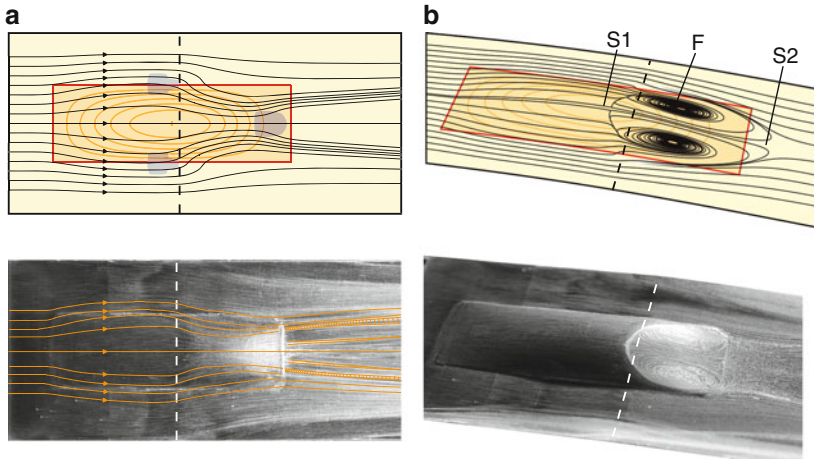


Fig. 9 Comparison of computed surface streamlines with experimental surface flow visualisation: (a) on-design, (b) off-design. The *dashed line* shows the shock position; in (a) the *blue regions* in the CFD image indicate low shear stress

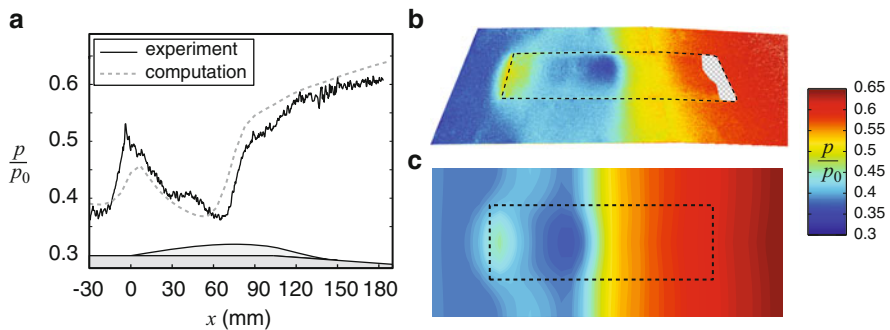


Fig. 10 Comparison of surface pressure fields obtained from experiment and CFD for on-design shock position: (a) centreline distribution; (b) experimental result obtained using pressure sensitive paint; (c) computational result

is two bump widths or less, in this case allowing a more aggressive tail design than would otherwise be possible without inducing separation.

A further point of note is that the separation in the off-design case is better predicted in this case than the corresponding results presented by Nübler et al. [11]. It is thought that this may have been influenced by the increased post-shock adverse pressure gradient imposed by the new wind tunnel set-up, which was not present for the experiments quoted in that paper.

In summary, a good agreement between computational and experimental results is achieved, which confirms that the method of adapting both wind tunnel and

CFD to model the same flow but with different approaches is viable and produces satisfactory and useful results.

3.3 Some Elements of SCB Physics

The validated computational data can be used to improve understanding of flow structures found by wind tunnel studies for which certain parameters would be difficult to compute from the experimental data. The following section presents some preliminary results of an ongoing study on the effects of the tail geometry.

Experimental measurements of total pressure in the wake are presented in Fig. 11. These show that just downstream of the bump, there is a high total pressure loss relative to the no control case, and thus using just these measurements (as has traditionally been done in experiments, [5, 12]) it would be concluded that the SCB incurs a severe penalty in the boundary layer. However, further downstream (towards the end of the diffuser) the total pressure is seen to have recovered such that there is a net *gain* in the lower boundary layer. Using these results, the conclusions on SCB viscous penalty would therefore be the opposite to before.

Similar behaviour is evident in the computational solution. Figure 12 shows the development of the total pressure deficit, defined as

$$\Pi_{\text{SCB}}(x) = \frac{1}{A(x)} \int_{A(x)} \frac{p_0 - p_{\text{total}}(x, y, z)}{p_0} dydz$$

where A is the area spanning the height of the computational domain and the width of the bump. A region of clear improvement in the deficit (evidenced by a decreasing value of Π) is seen which persists until around 100 mm downstream of the bump. Further downstream the deficit increases again, and this is thought to be due to the boundary layer growth in the increasing adverse pressure gradient towards the trailing edge of the wing.

It is thought that this is brought about by streamwise vorticity, [15], the existence of which was postulated from velocity measurements in a plane 20 mm downstream of the SCB (reproduced in Fig. 13) which show a region of clear downwash. Experimental determination of vorticity was not possible due to limited optical access preventing measurement of the spanwise component of velocity. However, the presence of this structure is confirmed by the computational results, also shown in Fig. 13, where clear regions of vorticity are apparent.

The vortical structure does not represent vortices in the traditional sense, but is a simple spanwise variation in the flow direction. That this effect can bring about the total pressure recovery is encouraging because, unlike with ‘traditional’ vortices, there is no corresponding upwash region and therefore the presence of a zone of increased total pressure deficit is minimised.

Figure 14 shows how the magnitude of the vorticity varies in the streamwise direction via the circulation, defined as

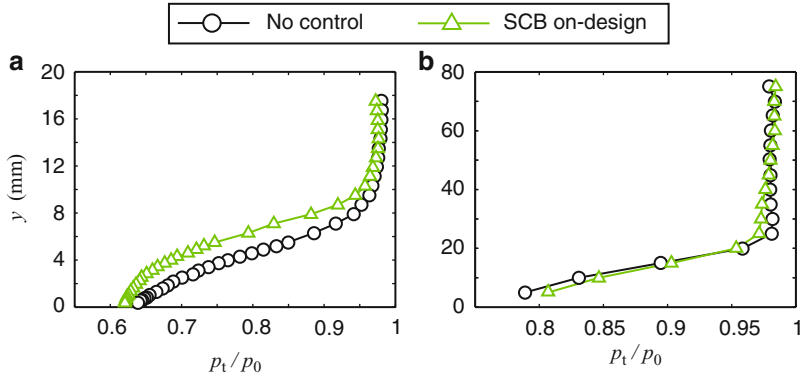



Fig. 11 Total pressure measurements downstream of SCB in wind tunnel: (a) just downstream of the bump at $x = 170$ mm; (b) far downstream at $x = 485$ mm, corresponding to the wing trailing edge

Fig. 12 Total pressure deficit variation in SCB wake – note that $x = 150$ mm corresponds to the back of the bump



$$\Gamma(x) = \int_{A(x)} \omega \cdot d^2\mathbf{x}$$

The area $A(x)$ in this case is the semi-span of the computational domain. This is non-dimensionalised against the skin friction velocity u_τ ($\equiv \sqrt{\tau_w/\rho}$) and h , the height of the bump, following a similar treatment by Ashill et al. [17] to enable the strength of the vorticity generated by the bump to be compared to that of vortex generators.

Two peaks of circulation are observed, which correspond to vorticity generated by the curved shock structure; this curvature is evident in the surface pressure results presented in Fig. 10. These represent the peaks of vorticity found anywhere in the flow, although they are seen to be relatively transient, decaying quickly downstream. A particularly interesting feature is the increase in circulation between $x = 100$ and 225 mm, suggesting that the vorticity observed in the wake which is responsible for

Fig. 13 Wake structure in plane 20 mm downstream of SCB for the on-design shock position. The velocities are LDA measurements made in the wind tunnel; streamwise vorticity is calculated from the CFD solution





Fig. 14 Distribution of circulation generated by SCB, calculated from CFD data



the total pressure recovery is generated on the bump tail, confirming the suggestion by Colliss et al. [15].

The peak wake circulation generated is seen to be $\Gamma / (u_\tau h) \approx 4$. This is very weak compared to the typical vorticity generated by a vortex generator for devices of this height (where $\Gamma / (u_\tau h)$ is approximately 16 for VGs run subsonically, [17]). This is perhaps not surprising, since the SCB is not explicitly designed to generate strong vortices. However, the fact that this flow structure brings about measureable benefits is encouraging, suggesting the SCB’s viability as a potential method to control the boundary layer on the aft sections of wings. This remains subject to further research.

4 Computational Resources

The FLOWer code is very suitable for vector computing. It was however decided to use the HLRS’ NEC Nehalem and Cray XE6 clusters for this study, although good experience with the NEC SX machines was made within this project earlier.

Table 5 Grid size on levels for convergence study

Mesh size	CPUs	Nodes	Iterations	Job TaT [h]	Iter/CPU/h
5.78 mio.	7	1	4,000	12.25	46.6
5.78 mio.	28	4	3,250	2.67	43.5

The reasons are the high availability of the clusters and a cost efficient and flexible computing due to the lower CPU rates on the clusters.

$$\text{balancing ratio} = \frac{\text{point number of biggest block}}{\text{point number of total mesh/number of CPUs}}$$

Computation details and parallelisation performance is shown in Table 5 for a level 1 grid as described in Sect. 2.1 (Meshing and acceleration). These computations were conducted on the NEC Nehalem cluster. As a structured mesh is used, reasonable block splitting is necessary to balance the load of the different CPUs. Two different splitting levels are compared showing a parallelisation loss of around 7.2 % for the higher splitted mesh. Note, that both meshes had a similar balancing ratio of 1.089. The lower number of iterations for the higher splitted mesh was found to be sufficient for satisfying convergence of the computation in the considered case and is independent of the block splitting.

To identify good bump designs, automated shape optimizations are performed, where around 80 consecutive computations are necessary. The overall turn around time for such an optimization can be reduced from around 1 month to around 1 week with the higher splitting of the mesh.

5 Conclusions

A joint CFD and wind tunnel study was conducted to develop a method of performing detailed studies of the flow physics of shock control bumps. In order to make the comparison possible, the free stream conditions in CFD were adapted so that the boundary layer properties upstream of the shock/boundary layer interaction on the clean airfoil matched those of the wind tunnel inflow. Additionally an additional adverse pressure gradient was imposed in the wind tunnel using a post-shock diffuser, modelling the effect of airfoil curvature. Good agreement between the methods was achieved with respect to both considerations in the region of interest.

A contour SCB was tested both experimentally and computationally using the new set-up for two different shock positions relative to the bump crest. CFD and wind tunnel results showed good agreement with each other, including the separation behaviour for the ‘off-design’ shock position. This validation of the CFD results, especially the separation prediction, gives confidence in the computational results for varying flow conditions, allowing future studies of parametric variation

effects on the flow to be studied quickly and relatively inexpensively. Additionally, it was shown that typical flow conditions on a transonic airfoil can be replicated in the wind tunnel, allowing relevant further research into detailed flow physics.

Finally, computational results were used to examine the behaviour of the flow and confirm some suggestions made by previous publications from the current work. In particular, the vortical wake structure produced by the bump was examined. The vorticity was shown to be generated by the bump tail and, although weak compared to other forms of boundary layer control, is seen to bring about a measureable total pressure recovery in the lower boundary layer further downstream. This suggests that the SCB could be used as an effective form of boundary layer control for wings, although further research is required into this concept at present.

Acknowledgements The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007–2013) under grant agreement number 271843 as part of the NextWing project, a component of the Smart Fixed Wing Aircraft initiative.

References

1. Stanewsky, E., Déleroy, J., Fulker, J. and de Matteis, P., *Drag reduction by shock boundary layer control*, Vol 80 of: Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Springer Verlag, 2002.
2. Pätzold, M., *Auslegungsstudien von 3-D Shock-Control-Bumps mittels numerischer Optimierung (Design of 3D shock control bumps by numerical optimisation)*, PhD Thesis, Institute for Aerodynamics and Gas Dynamics, University of Stuttgart, 2009.
3. König, B., *Numerical and experimental validation of three-dimensional shock control bumps*, AIAA Paper 2008-4001, 4th AIAA Flow Control Conference, Seattle, Washington, 2008.
4. Ogawa, H. and Babinsky, H., *Wind-tunnel set up for investigations of normal shock wave/boundary layer interaction control*, AIAA J., vol 44, pp 2803–2805, 2006.
5. Ogawa, H., *Experimental and analytical investigation of transonic shock/boundary layer interaction control with three-dimensional bumps*, PhD Thesis, University of Cambridge, 2006.
6. Menter F., *Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications*, AIAA Journal, 32:1598–1605, 1994
7. Kroll, N., and Fassbender, J. K., *MEGAFLOW - Numerical Flow Simulation for Aircraft Design*, Springer Verlag Berlin/Heidelberg/New York, ISBN 3-540-24383-6.
8. Streit, T., Horstmann, K.-H., Schrauf, G., Hein, S., Fey, U., Egami, Y., Perraud, J., El Din, I., Cella, U., Quest, J.: *Complementary Numerical and Experimental Data Analysis of the ETW Telfona Pathfinder Wing Transition Tests*, 49th AIAA Aerospace Sciences Meeting, AIAA 2011-881.
9. Sun, C.-C. and Childs, M.E., *A modified wall wake velocity profile for turbulent compressible boundary layers*, J. Aircraft vol 10, pp. 381–383, 1973.
10. Drela, M. and Giles, M.B., *Viscous-Inviscid Analysis of Transonic and Low Reynolds Number Airfoils*, AIAA Journal, Vol. 25, No. 10, pp. 1347–1355, 1987.
11. Nübler, K., Colliss, S., Lutz, T., Krämer, E., Babinsky, H., *Shock Control Bump Robustness Enhancement*, 50th AIAA Aerospace Sciences Meeting Nashville, Tennessee, AIAA 2012-0046, 2012.
12. Bruce, P.J.K. and Babinsky, H., *An experimental investigation of the flow physics of three-dimensional shock control bumps*, 49th AIAA Aerospace Sciences Meeting, Orlando, Florida, AIAA 2011-855, 2011.

13. Wong, W.S., Qin, N., Sellars, N., Holden, H., Babinsky, H., *A combined experimental and numerical study of flow structures over three-dimensional shock control bumps*, Aerospace Science and Technology, 12, pp. 436–447, 2008.
14. Eastwood, J.P. and Jarrett, J.P., *Towards designing with 3D bumps for wing drag reduction*, 49th AIAA Aerospace Sciences Meeting, Orlando, Florida, AIAA 2011–1168, 2011.
15. Colliss, S.P., Babinsky, H., Bruce, P.J.K., Nübler, K. and Lutz, T., *An experimental investigation of three-dimensional shock control bumps applied to transonic airfoils*, 50th AIAA Aerospace Sciences Meeting, Nashville, Tennessee, AIAA 2012–43, 2012.
16. Perry, A.E. and Hornung, H., *Some aspects of three-dimensional separation, part II: Vortex Skeletons*, Z. Flugwiss. Weltraum., vol. 8, pp 155–160, 1984.
17. Ashill, P.R., Fulker, J.L. and Hackett, K.C., *A review of recent developments in flow control*, The Aeronautical Journal, vol 109, no 1, pp 205–232, 2005.

Water Droplet Flow Paths and Droplet Deposition in Low Pressure Steam Turbines

J. Starzmann, M.V. Casey, and J.F. Mayer


Abstract The complex three-dimensional two-phase flow in a low pressure steam turbine is investigated with comprehensive numerical flow simulations. In addition to the condensation process, which already takes place in the last stages of steam turbines, the numerical flow model is enhanced to consider the drag forces between the droplets and the vapour phase. The present paper shows the differences in the flow path of the phases and investigates the effect of an increasing droplet diameter. For the flow simulations a performance cluster is used because of the high effort for such multi-momentum two-phase flow calculations. In steam turbines the deposition of small water droplets on the stator blades or on parts of the casing is responsible for the formation of large coarse water droplets and these may cause additional dissipation as well as damage due to blade erosion. A method is presented that uses detailed CFD data to predict droplet deposition on turbine stator blades. This simulation method to detect regions of droplet deposition can help to improve the design of water removal devices.

1 Introduction

The development of steam turbines reaches back to the nineteenth century where Parsons and at the same time de Laval worked on the first steam turbine concept for industrial use. Today steam turbines are very important in energy conversion, because for 70 % of world wide electricity generation a steam turbine is used, see e.g. [17]. Especially in new industrialised countries the need of steam power plants will even rise in order to meet the high demand of energy in these countries. Classically fossil fuels are used to generate steam which is then expanded in the

J. Starzmann (✉) · M.V. Casey · J.F. Mayer
ITSM – Institute of Thermal Turbomachinery and Machinery Laboratory, Universität Stuttgart,
Pfaffenwaldring 6, D-70569 Stuttgart, Germany
e-mail: starzmann@itsm.uni-stuttgart.de

Fig. 1 Wetness formation in steam turbines and the corresponding expansion line



steam turbine. But besides the fossil fuels other heat sources such as biomass, geothermal or solar-thermal energy can also be used to vaporise the working fluid, so that steam turbines will retain their important role in energy conversion, even in an environment with renewable energy sources.

The present paper investigates a special but well-known issue in steam turbines. Steam expansion in almost all low pressure steam turbines reaches saturated steam conditions even before the condenser is reached. At the outlet of the turbine already 8 to 16 % wetness exists which leads to different kinds of problems. In addition to enhanced corrosion and energy dissipation a serious danger for operational safety is given by droplet erosion of the blades. The wet steam flow is investigated numerically with the aim to improve modelling methods for this flow and to enhance the understanding of the complex three-dimensional two-phase flow phenomena that take place in low pressure steam turbines.

In Fig. 1 the process of wetness formation in a steam turbine is illustrated on the left, and on the right the associated expansion process is shown in a enthalpy-entropy diagram. The most important part of the wetness losses is called thermodynamic wetness loss because it is caused by an irreversible heat transfer during the non-equilibrium condensation itself. The considerable entropy increase after droplet formation is visible in the h-s-diagram. The wetness losses that are generated downstream of this zone of first condensation are lower but also noticeable.

Condensation in low pressure steam turbines mainly happens in terms of droplet condensation. Thus a two-phase flow exists which consists of saturated steam and dispersed small water droplets. Between the droplets and the vapour a second source of loss occurs, known as drag loss, because large droplets are not able to follow the steam path exactly and thus exert a drag force on the steam.




Fig. 2 Velocity triangles of steam and coarse water (left), example for blade erosion [14] (right)

In principle the third important loss is caused by droplets which deposit on the stator blades. These droplets form thin water films or rivulets on the blade surface which are transported to the trailing edge of the blade due to the acting shear forces on the water surface. At the trailing edge this water film breaks up and coarse water droplets with diameters between 10 and 500 μm are generated, as also outlined in Fig. 1. The distance between the stator trailing edge and the rotor leading edge is too short that these big droplets can be accelerated from their initially very low velocities (similar to the water film velocity) to the considerable higher steam velocities. The velocity triangles of the steam flow and the droplets are compared to each other in Fig. 2. It is shown that the droplets strike under high negative incidence on the rotor blades. In addition to further drag losses and the momentum exchange with the rotor (known as braking losses), the coarse water droplets lead to erosion. An example of erosion damage in a low pressure steam turbine is given in Fig. 2. The picture indicates that droplet erosion mainly takes place at higher blade radius where the impact velocities on the rotor blades increase to about 300 to 600 m/s.

2 Objectives of the Project

Within the research initiative “Kraftwerke des 21. Jahrhunderts” a project has been initialised to investigate the condensation effects in the flow field of low pressure steam turbines. An important aim of this project is to enhance knowledge about the complicated physics in rapidly expanding and condensing flows and to improve the prospects of numerical flow modelling.

Results concerning the influence of the rotor-stator induced unsteadiness on the nucleation process has already been published in the last workshop, see [16]. This study as well as previous investigations of other authors assumes that the droplets follow the steam flow without any slip. The present paper shows currently achieved results concerning the influence of the friction between the droplets and the vapour phase on the flow field of a low pressure steam turbine. The present paper also focuses on the prediction of droplet deposition that causes large coarse water droplets which are liable for braking losses and erosion damage.

3 Developments in Numerical Wet Steam Flow Modelling

In recent decades specific wet steam models to calculate the droplet formation (nucleation) and droplet growth were implemented in one-, two- and three-dimensional CFD codes. The two-phase flow models differ in the way how wetness equations are considered. The early models follow a Euler-Lagrangian approach, where the conservation equations of the vapour phase is treated in an Eulerian frame of reference and the additional equations related to the wetness formation are integrated along the streamlines in a Lagrangian way. A non-viscous two-dimensional model was published by Bakhtar [1] and a three-dimensional method that also considers the viscous effects within the vapour phase was proposed by Gerber [8]. Modern three-dimensional investigations of the two-phase flow in steam turbines are realised with RANS solvers using an Euler-Euler multi-phase approach. The works of Heiler [13], who also spent much work in validation of the specific wet steam models and Wroblewski [19] can be referenced to here. A more detailed overview about theoretical treatments for wet steam flows is given by Bakhtar [2].

4 Multi-momentum Flow Model

Condensing flow in the investigated model low pressure steam turbine is calculated with a special two-phase model implemented in Ansys CFX. This model was developed by Gerber and a first three-dimensional simulation of a multi-stage turbine using this model was realised together with the Institute of Thermal Turbomachinery (ITSM), see [10]. For further details about the numerical model the interested reader should refer to the publications of Gerber et al. [8–10] and to [16] where the models and their implementation are described.

Recently the multi-phase RANS solver was enhanced to consider the drag between the droplets and the vapour phase, see [9]. In addition to the momentum equation for the vapour phase this model solves its own momentum equation for the droplet phase (multi-momentum model). The momentum equations for the continuous (c) vapour phase and the dispersed (d) droplet phase are given as follows, wherein α is the liquid volume fraction, u the velocity, p is the pressure and τ_{ij} the viscous stress tensor.

$$\frac{\partial (\alpha \rho u_i)_c}{\partial t} + \frac{\partial (\alpha \rho u_j u_i)_c}{\partial x_j} = -\alpha_c \frac{\partial p}{\partial x_i} + \frac{\partial (\alpha \tau_{ij})_c}{\partial x_j} - \sum_{d=1}^{nd} \dot{m}_{c/d} u_{i,d} + \sum_{d=1}^{nd} S_{F,d} \quad (1)$$

$$\frac{\partial (\alpha \rho u_i)_d}{\partial t} + \frac{\partial (\alpha \rho u_j u_i)_d}{\partial x_j} = -\alpha_d \frac{\partial p}{\partial x_i} + \frac{\partial (\alpha \tau_{ij})_d}{\partial x_j} + \sum_{d=1}^{nd} \dot{m}_{c/d} u_{i,d} - \sum_{d=1}^{nd} S_{F,d} \quad (2)$$

On the right hand side of equations (1) and (2) the momentum exchange between the phases due to the condensing mass flow $\dot{m}_{c/d}$ and the drag force acting between the phases is calculated. The source term $S_{F,d}$ is given by the drag force on a single droplet times the number of droplets N .

$$S_{F,d} = N \cdot \frac{1}{2} C_D A_p \rho_d |u_{i,c} - u_{i,d}| (u_{i,c} - u_{i,d}) \quad (3)$$

A_p is the projected surface of a droplet with the radius r and equates to $A_p = \pi r^2$. The drag force that is acting between the phases depends on the drag coefficient which is modelled by the well-known drag law of Schiller-Naumann [15] for spherical particles.

$$C_D = \frac{24}{Re_D} (1 + 0.15 Re_D^{0.687}) \quad (4)$$

For large particle Reynolds numbers (over 1,000) the drag value is limited to a value of 0.44 because in this regime the drag becomes independent of the Reynolds number, see e.g. [4].

5 Influence of Inter-phase Drag on the Flow Field

The ITSM steam turbine is a scale model of a modern three-stage turbine design. It is scaled by a factor of approximately 4. The focus of the present work is on the last stage (S3/R3) of the turbine because wet steam conditions already exist at the inlet (plane E30) of this stage, as Fig. 3 shows.

In the case of a model steam turbine the fog droplets that are generated by homogeneous nucleation are very small. The circumferential averaged mean droplet size at the inlet of the last stage is shown over the blade height in Fig. 4, where it can be seen that the typical droplet size is 0.1 to 0.3 μm . Such small droplets are able to follow the steam path very well and, as they travel at steam velocity, no drag loss is generated. This situation changes as the droplet sizes become larger and are less able to follow the steam flow. For this reason an additional investigation of the last stage of the turbine was made to clarify the size of the droplets for which the situation changes. Several calculations of the flow in the last stage are performed. The droplet diameter at the inlet is increased systematically, whereas the liquid mass flow is kept constant. The inlet conditions for the default inlet diameter ($D \times 1$) and the four times larger droplet diameter distribution ($D \times 4$) are shown in Fig. 4 together with the wetness fraction γ .

Results of simulations with the default inlet diameter distribution ($D \times 1$), a four ($D \times 4$), eight ($D \times 8$) and sixteen ($D \times 16$) times larger droplet diameter distribution are shown in Fig. 5. The meridional streamlines of the vapour and the droplets for

Fig. 3 Sketch of the investigated ITSM low pressure steam turbine





Fig. 4 Droplet diameter size and wetness fraction at the inlet of the last stage (S3/R3)



the last stage of the turbine show that no noticeable difference in the flow path occurs until the (D×8)-calculation.

In the following section the influence of the droplet size on the deposition of droplets on the stator blade S3 is studied. But already from the streamlines on a surface at constant blade height, see Fig. 6, a tendency can be seen. For the calculation with the default inlet diameter (D×1) even the strong turning of the flow within the blade channel does not lead to a different flow path of the vapour and the droplets. As the droplet sizes are increased the calculation (D×8), with 8-times bigger droplets, is the first that shows a considerable effect. The absolute averaged droplet diameter at 50 % span height at the inlet of the stator S3 is then 0.7 μm. This can be obtained from Fig. 4, if the (D×1)-value at 50 % span is multiplied by the factor of 8.




Fig. 5 Meridionale streamlines for steam (*black*) and droplet trajectories (*blue* with symbols)

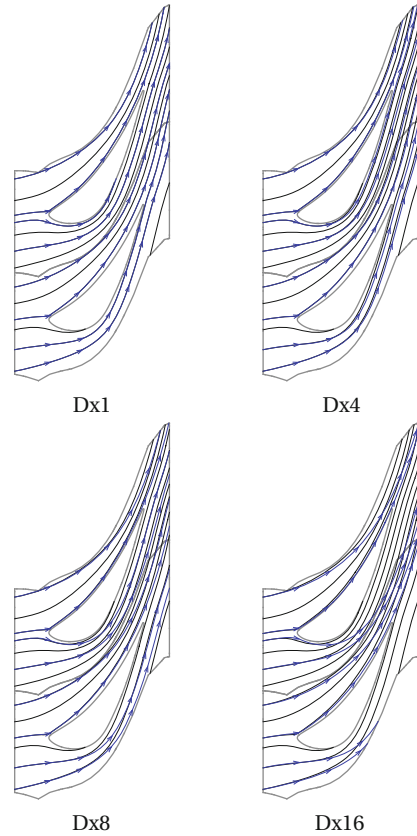
6 Droplet Deposition in Steam Turbines

For steam turbines that operate completely under wet steam conditions information about droplet deposition and coarse water formation are already needed during the design process. The turbine designer has to decide if special water removal devices at the casing or even hollow stator blades with water removal suction slots are needed to avoid erosion problems, see Fig. 7. In the following section a short overview about the mechanisms of droplet deposition is given. Subsequently the method to predict three-dimensional droplet deposition together with the numerical results are presented.

6.1 Theory of Droplet Deposition

Particle deposition mechanism in low pressure steam turbines can be divided into the laminar phenomenon “inertial depositio” caused by streamline curvature (see Fig. 2) and “turbulent diffusional deposition”. Droplet deposition by turbulent flow

Fig. 6 Vapour streamlines (black) and droplet trajectories (blue with symbols) at 50 % span



effects were investigated in numerous earlier studies where turbulent pipe flows were considered. An overview of the literature and the theoretical aspects can be taken from Guha [11], more specific information related to steam turbines are recapitulated by Crane [6].

In the following some theoretical principles about droplet deposition are given, as this is necessary to understand the results for a steam turbine stator blade presented later. A first important quantity is the relaxation time which is the time a droplet needs to be accelerated to the steam velocity. According to Gyarmathy [12] for small droplets with a diameter D the kinematic relaxation time τ can be obtained from Stokes drag law with some adjustments based on the Knudsen number Kn .

$$\tau = \frac{\rho_d D^2}{18\nu_c \rho_c} (1 + 2.7\text{Kn}) \quad (5)$$

A non-dimensional kinematic relaxation time can be defined with the friction velocity $u_* = \sqrt{\tau_w/\rho_c}$ whereas τ_w is the wall shear stress.




Fig. 7 Example for water removal devices in low pressure steam turbines (Taken from [7])

$$\tau_+ = \frac{u_*^2 \tau}{\nu_c} \tag{6}$$


The useful deposition velocity V is calculated by the droplet mass flux to the wall J_w per unit area divided by the bulk particle density γ (mass of droplets per volume). Again the friction velocity u_* can be used to get a non-dimensional deposition velocity V_+ .

$$V_+ = \frac{J_w}{\gamma u_*} \tag{7}$$

From experiments on turbulent pipe flows the behaviour of the non-dimensional deposition velocity with changing particle relaxation times can be reached, an example is shown in Fig. 8. Three different regions can be detected from the experimental data. The first is called “turbulent particle diffusion regime” because the transport mechanism of very small droplets across the viscous boundary to the wall is due to Brownian and also eddy diffusion. With rising particle diameter inertia becomes relevant and the interaction with the turbulent eddies in the near wall region results in a significant increase in droplet deposition. The second regime is therefore called “eddy-diffusion-impaction regime”. First investigations of Caporaloni [3] suppose that a mechanism called “turbophoresis” is responsible for the increase of deposition in this second regime. Turbophoresis means that particles are transported into regions with lower turbulence levels which is comparable to the general diffusion mechanism, where a diffusive flux of particles occurs towards lower particle concentrations. A further growth in particle size, and thus in relaxation time, leads to a dominance of inertial forces on the particle deposition. Larger particles are less affected by small-scale turbulent eddies and thus the third region in Fig. 8 is called “inertia-moderated” regime.

It has to be mentioned that prediction of droplet deposition in steam turbines is hindered by the fact that the droplet relaxation times can be expected within the

Fig. 8 Deposition velocity for different relaxation times according to Guha [11]



second regime, and in this regime the droplet deposition is highly sensitive to small changes in droplet size.

6.2 Method to Predict Droplet Deposition

Different methods to predict droplet or particle deposition for simple pipe flows exist, see Guha [11]. First serious investigations on droplet deposition on turbine blades due to turbulent effects were carried out by Crane [5]. Yau and Young [20] adjusted a model from Wood [18] for use on turbine flows. The model of Wood was originally developed for the prediction of particle deposition in turbulent pipe flows. Experimental data used in this semi-empirical approach to fit a set of mathematical equations to achieve a better agreement to measured deposition rates.

For the present investigation the model of Yau and Young [20] is used to determine droplet deposition on the stator blades of the last stage turbine that is investigated here. The approach was programmed in an external routine and embedded in the post-processing framework of Ansys-CFX, which enables the prediction of droplet deposition without any extra parallel computing time. However, comprehensive multi-momentum CFD simulations are needed to provide the database which is needed for the prediction of droplet deposition.

A key aspect of the present work compared to the original approach of Yau and Young [20] from the year 1987 is that precise information about the near wall flow field is available from the simulation. In the late 1980s the friction velocity, for example, had to be calculated by 2D boundary layer computation methods which were the best available method at that time. The necessary input data for this method

were obtained from two-dimensional non-viscous blade-to-blade calculations. In the present analysis flow data from three-dimensional viscous two-phase flow simulations can be used, and as the simulations have an average y^+ of 4 more details of the boundary layers are resolved.

The present CFD simulations consider the friction between droplets and steam (see Sect. 4), thus the deposition due to inertial effects can be evaluated directly from the CFD results. In conclusion within the present approach a three-dimensional prediction of the droplet deposition due to inertial as well as turbulent diffusion is possible.

6.3 Results of Droplet Deposition

Droplet deposition on the stator S3 has been predicted for the flow situation given by calculations already presented in Sect. 5. These calculations only differ in the inlet droplet diameter distributions whereas the mass flow of the vapour and the liquid phase is kept constant. The diagrams in Fig. 9 show the deposited mass flow related to the overall liquid mass flow at the inlet of the stator S3. From the curves in the left picture it can be seen, that for small droplet sizes deposition due to turbulent diffusion dominates. However, if the inlet droplet size is considerably increased the inertial deposition becomes more and more important. Of course, the overall deposition is also increased for coarser droplet distributions. The right diagram of Fig. 9 gives the difference in droplet deposition between the suction and the pressure surface of the stator blade and shows that due to the inertial effects the major deposition occurs on the pressure surface.

Further details about the three-dimensional characteristics of droplet deposition are presented in Fig. 10. The two contours present the predicted absolute deposition mass flow per unit area for the pressure side (PS) and the suction side (SS) of the blade. First of all it can be seen that droplet deposition mainly occurs around the leading edge and also on the pressure surface near the trailing edge (TE). Above all the deposition around the leading edge is due to inertial deposition because especially large droplets are not able to follow the high curvature of the vapour flow path in this region and strike directly onto the blade.

In the diagrams on the right hand side of Fig. 10 a relative deposition δ_{rel} is defined over the blade height which also helps to identify the relationship between properties of the flow field and droplet deposition characteristics. Relative deposition means that the absolute deposited mass flow, which is detected along a constant blade height, is divided by the overall deposited mass flow. The white bars represent the sum of inertial and turbulent deposition and the grey bars are related to the inertial deposition only. Obviously the droplet deposition over the span height is related to the droplet diameter distribution. In addition to the diameter also the wetness fraction itself has to be taken into account, however, in the present case the wetness fraction is nearly constant, see Fig. 4.

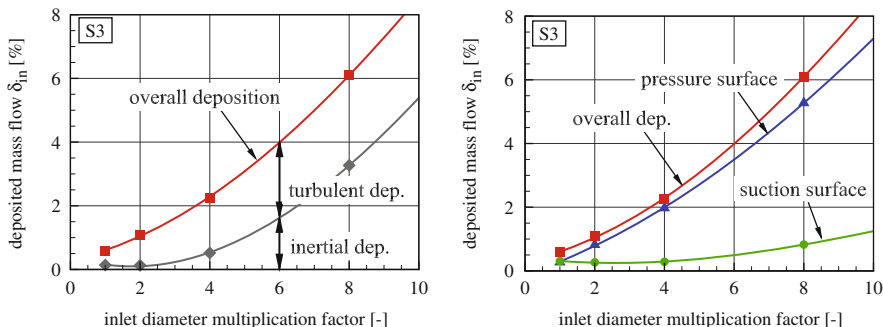


Fig. 9 Relative droplet deposition on S3 with increasing inlet droplet size

7 Computational Efforts

The high performance cluster NEHALEM of the HLRS in Stuttgart has been used for the current computations with the flow solver Ansys CFX. Several issues contribute to the need to use a computing cluster for the simulations. Firstly the standard non-equilibrium condensation model requires an additional numerical effort compared to a single-phase model because extra equations have to be solved. In addition, for the present investigation the original two-phase model was further enhanced to consider the drag between the phases which leads to a still more extensive computational model. Finally, in order to reach a sufficient residual level (RMS-residuals $\leq 10^{-5}$) the timestep has to be reduced compared to the standard non-equilibrium model. In this context also the numerical grid resolution has to be mentioned. A block structured hexahedral grid with an O-grid topology around the blades was used. For the calculation of droplet deposition an additional refinement in radial direction was necessary to reach a good quality of the results which leads to a grid with 3.3 million nodes for the last stage of the turbine.

A single multi-momentum calculation of the last stage can be realised at the ITSM PC-cluster with 24 CPUs (2.7 MHz) and takes about 3 days. An important speed-up by a factor of 5.5 has been achieved by using the NEHALEM cluster, and this is due to two effects. First the number of CPUs can be increased from 24 to typically 56 CPUs, but also the optimised computer architecture leads to a higher computing performance in this parallel simulation. Within the project about 30 simulations were made because different wet steam models and three different turbine operation conditions were included in the investigation. On this basis the pure numerical calculations would take between 2 and 3 months on a simple PC-cluster and a simulation time of only 2 weeks is needed on the high performance computing cluster.

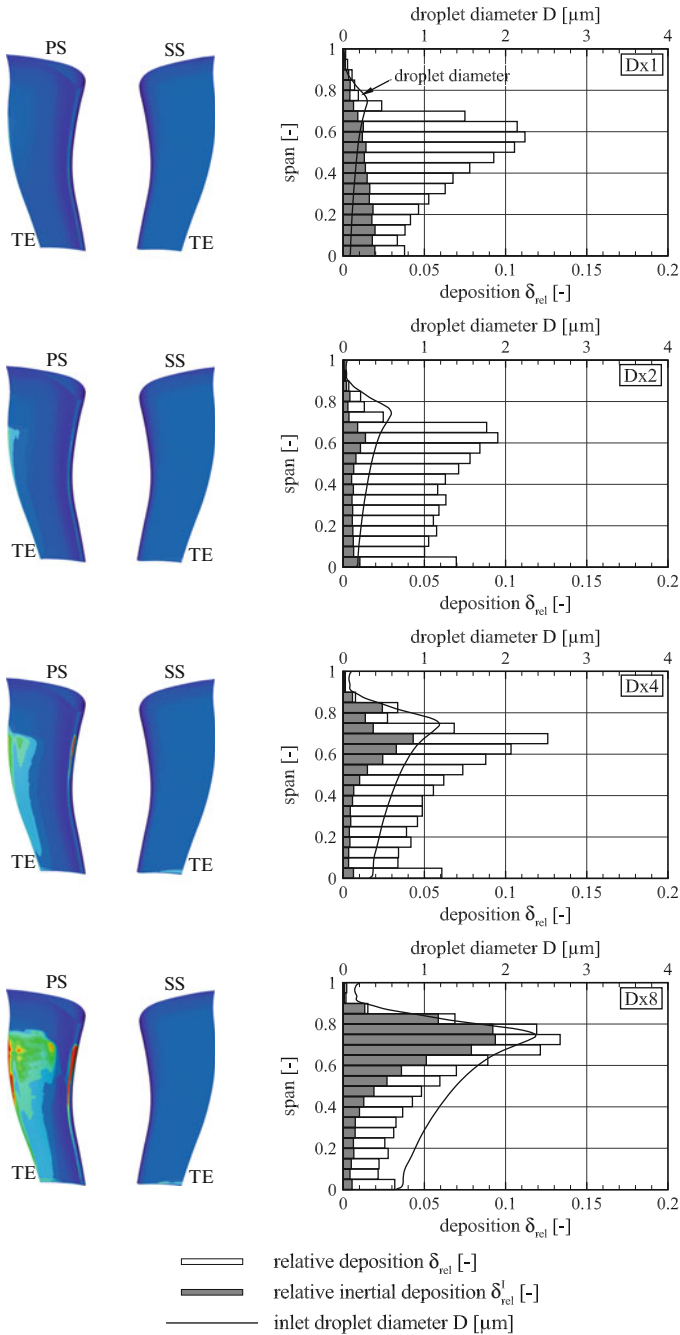


Fig. 10 Deposited mass flow rate J_w on S3 (left) and relative droplet deposition δ_{rel} (right)

8 Conclusion

The present work investigates the flow field in the last stage of a three stage model steam turbine by means two-phase flow simulations with Ansys CFX. The numerical method considers the condensation process that takes place in the flow path of the turbine. The main improvement compared to two-phase flow simulations in the past is that the drag between the vapour and the droplet phase is now also considered within the model. It has been shown that for normal operation conditions the condensed droplets are so small that almost no friction loss occurs. Furthermore the influence of increasing droplet diameter on the flow path of the liquid droplet phase compared to the general flow path of the vapour phase is discussed.

In steam turbines a certain percentage of the liquid droplets deposit on the stator blades which is responsible for the formation of large coarse water droplets. These big droplets provoke additional losses (braking and drag losses) and in addition they can lead to a reduced operational safety due to blade erosion. The work presented here provides for the first time a method that is able to predict droplet deposition due to inertial and turbulent diffusion effects for a three-dimensional turbine blade. Such numerical results can be helpful during the time-scale of the design process because the engineers need to decide if water extraction methods are needed and where they should be positioned.

References

1. Bakhtar, F.; Tochai, M. T. M.: An Investigation of Two-Dimensional Flows of Nucleating and Wet Steam by the Time-Marching Method, *Int. J. Heat and Fluid Flow*, 2(1):5–18, 1980
2. Bakhtar, F.; Young, J. B.; White, A. J.; Simpson, D. A.: Classical nucleation theory and its application to condensing steam flow calculations, *Proc. IMechE, Part C: J. Mech. Engrg. Science*, 219(12):1315–1333, 2005
3. Caporaloni, M.; Tampieri, F.; Trombetti, F.; Vittori, O.: Transfer of Particles in Nonisotropic Air Turbulence, *Journal of the atmospheric sciences*, 32:565–568, 1975
4. Clift, R.; Grace, J. R.; Weber, M. E.: *Bubbles, Drops and Particles*, Academic Press, 1978
5. Crane, R. I.: Deposition of Fog Drops on Low Pressure Steam Turbine Blades, *Int. J. Mech. Science*, 15(8):613–631, 1973
6. Crane, R. I.: Droplet deposition in steam turbines, *Proc. IMechE, Part C: J. Mech. Engrg. Science*, 218(8):859–870, 2004
7. Dooley, R. B.: *Steam Turbine Efficiency and Corrosion: Effects of Surface Finish, Deposits, and Moisture*, Tech. Rep. TR-1003997, Electric Power Research Institute, Palo Alto, California, USA, 2001
8. Gerber, A. G.: Two-Phase Eulerian/Lagrangian Model for Nucleating Steam Flow, *Trans. ASME, J. Fluids Engrg.*, 124(2):465–475, 2002
9. Gerber, A. G.: Inhomogeneous Multifluid Model for Prediction of Nonequilibrium Phase Transition and Droplet Dynamics, *Trans. ASME, J. Fluids Engrg.*, 130(3):1–11, March 2008
10. Gerber, A. G.; Sigg, R.; Völker, L.; Casey, M. V.; Sürken, N.: Predictions of Nonequilibrium Phase Transition in a Model Low Pressure Steam Turbine, *Proc. IMechE, Part A: J. Power and Energy*, 221(6):825–835, 2007

11. Guha, A.: Transport and Deposition of Particles in Turbulent and Laminar Flow, *J. Fluid Mech.*, 40:311–341, 2008
12. Gyarmathy, G.: *Grundlagen einer Theorie der Naßdampfturbine*, Dissertation, ETH Zürich, 1962
13. Heiler, M.: *Instationäre Phänomene in homogen/heterogen kondensierenden Düsen- und Turbinenströmungen*, Dissertation, Universität Karlsruhe (TH), 1999
14. Hesketh, J. A.; Walker, P. J.: Effect of wetness in steam turbines, *Proc. IMechE, Part C: J. Mech. Engrg. Science*, 219(12):1301–1314, March 2005
15. Schilller, L.; Naumann, A.: Über die grundlegenden Berechnungen bei der Schwerkraftaufbereitung, *Zeitschrift des Vereines Deutscher Ingenieure*, 77:318–320, 1933
16. Starzmann, J.; Casey, M. V.; Mayer, J. F.: Unsteady Numerical Study of Wet Steam Flow in a Low Pressure Steam Turbine, in: *High Performance Computing in Science and Engineering '11*, Transactions of the High Performance Computing Center Stuttgart (HLRS) 2011, pp. 437–450, Springer, 2012
17. VGB PowerTech: *Zahlen und Fakten zur Stromerzeugung 2010*, Sept. 2010, http://www.vgb.org/daten_stromerzeugung.html
18. Wood, N. B.: A simple method for the calculation of turbulent deposition to smooth and rough surfaces, *J. Aerosol Sci.*, 12(3):275–290, 1981
19. Wróblewski, W.; Dykas, S.; Gardzilewicz, A.; Kolovratnik, M.: Numerical and Experimental Investigations of Steam Condensation in LP Part of a Large Power Turbine, *Trans. ASME, J. Fluids Engrg.*, 131(4), 2009
20. Yau, K.; Young, J. B.: The Deposition of Fog Droplets on Steam Turbine Blades by Turbulent Diffusion, *Trans. ASME, J. Turbomachinery*, 109:429–435, July 1987

MPC and Coarray Fortran: Alternatives to Classic MPI Implementations on the Examples of Scalable Lattice Boltzmann Flow Solvers

Markus Wittmann, Georg Hager, Gerhard Wellein, Thomas Zeiser,
and Bettina Krammer

1 Introduction

In recent years, more and more parallel programming concepts have emerged as alternatives or improvements to the well established MPI concept. Arguments for all the new parallel languages or alternative communication frameworks are typically the increasing number of cores in modern systems and the hierarchical memory structure in clusters of multi-socket multi-core compute nodes.

Hybrid parallelization using OpenMP within a (ccNUMA) shared memory node and MPI between the compute nodes is a long known concept. However, correct placement of processes and efficient synchronization remain open issues. Today, there are only limited examples where a hybrid OpenMP/MPI parallelization shows a significant performance boost compared to pure MPI.

The MPC (MultiProcessor Computing) framework [9] is a unified parallel runtime for the MPI and OpenMP parallel programming models. Its unified design may provide better performances on hybrid MPI/OpenMP codes. The PGAS language Coarray Fortran (CAF) has the appeal to be part of the next Fortran standard.

Therefore, we started to evaluate both concepts, MPC and Coarray Fortran, using very simple (communication) benchmarks but also lattice Boltzmann flow solvers.

M. Wittmann (✉) · G. Hager · G. Wellein · T. Zeiser
Regionales Rechenzentrum Erlangen (RRZE), Universität Erlangen-Nürnberg, Martensstraße 1,
91058 Erlangen, Germany
e-mail: hpc@rrze.uni-erlangen.de

B. Krammer
Exascale Computing Research Center, University of Versailles Saint-Quentin-en-Yvelines, 78000
Versailles, France
e-mail: bettina.krammer@uvsq.fr

2 Related Work

Application performance optimization on shared-memory systems through the MPI library has been researched for a long time. A strategy is to execute MPI ranks inside threads instead of processes, which should make scheduling more lightweight and introduce potential for optimizations due to the shared address space. Besides MPC, this is investigated in TOMPI [3], TMPI [13], and AzequiaMPI [2, 8]. In contrast to MPC only MPI was targeted, whereas MPC unifies the runtime system of MPI, OpenMP, and POSIX threads.

An evaluation of Coarray Fortran in the context of a lattice Boltzmann flow solver has also been carried out by HASERT et al. [6]. They compare different parallelization approaches with CAF and MPI. However, they come to the conclusion that on current systems the MPI parallelization yields the best performance.

3 MPC: The MultiProcessor Computing Framework

The design goals of the MPC framework are to improve the scalability and performance of applications running on clusters of (large) multi-processor/multi-core NUMA nodes and to reduce the memory footprint of the runtime system. MPC provides a unified runtime with its own implementations of the POSIX threads, OpenMP 2.5, and MPI 1.3 (with support for `MPI_THREAD_MULTIPLE` level) standards [1, 9, 10]. MPC uses a patched version of the GNU compiler to circumvent the GOMP runtime and call the MPC OpenMP runtime instead. Thus, existing MPI or OpenMP codes only have to be recompiled with MPC.

MPC uses its own threading library, and its scalable and NUMA-aware memory allocator. The main difference between MPC and other typical MPI libraries is the process virtualization technique: MPI ranks are executed inside threads instead of processes. The fact that MPI ranks on the same node share the same virtual memory space can provide performance benefits for example for intra-node MPI communications.

MPI-style wrappers (`mpc_cc`, `mpc_cxx`, or `mpc_f77`, and `mpc_run`) are provided to facilitate the use of MPC. Because of the process virtualization technique, global variables have to be privatized so that existing MPI codes behave as expected. Without this modification, global variables would be shared by all MPI ranks on the same node leading to an incorrect behavior in most cases. A patched version of GCC 4.4 can automatically do the required transformations using the command line flag `-fmpc-privatize`. MPC can also be used with any compiler for MPI-only code. In that case, the user may have to privatize global variables manually. However, for OpenMP or hybrid MPI+OpenMP codes, MPC can only be used with its patched GNU compiler.

The startup mechanism of MPC currently only supports a limited number of batch systems (originally only Slurm; in the meantime also Torque).

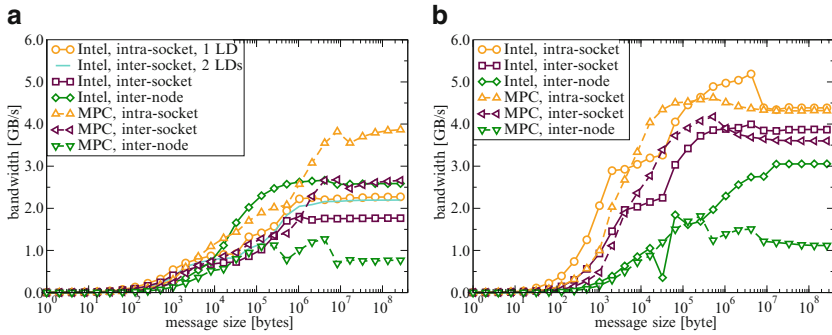


Fig. 1 PingPong bandwidth measured with Intel MPI Benchmarks (with `-off_cache` option set). “LD” means “locality domain”, i.e. cores sharing a memory interface. In the legend, “Intel” denotes Intel compiler with Intel MPI and “MPC” denotes MPC with the patched GCC 4.4. (a) LRZ’s MPP cluster (b) RRZE’s Westmere cluster

For the benchmarks the Intel compiler 12.1 in combination with Intel MPI 4.0.3 (labeled as “Intel”) and MPC 2.3.1 with the patched GCC 4.4 (labeled as “MPC”) were used.

Figure 1 shows results obtained on LRZ’s MPP cluster (2 socket nodes with octo-core AMD Opteron 6128HE processors and QDR Infiniband between the nodes) and on RRZE’s Westmere cluster (2 socket nodes with hexa-core Intel Xeon X5650 processors and QDR Infiniband between the nodes). As expected, the intra-node bandwidth as measured by the PingPong benchmark of the Intel MPI benchmarks (IMB, version 2.3.2) [7] is higher for MPC than with typical MPI implementations. The inter-node results of MPC are less convincing. For the PingPong benchmark the `-off_cache` option was used to avoid cache effects [5].

A severe limitation of common MPI implementations is that non-blocking communications (e.g. `MPI_Isend` and `MPI_Irecv`) are unfortunately not asynchronous, i.e. there is only progress while MPI library code is executing. On Cray XT systems with SeaStar interconnect asynchronous data transfer did occur [11]. The newer Cray XE6 systems with Gemini interconnect today no longer allow overlapping the data transfer for large messages. The overlapping capabilities were investigated by a simple overlap benchmark as described in [5]. Using `Isend-Recv`, MPC does asynchronous data transfer, however, there is no overlap of communication with computation in case of `Send-Irecv` or `Isend-Irecv` as shown in Fig. 2. Overlapping capabilities may improve in the next version of MPC (2.4.0) [4].

Applying hybrid MPI/OpenMP to our large-scale lattice Boltzmann flow solver did not give any benefit. Best performance is obtained with pure MPI as shown in Fig. 3. In pure MPI mode, Intel MPI with Intel Compiler (“Intel”) performs slightly better than MPC MPI with patched GCC (“MPC”), but this may be due to the compiler, as Intel MPI with GCC (“GCC”) shows almost the same performance as “MPC”. In MPI/OpenMP mode, Intel MPI coupled with Intel OpenMP is

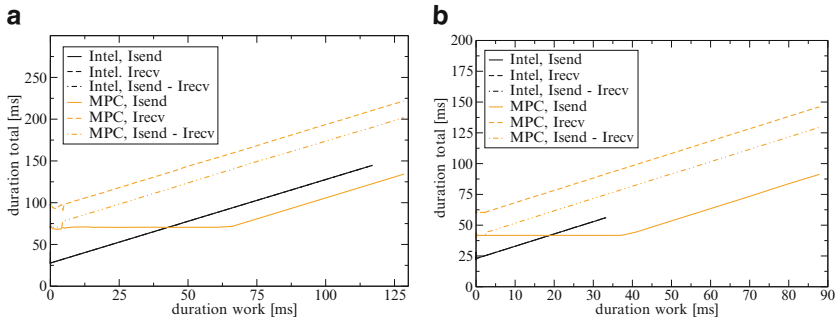


Fig. 2 Ability of MPI implementations for asynchronous communication with non-blocking point-to-point calls: A fixed amount of data (70 MB) is always transferred while the amount of work is increased. Ideally, small amounts of work should completely be hidden by the constant communication time. The Intel MPI curves for all three communication types collapse. In the legend “Intel” denotes Intel compiler with Intel MPI and “MPC” denotes MPC with the patched GCC 4.4. (a) LRZ’s MPP cluster. (b) RRZE’s Westmere cluster




Fig. 3 Performance of ILBDC flow solver simulating a randomly filled packed bed catalyst reactor on RRZE’s Westmere cluster using MPI and hybrid MPI/OpenMP. The simulation domain consists of roughly 150 million fluid nodes. In the legend “Intel” denotes Intel compiler with Intel MPI, “MPC” denotes MPC with the patched GCC 4.4, and “GCC” denotes a standard GCC 4.4 with Intel MPI

significantly better than MPC with the patched GCC compiler, and hybrid Intel MPI with GCC compiler is still better than hybrid MPC with GCC: the fact that the MPI and OpenMP runtimes are unified in MPC does not seem to offer a performance benefit in this case. Improving the OpenMP runtime and support for hybrid applications is an ongoing effort in MPC.

4 Coarray Fortran

Several Coarray Fortran implementations do exist, e.g. Cray, Intel (since the 12.x compiler series), Rice, Open64, g95.

We have not been able to get the Rice or Open64 compiler running in any reasonable way. Development of g95 seems to have stopped. The Intel compiler can be used for program development due to the high quality of the generated code; however, (multi-node) performance even on Infiniband clusters is not yet ready for productive use in the current version (12.1 update 4). The only platform with reasonable compilers and hardware support is the Cray XE6 with Gemini interconnect.

An often heard argument for Coarray Fortran is the ease of converting a non-parallel program into a distributed-parallel program. However, this simplicity is as ambivalent as with OpenMP. In our tests, good performance could only be obtained if an MPI-style parallelization concept was followed.

Details can be found in the Master Thesis of Klaus Sembritzki [12] and will be published elsewhere.

5 Conclusion and Future Work

There is still no real alternative to MPI on the horizon. But nevertheless, in the near future, we plan to investigate GPI (Global Address Space Programming Interface) from FHG/IWTM as an other PGAS approach as GPI shall be extended by IWTM within the BMBF project GASPI to include fault tolerance features which are important for the FETOL BMBF project we are participating in.

Acknowledgements This work was financially supported through the framework of the *Competence Network for Technical, Scientific High Performance Computing* in Bavaria (KONWIHR) and by BMBF under grant No. 01IH08003A (project SKALB).

It was also partly conducted at the Exascale Computing Research Center (ECR), with support provided by CEA, GENCI, Intel, and UVSQ. The hospitality of the Exascale Computing Research Center at Université de Versailles St-Quentin-en-Yvelines while working on the MPC benchmarks is gratefully acknowledged by Markus Wittmann.

Special thanks go to Prof. William Jalby for enabling this research visit at ECR, and to Marc Tchiboukdjian and Sylvain Didelot for their kind help with MPC.

The Coarray Fortran tests have been done by Klaus Sembritzki as part of his Master Thesis, which was also carried out in cooperation with ECR at UVSQ.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the CEA, GENCI, Intel or UVSQ.

References

1. P. Carribault, M. Pérache, and H. Jourden. Enabling low-overhead hybrid MPI / OpenMP parallelism with MPC. *Beyond Loop Level Parallelism in OpenMP Accelerators Tasking and More*, pages 1–14, 2010.
2. J. Dáz Martín, J. Rico Gallego, J. Álvarez Llorente, and J. Perogil Duque. An MPI-1 compliant thread-based implementation. In M. Ropo, J. Westerholm, and J. Dongarra, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, volume 5759 of *Lecture Notes in Computer Science*, pages 327–328. Springer Berlin / Heidelberg, 2009. doi:10.1007/978-3-642-03770-2_42.
3. E. Demaine. A threads-only MPI implementation for the development of parallel programs. In *Proceedings of the 11th International Symposium on High Performance Computing Systems*, pages 153–163, 1997.
4. S. Didelot, P. Carribault, M. Pérache, and W. Jalby. Improving MPI communication overlap with collaborative polling. EuroMPI, 2012. to appear.
5. G. Hager and G. Wellein. *Introduction to High Performance Computing for Scientists and Engineers*. CRC Press, 2010.
6. M. Hasert, H. Klimach, and S. Roller. CAF versus MPI – applicability of CoArray Fortran to a flow solver. In *Proceedings of the 18th European MPI Users' Group conference on Recent advances in the message passing interface*, EuroMPI'11, pages 228–236, Berlin, Heidelberg, 2011. Springer-Verlag.
7. Intel Corp. Intel MPI Benchmarks. <http://software.intel.com/en-us/articles/intel-mpi-benchmarks>, 2012.
8. Media Engineering Group (GIM) at the University of Extremadura. AzequiaMPI. <http://gim.unex.es/azequiampi>, 2012.
9. MPC. MPC project web site, 2012. <http://mpc.sourceforge.net/>.
10. M. Pérache, H. Jourden, and R. Namyst. MPC: A unified parallel runtime for clusters of NUMA machines. In *Proceedings of the 14th international Euro-Par conference on Parallel Processing*, Euro-Par '08, pages 78–88, Berlin, Heidelberg, 2008. Springer-Verlag. doi:10.1007/978-3-540-85451-7_9.
11. G. Schubert, H. Fehske, G. Hager, and G. Wellein. Hybrid-parallel sparse matrix-vector multiplication with explicit communication overlap on current multicore-based systems. *Parallel Processing Letters*, 21(03):339–358, 2011. doi:10.1142/S0129626411000254.
12. K. Sembritzki. Evaluation of the CoArray Fortran programming model on the example of a lattice Boltzmann code. Master's thesis, Erlangen Regional Computing Center, University of Erlangen-Nuremberg, 2012. <http://blogs.fau.de/hager/files/2012/07/MASembritzki.pdf>.
13. H. Tang, K. Shen, and T. Yang. Compile/run-time support for threaded MPI execution on multiprogrammed shared memory machines. *SIGPLAN Not.*, 34(8):107–118, May 1999. doi:10.1145/329366.301114.

Part V

Transport and Climate

Prof. Dr. Christoph Kottmeier

The HPC requirements for simulations of natural systems are rapidly increasing. This is reflected by, e.g., the large storage amount and CPU-times of regional climate simulations in the projects HRCM: “Modelling Near Future Regional Climate Change for Germany and Africa” (KIT Karlsruhe) and LUCCI: “Regional Climate Simulations for Southeast Asia” (KIT Garmisch-Partenkirchen). These studies provide assessments of the capabilities of regional climate models in simulating the observed climate of the last decades and the future in various regions on Earth.

Regional climate change projections of this kind are used more and more to enable estimates of climate change consequences in various economic and social sectors. HRCM addresses the needs of hydrology and flood management in medium size river catchments of basic land-atmosphere exchange studies and climate predictability in Europe and Africa. LUCCI aims at providing information to assess changes in future land-use and agricultural productivity in Central Vietnam.

Other projects like WRFCLIM (University Hohenheim) or MIPAS (KIT Karlsruhe) also reflect very well the high importance of the HLRS and SSC computing facilities for highly visible research programmes in actual research in meteorology and oceanography. They were not chosen for presentation, since they were described in HLRS-reports before.

The project AGULHAS (“The Agulhas System as a Key Region of the Global Oceanic Circulation”, IFM-Geomar Kiel) applies a model hierarchy which is based on the proven oceans model AG01 and the tested high resolutions model INALT01. For a full representation of the AGULHAS dynamics a high spatial resolution with grid scales less than 10km is achieved. The new focus of recent runs was on the climate change signal on the Agulhas leakage of Indian Ocean water to the Atlantic.

The simulations show an intensification of this process causing a salinification of upper waters in the South Atlantic.

At much higher model resolution and for much smaller domains, the project TIGRA (“Numerical Investigation of stratified turbulence”) addresses Direct Numerical Simulation (DNS) and Large Eddy Simulation (LES) of solidly stratified turbulent fluids. Simulation of geophysical turbulent flows as mentioned above requires a robust and accurate subgrid-scale turbulence modeling.

It could be shown that the implicit turbulence model ALDM correctly predicts the turbulence energy budget and the energy spectra of stratified turbulence. This is surprising, since dissipative structures are not resolved on the computational grid.

Modelling Near Future Regional Climate Change for Germany and Africa

H.-J. Panitz, G. Fossier, R. Sasse, A. Sehlinger, H. Feldmann, and G. Schädler

Abstract The modelling of future regional climate change for Germany and Africa using the regional climate model COSMO-CLM (CCLM) comprises basic studies on how temperature and precipitation are affected in general, but also specific impact studies whose results can be used for the planning of adaptation and mitigation measures.

For Africa ERA-Interim driven simulations have been carried out within the CORDEX framework. Evaluation studies using two different horizontal grid spacings (0.44° and 0.22°) did not show significant differences in the results. Furthermore, these simulations have been used to perform the transition from NEC computing systems to the new CRAY XE6, and to investigate the impact of this change on the model results, which is very small.

The impact of likely heavier summer rainfall on soil erosion in southern Germany is investigated within the KLIWA project “Bodenabtrag durch Wassererosion in Folge von Klimaveränderungen”. The corresponding simulations are performed with very high horizontal resolution (7, 2.8, and 1 km). Results show an added-value of the convection-permitting scale (2.8 km) in comparison with coarser spatial resolution (7 km). The high-resolution simulation also represents well the frequency of high and very high precipitation days.

In order to answer the important question how to account for uncertainties in regional climate projections, it is necessary to understand the sensitivity of the model results to various processes, e.g. the initialisation of soil moisture, and to create multi-member ensembles of climate simulations. The latter aspect is tackled within the Helmholtz Climate Initiative REKLIM (Regional Climate Change). In order to capture uncertainties related to the positioning of synoptic systems, a multi-member ensemble of climate simulations is generated by introducing small

H.-J. Panitz (✉) · G. Fossier · R. Sasse · A. Sehlinger · H. Feldmann · G. Schädler
Institut für Meteorologie und Klimaforschung, Karlsruher Institut für Technologie (KIT),
Eggenstein-Leopoldshafen, Germany
e-mail: hans-juergen.panitz@kit.edu

shifts to the large-scale atmospheric forcing provided by low-resolution global climate models (GCMs). The shifted atmospheric fields are then used to drive CCLM simulations at 50 km resolution. These shifts have a considerable effect on the CCLM results, in particular during hydrological summer. Thus, the ensemble generation using the shifting method and, moreover, the usage of different GCMs are valuable for an improved representation of present climate conditions and projecting regional climate change. Soil moisture is a crucial component in the atmospheric water cycle, due to the long-term memory effect of the deep soil and its feedbacks with the atmosphere. Although observational data on the temporal and heterogeneous 3-dimensional spatial distribution of soil moisture is sparse, such information is necessary to initialise climate models—especially for climate forecasts. Sensitivity studies varying the initial soil moisture distribution have been performed with COSMO-CLM. The model was able to reproduce the observations with respect to the temporal evolution of a drought index for the major European regions. Even after several years effects of the soil initialisation could be found. The effect was most pronounced in areas with continental characteristics or at high latitudes (Scandinavia, Eastern Europe or the Mediterranean) and less towards the Atlantic.

1 Introduction

The investigation of regional climate change comprises basic studies on how temperature and precipitation are affected in general, but also specific impact studies whose results can be used for the planning of adaptation and mitigation measures by the authorities. The regional climate simulations efforts carried out at the Institute for Meteorology and Climate Research (IMK) of the Karlsruhe Institute of Technology (KIT) using the regional climate model (RCM) COSMO-CLM (CCLM) consider both aspects.

The simulations for Africa [1] carried out within the CORDEX framework [2] (CORDEX: Coordinated Regional climate Downscaling Experiment), which are part of the basic studies, have been further elaborated. Furthermore, these simulations have been used to perform the transition from NEC computing systems to the new CRAY XE6, and to investigate the impact of this change on the model results.

The KLIWA (Klimaveränderung und Wasserwirtschaft; www.kliwa.de) project “Bodenabtrag durch Wassererosion in Folge von Klimaveränderungen” aims to assess the impact of climate change on soil erosion in southern Germany. Within this project precipitation data relevant for soil erosion will be provided for the recent past 1971–2000 and near future 2021–2050 using CCLM. Various project-partners will use these data as input for an erosion model. The work will focus on modelling extreme precipitation events at higher spatial and temporal resolution (2.8 km, 1 km; 1 h, 15 min). The very high spatial and temporal resolutions (1 km and 15 min) are required by erosion models.

To reach the goals of all projects and to assess the uncertainties of the climate projections, it is necessary to understand the sensitivity of the model results to various processes, e.g. the initialisation of soil moisture, and to create multi-member ensembles of climate simulations. The latter aspect is tackled within the Helmholtz Climate Initiative REKLIM (Regional Climate Change), where uncertainties related to the RCM and the large-scale atmospheric forcing from low-resolution global climate models (GCMs), which drive the RCM at its lateral boundaries, are assessed to draw reliable conclusions about the predictability of regional climate change.

2 The CCLM Model

The regional climate model CCLM is the climate version of the operational weather forecast model COSMO (Consortium for Small-scale Modeling) of the German Weather Service (DWD). It is a three-dimensional non-hydrostatic model which means that spatial resolutions below 10 km (which is considered the limit for hydrostatic models) are possible. The model solves prognostic equations for wind, pressure, air temperature, different phases of atmospheric water, soil temperature and soil water content.

Further details on COSMO and its application as a RCM can be found in [3] and [4], on the web-page of the COSMO consortium (<http://www.cosmo-model.org>), and in [5, 6], and [7].

The model is coded in Fortran 90, making extensive use of the modular structures provided in this language. Code parallelisation is done via MPI (message passing interface) on distributed memory machines using horizontal domain decomposition with a two-grid halo.

3 Regional Climate Simulations Using the HLRS Facilities

3.1 *The Soil Erosion Project*

For the near future, recent research predicts an increase in precipitation in winter coupled with a slight decrease and more variability in summer in Central Europe. Despite high regional variability, studies indicate that there will be a rise in the frequency and intensity of summer extreme precipitation events. Such increased heavy summer rainfall is likely to lead to enhanced erosion risk.

In this context, the KLIWA project “Bodenabtrag durch Wassererosion in Folge von Klimaveränderungen” was initiated to assess the impact of climate change on soil erosion in Southern Germany. Since soil erosion is a process acting on very small spatial and temporal scales, precipitation data and statistics in very high

resolution are necessary. Therefore, it is intended to downscale climate projections to regional scales using CCLM in order to reach a spatial and temporal resolution of 1 km and 15 min respectively for selected locations with high erosion risk (namely Weiherbach in Baden-Württemberg, which will be in the focus of this report, Mertesdorf in Rheinland-Pfalz, and Scheyern in Bayern). In addition, it is evaluated how the convection-permitting scale affects the precipitation statistics in comparison with coarser resolution, the changes in the climate change signal, and the erosion pattern due to extreme precipitation.

For the high-resolution simulations a multiple nesting strategy is used; this means that in a first RCM step results of a global model are downscaled to a resolution of 50 km, whose results are then used to force the simulations with 7 km resolution, and so on for the runs with 2.8 and 1 km resolution. For all of these nesting steps the same CCLM version is used. Long-term simulations at 50 and 7 km resolution have already been available from the CEDIM project “Flood Hazard in a Changing Climate” [1, 8]. Two sets of downscaling experiments have been carried out from the global scale to 1 km resolution. The first uses ERA-40 reanalysis data [9] to drive the 50 km simulation within the first nesting step. Its purpose is the validation of the results of the whole downscaling procedure (from 50 to 1 km resolution) for the recent past, 1971–2000. In the second set the results of the GCM ECHAM5 [10] are used as forcing data for the past (1971–2000) and the future (2021–2050). In both sets of experiments the 1 km resolution simulations will be applied only to certain episodes because of constraints in simulation time and memory storage capacity.

It is not known a priori whether the model is sensitive to the domain size and location as well as to internal setup. Therefore, sensitivity studies are needed in order to test the model behaviour to these parameters.

To establish the best domain size and location, four domains have been defined (Fig. 1). For each of them a 1-year simulation (the year 1986 was chosen) has been carried out, using a horizontal resolution of 2.8 km. The number of grid points within the domains vary between 60*60 and 140*140. Domain S5 represents the innermost simulation area, which encompasses the major area of interest, Weiherbach, for which the results of all sensitivity studies are evaluated. The larger domain S3 includes the entire upper and middle Rhine valley, the Black Forest, the Swabian Jura, Eifel and Hunsrück. Due to this selection regional atmospheric features induced by the local orography, e.g. the wind channeling in the Rhine Valley, can be represented more accurately. The largest domain, S4, covers the southwestern part of Germany, Luxemburg and part of France.

For the sensitivity studies, the HYRAS [11] dataset has been used to validate the model results. This daily dataset, created by the German Weather Service (DWD), provides, for the period 1951–2006, the most important parameters for hydrological applications on a regular 1 km grid.

This comparison revealed that the RCM results, especially precipitation averages, did not depend significantly on the chosen model domain (not shown). Therefore,

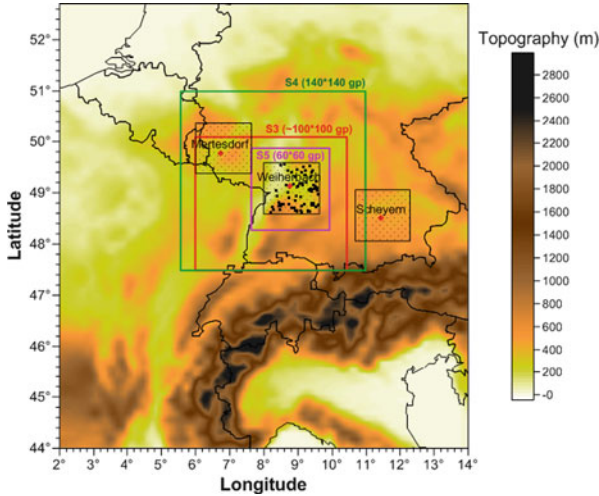


Fig. 1 Map of Central Europe showing the locations of Mertesdorf, Weiherbach, and Scheyern catchments, the investigation area for Weiherbach (*black box*), and the three model domains (S3, S4, and S5) used in the sensitivity study

the smallest domain S5 has been chosen for further sensitivity studies on different model settings, and also for the long-term simulations.

To test the sensitivity of the model to internal settings, two setups have been tested, one consistent with the setup of the simulation using the coarser resolution (7 km), and a second one, which is consistent with the suggestions of DWD for high-resolution simulations (COSMO-DE) [3]. A further test took into account a more frequent update of the forcing data (3-hourly versus 6-hourly for the two reference setups).

The studies have been performed for a 5-year period (1980–1984). They showed that the first setup together with the 3-hourly update of forcing data provided the most realistic results. Thus, using the first setup allows a more coherent comparison of simulation results using two different horizontal resolutions (7 vs. 2.8 km). It will also be used for the long-term studies (using 2.8 km resolution) and the episodic studies (using 1 km resolution) of the Weiherbach catchment.

Although the research is not yet finished, some insight can be gained about the added-value of the convection-permitting scale (2.8 km) in comparison with coarser spatial resolution (7 km).

Simulations at 7 km resolution are known to have an excess of drizzle resulting in an overestimation of both frequency and amount of low precipitation intensities ($RR < 1$ mm per day). On the other hand, runs at 2.8 km systematically underestimate the number of wet days ($RR > 1$ mm per day). This dryness of higher resolution simulations can be compensated by a more frequent update of the forcing data, because the characteristics of the coarse resolution get a stronger impact

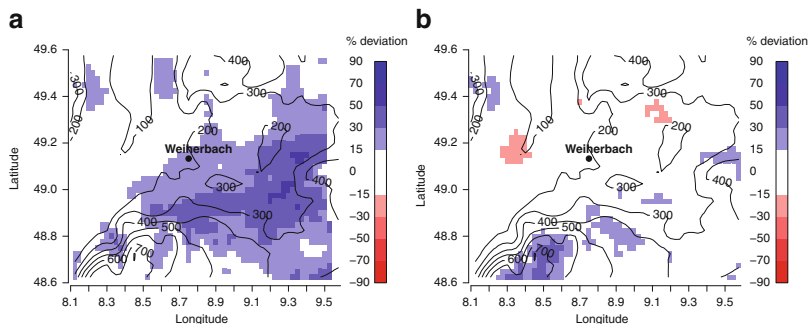


Fig. 2 Difference between CCLM results and HYRAS data for mean daily precipitation for hydrological summers 1980–1984. **(a)** 7 km horizontal resolution. **(b)** 2.8 km horizontal resolution

on the nested simulation leading to a substantial decrease of the dry days. The high-resolution simulation also represents well the frequency of high and very high precipitation days ($RR > 10$ mm and $RR > 20$ mm per day) over the whole investigation domain.

For the evaluation area around Weiherbach, Fig. 2 displays the difference in mean daily precipitation amount between HYRAS data and model results. The daily precipitation amount is averaged over all days of the hydrological summer months (May to October) of the period 1980 until 1984. It is evident that higher resolution (Fig. 2, right) reduces the bias substantially.

3.2 The Helmholtz Climate Initiative REKLIM

Within the Helmholtz Network REKLIM (Regional Climate Change), high-resolution simulations from regional climate models are used to determine climate variability and change on the regional scale for attribution and impact studies (www.reklim.de). Uncertainties related to the RCM and the large-scale atmospheric forcing from low-resolution global climate models, which drives the RCM at its lateral boundaries, have to be assessed to draw reliable conclusions about the predictability of regional climate change. In order to capture these uncertainties, an ensemble of RCM simulations is generated by means of an innovative technique called Atmospheric Forcing Shifting (AFS).

The RCM simulations are conducted using COSMO-CLM with a horizontal resolution of 50 km for Europe. ERA-40 reanalysis [9] provide the initial and boundary conditions which are updated every 6 h. The CCLM simulations span the period 1979–1984 including 1 year spin-up time which is disregarded in the further investigations. Moreover, the CCLM output is analysed for Central Europe considering only land points. Ensemble generation via AFS is realised by shifting

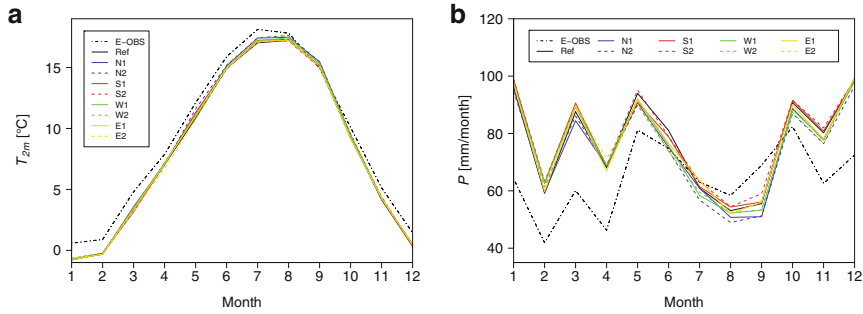


Fig. 3 Observed and modelled mean annual cycle of (a) 2m temperature ($^{\circ}\text{C}$) and (b) precipitation (mm/month) over Central Europe for the period 1980–1984

the atmospheric forcing from the ERA-40 reanalyses with respect to topography to the North (N), South (S), West (W) and East (E) by 25 km (1) and 50 km (2), respectively [12]. The shifted atmospheric fields are then used to drive the 50 km CCLM runs and, thus, eight shifting scenarios (N1, N2, S1 etc.) are performed, plus the reference simulation (Ref) which is forced by the non-shifted atmospheric fields.

On the NEC SX-8 at the HLRS facilities, the computing requirements for the domain size of $118 \times 110 \times 40$ grid points were about 48 node-hours per simulation year at 50 km resolution. Consequently, 108 node-days were required for the 6-year simulation period and nine model experiments. A node-hour is defined as the CPU time in hours that one node needs for the simulations, using all its available cores [1]. The necessary storage capacity per simulation year is in the order of 0.05 Tb and, thus, amounts to 2.7 Tb for the nine CCLM runs spanning the whole simulation period.

In order to investigate the effect of the AFS on the CCLM simulations, the mean annual cycle of 2m temperature and precipitation (Fig. 3) are determined from the model output and compared to E-OBS observations [13]. The mean absolute deviation between Ref and E-OBS is 1.0°C for 2m temperature and 15.9 mm/month for precipitation. In average, the spread between the COSMO-CLM ensemble is in the order of 0.4°C for 2m temperature and 5.4 mm/month for precipitation which amounts to 40 % (2m temperature) and 34 % (precipitation) in comparison to the deviation between Ref and E-OBS. For 2m temperature, the minimum ensemble spread can be found in February (0.1°C) whereas the maximum is in May (0.9°C). For precipitation, the ensemble spread is lowest in December (3.1 mm/month) and largest in September (8.0 mm/month).

In comparison to an ensemble of COSMO-CLM simulations driven by different GCMs [1] (Fig. 4), the spread of the AFS ensemble is clearly smaller and, in average, amounts to 15 % (2m temperature) and 27 % (precipitation) of the GCM ensemble. In particular for precipitation, the variability of the AFS ensemble spread during the mean annual cycle is much larger than for 2m temperature. The combination of both ensemble approaches, AFS and forcing from different GCMs, leads to

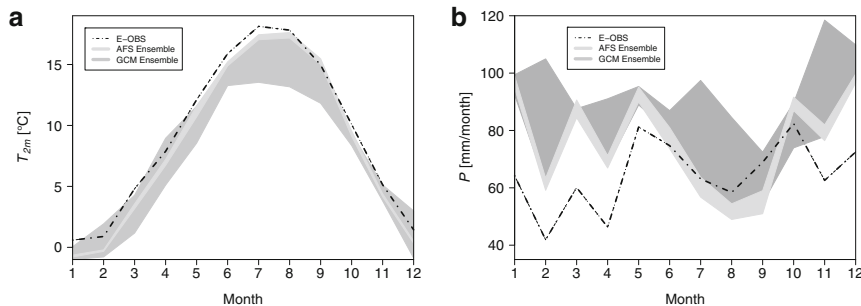


Fig. 4 Mean annual cycle of (a) 2m temperature (°C) and (b) precipitation (mm/month) from E-OBS, AFS ensemble and GCM ensemble over Central Europe for the period 1980–1984

an increase of the ensemble spread especially in summer. The resulting improved reproduction of observations goes along with an advanced representation of present climate conditions.

Moreover, the AFS effect on the spatial distribution of 2m temperature and precipitation is exemplified for the W2 and E2 scenarios (Fig. 5). In comparison to Ref, the W2 scenario leads to temperature increases of up to 0.6 °C (Fig. 5a). This coincides with a predominant precipitation decrease of up to 20.4 mm/month with large changes over the Alps (Fig. 5b). Precipitation enhancement (up to 15.4 mm/month) occurs in regions where the temperature increases are lower than in the surrounding areas. In contrast, the E2 scenario leads to temperature reduction of up to 0.3 °C in France, Germany, Poland and the Balkan States whereas regions with rising temperature (up to 0.2 °C) are rather small (Fig. 5c). The predominant temperature decrease is associated with enhanced precipitation of up to 16.9 mm/month (Fig. 5d). Precipitation decreases of up to 23.9 mm/month occur again over the Alps which might be an indicator for the sensitivity of precipitation to AFS over orographically structured terrain. As it can be seen in Fig. 5, the temperature and precipitation changes resulting from AFS are related to each other. In particular, warming coincides with reduced precipitation whereas cooling is associated with precipitation increase. Furthermore, the comparison to E-OBS (not shown) indicates that deviations between Ref and E-OBS can be partly compensated through AFS.

The model experiments show that AFS is a valuable ensemble generation technique for capturing uncertainties related to RCMs, the large-scale atmospheric forcing, and extreme events. In particular, the synergy of combining different ensemble approaches such as AFS and forcing from different GCMs fosters the assessment of uncertainties in climate modelling and of the predictability of climate change on regional scales. For the generation of large high-resolution ensembles, the HLRS resources thus remain crucial. In the future, the CRAY XE6 at HLRS will be used for generating a four-member ensemble of COSMO-CLM simulations based on AFS at horizontal resolutions of 50 and 7 km for 30-year periods under present and future climate conditions.

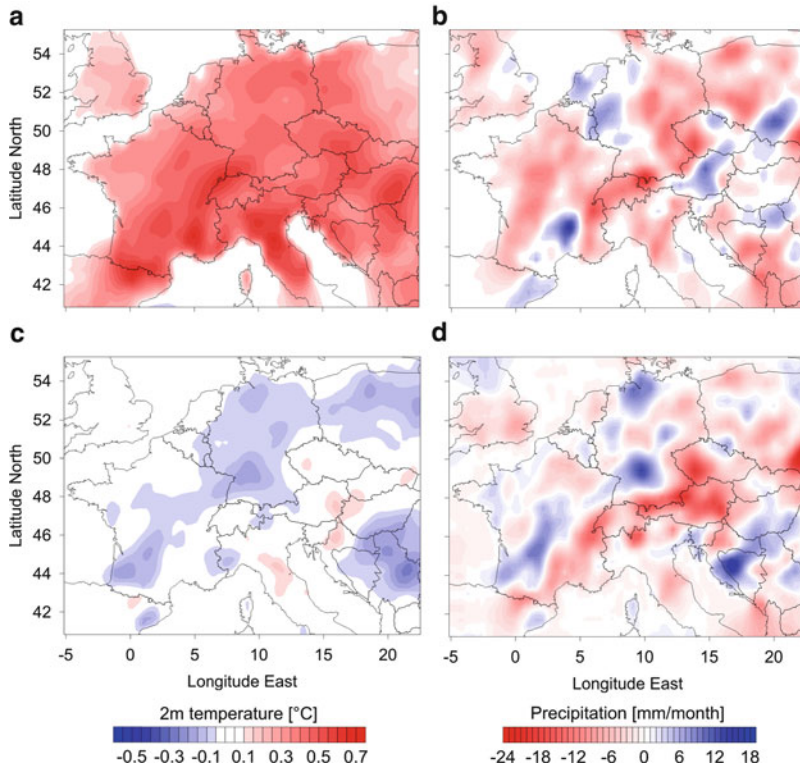


Fig. 5 Difference between the simulations W2-Ref (a and b) and E2-Ref (c and d) for 2m temperature (°C) (a and c) and precipitation (mm/month) (b and d) for the period 1980–1984 (only hydrological summer)

3.3 Sensitivity to Soil Moisture Initialization

Soil moisture is a crucial component in the atmospheric water cycle. It provides a longer-term memory compared to the atmosphere, due to its storage capacity and the longer response times of deeper layers. The feedbacks between soil and atmosphere include precipitation as the primary source for the soil water content and the evapotranspiration of water vapour from soil and vegetation towards the atmosphere [14].

Climate models need reliable information about the 3-dimensional distribution of the soil moisture at the initialisation stage. Errors in the soil moisture fields will cause artificial trends in the model simulations overlaying the climate trends. The observational base to derive soil moisture distributions for large areas at a given time is insufficient and the soil is very heterogeneous even on small scales. The few ground-based observations are therefore likely not representative for larger regions and usually do not cover climatological time scales or large areas. Therefore, remote

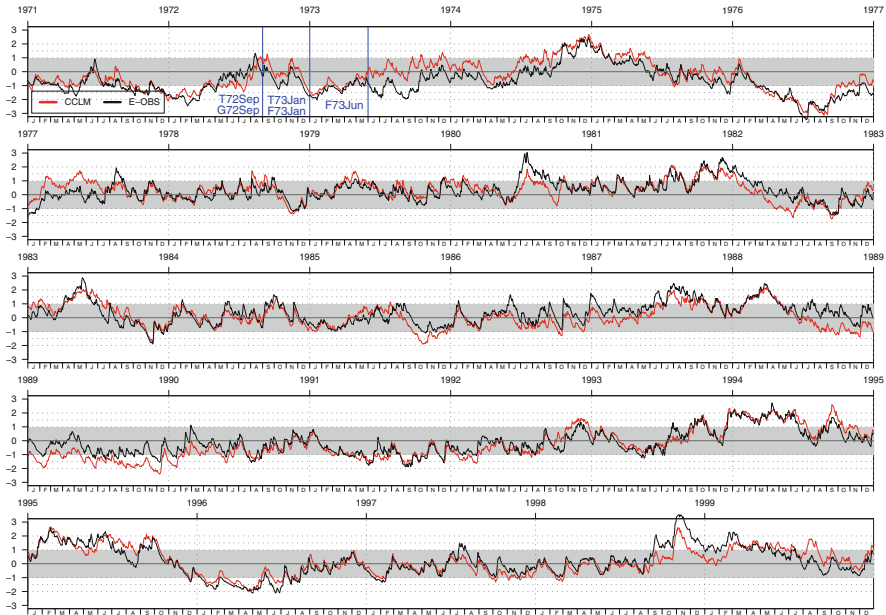


Fig. 6 Time series of the effective drought index (EDI) for Mid-Europe (Data derived from CCLM (red line) and from E-OBS (black line))

sensing is likely to be the only method to provide large-scale information, although these measurement techniques only cover the uppermost layers of the soil [15].

Under such circumstances regional climate models use a very pragmatic approach. They generate their own balanced soil moisture distribution by a pre-analysis phase simulation starting from simple vertical moisture profile. Usually, about 1–3 years spin-up time is used as a compromise between computing time requirements and the need for balanced soil moisture distribution.

To test the uncertainties and sensitivities induced by this initialisation problem, several simulation experiments with CCLM have been performed. The model domain covers Europe with a resolution of 50 km. The 30-year base simulation covers the period 1968 until 2000 and is driven by the ERA-40 reanalysis [9] to allow for a direct comparison with observations.

To identify drought and moist periods the Effective Drought Index (EDI) [16] is calculated for simulated and observed precipitation (Fig. 6). The CCLM simulation is able to reproduce the temporal development of EDI very closely. The duration and intensity of extreme periods, like the long-term drought in 1976, are well matched.

From this basic analysis a test period for the sensitivity experiments has been derived. The concept of the experiments is to examine the uncertainty induced by the insufficient information of the soil moisture in connection with the long-term memory of the soil. With all other conditions fixed, the initial soil moisture conditions have been altered for different starting dates September 1st, 1972,




Fig. 7 Monthly NRMSD of total precipitation, evapotranspiration and soil moisture in different soil layers (Data extracted from the experiment with reduced soil moisture by 50 % on January 1st, 1973)

January 1st, 1973 and June 1st, 1973 (blue lines in Fig. 6) to study the effect of initialisation in different seasons and changing soil conditions. The soil conditions have been altered by up to 50 % towards wetter and drier conditions.

Figure 7 shows the normalized root mean square difference (NRMSD) of the sensitivity experiment compared to the base simulation for different soil layers, precipitation and evapotranspiration in Central Europe. The time series of NRMSD for an experiment starting on January 1st, 1973 with reduced soil moisture by 50 % shows an initial relaxation of the disturbed soil moisture conditions towards the bases simulation. The response time is much longer in deeper soil layers. But even after several years a memory of the starting conditions is visible. This is also the case not just within the soil but also for precipitation and evapotranspiration. The differences are most pronounced in summer with typically drier soil conditions than in winter.

The experiments cover different European regions with different climatological characteristics from the Iberian Peninsula to Scandinavia. The regions show different sensitivities with respect to the soil initialisation especially in the deep soil (Fig. 8). In the uppermost layers (above 20–30 cm depth) all regions show a relaxation time within typical time for a spin-up simulation. But the lower soil layers need considerably more time to get to a state close to the base simulation. The longest relaxation periods were found in Scandinavia, where frozen soils may occur, followed by Eastern Europe with more continental conditions and the dry Mediterranean area. The westerly regions, close to the Atlantic, like the British Isles and France need somewhat shorter times.




Fig. 8 Dependence on spin-up time of European regions and soil layers (Data extracted from the experiment with reduced soil moisture by 50% on January 1st, 1973)

To conclude:

- The highest soil moisture sensitivities were found for dry conditions and when frozen soils are involved.
- Under moist conditions, when the soil is close to saturation, the memory is shorter than for dry conditions. Therefore, a model initialisation in winter or other moist periods is preferable.
- Even a several year spin-up period of climate models might not be sufficient. An initial soil moisture initialisation based on long-term prior simulations might reduce the problem but further studies and a better observational base are needed to overcome this initialisation problem.

3.4 The CORDEX Framework

Simulations with CCLM had been carried out for Africa within the CORDEX framework [2]. Africa had been chosen as the first CORDEX target region because of its vulnerability to climate change in terms of impacts on temperature and precipitation patterns. The ERA-Interim [17] driven evaluation results, which used the demanded horizontal resolution of 0.44° (AFR44), had been briefly described in [1]. Besides this evaluation run a simulation with a higher resolution of 0.22° (AFR22) had been performed, which uses also the direct forcing from the ERA-Interim reanalyses. A comparison of the results with climatological observations

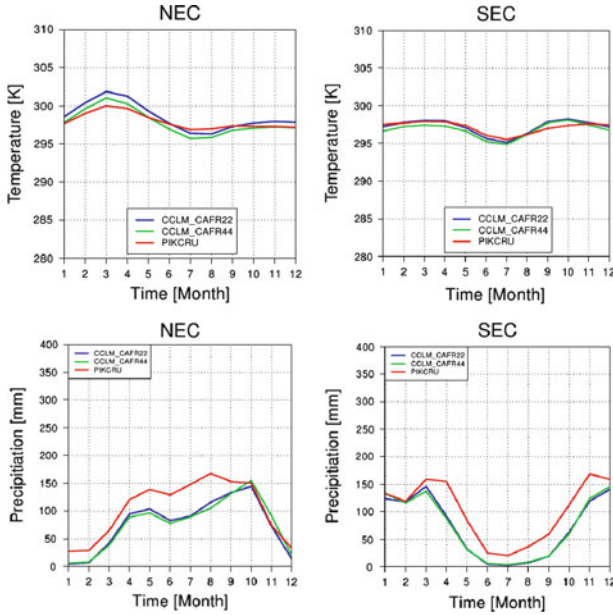


Fig. 9 Annual cycle of monthly means of temperature at 2 m height (*upper row*) and mean monthly sums of precipitation (*lower row*) for African land regions north (NEC, *left*) and south (SEC, *right*) of the Equator for the period 1991–2000; *blue curves*: results of AFR22 CCLM simulation; *green curves*: results of AFR44 CCLM simulation; *red curves*: climatological observations (M. Büchner, PIK, Potsdam, pers. comm.)

did neither show a significant improvement of the model results due to the finer resolution nor a significant deterioration (Fig. 9). In general, the higher resolution led to a slight increase of mean monthly temperatures by about 0.4 °C, and to slightly more precipitation of about 40 mm/month.

3.5 Implementation on CRAY XE6

After some technical changes in the code, the RCM COSMO-CLM could be successfully implemented on the CRAY XE6 at HLRS. With the help of scaling tests a suitable domain decomposition into 32*32 sub-domains (= 1,024 cores) had been found as appropriate for all simulation activities. Much effort had been spent to find compiler options that guarantee identical results when changing only the domain decomposition, and thus the number of cores. Table 1 summarizes the CRAY and INTEL compiler options with highest optimisation level that fulfill this requirement.

The next step was to study the impact of the change of the computing system (from NEC SX8 to CRAY XE6) on the model results. For this purpose the CORDEX

Table 1 Recommended option to compile the RCM COSMO-CLM using the CRAY or INTEL programming environments on CRAY XE6

Programming environment	Compiler options
CRAY	-c -O1 -O fp1 -eF
INTEL	-c -cpp -Os -no-vec

Table 2 Performance of CCLM CORDEX Africa simulation on CRAY XE6 and NEC SX8. All values are valid for one simulation year; domain size of 0.44° simulation: 214*221*35; unit nh/year: node hours per simulation year

Computing platform	No. of cores	No. of nodes	Wall-clock time (h)	CPU-time (nh/year)
CRAY XE6	1,024	32	2.8	90.1
NEC SX8	16	2	8.8	17.0

Africa evaluation simulations have been repeated on the CRAY for the period 1989 till 1996, using the identical forcing data as on the NEC SX8. Fortunately, the differences between the results are negligible (not shown).

This test gave also the opportunity to compare the performances of the simulations on both computing systems. Table 2 summarizes the results of this comparison. Due to the large number of cores that can be used on the CRAY the wall-clock time is much less than on the NEC SX8 for a 1-year simulation. On NEC SX8 there was no performance gain when using more than 16 cores (2 nodes). In contrast, the CPU time per simulation year increases considerably when using the CRAY. For this test all cores of a node had been used. When using only half the number of cores per node on the CRAY without changing the total number of cores (thus, doubling the number of used nodes), wall-clock time and CPU time could be reduced by about 25 % due to optimized cache usage.

4 Future Work and HPC Demands

Between the timescales of short-term and seasonal weather prediction on the one side, and long-term climate projections on the other side, there is a gap in the timescales of several years to several decades, for which climate predictions are not yet possible. On these timescales, weather and climate patterns depend not only on the anthropogenic rise of greenhouse gas concentrations, but also strongly on natural climate variability, induced by variations in the atmosphere and oceans (internal natural forcing), and by solar variability and volcanic eruptions (external natural forcing). The aim of the research program MIKLIP (Mittelfristige Klimaprognosen), which has been initiated by the Federal Ministry for Science

and Education (BMBF), is to close this gap, and to develop a system for climate predictions for up to a decade ahead. The IMK contributes to MIKLIP in joint projects which aim at regional decadal predictions of climate for Central Europe and the West African monsoon region.

In addition, the works within the CORDEX framework, the KLIWA erosion project, and the REKLIM research initiative will be continued.

The new project KLIMOPASS continues the goals of the former project “Herausforderung Klimawandel”. In this project, new ensembles of high-resolution climate simulations (7 and 2.8 km), driven by new GCM data, will be established and analysed. New ensemble generating techniques, taken e.g. from the REKLIM project, will also be used. The projection period will include 2021–2050, and possibly also 2071–2100.

For all these activities large demands of High Performance Computing are necessary. A rough estimate amounts to about 400.000 CPU-h for the next 2 years, and 160 Tb of storage capacity, which have to be available online in order to perform the CCLM simulations and the subsequent analysis of the model results (post-processing).

References

1. Panitz, H.-J., P. Berg, G. Schädler, and G. Fossler (2012): Modelling Regional Climate Change for Germany and Africa, In: High Performance Computing in Science and Engineering '11 [W. E. Nagel, D. Kröner, M. Resch (Eds.)]. DOI: 10.1007/978-3-642-23869-7, Springer Berlin Heidelberg New York 2012, pp 503–512.
2. Giorgi, F., C. Jones, and G. R. Asrar (2009): Addressing climate information needs at the regional level: the CORDEX framework. WMO Bulletin 58 (3)- July 2009, 175–183.
3. Baldauf, M., A. Seifert, J. Förstner, D. Majewski, and M. Raschendorfer (2011): Operational Convective-Scale Numerical Weather Prediction with the COSMO Model: Description and Sensitivities. Mon. Wea. Rev. 139, 2011, 3887–3905, DOI: 10.1175/MWR-D-10-05013.1.
4. Doms, G. and U. Schättler (2002): A description of the nonhydrostatic regional model LM, Part I: Dynamics and Numerics. COSMO Newsletter, 2, 225–235.
5. Meissner, C. and G. Schädler (2007): Modelling the Regional Climate of Southwest Germany: Sensitivity to Simulation Setup. In: High Performance Computing in Science and Engineering '07 [W. E. Nagel, D. Kröner, M. Resch (Eds.)]. ISBN 978-3-540-74738-3, Springer Berlin Heidelberg New York.
6. Rockel, B., A. Will, and A. Hense (2008): Regional climate modelling with COSMO-CLM (CCLM). Meteorologische Zeitschrift, 17, 4, 2008, 347–348, DOI: 10.1127/0941-2948/2008/0309, special issue, ISSN 0941–2948.
7. Meissner, C., G. Schädler, H.-J. Panitz, H. Feldmann, and Ch. Kottmeier (2009): High resolution sensitivity studies with the regional climate model COSMO-CLM. Meteorologische Zeitschrift, 18, 543–557, DOI: 10.1127/0941-2948/20090400.
8. Schädler, G., P. Berg, D. Dütthmann, H. Feldmann, J. Ihringer, H. Kunstmann, J. Liebert, B. Merz, I. Ott, and S. Wagner (2012): Flood Hazards in a Changing Climate. Project Report, pp 83. Centre for Disaster Management and Risk Reduction Technology (CEDIM), http://www.cedim.de/download/Flood_Hazards_in_a_Changing_Climate.pdf

9. Uppala, S. M., P. W. Kållberg, A. J. Simmons, U. Andrae, V. Da Costa Bechtold, M. Fiorino, J. K. Gibson, J. Haseler, A. Hernandez, G. A. Kelly, X. Li, K. Onogi, S. Saarinen, N. Sokka, R. P. Allan, E. Andersson, K. Arpe, M. A. Balmaseda, A. C. M. Beljaars, L. Berg, J. Van De Bidlot, N. Bormann, S. Caires, F. Chevallier, A. Dethof, M. Dragosavac, M. Fisher, M. Fuentes, S. Hagemann, E. Hólm, B. J. Hoskins, L. Isaksen, P. A. E. M. Janssen, R. Jenne, A. P. McNally, J.-F. Mahfouf, J.-J. Morcrette, N. A. Rayner, R. W. Saunders, P. Simon, A. Sterl, K. E. Trenberth, A. Untch, D. Vasiljevic, P. Viterbo, and J. Woollen (2005): The ERA-40 re-analysis. *Q. J. R. Meteorol. Soc.*, 131, 2961–3012.
10. Roeckner, G., G. Baeuml, L. Bonaventura, R. Brokopf, M. Esch, M. Giorgetta, S. Hagemann, I. Kirchner, L. Kornbluh, E. Manzini, A. Rhodin, U. Schlese, U. Schulzweida, and A. Tompkins (2003): The atmospheric general circulation model ECHAM 5. PART I: Model description. Technical Report 349, pp 127, Max-Planck-Institut für Meteorologie, Bundesstr. 55, D-20146 Hamburg, Germany.
11. Steiner H., U. Riediger, and A. Gratzki (2011): The HYRAS data set - A high resolution gridded reference data set covering Germany and neighbouring river basins. EMS Annual Meeting Abstracts, Vol. 8, EMS2011-589, 2011.
12. Schlüter, I. and G. Schädler (2010): Sensitivity of Heavy Precipitation Forecasts to Small Modifications of Large-Scale Weather Patterns for the Elbe River. *J. Hydrometeor.*, 11, 770–780.
13. Haylock, M. R., N. Hofstra, A. M. G. Klein Tank, E. J. Klok, P. D. Jones, and M. New (2008): A European daily high-resolution gridded dataset of surface temperature and precipitation. *J. Geophys. Res.*, 113, D20119.
14. Seneviratne, S., E. Corti, E. L. Davin, M. Hirschi, E. B. Jaeger, I. Lehner, B. Orlowsky, and A. J. Teuling (2010): Investigating soil moisture-climate interactions in a changing climate: A review. *Earth-Sci. Rev.*, 99, 125–161.
15. Kerr, Y. H. (2007): Soil moisture from space: Where are we? *Hydrogeol. J.*, 15, 117–120.
16. Byun, H. R. and D. A. Wilhite (1999): Objective Quantification of Drought Severity and Duration. *J. Climate*, 12, 2747–2756.
17. Simmons, A., S. Uppala, D. Dee, and Sh. Kobayashiera (2006): New ECMWF reanalysis products from 1989 onwards. ECMWF Newsletter No. 110, Winter 2006/07, 25–35.

Setting Up Regional Climate Simulations for Southeast Asia

Patrick Laux, Van Tan Phan, Christof Lorenz, Tran Thuc, Lars Ribbe, and Harald Kunstmann

Abstract Climate change and climate variability are main drivers for land–use, especially for regions dominated by agriculture. Within the framework of the project *Land–Use and Climate Change Interactions in Central Vietnam* (LUCCI) regional climate simulations are performed for Southeast Asia in order to estimate future agricultural productivity and to derive adaptive land–use strategies for the future. Focal research area is the Vu Gia–Thu Bon (VGTB) river basin of Central Vietnam. To achieve the goals of this project reliable high resolution climate information for the region is required. Therefore, the regional non-hydrostatic *Weather Research and Forecasting* (WRF) model is used to dynamically downscale large-scale coupled atmosphere–ocean general circulation model (AOGCM) information. WRF will be driven by the ECHAM5-GCM data and the business-as-usual scenario A1B for the period 1960–2050. The focus of this paper is on the setup of WRF for East Asia. Prior to running the long-term climate simulation in operational

P. Laux (✉) · C. Lorenz · H. Kunstmann

Karlsruhe Institute of Technology (KIT), Institute for Meteorology and Climate Research, Atmospheric Environmental Research (IMK-IFU), Kreuzeckbahnstrasse 19, 82467 Garmisch-Partenkirchen, Germany

e-mail: patrick.laux@kit.edu; christof.lorenz@kit.edu; harald.kunstmann@kit.edu

V.T. Phan

Faculty of Hydrology, Meteorology and Oceanography, Hanoi University of Science, 334 Nguyen Trai, Thanh Xuan, Hanoi, Vietnam

e-mail: tanpv@vnu.edu.vn

T. Thuc

Ministry of Natural Resources and Environment (MONRE), Institute of Meteorology, Hydrology and Environment (IMHEN), Nguyen Chi Thanh street, Hanoi, Vietnam

e-mail: thuc@netnam.vn

L. Ribbe

Cologne University of Applied Sciences, Institute for Technology and Resources Management in the Tropics and Subtropics (ITT), Betzdorfer Str. 2, 50679 Köln, Germany

e-mail: lars.ribbe@fh-koeln.de

mode, experimental simulations using different physical parameterizations have been conducted and analyzed. Different datasets have been used to drive the WRF model and to validate the model results. For the evaluation of the parameterization combination special emphasis is given to the representation of the spatial patterns of rainfall and temperature. In total, around 1.7 Mio CPUh are required to perform the climate simulations. The required computing resources have been approved from the Steinbuch Centre for Computing (KIT, SCC).

1 Introduction

Climate change and climate variability are of major concern for Central Vietnam's environment and people's well-being. Increasing frequency and severities of extreme events like floods, droughts, hurricanes but also increasing temperatures, sea level rise and salt water intrusion in the coastal areas are expected to have dramatic consequences for agricultural productivity and thus food security in the VGTB basin. These challenges demand for informed stakeholders and a land management strategies to increase the resilience of the ecosystems.

Vietnam's lowlands and midlands are predominantly characterized by rice landscapes. Rice is the pillar of food security for many million households. Rice production is complexly linked with land management, water and environment. Judicious management of rice ecosystems is seen today – in a post-green revolution age – as a major strategy for raising rice productivity, protect the environment and achieve long term food security for rural and urban populations in Vietnam and all over rice producing Asia. Many problems related to rice production and climate change/climate variability in Vietnam became obvious during the last decades. Temperature stress, especially during sensitive rice development stages, negatively affects crops development and yields. Sea level rise and salt water intrusion in lowland coastal areas are affecting rice cropping, while the magnitude of effects depend on the complex interactions between the cropping calendar and the hydrological situation. Although rice is a very salt resistant crop, salinity levels beyond threshold levels will eventually decrease yields. Excessively high water level and prolonged inundation periods can severely affect yields.

The major goals of the LUCCi project is to provide a sound future land use management framework that considers socio-economic development, population growth and expected impacts of climate change on land and water resources. This framework links climate change mitigation – through the reduction of GHG emissions – with adaptation strategies to secure food supply in a changing environment. As a basis, present and past land use practices and the use of water resources in the VGTB basin are analyzed with special emphasis on possible climate change impacts. This in-depth analysis will allow deriving carbon-optimized land and water use strategies for the VGTB basin as well as for the larger region of Central Vietnam.

The *Dynamical Downscaling* (DDS), which is performed by KIT, IMK-IFU will contribute to assess future land-use and quantify agricultural productivity in the VGTB basin of Central Vietnam. This requires reliable and high resolution information about the climatology for the present and past, but also future regional

climate projections. It is widely accepted that present-day General Circulation Models (GCMs) are able to simulate the large-scale state of the atmosphere in a realistic manner, and predict large-scale climate change based on assumptions about future greenhouse gas emissions (AR4, IPCC). Their implications on regional and local scales, however, are inadequate mainly due to the limited representation of mesoscale atmospheric processes, topography, and land-sea distribution in GCMs (e.g. [2, 14]). A direct application of GCM output for regional and local impact studies would lead to inconsistencies in frequency statistics, such as the occurrence probabilities of rainfall events (e.g. [7, 12]). Within the LUCCi framework, both *Dynamical* as well as *Statistical Downscaling* (SDS) approaches will be followed and combined in that way, that the advantages of each approach are capitalized: The scarcity of observed climate data in this region, and the most probably non-stationary climatic processes for the future period of interest (2010–2050) demand for a DDS approach. Therefore, the regional non-hydrostatic *Weather Research and Forecasting* (WRF) model will be driven in three nesting steps with resolutions of 45, 15, and 5 km to obtain transient and consistent climate simulations for the period 1960–2050. Due to the high computational demands, WRF will solely be driven by the coupled atmosphere–ocean general circulation model ECHAM5 and the business-as-usual A1B scenario. The underlying assumptions for the chosen scenario are a future world of very rapid economic growth, low population growth and rapid introduction of new and more efficient technology. In addition to the dynamical downscaling approach, a multi-model and multi-scenario ensemble-based SDS technique will be developed which allows for quantifying uncertainties inherent to the downscaling approaches, the GCMs and the chosen scenarios.

The main goal of this paper is to identify a suitable setup of WRF physical parameterizations for Southeast Asian, which is the prerequisite to conduct long-term climate simulations. Based on the identified setup, the long-term climate simulations are conducted. Albeit few regional climate simulation efforts have been done for Southeast Asia (e.g. [9]), to the best of our knowledge these simulations represent the first efforts (i) to identify the best physical parameterization combination in a systematic manner, and (ii) to conduct transient long-term climate simulations for the VGTB river basin with this detailed spatio-temporal resolution.

2 The Regional Climate Model Simulations

2.1 The Regional Climate Model WRF

WRF is a next-generation mesoscale numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs. The WRF Software Framework (WSF) provides the infrastructure that accommodates the dynamics solvers, physics packages that interface with the solvers, programs for initialization, WRF-Var, and WRF-Chem. There are two dynamics solvers in the WSF. The one applied in this project is the Advanced Research WRF (ARW)

solver which was primarily developed at NCAR (National Centre for Atmospheric Research, USA). The ARW dynamics solver integrates the compressible, non-hydrostatic Euler equations. The equations are cast in flux form using variables that have conservation properties. The equations are formulated using a terrain-following mass vertical coordinate. The flux form equations in Cartesian space are extended to include the effects of moisture in the atmosphere and projections to the sphere.

For the temporal model discretization the ARW solver uses a time-split integration scheme. Generally speaking, slow or low-frequency (meteorologically significant) modes are integrated using a third-order Runge-Kutta (RK3) time integration scheme, while the high-frequency acoustic modes are integrated over smaller time steps to maintain numerical stability. The horizontally propagating acoustic modes (including the external mode present in the mass-coordinate equations using a constant-pressure upper boundary condition) and gravity waves are integrated using a forward-backward time integration scheme, and vertically propagating acoustic modes and buoyancy oscillations are integrated using a vertically implicit scheme (using the acoustic time step). The time-split integration for the flux-form equations is described and analyzed in [5].

The spatial discretization in the ARW solver uses a C grid staggering for the variables. That is, normal velocities are staggered one-half grid length from the thermodynamic variables. The grid lengths Δx and Δy are constants in the model formulation; changes in the physical grid lengths associated with the various projections to the sphere are accounted for using map factors. The vertical grid length $\Delta 2\eta$ is not a fixed constant; it is specified in the initialization. The user is free to specify the η values of the model levels subject to the constraint that $\eta = 1$ at the surface, $\eta = 0$ at the model top, and η decreases monotonically between the surface and model top.

2.2 WRF Domain Setup

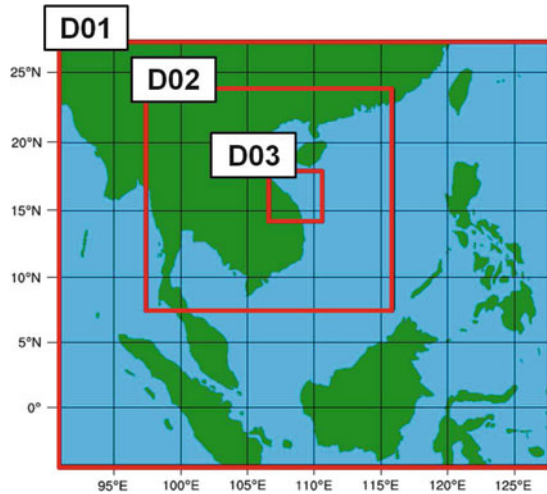
For the research project LUCci climate simulations with a target resolution of 5 km for Central Vietnam for 1960–2050 shall be performed. For this purpose, WRF is nested in the general circulation model ECHAM5 using the following setup and location of the domains (Fig. 1).

1. Domain1:

- Horizontal: 99×99 grid points with a resolution of 45 km
- Vertical: 50 layers up to 5,000 Pa
- Time step: 180 s (adaptive time-step option enabled)

(continued)

Fig. 1 Domains to be modelled by the regional climate model WRF using nesting strategy



(continued)

2. Domain2:

- Horizontal: 142×145 grid points with a resolution of 15 km
- Vertical: 50 layers up to 5,000 Pa
- Time step: 180 s (adaptive time-step option enabled)


3. Domain3:

- Horizontal: 66×75 grid points with a resolution of 5 km
- Vertical: 50 layers up to 5,000 Pa
- Time step: 180 s (adaptive time-step option enabled)

2.3 Required HPC Ressources

Due to the limitations of the CFL stability condition, the option adaptive time-step has been enabled resulting in an indefinable number of required integration steps. Calculating with the predefined time steps of 180 s, more than 15 Mio integration steps are required respectively on 490,050 grid cells for Domain 1, 1,029,500 grid cells for Domain 2, and 247,500 grid cells for Domain 3 with more than 10 degrees of freedom on each grid cell (momentum, mass, pressure and various mixing ratios for moisture like water vapor, cloud water, ice water, rainwater,

Fig. 2 Performance of WRF on HPC XC4000 in Karlsruhe



snow etc.). The resulting size of this task makes it necessary to move to a suitable high performance computing environment.

Software development, testing, benchmarking, and required preprocessing are performed on the KIT, IMK-IFU linux cluster. KIT, IMK-IFU is managing an Infiniband based Linux-Cluster with 116 Opteron, 96 Istan, and 192 Magny processors.

The preprocessing is performed on annual basis. The generated files, which are required to drive WRF, are then transferred to the HPC environment XC4000 via scp. Based on the performed preprocessing WRF is also run in annual time slices, however the WRF restart option enables transient climate simulations, i.e. without initializing WRF every year. WRF and other required software such as netcdf have been successfully installed at the cluster architecture HPC XC4000 using a test account. WRF test runs have been performed using shared memory parallelism (OpenMP).

It is found that 128 processors (32 CPU cores) show the best performance (see Fig. 2). Computing time for the three domains for 1 month was approximately 8 h which results in $8 \times 128 = 1,024$ CPUh.

In order to derive the signal expected from future climate change and climate variability, the future climate projection must be compared against the control run. Beside that, a climate simulation driven by ERA40 reanalysis will be performed to assess the quality of the control run. The simulation efforts can be subdivided into three blocks of 41 years of simulation time plus 4 years of spin-up time.

1. Climate simulations of the control period (1960–2001)
2. Climate simulations for the future scenario A1B (2001–2050)
3. ERA40 reanalysis simulations of the control period (1960–2001)

This means that for each block of 45 years $45 \times 12 \times 1,024 = 552,960$ CPUh are required, which we extend to 553,000 CPUh to have some additional capacities.

For the proposed climate simulations for LUCCi, in total $3 \times 553,000 = 1,659,000$ CPUh are required for the ERA40 reanalysis and control plus the future climate simulations (A1B). With the proposed 128 CPUs this means approximately 180 days pure computing time per 45 years simulation. Thus, we assess around 8 month total simulation time per block (6 month pure computation time and queue/waiting time), which results in a total time of 2 year for the proposal.

For the intermediate storage of the required input data and the WRF output of a short time period we estimate a storage capacity of 2 TB permanent and 2 TB temporary disc space at HP XC2. The results will be transferred to KIT, IMK-IFU. To permanently archive the simulation results a storage capacity of about 30 TB is required.

3 Experimental Parameterization Combinations of WRF

Due to the fact that long-term WRF simulations for Southeast Asia are applied for the first time by KIT, IMK-IFU the performance of climate simulation at a HPC XC4000 requires several preparatory works. These mainly include test simulations with WRF in order to find an optimal setup of the physical parameters, i.e. parameterization for the microphysics, planetary boundary layer, and also cumulus convection. Please note that the cumulus convection scheme is switched off for the 5 km resolution (here: Domain 3). The different parameter combinations applied for the year 2000 are shown in Table 1.

The experimental climate simulations are performed for the year 2000 and validated using gridded observation data for rainfall and temperature. The validation of simulated precipitation fields has been performed using *Asian Precipitation – Highly-Resolved Observational Data Integration Towards Evaluation of the Water Resources* in 0.25° resolution [15], hereinafter referred to as Aphrodite data, and for temperature, CRU TS 2.1 data in 0.5° resolution [8], referred to as CRU data, have been used. As boundary conditions for the RCM, both NCEP/NCAR [3,4] as well as ERA40 [13] reanalysis data have been applied. As the first directive of the regional climate simulations is to match the spatial patterns of precipitation and temperature, the results of the climate simulations of Domain 2 are validated, because they match best with the resolution of the gridded observations. Nevertheless, the climate simulation outputs are regridded to the resolution of gridded observations, i.e. 0.25° and 0.5° for Aphrodite and CRU, respectively. In the following section the climate simulation results of Domain 2 driven by the NCEP/NCAR and ERA40 reanalysis data are validated against the gridded observation data for the year 2000. The results for precipitation and temperature are presented in the sequel. The Taylor diagrams provide a visual framework to compare the simulated (WRF) results against a reference (here: the gridded precipitation product *Aphrodite*).

Table 1 Physical parameterization combinations of the different WRF experimental runs. **Microphysics:** 2 – Lin, 3 – WRF Single-Moment-3-class, 13 – Ston Brook University; **Planetary Boundary Layer (PBL):** 1 – Yonsei University, 5 – Mello-Yamada Nakanishi and Nino Level 2.5 PBL; **Cumulus:** 2 – Betts-Miller-Janjic, 8 – New Simplified Arakawa-Schubert

	Microphysics D1-D3	PBL D1-D3	Cumulus D1-D2
B	2	1	2
C	2	5	2
D	2	5	14
E	2	1	14
F	3	1	2
G	3	5	2
H	3	1	14
I	3	5	14
J	13	1	2
K	13	5	2
L	13	5	14
M	13	1	14

4 Validation Results of WRF Simulations Driven by NCEP/NCAR and ERA40

Figure 3 shows the Taylor diagrams for precipitation (WRF-NCEP/NCAR compared to Aphrodite) obtained for the different physical parametrization setups. Partly high differences between the different setup can be found for precipitation. The spatial correlation patterns strongly depend on the season: the correlation is higher for winter (DJF) and fall (SON), the periods in which most of the rain falls in Southeast Asia. The correlation coefficient is lowest for summer (JJA), the season in which convective rainfall dominates. This season the standard deviation as well as the RMS differences are greatest.

In general, the spatial correlations for temperatures between NCEP/NCAR reanalysis driven WRF simulations and observations (CRU) are high. They are in the order of 0.8, and the different physical parameterizations are not very sensitive for temperature (not shown here).

Comparing the ERA40 driven simulation results, illustrated in Fig. 4, with the results obtained using NCEP/NCAR reanalysis data, one can clearly observe a better representation of the spatial correlations for the summer season while the skill of the fall and winter seasons are reduced. The deviations of single WRF experiments compared to the observations, such as B, C, J, and K (see Table 2) illustrated as RMS differences and standard deviations are drastically increased.

However, the temperature bias is reduced compared to the NCEP/NCAR driven WRF simulation results.

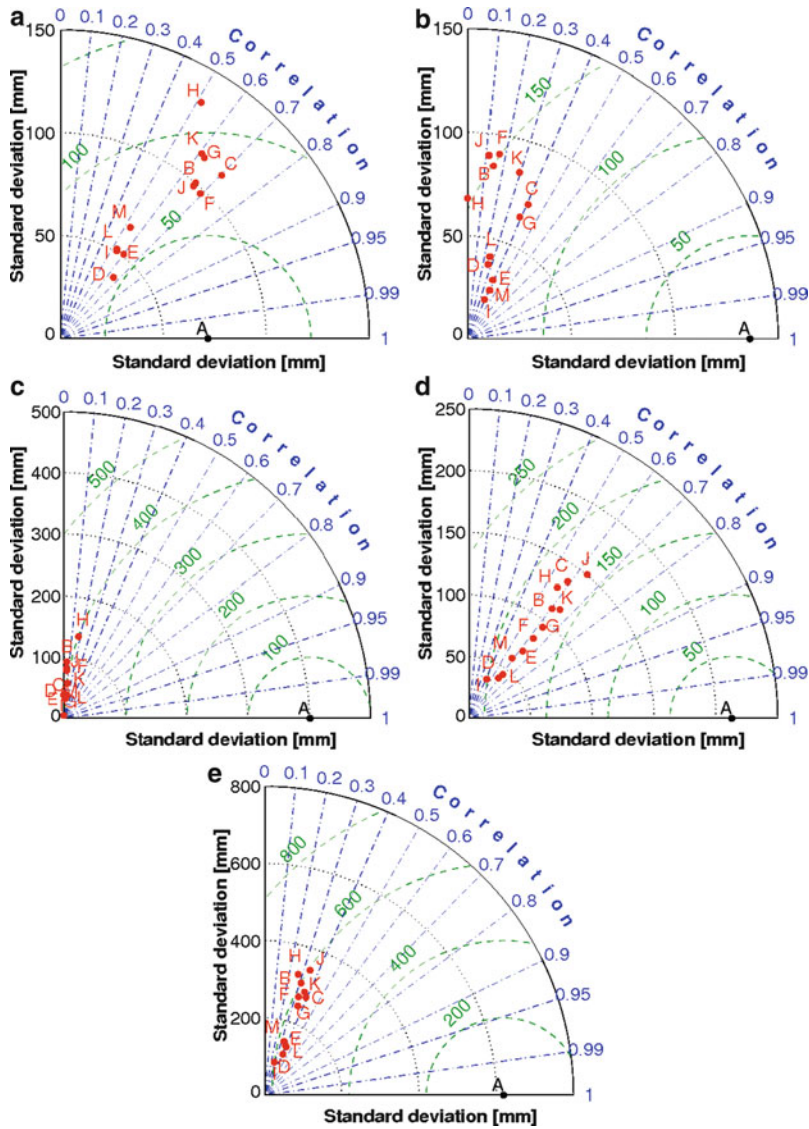


Fig. 3 Taylor diagram of NCEP/NCAR driven WRF simulations and observed (Aphrodite) precipitation amounts for (a) winter 2000 (DJF), (b) spring 2000 (MAM), (c) summer 2000 (JJA), (d) fall 2000 (SON), and (e) the whole year of 2000. The observation data (here: Aphrodite) is shown as A, the coding for the parameter combinations can be obtained from Table 1. The similarity of two patterns is quantified in terms of their correlation (*blue*), their root-mean-square (RMS) difference (*green*), and their standard deviation (*black*)

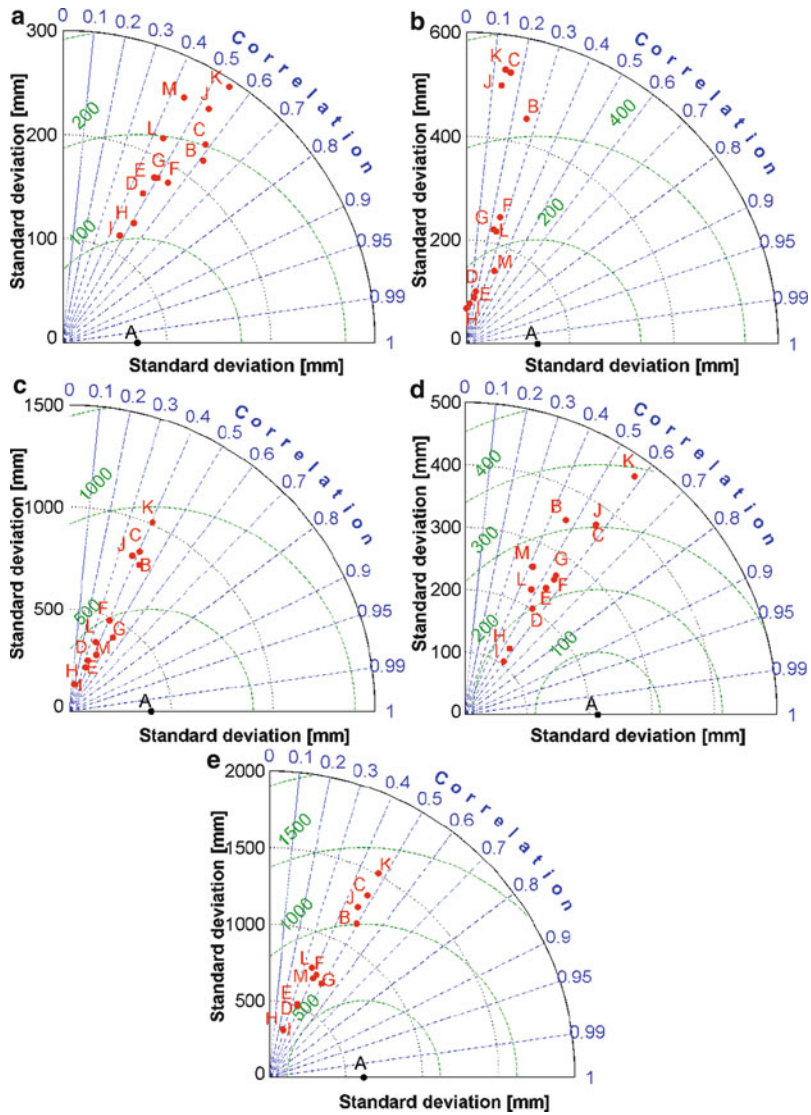


Fig. 4 Taylor diagram of ERA40 driven WRF simulations and observed (Aphrodite) precipitation amounts for (a) winter 2000 (DJF), (b) spring 2000 (MAM), (c) summer 2000 (JJA), (d) fall 2000 (SON), and (e) the whole year of 2000

Table 2 Bias of mean temperatures between NCEP/NCAR reanalysis driven WRF simulations and gridded observations CRU for winter, spring, summer, and fall of the year 2000 and the whole Domain 2. The observation data (here: CRU) is shown in A, the coding for the parameter combinations can be obtained Table 1

	Winter			Spring			Summer			Fall					
	\bar{x}	σ	<i>RMSc</i>	<i>r</i>	\bar{x}	σ	<i>RMSc</i>	<i>r</i>	\bar{x}	σ	<i>RMSc</i>	<i>r</i>			
A	28.22	5.19	-	-	25.63	2.82	-	-	26.71	1.87	-	-	24.71	2.24	-
B	13.01	6.67	4.33	0.76	22.95	3.10	1.87	0.81	24.99	3.00	1.94	0.78	19.11	3.99	2.38
C	11.11	7.29	4.77	0.76	22.38	3.05	2.07	0.75	24.66	3.21	2.09	0.79	19.28	3.99	2.39
D	13.31	6.95	4.22	0.80	22.50	3.16	1.94	0.80	24.76	3.11	2.07	0.76	18.30	4.34	2.83
E	14.66	6.89	4.01	0.81	20.11	3.16	1.98	0.79	20.91	3.08	2.21	0.70	19.44	3.61	2.05
F	10.46	7.18	5.11	0.70	22.07	3.12	1.93	0.79	24.78	3.02	1.94	0.78	17.61	4.39	2.89
G	14.67	6.84	3.93	0.82	22.28	3.21	2.07	0.77	22.56	3.05	2.05	0.75	18.53	4.08	2.52
H	14.93	6.54	3.69	0.83	20.37	3.21	1.91	0.81	17.81	3.16	2.28	0.70	16.75	3.83	2.38
I	13.41	6.86	4.14	0.80	19.88	3.33	2.08	0.78	17.79	3.43	2.54	0.69	16.20	4.29	2.66
J	12.18	6.56	4.33	0.75	22.70	3.13	1.89	0.80	25.69	3.20	2.08	0.79	19.64	3.82	2.22
K	13.96	7.06	4.47	0.78	22.82	3.07	2.00	0.77	24.04	3.08	2.00	0.78	17.61	4.25	2.77
L	13.45	6.89	4.20	0.79	22.33	3.18	1.92	0.80	24.44	3.15	2.08	0.77	18.48	4.15	2.63
M	14.54	6.99	4.13	0.81	20.30	3.31	1.95	0.81	23.60	3.38	2.36	0.74	16.03	4.14	2.67

Table 3 ERA40-CRU (winter, spring, summer, and fall). The observation data (here: CRU) is shown in A, the coding for the parameter combinations can be obtained Table 1

	Winter			Spring			Summer			Fall			
	\bar{x}	σ	<i>RMSc</i>	\bar{x}	σ	<i>RMSc</i>	\bar{x}	σ	<i>RMSc</i>	\bar{x}	σ	<i>RMSc</i>	<i>r</i>
A	28.22	5.19	-	25.63	2.82	-	26.71	1.87	-	24.71	2.24	-	-
B	17.89	6.83	3.37	23.12	2.95	1.66	24.76	2.24	1.38	21.24	3.08	1.60	0.87
C	17.16	6.71	3.41	23.07	2.90	1.67	24.67	2.23	1.38	20.72	3.14	1.67	0.86
D	17.34	6.61	3.27	22.88	2.95	1.62	24.16	2.40	1.41	21.06	3.30	1.73	0.87
E	18.17	6.82	3.29	22.31	3.12	1.79	24.96	2.47	1.37	20.89	3.36	1.81	0.87
F	17.38	6.89	3.59	23.24	2.96	1.66	24.81	2.32	1.36	21.10	3.25	1.73	0.87
G	17.18	6.62	3.31	22.94	2.94	1.69	24.80	2.24	1.31	20.60	3.17	1.70	0.86
H	17.31	6.77	3.52	23.20	3.05	1.54	26.99	3.00	1.72	21.92	3.40	1.76	0.89
I	17.44	6.52	3.16	24.05	3.18	1.61	26.02	2.75	1.52	21.44	3.34	1.73	0.88
J	18.10	6.80	3.31	23.39	2.96	1.65	24.88	2.16	1.35	21.37	3.06	1.57	0.87
K	17.59	6.67	3.28	23.15	2.87	1.68	24.73	2.15	1.35	20.85	3.08	1.61	0.86
L	17.48	6.74	3.40	23.16	3.12	1.87	24.38	2.17	1.28	18.96	3.35	2.06	0.80
M	17.44	6.52	3.16	24.05	3.18	1.61	26.02	2.75	1.52	21.44	3.34	1.73	0.88

Amongst all test parameter combinations it is found that combination G (see Table 1), i.e. WRF Single-Moment-3-class, Mello-Yamada Nakanishi and Nino Level 2.5 PBL, and Betts-Miller-Janjic is leading to reasonable results for both precipitation and temperature values. Due to the lower biases in temperature (especially during the winter months) it is decided for the ERA40 reanalysis as driving dataset.

5 Analysis of the Temperature Biases

Albeit the spatial patterns for temperature between observed and simulated temperatures are matched well, the large biases (Table 2 and Table 3) for the winter and fall season require further investigations. As all parameterization combinations are leading to underestimated temperatures of the WRF simulations within Domain 2, it has been speculated that the CRU data might overestimate temperature for this region. Figure 5 shows the differences between the gridded observation datasets CRU [8], DEL [6], and GLDAS [10] and the ERA Interim Reanalysis [1, 11] for winter of the year 2000. Comparing CRU with alternative gridded observation datasets it is observed that CRU significantly overestimates the temperature during winter. Especially for Vietnam, this overestimation is obvious. Figure 6 illustrates the differences of the different datasets averaged over the grid cells corresponding to Vietnam. It is speculated that the overestimated temperatures of the CRU data result from the coarse network of observation stations within this topographically complex region. These observations are interpolated based on the surrounding eight stations. First inspection of the CRU derived interpolated elevation model based on the heights of the eight surrounding observation stations shows party strong deviations compared to the elevation model as used for the WRF model in Domain 2 (not shown). As a consequence, the temperature estimates of the CRU grid cells can deviate significantly from model results. Comparing CRU with other observation datasets in regions with similar complex terrain and low observation densities could help to judge the quality of the CRU dataset. However, it remains unclear why this deviation is stronger during winter and fall than during the other seasons. A more detailed analysis is required.

6 Ongoing Activities and Future Work

The operational long-term climate simulations are performed on monthly time slices using the restart option of WRF. This means that the results of the last time step of the previous month will be temporarily stored and used as input for the

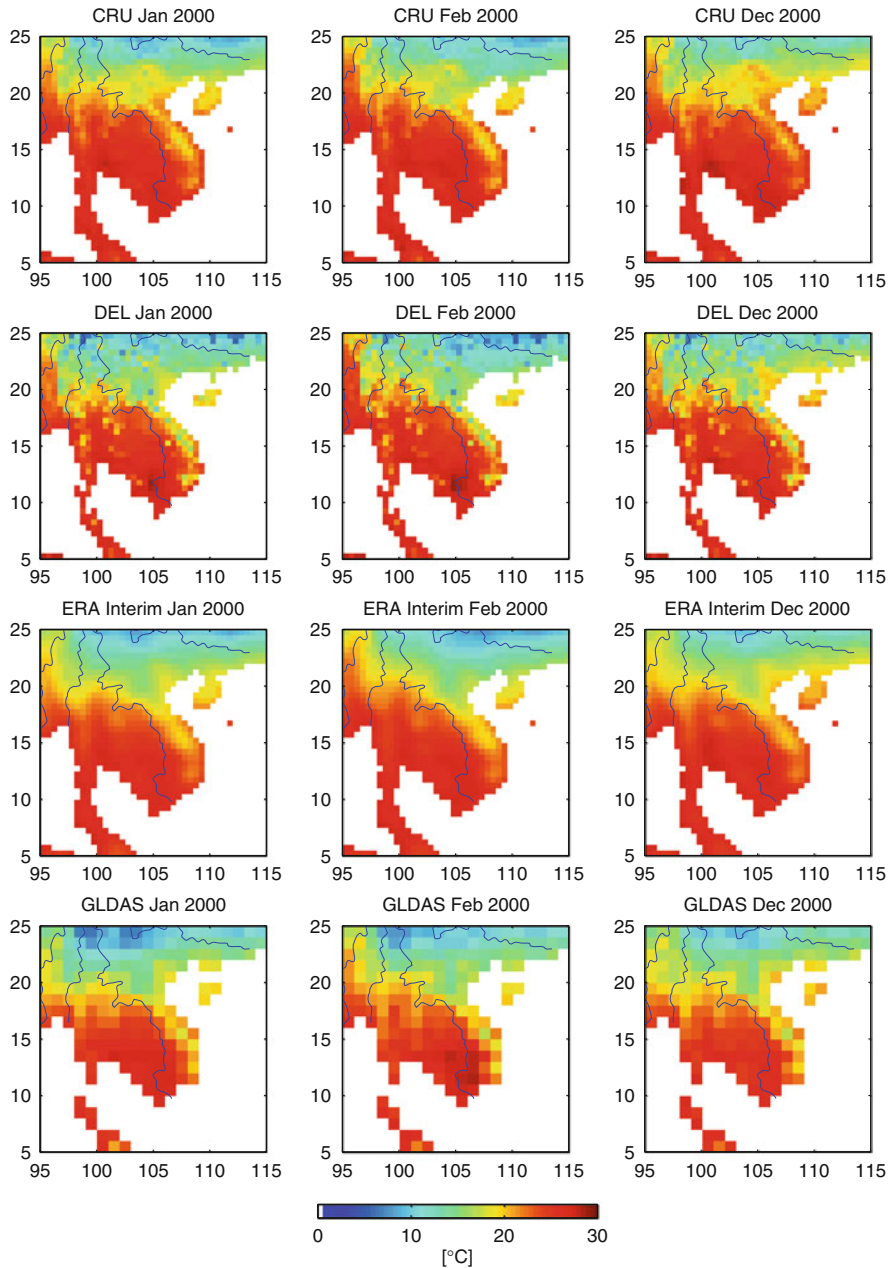


Fig. 5 Mean temperature ($^{\circ}$ C) for January, February, and December (2000) for the datasets CRU, DEL, ERA Interim, and GLDAS




Fig. 6 Time series of monthly temperature for Vietnam (average values for the grid cells corresponding to Vietnam) using the different datasets CRU, DEL, ERA Interim and GLDAS

subsequent month. For the continuous performance of the climate simulations batch processing will be applied. The model output will be stored in a 6 hourly resolution. Many surface climate variables are additionally stored in hourly resolution to meet specific requirements of the LUCCi project consortium. Besides the instantaneous variables, retained in a 6 hourly and hourly resolution, additional meteorological surface variables containing the magnitude and timing of the actual minimum and maximum values of a day are retained. This will allow for detailed analysis of extreme events lateron.

It is expected that the climate simulation will be finalized by end of 2012. The results will be provided to the LUCCi project consortium, but results can be provided on request to the climate change and climate impact community.

Acknowledgements This research is funded by the Federal Ministry of Education and Research (research project: Land Use and Climate Change Interactions in Central Vietnam (LUCCi), reference number 01LL0908C). The provision of CPU and storage resources at Karlsruhe Institute of Technology (KIT), Steinbuch Centre for Computing (SCC) and Karlsruhe Institute of Technology (KIT), Institute of Meteorology and Climate Research (IMK-IFU) is highly acknowledged. Technical support has been given by Hartmut Häfner, Johannes Werhahn, Sven Wagner and Benjamin Fersch.

References

1. Berrisford P, Dee D, Fielding K, Fuentes M, Kallberg P, Kobayashi S and Uppala S (2009) The ERA Interim archive Version 1.0. Era report series, ECMWF, 16 pp. Available electronically at www.ecmwf.int.
2. Cohen SJ (1990) Bringing the global warming issue closer to home: the challenge of regional impact studies. *Bulletin of the American Meteorological Society* 71:520–526
3. Kalnay et al. (1996) The NCEP/NCAR 40-year reanalysis project, *Bull. Amer. Meteor. Soc.* 77:437–470
4. Kistler et al. (2001) The NCEP-NCAR 50-Year Reanalysis: Monthly means CD-ROM and Documentation. *Bull. Amer. Meteor. Soc.* 82:247–267.

5. Klemp JB, Skamarock WC, Dudhia J (2007) Conservative split-explicit time integration methods for the compressible nonhydrostatic equations. *Mon. Wea. Rev.* 135:2897–2913
6. Matsuura K and Willmott CJ (2009) Terrestrial Air Temperature: 1900–2008 Gridded Monthly Time Series (Version 2.01). Center for Climatic Research, University of Delaware, available electronically at <http://climate.geog.udel.edu/~climate/>.
7. Mearns LO, Rosenzweig C, Goldberg R (1997) Mean and Variance Change in Climate Scenarios: Methods, Agricultural Applications, and Measures of Uncertainty. *Climatic Change* 35:367–396
8. Mitchell TD, Jones PD (2005) An improved method of constructing a database of monthly climate observations and associated high-resolution grids. *International Journal of Climatology* 25:693–712
9. Phan VT, Ngo-Duc T, Ho TMH (2009) Seasonal and interannual variations of surface climate elements over Vietnam. *Journal of Climate Research* 40(1):49–60
10. Rodell M, Houser PR, Jambor U, Gottschalck J, Mitchell K, Meng C-J, Arsenault K, Cosgrove B, Radakovich J, Bosilovich M, Entin JK, Walker JP, Lohmann D and Toll D (2004) The Global Land Data Assimilation System. *Bull. Amer. Meteor. Soc.*, 85(3):381–394
11. Simmons A, Uppala S, Dee D and Kobayashi S (2006) ERA-Interim: New ECMWF reanalysis products from 1989 onwards. *ECMWF Newsletter*, 110, 25–35, available electronically at www.ecmwf.int.
12. Smiatek G, Kunstmann H, Knoche R, Marx A (2009) Precipitation and temperature statistics in high-resolution regional climate models: Evaluation for the European Alps. *Journal of Geophysical Research*, 114:D19107. doi:10.1029/2008JD011353
13. Uppala et al. (2005) The ERA-40 re-analysis. *Q.J.R. Meteorol. Soc.*, 131:2961–3012. doi:10.1256/qj.04.176
14. Von Storch H, Zorita E, Cubasch U (1993) Downscaling of global climate change estimates to regional scales: An application to Iberian rainfall in wintertime, *J. Climate*, 6:1161–1171
15. Yatagai A, Arakawa O, Kamiguchi K, Kawamoto H, Nodzu MI, Hamada A (2009) A 44-Year Daily Gridded Precipitation Dataset for Asia Based on a Dense Network of Rain Gauges. *SOLA* 5:137–140

The Agulhas System as a Key Region of the Global Oceanic Circulation

J.V. Durgadoo and A. Biastoch

1 Introduction

The oceans around southern Africa form a unique system, impacting the regional and global climate [1]. From the Indian Ocean to the Atlantic Ocean vigorous interoceanic exchange of warm and saline waters takes place that is subject to a complicated interplay between local dynamics and global embedment. Central element of the circulation around South Africa is the Agulhas Current [2] that flows poleward along the east coast, closely bound to the shelf at first, and subsequently overshoots the southern tip of Africa to abruptly turn back into the Indian Ocean. Part of the warm and saline waters with tropical Indian Ocean origin, the “Agulhas leakage” [3], flows into the Atlantic and forms the surface return flow of the global thermohaline circulation towards the North Atlantic [4]. The exchange takes place in a highly nonlinear manner, with mesoscale eddies being separated from the retroflecting Agulhas Current, which then strongly interact in the Cape Basin [5]. West of the Cape Basin, large Agulhas rings that have been formed [6] transport the anomalous warm and saline waters into the South Atlantic. In addition to its own dynamics, the Agulhas Current system is influenced by nonlinearities in the source regions: mesoscale eddies originating from the Mozambique Channel and east of Madagascar [7, 8] drift southward and cause the Agulhas Current to be displaced offshore of its mean position by more than 100 km. These solitary meanders (a.k.a. “Natal Pulses”) [9] rapidly propagate downstream triggering the timing of Agulhas rings [10, 11].

Owing to strong eddy mean flow interactions, the Agulhas leakage is a highly intermittent process. The dynamic imprint of the Agulhas mesoscale variability crosses the South Atlantic via baroclinic Rossby waves and is then rapidly

J.V. Durgadoo (✉) · A. Biastoch
GEOMAR Helmholtz Zentrum für Ozeanforschung Kiel, Düsternbrooker Weg 20, 24105 Kiel,
Germany
e-mail: jdurgadoo@geomar.de

communicated via topographic waves along the South American shelf into the North Atlantic, an effect that is visible in the Atlantic Meridional Overturning Circulation (AMOC) [12]. In contrast, the advection of the anomalous water mass characteristics takes longer [13]. Experiments with hindcast experiments over the past 40 years have demonstrated that the volumetric amount of Agulhas leakage is subject to increase due to changing wind systems in the Indian and Southern oceans [14, 15]. This has an important consequence since the additional salt import into the Atlantic will find its way to the north and has the general potential to stabilize a declining Gulf Stream system [16].

Despite its global importance, the dynamics of the Agulhas system is not yet fully understood and, owing to the strong nonlinear interactions and temporal variability, is far from being properly quantified. The notorious undersampling of the oceans around South Africa calls for dedicated studies within ocean general circulation models and emphasizes the importance of a reasonable representation of the current system in both modern and future coupled-climate models.

2 INALT01

The scarcity of continuous direct observations presents a general challenge to oceanographic research and this is particularly true in the Agulhas region. Consequently a modelling approach remains the most useful way to study this important area of the world's ocean. The model used is based on the NEMO (Nucleus for European Modelling of the Ocean, v.3.1.1, [17]) code within the DRAKKAR [18] framework. NEMO is derived from the Navier-Stokes set of primitive equations as well as the nonlinear equation of state, coupling temperature and salinity to the velocity of the field. Assumptions are made (e.g. incompressibility of flow, hydrostatic hypothesis, Boussinesq fluid, among others) based on scale analysis. Being an ocean/sea-ice only model, boundary conditions at the surface need to be specified. This is achieved through bulk-formulae, where the pseudo-atmosphere and ocean exchange the necessary horizontal momentum and heat to drive the ocean. The sea-ice component employs a LIM2 two-level thermodynamic-dynamic model [19]. The configuration of NEMO used in this project (ORCA05, AG01 and INALT01) consists of a tri-polar horizontal grid (poles over Antarctica, Canada and Russia) with variable arranged on an Arakawa C-Grid. In the vertical, 46 z-levels are used with partial cells for a better representation of bottom flow.

The high-resolution AG01 model [11, 12, 14, 20] has been amply used over the last years and has proven a success (viz. Sect. 3). However, it does have some limitations. The western extent of the refinement of the circulation only spans half of the South Atlantic basin. Because of this restriction, it is impossible, within this framework, to clearly distinguish the far-reaching impact of the Agulhas Leakage – a task that is of key relevance and has become of subject of active research. Furthermore, in addition to the limited time period of the hind-cast simulation (which is determined by




Fig. 1 Snapshot from INALT01 showing temperature (°C, color-scheme) and currents (gradients) at 450 m. The INALT01 configuration consists of an ocean/sea-ice global coarse resolution (0.5°) base model with a nest model embedded within, refining the horizontal grid (0.1°) over the greater Agulhas Current region, the South and Tropical Atlantic Ocean

available atmospheric fields), AG01 is based on a previous version of the numerical code. In order to effectively address the above-mentioned limitations, INALT01 (Fig. 1) was conceived. INALT01, inheriting from the experience of AG01, provides a high-resolution nest over the entire South Atlantic basin and the tropical Atlantic, with a hind-cast simulation spanning six decades.

INALT01 comprises of two ocean models that are nested together; a global (base) model that represent the major oceanic components of the circulation with a nominal horizontal resolution of 50 km and a refined (nest) model that regionally enhances the horizontal resolution to 10 km (Fig. 2). The latter model is crucial for the correct representation of the Agulhas dynamics and being embedded within the former offers the possibility to diagnose any possible impact the Agulhas system may have on region of the oceans where climatically important changes are currently being observed.

Running such a model configuration is computationally extremely demanding. With over 62 million combined grid points, both base and nest models of INALT01 run concurrently by decomposing each respective domains horizontally onto 16 processors. Including post-processing, a model year requires computation time of approximately 12–15h on the NEC SX-9 at HLRS, depending on the time-




Fig. 2 Bathymetry (colorbar in meters) illustrating the refinement of the nest over the greater Agulhas region and South Atlantic compared to the global base host

resolution of the output fields. A typical model year with 5-daily output fields generate 150 GB of raw data, a full 60-year simulation requiring up to 9 TB. The high spatial and temporal resolutions are essential for a correct representation of the range of dynamics of the Agulhas Current system.

3 Results

3.1 *On the Discontinuous Nature of the Mozambique Current*

The concept of a spatially continuous western boundary current in the Mozambique Channel has historically been based on erroneous interpretations of ships' drift. Recent observations have demonstrated that the circulation in the Channel is instead dominated by anti-cyclonic eddies drifting poleward. It has therefore been suggested that no coherent Mozambique Current exists at any time. However, satellite and other observations indicate that a continuous current not necessarily an inherent part of Mozambique Eddies may at times be found along the full Mozambican shelf break. Using the high-resolution AG01 it has been demonstrated how such a feature may come about [21]. In the model, a continuous current is a highly irregularly occurring event, occurring about once per year, with an average duration of only 9 days and with a vertical extent of about 800 m (Fig. 3). Surface speeds may vary from 0.5 to 1.5 m/s and the volume flux involved is about 10 Sv. The continuous current may occasionally be important for the transport of biota along the continental shelf and slope.

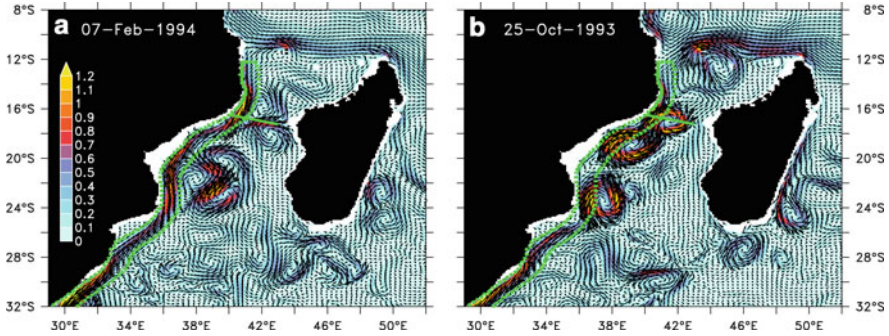


Fig. 3 Simulated velocities in the Mozambique Channel at a depth of 93 m in the model for (a) 07 February 1994 and (b) 25 October 1993. (b) Shows the normal current configuration in which three strong anti-cyclonic eddies are evident, all heading poleward. In contrast, (a) presents the more unusual case of a continuous current along the full length of the Mozambican shelf; this current on this occasion stayed intact for about 5 days. *Arrows* indicate current directions [21]




Fig. 4 Wind shift and Agulhas leakage response. (a) Ensemble mean of zonally averaged wind stress (zonal component only) in the climate model under present-day (*black line*) and global warming (*blue*) conditions. The difference function (*thick*) was applied to the sensitivity experiments. (b) Annual Agulhas leakage for reference (*dark bars*) and sensitivity (*light bars*) experiments

3.2 Anthropogenic Impact on Agulhas Leakage

Climate model projections for the 21st century predict a progressive southward migration of the Southern Hemisphere westerlies, associated with a southward shift in the latitude of zero wind stress curl [22]. The potential effects on the ocean circulation of such an anthropogenic trend in wind stress are studied here with a sequence of experiments with the high-resolution ocean model. The model suggests an increase of 5 Sv in Agulhas leakage (Fig. 4) in response to a 2°-shift in the

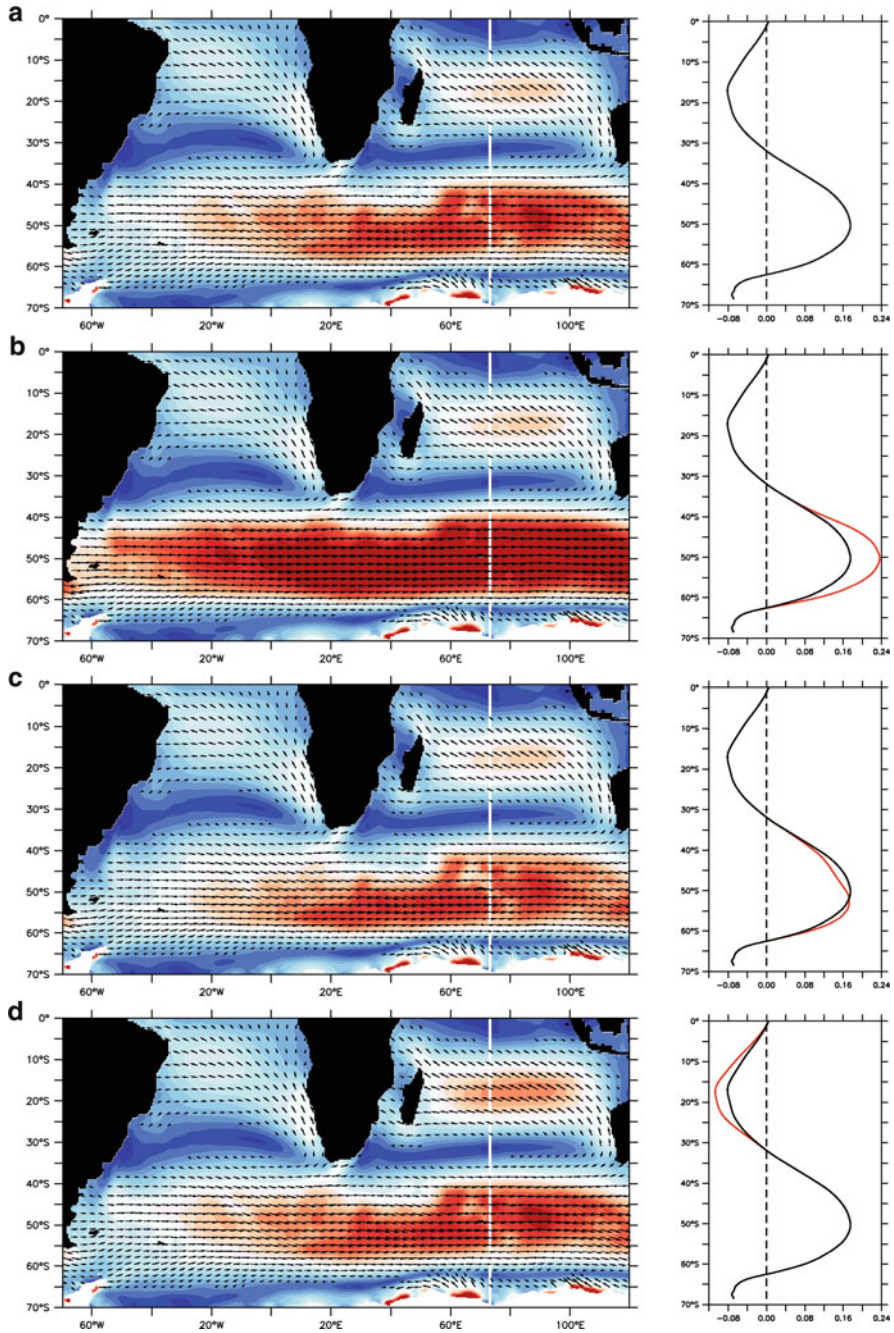


Fig. 5 Wind stress patterns (*left panels*) over the Atlantic and Indian Ocean sectors of the Southern Hemisphere and their corresponding zonal averages (*right panels*). Examples of sensitivity cases are shown whereby the reference winds (*top panel*) are independently nudged towards an increase (*lower top panel*) or a shift (*upper bottom panel*) of the Westerly winds or an increase of the Trade winds (*bottom panel*); black and red curves representing reference and sensitivity respectively

zero wind stress curl line, associated with a southward expansion and spinup of the subtropical supergyre. The integral northward transport in the upper branch of the AMOC gradually increases by up to 1.5 Sv at 20–25S; this dynamic signal extends north across the equator and fades out in the subtropical North Atlantic. A main effect of the increasing inflow of Indian Ocean waters is the salinification of upper-thermocline waters in the South Atlantic which extends into the North Brazil Current regime within just one and two decades.

4 On-Going Work

The importance of the Agulhas Leakage has been highlighted in various published research [1, 4]. Over the last few decades, there has been a noteworthy increase in the magnitude of Indian Ocean waters leaking in to Atlantic [14, 15]. The cause of this increase has been and remains under debate owing to the intricate dynamics resulting in Leakage south of Africa. One of the external factors potentially determining the magnitude of Agulhas Leakage on various timescales is the Southern Hemisphere wind patterns, namely the Trades winds and the Westerlies. On-going work seeks to address this complex issue by applying smooth anomalies to the current wind pattern to effectively nudge the average winds to a different state and so-doing, systematically with each anomaly applied in isolation (Fig. 5). Monitoring changes in Agulhas Leakage as a result of this detangling of individual impacts of these two wind belts helps in understanding the underlying dynamics.

References

1. L. M. Beal, W. P. M. de Ruijter, A. Biastoch, R. Zahn, and members of SCOR/WCRP/IAPSO Working Group 136. On the role of the Agulhas system in global climate. *Nature*, 472:429–436, doi: 10.1038/nature09983, 2011.
2. J. R. E. Lutjeharms. *The Agulhas Current*. Springer, Berlin, 2006.
3. W. P. M. De Ruijter, A. Biastoch, S. S. Drijfhout, J. R. E. Lutjeharms, R. Matano, T. Pichevin, P. J. van Leeuwen, and W. Weijer. Indian-Atlantic inter-ocean exchange: Dynamics, estimation and impact. *J. Geophys. Res.*, 104:20,885–20,910, 1999.
4. A. L. Gordon. Oceanography: The brawniest retroreflection. *Nature*, 421:904–905, 2003.
5. O. Boebel, J. Lutjeharms, C. Schmid, W. Zenk, T. Rossby, and C. Barron. The Cape Cauldron: a regime of turbulent inter-ocean exchange. *Deep-Sea Res. II*, 50:57–86, 2003.
6. D. B. Olson and R. H. Evans. Rings of the Agulhas Current. *Deep-Sea Res.*, 33:27–42, 1986.
7. A. Biastoch and W. Krauss. The role of mesoscale eddies in the source regions of the Agulhas Current. *J. Phys. Oceanogr.*, 29:2303–2317, 1999.
8. M. W. Schouten, W. P. M. de Ruijter, P. J. van Leeuwen, and H. Ridderinkhof. Eddies and variability in the Mozambique channel. *Deep-Sea Res. II*, 50:1987–2004, 2003.
9. W. P. M. De Ruijter, P. J. van Leeuwen, and J. R. E. Lutjeharms. Generation and evolution of Natal Pulses, solitary meanders in the Agulhas Current. *J. Phys. Oceanogr.*, 29:3043–3055, 1999.

10. M. W. Schouten, W. P. M. de Ruijter, and P. J. van Leeuwen. Upstream control of Agulhas Ring shedding. *J. Geophys. Res.*, 107(10.1029), 2002.
11. A. Biastoch, J. R. E. Lutjeharms, C. W. Böning, and M. Scheinert. Mesoscale perturbations control inter-ocean exchange south of Africa. *Geophys. Res. Lett.*, 35:L20602, 2008.
12. A. Biastoch, C. W. Böning, and J. R. E. Lutjeharms. Agulhas leakage dynamics affects decadal variability in Atlantic overturning circulation. *Nature*, 456:489–492, 2008.
13. E. Van Sebille, L. M. Beal, and W. E. Johns. Advective time scales of Agulhas leakage to the North Atlantic in surface drifter observations and the 3D OFES model. *J. Phys. Oceanogr.*, 41(5):1026–1034, 2011.
14. A. Biastoch, C. W. Böning, F. U. Schwarzkopf, and J. R. E. Lutjeharms. Increase in Agulhas leakage due to poleward shift of the Southern Hemisphere westerlies. *Nature*, 462:495–498, 2009.
15. M. Rouault, P. Penven, and B. Pohl. Warming in the Agulhas Current system since the 1980's. *Geophys. Res. Lett.*, 36:L12602, 2009.
16. J. M. Gregory, K. W. Dixon, R. J. Stouffer, A. J. Weaver, E. Driesschaert, M. Eby, T. Fichefet, H. Hasumi, A. Hu, J. H. Jungclaus, I. V. Kamenkovich, A. Levermann, M. Montoya, S. Murakami, S. Nawrath, A. Oka, A. P. Sokolov, and R. B. Thorpe. A model intercomparison of changes in the Atlantic thermohaline circulation in response to increasing atmospheric CO₂ concentration. *Geophys. Res. Lett.*, 32:doi: 10.1029/2005GL023209, 2005.
17. G. Madec. Nemo ocean engine. Technical Report 27, Note du Pole de modelisation, Institut Pierre Simon Laplace (IPSL), France, 2006.
18. The DRAKKAR Group. Eddy-Permitting Ocean Circulation Hindcasts of Past Decades. *Clivar Exchanges*, 12:8–10, 2007.
19. T. Fichefet and M. A. Morales Maqueda. Sensitivity of a global sea ice model to the treatment of ice thermodynamics and dynamics. *J. Geophys. Res.*, 102(C6):12609–12, 1997.
20. A. Biastoch, L. Beal, T. G. D. Casal, and J. R. E. Lutjeharms. Variability and coherence of the Agulhas Undercurrent in a high-resolution ocean general circulation model. *J. Phys. Oceanogr.*, 39:2417–2435, 2009.
21. J. R. E. Lutjeharms, A. Biastoch, P. M. van der Werf, H. Ridderinkhof, and W. P. M. de Ruijter. On the discontinuous nature of the Mozambique Current. *S. Afr. J. Sci.*, 108(1/2):5–pages, 2012.
22. A. Sen Gupta, A. Santoso, A. S. Taschetto, C. C. Ummenhofer, M. H. England, and J. Trevena. Projected changes to the Southern Hemisphere ocean and sea-ice in the IPCC AR4 climate models. *J. Clim.*, 22:3047–3078, 2009.

Numerical Investigation of Stratified Turbulence

S. Remmler and S. Hickel

Abstract In order to evaluate turbulence subgrid-scale models for stably stratified flows, we performed direct numerical simulations (DNS) of homogeneous stratified turbulence with large-scale horizontal forcing. In these simulations we found that energy dissipation is concentrated within thin layers of horizontal tagliatelle-like vortex sheets between large pancake-like structures. For large eddy simulation (LES), we use an implicit subgrid-scale model, based on the Adaptive Local Deconvolution Method (ALDM). Our analysis proves that the implicit turbulence model ALDM correctly predicts the turbulence energy budget and the energy spectra of stratified turbulence, even though dissipative structures are not resolved on the computational grid.

1 Introduction

To predict atmospheric and oceanic mesoscale flows, we need to understand and parametrize small scale turbulence that is strongly affected by the presence of stable density stratification. The stratification suppresses vertical motions and thus makes all scales of the velocity field strongly anisotropic. Using aircraft observations, the horizontal velocity spectrum in the atmosphere was analyzed by Nastrom and Gage [23]. They found a power-law behavior in the mesoscale range with an exponent of $-5/3$. In the vertical spectrum, on the other hand, Cot [4] observed an exponent of -3 in the inertial range.

There has been a long and intensive discussion whether the observed spectra are due to a backward cascade of energy [8, 10, 18] as in two-dimensional turbulence [15], or due to breaking of internal waves, which means that a forward cascade is the

S. Remmler (✉) · S. Hickel
Technische Universität München, Institute of Aerodynamics and Fluid Mechanics,
85747 Garching bei München, Germany
e-mail: remmler@tum.de

dominant process [5, 32]. In different numerical and theoretical studies, ambiguous or even conflicting results were obtained [19].

During the last decade, a number of new simulations and experiments addressed the issue. Smith and Waleffe [28] observed a concentration of energy in the lowest modes in their simulations. Other studies [16, 33] suggested that the character of the flow depends on the Reynolds number. Apparently, high Reynolds numbers are associated with stronger three-dimensionality and a forward cascade of energy. Riley and de Bruyn Kops [25] suggested that the flow can be strongly stratified but still turbulent if $\text{Fr}^2\text{Re} > 1$. Lindborg [20] presented a scaling analysis of the Boussinesq equations for low Froude and high Reynolds number. His theory of strongly anisotropic, but still three-dimensional, turbulence explains the horizontal $k_h^{-5/3}$ spectrum as well as the vertical k_v^{-3} spectrum. On the basis of these findings, Brethouwer et al. [2] showed that the relevant non-dimensional parameter controlling stratified turbulence must indeed be the buoyancy Reynolds number $\mathcal{R} = \text{Fr}^2\text{Re}$. For $\mathcal{R} \gg 1$, they predict stratified turbulence including local overturning and a forward energy cascade. In the opposite limit, for $\mathcal{R} \ll 1$, the flow is controlled by viscosity and does not contain small-scale turbulent motions. A detailed analysis of the spectral structure and spectral energy budget of homogeneous stratified turbulence based on direct numerical simulations is provided by Remmler and Hickel [24].

Since a full resolution of all turbulence scales is only possible for very low Reynolds numbers, many groups performed large eddy simulations (LES), which rely on a subgrid-scale model to represent effects of unresolved small-scale turbulence. For example, Metais and Lesieur [22] used a spectral eddy viscosity model, based on the eddy damped quasi-normal Markovian (EDQNM) theory. This required a flow simulation in Fourier space and the cut-off wavenumber to be in the inertial range. For LES in physical space, Smagorinsky models are widely used, either in the classical formulation [14] or with certain modifications for stratified turbulence that are usually based on the local Richardson number [6]. Better results can be obtained if the model constant of a Smagorinsky model is not prescribed, but computed by the dynamic procedure of Germano et al. [9]. This approach was successfully applied to stably stratified turbulent channel flow by Taylor et al. [30] and others. Staquet and Godeferd [29] presented a two-point closure statistical EDQNM turbulence model, which was adapted for axisymmetric spectra about the vertical axis. Recently, many groups presented regularized direct numerical simulations (DNS) of stratified turbulence, which means rather pragmatically stabilizing under-resolved DNS by removing the smallest resolved scales. This is usually achieved by a hyperviscosity approach [20] or by de-aliasing in spectral methods using the “2/3-rule” [1, 7].

In practice, all SGS turbulence models suffer from the problem that the computed SGS stresses are of the same order as the grid truncation error. This typically leads to interference between SGS model and numerical scheme, that can manifest in instability and lack of grid convergence. This issue can be solved by combining discretization scheme and SGS model in a single approach. This is usually referred

to as “implicit” LES (ILES) in contrast to the traditional “explicit” SGS models. The idea of physically consistent ILES was realized by Hickel et al. [11] in the Adaptive Local Deconvolution Method (ALDM) for neutrally stratified fluids. ALDM provides a framework for the design, analysis and optimization of numerical discretizations with an implicit SGS model that is consistent with turbulence theory. Based on this method and ALDM for passive scalar transport [12], we developed an implicit SGS model for Boussinesq fluids. In the present paper, we evaluate the applicability of ALDM for stably stratified turbulence.

We simulated forced homogeneous stratified turbulence in a triple-periodic box at different Froude and Reynolds numbers. We present not only ILES, but also LES with a standard Smagorinsky model (SSM) and a dynamic Smagorinsky model (DSM) as well as high-resolution DNS as benchmark solutions.

2 Governing Equations

The non-dimensional Boussinesq equations for a stably stratified fluid in Cartesian coordinates read

$$\nabla \cdot \mathbf{u} = 0 \quad (1a)$$

$$\partial_t \mathbf{u} + \nabla \cdot (\mathbf{u}\mathbf{u}) = -\nabla p - \frac{\rho}{\text{Fr}_0^2} \mathbf{e}_z + \frac{1}{\text{Re}_0} \nabla^2 \mathbf{u} \quad (1b)$$

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = -w + \frac{1}{\text{Pr Re}_0} \nabla^2 \rho \quad (1c)$$

where velocities $\mathbf{u} = [u, v, w]$ are made non-dimensional by \mathcal{U} , all spatial coordinates by the length scale \mathcal{L} , pressure by \mathcal{U}^2 , time by \mathcal{L}/\mathcal{U} , and density fluctuation $\rho = \rho^* - \bar{\rho}$ (ρ^* : local absolute density, $\bar{\rho}$: background density) by the background density gradient $\mathcal{L}|d\bar{\rho}/dz|$. The non-dimensional parameters are

$$\text{Fr}_0 = \frac{\mathcal{U}}{N\mathcal{L}}, \quad \text{Re}_0 = \frac{\mathcal{U}\mathcal{L}}{\nu}, \quad \text{Pr} = \frac{\nu}{\alpha}, \quad (2)$$

where $N = \sqrt{g/\rho_0 d\bar{\rho}/dz}$ is the Brunt-Väisälä frequency, ν is the kinematic viscosity and α is the thermal diffusivity. We chose a Prandtl number of $\text{Pr} = 0.7$, corresponding to values found in the atmosphere. Froude and Reynolds number are the parameters that control the flow regime.

With the instantaneous values of kinetic energy E_k and kinetic energy dissipation ε_k , we find the local Froude and Reynolds number as well as the buoyancy Reynolds number \mathcal{R} , defined by Brethouwer et al. [2]:

$$\text{Fr} = \frac{\text{Fr}_0 \mathcal{L}}{\mathcal{U}} \frac{\varepsilon_k}{E_k}, \quad \text{Re} = \frac{\text{Re}_0}{\mathcal{U} \mathcal{L}} \frac{E_k^2}{\varepsilon_k}, \quad \mathcal{R} = \text{Re Fr}^2 \quad (3)$$

In ILES by construction we do not have direct access to the value of ε_k , as only a small part of it is resolved. For LES, we thus estimate ε_k from the total energy balance

$$\partial_t \langle E_t \rangle = \partial_t \langle E_k \rangle + \partial_t \langle E_p \rangle = P - \varepsilon_k - \varepsilon_p = P - (1 + \Gamma)\varepsilon_k, \quad (4)$$

where the temporal change of total energy in the flow $\partial_t \langle E_t \rangle$ can be computed from the energy levels at subsequent time steps, P is the power inserted into the system by the external forcing and $\Gamma = \varepsilon_p/\varepsilon_k$ is the mixing ratio assumed to be constant $\Gamma = 0.4$, which is an acceptable approximation for a wide range of parameters [24]. Here, ε_p is the dissipation rate of potential energy. The kinetic energy dissipation rate can then be computed from

$$\varepsilon_k = \frac{1}{1 + \Gamma} (P - \partial_t \langle E_t \rangle). \quad (5)$$

3 Numerical Method

3.1 Flow Solver

Our finite-volume solver INCA offers different discretization schemes depending on the application. For DNS and LES with SSM and DSM, we used a non-dissipative central difference scheme with fourth order accuracy for the convective terms and second order central differences for the diffusive terms and the continuity equation (Poisson equation for pressure). For implicit LES, we replaced the central difference scheme for the convective terms by the implicit turbulence model ALDM. All computations were run on Cartesian staggered grids with uniform cell size.

For time integration, we used an explicit third-order accurate Runge-Kutta scheme, proposed by Shu [26]. The time step was dynamically adjusted to keep the CFL number smaller than unity.

The Poisson equation for the pressure is solved in every Runge-Kutta sub step. The Poisson solver employs fast Fourier-transform (FFT) in the vertical direction and a Stabilized Bi-Conjugate Gradient (BiCGSTAB) solver [31] in the horizontal plane. By the FFT, the three-dimensional problem is transformed into a set of independent two-dimensional problems, which can be solved in parallel.

Our code INCA is parallelized both for distributed and shared memory usage. Additionally, it is optimized for running on vector computer systems. For the computations presented here, we use the single domain shared memory approach for efficient computation of Fourier transforms of the whole data set. On this single domain, we use the openMP shared memory parallelization capabilities of INCA. The limiting factor for the single-domain approach is the size of the shared memory of the computer. We found excellent conditions on the NEC SX-9 vector computer

at the High Performance Computing Center Stuttgart (HLRS). One node of the SX-9 provides 510 GB shared memory for 16 vector CPUs. This enabled us to run DNS with up to approximately one billion cells. The computation time was about 70 ns per time step and cell. The computational performance of the pressure Poisson solver reached approximately 19 GFLOP/s. The computation of the numerical flux function reached 35 GFLOP/s and the reconstruction routine 49 GFLOP/s, which is half the nominal peak performance of the SX-9.

3.2 *The Adaptive Local Deconvolution Method*

Our approach to ILES is based on a nonlinear finite-volume scheme involving a solution adaptive reconstruction (deconvolution) of the numerical solution. The nonlinear discretization scheme generates a certain controllable spectral numerical viscosity. Using an evolutionary optimization algorithm, free parameters of the discretization scheme have been calibrated in such a way that the effective spectral numerical viscosity is identical to the spectral eddy viscosity from turbulence theory for asymptotic cases [11]. For a detailed description of ALDM, we refer to Hickel et al. [11]. The extension of ALDM to stably stratified media is provided by Remmler and Hickel [24].

3.3 *Standard Smagorinsky Model*

If we apply a generic spatial filter (denoted by an overbar) to the dimensional momentum equation, we obtain

$$\partial_t \bar{u}_i + \partial_{x_j} (\bar{u}_i \bar{u}_j) + \partial_{x_i} \bar{p} = \partial_{x_j} (2\nu \bar{S}_{ij}) - \partial_{x_j} \tau_{ij}. \quad (6)$$

where $\bar{S}_{ij} = 0.5 (\partial_{x_i} \bar{u}_j + \partial_{x_j} \bar{u}_i)$ is the filtered strain rate tensor and $\tau_{ij} = \overline{u_i u_j} - \bar{u}_i \bar{u}_j$ is the unknown SGS stress tensor, which has to be modeled. With a Boussinesq approach the SGS stress tensor is modeled as

$$\tau_{ij}^{mod} = -2\nu_t \bar{S}_{ij}. \quad (7)$$

The eddy viscosity concept is common to many SGS turbulence models. Smagorinsky [27] estimated the unknown eddy viscosity ν_t from

$$\nu_t = (C_S \bar{\Delta})^2 |\bar{S}|; \quad |\bar{S}| = \sqrt{2\bar{S}_{ij}\bar{S}_{ij}}, \quad (8)$$

where $\bar{\Delta} = (\Delta_x \Delta_y \Delta_z)^{1/3}$ is the grid or filter size, respectively. In this formulation, the unknown SGS fluxes can be computed directly from the resolved

velocity field. There is no universal value for the model constant C_S ; for different flow configurations different values of the constant have been found to be optimal. In our simulations we use a value of $C_S = 0.18$, which follows from theory of isotropic turbulence [17] and has been found to yield good results in practice [3]. The buoyancy equation is closed analogously by an eddy diffusivity model with $\alpha_t = \nu_t / \text{Pr}_t$.

3.4 Dynamic Smagorinsky Model

The case-dependence of the value for the model constant in the standard Smagorinsky model led to the idea of replacing the constant by a dynamic parameter, which automatically adjusts to the flow conditions. Germano et al. [9] presented a general dynamic procedure for eddy viscosity models and applied it to the Smagorinsky model. The basic idea is a similarity between the interactions of the smallest resolved scales and unresolved scales compared to the interactions between medium scales and the smallest resolved scales.

The solution is available in its filtered form \bar{u} with a grid filter width $\bar{\Delta}$. This filtered velocity field is explicitly filtered by a test filter with a larger filter width $\hat{\Delta}$. As a test filter, we use a top-hat filter with $\hat{\Delta} = 2\bar{\Delta}$. The subfilter-scale stress tensor is $T_{ij} = \widehat{\bar{u}_i \bar{u}_j} - \widehat{\bar{u}_i} \widehat{\bar{u}_j}$. It cannot be computed directly from the filtered velocity field, but one can compute the Leonard stress tensor $L_{ij} = \widehat{\bar{u}_i \bar{u}_j} - \widehat{\bar{u}_i} \widehat{\bar{u}_j}$. Using the Germano identity

$$T_{ij} = L_{ij} + \widehat{\tau}_{ij} \quad (9)$$

and the standard Smagorinsky model for τ_{ij} and T_{ij} , we can minimize the difference between L_{ij} and

$$\begin{aligned} L_{ij}^{mod} &= T_{ij}^{mod}(C, \hat{\Delta}, \widehat{\mathbf{u}}) - \widehat{\tau}_{ij}^{mod}(C, \bar{\Delta}, \bar{\mathbf{u}}) \\ &\doteq -2C \hat{\Delta}^2 |\widehat{S}| \widehat{S}_{ij} + 2C \left(\overline{\bar{\Delta}^2 |\bar{S}| \bar{S}_{ij}} \right) \doteq 2CM_{ij} \end{aligned} \quad (10)$$

by a least-squares procedure

$$C = \frac{1}{2} \frac{L_{ij} M_{ij}}{M_{ij} M_{ij}}. \quad (11)$$

We apply this dynamic procedure in every time step to obtain the model parameter $C_S^2 = C$ for the flow. Since the presently investigated flows are homogeneous in all spatial directions, numerator and denominator of Eq. (11) are averaged in space before evaluating Eq. (11). So finally, the parameter C_S is spatially constant, but can vary in time.




Fig. 1 Regime diagram for our simulations of stratified turbulence

4 Results and Discussion

We investigated homogeneous stratified turbulence in a statistically steady state. The flow is maintained at an approximately constant energy level by a large scale vertically uniform forcing of the horizontal velocity components. This approach models the forcing by the synoptic scale flow in the atmosphere and was successfully applied by several authors before [20, 21, 33].

We ran two series of DNS, series A with $Re_0 = 6,500$ and series B with $Re_0 = 13,000$. The domain size was 320^3 cells for series A and 640^3 cells for series B. Within the single series, the Froude number was varied to cover different buoyancy Reynolds numbers. The basic domain size again was $2\pi\mathcal{L}$. For low Froude numbers, we used a flat domain with a height of only $\pi\mathcal{L}$, but keeping cubical cells. This is permitted since in stratified turbulence there is only a very small amount of energy contained in the large scale vertical modes.

For both series, we performed LES, both with implicit ALDM and explicit SSM and DSM. For all these simulations, we used grid boxes with 64^3 cells. For the low Froude number simulations, the domain was flattened as well, leading to a doubled resolution in vertical direction. Figure 1 shows the local Froude and Reynolds number of the simulations.

Most important for the assessment of a parametrization scheme for stratified turbulence is its ability to correctly predict the amount of energy converted from horizontal kinetic energy to vertical kinetic energy and available potential energy before the energy is finally dissipated on the smallest represented scales. In Fig. 2, we show the ratio E_v/E_p as a function of local Froude number as predicted by DNS and LES with ALDM. The ratio E_v/E_p is not influenced by the forcing and can thus freely develop according to the dynamic interaction of convective, pressure and buoyancy term.




Fig. 2 Ratio of vertical to potential energy in HST as a function of local Froude number

The vertical to potential energy ratio increases almost linearly with Froude number in the DNS. We find the same trend in our LES with ALDM. The agreement between DNS and LES is best in the region of high Froude numbers (weakly stratified turbulence), whereas for low Froude numbers the vertical kinetic energy is slightly underpredicted. Note that the results from ALDM and SSM differ from each other most at the lowest Froude number. This is an indication for ALDM being better capable of handling the strong turbulence anisotropy in strongly stratified flows.

For a visual comparison of neutrally and stably stratified turbulence, we show the results of two computations from series B in Figs. 3 and 4. Both show snapshots of the developed turbulent flow at comparable Reynolds numbers. The turbulent structures are visualized by iso-surfaces of the Q -criterion [13]. For presentation purposes, the visualization includes only slices of finite thickness at the domain boundaries.

Figure 3 shows a case with neutral stratification ($Re = 8,440$). For better comparability with the stratified case, we show only the lower half of the cubical computational domain. The visualization shows that the turbulent structures have no preferred orientation. This proves that isotropic turbulence can be generated by the horizontal large-scale forcing that we used. Furthermore, we observe a spatially intermittent field of turbulence which reflects the remaining large-scale anisotropy. There are regions of higher and lower density of turbulent vortices. The regions with strong turbulence activity are associated with much higher values of molecular dissipation rate. Note that the color scale of normalized dissipation rate is logarithmic. In the (rare) red regions, the dissipation rate is more than 30 times higher than the instantaneous spatial average in the whole domain.

In case of strongly stratified turbulence ($Re = 7,110$, $Fr = 0.017$, $\mathcal{R} = 2.1$, cf. Fig. 4), the turbulence structures look completely different. Although Reynolds

Fig. 3 DNS of neutrally stratified turbulence ($Re = 8,440$). Iso-surface of Q , color of iso-surface and shading of domain faces by local molecular dissipation rate (normalized by the spatial average)





Fig. 4 DNS of strongly stratified turbulence ($Re = 7,110$, $Fr = 0.017$). Coloring as in Fig. 3



number and mean total energy dissipation are similar to the neutrally stratified case, the smallest eddies are much larger. Additionally, all eddies are aligned horizontally. Despite the larger vortices, the molecular dissipation rates are comparable to the neutral case. This is due to the intensified shear between the horizontal layers.

We present vertical and horizontal cuts of the DNS domain in Figs. 5 and 6. In the neutrally stratified case, we can still see the signature of the horizontal forcing in the velocity magnitude plots (left panels of Fig. 5). There are large-scale column-like regions of higher or lower velocity, superposed by a lot of small-scale variation. Apparently, this does not affect the behavior of the small scales. The signature of the column structures is missing in the contour plot of molecular dissipation




Fig. 5 Vertical ($x = 0$) and horizontal ($z = 2.4$) cuts through a DNS result of neutrally stratified turbulence. *Left panels:* velocity magnitude, *right panels:* molecular dissipation

(right panels). There is hardly any correlation between velocity magnitude and molecular dissipation. Small turbulence scales “forget” about the anisotropic large-scale forcing.

In case of strong stratification (Fig. 6), we find a strong horizontal layering, as described in many studies before. The flow is separated into thin layers of very different kinetic energy. The flow laminarization due to stratification reduces the amount of small-scale variations of velocity in horizontal planes compared to the neutrally stratified case. Neighboring layers can have strongly different (horizontal) velocities, which results in the formation of unstable shear layers. We observe a lot of Kelvin-Helmholtz-like structures all over the domain (cf. upper right panel of Fig. 6), which are responsible for the major part of energy dissipation. The signature of these Kelvin-Helmholtz layers can also be observed in the horizontal cut of the domain.

Many authors refer to the structures in stratified turbulence as “pancake” vortices. This popular image is used to describe large flat vortices rotating around a vertically oriented core. We observe large flat structures in the flow, but they are far from being coherent quasi-two-dimensional vortices. Instead, we find only very weak rotation around the vertical axis. The dominant turbulence structures are small-scale




Fig. 6 Vertical ($x = 0$) and horizontal ($z = 0.1$) cuts through a DNS result of strongly stratified turbulence. *Left panels:* velocity magnitude, *right panels:* molecular dissipation

Kelvin-Helmholtz vortices. These vortices are dominating the energy dissipation in the stratified turbulent flow. They appear in horizontal patches between two layers of strongly differing velocities, but their axis of rotation is basically horizontal. To stay within the culinary images, these vortices may better be described as “tagliatelle” rather than pancakes.

In Fig. 7 we present the contours of velocity magnitude for ALDM computations of the two cases presented before. The LES resolution was 64^3 , so the total number of cells was three orders of magnitude lower than in the DNS. The global flow structure resembles well the DNS result, of course without the unresolved small scale content. In the strongly stratified case, we observe the same layering with large horizontal scales as in the DNS. Even the number and thickness of horizontal layers is similar. The major difference between LES and DNS consists in the horizontal Kelvin-Helmholtz vortices. Their vertical extension, as found in the DNS, is of the same order of magnitude as the vertical resolution in the LES. Hence, they form a typical subgrid-scale problem and their dissipative effect is accounted for by the implicit turbulence model.

For comparison of kinetic energy spectra, we selected one DNS in the weakly stratified regime ($\mathcal{R} = 41$, $Fr = 0.07$, $Re = 9,300$; cf. Fig. 8) and one DNS in the




Fig. 7 Vertical ($x = 0$) and horizontal (*left panel*: $z = 1.6$, *right panel*: $z = 0.9$) cuts through ALDM results (velocity magnitude) of neutrally and strongly stratified turbulence




Fig. 8 Weakly stratified turbulence kinetic energy spectra ($\mathcal{R} = 41$). **(a)** Horizontal spectra. **(b)** Vertical spectra

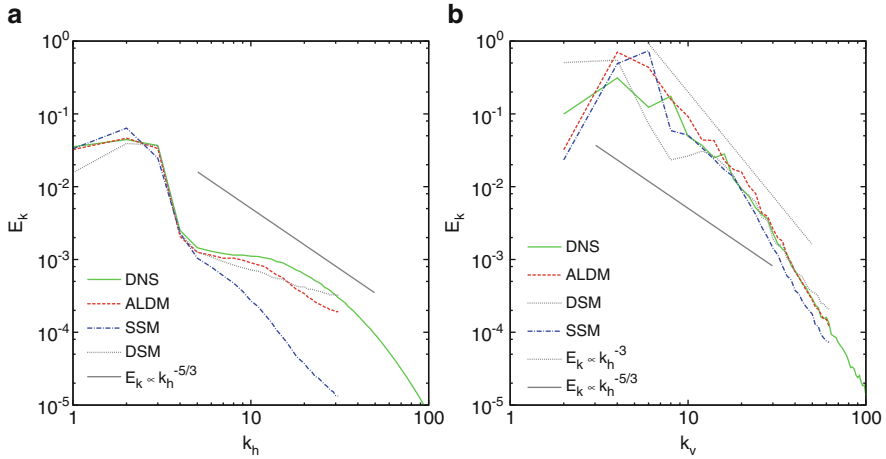


Fig. 9 Strongly stratified turbulence kinetic energy spectra ($\mathcal{R} = 6.3$). (a) Horizontal spectra. (b) Vertical spectra

strongly stratified regime ($\mathcal{R} = 6.3$, $Fr = 0.03$, $Re = 8,300$; cf. Fig. 9), both from series B. A more detailed analysis of these computations in spectral space is provided by Remmler and Hickel [24]. The corresponding LES have similar local Froude and Reynolds numbers.

In the horizontal spectra of kinetic energy, the differences between ALDM and the explicit SGS models are most obvious. In the weakly stratified case ($\mathcal{R} = 41$), the horizontal spectrum is still quite similar to the Kolmogorov spectrum of isotropic turbulence. In this case, all three SGS models predict the inertial range spectrum fairly well. The SSM and DSM are slightly too dissipative, but the difference to the DNS spectrum is acceptable. Things completely change for the stronger stratified case ($\mathcal{R} = 6.3$). The SSM dissipates too much energy and thus underpredicts the inertial range spectrum by more than one order of magnitude. Additionally, the predicted power-law exponent is significantly lower than $-5/3$. With the DSM, the overall prediction of the spectrum is much better than with the SSM, but the power-law exponent near the cut-off wavenumber is greater than $-5/3$. The spectrum predicted by ALDM agrees better with the DNS than results for both Smagorinsky models. Only ALDM correctly predicts the characteristic plateau region between the forcing scales and the inertial scales. Moreover, ALDM produces a power-law decay with an exponent of $-5/3$, corresponding to the DNS and theory derived from scaling laws [2].

We note that the dynamic model coefficient in the DSM is approximately constant in time as soon as the statistically steady state is reached. The dynamic procedure is efficient in choosing a proper coefficient depending on the strength of the stratification, but cannot cure the more structural weakness of the isotropic eddy viscosity approach.

In the vertical spectra of kinetic energy, the inertial range decay exponent changes from $-5/3$ in neutrally stratified fluid to -3 in strongly stratified turbulence. We find this change in the DNS and it is well reproduced by the LES. All three SGS models predict the turbulence inertial range decay well. At strong stratification, the ALDM result perfectly agrees with the DNS. The SSM result is slightly too dissipative in this region.

5 Conclusion

We presented a numerical investigation of homogeneous turbulence in a stably stratified fluid to proof the reliability of implicit turbulence modeling with the Adaptive Local Deconvolution Method (ALDM). As benchmark results, we used high resolution DNS data and LES results obtained with an explicit standard (SSM) Smagorinsky model and a dynamic Smagorinsky model (DSM). In most simulations, the buoyancy Reynolds number was larger than unity. The Froude and Reynolds number were chosen to cover the complete range from isotropic Kolmogorov turbulence up to strongly stratified turbulence.

We analyzed results from DNS of homogeneous turbulence with and without stable stratification. As in previous studies, we found a strong horizontal layering in the strongly stratified cases. Energy dissipation is concentrated within thin layers of Kelvin-Helmholtz vortices. Although these dominant vortices are not resolved in the LES, the LES results (with ALDM) agree well with the reference DNS, both in integral flow properties and energy spectra. This applies to the whole Froude number range from infinity down to very low values. Especially in the strongly stratified regime, ALDM performs better than the SSM. While the SSM is far too dissipative in this case, ALDM spectra agree very well with the reference DNS. With the DSM, the excessive dissipation of the SSM can be avoided, but the spectral slope near the cut-off wavenumber is still not correct. Among the investigated SGS models, only ALDM predicted the correct exponent of $-5/3$.

The results presented here were obtained without recalibrating the ALDM model constants for stratified turbulence. The good agreement with DNS data shows the ability of ALDM to automatically adapt to strongly anisotropic turbulence. Within the continuation of the project, we will investigate to which extent the results can be further improved by a recalibration of the model coefficients for stratified turbulence. But even without this possible improvements, ALDM provides a suitable parametrization of geophysical turbulence. After having applied ALDM successfully to rather simple model flows, we will continue working on more complex geophysical problems.

Acknowledgements This work was funded by the German Research Foundation (DFG) under the grant HI 1273-1 in the framework of the MetStröm priority program (SPP 1276). Computational resources were provided by the High Performance Computing Center Stuttgart (HLRS) under the grant TIGRA.

References

1. P. Bouruet-Aubertot, J. Sommeria, and C. Staquet. Stratified turbulence produced by internal wave breaking: two-dimensional numerical experiments. *Dyn. Atmos. Oceans*, 23(1–4):357–369, 1996. Stratified flows.
2. G. Brethouwer, P. Billant, E. Lindborg, and J.-M. Chomaz. Scaling analysis and simulation of strongly stratified turbulent flows. *J. Fluid Mech.*, 585:343–368, 2007.
3. R. A. Clark, J. H. Ferziger, and W. C. Reynolds. Evaluation of subgrid-scale models using an accurately simulated turbulent flow. *J. Fluid Mech.*, 91(01):1–16, 1979.
4. C. Cot. Equatorial mesoscale wind and temperature fluctuations in the lower atmosphere. *J. Geophys. Res.*, 106(D2):1523–1532, 2001.
5. E. M. Dewan. Stratospheric wave spectra resembling turbulence. *Science*, 204(4395):832–835, 1979.
6. A. Dörnbrack. Turbulent mixing by breaking gravity waves. *J. Fluid Mech.*, 375:113–141, 1998.
7. D. C. Fritts, L. Wang, J. Werne, T. Lund, and K. Wan. Gravity wave instability dynamics at high reynolds numbers. Part I: Wave field evolution at large amplitudes and high frequencies. *J. Atmos. Sci.*, 66(5):1126–1148, 2009.
8. K. S. Gage. Evidence for a $k^{-5/3}$ law inertial range in mesoscale two-dimensional turbulence. *J. Atmos. Sci.*, 36:1950–1954, 1979.
9. M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. A dynamic subgrid-scale eddy viscosity model. *Phys. Fluids A*, 3(7):1760–1765, 1991.
10. J. R. Herring and O. Métais. Numerical experiments in forced stably stratified turbulence. *J. Fluid Mech.*, 202(1):97–115, 1989.
11. S. Hickel, N. A. Adams, and J. A. Domaradzki. An adaptive local deconvolution method for implicit LES. *J. Comput. Phys.*, 213:413–436, 2006.
12. S. Hickel, N. A. Adams, and N. N. Mansour. Implicit subgrid-scale modeling for large-eddy simulation of passive scalar mixing. *Phys. Fluids*, 19:095102, 2007.
13. J. C. R. Hunt, A. Way, and P. Moin. Eddies, stream, and convergence zones in turbulent flows. Center for Turbulence Research Report CTR-S88, Center for turbulence research, Stanford University, 1988.
14. H.-J. Kaltenbach, T. Gerz, and U. Schumann. Large-eddy simulation of homogeneous turbulence and diffusion in stably stratified shear flow. *J. Fluid Mech.*, 280(1):1–40, 1994.
15. R. H. Kraichnan. Inertial ranges in two-dimensional turbulence. *Phys. Fluids*, 10(7):1417–1423, 1967.
16. J.-P. Laval, J. C. McWilliams, and B. Dubrulle. Forced stratified turbulence: Successive transitions with Reynolds number. *Phys. Rev. E*, 68(3):036308, 2003.
17. D. K. Lilly. The representation of small scale turbulence in numerical simulation experiments. In *IBM Scientific Computing Symposium on environmental sciences*, pages 195–210, 1967.
18. D. K. Lilly. Stratified turbulence and the mesoscale variability of the atmosphere. *J. Atmos. Sci.*, 40(3):749–761, 1983.
19. D. K. Lilly, G. Bassett, K. Droegemeier, and P. Bartello. Stratified turbulence in the atmospheric mesoscales. *Theor. Comput. Fluid Dyn.*, 11:139–153, 1998.
20. E. Lindborg. The energy cascade in a strongly stratified fluid. *J. Fluid Mech.*, 550(1):207–242, 2006.
21. O. Métais and J. R. Herring. Numerical simulations of freely evolving turbulence in stably stratified fluids. *J. Fluid Mech.*, 202(1):117–148, 1989.
22. O. Métais and M. Lesieur. Spectral large-eddy simulation of isotropic and stably stratified turbulence. *J. Fluid Mech.*, 239:157–194, 1992.
23. G. D. Nastrom and K. S. Gage. A climatology of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft. *J. Atmos. Sci.*, 42(9):950–960, 1985.
24. S. Remmler and S. Hickel. Spectral structure of stratified turbulence: Direct numerical simulations and predictions by large eddy simulation. *Theor. Comput. Fluid Dyn.*, in press, doi:10.1007/s00162-012-0259-9, 2012.

25. J. J. Riley and S. M. de Bruyn Kops. Dynamics of turbulence strongly influenced by buoyancy. *Phys. Fluids*, 15(7):2047–2059, 2003.
26. C.-W. Shu. Total-variation-diminishing time discretizations. *SIAM J. Sci. Stat. Comput.*, 9(6): 1073–1084, 1988.
27. J. Smagorinsky. General circulation experiments with the primitive equations. I: The basic experiment. *Mon. Wea. Rev.*, 91:99–164, 1963.
28. L. M. Smith and F. Waleffe. Generation of slow large scales in forced rotating stratified turbulence. *J. Fluid Mech.*, 451(1):145–168, 2002.
29. C. Staquet and F. S. Godeferd. Statistical modelling and direct numerical simulations of decaying stably stratified turbulence. Part 1. Flow energetics. *J. Fluid Mech.*, 360:295–340, 1998.
30. J. R. Taylor, S. Sarkar, and V. Armenio. Large eddy simulation of stably stratified open channel flow. *Phys. Fluids*, 17(11):116602, 2005.
31. H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 13(2):631–644, 1992.
32. T. E. van Zandt. A universal spectrum of buoyancy waves in the atmosphere. *Geophys. Res. Lett.*, 9(5):575–578, 1982.
33. M. L. Waite and P. Bartello. Stratified turbulence dominated by vortical motion. *J. Fluid Mech.*, 517:281–308, 2004.

Part VI

Miscellaneous Topics

Univ.-Prof. Dr.-Ing. Wolfgang Schröder, Prof.Dr.-Ing. Peter Wriggers,
and Univ.-Prof. Dr. Hans-Joachim Bungartz

The following contributions will show that numerical simulations represent useful tools to gain novel results and as such to improve the scientific knowledge not only in physics and/or solid and fluid mechanics. The papers evidence the close link between applied and fundamental research. In other words, the close relationship between mathematics, computer science and the ability to develop scientific models is clearly shown. Compact mathematical descriptions will be solved by highly sophisticated and efficient algorithms on high-end computers. This interdisciplinary collaboration between several scientific fields defines the extremely intricate numerical challenges and as such drives the progress in fundamental and applied research. The subsequent articles represent an excerpt of the vast amount of projects being linked with HLRS. The computations are not only used to obtain some quantitative results but to confirm fundamental physical models and to even derive new theoretical approaches. Nevertheless, it has to be emphasized that it will always make sense to substantiate numerical simulations by experimental investigations and analytical solutions.

In the first contribution of the Goethe University Frankfurt a software framework is discussed. The modeling of physical phenomena in a variety of fields of scientific interest leads to a formulation in terms of partial differential equations. Especially when complex geometries as the domain of definition are involved, a direct and exact solution is not accessible such that numerical schemes are to be

W. Schröder (✉)

Institute of Aerodynamics, RWTH Aachen University, Wüllnerstr. 5a, 52062, Aachen, Germany
e-mail: office@aia.rwth-aachen.de

P. Wriggers

Institut für Kontinuumsmechanik, Leibniz Universität Hannover, Appelstr. 11, 30167,
Hannover, Germany

H.-J. Bungartz

Fakultät für Informatik, Technische Universität München, Boltzmannstr. 3, 85748,
Garching, Germany

used to compute an approximate discrete solution. The focus is on elliptic and parabolic equations that include spatial operators of second order. Then discretising such problems using commonly known discretization schemes like finite element methods or finite volume methods, large systems of linear equations arise naturally. Their solution takes the largest amount of the overall computing time. A powerful algorithm for the solution of systems of linear equations is the multi-grid method. The discussion deals with the scalability of several multi-grid approaches to show that the presented methods are well-suited to be used on parallel supercomputers. All methods are included in the simulation framework called ug4. This framework is designed for the solution of differential equations on unstructured hybrid grids in one, two, and three space dimensions on systems, ranging from laptops to massively parallel computers.

In the contribution of the University of Paderborn the capability of molecular simulation to determine fluid phase equilibria is examined, especially with respect to the problems of the standard approach in which the data for the design of such thermodynamic processes are aggregated by empirical correlations, which are based on experimental data. Molecular modeling and simulation is a modern route for the prediction of thermodynamic properties. Being based on mathematical representations of the intermolecular interactions, it has strong predictive capabilities as it adequately covers structure, energetic and dynamics on the microscopic scale that govern the fluid behavior on the macroscopic scale. The capability of previously discussed models to predict the liquid-liquid equilibrium (LLE) data is studied for the binary mixture Nitrogen and Ethane. The decomposition of a randomly distributed mixture of these two components in their LLE phases is investigated. Subsequently, the composition of the phases in equilibrium is compared to experimental data. It is shown that vapor-liquid equilibrium (VLE) and LLE can be predicted by molecular simulation with good agreement to experiments.

The next article of the German Aerospace Center, Berlin and the University of Münster discusses a particle-in-cell-method to simulate the impact of partial melt on mantle convection. Solid-state convection is the principal mechanism that controls the global dynamics and thermal evolution of the terrestrial planets. Observations such as seismology and mission data from geological structures at the planetary surfaces show important constraints for the interior dynamics. However, the main knowledge stems from laboratory experiments and in particular from computer models. In the last years, due to significant improvements in high performance computing, computer simulations have offered a powerful access to this fluid-dynamical problem by approximately solving partial differential equations to describe the flow in space and time. Results will be presented from the cylindrical/spherical code GAIA which is based on a particle-in-cell method to account for compositional changes due to partial melting of the mantle.

The fourth contribution of the Friedrich-Schiller University, Jena also deals with mantle convection. Some essential features of Andean orogenesis cannot be explained only by a dynamic regional model since there are essential influences across its vertical boundaries. A dynamic regional model of the Andes should be embedded in a 3-D spherical-shell model. Because of the energy distribution on

the poloidal and toroidal parts of the creep velocity and because of geologically determined mass transport alongside the Andes, both models have to be three-dimensional. A new viscosity profile of the mantle with very steep gradients at the lithospheric-asthenospheric boundary was developed. Based on the new viscosity profile a new forward spherical-shell model was introduced. For this model a new extended acoustic Grüneisen parameter, γ_{ax} , new profiles of the thermal expansivity, α , and of the specific heat, c_v , at constant volume as well as a solidus depending on both the pressure and the water abundance were derived. A regional model of the Andean orogenesis with the same new viscosity profile was introduced. In connection with another spherical-shell model the regional model is supposed to numerically explain why a plateau-type orogen evolved at an oceanic-continent plate boundary.

In the project of the University of Stuttgart the gravity gradiometry data with near global coverage from a single source, namely the satellite mission GOCE (Gravity field and steady-state Ocean Circulation Explorer) launched on 17 March 2009, are investigated. The question to be addressed is whether geodesy can benefit from Euler deconvolution, e.g., to retrieve global gravity models, or is able to contribute its methodologies to enhance Euler deconvolution. Until now the project is still in preparatory stage, mainly, because the GOCE gradiometry data need to be preprocessed.

In the article from the University of Hohenheim in Stuttgart economic capital allocation in banking is considered. The model was described as a mixed integer nonlinear program (MINLP) and an appropriate solving algorithm in form of threshold accepting was introduced. In the current contribution the parameterization of the algorithm is addressed. The implementation of threshold accepting requires certain model modifications which also affect the parameterization procedure. Using an adequate parameterization the model provides an indication of optimum economic capital allocation's superiority compared to alternative allocation methods.

Structural mechanics and material mechanics are areas that include demanding problems as well from the theoretical side as from the computational point of view. On one hand, based on the computing power, physical models can be refined and enhanced such that a virtual testing of complex behaviour is possible on different length and time scales. On the other hand, algorithms become more and more complex and need to be stabilized in order to obtain robust and reliable results. Challenges lie in the development of new computational methodologies and the reproduction of correct physical behaviour.

The two contributions that relate to structural mechanics are concerned with the virtual testing of heterogeneous materials including damage effects using multiscale approaches. The work by Schrader and Könke provides insight in the behavior of different parallel computer architectures when applied to large scale problems. The implementation of solution algorithms was investigated by comparing algorithms for a linear solve within a nonlinear solution algorithm on the hybrid CPU-GPU NEC Nehalem cluster and the CRAY XE6 system. While this contribution was concerned with linear solvers, the second contribution of Eck et al. devoted its attention to numerical sensitivities in complex impact problems, e.g. crash

simulations. The question was how can these nonlinear simulations be done in a robust manner, meaning that the results have to be reliable and reproducible for the same mesh in different runs when large number of unknowns are present. Different solutions are caused either by physical or by numerical instabilities. The aim of the project was to reduce scatter of the responses by identification of the sources of scatter through parameter studies and to improve the robustness by an enhancement of existing mathematical formulations and algorithms. This is especially needed when these simulations are executed on parallel computing systems with large numbers of unknowns.

There have been two projects with an application background from informatics: “Characterization of Carrier Sense Multiple Access in Vehicular Propagation Channels (CAR2X)” and “HMDB51 – A Large Video Database for Human Motion Recognition (HMDB)”. Both projects are computationally intense, but not primarily focused on parallel computing or HPC issues.

CAR2X is a successor of previous related projects. Wireless communication between vehicles is considered to be important to increase the safety level of future smart transportation systems. While it sounds intuitively clear that a periodic exchange of status information may help to avoid dangerous traffic situations, it is not clear whether the envisioned communications system is sufficiently reliable and robust, whether the employed Carrier Sense Multiple Access (CSMA) mechanism is able to coordinate concurrent access by multiple network nodes in a highly dynamic environment as intended. Therefore, the performance of a CSMA-based coordination mechanism is evaluated, based on a network simulation framework that emulates the signal processing steps of a transceiver and accurately models the multi-path propagation effects of the wireless vehicular radio channel. Due to this accuracy requirements, the execution of such high-fidelity simulations is computationally highly expensive.

The project HMDB addresses recognition and search in videos, a highly relevant topic against the background of millions of online videos viewed every day. While much attention has been paid to the collection and annotation of large static image datasets containing thousands of image categories, human action datasets lag far behind: Current action recognition databases contain on the order of ten different action categories. To address the important issues of collecting and benchmarking, the largest action video database to-date with 51 action categories, which in total contain around 7,000 manually annotated clips extracted from a variety of sources ranging from digitized movies to YouTube, was collected. The goal is to provide a tool to evaluate the performance of computer vision systems for action recognition and to explore the robustness of these methods under various conditions such as camera motion, viewpoint, video quality, and occlusion.

Software Framework ug4: Parallel Multigrid on the Hermit Supercomputer

Ingo Heppner, Michael Lampe, Arne Nägel, Sebastian Reiter, Martin Rupp, Andreas Vogel, and Gabriel Wittum

1 Introduction

The modeling of physical phenomena in a variety of fields of scientific interest lead to a formulation in terms of partial differential equations. Especially when complex geometries as the domain of definition are involved, a direct and exact solution is not accessible, but numerical schemes are used to compute an approximate discrete solution. In this report, we focus on elliptic and parabolic types of equations that include spatial operators of second order. When discretizing such problems using commonly known discretization schemes such as finite element methods or finite volume methods, large systems of linear equations arise naturally. Their solution takes the largest amount of the overall computing time.

A powerful algorithm for the solution of systems of linear equations is the multigrid method. A general introduction to this topic can be found, e.g., in [3]. It is known to have optimal complexity $O(N)$, where N is the number of unknowns of the discrete system. In this report, we focussed on the scalability of several multigrid approaches to show that the presented methods are well-suited for the usage on parallel supercomputers. All methods are included in the ug4 simulation framework [9]. ug4 is designed for the solution of differential equations on unstructured hybrid grids in 1, 2 and 3 space dimensions on systems, ranging from laptops to massively parallel computers. It is written in C++, striving for a flexible, yet fast and robust simulation environment.

I. Heppner (✉) · M. Lampe · A. Nägel · S. Reiter · M. Rupp · A. Vogel · G. Wittum
G-CSC, Goethe-Universität Frankfurt, Kettenhofweg 139, 60325, Frankfurt (M.), Germany
e-mail: ingo.heppner@gcsc.uni-frankfurt.de; michael.lampe@gcsc.uni-frankfurt.de;
arne.naegel@gcsc.uni-frankfurt.de; sebastian.reiter@gcsc.uni-frankfurt.de;
martin.rupp@gcsc.uni-frankfurt.de; andreas.vogel@gcsc.uni-frankfurt.de;
gabriel.wittum@gcsc.uni-frankfurt.de




Fig. 1 Schematic overview of the multigrid setup

All computations were performed on the Cray XE6 “Hermit” at HLRS Stuttgart. The XE6 is a 3,552 node cluster. Each node is a Dual Socket AMD Opteron 6276 (Interlagos) @ 2.3 GHz 16 cores each, which results in 113.664 cores in total. The nodes we were using have 32 GB RAM, resulting in 1 GB RAM per core when using the maximum of 32 cores/node.

The remainder is organized as follows: Sect. 2 briefly summarizes ingredients of the multigrid methods. Section 3 is devoted to simple scalability benchmark tests. Section 4 presents an application of the methods for solving problems from density driven flow.

2 Methods

2.1 Geometric Multigrid

In this section, we present the idea of our implementation of a parallel geometric multigrid method. A detailed description is given in [7]. The parallelization has been implemented using MPI.

The main aspect of the parallelization is the distribution of the multigrid hierarchy to the processes using a hierarchical approach. See Fig. 1 for an illustration of a possible assignment of a 1d problem to two processes. For the finest grid level of the multigrid, this can be thought of as a partition of the domain with respect to the processes. At the process borders within each grid level we build up horizontal interfaces that allow the communication between nodes that are representing the same global object but are a local copy for each process. Going down in the multigrid hierarchy, at a certain point the grid levels become very sparse and therefore we start to gather parts of the grids on some processes while the other processes have no grid part on the coarser levels. When traversing the multigrid in vertical direction during a multigrid cycle, this will involve communication of data between the gathering processes and those that have no coarser grid level. To this aim, we introduce vertical interfaces that allow the communication in this direction.

The multigrid cycle is performed as usual: starting on the finest grid, where all processes are involved, some smoothing is performed on each grid level using the horizontal interfaces. Going down in the multigrid hierarchy the restriction is performed process-locally until a vertical cut is encountered. At this point, the vertical interfaces are used to transfer the data to the gathering process and the sending processes will idle. Those processes that still have a multigrid part will continue with the multigrid cycle. Once at the bottom, this procedure is performed the way up, again using vertical interfaces, to bring all processes back to work at the end.

The success of this algorithm relies on the fact that on coarser grids the amount of computational work is very small compared to the smoothing effort on the finer grids. Therefore, the part of the algorithm where idle processes are present is negligible compared to the smoothing parts on finer grid levels where all processes are involved. When dealing with very large number of processes, i.e. thousands and more, we introduce several vertical cuts at certain points of the multigrid cycle. This leads to a tree-structure of gathering during the multigrid cycle and is important to keep the costs for the one-to-many communication using vertical interfaces within a reasonable size.

2.2 Filtering Algebraic Multigrid

Geometric multigrid methods, as introduced in the previous subsection, employ a hierarchy which is built up from coarse to fine. Thus, the distribution of the grids can be controlled by a grid manager from the start. In contrast to this, Algebraic Multigrid methods, see e.g., [8], are building up the multigrid hierarchy from fine to coarse grids. Ideally, methods of this type only rely on matrices and graph structures provided by the user. This work focuses on the Filtering Algebraic Multigrid method (FAMG, [6, 11]). The key ideas can be summarized as follows:

An essential ingredient of any multigrid method is an efficient interplay between smoother and coarse grid correction, i.e., those components of the error which are not reduced well by the smoother S on the fine grid must be treated by a coarse grid correction. We call those components *algebraically smooth* vectors. For this we need a restriction operator R (mapping defects from finer to coarser grids), an interpolation operator P (mapping corrections from coarser to finer grids), and a coarse grid operator A_H .

In FAMG, we achieve this in the following way: Let t be a vector, which is known to be characteristic for the subspace in which the smoother does not converge well. Now the algorithm constructs the operators P and R so that

- The vector t is contained in $\text{Range}(P)$
- The two grid correction $T = (I - PA_H^{-1}RA_h)S$ is efficient, i.e., $\min_P \|T\|$

For the sake of illustration, we restrict ourselves to the symmetric case, i.e., $A = A^T$, which leads to the letting $R^T = P$. Note that there are extensions to FAMG for non-symmetric matrices, systems of equations and multiple testvectors.

Global and Local Minimization Problems

As shown in [11], it suffices to use the injection operator $R = R^{\text{inj}}$, so that we obtain the following minimization problem:

$$\begin{aligned} \min_P \|(I - PR^{(\text{inj})})S\|_{D^{-1}}, \\ \text{s.t. } (I - PR^{(\text{inj})})St = 0 \end{aligned} \quad (1)$$

Here $D = \text{diagonal of } A_h$ and $\|\cdot\|_{D^{-1}}$ is the induced operator norm of the vector norm $\|x\|_{D^{-1}} = \sqrt{\langle D^{-1}x, x \rangle}$.

In order to express this locally, we assume that a fine node i is interpolated by coarse nodes which are neighbors of i in the adjacency graph of the matrix A_h . Given a partitioning of all nodes in coarse (C) and fine (F) and $q_i = i$ -th row of $(I - PR^{\text{inj}})$, we can rewrite the problem 1 as a sum over *local* (i.e. independent) minimization problems:

$$\begin{aligned} \sum_{i \in F} \min_{q_i} |a_{ii}| \|S'^T q_i\|_{D^{-1}}^2 \\ \text{s.t. } \langle q_i, S't \rangle = 0 \quad \forall i \in F \end{aligned} \quad (2)$$

This leads to Algorithm 1.

Algorithm 1 GetPossibleParentNodes

- 1: **for** all nodes i **do**
- 2: get t_i : (local) representation of the testvector
- 3: Calculate for all neighbor pairs $n \neq m \in N_i$

$$\begin{aligned} F_{i,nm} = \min_{q_{i,nm}} |a_{ii}| \|S'^T q_{i,nm}\|_{D^{-1}}^2 \\ \text{s.t. } \langle q_{i,nm}, S't \rangle = 0 \end{aligned} \quad (3)$$

- 4: **if** $\theta F_{i,nm} \leq \min_{a \neq b \in N_i} F_{i,ab}$ and $F_{i,nm} \leq \delta$, $\delta < 1$ **then**
 - 5: save the pair (n, m) in PN^i and their quality value $F_{i,nm}$.
 - 6: **end if**
 - 7: **end for**
-

After that, a set of possible parent nodes PN^i is assigned to every node i .

Used Smoother S'

In the construction of the interpolation operator, we only need to focus on interpolating smooth vectors. For this, we use a special smoothing operator S' . S'

consists of one Jacobi-Step $S_{\text{jac}} = \mathbb{I} - \omega D^{-1} A$, followed by a Jacobi-Step updating only fine nodes (F-smoothing):

$$S_{\text{jac},F} = \mathbb{I} - \sum_{i \in F} e_i e_i^T D^{-1} A \quad (4)$$

so we obtain

$$S' = S_{\text{jac},F} S_{\text{jac}} \quad (5)$$

Coarsening Algorithm

The FAMG coarsening algorithm chooses nodes which are interpolated well, and sets the parent nodes coarse. Since we want to reduce the number of coarse nodes, we prefer those nodes which create the least amount of new coarse nodes. Additionally, we want to take the best interpolating pair. The rating for each node is calculated by

$$R_i = \min_{M \in PN^i \setminus F} |M| - |M \cap C| + F_{i,nm}/\delta \quad (6)$$

This leads to Algorithm 2. Note that $F_{i,nm}/\delta < 1$, so that reducing the number of coarse nodes has higher priority. Since a node can only be either coarse or fine, it is important that we have calculated more than one possible parent pair for each node in Algorithm 1, so that we have greater flexibility in the coarsening process.

Algorithm 2 FAMGCoarsening

- 1: calculate ratings R_i for all nodes.
 - 2: **while** interpolateable nodes left **do**
 - 3: Get interpolateable node i which is neither coarse nor fine with lowest rating
 - 4: Get best available parent nodes n, m for i
 - 5: Set i fine
 - 6: Set parent nodes n, m coarse.
 - 7: **for** all neighbors $j \in N_i$ **do**
 - 8: Since i cannot be coarse anymore:
 - 9: Remove all parent pairs in PN_j which contain i .
 - 10: Update ratings R_j
 - 11: **end for**
 - 12: **end while**
-

Parallelization

For (3) we need to know the matrix A in the neighbors of i (N_i^1) and the neighbors of the neighbors of i (N_i^2), i.e., in parallel, we require an *overlap* of 2 of the matrix A . To ensure a consistent coarsening we also need a parallelization of the coarsening process. In FAMG, we use a graph coloring algorithm for this purpose so that no two cores which could set the same node coarse or fine are coarsening at the same time (Algorithm 3).

Algorithm 3 FAMGParallelCoarsening

- 1: Calculate a graph coloring so that no two cores which could set the same node coarse or fine have the same color
 - 2: Receive coarsening data from cores with lower color
 - 3: call FAMGCoarsening (Algorithm 2)
 - 4: Send coarsening data to cores with higher color
-

Putting all this together, we obtain Algorithm 4:

Algorithm 4 ParallelFAMG

- 1: Calculate Overlap 2 of A .
 - 2: Call GetPossibleParentNodes (Algorithm 1)
 - 3: Call FAMGParallelCoarsen (Algorithm 3)
 - 4: Send and receive prolongation on border nodes
 - 5: Calculate $R = P^T$, $A_H = RA_hP$
-

Agglomeration

Parallel algorithms work best if the time spent for computation and the time spent for communication is well balanced. Since the coarsening Algorithms 2 and 3 decrease the problem size, we eventually come to the point where the problem size on one core is so small that communication becomes the bottleneck. Let N_{\min} be the minimum number of unknowns on one core. If a core has less unknowns than N_{\min} on the current level, all participating processors send the following information to one processor:

- The number of unknowns
- The cores which they are connected to via interfaces, and the size of the interfaces

Now an agglomeration algorithm is run, which calculates a heuristic graph partitioning so that

- All participating cores have more than N_{desired} unknowns
- The number of participating cores is maximized
- The maximum size of the interfaces between two cores is minimized

After that, some cores will be agglomerated and become idle. Note that the general idea is similar to the algorithm used in geometric multigrid: While geometric multigrid uses vertical cuts, FAMG employs a vertical agglomeration strategy.

2.3 FETI-DP

The FETI-DP method (“Dual-Primal Finite Element Tearing and Interconnecting”) is a domain decomposition method that belongs to the class of non-overlapping domain decomposition methods [1]. In this method, global continuity of the solution




Fig. 2 Speedup on Hermit for 2d Laplace problem, solved with GMG

at subdomain interfaces is initially only enforced for the subdomain corners (a.k.a. “subdomain vertices”, “cross points”), the so-called *primal variables*. For the remaining unknowns on the subdomain edges, the so-called *dual variables*, continuity is enforced by Lagrange multipliers, which enter as additional unknowns in a CG iteration. In each step of the CG-iteration systems of equations must be inverted on the subdomains, which can be done independently in parallel. Additionally, at the cross points a Schur complement equation has to be assembled and solved for the primal variables which serves as a “coarse problem” that ensures a global exchange of information between all subdomains necessary for a good parallel scalability.

While in a standard FETI-DP implementation usually only one PE per subdomain is utilised to solve the sub problems here defined our implementation allows also a multiple of processes per subdomain. We use MPI communicators to describe the group of PE’s working in parallel on a FETI subdomain.

3 Numerical Results

3.1 Geometric Multigrid

In order to test the implementation we show the result for the Laplace equation on the unit cube

$$-\Delta u = f \quad \text{on } [0, 1]^d, \tag{7}$$

where $d = 2, 3$ denotes the physical world dimension. The equation is discretized using the vertex-centered finite volume scheme on a regular grid consisting of squares and cubes respectively. Although this model problem is simple, it provides a good example for a variety of problems that arise in more involved real-world situations.

Table 1 Laplace 2d on Hermit, solved with GMG, 7–14 refinements

PE	level	DoF	T_r (s)	T_a (s)	T_s (s)	T_{a+s} (s)	E_{a+s} (%)	S_{a+s}	S_{ideal}
4	6	263,169	3.345	0.469	0.684	1.153	100.0	1.0	1
16	7	1,050,625	4.249	0.506	1.049	1.555	74.2	3.0	4
64	8	4,198,401	4.376	0.529	1.075	1.604	71.9	11.5	16
256	9	16,785,409	4.592	0.529	1.100	1.629	70.8	45.3	64
1,024	10	67,125,249	4.991	0.518	1.206	1.724	66.9	171.3	256
4,096	11	268,468,225	5.108	0.521	1.140	1.662	69.4	710.8	1,024
16,384	12	1,073,807,361	7.985	0.530	1.248	1.779	64.9	2656.6	4,096
65,536	13	4,295,098,369	18.856	0.528	1.341	1.869	61.7	10112.0	16,384

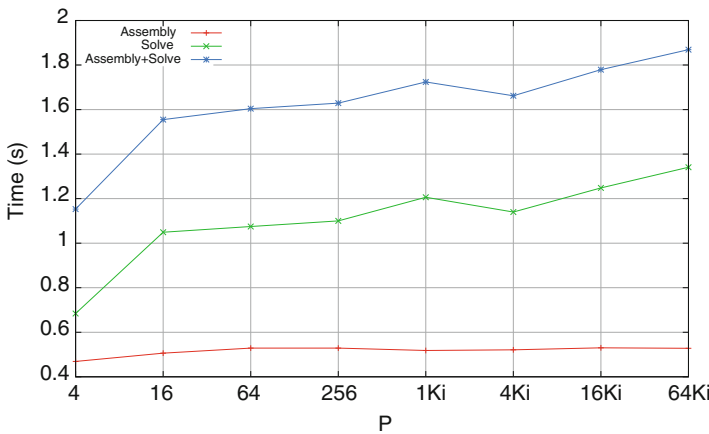


Fig. 3 Times on Hermit for 2d Laplace problem, solved with GMG

Table 1 shows the result of the weak scalability test for the 2d problem up to 65'536 processes (“processing entities” or “PE” in short) on Hermit. The assembling of the matrix is trivially parallelizable by adding contributions of each element to the process-local matrix and the measured scaling is nice as expected. As displayed in Fig. 3, the times for the multigrid solver roughly stays constant over a long range of processes.

Table 2 shows corresponding result for the 3d problem up to 4,096 processes on Hermit. Results are similar to the 2d case and shown in Figs. 4 and 5.

3.2 Algebraic Multigrid

Table 3 shows the results of FAMG solving a 2d-Laplace equation (7) on Hermit. The equation was discretized with a finite element method on triangles resulting in a 5-point-operator and solved with FAMG as preconditioner for the CG method.

Table 2 Laplace 3d on Hermit, solved with GMG

PE	level	DoF	T_t (s)	T_a (s)	T_s (s)	T_{a+s} (s)	E_{a+s} (%)	S_{a+s}	S_{ideal}
1	4	35,937	3.082	0.585	0.438	1.023	100.0	1.0	1
8	5	274,625	4.619	0.766	1.136	1.903	53.8	4.3	8
64	6	2,146,689	6.087	0.811	1.204	2.015	50.8	32.5	64
512	7	16,974,593	5.703	0.820	1.268	2.088	49.0	250.7	512
4,096	8	135,005,697	6.279	0.807	1.285	2.091	48.9	2003.2	4,096
32,768	9	1,076,890,625	16.407	0.813	1.490	2.303	44.4	14554.1	32,768

PE number of processing entities (cores), level refinement level, T_t total time spent, T_a time spent for assembly, T_s time spent for solving, T_{a+s} , E_{a+s} , S_{a+s} time spent, efficiency and speedup for assembly and solve, S_{ideal} ideal speedup




Fig. 4 Speedup on Hermit for 3d Laplace problem, solved with GMG




Fig. 5 Times on Hermit for 3d Laplace problem, solved with GMG

Table 3 Scaling of Laplace problem in 2d on Hermit solved with FAMG

PE	level	DoF	T_{total} (s)	T_{setup} (s)	AMG levels	T_{solve} (s)	N_{iter}	$\frac{T_{\text{solve}}}{N_{\text{iter}}}$
4	8	263,169	6.87	3.47	9	1.20	9	0.13
16	9	1,050,625	9.92	4.17	11	2.26	11	0.21
64	10	4,198,401	12.11	4.92	13	2.40	11	0.22
256	11	16,785,409	14.02	5.92	15	2.78	12	0.23
1,024	12	67,125,249	16.60	6.95	18	3.37	14	0.24
4,096	13	268,468,225	20.39	8.53	20	3.90	15	0.26

T_{total} total time for program, T_{setup} time for FAMG setup, T_{solve} time for CG solver using FAMG multigrid as preconditioner, $N_{\text{iterations}}$ number of CG iterations

Table 4 Number of participating cores for parallel FAMG

AMG level	PEs: 4	PEs: 16	PEs: 64	PEs: 256	PEs: 1,024	PEs: 4,096
	Used cores	Used cores	Used cores	Used cores	Used cores	Used cores
0–7	4	16	64	256	1,024	4,096
8	2	8	30	120	479	1,883
9	1	4	15	62	249	1,044
10		2	8	31	124	533
11		1	4	16	72	291
12			2	8	44	188
13			1	4	24	110
14				3	13	61
15				1	7	37
16					3	17
17					2	11
18					1	5
19						2
20						1

Table 4 lists the number of participating cores on each level for $N_{\text{min}} = N_{\text{desired}} = 1,000$ generated by the agglomeration strategy from Sect. 2.2.

First results of FAMG on Hermit are encouraging: Over a large range of core numbers, the total running time grows, but still stays within an acceptable range. For large number of processes, both the setup time and solution time are increasing. Reasons may be that the coarsening rate of the standard FAMG method is only 50 %, resulting in a larger number of levels than geometric multigrid and that the number of iterations does not remain constant. The time per iteration, however, is bounded, and even remains constant if the increase in operator complexity is considered as well.

For future work, we hope to address the previously mentioned shortcomings by improving coarsening strategies and graph coloring algorithms. We expect this to improve coarsening rates and efficiency as well.

Table 5 Laplace 2d on Hermit, solved with FETI-DP, and FAMG as subproblem solvers, 1 PE per FETI subdomain – weak scalability relative to 4 PE

PE	level	DoF	T_t (s)	T_a (s)	T_s (s)	T_{a+s} (s)	E_{a+s} (%)	S_{a+s}	S_{ideal}
4	5	66,049	7.484	0.121	1.883	2.004	100.0	1.0	1
16	6	263,169	13.717	0.135	6.885	7.019	28.5	1.1	4
64	7	1,050,625	18.008	0.139	10.419	10.558	19.0	3.0	16
256	8	4,198,401	18.812	0.138	10.924	11.062	18.1	11.6	64
1,024	9	1,6785,409	19.765	0.140	10.920	11.060	18.1	46.4	256

Table 6 Laplace 2d on Hermit, solved with FETI-DP, and FAMG as subproblem solvers, 4 PE per FETI subdomain – weak scalability relative to 16 PE

PE	level	DoF	T_t (s)	T_a (s)	T_s (s)	T_{a+s} (s)	E_{a+s} (%)	S_{a+s}	S_{ideal}
16	6	263,169	10.652	0.133	3.554	3.688	100.0	1.0	1
64	7	1,050,625	18.346	0.137	10.004	10.141	36.4	1.5	4
256	8	4,198,401	23.019	0.140	14.313	14.453	25.5	4.1	16
1,024	9	16,785,409	24.598	0.142	15.230	15.372	24.0	15.4	64
4,096	10	67,125,249	26.904	0.142	15.980	16.123	22.9	58.6	256

Table 7 Laplace 2d on Hermit, solved with FETI-DP, and FAMG as subproblem solvers, 4 PE per FETI subdomain – weak scalability relative to 64 PE

PE	level	DoF	T_t (s)	T_a (s)	T_s (s)	T_{a+s} (s)	E_{a+s} (%)	S_{a+s}	S_{ideal}
64	7	1,050,625	13.651	0.141	5.333	5.474	100.0	1.0	1
256	8	4,198,401	24.151	0.140	14.379	14.519	37.7	1.5	4
1,024	9	16,785,409	30.729	0.143	20.433	20.576	26.6	4.3	16
4,096	10	67,125,249	33.661	0.145	21.981	22.126	24.7	15.8	64
16,384	11	268,468,225	42.200	0.144	22.268	22.412	24.4	62.5	256

3.3 FETI-DP

To investigate and to demonstrate the potential of the FETI-DP implementation in ug4, we solved the model problem (7) from Sect. 2.1 and FAMG (Sect. 2.2) as subproblem solvers. This problem (with only one process per FETI subdomain, i.e., “standard” FETI-DP) has previously been used as a benchmark, e.g., in [4].

Tables 5–7 show the results of the weak scalability tests performed on Hermit for the 2d problem up to 16 Ki (16*384) PE, for three series with one, four and 16 PE per FETI subdomain. Please note that since the number of available MPI communicators is limited, every series is only up to 1 Ki FETI subdomains.

As for the other methods described in this paper the assembly times scale very well. In contrast the solve times show a distinct jump at the beginning of each series. But when the speedups are computed relative to a job of a series with a larger number of PE’s, i.e. 64 PE, our FETI-FAMG method scales quite nicely as can be seen from Fig. 6 where the curves become parallels to the line denoting the ideal speedup. This behavior has to be investigated yet.




Fig. 6 Speedup on Hermit for the 2d Laplace problem, solved with FETI-DP, 16 PE per FETI subdomain

4 Density Driven Flow

The transport of dissolved salt in flowing groundwater in porous media can be described by two nonlinear, coupled, time-dependent differential equations, describing the balance of the fluid-phase as a whole and the balance of the mass of brine (cf. [5]):

$$\begin{aligned} \partial_t(\phi\rho) + \nabla \cdot (\rho\mathbf{q}) &= 0, \\ \partial_t(\phi\rho\omega) + \nabla \cdot (\rho\omega\mathbf{q} - \rho D\nabla\omega) &= 0, \\ \mathbf{q} &= -\frac{K}{\mu}(\nabla p - \rho\mathbf{g}) \end{aligned} \quad (8)$$

where ω , p are the unknowns brine mass fraction and pressure, ϕ is porosity, K the permeability, μ the viscosity, ρ the density, \mathbf{g} the gravity field, and D the diffusion-dispersion tensor.

One well known benchmark is the so-called *Elder problem* [10]. In its 2d formulation it is defined on a domain $\Omega := [0,600 \text{ m}] \times [0,150 \text{ m}]$. Initially the brine mass fraction is set to zero in the whole domain. At the middle top, a dirichlet boundary conditions of 1 is set for the brine mass fraction while it is fixed to zero at the bottom. All other boundary parts are noflow boundaries. Due to density differences and the gravity the evolution of the system in time will lead to a downward flow of brine. See Fig. 7 for a computed solution. For 3d we use a corresponding setting on the domain $\Omega := [0,600 \text{ m}] \times [0,600 \text{ m}] \times [0,150 \text{ m}]$. Results are shown in Fig. 8.

This system of equations is discretized using a vertex-centered finite volume scheme [2]. As time stepping scheme the implicit Euler method is used.



Fig. 7 Computed solution of the Elder problem in 2d at an early (*left*) and a more evolved (*right*) time point




Fig. 8 Computed solution of the Elder problem in 3d at an early (*left*) and a more evolved (*right*) time point

Table 8 Scaling of Elder problem in 2d on Hermit

PE	level	DoF	$T_T(s)$	$T_N(s)$	N_N	$\frac{T_N}{N_N}(s)$	$T_{GMG}(s)$	N_{GMG}	$\frac{T_{GMG}}{N_{GMG}}(s)$
64	9	8,398,850	407.251	396.775	20	19.839	195.142	352	0.554
256	10	33,574,914	577.662	565.889	23	24.604	349.564	636	0.550
1,024	11	134,258,690	1424.820	1412.930	39	36.229	1099.180	1,956	0.562
4,096	12	536,952,834	1138.300	1088.380	40	27.210	793.910	1,413	0.562
16,384	13	2,147,647,490	1126.930	1036.010	40	25.900	775.619	1,380	0.562

T_N time for whole newton including assembling of jacobian, N_N total number of newton steps, T_{GMG} overall time spent in multigrid cycle, N_{GMG} number of multigrid cycles

The resulting discrete fully-coupled non-linear equations are solved by a Newton method. Within each Newton step the Jacobi-matrix of the problem is computed exactly. For this linearized matrix system the linear geometric multigrid solver is applied. As smoother a block-Jacobi smoother with inner ILU decomposition is used, where the Jacobi blocks correspond to the parts of the matrix that is owned by each parallel process.

Results

Table 8 shows the results of the weak scalability test for the 2d problem. It is well observed that the time for each multigrid cycle is constant. Since this problem is non-linear, the number of Newton steps does not remain constant, but is still in a reasonable range. Table 9 displays the corresponding results for the 3d problem.

Table 9 Scaling of Elder problem in 3d on Hermit

PE	level	DoF	$T_T(s)$	$T_N(s)$	N_N	$\frac{T_N}{N_N}(s)$	$T_{GMG}(s)$	N_{GMG}	$\frac{T_{GMG}}{N_{GMG}}(s)$
64	4	1,098,306	240.184	196.331	20	9.817	86.033	457	0.188
512	5	8,586,370	480.574	436.349	20	21.817	132.401	679	0.195
4,096	6	67,897,602	551.408	500.743	20	25.037	233.181	1,170	0.199
32,768	7	540,021,250	784.101	602.627	39	15.452	307.558	1,551	0.198

T_N time for whole newton including assembling of jacobian, N_N total number of newton steps, T_{GMG} overall time spent in multigrid cycle, N_{GMG} number of multigrid cycles

5 Conclusion and Outlook

This report presented some scalability studies for solvers in the ug4 library. Although the results are still preliminary, we showed that multigrid methods are scalable and efficient tools for solving partial differential equations on parallel supercomputers. On Hermit, the tests show scalability for up to 65,536 cores. Future work will be dedicated to the consolidation and enhancement of ug4 and its application to more real-world problems.

Acknowledgements This work has been supported by the Goethe-Universität Frankfurt and by the German Ministry of Economy and Technology (BMWi) via grant 02E10326 and grant 02E10568. We thank the HLRS for the opportunity to use Hermit and their kind support.

References

1. C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen. Feti-dp: a dual-primal unified feti method - part i: A faster alternative to the two-level feti method. *International Journal for Numerical Methods in Engineering*, 50(7):1523–1544, 2001.
2. P. Frolkovic. Finite volume discretizations of density driven flows in porous media. *Vilsmeier R. Benkhaldoun F., editor, Finite Volumes for Complex Applications*, pages 433–440, 1996.
3. W. Hackbusch. Multi-grid methods and applications, vol. 4 of springer series in computational mathematics, 1985.
4. A. Klawonn and O. Rheinbach. Highly scalable parallel domain decomposition methods with an application to biomechanics. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 90(1):5–32, 2010.
5. A. Leijnse. Three-dimensional modeling of coupled flow and transport in porous media. 1992.
6. A. Naegel, R. D. Falgout, and G. Wittum. Filtering algebraic multigrid and adaptive strategies. *Computing and Visualization in Science*, 11(3):159–167, 2008.
7. S. Reiter, A. Vogel, I. Heppner, M. Rupp, and G. Wittum. A massively parallel geometric multigrid solver on hierarchically distributed grids. *Computing and Visualization in Science*, 2012, in press.
8. J. W. Ruge and K. Stüben. *Multigrid Methods*, volume 3 of *Frontiers in Applied Mathematics*, chapter Algebraic multigrid (AMG), pages 73–130. SIAM, Philadelphia, PA, 1987.
9. A. Vogel, S. Reiter, M. Rupp, A. Nägel, and G. Wittum. Ug 4 - a novel flexible software system for simulating pde based models on high performance computers. *Computing and Visualization in Science*, 2012, in press.

10. C.I. Voss and W.R. Souza. Variable density flow and solute transport simulation of regional aquifers containing a narrow freshwater-saltwater transition zone. *Water Resources Research*, 23(10):1851–1866, 1987.
11. C. Wagner. On the algebraic construction of multilevel transfer operators. *Computing*, 65:73–95, 2000.

Simulation of Liquid-Liquid Equilibria with Molecular Models Optimized to Vapor-Liquid Equilibria and Model Development for Hydrazine and Two of Its Derivatives

Stefan Eckelsbach, Thorsten Windmann, Ekaterina Elts, and Jadran Vrabec

1 Introduction

In the chemical industry, knowledge on fluid phase equilibria is crucial for design and optimization of many technical processes. In a chemical plant, the costs for separation facilities constitute one of the highest investment outlays, typically in the order of 40–80 % [15]. Not only vapor-liquid equilibrium (VLE) data are of interest, e.g. for distillation columns, but also other types of phase equilibria. For example, liquid-liquid equilibrium (LLE) data provide the basis for extraction processes.

Classically, thermodynamic data for the design of such processes have to be measured experimentally and have to be aggregated by empirical correlations. For practical applications this leads to problems. For example, it is not possible to describe the entire fluid phase behavior consistently with a single model and set of parameters. Thus LLE data cannot be predicted reliably from VLE data (or vice versa) based on such correlations. Furthermore, the effort for measurements in the laboratory is very high, because every single fluid system of interest has to be measured individually. This approach particularly reaches its limits when multi-component fluids or systems with multiple phases are of interest due to the sheer amount of independent variables. In a recent study by Hendriks et al. [10] about the demand of thermodynamic and transport properties in the chemical industry, the urgent need for a reliable and predictive approach to describe VLE as well as LLE with a single model and parameter set is pointed out.

In this work, the capability of molecular simulation to determine fluid phase equilibria is examined, especially with respect to the problems of the classical approach described above. Molecular modeling and simulation is a modern route

S. Eckelsbach · T. Windmann · E. Elts · J. Vrabec (✉)
Lehrstuhl für Thermodynamik und Energietechnik (ThEt), Universität Paderborn,
Warburger Str. 100, 33098, Paderborn, Germany
e-mail: jadran.vrabec@upb.de

for the prediction of thermodynamic properties. Being based on mathematical representations of the intermolecular interactions, it has strong predictive capabilities as it adequately covers structure, energetics and dynamics on the microscopic scale that govern the fluid behavior on the macroscopic scale. In preceding work, molecular models (force fields) were developed to accurately describe pure substance VLE data [21] and were successfully assessed with respect to VLE data of binary and ternary mixtures [18]. In the present work, the capability of those models to predict LLE data is studied for the binary mixture Nitrogen + Ethane. The decomposition of a randomly dispersed mixture of these two components in their LLE phases is investigated. Subsequently, the composition of the phases in equilibrium is compared to experimental data. It is shown that VLE and LLE can be predicted by molecular simulation in good agreement with experimental results.

Furthermore, three new models for Hydrazine, Monomethylhydrazine and Dimethylhydrazine were developed. They are based on quantum chemical information and were optimized to experimental data on saturated liquid density and vapor pressure. Thereafter, the models were assessed by comparing simulated VLE to experimental reference data, which they are able to reproduce.

2 Phase Equilibria

2.1 Vapor-Liquid Equilibria

The molecular model for the mixture Nitrogen + Ethane was developed in a preceding work of our group [20]. The application to mixtures can be done straightforwardly by assigning the unlike interaction parameters of the two components A and B. They are defined by the modified Lorentz-Berthelot combination rules. It was already shown in [20] that one additional parameter in the equation for calculating the unlike energy parameter ϵ_{AB} , which allows to adjust the simulation results to the VLE data of the mixture, leads to an improvement of the predictive quality of the model.

$$\sigma_{AB} = \frac{\sigma_A + \sigma_B}{2} \quad (1)$$

$$\epsilon_{AB} = \xi \sqrt{\epsilon_A + \epsilon_B} \quad (2)$$

The binary parameter of the mixture Nitrogen + Ethane was defined as $\xi = 0.974$ [20]. Figure 1 shows the VLE phase diagram at 200 and 290 K as determined on the basis of this mixture model in comparison to experimental data. The Peng-Robinson equation of state is given as an example of a classical correlation approach. These simulations coincide closely with the experimental reference values for pressures below 7 MPa. With raising pressure, the deviations also increase, but the values

Fig. 1 Vapor-liquid equilibria of the mixture Nitrogen and Ethane: simulation results (●); experimental data (+) [4]; Peng-Robinson equation of state (—). The statistical uncertainties of the present data are within symbol size





Fig. 2 Progress of the mole fraction of Nitrogen in phase 1 (▲) and phase 2 (■) over the duration of the simulation t : The simulation was performed for a temperature of 128 K and a pressure of 11.03 MPa



are closer to the experiment than the Peng-Robinson equation of state, which overestimates the critical region and shows significant deviations particularly on the saturated liquid line.

2.2 Liquid-Liquid Equilibria

Molecular dynamics simulations were performed for a mixture of 20,000 molecules, consisting of 40 mol-% Ethane and 60 mol-% Nitrogen. The two components were randomly dispersed in the initial configuration. The simulations were carried out with our molecular dynamics code *ls1 mardyn*. They were performed in a canonical ensemble and the length of one time step was set to 2 fs.

Starting from a randomly dispersed mixture, it was found that the system decomposes spontaneously into two coexisting liquid phases. As an example, the progress of a simulation at the pressure 11.03 MPa and the temperature 128 K is plotted in Fig. 2. The simulation required about $2.4 \cdot 10^7$ time steps to lead to an equilibrated state, which represents a typical duration of the present simulations of around 50 ns. Thereafter, the two phases can be clearly identified by the mole fraction over

Fig. 3 Mole fraction of Nitrogen over the length of the simulation volume in z direction: the simulation was carried out at a temperature of 128 K and a pressure of 11.03 MPa





Fig. 4 Liquid-liquid equilibria of the mixture Nitrogen + Ethane: simulation results (●) for a temperature of approximately 128 K; experimental data (+) for a temperature of 126.7–129 K [4]



the length of the simulation volume, which is presented in Fig. 3. This provides the ability of predicting the LLE phase behavior for different thermodynamic conditions. The comparison of the present results with experimental data is shown in Fig. 4. The simulated mole fractions agree well with the experimental data and so they also reproduce the pressure dependence of the LLE.

3 Hydrazines

3.1 Molecular Model Class

The present molecular models include three groups of parameters. These are (1) the geometric parameters, specifying the positions of different interaction sites, (2) the electrostatic parameters, defining the polar interactions in terms of point charges, dipoles or quadrupoles, and (3) the dispersive and repulsive parameters,

determining the attraction by London forces and the repulsion by electronic orbital overlaps. Here, the Lennard-Jones (LJ) 12-6 potential [11, 12] was used to describe the dispersive and repulsive interactions. The total intermolecular interaction energy thus writes as

$$\begin{aligned}
 U = & \sum_{i=1}^{N-1} \sum_{j=i+1}^N \left\{ \sum_{a=1}^{S_i^{\text{LJ}}} \sum_{b=1}^{S_j^{\text{LJ}}} 4\varepsilon_{ijab} \left[\left(\frac{\sigma_{ijab}}{r_{ijab}} \right)^{12} - \left(\frac{\sigma_{ijab}}{r_{ijab}} \right)^6 \right] + \right. \\
 & \sum_{c=1}^{S_i^e} \sum_{d=1}^{S_j^e} \frac{1}{4\pi\epsilon_0} \left[\frac{q_{ic}q_{jd}}{r_{ijcd}} + \frac{q_{ic}\mu_{jd} + \mu_{ic}q_{jd}}{r_{ijcd}^2} \cdot f_1(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j) + \frac{q_{ic}Q_{jd} + Q_{ic}q_{jd}}{r_{ijcd}^3} \cdot f_2(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j) + \right. \\
 & \left. \left. \frac{\mu_{ic}\mu_{jd}}{r_{ijcd}^3} \cdot f_3(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j) + \frac{\mu_{ic}Q_{jd} + Q_{ic}\mu_{jd}}{r_{ijcd}^4} \cdot f_4(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j) + \frac{Q_{ic}Q_{jd}}{r_{ijcd}^5} \cdot f_5(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j) \right] \right\}, \quad (3)
 \end{aligned}$$

where r_{ijab} , ε_{ijab} , σ_{ijab} are the distance, the LJ energy parameter and the LJ size parameter, respectively, for the pair-wise interaction between LJ site a on molecule i and LJ site b on molecule j . The permittivity of the vacuum is ϵ_0 , whereas q_{ic} , μ_{ic} and Q_{ic} denote the point charge magnitude, the dipole moment and the quadrupole moment of the electrostatic interaction site c on molecule i and so forth. The expressions $f_x(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j)$ stand for the dependence of the electrostatic interactions on the orientations $\boldsymbol{\omega}_i$ and $\boldsymbol{\omega}_j$ of the molecules i and j , cf. [1, 7]. Finally, the summation limits N , S_x^{LJ} and S_x^e denote the number of molecules, the number of LJ sites and the number of electrostatic sites, respectively.

3.2 Molecular Pure Substance Models

For all molecular models developed in the present work, the internal degrees of freedom were neglected and the models were chosen to be rigid. As a first step, the geometric data of the molecules, i.e. bond lengths, angles and dihedrals, were determined by QC calculations. Therefore, a geometry optimization was performed via an energy minimization using the GAMESS (US) package [17]. The Hartree-Fock level of theory was applied with a relatively small (6–31 G) basis set. All LJ parameters and the charge magnitudes were initially taken from prior models and fine tuned during the model parameter optimization to vapor pressure and saturated liquid density.

Figure 5 shows the developed molecular models.

The results for saturated densities obtained with the present model are compared to the available experimental data [4, 5, 8] and to the simulation results by Gutowski et al. [8] in Fig. 6.

Fig. 5 Snapshot of Hydrazine (*top, left*), Monomethylhydrazine (*top, right*) and Dimethylhydrazine





Fig. 6 Saturated densities of Hydrazine (●), Monomethylhydrazine (▲) and Dimethylhydrazine (■): (*striped symbols*) experimental critical data [4, 5, 8, 14]; (+) experimental saturated liquid densities [4, 5]; (*full symbols*) simulation data on the basis of the present models; (*empty symbols*) simulation data (saturated liquid only) by Gutowski et al. [8]; (—) correlation of experimental data [5]; (- - -) correlation of present simulation data [9]. The statistical uncertainties of the present data are within symbol size

3.3 Binary Vapor-Liquid Equilibria

Based on the discussed three molecular hydrazine models, VLE data were predicted for all three binary Hydrazine mixtures with Water as well as for the mixture Dimethylhydrazine + Hydrazine.

3.3.1 Water + Hydrazine

Figure 7 shows the isobaric VLE of Water + Hydrazine at 0.1013 MPa from experiment [4, 13, 19], simulation and Peng-Robinson EOS. The mixture is azeotropic, having a temperature maximum. The azeotropic point is at $x_{\text{H}_2\text{O}} \approx 0.41$ mol/mol. The experimental vapor pressure by Uchida et al. [19] at 388.25 K and

Fig. 7 Isobaric vapor-liquid phase diagram of Water + Hydrazine at 0.1013 MPa: (×) experimental data by Lobry de Bruyn and Dito [13]; (+) experimental data by Uchida et al. [19]; (●) present simulation data with $\xi = 1.3$; (—) Peng-Robinson EOS with $k_{ij} = -0.1325$





Fig. 8 Isobaric vapor-liquid phase diagram of Monomethylhydrazine + Water at 0.1013 MPa: (×) experimental data by Ferriol et al. [6]; (+) experimental data by Cohen-Adad et al. [2]; (●) present simulation data with $\xi = 1.3$; (—) Peng-Robinson EOS with $k_{ij} = -0.197$



$x_{\text{H}_2\text{O}} = 0.6925$ mol/mol was taken to adjust the binary parameter of the molecular model ($\xi = 1.3$). Considering the substantial experimental uncertainties, both data sets agree very favorably.

3.3.2 Monomethylhydrazine + Water

Figure 8 depicts the VLE of Monomethylhydrazine + Water at ambient pressure. Like aqueous Hydrazine, this mixture is azeotropic, having a temperature maximum. In this case, the azeotropic point lies at $x_{\text{CH}_3-\text{N}_2\text{H}_3} \approx 0.25$ mol/mol. The experimental data by Ferriol et al. [6] at 372.55 K and $x_{\text{CH}_3-\text{N}_2\text{H}_3} = 0.476$ mol/mol were taken to adjust the binary parameter of the molecular model ($\xi = 1.3$). In the Water-rich region, to the left of the azeotropic point in Fig. 8, VLE simulations were not feasible, because of sampling problems. Considering the substantial experimental uncertainties, the data sets from all three approaches agree very favorably.

Fig. 9 Isobaric vapor-liquid phase diagram of Dimethylhydrazine + Water at 0.1013 MPa: (\times) experimental data by Carleton [3]; (+) experimental data by Ferriol et al. [6]; (\bullet) present simulation data with $\xi = 1.3$; (—) Peng-Robinson EOS with $k_{ij} = -0.285$





Fig. 10 Isobaric vapor-liquid phase diagram of Dimethylhydrazine + Hydrazine at 0.1013 MPa: (+) experimental data by Pannetier and Mignotte [16]; (\bullet) present simulation data with $\xi = 1.01$; (—) Peng-Robinson EOS with $k_{ij} = -0.1$



3.3.3 Dimethylhydrazine + Water

Figure 9 shows the isobaric VLE of Dimethylhydrazine + Water at 0.1013 MPa from experiment, simulation and the Peng-Robinson equation of state. In contrast to the the previous two binary systems, this mixture is zeotropic. The experimental vapor pressure by Ferriol et al. [6] at 345.17 K and $x_{(\text{CH}_3)_2\text{-N}_2\text{H}_2} = 0.571$ mol/mol was taken to adjust the binary parameter of the molecular model ($\xi = 1.3$). It can be seen in the Fig. 9 that the results obtained by molecular simulation agree well with the experimental results on the saturated liquid line, but overestimate the Dimethylhydrazine content on the saturated vapor line for intermediate compositions.

3.3.4 Dimethylhydrazine + Hydrazine

The VLE of Dimethylhydrazine + Hydrazine is presented in Fig. 10 at ambient pressure. This system is zeotropic. The experimental vapor pressure by Pannetier and Mignotte [16] at 346.35 K and $x_{(\text{CH}_3)_2\text{-N}_2\text{H}_2} = 0.4717$ mol/mol was taken to

adjust the binary parameter of the molecular model ($\xi = 1.01$). It can be seen that simulation results match almost perfectly with the experimental data on the saturated liquid line. On the saturated vapor line, the experimental data and the simulation results exhibit some scatter, but the agreement is reasonable.

4 Conclusion

It was shown that molecular models that were adjusted to VLE data provide the option to simulate LLE and predict their pressure dependence. The simulations yield reliable data that are in good agreement with experimental values. In contrast to the classical approach to provide phase equilibrium data, this route has the capability of predicting VLE and LLE data with a single model and set of parameters.

Furthermore, molecular modeling and simulation was applied to predict the VLE phase behavior of pure fluids and binary mixtures containing Hydrazine and two of its derivatives. New molecular models were developed for Hydrazine, Monomethylhydrazine and Dimethylhydrazine, partly based on quantum chemical information on molecular geometry and electrostatics. Experimental data on the saturated liquid density and the vapor pressure were taken into account to optimize the pure substance models. These pure substance properties were represented accurately from the triple point to the critical point.

For an optimized description of the binary VLE, the unlike dispersive interaction was adjusted for all studied binary systems to a single experimental vapor pressure of the mixture in the vicinity of ambient conditions. With these binary mixture models, VLE data were predicted for a temperature and composition range. The predictions show a good agreement with experimental binary VLE data that were not considered in the model development.

In this work, molecular modeling and simulation was used to predict the VLE phase behavior and the thermodynamic properties of pure hydrazines and binary aqueous hydrazine mixtures, for which experimental data were available for comparison. The presented molecular models were able to well reproduce the experimental data that were not considered in the model development. Thus, these new models could also be valuable for the prediction of properties under different conditions and for systems, where no experimental data are available.

Acknowledgements We gratefully acknowledge support by Deutsche Forschungsgemeinschaft. This work was carried out under the auspices of the Boltzmann-Zuse Society (BZS) of Computational Molecular Engineering. The simulations were performed on the NEC SX-9, the NEC Nehalem Cluster and the Cray XE6 (Hermit) at the High Performance Computing Center Stuttgart (HLRS).

References

1. Allen, M. P.; Tildesley, D. J.: Computer simulations of liquids. *Oxford University Press, Oxford* **1987**.
2. Cohen-Adad, M. T.; Allali, I.; Getzen F. W. *J. Solution Chem.* **1987** *16*, 659.
3. Carleton, L. T. *Ind. Eng. Chem. Chem. Eng. Data Series* **1956** *1*, 21.
4. Dortmundur Datenbank, Mixture Properties. Version 6.3.0.384 **2010**.
5. Rowley, R. L.; Wilding, W. V.; Oscarson, J. L.; Yang, Y.; Zundel, N. A.; Daubert, T. E.; Danner, R. P.: DIPPR Data Compilation of Pure Chemical Properties. *Design Institute for Physical Properties, AIChE, New York* **2011**.
6. Ferriol, M.; Laachach, A.; Cohen-Adad, M. T.; Getzen, F. W.; Jorat, L.; Noyel, G.; Huck, J.; Bureau, J. C. *Fluid Phase Equilib.* **1992** *71*, 287.
7. Gray, C. G.; Gubbins K. E.: Theory of molecular fluids, 1. Fundamentals. *Clarendon Press, Oxford* **1984**.
8. Gutowski, K. E.; Gurkan, B.; Maginn E. J. *Pure Appl. Chem.* **2009** *81*, 1799.
9. Guggenheim, E. A. *J. Chem. Phys.* **1945** *13*, 253.
10. Hendriks, E.; Kontogeorgis, G. M.; Dohrn, R.; de Hemptinne, J. C.; Economou I. G.; Žilnik, L. F.; Vesovic, V. *Ind. Eng. Chem. Res.* **2010** *49*, 11131.
11. Jones, J. E. *Proc. Roy. Soc.* **1924** *106A*, 441.
12. Jones, J. E. *Proc. Roy. Soc.* **1924** *106A*, 463.
13. Lobry de Bruyn, C. A.; Dito, J. W. *Proc. Sec. Sci./K. Akad. Wet. Amsterdam* **1902-1903** *5*, 171.
14. Hradetzky, G.; Lempe, D. A.: Merseburg Data Bank MDB for thermophysical data of pure compounds. Revision 7.1.0 **2010**.
15. Prausnitz, J. M.; Lichtenthaler, R. N.; de Azevedo, E. G.: Molecular Thermodynamics of Fluid-Phase Equilibria. *Prentice-Hall: Upper Saddle River, New Jersey* **1999**.
16. Pannetier, G.; Mignotte, P. *Bull. Soc. Chim. Fr.* **1961** *143*, 985.
17. Schmidt, M. W.; Baldrige, K. K.; Boatz, J. A.; Elbert, S. T.; Gordon, M. S.; Jensen, J. H.; Koseki, S.; Matsunaga, N.; Nguyen, K. A.; Shujun, S.; Windus, T. L.; Dupuis, M.; Montgomery, A. M. *J. Comput. Chem.* **1993** *14*, 1347.
18. Stoll, J.; Vrabec, J.; Hasse, H. *AIChE J.* **2003** *49*, 2187.
19. Uchida, S.; Ogawa, S.; Yamaguchi, M. *Jap. Sci. Rev. Eng. Sci.* **1950** *1*, 41.
20. Vrabec, J.; Huang, Y.-L.; Hasse, H. *Fluid Phase Equilib.* **2009** *279*, 120.
21. Vrabec, J.; Stoll, J.; Hasse, H. *J. Phys. Chem. B* **2001** *105*, 12126.

A Particle-in-Cell Method to Model the Influence of Partial Melt on Mantle Convection

Ana-Catalina Plesa, Doris Breuer, and Tilman Spohn

Abstract Solid-state convection is the principal mechanism that controls the global dynamics and thermal evolution of the terrestrial planets. Observations such as seismology and mission data from geological structures at the planetary surfaces offer important constraints for the interior dynamics. However, our main knowledge stems from laboratory experiments and in particular from computer models. In the last years, due to the significant improve of high performance computing, computer simulations have become the most powerful access to this fluid-dynamical problem by solving partial differential equations in a discrete formulation to describe the flow in space and time. In the present work we will present results obtained using the cylindrical/spherical code GAIA with a particle-in-cell method to account for compositional changes due to partial melting of the mantle.

1 Introduction

Thermal and compositional buoyancy drive the slow creeping flow of rocky planetary mantles that is ultimately responsible for the heat transport efficiency, magnetic field generation and the formation of surface geological structures such as volcanoes

A.-C. Plesa (✉)

German Aerospace Centre, Institute of Planetary Research, Berlin and WWU, Institute of Planetology, Muenster, Germany
e-mail: aples01@uni-muenster.de

D. Breuer

German Aerospace Centre, Institute of Planetary Research, Berlin, Germany
e-mail: doris.breuer@dlr.de

T. Spohn

WWU, Institute of Planetology, Muenster, Germany
e-mail: spohn@uni-muenster.de

and tectonic plates. Besides thermal convection, which is driven by secular cooling, by the heat from the metallic core and the internal heat due to the decay of radioactive isotopes, heterogeneities due to different chemical compositions can strongly affect the planform of mantle convection and thus the heat transport and planetary evolution. Compositional heterogeneities in planetary mantles can occur on all spatial scales. On Earth, seismic data suggests that the region historically known as D'' occupying the lowermost few 100 km of the mantle atop the core-mantle boundary is chemically distinct. On Mars, the so-called SNC meteorites likely originate from separate chemical reservoirs that have been preserved for the entire evolution of the planet. These meteorites are basalts and basaltic cumulates samples which hold a record of volcanism on Mars over most of its history since accretion and differentiation, to recent times [21]. Geochemical analysis of the SNC meteorites implies the presence of at least three isotopically distinct reservoirs, two of them being depleted in lithophile elements and therefore most likely situated in the mantle and a third one enriched and assumed to lie in the crust (e.g. [5, 18]).

Until now, due to computational reasons, most of the geodynamic models are highly simplified and use only thermal buoyancy to compute the slow creep flow inside the planetary mantles, not taking into account partial melting and its associated effects. However, to investigate the above mentioned features, compositional buoyancy due to partial melting is one of the most important ingredients.

In the present project, we present a particle-in-cell method which has been adapted to investigate the influence of partial melt on the global dynamic of planetary mantles using numerical simulations using the 2D cylindrical and 3D spherical convection code GAIA [12, 13]. The model was applied in particular to Mars using constraints inferred by the data analyzed from the SNC meteorites [21].

2 Mantle Convection Model

The non-linear nature of convection enables analytical solutions to be found only for strongly simplified scenarios. In addition, laboratory experiments cover only a limited range of parameters, not always relevant for planetary evolution. Over the years due to the massive increase in computational power, numerical simulations have grown to be one of the most powerful tools in solving fluid dynamical problems. Today, parallel computing has established as a standard procedure in solving the partial differential equations describing the temporal evolution of the mantle flow with a complex rheology in different geometries. To model thermal mantle convection, conservation equation of mass, momentum, energy and composition are solved in a numerical sense [26]. These equations are scaled with the mantle thickness D as a length scale and thermal diffusivity κ as a time scale. Their nondimensional formulation assuming an incompressible fluid in a Boussinesq approximation with a Newtonian rheology and an infinite Prandtl number is [7]:

$$\nabla \cdot \mathbf{u} = 0 \quad (1)$$

$$\nabla \cdot [\eta(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)] + (RaT + Ra_C F)\mathbf{e}_r - \nabla p = 0 \quad (2)$$

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T - \nabla k \nabla T - \frac{Ra_Q}{Ra} + (T + T_{surf}) \frac{\Delta S}{c_p} \frac{\partial F}{\partial t} = 0 \quad (3)$$

$$\frac{\partial C}{\partial t} + \mathbf{u} \cdot \nabla C + \Gamma(F) = 0 \quad (4)$$

where \mathbf{u} is the velocity field, η is the viscosity, T is the temperature, C is the chemical component, \mathbf{e}_r is the unity vector in radial direction, p is the pressure, T_{surf} is the surface temperature, ΔS is the entropy change upon melting, c_p is the heat capacity at constant pressure, k is the thermal conductivity, t is the time, F is the melt fraction and $\Gamma(F)$ is a function, which describes the changes of the chemical component depending on the melt fraction.

In the Eqs. (2) and (3) Ra is the thermal Rayleigh number and Ra_Q is the Rayleigh number for internal heat sources. Both Ra and Ra_Q are related to thermal buoyancy. Ra_C is the compositional Rayleigh number which accounts for the buoyancy due to chemical heterogeneities. These Rayleigh numbers are defined as follows [1, 4]:

$$Ra = \frac{\rho g \alpha \Delta T D^3}{\kappa \eta} \quad (5)$$

$$Ra_Q = \frac{\rho^2 g \alpha Q_m D^5}{\kappa k \eta} \quad (6)$$

$$Ra_C = \frac{\Delta \rho g D^3}{\kappa \eta} \quad (7)$$

with ρ the mantle density, g gravitational acceleration, ΔT temperature difference between surface and core-mantle boundary, D mantle thickness, κ thermal diffusivity, η reference viscosity, Q_m mantle radioactive heat sources, k thermal conductivity and $\Delta \rho$ density difference upon mantle differentiation.

The viscosity is calculated according to the Arrhenius law for diffusion creep [15]. The non-dimensional formulation of the Arrhenius viscosity law for temperature and depth dependent viscosity [25] is given by:

$$\eta(r, T) = \exp\left(\frac{E + (r_p - r)V}{T + T_{surf}} - \frac{E + (r_p - r_{ref})V}{T_{ref} + T_{surf}}\right) \quad (8)$$

where E is the activation energy, T_{surf} is the surface temperature, V is the activation volume, r is the radius, and r_p is the planet radius. T_{ref} and r_{ref} are reference values for temperature and radius.

In our simulations we consider a one-plate planet with cooling boundary conditions and decaying radioactive elements. Radioactive elements, namely the

uranium isotopes ^{235}U and ^{238}U , the thorium isotope ^{232}Th , and the potassium isotope ^{40}K , are considered [1]:

$$Q_m = 0.9928C_0^U H^{U^{238}} \exp\left(\frac{t \ln 2}{\tau_{1/2}^{U^{238}}}\right) + 0.0071C_0^U H^{U^{235}} \exp\left(\frac{t \ln 2}{\tau_{1/2}^{U^{235}}}\right) + C_0^{Th} H^{Th} \exp\left(\frac{t \ln 2}{\tau_{1/2}^{Th}}\right) + 1.1910^{-4}C_0^K H^{K^{40}} \exp\left(\frac{t \ln 2}{\tau_{1/2}^{K^{40}}}\right) \quad (9)$$

where H^i are the rates of heat production, $\tau_{1/2}^i$ the half lives of the isotopes, C_0^i the initial concentrations and t the time.

During the melting process, the extraction of melt and formation of crust redistributes the radioactive heat producing elements (HPEs). Being highly incompatible, HPEs are enriched in the melt and by extracting the melt to form crust, the HPEs become depleted in the residual mantle. The radioactive heat sources are calculated using the accumulated fractional melting formula which allows us to obtain their concentration in the melt [6, 20]:

$$\Gamma_F = \frac{\Gamma_{bulk}}{F} \left(1 - (1 - F)^{1/\delta}\right) \quad (10)$$

The term Γ_{bulk} is the initial concentration in the mantle and δ the partitioning coefficient [16]. Typical values of $\delta = 0.002$ for the radioactive heat sources are used. These values are obtained when averaging the partitioning coefficients by their mineral percentage in the mantle. The lower the partitioning coefficient the faster the mantle depletion in the incompatible elements. Assuming accumulated fractional melting and the partitioning coefficients from [6], nearly complete extraction of heat-producing elements for melt fractions greater than 1 %.

3 Technical Realization

We consider the mantle convection in 2D cylindrical [23] and 3D spherical shells using the code GAIA [12, 13]. GAIA uses BICGStab with a Jacobi preconditioner to solve the resulting linear systems of equations. For the matrix storage a Harwell-Boeing sparse matrix class has been chosen.

The discretization of the governing equations is based on the finite-volume method with the advantage of utilizing fully irregular grids in three and two dimensions, efficiently parallelized for up to 396 CPUs [12, 13]. The space is discretized by a fixed grid while for the temporal discretization a fully implicit 2nd-order method, also called an implicit three-level scheme, after [9] has been used. In contrast to spatial discretization, the temporal discretization is flexible and can adapt with a varying time step Δt to the situation. A method proposed by Caretto et al. in [2] and Patankar in [22] called SIMPLE was adopted to solve the coupling of the continuity equation with the momentum equation.

The model was validated by a comparison with analytically known solutions as well as published numerical results [12, 23]. A comparison with a commercial product also yielded satisfying results. A convergence test with successively refined grids proved the convergence of global quantities towards an extrapolated solution [14].

When modelling thermo-chemical convection, one of the most challenging parts is the handling of Eq. (4). In [24] two methods for modelling active compositional fields have been compared. The grid based method substitutes Eq. (4) by a partial differential equation alike the heat transport equation with a marginal diffusivity:

$$\frac{\partial C}{\partial t} + \mathbf{u} \cdot \nabla C - \frac{1}{Le} \nabla^2 C + \Gamma(F) = 0 \quad (11)$$

where Le is the Lewis number, which is the ratio of thermal diffusivity to chemical diffusivity. Major disadvantages of this method are the non-negligible numerical diffusivity and the simultaneous advection of multiple materials with different physical properties (e.g. [8]), for which further equations have to be solved [24].

A better approach to model Eq. (4) is to use a particle-in-cell method [24]. In this case, massless particles are advected in every time step according to the velocity field. We have adapted this method to account for multiple physical properties which change during the melting process. In our simulations, particles can carry properties like density, thermal conductivity, radioactive heat sources, water concentration etc. At the beginning of a simulation the particles are distributed within each cell of the computational domain. The initial conditions are interpolated from the corresponding fields onto the particles. In this way complex initial conditions can easily be imposed on the particles. In every time-step of the simulation the particles are moved using the velocity field computed from the momentum equation (Eq. (2)). The next position of a particle is then computed by solving a trajectory equation using the Runge–Kutta 4th-order method. After the new position of every particle has been determined, the particles' physical values are interpolated back onto the grid to solve the Eqs. (1–3) in the next time-step.

The drawback of this method compared to the grid-based approach is the increase in both memory and computational time (see Fig. 1). However, this is natural keeping in mind that 15–20 particles per cell are needed in order to maintain the accuracy when interpolating back from the particles to the field [3]. The major advantage are the diffusivity free method and the fact that multiple different physical properties can be handled simultaneously.

At the moment, beside GAIA, only a small number of numerical codes worldwide [17, 28] can handle 3D spherical geometry using a particle-in-cell method to account for the active compositional fields.




Fig. 1 Slowdown factor depending on the number of particles used. A 2D cylindrical test with a grid resolution of 1.5×10^5 computational points and an increasing number of particles per cell has been used. The choice of 5 particles per cell results in a total of 7.4×10^5 particles in the whole computational domain, while 30 per cell correspond to a total of 4.4×10^6 particles. The benchmark has been performed using 96 computational cores at the SCC Karlsruhe on the XC4000 supercomputer

4 Results and Discussion

First we present a comparison between the grid-based and particle-in-cell methods to illustrate the advantages of the latter when modelling active compositional fields. For this we consider a test where a dense ball of composition 1 sinks into a less dense medium which has the composition 0. Both the ball and the surrounding medium have the same viscosity. During the movement of the ball, its shape will deform. For this test we impose free-slip boundary conditions for the velocity and turn off the solving of the energy equation since we are only interested in the buoyancy caused by the compositional heterogeneities.

The compositional Rayleigh number Ra_C is set to 1. The tests were performed using both grid-based and particle-in-cell methods with a 1.5×10^5 computational points grid. In the particle-based method case we use 30 particles per cell resulting in a total of 4.5×10^6 particles for the entire computational domain.

Figure 2 shows the superiority of the particle-in-cell method. While deformation takes place in both cases (i.e. grid-based method and particle-in-cell-method), we clearly observe the pronounced numerical diffusion in the grid-based case. Therefore for the rest of the presented results we use the particle-in-cell method to compute the changes in different physical properties due to melting.




Fig. 2 (a) Particle-in-cell case, (b) grid-based method, (c) difference image between (a) and (b). In the grid-based case the ball's shape diffuses whereas in the particle-in-cell case the ball in spite its deformation remains separate from the surrounding material

We compare next two cases including partial melting and its associated effects in a 2D cylindrical and a 3D spherical geometry. For the 2D cylindrical geometry case we use a geometry factor to rescale the inner to outer perimeter in order to match the inner to outer surface ratio in the 3D geometry [11]. Tests show that without using this scaling the heat-flow from the core into the mantle is overestimated in the 2D geometry resulting in about 100 K higher mantle temperatures. This can have a major influence, overestimating the melt fraction and the crustal volume produced during the thermal evolution.

For the 2D cylindrical geometry we use a 1.5×10^5 computational points grid resulting in a 10 km resolution and a total number of 2.2×10^6 particles for the entire computational domain. In the 3D spherical geometry we use a 2.8×10^6 computational points grid resulting in a 30 km resolution and a total number of 2.2×10^7 particles for the entire computational domain.

The melting effects considered in these cases are (1) the influence of the melting temperature due to the loss of low-melting point components during the melting process [19], (2) decrease of the density of the mantle material due to the mantle depletion in crustal components [4], (3) decrease of thermal conductivity in the crust [27] and (4) radioactive elements redistribution [10].

When melting occurs in the mantle, the crustal volume and therefore the corresponding crustal thickness are calculated. The physical properties stored on the particles are then changed depending of the particle position (crust or mantle). For the particles which lie in the crust, the thermal conductivity is lowered and the amount of radioactive elements is increased. The mantle particles become depleted in radioactive elements while in this test we don't consider a change in density between depleted and undepleted mantle material.

Figure 3 shows a good agreement between the 2D cylindrical geometry and 3D spherical geometry case keeping in mind the different resolutions applied for the 2D




Fig. 3 (a) Temperature slice (2D cylindrical geometry) – partial melt regions in white, (b) temperature iso-surface (3D spherical geometry) – partial melt regions in red, and (c) radial profiles (the dashed lines show the 2D and the full lines the 3D geometry respectively) at $t = 2.5$ Ga for a case with 1, 600 K initial mantle temperature and no compositional changes upon mantle differentiation

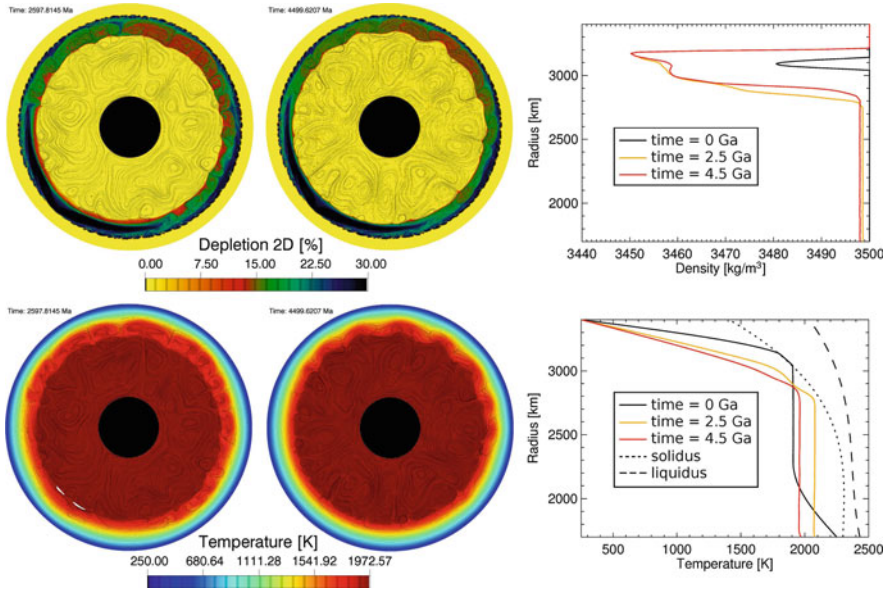


Fig. 4 Depletion slices and radial profiles translated as density (upper row) and temperature slices and temperature radial profiles (lower row) at $t = 2.5$ Ga and $t = 4.5$ Ga for a case with 1,900 K initial mantle temperature and compositional changes, which cause a difference in the mantle density of 60 kg/m^3 upon 30 % mantle differentiation

and 3D case (i.e. 10 km resolution in the 2D case whereas only 30 km resolution has been used in the 3D case).

For Mars the SNC meteorites, suggest that on Mars separate reservoirs have formed early in the planets’ evolution and did not mix since. To explain the formation of such reservoirs, we consider density variations due to the extraction of crustal components during the melting process. Density difference of 60 kg/m^3 upon 30 % differentiation (from peridotite – fertile mantle material – to residues like harzburgite) is considered. The composition for higher degrees of depletion like dunite (up to 60 % depletion) does not cause any changes in the residue density. Therefore the maximum density variations are reached between peridotite and harzburgite (0–30 % depletion).

The test in Fig. 4 was performed for a dry mantle rheology, which assumes a reference viscosity of 10^{21} Pa s assuming an initially depleted layer and neglecting the effects of partial melting on the thermal conductivity and radioactive heat source redistribution. Figure 4 shows density variations resulting in mantle inhomogeneities which form a depleted layer that prevents the lower mantle from efficient cooling. In this case hot material from the lower mantle rises to shallower depths where it can melt. The heterogeneities produced due to mantle depletion in this case remain stable during the entire thermal evolution and could explain the isotopic characteristics of the SNC meteorites.

5 Conclusions and Outlook

In this work we have presented results using 2D cylindrical and 3D spherical convection models which account for effects of partial melt during the thermal evolution of terrestrial planets. We have adapted a particle-in-cell method to account for the changes in physical properties due to melting and have applied our model in particular to Mars.

The results presented in Sect. 4 show (1) the superiority of the particle-in-cell method compared to a grid-based approach, (2) a good agreement of 2D cylindrical and 3D spherical geometry results when using a scaling of the inner to outer surfaces in 2D to match the surfaces ratio in 3D geometry and (3) the formation of chemical reservoirs which remain stable during the entire thermal evolution, as inferred by the SNC meteorites for Mars.

Over the past 2 years we have improved our mantle convection code to account for partial melting effects by using a particle-in-cell method in both 2D and 3D geometry. Satisfying results have been achieved when applying our methods to Mars-like parameters. In a follow-on project, we plan to investigate the effects of partial melting on terrestrial exo-planets. In the last several years, space missions like CoRoT and Kepler have provided us with promising candidates for terrestrial exo-planets (i.e. terrestrial planets outside the Solar System with masses less than 10 Earth masses and/or radii around or below 2 Earth radii).

Due to the measurements errors in mass and radius, the interior structure of an exo-planet and hence the mantle thickness is poorly constrained. This can have a major effect on the partial melt production in the mantle, since for a thin mantle, the planet is expected to cool faster and hence less amount of melt is produced during the planet's thermal evolution. Therefore, we plan to apply our model to investigate how the mantle thickness influences the thermal evolution and the partial melting in the mantle. For this we plan tests with (1) Mercury-like, i.e. thin mantle, (2) Earth-like and (3) Moon-like interior structure, i.e. thick mantle.

Due to the detection methods, most of the discovered exoplanets are closed to their stars and therewith tidally locked (i.e. during it's rotation around the star, the planet faces the star with the same side). In fact, this is also suggested for the two terrestrial exoplanets CoRoT 7b and Kepler 10b. We will use our model to investigate how the surface temperature variations influence the convection structure and the partial melting in the mantle. For this study tests with (1) uniform surface temperature, (2) day-night scenario and (3) cold poles and warm equator will be performed.

Both mantle thickness and surface temperature variations can have a major effect on the partial melt production and therefore on the entire thermal evolution of a terrestrial exo-planet. Therefore we plan to apply the GAIA mantle convection code to investigate the thermal evolution of terrestrial exo-planets accounting for partial melting, various mantle thickness and surface temperature variations.

Acknowledgements This research has been supported by the Helmholtz Association through the research alliance “Planetary Evolution and Life”. The project is conducted at the SCC Karlsruhe on the XC4000 supercomputer.

References

1. Breuer, D.: Dynamics and thermal evolution. In: Landolt-Börnstein Astronomy and Astrophysics (Group VI), 4 - Astronomy, Astrophysics, and Cosmology (B - Solar System), Springer, Pages 254–270. ISBN 978 3 540 88054 7., (2009).
2. Caretto, L.S.; Gosman, A.D.; Patankar, S.V.; Spalding, D.B.: Two calculation procedures for steady, three-dimensional flows with recirculation. edn, Vol., pp. Pages, Proc. Third Int. Conf. Numer. Methods Fluid Dyn., Paris, (1972).
3. Christensen, U.; Hofmann, A.W.: Segregation of subducted oceanic crust in the convecting mantle, *Journal of Geophysical Research* **99**, 19,867–19,884, (1994).
4. De Smet, J. H.; Van Den Berg, A. P.; Vlaar, N. J.: The evolution of continental roots in numerical thermo-chemical mantle convection models including differentiation by partial melting. *Lithos*, **48**, 153–170, (1999).
5. Foley, C.N.; Wadhwa, M.; Borg, L.E.; Janney, P.E.; Hines, R.; Grove, T.L.: The early differentiation history of Mars from ^{182}W - ^{142}Nd isotope systematics in the SNC meteorites, *Geochim. Cosmochim. Acta*, **69**, 2005, 4557–4571, (2005).
6. Fraeman, A.; Korenaga, Y.: The influence of mantle melting on the evolution of Mars. *Icarus*, **210**, 43–57, (2010), doi:10.1016/j.icarus.2010.06.030.
7. Grasset, O.; Parmentier, E.M.: Thermal convection in a volumetrically heated, infinite Prandtl number fluid with strongly temperature-dependent viscosity: Implications for planetary thermal evolution. *J. Geophys. Res.*, **103**, 18,171–18,181, (1998).
8. Gerya, T.V.: *Introduction to Numerical Geodynamic Modelling*. Cambridge University Press, 345 pp., (2010).
9. Harder, H.; Hansen, U.: A finite-volume solution method for thermal convection and dynamo problems in spherical shells. *Geophysical Journal International*, **161**, 522–532, (1986).
10. Hauck, S.A.; Phillips, R.J.: Thermal and crustal evolution of Mars. *J. Geophys. Res.*, 2002, **107**, NO. E7, 10.1029/2001JE001801, (2007).
11. Hernlund, J. W.; Tackley, P.J.: Modeling mantle convection in the spherical annulus, *Phys. Earth Planet. Inter.*, **171** (1–4), 48–54. doi: 10.1016/j.pepi.2008.07.037, (2008).
12. Huettig, C.; Stemmer, K.: Finite volume discretization for dynamic viscosities on Voronoi grids. *Phys. Earth Planet. Interiors* (2008), doi: 10.1016/j.pepi.2008.07.007.
13. Huettig, C.; Stemmer, K.: The spiral grid: A new approach to discretize the sphere and its application to mantle convection. *Geochem. Geophys. Geosyst.*, **9**, (2008), Q02018, doi:10.1029/2007GC001581.
14. Huettig, C.: *Scaling Laws for Internally Heated Mantle Convection*, Ph. D. Thesis, (2009).
15. Karato, S.; Paterson, M.S.; Fitz Gerald, J.D.: Rheology of synthetic olivine aggregates: influence of grain size and water. *J. Geophys. Res.*, 8151–8176, **91**, (1986).
16. Katz, R. F.; Spiegelman, M.; Langmuir, C. H.: A new parameterization of hydrous mantle melting. *Geochemistry, Geophysics, Geosystems* **4** (9), art. no. 1073, (2003).
17. Keller, T. and Tackley, P. J.: Towards self-consistent modelling of the Martian dichotomy: The influence of low-degree convection on crustal thickness distribution, *Icarus* **202** (2), 429–443 (2009).
18. Lee, D.C.; Halliday, A.N.: Core formation on Mars and differentiated asteroids, *Nature*, **388**, 6645, 854–857, (2004).
19. Maaløe, S.: The solidus of harzburgite to 3 GPa pressure: The compositions of primary abyssal tholeiite. *Mineralogy and Petrology* **81** (12), pp. 117 (2004).

20. Morschhauser, A.; Grott, M.; Breuer, D.: Crustal recycling, mantle dehydration, and the thermal evolution of Mars. *Icarus*, Volume **212(2)**:541–558, doi:10.1016/j.icarus.2010.12.028. (2011).
21. Papike, J.J.; Karner, J.M.; Shearer, C.K.; Burger, P.V.: Silicate mineralogy of martian meteorites. *Geochimica et Cosmochimica Acta* **73**, 7443–7485, (2009), doi:10.1016/j.gca.2009.09.008.
22. Patankar, S.V.: Numerical heat transfer and fluid flow. edn, Vol., pp. Pages, McGraw-Hill, New York (1980).
23. Plesa, A.-C.:Mantle Convection in a 2D Spherical Shell. Proceedings of the First International Conference on Advanced Communications and Computation (INFOCOMP 2011), October 23–29, 2011, Barcelona, Spain, pp. 167–172, Rckemann, C.-P., Christmann, W., Saini, S., & Pankowska, M. (eds.), ISBN: 978-1-61208-161-8. Retrieved on November 3, 2011, from http://www.thinkmind.org/download.php?articleid=infocomp_2011_2_10_10002, (2011).
24. Plesa, A.-C.; Tosi, N.; Huettig, C.: Thermo-chemical convection in planetary mantles: advection methods and magma ocean overturn simulations. *Integrated Information and Computing Systems for Natural, Spatial, and Social Sciences*, ed. Claus-Peter Rueckemann, (2012).
25. Roberts, J.H.; Zhong, S.: Degree-1 convection in the Martian mantle and the origin of the hemispheric dichotomy. *Journal of Geophysical Research E: Planets*, **111**, (2006).
26. Schubert, G.; Turcotte, D.L.; Olson, P.: Mantle convection in the Earth and planets. Cambridge University Press, (2001).
27. Schumacher, S.; Breuer, D.: Influence of a variable thermal conductivity on the thermochemical evolution of mars. *J. Geophys. Res.*, 111(E2), CiteID E02006, (2006).
28. Zhong, S.; McNamara, A.; Tan, E.; Moresi, L.; and Gurnis, M.: A Benchmark study on mantle convection in a 3-D spherical shell using CitcomS, *Geochemistry, Geophysics, Geosystems*, **9**, Q10017, doi:10.1029/2008GC002048, 32 pp., (2008).

A Forward Model of Mantle Convection with Evolving Continents and a Model of the Andean Subduction Orogen

Uwe Walzer, Roland Hendel, Christoph Köstler, Markus Müller, Jonas Kley, and Lothar Viereck-Götte

Abstract Some essential features of Andean orogenesis cannot be explained *only* by a dynamic *regional* model since there are essential influences across its vertical boundaries. A dynamic regional model of the Andes should be embedded in a 3-D spherical-shell model. Because of the energy distribution on the poloidal and toroidal parts of the creep velocity and because of geologically determined mass transport alongside the Andes, both models have to be three-dimensional. Furthermore, we developed a new viscosity profile of the mantle with very steep gradients at the lithospheric-asthenospheric boundary and at a depth of 410 and 660 km. Therefore, the challenges to the code Terra are now essentially larger. In the last 3 years we have resolved these problems in an international cooperation (see Sect. 2.2). Based on the new viscosity profile and on an improved Terra, we computed a new forward spherical-shell model (Walzer and Hendel, J Geophys Res submitted, 2012b). For this model, we derived also a new extended acoustic Grüneisen parameter, γ_{ax} , new profiles of the thermal expansivity, α , and of the specific heat, c_v , at constant volume as well as a solidus depending on both the pressure and the water abundance. These innovations are essential to incorporate a chemical-differentiation mechanism into the model. We arrived at rather realistic episodes of continental growth interrupted by magmatically quiet time spans distributed over the whole time axis. Nevertheless, the model shows a main magmatic event at the very beginning of the Earth's evolution. Papers on the improvement of Terra (Köstler et al. Comput Geosci submitted, 2012; Müller and Köstler, Int J Numer Methods Eng submitted, 2012) have been written. We conceived a regional model of the Andean orogenesis (Sect. 3.2.1) with the same new viscosity profile. We want to investigate why there is flat-slab subduction in some segments of the Andes and why deformation of the crust and volcanism migrate eastward. The evolution of

U. Walzer (✉) · R. Hendel · C. Köstler · M. Müller · J. Kley · L. Viereck-Götte
Institut für Geowissenschaften, Friedrich-Schiller-Universität, Humboldtstr. 11,
07743 Jena, Germany
e-mail: u.walzer@uni-jena.de

the abundances of incompatible elements indicate a cycle which was finished by a fast process, perhaps by a large-scale delamination of the lower plate, perhaps also by another type of delamination. In connection with another spherical-shell model (with prescribed plate boundaries), the regional model should numerically explain why a plateau-type orogen evolved at an oceanic-continental plate boundary.

1 Introduction

The papers which refer to our topic can be classified by seven subjects.

- (a) Geological description of dynamic problems of thermal evolution of the Earth, plate tectonics and chemical differentiation as well as of the Andean orogeny,
- (b) Models which are partially kinematic and partially dynamic. In this kind of models, essential features are prescribed in order to gain a large adaptation to geological and geophysical observations.
- (c) Geochemical models of growth and differentiation of continents which do not contain any dynamic modeling,
- (d) Dynamic models of the subduction process to understand the physical mechanism behind subduction,
- (e) Circulation models,
- (f) Fully dynamic models of subduction in a diamond (cf. Fig. 8), i.e. in a certain 3-D sector of the spherical shell, which represents the mantle, where this diamond is embedded into a realistic 3-D spherical-shell solution,
- (g) Fully dynamic forward models of spherical-shell convection with chemical differentiation and generation of continents.

We systematically described the papers of types (a)–(e) in [110]. Therefore, we will not repeat it here. Up to now, there is no paper of type (f). Some supplements will follow: To obtain a more profound analysis of the Andean orogeny, it is important to understand why a plateau-type orogen formed between a purely oceanic lithospheric plate (Nazca plate) and a continent (South America) [71]. The Andean mountain belt belongs to the non-collisional type. Kley and Monaldi [42,45,46] found a Cenozoic shortening of 250–350 km whereas Arriagada et al. [4] derived 400 km for the central Andes. In other areas of the Earth, however, the subduction zones at an oceanic-continental plate-boundary site have only little or no shortening and do not show any elevated plateaus. In most cases, the upper plate is characterized by backarc extension. Schellart and Rawlinson [82] discuss some hypotheses of which it is claimed that they explain this exceptional behavior of the Andes.

- In [66], it is proposed that the young age of the Nazca plate and low negative buoyancy cause this phenomenon.

- Climatic conditions, high friction and subduction erosion is thought to be the decisive factor [57, 58].
- Heuret and Lallemand [30] emphasize the eminent role of acceleration of the westward movement of the upper, South American plate. Sobolev et al. [84] find out that the accelerating westward movement of the South American plate is the most important factor for the Andean orogeny.

The third hypothesis generates the question of the mechanism which drives South America westward. The ridge push at the mid-Atlantic ridge and the slab pull at the Lesser Antilles and the Scotia arc have been proposed but it could be that these contributions are too small. Schellart and Rawlinson [82] remarkably mention the slab pull of the Nazca plate. However, we propose that a downwelling of the bulk convection beneath South America caused by the thermal screening of the thick continental lithosphere could play a role. If there is a large upwelling east of South America and a large downwelling of the bulk convection under South America we could understand why the bulk convection current would move the South American plate westward since, because of the thick continental lithosphere, South America is *not* decoupled from the bulk convection by the asthenosphere. Additionally, we observe a large upwelling of the bulk convection beneath the Pacific. So we can expect that another current of the bulk convection will go eastward, producing the arcs of the Lesser Antilles and Scotia. It is unclear if these speculative arguments are appropriate. However, they show not only the necessity of dynamic, numerical regional models but also that the regional model must be nestled among time dependent boundary conditions determined by a global dynamic model. Evidently, some *essential* features of the Andean orogenesis cannot be explained by an isolated regional model.

We discuss further *new* geological and geophysical papers in Sect. 3.2.1. We did that on purpose in order to substantiate why we used certain details in our development of two special (alternative) mechanisms for a regional model of the South American subduction zone and Andean orogenesis. Further *new* papers are mentioned in Sect. 3.2.2 in order to select an appropriate set of prescribed plate movements in the surrounding spherical-shell convection model which is necessary to embed the regional model.

2 A Spherical-Shell Forward Model and Other Results

Under Sects. 2.1–2.3, we describe what we have done using the system HP XC4000 of the Steinbuch Center for Computing in Karlsruhe. Under Sect. 2.4, other efforts of us which have some relation to the topic are reported. There is a direct relationship between the Sect. 2.1–2.3 and the running works which we describe in Sect. 3.2.1 in rather distinct outlines, i.e. we discuss our specific South American dynamical model which is embedded in a second spherical-shell model.

2.1 Spherical-Shell Model: Forward Model

We developed two spherical-shell convection models. The forward model is described here; another model with prescribed plate movements is outlined under Sect. 3.2.2. For both spherical-shell models as well as for the regional models of the Andes (Sect. 3.2.1), we newly derived, in some cases adopted new radial profiles of the relevant physical parameters of the mantle [101]. As a Grüneisen parameter, we calculated an extended acoustic gamma, γ_{ax} , using seismic observations. These observed values are the bulk modulus K , the shear modulus μ , dK/dP and $d\mu/dP$, where P is the pressure. This procedure has the advantage to be based on observable quantities without any assumptions on mineralogy. A further calculated property is the adiabatic gradient. We derived new profiles for the thermal expansivity, α , and the specific heat, c_v , at constant volume. For the chemical differentiation we used a simple melting criterion, $T > f_3 \cdot T_m$, where T is the temperature and T_m is the solidus. We estimated the lower-mantle solidus using values of [60] and of other authors continuing them by using our γ_{ax} . In the upper mantle and the transition layer, we took into consideration the water dependence of the solidus. Litsov [61] investigated the influence of different water concentrations on the solidus of peridotite and we took the derived depressions of the solidi into account in the computation of our convection-differentiation model of the mantle's evolution. So in our new model, the melting temperature is a function of time and position. The most important innovation of the new convection model concerns our newly derived viscosity profile which is based on solid-state physics and seismological results [100, 101]. This viscosity distribution has a resemblance to the viscosity model of [65] although its derivation is totally different. The full set of convection-differentiation equations has been solved using the improved code Terra (see Sect. 2.2). For each run, we obtained lots of parameters, some of which can be compared with observational quantities. For ages from 4,490 Ma to the present time, we received the curves of the laterally averaged heat flow density q_{ob} at the surface, the converted continent-tracer mass per Ma, the Urey number Ur , the Rayleigh number Ra , the Nusselt number Nu , the volumetrically averaged mean temperature T_{mean} of the mantle, showing a very realistic temperature drop when compared with Archean komatiite temperatures, the integrated mass of continents, the kinetic energy of the mantle flow E_{kin} , the radiogenic heat production Q_{bar} and the laterally averaged heat flow density q_{cmb} at the core-mantle boundary (CMB). Further results are the vector field of the creeping velocity and the temperature distribution for every time step of the Earth's evolution. Up to now, we varied the melting parameter f_3 and the thermal conductivity k . There are $f_3 - k$ clusters of realistic totalities of solutions. Figure 1 represents an episodic distribution of juvenile additions to the continents which is connected with the episodicity of orogenic epochs. Figure 2 shows the present-day distribution of continents of the same run. In this example, 42.2 % of the Earth's surface is covered by continents. The present-day surface heat flow of this run is 81.39 mW/m², comparable with the observed value of 90.185 mW/m² [12]. The present-day value of q_{cmb} of this run is 20.20 mW/m² which seems to be realistic, too.




Fig. 1 Juvenile additions to the sum of continental masses acc. to the *new* convection-differentiation model Run 498, [101]




Fig. 2 The distribution of continents (red), oceanic lithosphere (yellow) and oceanic plateaus (black dots) for the present time according to the *new* convection-differentiation model [101], Run 498, $r_n = 0.5$, $\sigma_y = 120$ MPa, continental percentage = 42.2%

We computed the shear viscosity, η , using

$$\eta(r, \theta, \phi, t) = 10^{r_n} \cdot \frac{\exp(c \frac{T_m}{T_{av}})}{\exp(c \frac{T_m}{T_{st}})} \cdot \eta_4(r) \cdot \exp \left[c_l \cdot T_m \left(\frac{1}{T} - \frac{1}{T_{av}} \right) \right] \quad (1)$$

where r is the distance from the Earth's mass center, θ the colatitude, ϕ the longitude, t the time and r_n the viscosity-level parameter. The quantity r_n has been varied to shift the viscosity profile to the left or to the right. So we generated different time-averaged Rayleigh numbers, Ra , varying from run to run. T_m is the newly invented melting temperature [101] which additionally depends on water abundance. So T_m is a function of time, too. T_{av} is the laterally averaged temperature, T_{st} the initial temperature profile. The quantity η_4 denotes the new viscosity profile [101] for the initial temperature and $r_n = 0$. For MgSiO_3 perovskite we should insert $c = 14$, for MgO wüstite $c = 10$ according to [113]. Therefore, the lower-mantle value of c should be somewhere between these two values. For numerical reasons, however, we are able to use only $c = 7$. In the lateral-variability term we use $c_l = 1$. The temperature is denoted by T .




Fig. 3 We compare different runs which are distinguished only by the viscoplastic yield stress, σ_y , and the viscosity-level parameter r_n . The *symbols* represent the magnitude of the difference, d_c^* , between the computed and observed present-day surface percentage of continents, expressed in percent

Figure 3 shows the difference between computed and observed present-day continental surface percentage. The central run with $r_n = 0.5$ (corresponding to $Ra \approx 10^8$) and $\sigma_y = 120$ MPa has the minimum difference between theory and observation, namely 1.85%. This run shows also other optimal results, e.g. regarding the temporal distribution of the magmatic-activity episodes (cf. Fig. 1) and the (spectral) distribution of continents (cf. Fig. 2). Further variations of the parameters refer to the factor f_3 of the melting criterion and to the thermal conductivity, k , keeping r_n and σ_y fixed at the mentioned two values. At a first glance on Fig. 4, we could have the impression that there is a trade-off between k and f_3 since optimal solutions (black circles with an outer ring) cluster along a certain curve. But all the other observable quantities in other $f_3 - k$ plots show that only the three optimum solutions with $k = 5.0$ W/(m·K) in the upper right corner of Fig. 4 are realistic. Figure 5 demonstrates, e.g., that only thermal-conductivity values around $k = 5.0$ W/(m·K) lead to solutions which are satisfactory for *all* observables. This value is acceptable also from the physical point of view [115]. Therefore, we varied f_3 in small steps from 1.000 downward, keeping $k = 5.0$ W/(m·K). As expected from Figs. 4 and 5 and similar $f_3 - k$ plots we obtained very realistic solutions down to $f_3 = 0.985$. However, already $f_3 = 0.983$ and $f_3 = 0.981$ generates less convincing results. Figure 6, e.g., shows the corresponding present-day distribution of continents.




Fig. 4 Keeping $r_n = 0.5$ and $\sigma_y = 120$ MPa fixed, we vary the melting-criterion factor, f_3 , and the thermal conductivity, k . The symbols represent intervals of the theoretical present-day continental surface, p_c , in percent. Note for comparison that 40.35% of the real present-day Earth is covered by continents and epicontinental seas. The optimal run, shown in Figs. 1 and 2, corresponds to $k = 5.0$ W/(m·K) and $f_3 = 0.995$




Fig. 5 A $f_3 - k$ plot for constant $r_n = 0.5$ and $\sigma_y = 120$ MPa. The laterally averaged surface heat flow is denoted by qob . The temporal average of qob over the last 900 Ma is called qob_9 and is expressed in mW/m^2




Fig. 6 Less realistic distributions of continents (*red*), oceanic lithosphere (*yellow*) and oceanic plateaus (*black dots*) for the present time. $r_n = 0.5$, $\sigma_y = 120$ MPa, $k = 5.0$ W/(m·K), $f_3 = 0.983$ and $f_3 = 0.981$, respectively

2.2 Numerical Improvements

Two submitted publications of our group [55,69] deal with *numerical improvements*. Their conclusions will not be discussed here. However, their results are very important for the realization of our Andean model (Sect. 3.2.1).

A big part of our efforts was concentrated on the creation of an essentially improved Terra code, which is necessary to resolve the numerical problems with the newly derived viscosity profiles. We equally apply this new viscosity profile [101] to a spherical-shell mantle convection model [100, 101] as well as to the regional model of the South American subduction slabs (Sect. 3.2.1). In 2009, the Terra developers of the universities Munich (H.-P. Bunge, M. Mohr), Cardiff (H. Davies), Leeds (P. Bollada), Jena (M. Müller, C. Köstler) and of the Imperial College London (R. Davies) and others started a close cooperation in the further development of the Terra code and intensified the collaboration with J. Baumgardner (San Diego, USA), the inventor of Terra. At a first joint meeting in Munich in 2009, the group decided

to set up a community svn-repository for further code development, supplemented by trac, a web-based project management and bug-tracking tool, and automated compile/test cycles using BuildBot. From then on the group worked on a common code base using automated tests for every revision of the code. There have been three successive joint meetings in Cardiff in 2010, in Jena in 2011 (see also <http://www.igw.uni-jena.de/geodyn/terra2011.html>) and in Munich in 2012. The progress being made since the first meeting includes the following items.

- Enhancement of the code to increase global resolution and maximum number of MPI processes.
- Further development and integration of the Ruby test framework [68] into the automated BuildBot tests.
- Implementation of a finite-element inf-sup stabilization using pressure-polynomial projections proposed in [15].
- Development and implementation of an efficient preconditioner for the variable-viscosity Stokes system [54].
- Refinement of the Pressure Correction algorithm [54], giving more robust convergence.
- Restructuring of the code to use language features of Fortran95 and Fortran2003 where possible.
- Integration of automated code documentation using doxygen.
- Integration of VTK-support and automated visualization.
- Significant improvements in the formulation of the free-slip boundary condition on the spherical surface.

Continued effort is spent on several numerical and technical topics as well as on including more realistic physical models. In the following, some important details are given.

- In [97, 100, 101], two pressure- and temperature-dependent viscosity profiles of the Earth's mantle are developed and used for the computation of two mantle convection models with chemical differentiation of oceanic plateaus and generation of continents. The new viscosity model includes very strong viscosity gradients at lithosphere-asthenosphere boundary and at 410- and 660-km phase boundaries. Therefore, it is very important to make the code Terra fit for such a strong challenge. Regarding a physically consistent variable viscosity momentum operator, J. Baumgardner, P. Bollada and C. Köstler figured out in which way the code has to be changed to apply a physically consistent A-operator using cell-averaged viscosities. The most significant code change is the switch from nodal based to triangle based operator parts on the sphere. The viscosity-weighted summation over triangular integrals is then done in the application of the operator. We expect that the cost for applying Au will be doubled but a consistent formulation on all grid-levels could pay off for this, especially if we get a better convergence rate of the multigrid algorithm. The implementation of the triangle-based operator formulation is now under way.

- Exporting Terra-operators in sparse matrix format: When exporting the FE-matrices for the whole grid, they can be analyzed and PETSc or other parallel solver packages can be applied to it. This is intended to provide flexibility and to ensure reliability of future code changes.
- Multigrid(MG)-implementation: As mentioned in [90], the current MG-implementation in Terra does not fulfill the expectations raised by the performance of a 2-D Cartesian version of Terra, documented in [114]. In [54] it is also identified to be the worst performing part of the iterative solver in Terra. It is only poorly analyzed, and it is not satisfactory documented. M. Mohr is going to analyze and document the currently used matrix-dependent transfer multigrid in detail. The MG-implementation is also to be changed to using cell-based viscosity averages.
- Free-slip boundary condition and propagator matrix benchmark tests: P. Bollada and R. Davies showed that adding boundary terms to the right hand side of the momentum equation reduces some sort of errors while other kinds of errors still exist. They will continue to figure out the exact cause of that behavior and work on fixing this. With direct access to the radial velocity component, the local spherical coordinate system offers a way to straightforward implement the free-slip boundary condition. J. Baumgardner implemented this in a local copy of Terra, and it is ready to be used. The group has agreed to create a repository branch to continue working on that. If successful, the local spherical coordinate system-version will be merged into the trunk after some months of testing.
- Adding Ruby tests: The Ruby test framework is ready to be used extensively in testing Terra's subroutines as individually as possible. It can also be used for debugging by application of subroutines to predefined scalar and vector fields. Still the test coverage of the Terra by Ruby tests is very low and needs to be extended.

2.3 *Andean Model*

A large expenditure of time of U. Walzer in the last 3 years was the analysis and synopsis of geophysics, geology and geochemistry of the Andean orogenies. This has been done in close cooperation with J. Kley and L. Viereck-Götte. Some geophysical and geological considerations are outlined in Sect. 3.2.1.

2.4 *Long-Term Related Works of Our Group*

Kley investigated the structural geology of some areas of the central Andes [40] and gave a regional structural analysis and kinematic restoration [41,47]. He participated in an effort to extend quantitative structural analysis to a transect right across the backarc area. First steps were also taken towards constraining the evolution of

strain rates over time [49], leading to the suggestion that the continental strain rates had increased during the Andean orogeny. Kley began to extend kinematic analysis to the entire orogen, employing serial balanced sections and map view restoration techniques [44]. A map view kinematic model of the central Andes [45] formed the basis for comparison of geologically derived, orogen-scale kinematics with GPS and seismologic data [31, 32, 52]. It could be shown that the present-day strain field from satellite geodetic data closely matches the strain field for the last 10 Ma as inferred from geologic evidence. Additional evidence was presented that the strain rate in the South American plate becomes independent of plate convergence rate in the later stages. Using the map-view strain field as input for a numerical model, an attempt was also made to constrain crustal thickness evolution and the flux of crustal material during the Andean orogeny [33]. The studies on variations in structural style along the Andes [42] also triggered a second line of research dealing with the influence of inherited lithospheric heterogeneities on the spatial strain distribution. One important factor here is the widespread occurrence of Mesozoic rift basins [50] that were partially inverted as the thrust front migrated across them. Several case studies from a particularly well-exposed rift system in northern Argentina helped to clarify the importance of fault reactivation and stratigraphic discontinuities in conditioning the mechanical behavior of the upper crust in contraction [43, 46, 51, 67]. The results of these studies were incorporated in [48, 71].

Walzer and cooperators worked on convection-fractionation problems. The thermal evolution of the mantle and the chemical evolution of the principal geochemical reservoirs have been modeled simultaneously by a fractionation mechanism plus 2D-FD thermal convection [93–96]. Oceanic plateaus, enriched in incompatible elements, develop leaving behind the depleted parts of the mantle. The resulting inhomogeneous heat-source distribution generates a first feed-back mechanism. The lateral movability of the growing continents causes a second feed-back mechanism [95]. Effects of the viscosity stratification on convection and thermal evolution of a 3D spherical-shell model have been investigated and a viscosity profile of the mantle was developed [103, 104]. The paper [106] presents 2D and 3D thermochemical models of mantle evolution where a self-consistent theory is included using the Helmholtz free energy, the Ullmann-Pan'kov equation of state, the free-volume Grüneisen parameter and Gilvarry's formulation of Lindemann's law. In order to obtain the relative variations of the radial factor of the shear viscosity, the pressure, P , the bulk modulus, K , and $\partial K / \partial P$ from the seismic model PREM have been used. The publications [97, 104, 105, 107] present models of self-consistent generation of stable, but time-dependent plate tectonics on a 3D spherical shell. Different types of solutions have been found for different models by systematic variation of parameters [97, 102, 104, 108, 109]. Stirring effects are investigated in [22]. A 3D spherical-shell mantle convection and evolution model with growing continents [97, 100, 101, 108, 109] has been developed. The evolution model equations guarantee conservation of mass, momentum, energy, angular momentum, and of four sums of the numbers of atoms of the pairs ^{238}U - ^{206}Pb , ^{235}U - ^{207}Pb , ^{232}Th - ^{208}Pb , and ^{40}K - ^{40}Ar . The pressure- and temperature-dependent viscosity is supplemented by a viscoplastic yield stress.

The lithospheric viscosity is partly imposed to mimic the viscosity increase by chemical layering and devolatilization. Stochastic effects [98] are shown to exist especially in the chemical differentiation. Although the convective flow patterns and the chemical differentiation of the oceanic plateaus are coupled, the evolution of the time-dependent Rayleigh number, Ra_t , is relatively well predictable from run to run and the stochastic parts of the $Ra_t(t)$ – curves are small [97].

Viereck-Götte and cooperators worked not only about the genesis of the Jurassic Ferrar large igneous province in Antarctica but also about the connection between Antarctica, South America and Africa. Plateau forming lavas in the Karoo province of South Africa and in the Ferrar province in Antarctica were emplaced synchronously at about 180 Ma. They thus seem to originate from the same dynamic mantle process. However, both are distinguished by their isotopic characteristics with respect to the Rb/Sr- and Sm/Nd systems: while the Karoo magmas show mantle values, the Ferrar magmas exhibit enriched upper crustal values. However, the boundary between both Jurassic igneous provinces is marked by a large transpressional shear zone (Heimefront SZ) of Pan-African age (600–500 Ma), crossing Dronning Maud Land (S-African side of Antarctica) on the continent side of the Grunehogna craton, a fragment of the W-Gondwana Transvaal craton. This shear zone is interpreted as a reactivated suture of Grenvillean age (1.1 Ga). If Jurassic magmas on either side of this boundary are isotopically different, it must be concluded that

1. This is not a signature in a lower mantle plume.
2. This must be a signature within the subcontinental lithospheric mantle.
3. This signature must be older than Grenvillean in age.
4. It must be introduced into a paleo supra-subduction mantle wedge by subduction processes if it is a crustal isotopic signature.

Their studies concentrated on the timing of the initiation of the Ferrar as a large igneous province as well as on the physicochemical characterization of the melt source region conditions during melt differentiation. Studying the intrusive, extrusive and volcanoclastic rocks:

- They reconstructed the initiation of a large igneous province to have occurred in several steps of melt pulses within 5 Ma (189–183 Ma ago).
- They showed initiation to have started with large-volume shallow-level intrusions of low-Ti andesitic melts into wet fluvial sediments (Triassic/Jurassic) associated with diatreme-forming Taalian-type eruptions.
- The eruptions followed by basaltic andesites in small volume eruptions of partly pillowed lavas from local eruptive centers prior to large volume plateau forming lava extrusions from feeder dikes followed by a final pulse of a large volume andesites high in Ti that had differentiated from a common primary melt under lower pressure, oxygen fugacity and water activity.

All melts belong to the tholeiitic differentiation series, however with orthopyroxene instead of olivine as early fractionating mafic solidus phase. REE, Sr-Nd-isotope and PGE characteristics indicate generation of the primary melts within the

spinel-lherzolite zone of an isotopically enriched and sulfur-undersaturated subcontinental lithospheric mantle. Crustal silica enrichment of this source due to an overprint in a supra-subduction environment had already been concluded from Re-Os isotopy. Our Sr-isotope data in plagioclase phenocrysts, however, exhibit decreasing radiogenic character in subsequent melt pulses indicating additional assimilation of crust to decreasing extents during ascent and differentiation. Due to a Cretaceous thermal event, $^{40}\text{Ar}/^{39}\text{Ar}$ -ages (ranging from 235 to 90 Ma) and S-isotope characteristics are heavily disturbed, only a few samples exhibit the relict primary $\delta^{34}\text{S}$ -value of -19.5 .

3 The Andean Subduction Model and Its Embedding

3.1 *Some Fundamental Remarks*

Our modeling of the Andean slab should be compatible with our basic assumption that plate tectonics is an integral part of the convection of the entire mantle. Therefore, we should aim at a regional model which is embedded into a realistic convective spherical-shell model. It is possible that the mantle convection is characterized by top-down control. But it is also possible that the larger part of the mantle mass essentially determines the movements at the surface since the most important share of the primordial energy is stored there and also the principal share of the radiogenic energy is released there. Here we refer to the absolute value, not to the density of these quantities. On the other hand, for reasons of energy, it is evident that the Earth's core cannot play a prominent part in controlling mantle convection. It is exactly the reverse. The mantle convection determines the boundary conditions of the hydromagnetic convection in the outer core. For example, in time spans of high activity of mantle convection, the latter one generates lateral temperature differences in D'' which reduce the number of magnetic reversals or even make them impossible because of the anti-dynamo theorems.

Furthermore, it is well known at the present time that the oceanic lithosphere consists of three (or more) layers which are chemically different and that its lower boundary is characterized by a sharp viscosity jump. At the same time the oceanic lithosphere is also a thermal boundary layer. So, we do not intend to use the old simplified approach that the existence of an oceanic lithosphere can *exclusively* be explained by a thermal boundary layer. In this case we should expect a gradual decrease of the viscosity at the lower boundary. Furthermore, we conclude that the principal part of the buoyancy of the slab heavily depends on chemistry and that phase transitions, especially the basalt-eclogite transition, play an eminent role.

So, we want to combine a global 3-D spherical-shell convection model with a 3-D regional convection model of South America and the surrounding plates and a model of the orogeny of the Cordilleras. The Andes are an ideal test case for various reasons:

- They are very large, thus keeping the inevitable problems of different scales at a minimum.
- They are active and well-studied. A wealth of information constrains the processes of orogenesis.
- They have a simple large-scale geometry. The plate margin is gently curved and was even straighter in the past. Convergence is orthogonal to the mean trend of the margin. The age structure of the Nazca plate near the contact with South America is symmetric, with the oceanic lithosphere oldest in the center and younging to both sides.
- Their plate kinematic framework has remained nearly unchanged for the last 50 Ma. On the other hand, there was a prominent switch in the mode of subduction earlier, from presumably steep with backarc extension to low-angle with backarc shortening. In the Andes, both end-member settings can therefore be studied in the same place.
- There is a clear compositional and rheological contrast between the two converging plates. It is unlikely that large volumes of material are transferred across the plate contact, in strong contrast to continental collision zones.
- The Andean substrate has a simple geologic history. Much of the South American plate was in place before subduction started some 200 Ma ago. Except for the northern Andes, no terranes were accreted after the Paleozoic.

We aim at a numerical, nearly purely dynamic model with a minimum number of restrictions and additional assumptions. This model is to explain the physical mechanism of the essential features of Andean orogeny.

3.2 A Model of the Andean Subduction

Outlines of the Model

We do not develop an exclusively regional model of the Andean orogenesis since, in this case, the temporally varying boundary conditions are unknown. Therefore, it is often assumed, for reasons of simplicity, that there are no effects from outside of the regional computational domain. However, some changes in the arc volcanism and in the tectonic shortening of the Andes suggest a connection with the 30-Ma Africa-Eurasia collision. Therefore, we embed a regional 3D model into a 3D spherical-shell model. So, we solve the balance equations of momentum, energy and mass in the spherical-shell model using somewhat larger time steps on a coarser whole-mantle grid, coarser than in the regional model. The values of creeping velocity, temperature and pressure, determined in that way and lying at the boundaries of the *regional* computational domain, then serve as boundary conditions for a computation with smaller time steps for which the balance equations are solved in the regional computational domain.

3.2.1 The Regional Model

Meanwhile, we have developed a numerical tool for several classes of regional models, and it was indeed a great effort. *So, we are in no way restricted to the model developed in the following.* We could as well compare the effects of different geological proposals. To design the regional model, we developed not only extensive numerical improvements (see Sect. 2.2) in the Terra code but studied also the latest geophysical, geological and geochemical results taking them into consideration in the draft of a physically reliable mechanism which is not only geodynamically probable but also numerically feasible (U. Walzer, J. Kley, L. Viereck-Götte). There are several geologically descriptive model proposals. We mention only [13, 28, 29, 37, 71, 76]. Our designed computable model system ought to solve the following problems.

- Why do we observe today in some segments of the Andes flat subduction and magmatic lull {Bucaramanga, Peruvian (2–15°S) and Pampean (27–33°S) flat slab [75]}, but in other segments normal subduction with dip angles between 30° and 40°? This is amazing since the westward velocity of the South American plate and the eastward velocity of the Nazca plate do not *essentially* vary alongshore.
- Why do the flat-slab segments migrate in Cenozoic times along-side the Andes [76]? Ramos and Folguera [77] found an almost continuous belt of flat slabs which migrate southward.
- Why is the present-time volcanism restricted to segments with a 30–40° dipping slab [1, 36, 48, 83]?
- Why does the deformation essentially start in the West and migrate to and finish in the East [71]? We derived some estimations on some other hypotheses. After that we consider the assumption as most promising to answer the mentioned questions by the assumption that oceanic plateaus and aseismic ridges are carried by the conveyor belt, i.e. by the Nazca plate. The plateaus and ridges generate additional positive buoyancy [27]. So the hinge of the slab migrates to the East until the volcanism totally vanishes. However, if we scrutinize a good geological map of South America we notice that the relation between flat-slab segments and ridges are not simple. We can assign the Pampean flat slab to the Juan Fernandez Rigde [2]. Even the southward migration of the Pampean flat-slab zone can be explained by form and movement of the Juan Fernandez Rigde [37]. The southern part of the Peruvian flat slab can be referred to the Nazca Rigde. For the northern part of the Peruvian flat slab we have to introduce the hypothesis of an immersed Inca Plateau [28]. However, east of the Carnegie Rigde there is an abundant volcanism in Ecuador. Michaud et al. [64] show a detailed reconstruction of the eastward movement of the Carnegie Rigde which extends at least 60 km below the South American plate with a continuous plunging slab down to a depth of 200 km. The adakitic signal is proposed to be ridge-induced. In the case of the Iquique rigde we have to assume that there is no eastern continuation of it. Pindell

and Kennan [74] show that the flat-slab area in northern Colombia and Venezuela might be considerably larger than the area assumed by Ramos [76].

- *Why should our mechanism of the generation of the Andes be three-dimensional?* Several observations suggest the idea that essential features of the Andean orogenesis cannot be explained by 2-D dynamical models:
 - (a) Hindle et al. [33] conclude that a mass balance which is restricted to a cross-section through the Andes leads into contradictions. They show that there has to exist an essential mass transport alongside the Andes. In particular, the displacement of material toward the axis of the bend in the central Andes leads to a significant crustal thickening. This cannot be explained by a two-dimensional model, neither kinematically nor dynamically.
 - (b) Anderson et al. [3], Kneller and van Keken [53] and Barnes and Ehlers [5] show and discuss, for the southern Andean subduction zones, trench-parallel high seismic shear velocities near to the trench and an abrupt transition to trench-perpendicular high seismic shear velocities in the back arc. The Brazilian subcrustal lithosphere beneath the eastern Cordillera, the Interandean and the Subandes are East-West fast whereas the shear-wave velocity under the Altiplano and Puna has maximum values in the North-South direction. This is a hint that a significant three-dimensional flow might be involved in the mechanism.
 - (c) *The toroidal component.* A pertinent argument for the necessity of 3-D models results from the following considerations and calculations. Already Gable et al. [20] and O'Connell et al. [70] showed the relevance of the toroidal-poloidal partitioning for lithospheric plate motions. The lateral subducting slab movement induces slab-parallel flows and a rollback-generated flow around the slab [81]. Stegman et al. [86] demonstrated that in the case of non-vanishing rollback of the subduction slab for some typical cases, 69 % of the energy of the negative buoyancy of the slab is converted into the toroidal component of the rollback-induced flow whereas 18 % are consumed for the weakening of the plate. These numbers show the importance of a 3-D modeling in a striking way. Only in the very beginning of the Andean-specific modeling we prescribe the velocity of the migrating subduction hinges as a function of time according to [71] in order not to overburden the model. But the other degrees of freedom of the flexible slab should be determined by the differential equations of the model. We vary the lateral extend of the individual slab between 200 and 5,000 km. That is, in a first type of numerical experiments we introduce the individual lobes shown by the distribution of seismicity. In a second numerical experiment we assume an undivided slab for the whole South American continent (Fig. 7), at best with a disconnection at the Chile Rise. For all versions we cut out a spherical diamond (Fig. 8) from the newly improved and inf-sup stable Terra code which is able to solve the convection differential equations in a spherical-shell mantle. Essential parts of the South American plate and the Nazca plate fit into this spherical diamond. The vertical boundary conditions are




Fig. 7 The geometric starting configuration in our second Andean regional model (Taken from [92])




Fig. 8 The diamond (red), embedded in the global grid, which will be dyadically refined to the model resolution (Modified from [6, Fig. 1]). Of course, we use a further refinement of the grid

iteratively taken from a whole-mantle convection model. Cf. Sect. 3.2.2. In contrast with [86], we do not neglect the energy equation. It is evident that the subduction mechanism is possible only assuming a low-viscosity asthenosphere [18] which is less dense than the average oceanic lithosphere. Furthermore, the rheology may not be purely viscous. For this purpose, we [97] prefer a viscoplastic yield stress. The lithospheric-asthenospheric boundary for the viscosity is relatively sharp, also for the oceanic lithosphere. This is a special challenge for our code.

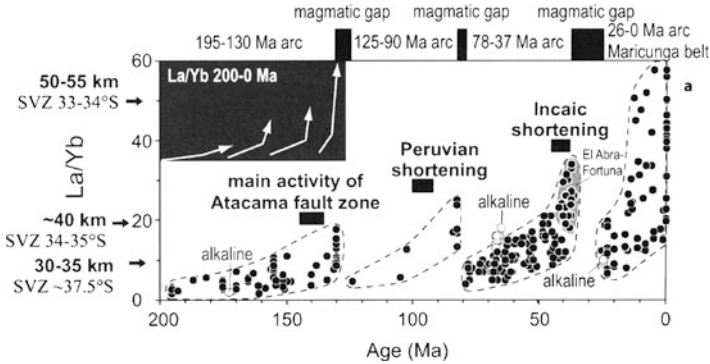


Fig. 9 The La/Yb ratio as a function of time for the igneous rocks of the north Chilean arc (21–26°S) acc. to Haschke et al. [29]. Note that *after* the flat-slab stage, the La/Yb *suddenly* decreases to a low starting value. The periods of orogenetic activity are *before* these drops

- *By the model, it should be possible to understand essential geochemical observations relevant for South America.* The difficult problem of delamination is very probably connected with this issue. It is necessary to clarify what kind of delamination is dominant. Figure 9 shows the gradual increase of the flow of incompatible elements. After each flat-slab period, this rise is abruptly interrupted. Such kind of plots exist not only for the mass ratio La/Yb but also for $^{87}\text{Sr}/^{86}\text{Sr}$ and $^{144}\text{Nd}/^{143}\text{Nd}$. These curves describe a slow growth of the abundances of elements with large ionic radii within one cycle since ^{87}Sr is the daughter isotope of ^{87}Rb , ^{143}Nd is the daughter isotope of ^{147}Sm and the other mentioned isotopes are stable. The fluctuations of the whole-rock initial ϵ_{Nd} of the central Andean arc by DeCelles et al. [13] are evidently related to Fig. 9. DeCelles et al. [13] emphasize that the cyclical changes of the isotopic compositions of arc magmas cannot be explained by changes in the convergence rates. They and also we expect that these changes have to be explained by episodic gravitational foundering. There are two principal possibilities.

(A) According to [13], below arc and hinterland, i.e. below the western parts of the South American continent, the eclogitization of the thickening lower continental crust and of the lithospheric mantle causes a density increase and therefore a delamination so that these units sink into the mantle wedge driven by their own weight. Carlson et al. [10] discuss the possibility that the continental mantle above the wedge of the mantle overlying a subducting oceanic plate can become unstable. The detachment from the overlying continental crust can cause major orogenetic episodes. Davidson and Arculus [11] propose a delamination of the cumulate layers below the seismological Moho back into the mantle of the sub-arc wedge. The two-dimensional numerical model by Sobolev et al. [84, 85] is compatible with the geological models mentioned under a) and belongs to b)-type of papers mentioned in the Introduction. They use a viscoelastic rheology supplemented by Mohr-




Fig. 10 A calculated marble-cake mantle model acc. to Walzer and Hendel [99]. This is an equatorial section showing the present-time state of the chemical evolution of incompatible elements of the Earth's mantle. We use a modernized reservoir theory. The depleted MORB mantle (DMM) and a mantle which is rich in incompatible elements, yet, are strongly intermixed. Strongly depleted parts of the mantle which include more than 50 % DMM are represented by *yellow areas*. Relatively rich mantle parts with less than 50 % DMM are *orange-colored*. In general, the *yellow-orange* boundary does *not* correspond to a discontinuity of the abundances of U,Th,K, etc. The cross sections through the continents are *red*. *Black dots* represent the oceanic plateaus. The yield stress is 125 MPa, the viscosity-level parameter is -0.50

Coulomb plasticity for the layered lithosphere. The drift of the overriding plate and the pulling of the slab is prescribed by the velocities at the boundaries of the 2-D model area and it is not calculated by solution of the balance equations though. What drives Andean orogeny? Sobolev et al. [84, 85] answer this question by numerical experiments using their 2-D model and varying only one influence parameter each. They conclude that the major factor is the westward drift of the South American plate. Paragraph (A) outlines only *one* way of thinking which we intend to test.

- (B) Here we propose a second hypothesis to understand the mentioned geochemical observations of [29] and of [13, Fig. 5]. This hypothesis is patronized by the idea of a geochemical marble-cake mantle [97, 99] (see Fig. 10) but composed of irregularly formed parts of a depleted MORB mantle (DMM) with 80–180 ppm H_2O and 50 ppm C and another, richer reservoir with 550–1,900 ppm H_2O and 900–3,700 ppm C [23, 111, 112]. The mid-oceanic ridge basalt is denoted by MORB. It is possible that the reservoirs are

intermixed in smaller quantities so that there is no sharp boundary between the different parts of the mantle [34, 89]. DMM dominates in the smaller depths of the mantle. The deeper the slab dives into the mantle, the higher is the probability to touch regions with a high abundance of ^3He and incompatible elements. Pilz [73] investigated $^3\text{He}/^4\text{He}$ ratios in the Puna plateau and at the volcano Tuzgle. He found that these $^3\text{He}/^4\text{He}$ values are higher than in the western Cordillera and in the Salta Basin east of it. Pilz and also we conclude that a provenance by degassing of the slab is not feasible since the mantle's ^3He is primordial. The ^3He of the atmosphere cannot be subducted in appreciable quantities.

There are different, but related models of chemical layering inside the subducting oceanic lithosphere [21, 80]. There is a water-rich subduction channel above the slab. Because of the sediment dehydration and the oceanic crustal dehydration it is not entirely clear how deep this hydrous channel is extended. Not only the $^3\text{He}/^4\text{He}$ ratio of the Puna plateau [73] but also La/Yb, $^{87}\text{Sr}/^{86}\text{Sr}$ and $^{144}\text{Nd}/^{143}\text{Nd}$ increase from an age of $\tau = 4.5$ Ma until the present time [29]. There are different explanations for this phenomenon. One of them would be a rise in an antiparallel direction in hot fingers immediately or in some distance above the subducting slab. Such a suggestion has been offered by Tamura et al. [91] for the NE Japanese arc. Marsh [63] proposed a similar idea of a hydrothermal flow field, in this case immediately above the upper surface of the down-going slab. In this way he explains also the sharp line of volcanoes or the volcanic front which can be observed in the Andes. Pilz [73] concludes from seismic observations that such hot fingers are also in the wedge of the southern central Andes and that these fingers are near the surface of the slab. So we want to develop a numerical model with fluid pathways of hydrothermal fluids, described by a particle approach, not very far from the slab but with an antiparallel flow direction. This idea is corroborated also by Furukawa [19].

We are working about the numerical problem to cut out a spherical diamond, out of the dynamical 3-D spherical-shell model, and to determine the temporally changing boundary conditions at the vertical side walls of the 3-D diamond by the solution of the convection in a spherical shell, but now with prescribed plates at the surface of the sphere. The Nazca plate moves with a (in the first approach) prescribed angular velocity through the margin into the computational domain and dives under the South American continent because of a Rayleigh-Taylor instability which is induced *mainly* by the transition to eclogite. The oceanic plate has a given sandwich structure of zero-pressure densities, ρ_0 , [21] and viscosities, η_0 . The slab is supposed to be freely movable at the surface and floating inside the mantle. It is well-known that it is difficult to detach a spherical-shell plate from the surface of the shell into the mantle in a slab-like manner [8]. As in [97], we want to introduce a viscoplastic yield stress at the near-surface lithospheric region and compare the creeping viscosity, η_c , at each position and time with the plastic viscosity, η_p , and use the minimum value in the model. This procedure is sensible from the physical point of view. If a piece of the Nazca plate which is thickened by an oceanic plateau

or a passive ridge approaches to South America then it pushes the hinge back under the South American plate because of its positive buoyancy. Therefore the volcanic front migrates to the East. The residual slab will become steeper and its tip will touch deeper and deeper parts of the mantle. Therefore, more and more atoms of incompatible elements will rise using the fluid pathways explaining the rising parts of the curve of Fig. 9. We are able to describe the movements of the hydrothermal fluid by a tracer modulus. These particles, describing the migration velocities of higher abundances of incompatible elements, ^3He , metals like Cu, Ag, etc., are *not* carried along with the creeping rock. They can move along a surface antiparallel to the slab which represents a thin layer. A serpentinized layer-like part of the mantle wedge just above the subducting oceanic crust could serve in reality as such a thin layer [25]. The slab movement is (in this model as in reality) an integral part of the solid-state mantle convection. The mentioned quantities ρ_0 and η_0 of the slab will be additionally described by other tracers which, in this case, will be entrained and carried along by the solid-state creep down to the interior of the mantle. The generation and eastward migration of major Andean ore deposits [29,38,39,79], the eastward migration of the volcanic front, the eastward movement of the deformation ages across the southern central Andes (21°S) [17,71] as well as the eastward migration of deformation in the Interandean and Subandean [16,41] can be dynamically modeled by our approach. When the subducting movement of the down-hanging part of the slab (and the flat-slab part behind) is accelerated, then an elevated activity of the Atacama fault zone, the Peruvian shortening, the Incaic shortening etc. are induced. The sudden interruption of the supply of highly incompatible elements (Fig. 9) must be induced by a radical event. Using the analogous, geologically describing models of [59, Fig. 4] and [35, Fig. 7], it cannot be concluded that many unusual features of the Permian-Jurassic South China fold belt can be explained by shallow subduction with an extensive final foundering of the lower plate combined with a roll-back of a small remnant subducting slab in the Mid-Jurassic. Humphreys [35] concludes for the Mid-Tertiary that an extensive sinking of the flat Farallon slab occurred which caused an uplifting of the continent covering a large area of the western North America. We propose an explanation by an extensive generalized eclogitization of the flat-slab part of the oceanic plate and an *extensive* delamination due to a Rayleigh-Taylor instability. The latter process can be simulated for the former and present South American flat slabs using a particle approach. The extensive tear-off would entirely interrupt the supplies of incompatible elements, ^3He , etc. since the tracer transport path near the upper surface of the low-hanging part of the slab and the flat part of the slab is ripped.

At first glance, (A) and (B) seem to be competing models, but in reality they do not entirely exclude one another. It is evident that the two roughly sketched numerical regional models are very ambitious. Therefore, we could report here only on the numerical and mathematical results of the preparatory efforts, the further development of the code and our geodynamical conception. However, a spherical-shell model is finished (cf. 2.1).

3.2.2 Spherical-Shell Model: Prescribed Angular Velocities of the Lithospheric Plates

To define the time-dependent boundary conditions of the vertical side walls of the spherical diamond containing South America and the Nazca plate, we introduce a spherical-shell convection model with the same radial profiles of the relevant physical parameters as in [100, 101] but with prescribed angular velocities of the plates for the last 200 Ma. This desirable time span is essentially determined by the available observational data, e.g. by Fig. 9. But the most plate-motion models do not go back so far. However, the main difficulty refers to another item. Even though we know all velocities of neighboring plates between each other, we do not know the “absolute” velocities of the plates relative to the highly viscous mid part of the lower mantle. But the net rotation of the averaged lithosphere is important for the kinematic analysis and the dynamic modeling of the slabs, especially the slabs at the margin of South America. The “global tectonic map” of [82] is based on the Indo-Atlantic hotspot reference frame by O’Neill et al. [72] and on the relative plate motion model by DeMets et al. [14]. This map shows a large eastward motion of the Nazca plate, large in comparison to the magnitude of the velocity of the South American plate. In relation to this feature, this map is similar to the deforming, no-net-rotation reference frame model GSRM by Kreemer et al. [56]. In contrast to these two models, the hotspot reference model HS-3 [24] shows a high westward velocity of the South American plate, high in comparison with the amount of the velocity vectors of the Nazca plate. HS-3 is based on the age progressions of ten Pacific ocean islands. Becker [7] and Becker and Faccenna [8] wrote a very good analysis of the problem of “absolute” plate velocities. Becker [7] and Long and Becker [62] try to determine the present-day plate velocities from the convective shearing movements and the seismic anisotropy of the upper mantle. But this procedure is more sensitive to the direction of the velocity vector than to its magnitude. HS-3 contains a large net rotation of the laterally averaged lithosphere relative to the high-viscosity parts of the lower mantle. The majority of the authors derived a smaller real net rotation which is defined as the spherical harmonic degree $l = 1$ component of the toroidal part of the plate velocities. Ricard et al. [78] estimated 30 % of HS-3, Steinberger et al. [88] calculated 38 %. The azimuthal seismic anisotropy is compatible only with values less than 50 % of HS-3 [7].

Already early, Steinberger and O’Connell [87] linked the hot spot tracks and the movement of oceanic lithospheric plates. Gurnis et al. [26] report on a very practicable open-source system which contains the angular velocities for the plates from 140 Ma to the present time. Each plate has a time-dependent Euler pole. The plates are described by time-dependent closed plate polygons. Each of these plate boundaries has its own, time-dependent Euler pole. The code allows to introduce new interactive plate boundaries [9]. We intend to use two or three seemingly realistic spherical-shell plate-motion models to define the boundary conditions of our regional convective system of South America and its surrounding. The repercussions of the different spherical-shell systems on the regional model should be investigated and compared.

3.2.3 Future Numerical and Technical Improvements

In addition to the joint work within the group of Terra developers, further developments are to be continued regarding the following items.

- We are going to further improve the parallelization of the particle tracking routines in Terra. Compared to previous Terra versions, there is an extra need for communication among several MPI-processes to figure out connected regions of partial melting in the mantle from which incompatible elements are extracted and transported to the surface. A similar communication is required to define the extent of continental lithospheric plates. With the high number of tracers, it is crucial to compress the required global information locally before it is exchanged among neighboring processors. R. Hendel will continue to reduce the communication overhead for tracking globally connected regions, so that the scalability of the particle routines will be extended to 500 and more processors. Such an optimized way of communicating global regions and features is also needed in modeling the elasticity of the subducting plates in the regional Andean model.
- We also plan to bring the elasticity model in Terra to work. It has also to be chosen carefully how elasticity is dealt with in the solution of mass and momentum equations. It could be necessary to iterate over the whole Stokes within every time step until an equilibrium between elastic and viscous forces is reached.
- To run both, regional and global models, with time-dependent boundary conditions at the surface, plate reconstruction data will be imported from the GPlates code [9] (www.gplates.org). The development of the interface between GPlates and Terra will be done together with L. Quevedo, Sydney.
- Furthermore, the documentation of the code, which is build from source code comments automatically with doxygen, will be enhanced to make it easier for new developers to work on Terra.

Acknowledgements We gratefully acknowledge the stimulating discussions with J. Baumgardner, H.-P. Bunge, P. Bollada, H. Davies, R. Davies and M. Mohr in the group of Terra developers. We thank the Steinbuch Center for Computing, Karlsruhe, for the supply of computational time under grant sphshell. This work was partly supported by the Deutsche Forschungsgemeinschaft under grant KL 495/16-1.

References

1. R. W. Allmendinger, T. E. Jordan, S. M. Kay, and B. L. Isacks. The evolution of the Altiplano-Puna plateau of the Central Andes. *Annu. Rev. Earth Planet. Sci.*, 25: 139–174, 1997.
2. P. Alvarado, M. Pardo, H. Gilbert, S. Miranda, M. Anderson, M. Saez, and S. Beck. Flat-slab subduction and crustal models for the seismically active Sierras Pampeanas region of Argentina. In S. M. Kay and V. A. Ramos, editors, *Backbone of the Americas: Shallow subduction, Plateau Uplift, and Ridge and Terrane Collision*, pages 261–278. Geol. Soc. Am., 2009.

3. M. L. Anderson, G. Zandt, E. Triep, M. Fouch, and S. Beck. Anisotropy and mantle flow in the Chile-Argentina subduction zone from shear wave splitting analysis. *Geophys. Res. Lett.*, 31: L23608, 2004. doi: 10.1029/2004GL020906.
4. C. Arriagada, P. Roperch, C. Mpodozis, and P. R. Cobbold. Paleogene building of the Bolivian orocline: Tectonic restoration of the central Andes in 2-D map view. *Tectonics*, 27: 1–29, 2008. doi: 10.1029/2008TC002269.
5. J. B. Barnes and T. A. Ehlers. End member models for Andean Plateau uplift. *Earth Science Reviews*, 97: 105–132, 2009.
6. J. R. Baumgardner and P. O. Frederickson. Icosahedral discretization of the two-sphere. *SIAM J. Numer. Anal.*, 22: 1107–1115, 1985.
7. T. W. Becker. Azimuthal seismic anisotropy constrains net rotation of the lithosphere. *Geophys. Res. Lett.*, 35: L05303, 2008. doi: 10.1029/2007GL032928.
8. T. W. Becker and C. Faccenna. A review of the role of subduction dynamics for regional and global plate motions. In S. Lallemand and F. Funiciello, editors, *Subduction Zone Geodynamics*, pages 3–34. Springer, 2009. doi: 10.1007/978-3-540-87974-9.
9. J. Boyden, R. D. Müller, M. Gurnis, T. Torsvik, J. Clark, M. Turner, H. Ivey-Law, R. Watson, and J. S. Cannon. Next-generation plate-tectonic reconstructions using GPlates. In G. R. Keller and C. Baru, editors, *Geoinformatics: Cyberinfrastructure for the Solid Earth Sciences*, pages 95–113. Cambridge Univ. Press, 2011.
10. R. W. Carlson, D. G. Pearson, and D. E. James. Physical, chemical, and chronological characteristics of continental mantle. *Rev. Geophys.*, 43: RG1001, 24 pp., 2005. doi: 10.1029/2004RG000156.
11. J. P. Davidson and R. J. Arculus. The significance of Phanerozoic arc magmatism in generating continental crust. In M. Brown and T. Rushmer, editors, *Evolution and Differentiation of the Continental Crust*, pages 135–172. Cambridge Univ. Press, Cambridge, UK, 2006.
12. J. H. Davies and D. R. Davies. Earth’s surface heat flux. *Solid Earth*, 1: 5–24, 2010.
13. P. G. DeCelles, M. N. Ducea, P. Kapp, and G. Zandt. Cyclicity in Cordilleran orogenic systems. *Nature Geosci.*, 2: 251–257, 2009. doi: 10.1038/NGEO469.
14. C. DeMets, R. Gordon, D. Argus, and S. Stein. Effect of recent revisions to the geomagnetic reversal time scale on estimates of current plate motions. *Geophys. Res. Lett.*, 21: 2191–2194, 1994.
15. C. Dohrmann and P. Bochev. A stabilized finite element method for the Stokes problem based on polynomial pressure projections. *Int. J. Num. Meth. Fluids*, 46: 183–201, 2004.
16. H. Ege. *Exhumations-und Hebungsgeschichte der zentralen Anden in Südbolivien (21° S) durch Spaltspur-Thermochronologie an Apatit*. PhD thesis, Freie Univ. Berlin, 2004.
17. K. Elger, O. Oncken, and J. Glodny. Plateau-style accumulation of deformation: Southern Altiplano. *Tectonics*, 24: TC4020, 2005. doi: 10.1029/2004TC001675.
18. K. M. Fischer, H. A. Ford, D. L. Abt, and C. A. Rychert. The lithosphere-asthenosphere boundary. *Annu. Rev. Earth Planet. Sci.*, 38: 551–575, 2010.
19. Y. Furukawa. Convergence of aqueous fluid at the corner of the mantle wedge: Implications for a generation mechanism of deep low-frequency earthquakes. *Tectonophysics*, 469: 85–92, 2009.
20. C. W. Gable, R. J. O’Connell, and B. J. Travis. Convection in three dimensions with surface plates: Generation of toroidal flow. *J. Geophys. Res.*, 96: 8391–8405, 1991.
21. J. Ganguly, A. M. Freed, and S. K. Saxena. Density profiles of oceanic slabs and surrounding mantle: Integrated thermodynamic and thermal modeling, and implications for the fate of slabs at the 660 km discontinuity. *Phys. Earth Planet. Int.*, 172: 257–267, 2009.
22. K.-D. Gottschaldt, U. Walzer, R. Hendel, D. R. Stegman, J. R. Baumgardner, and H.-B. Mühlhaus. Stirring in 3-D spherical models of convection in the Earth’s mantle. *Philosophical Magazine*, 86: 3175–3204, 2006.
23. D. H. Green and T. J. Falloon. Pyrolyte: a Ringwood concept and its current expression. In I. Jackson, editor, *The Earth’s Mantle. Composition, Structure and Evolution*, pages 311–378. Cambridge Univ. Press, Cambridge, UK, 1998.

24. A. E. Gripp and R. G. Gordon. Young tracks of hotspots and current plate velocities. *Geophys. J. Int.*, 150: 321–361, 2002.
25. S. Guillot, K. Hattori, P. Agard, S. Schwartz, and O. Vidal. Exhumation processes in oceanic and continental subduction contexts: a review. In S. Lallemand and F. Funiello, editors, *Subduction Zone Geodynamics*, pages 175–205. Springer, 2009. doi: 10.1007/978-3-540-87974-9.
26. M. Gurnis, M. Turner, L. DiCaprio, S. Spasojević, R. D. Müller, J. Boyden, M. Seton, V. C. Manea, and D. J. Bower. Plate tectonic reconstructions with continuously closing plates. *Computers & Geosciences*, 38: 35–42, 2012.
27. M. A. Gutscher. Andean subduction styles and their effect on thermal structure and interplate coupling. *J. South Amer. Earth Sci.*, 15: 3–10, 2002.
28. M. A. Gutscher, W. Spakman, H. Bijwaard, and E. Engdahl. Geodynamics of flat subduction: Seismicity and tomographic constraints from the Andean margin. *Tectonics*, 19: 814–833, 2000.
29. M. Haschke, A. Günther, D. Melnick, H. Ehtler, K. Reutter, E. Scheuber, and O. Oncken. Central and southern Andean tectonic evolution inferred from arc magmatism. In O. Oncken et al., editors, *The Andes*, pages 337–353. Springer, Berlin, 2006.
30. A. Heuret and S. Lallemand. Plate motions, slab dynamics and back-arc deformation. *Phys. Earth Planet. Int.*, 149: 31–51, 2005.
31. D. Hindle and J. Kley. Displacements, strains and rotations in the Central Andean plate boundary zone. In S. Stein and J. T. Freymuller, editors, *Plate boundary zones*, Geodynamics Series 30, pages 135–144. AGU, Washington, D. C., 2002.
32. D. Hindle, J. Kley, E. Klosko, S. Stein, T. Dixon, and E. Norabuena. Consistency of geologic and geodetic displacements during Andean orogenesis. *Geophys. Res. Lett.*, 29: 1188, 2002. doi: 10.1029/2001GL013757.
33. D. Hindle, J. Kley, O. Oncken, and S. Sobolev. Crustal balance and crustal flux from shortening estimates in the Central Andes. *Earth Planet. Sci. Lett.*, 230: 113–124, 2005.
34. A. W. Hofmann. Sampling mantle heterogeneity through oceanic basalts: Isotopes and trace elements. In R. W. Carlson, editor, *Treatise on Geochemistry, Vol.2: The Mantle and the Core*, pages 61–101. Elsevier, Amsterdam, 2003.
35. E. Humphreys. Relation of flat subduction to magmatism and deformation in the western United States. In S. M. Kay and V. A. Ramos, editors, *Backbone of the Americas: Shallow subduction, Plateau Uplift, and Ridge and Terrane Collision*, pages 85–98. Geol. Soc. Am., 2009.
36. S. M. Kay and J. M. Abbruzzi. Magmatic evidence for Neogene lithospheric evolution of the central Andean “flat-slab” between 30S and 32S. *Tectonophysics*, 259 (1–3): 15–28, 1996.
37. S. M. Kay and B. L. Coira. Shallowing and steepening subduction zones, continental lithospheric loss, magmatism, and crustalflow under the Central Andean Altiplano-Puna Plateau. In S. M. Kay and V. A. Ramos, editors, *Backbone of the Americas: Shallow subduction, Plateau Uplift, and Ridge and Terrane Collision*, pages 229–259. Geol. Soc. Am., 2009.
38. S. M. Kay and C. Mpodozis. Central Andean ore deposits linked to evolving shallow subduction systems and thickening crust. *GSA Today*, 11 (3): 4–9, 2001.
39. S. M. Kay, C. Mpodozis, and B. Coira. Neogene magmatism, tectonism, and mineral deposits of the Central Andes (22 degrees to 33 degrees S latitude). In B. J. Skinner, editor, *Geology and Ore Deposits of the Central Andes. Soc. Econom. Geol. Spec. Pub. 7*, pages 27–59. 1998.
40. J. Kley. Der Übergang vom Subandin zur Ostkordillere in Südbolivien (21 15–22 S): Geologische Struktur und Kinematik. *Berliner geowiss. Abh. (A)*, 156: 1–88, 1993.
41. J. Kley. Transition from basement-involved to thin-skinned thrusting in the Cordillera Oriental of southern Bolivia. *Tectonics*, 15 (4): 763–775, 1996.
42. J. Kley. Structural styles of foreland deformation in the Andes. *Z. dt. geol. Ges.*, 149: 13–26, 1998.
43. J. Kley. Geologic and geometric constraints on a kinematic model of the Bolivian orocline. *J. South Amer. Earth Sci.*, 12: 221–235, 1999.

44. J. Kley and C. R. Monaldi. Tectonic shortening and crustal thickness in the Central Andes: How good is the correlation? *Geology*, 26: 723–726, 1998.
45. J. Kley and C. R. Monaldi. Estructura de las Sierras Subandinas y del Sistema de Santa Bárbara. In G. González Bonorino, R. Omarini, and J. Viramonte, editors, *Geología del Noroeste Argentino*, volume 1, pages 415–425. Salta, 1999.
46. J. Kley and C. R. Monaldi. Tectonic inversion in the Santa Barbara System of the central Andean foreland thrust belt, northwestern Argentina. *Tectonics*, 21: 1061, 2002. doi: 10.1029/2002TC902003.
47. J. Kley and M. Reinhardt. Geothermal and tectonic evolution of the Eastern Cordillera and the Subandean Ranges of southern Bolivia. In K.-J. Reutter, E. Scheuber, and P. J. Wigger, editors, *Tectonics of the southern Central Andes*, pages 155–170. Springer, Berlin, 1994.
48. J. Kley and T. Vietor. Subduction and mountain building in the central Andes. In T. Dixon and J. Moore, editors, *The Seismogenic Zone of Subduction Thrust Faults*, volume 2 of *MARGINS Theoretical and Experimental Earth Science Series*, pages 624–659. Columbia University Press, 2007.
49. J. Kley, J. Müller, S. Tawackoli, V. Jacobshagen, and E. Manutsoglu. Pre-Andean and Andean-age deformation in the eastern Cordillera of southern Bolivia. *J. South Amer. Earth Sci.*, 10 (1): 1–19, 1997.
50. J. Kley, C. R. Monaldi, and J. A. Salfity. Along-strike segmentation of the Andean foreland: Causes and consequences. *Tectonophysics*, 301: 75–94, 1999.
51. J. Kley, E. A. Rossello, C. R. Monaldi, and B. Habighorst. Seismic and field evidence for selective inversion of Cretaceous normal faults, Salta rift, Northwest Argentina. *Tectonophysics*, 399 (1–4): 155–172, 2005.
52. E. R. Klosko, S. Stein, D. Hindle, J. Kley, E. Norabuena, T. Dixon, and M. Liu. Comparison of GPS, seismological, and geological observations of Andean mountain building. In S. Stein and J. T. Freymuller, editors, *Plate boundary zones*, Geodynamics Series 30, pages 123–133. AGU, Washington, D. C., 2002.
53. E. A. Kneller and P. E. van Keken. Trench-parallel flow and seismic anisotropy in the Mariana and Andean subduction systems. *Nature*, 450: 1222–1225, 2007. doi: 10.1038/nature06429.
54. C. Köstler. *Iterative solvers for modeling mantle convection with strongly varying viscosity*. PhD thesis, Friedrich-Schiller-Univ. Jena, <http://www.igw.uni-jena.de/geodyn>, 2011.
55. C. Köstler, M. Müller, U. Walzer, and J. Baumgardner. Krylov solvers and preconditioners for variable-viscosity convection models. *Comp. Geosci.*, submitted, 2012.
56. C. Kreemer, W. Holt, and A. Haines. An integrated global model of present-day plate motions and plate boundary deformation. *Geophys. J. Int.*, 154: 8–34, 2003.
57. N. Kukowski and O. Oncken. Subduction erosion - the “normal” mode of fore-arc material transfer along the Chilean margin? In O. Oncken et al., editors, *The Andes*, pages 217–236. Springer, Berlin, 2006.
58. S. Lamb and P. Davis. Cenozoic climate change as a possible cause for the rise of the Andes. *Nature*, 425: 792–797, 2003.
59. Z. X. Li and X. H. Li. Formation of the 1300-km-wide intracontinental orogen and postorogenic magmatic province in Mesozoic South China: A flat-slab subduction model. *Geology*, 35 (2): 179–182, 2007. doi: 10.1130/G23193A.1.
60. K. Litasov and E. Ohtani. Phase relations and melt compositions in CMAS-pyrolite–H₂O system up to 25 GPa. *Phys. Earth Planet. Int.*, 134: 105–127, 2002.
61. K. D. Litasov. Physicochemical conditions for melting in the Earth’s mantle containing a C-O-H fluid (from experimental data). *Russian Geology and Geophysics*, 52: 475–492, 2011.
62. M. D. Long and T. W. Becker. Mantle dynamics and seismic anisotropy. *Earth Planet. Sci. Lett.*, 297: 341–354, 2010.
63. B. D. Marsh. Magmatism, magma, and magma chambers. In G. Schubert, editor, *Treatise on Geophysics*, volume 6, pages 276–333. Elsevier, 2007.
64. F. Michaud, C. Witt, and J. Y. Royer. Influence of the subduction of the Carnegie volcanic ridge on Ecuadorian geology: Reality and fiction. In S. M. Kay and V. A. Ramos, editors, *Backbone of the Americas: Shallow subduction, Plateau Uplift, and Ridge and Terrane Collision*, pages 217–228. Geol. Soc. Am., 2009.

65. J. X. Mitrovica and A. M. Forte. A new inference of mantle viscosity based upon joint inversion of convection and glacial isostatic adjustment data. *Earth Planet. Sci. Lett.*, 225 (1): 177–189, 2004.
66. P. Molnar and T. Atwater. Interarc spreading and Cordilleran tectonics as alternates related to the age of subducted oceanic lithosphere. *Earth Planet. Sci. Lett.*, 41: 330–340, 1978.
67. C. R. Monaldi, J. A. Salfity, and J. Kley. Preserved extensional structures in an inverted Cretaceous rift basin, northwestern Argentina: Outcrop examples and implications for fault reactivation. *Tectonics*, 27: C1011, 2008. doi: 10.1029/2006TC001993.
68. M. Müller. *Towards a robust Terra code*. PhD thesis, Friedrich-Schiller-Univ. Jena, <http://www.igw.uni-jena.de/geodyn>, 2008.
69. M. Müller and C. Köstler. Stabilization of the 3-D spherical convection code Terra. *Int. Jour. Num. Meth. Engng.*, submitted, 2012.
70. R. J. O'Connell, C. W. Gable, and B. H. Hager. Toroidal-poloidal partitioning of lithospheric plate motions. In R. Sabadini and other, editors, *Glacial Isostasy, Sea-Level and Mantle Rheology*, pages 535–551. Kluwer, Norwell, MA, 1991.
71. O. Oncken, D. Hindle, J. Kley, K. Elger, P. Victor, and K. Schemmann. Deformation of the central Andean upper plate system - facts, fiction, and constraints for plateau models. In O. Oncken et al., editors, *The Andes*, pages 3–27. Springer, Berlin, 2006.
72. C. O'Neill, D. Müller, and B. Steinberger. On the uncertainties in hot spot reconstructions and the significance of moving hot spot reference frames. *Geochem. Geophys. Geosys.*, 6: Q04003, 2005. doi: 10.1029/2004GC000784.
73. P. Pilz. *Ein neues magmatisch-tektonisches Modell zur Asthenosphärendynamik im Bereich der zentralandinen Subduktionszone Südamerikas*. PhD thesis, Universität Potsdam, 2008.
74. J. Pindell and L. Kennan. Tectonic evolution of the Gulf of Mexico, Caribbean and northern South America in the mantle reference frame: An update. *Geol. Soc. London*, 328: 1–55, 2009.
75. V. Ramos, E. Cristallini, and D. Pérez. The Pampean flat-slab of the Central Andes. *J. South Amer. Earth Sci.*, 15: 59–78, 2002.
76. V. A. Ramos. Anatomy and global context of the Andes: Main geologic features and the Andean orogenic cycle. In S. M. Kay and V. A. Ramos, editors, *Backbone of the Americas: Shallow subduction, Plateau Uplift, and Ridge and Terrane Collision*, pages 31–65. Geol. Soc. Am., 2009.
77. V. A. Ramos and A. Folguera. Andean flat-slab subduction through time. In B. Murphy et al., editors, *Ancient Orogens and Modern Analogues*, pages 31–54. Geol. Soc. London, 2009.
78. Y. Ricard, C. Doglioni, and R. Sabadini. Differential rotation between lithosphere and mantle: A consequence of lateral mantle viscosity variations. *J. Geophys. Res.*, 96: 8407–8415, 1991.
79. G. Rosenbaum, D. Giles, M. Saxon, P. G. Betts, R. F. Weinberg, and C. Duboz. Subduction of the Nazca Ridge and the Inca Plateau: Insights into the formation of ore deposits in Peru. *Earth Planet. Sci. Lett.*, 239: 18–32, 2005.
80. L. H. Rüpke, J. Phipps Morgan, M. Hort, and J. A. D. Connolly. Serpentine and the subduction zone water cycle. *Earth Planet. Sci. Lett.*, 223: 17–34, 2004.
81. W. P. Schellart. Kinematics of subduction and subduction-induced flow in the upper mantle. *J. Geophys. Res.*, 109: B07401, 2004. doi: 10.1029/2004JB002970.
82. W. P. Schellart and N. Rawlinson. Convergent plate margin dynamics: New perspectives from structural geology, geophysics and geodynamic modelling. *Tectonophysics*, 483: 4–19, 2010.
83. M. Sébrier and P. Soler. Tectonics and magmatism in the Peruvian Andes from late Oligocene time to the present. In R. S. Harmon and C. W. Rapela, editors, *Andean Magmatism and its Tectonic Setting*, volume 265 of *Geol. Soc. Am. Spec. Paper*, pages 259–278. Geol. Soc. of America, Boulder, Col., 1991.
84. S. Sobolev, A. Babeyko, I. Koulakov, and O. Oncken. Numerical study of weakening processes in the central Andean back-arc. In O. Oncken et al., editors, *The Andes*, pages 513–535. Springer, Berlin, 2006.
85. S. V. Sobolev and A. Y. Babeyko. What drives orogeny in the Andes? *Geology*, 33 (8): 617–620, 2005.

86. D. Stegman, J. Freeman, W. Schellart, L. Moresi, and D. May. Influence of trench width on subduction hinge retreat rates in 3-D models of slab rollback. *Geochem. Geophys. Geosys.*, 7: Q03012, 2006. doi: 10.1029/2005GC001056.
87. B. Steinberger and R. J. O'Connell. Advection of plumes in mantle flow: Implications for hotspot motion, mantle viscosity and plume distribution. *Geophys. J. Int.*, 132: 412–434, 1998.
88. B. Steinberger, R. Sutherland, and R. J. O'Connell. Prediction of Emperor-Hawaii seamount locations from a revised model of global plate motion and mantle flow. *Nature*, 430: 167–173, 2004.
89. A. Stracke, A. W. Hofmann, and S. R. Hart. FOZO, HIMU and the rest of the mantle zoo. *Geochem. Geophys. Geosys.*, 6: Q05007, 2005. doi: 10.1029/2004GC000824.
90. P. J. Tackley. Modelling compressible mantle convection with large viscosity contrasts in a three-dimensional spherical shell using the yin-yang grid. *Phys. Earth Planet. Int.*, 171: 7–18, 2008.
91. Y. Tamura, Y. Tatsumi, D. Zhao, Y. Kido, and H. Shukuno. Hot fingers in the mantle wedge: New insights into magma genesis in subduction zones. *Earth Planet. Sci. Lett.*, 197: 105 – 116, 2002.
92. T. Vietor and H. Echtler. Episodic Neogene southward growth of the Andean subduction orogen between 30° S and 40° S - Plate motions, mantle flow, climate, and upper-plate structure. In O. Oncken et al., editors, *The Andes*, pages 375–400. Springer, Berlin, 2006.
93. U. Walzer and R. Hendel. Time-dependent thermal convection, mantle differentiation, and continental crust growth. *Geophys. J. Int.*, 130: 303–325, 1997a.
94. U. Walzer and R. Hendel. Tectonic episodicity and convective feedback mechanisms. *Phys. Earth Planet. Int.*, 100: 167–188, 1997b.
95. U. Walzer and R. Hendel. A new convection-fractionation model for the evolution of the principal geochemical reservoirs of the Earth's mantle. *Phys. Earth Planet. Int.*, 112: 211–256, 1999.
96. U. Walzer and R. Hendel. A new convection-fractionation model for the evolution of the principal geochemical reservoirs of the Earth's mantle. *EOS Transactions*, 80 (46): F1171, 2000.
97. U. Walzer and R. Hendel. Mantle convection and evolution with growing continents. *J. Geophys. Res.*, 113: B09405, 2008. doi: 10.1029/2007JB005459.
98. U. Walzer and R. Hendel. Predictability of Rayleigh-number and continental-growth evolution of a dynamic model of the Earth's mantle. In S. Wagner, M. Steinmetz, A. Bode, and M. Brehm, editors, *High Perf. Comp. Sci. Engng. '07*, pages 585–600. Springer, Berlin, 2009.
99. U. Walzer and R. Hendel. A geodynamic model of the evolution of the Earth's chemical mantle reservoirs. In W. E. Nagel, D. B. Kröner, and M. M. Resch, editors, *High Perf. Comp. Sci. Engng. '10*, pages 573–592. Springer, Berlin, 2011.
100. U. Walzer and R. Hendel. Silicate Earth differentiation, 3-D spherical-shell mantle convection, and evolution of continents. *Earth Planet. Sci. Lett.*, submitted, 2012a.
101. U. Walzer and R. Hendel. Real episodic growth of continental crust or artefact of preservation? A 3-D geodynamic model. *J. Geophys. Res.*, submitted, 2012b.
102. U. Walzer, R. Hendel, and J. Baumgardner. Variation of non-dimensional numbers and a thermal evolution model of the Earth's mantle. In E. Krause and W. Jäger, editors, *High Perf. Comp. Sci. Engng. '02*, pages 89–103. Berlin, 2003.
103. U. Walzer, R. Hendel, and J. Baumgardner. Viscosity stratification and a 3D compressible spherical shell model of mantle evolution. In E. Krause, W. Jäger, and M. Resch, editors, *High Perf. Comp. Sci. Engng. '03*, pages 27–67. Springer, Berlin, 2004a.
104. U. Walzer, R. Hendel, and J. Baumgardner. The effects of a variation of the radial viscosity profile on mantle evolution. *Tectonophysics*, 384: 55–90, 2004b.
105. U. Walzer, R. Hendel, and J. Baumgardner. Generation of plate-tectonic behavior and a new viscosity profile of the Earth's mantle. In D. Wolf, G. Münster, and M. Kremer, editors, *NIC Symposium*, volume 20 of *NIC Series*, pages 419–428, Jülich, 2004c. Springer.

106. U. Walzer, R. Hendel, and J. Baumgardner. Toward a thermochemical model of the evolution of the Earth's mantle. In E. Krause, W. Jäger, and M. Resch, editors, *High Perf. Comp. Sci. Engng. '04*, pages 395–454. Springer, Berlin, 2005.
107. U. Walzer, R. Hendel, and J. Baumgardner. Plateness of the oceanic lithosphere and the thermal evolution of the Earth's mantle. In W. E. Nagel, W. Jäger, and M. Resch, editors, *High Perf. Comp. Sci. Engng. '05*, pages 289–304. Springer, Berlin, 2006.
108. U. Walzer, R. Hendel, and J. Baumgardner. Continental growth and thermal convection in the Earth's mantle. In W. E. Nagel, W. Jäger, and M. Resch, editors, *High Perf. Comp. Sci. Engng. '06*, pages 473–497. Springer, Berlin, 2007.
109. U. Walzer, R. Hendel, and J. Baumgardner. Whole-mantle convection, continent generation, and preservation of geochemical heterogeneity. In W. E. Nagel, D. B. Kröner, and M. M. Resch, editors, *High Perf. Comp. Sci. Engng. '07*, pages 603–645. Springer, Berlin, 2008.
110. U. Walzer, R. Hendel, C. Köstler, and J. Kley. Andean Orogeny and Plate Generation. In W. E. Nagel, D. B. Kröner, and M. M. Resch, editors, *High Perf. Comp. Sci. Engng. '08*, pages 559–583. Springer, Berlin, 2009.
111. B. J. Wood. The effect of H₂O on the 410-kilometer seismic discontinuity. *Science*, 268: 74–76, 1995.
112. B. J. Wood and A. Corgne. Mineralogy of the Earth: Trace elements and hydrogen in the Earth's transition zone and lower mantle. In G. Schubert, editor, *Treatise on Geophysics*, volume 2, pages 63–89. Elsevier, 2007.
113. D. Yamazaki and S.-I. Karato. Some mineral physics constraints on the rheology and geothermal structure of the Earth's lower mantle. *Am. Min.*, 86: 385–391, 2001.
114. W.-S. Yang and J. R. Baumgardner. A matrix-dependent transfer multigrid method for strongly variable viscosity infinite Prandtl number thermal convection. *Geophys. Astrophys. Fluid Dyn.*, 92: 151–195, 2000.
115. A. Yoneda, M. Osako, and E. Ito. Heat capacity measurement under high pressure: A finite element method assessment. *Phys. Earth Planet. Int.*, 174: 309–314, 2009.

Euler Deconvolution of GOCE Gravity Gradiometry Data

M. Roth, N. Sneeuw, and W. Keller

Abstract Euler deconvolution is a standard tool of geophysical prospection. In the early 1980s, the beginning of its development, it was used for the evaluation of magnetic field data. However, since the 1990s, together with the increasing power of computers, research was intensified on the aspect that Euler deconvolution is also applicable to gravity gradiometry data. Now we are in the position that gravity gradiometry data with near global coverage from a single source are available, namely the satellite mission GOCE (Gravity field and steady-state Ocean Circulation Explorer), launched on 17 March 2009.

In this project we investigate the benefit of Euler deconvolution to geodesy, e.g. to retrieve global gravity models. We also assess whether geodetic methodology can contribute to enhance Euler deconvolution. Until now our project is still in preparatory stage, mainly, because the GOCE gradiometry data need to be preprocessed extensively.

1 Introduction

Euler deconvolution is a semi-automated method to estimate possible locations of magnetic field or gravity field sources. Since the paper of Thompson [10], who made Euler deconvolution applicable to magnetic data by the use of computers, this method became of great interest in research.

Marson and Klingel  [5] applied Euler deconvolution for the first time to gravity vertical gradients, Zhang et al. [12] applied it to full gravity gradient tensor data. Usually, such tensor data are measured on shipborne or airborne platforms for relatively small areas (several hundreds of square kilometers).

M. Roth (✉) · N. Sneeuw · W. Keller
University of Stuttgart, Institute of Geodesy, Geschwister-Scholl-Str. 24D,
D-70174 Stuttgart, Germany
e-mail: matthias.roth@gis.uni-stuttgart.de

The satellite mission GOCE provides a global geoid model in unprecedented precision [7]. The nominally planned operational phase of 12 months was finished on 2 March 2011 after 2 years in orbit. (The operational phase was interrupted by calibration phases and failures of the on-board computers which could be fixed.) However, the satellite's overall good health, as well as its excellent data quality led to the extension of the mission at least until the end of 2012 [3].

The data provided by GOCE are mainly the full gravity gradient tensor, as well as the satellite's coordinates for every observation epoch. However, satellite positions and data are given in different reference and time frames, which is one reason why a preprocessing of the data becomes necessary.

Up to the present, GOCE delivered a vast amount of data, the sample frequency being 1 Hz. The amount of data (around 70 GB at the moment), and the computational method of Euler deconvolution as presented in the following section, makes high performance computing (HPC) a necessity.

2 Euler Deconvolution

2.1 Euler's Homogeneous Function Theorem

The general form of Euler's homogeneous function theorem reads

$$\mathbf{x} \cdot \nabla f(\mathbf{x}) = n f(\mathbf{x}). \quad (1)$$

Here \mathbf{x} stands for the vector $(x_1, \dots, x_i, \dots, x_k)^\top$ and ∇ is the gradient operator $(\frac{\partial}{\partial x_1}, \dots, \frac{\partial}{\partial x_i}, \dots, \frac{\partial}{\partial x_k})^\top$. The dot product of both yields $\mathbf{x} \cdot \nabla = (x_1 \frac{\partial}{\partial x_1}, \dots, x_i \frac{\partial}{\partial x_i}, \dots, x_k \frac{\partial}{\partial x_k})^\top$.

If we now take the first-order derivative to one of the variables, e.g. x_i , we get for the left hand side

$$\begin{aligned} \frac{\partial}{\partial x_i} \left[\mathbf{x} \cdot \nabla f(\mathbf{x}) \right] &= \underbrace{\frac{\partial}{\partial x_i} x_i}_{=1} \frac{\partial}{\partial x_i} f(\mathbf{x}) + \mathbf{x} \cdot \nabla \frac{\partial}{\partial x_i} f(\mathbf{x}) \\ &= \frac{\partial}{\partial x_i} f(\mathbf{x}) + \mathbf{x} \cdot \nabla \frac{\partial}{\partial x_i} f(\mathbf{x}). \end{aligned}$$

The right hand side reads then

$$\frac{\partial}{\partial x_i} \left[n f(\mathbf{x}) \right] = n \frac{\partial}{\partial x_i} f(\mathbf{x}).$$

Putting both sides back together yields

$$\frac{\partial}{\partial x_i} f(\mathbf{x}) + \mathbf{x} \cdot \nabla \frac{\partial}{\partial x_i} f(\mathbf{x}) = n \frac{\partial}{\partial x_i} f(\mathbf{x}),$$

which can be simplified to

$$\mathbf{x} \cdot \nabla \frac{\partial}{\partial x_i} f(\mathbf{x}) = (n - 1) \frac{\partial}{\partial x_i} f(\mathbf{x}). \tag{2}$$

This expression now has the same form as (1). It shows, that the first-order partial derivatives $\frac{\partial}{\partial x_i} f(\mathbf{x})$ are homogeneous of degree $n - 1$.

2.2 Gravity Gradients

Because of the principle of superposition, we can consider Earth being assembled from numerous point masses. One of those point masses, located at the point $\mathbf{r}_0 = (x_0, y_0, z_0)$, produces at the measurement location $\mathbf{r} = (x, y, z)$ the potential

$$\delta V = -\frac{G \delta M}{r}, \tag{3}$$

where $r = \sqrt{(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2}$ is the distance between \mathbf{r}_0 and \mathbf{r} , $G = 6.67384 \cdot 10^{-11} \frac{\text{m}^3}{\text{kg s}^2}$ [2] is the Newtonian constant of gravitation and δM is the mass of the probe. We obtain the acceleration $\delta \mathbf{g}$ by taking the gradient of (3).

$$\delta \mathbf{g} = \nabla \delta V = \begin{pmatrix} \delta V_x \\ \delta V_y \\ \delta V_z \end{pmatrix} = \begin{pmatrix} \partial \delta V / \partial x \\ \partial \delta V / \partial y \\ \partial \delta V / \partial z \end{pmatrix} = \frac{G \delta M}{r^3} \begin{pmatrix} x - x_0 \\ y - y_0 \\ z - z_0 \end{pmatrix}. \tag{4}$$

By a further gradient operation we obtain the gravity gradient tensor

$$\begin{aligned} \nabla \delta \mathbf{g} = \nabla \otimes \nabla \delta V &= \begin{pmatrix} \delta V_{xx} & \delta V_{xy} & \delta V_{xz} \\ \delta V_{yx} & \delta V_{yy} & \delta V_{yz} \\ \delta V_{zx} & \delta V_{zy} & \delta V_{zz} \end{pmatrix} \\ &= \begin{pmatrix} \partial^2 \delta V / \partial x^2 & \partial^2 \delta V / \partial x \partial y & \partial^2 \delta V / \partial x \partial z \\ \partial^2 \delta V / \partial y \partial x & \partial^2 \delta V / \partial y^2 & \partial^2 \delta V / \partial y \partial z \\ \partial^2 \delta V / \partial z \partial x & \partial^2 \delta V / \partial z \partial y & \partial^2 \delta V / \partial z^2 \end{pmatrix}. \end{aligned} \tag{5}$$

Four gradients are linearly dependent because of the tensor's symmetry, i.e.

$$\begin{aligned}\delta V_{xy} &= \delta V_{yx}, \\ \delta V_{xz} &= \delta V_{zx}, \\ \delta V_{yz} &= \delta V_{zy},\end{aligned}\tag{6}$$

and the validity of Laplace's equation

$$\delta V_{xx} + \delta V_{yy} + \delta V_{zz} = 0.\tag{7}$$

Therefore it is sufficient to specify the following five independent tensor elements:

$$\begin{aligned}\delta V_{xx} &= \frac{G \delta M}{r^5} (-3(x - x_0)^2 + r^2), \\ \delta V_{xy} &= \frac{G \delta M}{r^5} (-3(x - x_0)(y - y_0)), \\ \delta V_{xz} &= \frac{G \delta M}{r^5} (-3(x - x_0)(z - z_0)), \\ \delta V_{yy} &= \frac{G \delta M}{r^5} (-3(y - y_0)^2 + r^2), \\ \delta V_{yz} &= \frac{G \delta M}{r^5} (-3(y - y_0)(z - z_0)).\end{aligned}\tag{8}$$

As an example, the gravity gradient tensor components of a single point mass are displayed in Fig. 1.

2.3 Standard Euler Deconvolution

This subsection portrays the standard approach for Euler deconvolution as presented e.g. by Thompson [10], Reid et al. [8], Wu [11] or Mushayandebvu et al. [6]. Euler's theorem (1) for the 3D case is given by

$$(x - x_0) \frac{\partial f}{\partial x} + (y - y_0) \frac{\partial f}{\partial y} + (z - z_0) \frac{\partial f}{\partial z} = nf,\tag{9}$$

again, (x_0, y_0, z_0) represents the source location, (x, y, z) the measurement location, f the field and n the degree.

Additionally, for the computation of the Euler deconvolution, we assume that the field f is the sum of a constant base field B and the difference Δf to the actual field




Fig. 1 Example of the gravity gradient tensor components T_{ij} of a point mass [9]

$$f = \Delta f + B .$$

We also benefit from the fact that (because of $B = \text{const.}$) it holds

$$\frac{\partial}{\partial z} (\Delta f + B) = \frac{\partial}{\partial z} \Delta f = \frac{\partial}{\partial z} f .$$

In consequence, (9) becomes

$$(x - x_0) \frac{\partial f}{\partial x} + (y - y_0) \frac{\partial f}{\partial y} + (z - z_0) \frac{\partial f}{\partial z} = n(\Delta f + B) . \tag{10}$$

To insert the gravity gradients, we set e.g. $f = \delta V_z$ and $\frac{\partial f}{\partial x} = \delta V_{xz}$, $\frac{\partial f}{\partial y} = \delta V_{yz}$ etc. and get

$$(x - x_0) \delta V_{xz} + (y - y_0) \delta V_{yz} + (z - z_0) \delta V_{zz} = (n - 1)(\Delta \delta V_z + B_z) . \tag{11}$$

The same procedure can be applied to the other gravity vector and tensor components. Hence, it is possible to use the complete gravity gradient tensor to determine the source of the gravity anomaly. However, this example sticks to the z -components of the gravity gradient tensor as done in (11).

In a next step, the observation data usually is either interpolated to regularly spaced points on lines or planar grids. A window of size m slides over the interpolated data and for each window a solution for the unknown parameters is found by a least squares estimation (Gauss–Markov model). However, if the measurement positions are uniformly distributed, this interpolation step might not be necessary. A reformulation of (11) for m sample points (setting $N = -(n - 1)$) yields

$$\begin{aligned} & \begin{pmatrix} \delta V_{xx,1} & \delta V_{yz,1} & \delta V_{xz,1} & -N \\ \vdots & \vdots & \vdots & \vdots \\ \delta V_{xx,m} & \delta V_{yz,m} & \delta V_{xz,m} & -N \end{pmatrix} \begin{pmatrix} x_0 \\ y_0 \\ z_0 \\ \mathbf{B}_z \end{pmatrix} \\ &= \begin{pmatrix} x_1 \delta V_{xz,1} + y_1 \delta V_{yz,1} + z_1 \delta V_{zz,1} + N \delta V_{z,1} \\ \vdots \\ x_m \delta V_{xz,m} + y_m \delta V_{yz,m} + z_m \delta V_{zz,m} + N \delta V_{z,m} \end{pmatrix}, \quad (12) \end{aligned}$$

[6]. N is the so called structural index and is also unknown per se (but needs to be chosen in the process of Euler deconvolution, e.g. $N = 2$ for point masses).

Now (12) is of the form

$$\boldsymbol{\eta} = \mathbf{A}\boldsymbol{\xi} + \mathbf{e}.$$

with the vector of “observations” $\boldsymbol{\eta}$ being the right hand side of (12), design matrix \mathbf{A} , vector of parameters $\boldsymbol{\xi}$ and the residual vector \mathbf{e} which holds model and measurement errors.

We need to choose an adequate window size that the window contains enough measurement positions and, in consequence, the problem becomes overdetermined. Then we are able to treat it by the least squares method to obtain the estimated parameters $\hat{\boldsymbol{\xi}}$.

$$\boldsymbol{\eta} = \mathbf{A}\boldsymbol{\xi} + \mathbf{e} \implies \mathbf{A}^T \boldsymbol{\eta} = \mathbf{A}^T \mathbf{A}\boldsymbol{\xi} \implies \hat{\boldsymbol{\xi}} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \boldsymbol{\eta},$$

In this way, for each window position one estimate of the source’s location is computed. Those estimated locations tend to cluster in zones of contrast in a field, hence might be of geological or geophysical interest [1]. Choosing the right window size is a bit troublesome. On the one hand, it has to be large enough that a single source gravitational effect is covered; on the other hand, it should be small enough that significant effects from multiple sources are not included [11]. A synthetic example of Euler deconvolution of multiple sources is given in Fig. 2 for the 2D case.

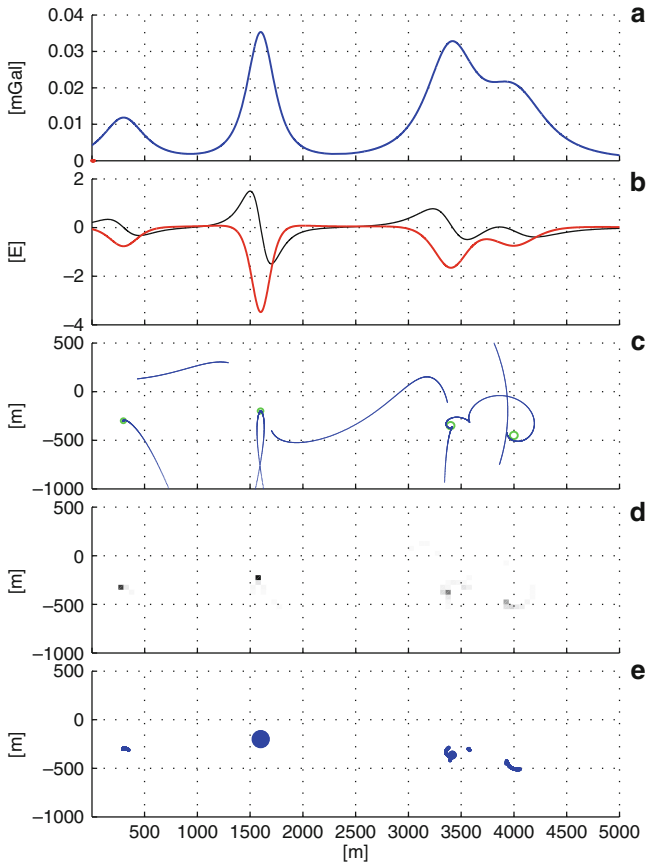


Fig. 2 Synthetic example of Euler deconvolution, 2D (along-track), multiple sources [9]. (a) δV_z (blue), window size (red); (b) δV_{xz} (black), δV_{zz} (red); (c) source (green circle) and solutions of Euler deconvolution (blue); (d) counting the amount of solution per grid cell (darker = more); (e) sum of reciprocal distances to the other solutions (blue circles)

2.4 Extended Euler Deconvolution

The standard Euler deconvolution, as presented above, disregards the fact that the gravity gradients are measured quantities. Measured quantities are *never* error free. As the design matrix consists mainly of such measured quantities, the included errors lead to an inaccurate estimation of the parameters. In such a case the Gauss–Helmert model is the better choice in terms of estimation precision. In this model also the measured quantities’ errors can be considered. However, the drawback is that matrices get several times bigger due to the additional conditions. We also want to retrieve variance-covariance information, hence we need one inversion of a matrix

per window. For the standard Euler deconvolution the size of that matrix depends only on the amount of unknown parameters, i.e. matrix size is 4×4 . The matrix size of the extended Euler deconvolution depends also on m (the amount of observations in each window), i.e. the matrix size increases to $4m + 4 \times 4m + 4$. This leads to an increase of computing time increase by a factor of 60.

Measured—and as such stochastic—quantities are $x, y, z, \delta V_{ij}, \Delta \delta V_i$; unknown quantities are x_0, y_0, z_0, B_i ($i, j = x, y, z$). Exemplarily, out of the three possible equations, let us examine (11), while taking into account the stochastic quantities (e.g. x becomes now $x + e_x$)

$$(x + e_x - x_0)(\delta V_{xz} + e_{\delta V_{xz}}) + (y + e_y - y_0)(\delta V_{yz} + e_{\delta V_{yz}}) + (z + e_z - z_0)(\delta V_{zz} + e_{\delta V_{zz}}) + N(\Delta \delta V_z + e_{\delta V_z} + B_z) = 0. \quad (13)$$

Equation (13) is not linear anymore, hence, linearization becomes necessary. Additionally, we can assume that the errors of (x, y, z) are very small (the position of GOCE is known from GPS positioning with high accuracy, i.e. we set $e_x = e_y = e_z = 0$). Rewritten in matrix-vector notation, we get

$$\begin{aligned} & \left[-(\delta V_{xz} + e_{\delta V_{xz}}) \quad -(\delta V_{yz} + e_{\delta V_{yz}}) \quad -(\delta V_{zz} + e_{\delta V_{zz}}) \quad N \right] \begin{bmatrix} \Delta x_0 \\ \Delta y_0 \\ \Delta z_0 \\ \Delta B_z \end{bmatrix} \\ & + \left[(x + e_x - x_0) \quad (y + e_y - y_0) \quad (z + e_z - z_0) \quad N \right] \begin{bmatrix} e_{\delta V_{xz}} \\ e_{\delta V_{yz}} \\ e_{\delta V_{zz}} \\ e_{\Delta \delta V_z} \end{bmatrix} \\ & + (x - x_0) \delta V_{xz} + (y - y_0) \delta V_{yz} + (z - z_0) \delta V_{zz} + N(\Delta \delta V_z + B_z) = 0. \quad (14) \end{aligned}$$

Here, 0 on top of a variable indicates, that this variable is evaluated at the Taylor point. Equation (14) is nearly in the form of the Gauss–Helmert model, so we can write

$$\mathbf{A} \Delta \boldsymbol{\xi} + \mathbf{B}^T \mathbf{e} + \mathbf{w} = 0 \quad (15)$$

with

$$\mathbf{B}^T_{m \times 4m} = \begin{bmatrix} X_1 & 0 & \dots & 0 & Y_1 & 0 & \dots & 0 & Z_1 & 0 & \dots & 0 & N & 0 & \dots & 0 \\ 0 & X_2 & 0 & \dots & 0 & Y_2 & 0 & \dots & 0 & Z_2 & 0 & \dots & 0 & N & 0 & \dots \\ \vdots & & \ddots & & \vdots & & \ddots & & \vdots & & \ddots & & \vdots & & \ddots & \\ 0 & \dots & 0 & X_m & 0 & \dots & 0 & Y_m & 0 & \dots & 0 & Z_m & 0 & \dots & 0 & N \end{bmatrix},$$

where $X_i = x_i + e_{x,i}^0 - x_{0,i}^0$, $Y_i = y_i + e_{y,i}^0 - y_{0,i}^0$ and $Z_i = z_i + e_{z,i}^0 - z_{0,i}^0$, $i = 1, \dots, m$, and

$$\mathbf{e}_{4m \times 1}^T = [e_{\delta V_{xz},1} \dots e_{\delta V_{yz},1} \dots e_{\delta V_{zz},1} \dots e_{\delta V_{z,1}} \dots].$$

Matrix \mathbf{A} (of size $m \times 4$) and vector \mathbf{w} (of length m) can be derived directly from (14). Minimizing the Legendre function

$$\mathcal{L}(\Delta \hat{\boldsymbol{\xi}}, \mathbf{e}, \boldsymbol{\lambda}) = \frac{1}{2} \mathbf{e}^T \mathbf{e} + \boldsymbol{\lambda}^T (\mathbf{A} \Delta \hat{\boldsymbol{\xi}} + \mathbf{B}^T \mathbf{e} + \mathbf{w}) \quad \longrightarrow \quad \min_{\Delta \hat{\boldsymbol{\xi}}, \mathbf{e}, \boldsymbol{\lambda}}$$

leads us to the linear system of equations

$$\begin{bmatrix} 0 & 0 & \mathbf{A}^T \\ 0 & \mathbf{I} & \mathbf{B} \\ \mathbf{A} & \mathbf{B}^T & 0 \end{bmatrix} \begin{bmatrix} \Delta \hat{\boldsymbol{\xi}} \\ \hat{\mathbf{e}} \\ \hat{\boldsymbol{\lambda}} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -\mathbf{w} \end{bmatrix} \tag{16}$$

whose solution is refined iteratively, until the vector of increments $\Delta \hat{\boldsymbol{\xi}}$ becomes small enough to meet the accuracy threshold. The initial values of the variables evaluated at the Taylor point are obtained by a preceding standard Euler deconvolution.

As a preliminary result, in Fig. 3 solutions of Euler deconvolution of the first 2 months of GOCE data over Africa are displayed. In this early stage, we assumed to detect point masses, i.e. $N = 2$, well knowing that the structures are not that simple in reality. At the moment we cannot explain the results of the depth estimates of the Euler deconvolution solutions which roughly lies around 4430 km distance from the geocentre.

3 Windowing of the Data

We already mentioned earlier, that Euler deconvolution is achieved by sliding a window over the data and estimating one set of parameters per window. However, the measurement data are ordered chronologically, i.e. we get the data along the satellite orbit. In conclusion we have to project a grid on those data to get the data windows for Euler deconvolution.

The simplest approach would be to divide the data according to Earth’s already present latitude-longitude grid. This has the disadvantage, that, depending on latitude, the window does not cover an equally sized area. Hence we follow the more sophisticated approach of a geodesic sphere (cf. Fig. 4).

Starting with a icosahedron whose top and bottom vertices coincides with the poles, we first sort the data to its 12 faces. Afterward, we divide each face into four new ones whose vertices are projected on the surface of the sphere. Again, we sort the data of the face to the four new faces. The last two steps are repeated

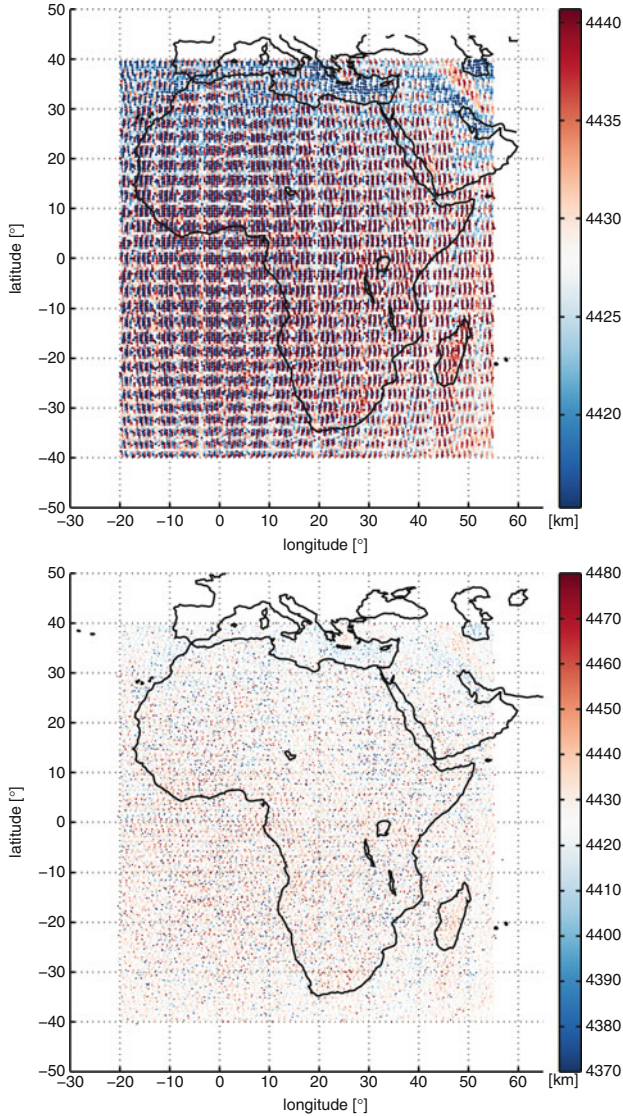


Fig. 3 Euler deconvolution of GOCE data of the first 2 months over Africa—still, interpretation is necessary. *Top*: latitude-longitude grid (with heavy aliasing), *bottom*: geodesic grid (with rather random errors)

until we reach the desired resolution. Hierarchical sorting speeds up the program enormously. For example, instead of checking all $12 \cdot 4^4 = 3072$ faces of a plain “step 4” geodesic sphere where an entry belongs (worst case scenario), for hierarchical sorting we need to check in total only $12 + 4 \cdot 4 = 28$ faces but at different step levels.




Fig. 4 Different steps of a geodesic sphere—starting with an icosahedron (step 0) on the left and quadrupling the number of faces in each next step




Fig. 5 Data structure in our “geodesic” library

We wrote a C-library, which is based on the ideas of [4] but additionally can handle our data and the subdivisions of the icosahedron. The whole data structure is realized by pointers like the scheme in Fig. 5 illustrates. Using pointers for the data gives us the benefit that the amount of copied data is much lower (e.g. 8 bytes for a pointer in comparison to around 100 bytes per complete entry).

4 Discussion, Conclusion and Outlook

We prepared the theoretical background of Euler deconvolution with the geodetic enhancement of the Gauss–Helmert model. We also applied Euler deconvolution to non-preprocessed GOCE data and found that the results are not yet interpretable. Hence, further research is necessary for a better understanding of the data, as well as a better interpretation and understanding of the results.

For our proposed “extended Euler deconvolution” via the Gauss–Helmert model, computation time increases exponentially: The observation equations become non-linear, hence, they must be linearized. Which in turn demands iterations to solve for the unknown parameters. Additionally, for our approach a $4m + 4 \times 4m + 4$ matrix has to be inverted (m being the number of observations within one window), compared to the 4×4 matrix of the standard Gauss–Markov model approach). The matrix inversion becomes necessary in order to retrieve variance covariance information. Altogether, this leads to a computation time increase by a factor of 60.

As the solutions for every Euler deconvolution window can be computed completely independently, we can parallelize our program and hence benefit from HPC.

At the moment, we are preprocessing the GOCE data and at the same time preparing a first implementation of our program in C as proof of concept with synthetic data.

The steps after this will be the parallelization of our program and its implementation on the CRAY XE6.

Acknowledgements The authors thank the High Performance Computing Center Stuttgart (HLRS) for the opportunity to use their computing facilities; furthermore, we gratefully acknowledge the helpful technical support. This work was supported by the European Space Agency (ESA) within the project UWB/GOCE-GDC – ITT AO/1-6367/10/NL/AF, STSE-GOCE+, Theme2.

References

1. Blakely RJ (1995), “Potential Theory in Gravity and Magnetic Applications”, Cambridge University Press.
2. CODATA (2010). Recommended values. Ed. by National Institute of Standards and Technology. URL: <http://physics.nist.gov/cgi-bin/cuu/Value?bg>, retrieved: 10 April 2012.
3. ESA (2011) GOCE website: <http://www.esa.int/goce>, retrieved: 10 April 2012.
4. Mantyla M (1988), “An introduction to solid modeling”, Principles of computer science series Vol. 13, Computer Science Press.
5. Marson I and Klingelée EE (1993), “Advantages of using the vertical gradient of gravity for 3-D interpretation”, *Geophysics*. Vol. 58(11), pp. 1588–1595.
6. Mushayandebvu MF, Lesur V, Reid AB and Fairhead JD (2004), “Grid Euler deconvolution with constraints for 2D structures”, *Geophysics*. Vol. 69(2), pp. 489–496.
7. Peil R, Bruinsma S, Migliaccio F, Förste C, Goiginger H, Schuh WD, Höck E, Reguzzoni M, Brockmann JM, Abrikosov O, Veicherts M, Fecher T, Mayrhofer R, Krasbutter I, Sansò F, Tscherning CC (2011), “First GOCE gravity field models derived by three different approaches”, *Journal of Geodesy*. Vol. 85, pp. 819–843.
8. Reid AB, Allsop JM, Granser H, Millett AJ and Somerton IW (1990), “Magnetic interpretation in three dimensions using Euler deconvolution”, *Geophysics*. Vol. 55(1), pp. 80–91.
9. Roth M (2009), “Marine full tensor gravity gradiometry data analysis and Euler deconvolution”, study thesis, Institute of Geodesy, University of Stuttgart.

10. Thompson DT (1982), "EULDPH: A new technique for making computer-assisted depth estimates from magnetic data", *Geophysics*. Vol. 47(1), pp. 31–37.
11. Wu P (2002), "Euler Deconvolution". URL: <http://www.geo.ucalgary.ca/~wu/Goph547/EulerDeconvolution.pdf>, retrieved: 10 April 2012.
12. Zhang C, Mushayandebvu MF, Reid AB, Fairhead JD and Odegard ME (2000), "Euler deconvolution of gravity tensor gradient data", *Geophysics*. Vol. 65(2), pp. 512–520.

Parameterization of Threshold Accepting: The Case of Economic Capital Allocation

H.-P. Burghof and J. Müller

Abstract The corresponding model of economic capital allocation was described as a mixed integer nonlinear program (MINLP) in Burghof and Müller (Allocation of economic capital in banking: a simulation approach. In: Nagel WE, Kröner DB, Resch M (eds) High performance computing in science and engineering. Springer, 2011). In this context an appropriate solving algorithm in form of threshold accepting was introduced. Now we address the algorithm's parameterization which is influenced by the model's specific implementation of threshold accepting. Without discussing the implementation in detail, an impression of the parameterization's handling and robustness in context with the model is given. On the basis of adequate parameterization the model provides a first indication of optimal economic capital allocation's superiority compared to alternative allocation methods.

1 Introduction

The model addresses the allocation process of economic capital in banks. Financial institutions hold economic capital in order to cushion unexpected losses and to prevent bankruptcy. Thereby, the capital at the same time mitigates the risk of contagion in the financial system.

The corresponding model of economic capital allocation was described as a mixed integer nonlinear program (MINLP) in [1]. In this context an appropriate solving algorithm in form of threshold accepting was introduced. Now we address the algorithm's parameterization which is influenced by the model's specific implementation of threshold accepting. Without discussing the implementation in detail, an impression of the parameterization's handling and robustness in context with

H.-P. Burghof (✉) · J. Müller

Lehrstuhl für Bankwirtschaft und Finanzdienstleistungen, Universität Hohenheim, 70599
Stuttgart, Germany

e-mail: Burghof@uni-hohenheim.de; Jan.Mueller@uni-hohenheim.de.

the model is given. On the basis of adequate parameterization the model provides a first indication of optimal economic capital allocation's superiority compared to alternative allocation methods.

Before concentrating on parameterization issues, we classify the underlying optimization problem. In this context we give an exemplary visual proof for the non convexity of the model cases' solution spaces. Furthermore, important parameters of threshold accepting are introduced. After describing the parameterization and the indication of the optimal allocation's superiority to alternative allocation methods, the last part provides technical information on typical computations occurring in context with actual and future analyses on the basis of the model.

2 Classification of the Underlying Optimization Problem

The model's underlying optimization problem belongs to the class of MINLPs. Similar to portfolio optimization, the present optimization of economic capital allocation has to face non linearity stemming from diversification effects among the securities' returns. Furthermore, the model features the consideration of influences originating from the decision makers' acting. Binary variables control value at risk (VAR) constraints.¹ These integer constraints cause the non convexity which represents the actual challenge of the present optimization.

An exemplary visual impression of the model's solution space surface gives Fig. 1 on the basis of three VAR-limits and decision makers² respectively. The x- and y-axis exhibit controlled limit quantities vl_1 and vl_2 concerning business unit one and two building a 142/120-grid through increasing the limits stepwise by 1.57 EUR. In contrast, the size of the third unit's limit vl_3 is adjusted in order to maximize the expected profit $\mu_{\text{bank}}(\mathbf{vl})$.

The example uses a total VAR-limit of $vl_{\text{bank}} = ec = 1,000$ EUR which corresponds at the same time to the bank's total amount of economic capital available. The only constraints therefore represent vl_{bank} , vl_j and the budget of investment capital $c_{\text{bank}} = 150,000$ EUR.³ Under these conditions assigning VAR-limits $vl_1 = 546.1$ EUR, $vl_2 = 393.4$ EUR and $vl_3 = 615.9$ EUR to the units induces the bank's maximum expected profit of $\mu_{\text{bank}}^*(\mathbf{vl}) = 55.4$ EUR.

In the example probabilities of success of $p_1 \approx 0.55$, $p_2 \approx 0.54$ and $p_3 \approx 0.56$ characterize the business units. Furthermore, the geometric Brownian motion (GBM) provides the multivariate returns used for modeling each business unit's

¹See binary variables $b_{\text{bank},i}$ and $b_{j,i}$ in context with the description of the model's MINLP in [1], pp. 542–543.

²The terms “decision maker” and “business unit” are used synonymously in the following.

³See Burghof and Müller [1], p. 542 for a notation of the used constraints.




Fig. 1 Extract from an exemplary solution space surface generated by the model

traded security.⁴ The corresponding standard deviations and correlations base on the return samples of three arbitrarily chosen S&P 500 stocks.⁵ The business units’ and securities’ characteristics enter the Monte Carlo trading simulation with 20,000 iterations which generates the bank’s profit and loss distribution $F_{\text{bank}}(pl)$ required for the valuation of each occurring limit allocation.

The surface of the resulting solution space exhibits many local extreme points. For finer grid structures, the consistency of the surface will change slightly but basically keep its many peaks and valleys. Pure gradient based optimization methods which first of all orientate themselves by the steepest ascent next to them will inevitably get stuck in the next best extreme point. In order to solve correspondent global problems, an algorithm has to be able to escape from extreme points which do not represent the global optimum. Furthermore, it has to feature probabilistic procedures in order to achieve a reasonable coverage of the entire solution space. Therefore, closed form solutions have to be combined with heuristics or even pure heuristic methods have to be applied. One appropriate heuristic method to solve the present financial optimization problem is threshold accepting.⁶

⁴However, the implementation of the model does not require any assumptions concerning the distribution of the returns. The model enables the use of historical data in combination with bootstrapping methods. It considers skewness and kurtosis if existent in the data.

⁵We use the stocks MMM, ABT and ANF with returns from August 9, 2010 to August 5, 2011.

⁶See e.g. [5–7] for introductions to threshold accepting and issues on heuristic optimization in finance in general.


```

1: parameterization
2: for  $i = 1$  to  $n_{\text{restarts}}$  do
3:   random current solution  $\mathbf{vl}^c$ ,  $\mathbf{vl} = vl_1, \dots, vl_{n \text{ units}}$ 
4:   for  $j = 1$  to  $n_{\text{rounds}}$  do
5:     for  $k = 1$  to  $n_{\text{steps}}$  do
6:       generate  $\mathbf{vl}^n \in N(\mathbf{vl}^c)$  and compute  $\Delta = \mu_{\text{bank}}(\mathbf{vl}^n) - \mu_{\text{bank}}(\mathbf{vl}^c)$ 
7:       if  $\Delta > t_j$  then  $\mathbf{vl}^c = \mathbf{vl}^n$ 
8:     end for
9:   end for
10:   $\mu_{\text{bank}, i} = \mu_{\text{bank}}(\mathbf{vl}^c)$  and  $\mathbf{vl}_j = \mathbf{vl}^c$ 
11: end for
12:  $\mathbf{vl}^* = \mathbf{vl}_j$  with  $j \mid \mu_{\text{bank}, j} = \max \{ \mu_{\text{bank}, 1}, \dots, \mu_{\text{bank}, n \text{ restarts}} \}$ 

```

Fig. 2 Threshold accepting pseudo code with key parameters

3 Key Parameters of Threshold Accepting

The functional principle of threshold accepting was already described in [1] on the basis of portfolio optimization under a VAR-constraint. Except for its final implementation, its basic principle stays the same with economic capital allocation optimization. The pseudo code gives an overview concerning the different parameters and parts of threshold accepting requiring specific setting (Fig. 2).⁷

There have to be determined the number of restarts n_{restarts} , the number of rounds n_{rounds} responsible for the number of used thresholds, the number of optimization steps n_{steps} , the neighborhood function $N(\mathbf{vl}^c)$ and the threshold sequence \mathbf{t} . Since the last two are less intuitive we introduce them more detailed by the following. Due to differences between the implementation of threshold accepting in context with portfolio optimization under a VAR constraint and economic capital allocation, we finally point out differences of the two fields of application of threshold accepting in finance.

The neighborhood function $N(\mathbf{vl}^c)$ defines how a new solution \mathbf{vl}^n is generated from the current solution \mathbf{vl}^c .⁸ In the model, an amount of one of two randomly chosen variables, here VAR-limits, is transferred to the other. The resulting solution is adjusted, if constraint violations occur. Alternatively, penalty terms could enforce the adherence to the constraints.⁹ Unnecessary strong modification of the current solution \mathbf{vl}^c hinders the successive development of the solutions across the iterations and should therefore be prevented.

However, the possible implementations of the neighborhood function $N(\mathbf{vl}^c)$ are various since the implementation depends on the context of the application of threshold accepting. The case of portfolio optimization under a downside risk

⁷See also e.g. [4], p. 9 for threshold accepting pseudo code.

⁸See e.g. [4], pp. 8–9 concerning neighborhood function issues.

⁹See [5], p. 22 for remarks on the usage of penalty terms.

constraint with the focus on further integer constraints, for example, requires the transfer of a constant percentage value.¹⁰ In contrast, investigations with the present model clearly suggest a non percentage transfer value drawn from a particular interval. We successively exponentially decrease this interval to zero across the iterations. Tests with percentage values remained unsatisfying. In order to ensure the quality and the efficiency of threshold accepting, we jointly calibrate the upper bound of the sequence's initial interval q with quantile p determining the starting point of threshold sequence \mathbf{t} which we introduce in Part IV.

The thresholds \mathbf{t} determine whether a current interim solution represents an appropriate basis for the generation of a new solution by the neighborhood function $N(\mathbf{v}|\mathbf{f})$.¹¹ The data for the threshold sequence \mathbf{t} is derived from a randomly initialized sequence of neighborhood solutions and their valuations' deltas' absolute values. Subsequently we choose an interval up to a certain quantile p from the resulting values' empirical distribution. The final n_{rounds} thresholds represent equidistant draws from the interval.¹² The corresponding procedure has the advantage of being data driven. Therefore, the thresholds orientate themselves by absolute deltas likely to occur in context with the current neighborhood function implementation which improves the convergence of the threshold accepting.

As already mentioned, an appropriate implementation of the neighborhood function strongly depends on the exact underlying application of threshold accepting. There are fundamental differences between the application in context with the referred to case of portfolio optimization under a VAR constraint and the current optimal economic capital allocation.

The portfolio optimization generates long and short position weights which are consequently restricted to 100%. In contrast, the economic capital allocation generates VAR-limits allowing decision makers to build long and short positions autonomously under the condition of full limit usage. Thereby, the exact budget these limits can be generated from is unknown until the optimization ends.

Finally, there can be the interpretation of a financial institution as one huge portfolio in order to be able to apply the common portfolio optimization under a downside risk constraint for economic capital allocation issues. Indeed, at first a position weight from portfolio optimization can be easily transformed into its corresponding VAR-limit. However, an autonomous decision maker provided with the limit would just be able to build the original position the VAR-limit was transformed from. Variation potentially causes violations of the total risk limit of the institution. Furthermore, variations are very likely to prevent the implementation of the highest profit expectations originally worked out by the portfolio optimization. These shortcomings of the VAR-limits from portfolio optimization originate from their lack in consideration of correlations' instability resulting from the autonomous decision making of the business units.

¹⁰See [5], p. 21 suggesting a percentage transfer value of 0.2%.

¹¹See [4], p. 9 introducing data-driven determined threshold sequence.

¹²Thereby, it might be advisable to choose the threshold of the last round being equal to zero.

Furthermore, implementation differences partly originate from the fact that the current model does not focus on further integer constraints since e.g. infinite divisibility of stocks is assumed. However, focusing on more realistic integer constraints with the current model would distract attention from the actual objective of investigating optimal economic capital allocation itself.

4 Exemplary Case of Parameterization

As an example the case from the surface diagram displayed in Fig. 1 is reused. Since threshold accepting's parameters depend on each other the determination of optimal values can be difficult. However, there are findings that accuracy of threshold accepting behaves relatively robust concerning the variation of at least some of them. In order to reduce the complexity, simple grid searches concerning certain pairs of parameters are performed. One pair consists of n_{restarts} and n_{steps} , the other of neighborhood function's q and threshold sequence's p . Among these pairs, the first one offers a save opportunity to at first set its values just high enough to prevent hindering optimal parameterization concerning the other pair. If the underlying optimization problem exceeds the available computational resources from the start, the highest possible levels for n_{restarts} and n_{steps} are recommended.

In order to get an idea which values should be set preliminarily for n_{restarts} and n_{steps} there can be performed a rough grid search to narrow down their possible preliminary parameterizations. During this search we keep the neighborhood function's q and the threshold sequence's p unmodified. They should be initialized through moderate values from the middle of the corresponding intervals in order to prevent a complete inappropriate parameterization. Their potential intervals' description below will give an idea of how to set reasonable dummy values. Furthermore we keep $n_{\text{rounds}} = 10$ since the relevant literature describes the parameter's variation as less decisive on the quality and the efficiency level of threshold accepting if $n_{\text{iters}} = n_{\text{steps}} \times n_{\text{rounds}}$ stays within a certain range.¹³

Results from the rough grid search for setting preliminary values for n_{restarts} and n_{steps} are displayed by Fig. 3. We find that $n_{\text{restarts}} = 100$ most certainly represents a sufficiently high rate since already optimizations using only $n_{\text{restarts}} = 50$ achieve no worse results. Additionally, the variation of n_{steps} hardly influences the quality of the found maximum values for μ_{bank} . However, the present extreme case of a very small model bank with only 3 business units almost consistently exhibits higher results for the lower $n_{\text{steps}} = 500$. There are only little restarts with $n_{\text{steps}} = 1,000$ achieving very good results and closing up to the results of $n_{\text{steps}} = 500$. The following test reveals that a higher number of business units leads to the expected superiority of higher iteration rates.

¹³See [6] for remarks on the required number of rounds n_{rounds} .




Fig. 3 Expected profits $\mu_{\text{bank}}(\mathbf{v}_i)$, $i = 1, \dots, n_{\text{restarts}}$, ordered by size for a model bank with 3 business units




Fig. 4 Expected profits $\mu_{\text{bank}}(\mathbf{v}_i)$, $i = 1, \dots, n_{\text{restarts}}$, ordered by size for a model bank with 50 business units

The case with 50 business units from Fig. 4 confirms the generally expected higher performance of high values for n_{iters} and n_{steps} respectively. Furthermore, the results fit and confirm the conclusion that 100 restarts can be assumed being more than sufficient to enable finding best solutions. This is because again already with $n_{\text{restarts}} = 50$ the best solution lies only marginally below the best one from the case where $n_{\text{restarts}} = 100$ is used. However, since the focus is on the special

case from Fig. 3 with only 3 business units the data suggests to set the preliminary values $n_{\text{restarts}} = 100$ and $n_{\text{steps}} = 500$. Certain verification could be done through performing the grid search once more under the use of moderately changed values for q and p .

After laying optimal foundations, the parameterization concerning neighborhood function $N(\mathbf{v}^c)$ and threshold sequence \mathbf{t} can now be performed through setting corresponding initialization values q and p . In case of the threshold sequence initialization value p represents a quantile and therefore is derived from the interval $[0, 1]$. In contrast, as mentioned above, the determination of neighborhood solutions bases on a sequence of absolute transfer values randomly drawn from a particular exponentially decreased interval. Initialization value q in this case denotes the upper bound of the corresponding sequence's first interval whereas zero denotes the lower bound. The interval $[0, \text{maximum VAR-limit}]$ defines the potential upper bounds.¹⁴ The maximum VAR-limit is the maximum limit allowed to be assigned to a single business unit and therefore is at the same time the maximum transfer value possibly occurring. In the present example 1,000 EUR determines the maximum VAR-limit meaning that there is virtually no maximum VAR-limit since this amount corresponds to vl_{bank} . In cases where each business unit is provided with a different maximum VAR-limit, consequently the highest of occurring limits determines the upper bound.

Figure 5 represents the findings of a 28/28-grid search on the described intervals of potential values for q and p . Each grid point was computed by using $n_{\text{restarts}} = 100$, $n_{\text{steps}} = 500$ and $n_{\text{rounds}} = 10$. While Fig. 5 exhibits the best solutions out of n_{restarts} Fig. 6 bases on the corresponding average values. The latter diagram is particularly helpful for the identification of the efficient parameter combinations which are promising concerning reductions of n_{restarts} , n_{steps} and therefore runtime without (too high) losses in solution quality.

The analysis of both diagrams roughly suggests choosing initialization values q and p corresponding to $500 \text{ EUR} < q < 1,000 \text{ EUR}$ and $0 < p < 1$.¹⁵ With $n_{\text{restarts}} = 100$ and $n_{\text{steps}} = 500$ finally any value combination is capable of finding very good solutions as displayed by Fig. 5. From efficiency aspects however there should be chosen $q > 500 \text{ EUR}$ suggested by Fig. 6. Combinations where $q < 500 \text{ EUR}$ can be expected to tolerate less reduction of n_{restarts} and n_{steps} until their best solutions suffer severe decline.

Let us assume the chosen parameters are $q = 750 \text{ EUR}$ and $p = 0.5$. Under the use of this parameterization now n_{restarts} , n_{steps} and the runtime could be reduced through a similar grid search in order to enhance efficiency. Since $n_{\text{restarts}} = 50$ and $n_{\text{steps}} = 500$ already proved appropriateness by Fig. 3, the investigation of lower

¹⁴In case there is also a minimum VAR-limit, maximum minus minimum VAR-limit represents this interval's upper bound.

¹⁵A very close look at the data actually suggests $0 < p < 0.5$ to further enhance the probability of finding very best solutions. Nevertheless, such high accuracy level is not required in context with the present application of threshold accepting.




Fig. 5 Findings for $\mu^*_{\text{bank}}(\mathbf{vl})$ of the parameterization grid search for initialization values q and p of neighborhood function $N(\mathbf{vl}^c)$ and threshold sequence \mathbf{t}




Fig. 6 Averages of $\mu_{\text{bank}}(\mathbf{vl}_i)$, $i = 1, \dots, n_{\text{restarts}}$, of the parameterization grid search for initialization values q and p of neighborhood function $N(\mathbf{vl}^c)$ and threshold sequence \mathbf{t}

value combinations e.g. on a 4/4-grind for a rough assessment appears advisable. In order to compare the performance of value combinations, now also the runtime has to be measured. The findings concerning n_{restarts} and n_{steps} depend on the number of used processors. Under the objectives of minimization of overall runtime

and simultaneous keeping of the solution quality the cluster size determines the advantageousness of reducing n_{restarts} in favor of n_{steps} and vice versa.¹⁶

Furthermore, there might be the requirement for the parameterization to fit a wider scope of possible model cases, e.g. different numbers of business units or different constraints. However, the described parameterization procedure's recommendations for n_{restarts} and n_{steps} can just be expected to be robust concerning slight variations from the model case which was originally parameterized for.

Increasing the number of business units increases the number of possible solutions exponentially and therefore increases the number of required restarts n_{restarts} and/or steps n_{steps} . If 3 as well as 200 business units belong to the scope of possible optimizations, a new parameterization procedure had to be performed (automatically) in front of every optimization. In this case the procedure had to slim down to become proportional in computational effort compared to the actual optimization without too much loss in the parameterization's quality.

Alternatively, particularly n_{restarts} and n_{steps} could be chosen just high enough so that most complex as well as simpler model cases could be faced by the same parameterization without being concerned about unacceptable low quality outcomes. The latter is accompanied by noticeable losses in the computations' overall efficiency but is safe concerning the solutions' quality and convenient to implement, sufficient computational resources provided.

In contrast to n_{restarts} and n_{steps} , there is the suspicion that the recommendation for q and p is quite robust concerning the variation of the model case, especially for p . The parameterization of q is expected to be at least robust if the interval for potential q -values reflects the scope of application. If there is the plan to process many different model cases under the use of the same q -value, its parameterization should consequently orientate by the highest used maximum VAR-limit among the occurring model cases.

5 Superiority of Optimal Economic Capital Allocation

The superiority of the optimal economic capital allocation is analyzed by comparing optimal allocation with two alternative allocation methods.¹⁷ We expect all three tested methods' performances to depend on the model bank's level of diversification. Therefore, we perform economic capital allocations for 10 different model banks with 10, 20, ..., 100 business units. The parameterization consists of $n_{\text{restarts}} = 100$, $n_{\text{steps}} = 1,000$, $n_{\text{rounds}} = 10$, $q = 750$ EUR and $p = 0.5$ which proved its appropriateness in the previous paragraph.

For a description of the optimal allocation method we refer to [1]. The simpler one of the alternative methods allocates the economic capital uniformly among the

¹⁶See [6], pp. 32–33 addressing cluster size issues in context with threshold accepting.

¹⁷We refrain from presenting more alternative methods here, to keep the remarks short.




Fig. 7 Indication of superiority of optimal economic capital allocation for differently diversified model banks and number of business units respectively

business units and exclusively among those exhibiting positive return expectations. The more sophisticated alternative uses the units’ weighted return expectations in order to derive limit assignments whereas the method also exclusively considers units with positive return expectations.

The total amounts of economic capital and investment capital remain constant at $e c_{\text{bank}} = v l_{\text{bank}} = 1,000$ EUR and $c_{\text{bank}} = 150,000$ EUR for every method and bank size to provide equal opportunities. The business units’ characteristics include certain probabilities of success and traded securities’ returns similar to the case from Fig. 1.¹⁸ These parameters enter a Monte Carlo trading simulation with 20,000 iterations.

While the left diagram of Fig. 7 displays the expected profits of the different methods and model banks, the right one plots the corresponding risk in form of value at risk (VAR). The results confirm the superiority of the comprehensive optimal allocation. Less sophisticated allocation methods are incapable of keeping up the exposure to risk across the stronger diversified banks entailing decreases in profit expectations. Across the displayed unrealistically low diversified banks, the optimal allocation only slightly outperforms the alternatives.

The alternative method orientating by the weighted business units’ profit expectations at least clearly outperforms the method which just assigns uniform limits. Both alternatives leave a noticeable part of the bank’s total VAR-limit $v l_{\text{bank}} = 1,000$ EUR unused which is inefficient.

6 Technical Information on Typical Computations

Typical computations result particularly from analyses where model parameters (not threshold accepting parameters) are varied in order to learn about the parameters’ impact on optimality of the economic capital allocation. A small exemplary case was represented in the previous Part V. Corresponding analyses require the computation of a new optimal economic capital allocation per variation. As a result typical computations consist in sequences of optimizations. In cases where several

¹⁸Exact characteristics of business units are irrelevant and therefore remain unlisted.

Table 1 Technical information on typical computations with the NEC Nehalem cluster

Computations	A ^a	B ^b
Model cases	10 ^c	1
$n_{\text{restarts/Jobs}}$	1,000	78,400
n_{rounds}	10	10
n_{steps}	1,000	500
p	0.5	0.5
q	750 EUR	750 EUR
Sequential	269.7 h	1,071 h
Parallel	8.7 h	4.2 h
Nodes	4	32
Processors	32	256
Servants	31	255

^aSuperiority of economic capital allocation

^bThreshold accepting parameterization

^cTen different bank structures with 10–100 business units

parameters are varied simultaneously the number of required optimizations per computation increases exponentially.

For example, we analyze the impact on the superiority of optimal allocation for different levels of information of the model bank's central management. The different levels result from varying the comprehensiveness of the management's Bayesian learning concerning the characteristics of the decision makers.¹⁹ A follow-on investigation, for example, additionally includes a further decision makers' property of herding and entering informational cascades which requires a rerun of all computations.²⁰ Finally, there are computations in context with threshold accepting parameterization issues according to paragraph 4. Table 1 gives some technical information about exemplary computations.

With the underlying implementation of threshold accepting one job corresponds to one restart. This job size represents the smallest possible without using inter-servant communication. So far, further decomposition remained disregarded. The implementation of decomposition requiring inter-servant communication would be challenging and not necessarily very advantageous because of potential parallel overhead through increased communication and idle times. The so far used master-and-servant-structure is a very classical one outlined by Fig. 8.

The experienced parallel overhead in context with the current implementation is marginal and insignificant compared to the runtime reductions gained from parallelization. Dynamic load balancing minimizes idle times through allocating the jobs efficiently among servants by sending a new job as soon as a servant has

¹⁹See [2], pp. 205–210 for the modeling of a Bayesian learning central planner introduced in context with a preliminary stage of the current model.

²⁰See [3] for issues on herding and informational cascades in context with economic capital allocation.




Fig. 8 Applied classical master-and-servant structure for parallel computing

finished one. We use synchronized communication between master and servants through blocking send and receive. The servants pass their output data to the master who selects and rearranges it immediately. As soon as the master has received the data of the last active servant he starts certain follow-on analyses before the computation ends. The model is written in C++ whereas the parallel computing parts are implemented on the basis of the Intel message passing interface (IMPI). For compilation the Intel compiler is used.

7 Conclusions

The present model of optimal economic capital allocation entails an effective implementation of the threshold accepting heuristic. Thereby, the appropriate solving concerning the model's underlying non convex optimization problem is enabled. Implementation of threshold accepting in context with the model requires certain modifications also affecting the heuristic's parameterization. Therefore, the parameterization's requirements were examined in order to ensure the validity of the model analyses' results. Besides the implementation of the heuristic itself the model's effectiveness stems from its ability to comprehensively apply parallel computing in order to obtain problem solutions.

The so far analyses indicate the superiority of the optimal economic capital allocation compared to the alternative allocation methods within the framework of the model. Thereby, the difficulties of economic capital allocation, for example in form of unstable correlations stemming from the decision makers' autonomous decision making, are considered.

Further analyses on the superiority of optimal economic capital allocation will address the case of an uninformed central management. During these investigations the central management uses Bayesian updating to learn about the characteristics of the decision makers and their trading activities. Another field of investigation will be the consideration of correlated decision making in context with herding and informational cascades and the corresponding consequences for optimal economic capital allocation.

References

1. Burghof, H.-P. and Müller, J. (2011). Allocation of Economic Capital in Banking: A Simulation Approach. In: High Performance Computing in Science and Engineering (eds. W. E. Nagel, D. B. Kröner, M. Resch). Springer.
2. Burghof, H.-P. and Müller, J. (2009). Allocation of Economic Capital in Banking: A Simulation Approach. In: Handbook of Value at Risk (ed. G. N. Gregoriou). McGraw-Hill.
3. Burghof, H.-P. and Sinha, T. (2005). Capital allocation with value-at-risk - the case of informed traders and herding. *Journal of Risk* 7(4), 47–73.
4. Gilli, M., Kellezi, E. and Hysi, H. (2006). A data-driven optimization heuristic for downside risk minimization, *Journal of Risk* 8(3), 1–19.
5. Gilli, M. and Schuman, E. (2010). Heuristic Optimization in Financial Modelling. Workingpaper. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1277114. Cited 20 Jul 2012.
6. Gilli, M. and Winker, P. (2008). A review of heuristic optimization methods in econometrics. Swiss Finance Institute, Research Paper Series No. 12. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1140655. Cited 21 Jul 2012.
7. Maringer, D. (2005). Portfolio management with heuristic optimization (eds. H. Amman and B. Rustem). Springer.

Distributed FE Analysis of Multiphase Composites Regarding 3D Elasticity Problems

Kai Schrader and Carsten Könke

Abstract Today the numerical simulation of damage effects in heterogeneous materials is done by adaption of multiscale approaches. A consistent modeling in three dimensions with a high discretization resolution on each scale based on a hierarchical or concurrent multiscale model still has issues. The algorithms have to be optimized in regards to the computational efficiency and the distribution among available hardware resources often based on parallel hardware architectures. In the last 5 years high performance computing (HPC) as well as GPU computation techniques were established for investigation of scientific aims. Consequently, in this work substructuring methods for partitioning of FE meshed specimens were implemented, tested and adapted to the HPC computing framework using several hundred CPU nodes. An memory-efficient iterative and parallelized equation solver combined with a special preconditioning technique for solving the underlying equation system was modified and adapted to the consideration of combined CPU and GPU based computations.

1 Introduction

Modern approaches of discretization methods such as the finite element method (FEM, [1]) approximates partial differential equations which have to be solved numerically. Today in (material) engineering science, many investigations of material behaviour in 3D, e.g. the damage initiation and propagation at different scales,


K. Schrader (✉) · C. Könke
Institute of Structural Mechanics, Bauhaus-University Weimar, Marienstrasse 15,
D-99423 Weimar, Germany
e-mail: kai.schrader@uni-weimar.de; carsten.koenke@uni-weimar.de

are based on complex and computationally expensive numerical simulations. Life-time assessment of engineering structures is relying on sophisticated material models, integrating all different aspects of damage initiation and deterioration over the expected life-time of a structure. Therefore, the current material models in engineering applications are integrating modern approaches from material science via multiscale methods. Especially for heterogeneous materials, these multiscale approaches allow a detailed insight into the material physics on appropriate scales [2]. A major drawback of multiscale methods is the tremendous increase in degrees of freedom (d.o.f.s.) in the resulting equation systems when studying models on meso- or microscale [3, 4]. In the framework of damage simulations, the incremental-iterative approach requires repeated solving of the linearized equation system [5], leading to an even more crucial computing time consumption. The key idea of the approach developed in this work is the application of an iterative distributed solver technique for the solution of the approximated Navier partial equations which are arising from these computational elasticity or inelasticity multiscale problems. In recent years, the distributed computing based on the message passing interface standard (MPI), [6] has been proven, which enables the distributed computation of linear equation systems utilizing as many computational nodes which are available in a high performance computing framework [7] with the widespread establishment of enormous hardware resources. Furthermore, during the past years, hybrid CPU-GPU architectures were developed, using graphics processing units (GPU) which enable the high-scalable implementation of routines of linear algebra. Therewith, the parallel code execution using different hardware architectures (based on CPU and GPU simultaneously is possible [8]. Consequently, a memory-advantageous iterative MPI-solver strategy based on the conjugate gradient method (CG) was chosen, and accelerated by an efficient preconditioning technique. The parallelization technique is based on a standard nonoverlapping domain decomposition method for FE problems combined with the elastic-inelastic domain split. Hence, the implementation takes different parallel hardware architectures into account, as well as the application in high performance computing centers. For that reason the developed algorithms for the linearized step have been evaluated considering the hybrid CPU-GPU NEC Nehalem cluster [9] or the new petaflop system Cray XE6 at the high performance computing center Stuttgart (HLRS).

2 Continuum Mechanics

The mechanical behaviour of a material continuum discretized into infinitesimal elements with bulk material properties can be described by fundamental formulations of the continuum mechanics (Fig. 1).

Fig. 1 Illustration of the motion of a material point from the reference to the momentary position



2.1 Kinematics

Therewith, the motion of a single material point of the continuum changing the position from the reference to the momentary position can be described by the equation

$$\mathbf{x} = \phi(\mathbf{X}; t) = \mathbf{x}(\mathbf{X}) \tag{1}$$

The corresponding material deformation gradient is defined by

$$\mathbf{F} = \frac{\partial \mathbf{x}}{\partial \mathbf{X}} = \mathbf{x}_{\cdot} \tag{2}$$

which causes the transformation of a material line element $\partial \mathbf{x}$ in the reference position to a material line element $\partial \mathbf{X}$ in the momentary position by

$$\partial \mathbf{x} = \mathbf{F} \partial \mathbf{X} \tag{3}$$

in which \mathbf{F} in general is an unsymmetric tensor. The material penetrability will be avoided by introduction of the condition

$$J = \det(\mathbf{F}) > 0 \tag{4}$$

The determinant J of the deformation gradient \mathbf{F} describes the volume ratio between reference and momentary configuration of a differential volume element deformed by \mathbf{F} with

$$J = \frac{\partial V_R}{\partial V_M} \tag{5}$$

Additionally, \mathbf{F} enables the derivation of different deformation measurements. The right Cauchy-Green deformation tensor yields to

$$\mathbf{C} = \mathbf{F}^T \mathbf{F} \quad (6)$$

as well as the Green-Lagrange strain tensor (with Identity \mathbf{I})

$$\mathbf{E} = \frac{1}{2}(\mathbf{C} - \mathbf{I}) \quad (7)$$

The left Cauchy-Green deformation tensor (Finger tensor) is defined by

$$\mathbf{b} = \mathbf{F} \mathbf{F}^T \quad (8)$$

With \mathbf{b} the Hencky strain tensor can be formulated by

$$\varepsilon = \frac{1}{2} \ln(\mathbf{b}) \quad (9)$$

2.2 Stress Tensors

The Cauchy stress tensor \mathbf{T} induces a infinitesimal surface force related to the deformed momentary configuration of a surface element. Based on \mathbf{T} several stress measurements can be defined. Therewith the Kirchhoff stress tensor results in

$$\tau = \mathbf{J} \mathbf{T} \quad (10)$$

The material stress tensor is defined by

$$\mathbf{S} = \mathbf{F}^{-1} \mathbf{T} \mathbf{F}^{-T} \quad (11)$$

and consequently

$$\mathbf{T} = \mathbf{F} \mathbf{S} \mathbf{F}^T \quad (12)$$

The first Piola-Kirchhoff stress tensor yields to

$$\mathbf{T}_{P_1} = \mathbf{J} \mathbf{T} \mathbf{F}^T \quad (13)$$

as well as the symmetric second Piola-Kirchhoff stress tensor is given by

$$\mathbf{T}_{P_2} = \mathbf{F}^{-1} \mathbf{T}_{P_1} = \mathbf{J} \mathbf{F}^{-1} \mathbf{T} \mathbf{F}^T \quad (14)$$

And finally the Mandel stress tensor results in

$$\mathbf{M} = \mathbf{S} \mathbf{C} = \mathbf{F}^{-1} \mathbf{T} \mathbf{F} \quad (15)$$

2.3 Constitutive Relations of Linear Elasticity

The constitutive relation between the Green-Lagrange strain tensor \mathbf{C} and the second Piola-Kirchhoff stress tensor is defined by

$$\mathbf{T}_{p_2} = \mathbb{C}\mathbf{E} \tag{16}$$

also known as Hook’s law for linear elastic materials with \mathbb{C} as the fourth order elasticity tensor. The linearization of the Green-Lagrange strain tensor results in the linearized strain tensor ε

$$\varepsilon = \text{lin}\mathbf{E} \tag{17}$$

The linearized stress tensor σ results in

$$\sigma = \mathbb{E}\varepsilon \tag{18}$$

with \mathbb{E} as the elastic material matrix. Considering stress and strain components Eq. (18) can be written as

$$\begin{pmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \tau_{xy} \\ \tau_{xz} \\ \tau_{yz} \end{pmatrix} = \mathbb{E} \begin{pmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \varepsilon_{zz} \\ \gamma_{xy} \\ \gamma_{xz} \\ \gamma_{yz} \end{pmatrix} \tag{19}$$

For the isotropic case the material elasticity matrix considering the material properties ν (poisson’s ratio) and E (Young’s modulus) can be formulated as

$$\mathbb{E} = \frac{E}{(1 + \nu)(1 - 2\nu)} \begin{pmatrix} 1 - \nu & \nu & \nu & 0 & 0 & 0 \\ \nu & 1 - \nu & \nu & 0 & 0 & 0 \\ \nu & \nu & 1 - \nu & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1-2\nu}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1-2\nu}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1-2\nu}{2} \end{pmatrix} \tag{20}$$

2.4 Extension to Material Inelasticity

According to the classical plasticity formulation [10] the linearized strain tensor will be splitted in an elastic and an inelastic part with

$$\varepsilon = \varepsilon^{\text{tot}} = \varepsilon^{\text{el}} + \varepsilon^{\text{pl}} \tag{21}$$

The linear elastic constitutive relation regarding Hook's law of Eq. (18) is modified to

$$\sigma = E(\varepsilon^{\text{tot}} - \varepsilon^{\text{pl}}) \quad (22)$$

The evaluation of the stress state (elastic or inelastic) will be performed by using a defined yield condition $f(\sigma)$ which depends on the actual stress state σ and the material specific yield stress σ_Y as a scalar value. In general f (without consideration of hardening effects) can be expressed as

$$f(\sigma) = |\sigma| - \sigma_Y = \begin{cases} < 0 & \text{elastic} \\ = 0 & \text{inelastic} \end{cases} \quad (23)$$

Since the evaluation of the $f(\sigma)$ for multiaxial stress states of the inelastic material behaviour does not result exactly in zero, mapping techniques, such as return mapping methods, were developed for the computation of the inelastic stress and strain components. The equivalent von-Mises stress is defined as

$$\sigma_V = \sqrt{3 \cdot I_2} \quad (24)$$

Hence, considering the second invariant of the stress tensor σ (with s_{ij} as components of deviatoric part σ^{dev}) which is given with

$$I_2 = \frac{1}{2} s_{ij} s_{ij} = \tau_{xy}^2 + \tau_{yz}^2 + \tau_{zx}^2 - (\sigma_{xx} \sigma_{yy} + \sigma_{yy} \sigma_{zz} + \sigma_{zz} \sigma_{xx}) \quad (25)$$

the von-Mises yield condition for isotropic material results in

$$f(\sigma) = |\sigma| - \sigma_Y = \sqrt{3 \cdot I_2} - \sigma_Y \quad (26)$$

in which σ_Y , e.g. describes the tensile strength as a material property used in case of uniaxial tensile loading conditions. The introduction into basic constructs of continuum mechanics will be complemented by some fundamental formulations of the finite element method (FEM) given in the following section.

3 Finite Element Method

The finite element method [1] describes an approximative solution technique for partial differential equations resulting from the discretization of different physical problems. Starting with a 3D elasticity problem \mathcal{V} which is bounded by domain Ω , this problem can be described by the equilibrium equation

$$\sigma_{ij,j} + b_i = 0 \quad (27)$$

with

$$\sigma_{ij} = C_{ijkl}\varepsilon_{kl}, \quad \varepsilon_{kl} = \frac{1}{2}(u_{k,l} + u_{l,k}) \equiv \nabla_s(u)_{kl} \quad (28)$$

The boundary conditions are subjected to

$$\Omega = \Omega_u + \Omega_t, \quad \Omega_u \cap \Omega_t = \emptyset \quad (29)$$

and

$$\sigma_{ij}n_j = t_i \quad \text{on } \Omega_t \quad u_i = \hat{u}_i \quad \text{on } \Omega_u \quad (30)$$

Here σ_{ij} are the components of the stress tensor, b_i the components of the body force and n_i the unit outward normal. U_i are the displacements and C_{ijkl} is the material tensor. The weak form of Eq. (27) with zero body force is given as

$$\int_{\mathcal{V}} \delta \nabla_s(u)_{ij} C_{ijkl} \nabla_s(u)_{kl} d\mathcal{V} - \int_{\Omega} \delta u_i t_i d\Omega = 0 \quad (31)$$

4 Preconditioning Techniques for the Conjugate Gradient Method

4.1 Preface

The conjugate gradient method (CG) can be used as an iterative and distributed solution procedure [11] to solve a symmetric, regular and positiv definite system of equations [12], which are arising from the FE approximation of the discretized problem. The preconditioning technique for CG [13] is crucial according to memory demand and computing time for the preconditioner as well as the computing time for the CG iteration procedure itself to solve a linear system of equations such

$$Ku = f \quad (32)$$

where K is the global stiffness matrix, u the nodal vector of the displacements and f the nodal vector of external structural forces. In this work a parallelized version of the preconditioned conjugate gradient method (PPCG) was implemented based on nonoverlapping domain decomposition [14, 15] without building explicitly the Schur complement system [16]. Secondly, the preconditioning technique will be restricted to a scaled maindiagonal precondition technique with a special scaling parameter based on the upper or lower bound of the spectral radius of the assembled submatrices. This reduces the time for building the preconditioning matrix as well as the memory demand and, additionally, the time for the execution of matrix-vector products involving the preconditioning matrix. Finally, the sparse matrix-vector product performed for each subdomain will take into account different matrix

storage formats. Due to the nodally storage scheme of the distributed FE data a nodal compressed row storage can be used to improve the performance compared to standard coordinate storage (coo) or compressed row storage (csr or crs).

4.2 Modified and Parallelized Jacobi Preconditioning

The easiest and most memory-efficient preconditioning technique which only involves the inverted main diagonal of the corresponding matrix, which has to be preconditioned, is known as the so-called Jacobi point preconditioner. The parallelization based on the domain decomposition method simultaneously considers the main diagonal of each domain matrix. The modified preconditioning technique to obtain $M^{(j)-1}$ per domain results from the scaled main diagonal or the scaled block diagonal of the global domain matrices $K^{(j)}$:

$$M^{(j)} = \begin{pmatrix} M_{ii} & 0 \\ 0 & M_{bb} \end{pmatrix}^{(j)} = \begin{pmatrix} \text{diag}[K_{ii}] & 0 \\ 0 & \text{diag}[K_{bb}] \end{pmatrix}^{(j)} \quad (33)$$

The necessary assembly (*MPI_Allreduce*) regarding the connecting boundary b results in

$$\widetilde{M}_{bb} = \sum_{j=1}^{nd} M_{bb}^{(j)} \quad (34)$$

Inversion and modified scaling yields to

$$M^{(j)-1} = \omega \begin{pmatrix} M_{ii}^{(j)} & 0 \\ 0 & M_{bb}^{(j)} \end{pmatrix}^{-1} = \omega \begin{pmatrix} \text{diag}[K_{ii}]^{-1} & 0 \\ 0 & \widetilde{M}_{bb}^{-1} \end{pmatrix}^{(j)} \quad (35)$$

where the scaling parameter ω is been approximated by using an efficient eigenvalue strategy for the main diagonal blocks of the stiffness matrices of all subdomains.

5 Numerical Results

5.1 NEC Nehalem Cluster: Voxel-Based Microstructural Specimen

The numerical analysis, based on the implemented parallel algorithms respecting a large-scale 3D microstructural bone specimen, which is illustrated in Fig. 2 (include several millions of degrees of freedom). Table 1 contains the information about




Fig. 2 3D porous bone specimen as FE structure based on voxel data: Perspective view (*left*), load-balanced decomposed FE structure in four nodally equal-sized domains, 8.9 million d.o.f.s

Table 1 Dimensions of the decomposed FE problem: Number of coupled nodes (ncn), number of nodes (nn), number of non-diagonal nodal fe blocks and memory demand for matrix storage in gigabyte by domain j

j	ncn	nn	n nb	mmem
1	5,108	488,783	5,418,438	0.547
2	5,896	487,927	5,548,558	0.546
3	4,399	489,546	5,545,580	0.554
4	3,198	488,472	5,570,790	0.550

the dimension of the decomposed problem such as number of elements, number of nodes, number of resulting d.o.f.s, number of nodal FE blocks and memory demand. The voxel model is converted to a regular grid based on hexahedral elements with linear shape functions. The nodal partitioning [17] of the hexahedral mesh, e.g. in four subdomains, is illustrated in Fig. 2 (right). The numerical tests including the building and assembling of the global stiffness matrices as well as the computation of the equation system in parallel with the preconditioned conjugate gradient method, scaled upto 64 MPI processes (equal to the number of subdomains).

5.1.1 Benchmarking: Multiple CPU-Nodes

Figure 3 presents the scaling of the computational time in relation to the assembly of the global stiffness matrix with a increasing number of subdomains including the numerical integration of the finite elements. Figure 4 contains the total times needed for solving the equation system with the parallelized PCG method, Fig. 5 illustrates the scaling of the computational time only considering the computation of the sparse matrix-vector product during the PCG iterations, respectively. Figure 6 shows the scaling of the iterative solver excluding the times for the sparse matrix-vector

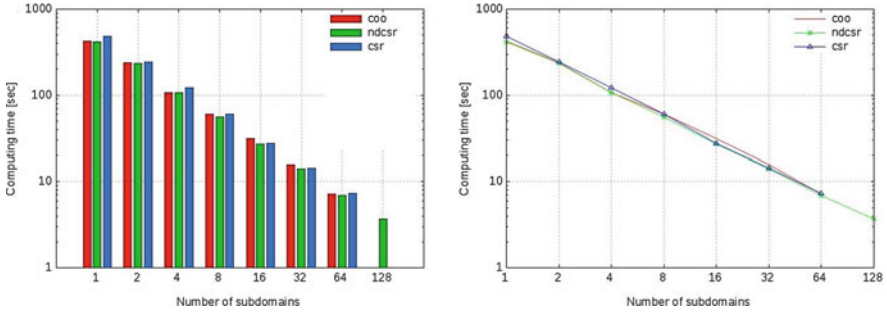


Fig. 3 NEC Nehalem cluster (CPU): Total computational time for the parallel assembling of global stiffness matrices (including the numerical integration) with increasing number of subdomains, 8.9 million d.o.f.s

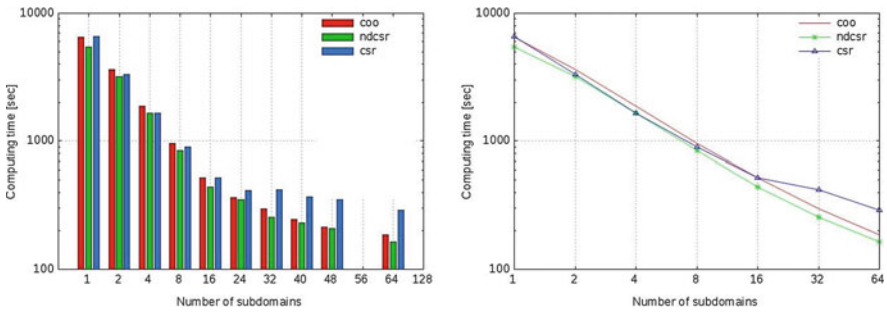


Fig. 4 NEC Nehalem cluster (CPU): Total computational time based on the parallelized preconditioned conjugate gradient method, 8.9 million d.o.f.s

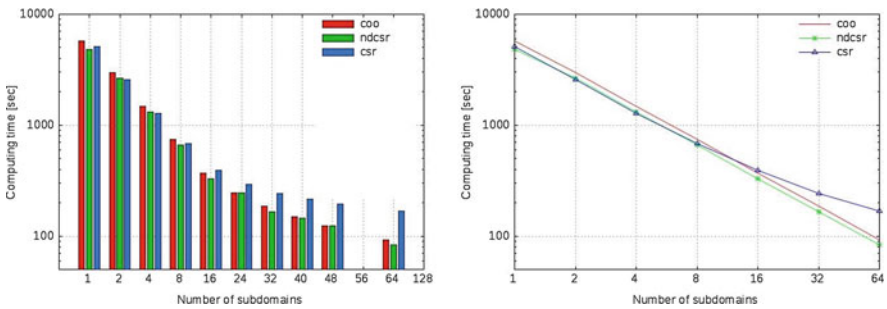


Fig. 5 NEC Nehalem cluster (CPU): Accumulated time for sparse matrix-vector operations of the PPCG method, 8.9 million d.o.f.s

product. Finally, Fig. 7 presents the computational times caused by the MPI communication overhead between different MPI processes which is mainly induced by the *MPI_Allreduce* operation during the PPCG iteration. There are moderately increased the more subdomains are being used.

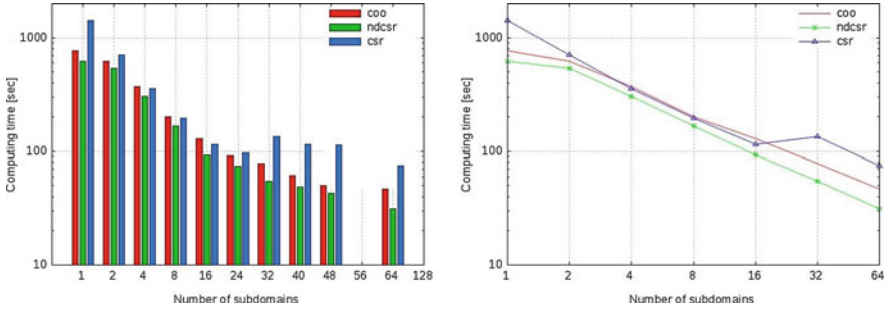


Fig. 6 NEC Nehalem cluster (CPU): Scaling of accumulated computational times for non-matrix-vector operations of the PPCG method, 8.9 million d.o.f.s

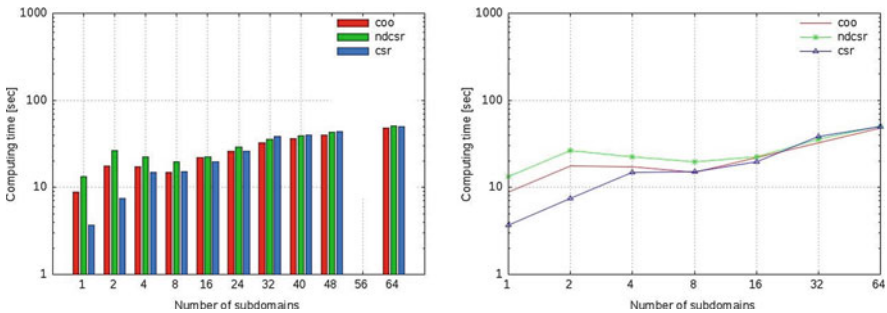


Fig. 7 NEC Nehalem cluster (CPU): Scaling of accumulated computational times for the MPI based communication of the PPCG method, 8.9 million d.o.f.s

5.1.2 Benchmarking: Hybrid Multiple CPU-GPU Nodes

Additionally, we have implemented a hybrid parallelization technique for the PCG method using the CPU and GPU in a combined fashion suitable on up to 16T nodes of the Nec Nehalem cluster. The graphical Tesla subsystem is based upon the Nvidia Tesla S1070 GPU. Each MPI thread will have access to and execute the sparse matrix-vector product of the PCG method on 1 T unit. The coo storage format for the MPI-distributed stiffness matrices of each subdomain and their allocation on the GPU has been used. Then, the subdomain matrix will be distributed over the maximum number of available GPU threads per GPU which enables a simultaneous sparse matrix-vector multiplication per matrix block and GPU thread. The results of the hybrid model are compared to the computational times for the spmv operation on the CPU (Fig. 5) nodes, as illustrated in Fig. 8 for different storage formats. and in Fig. 8 for different data transfer techniques from and to the host memory. The computational times of the hybrid model including the times for the memory transfer from CPU host to GPU device and from GPU

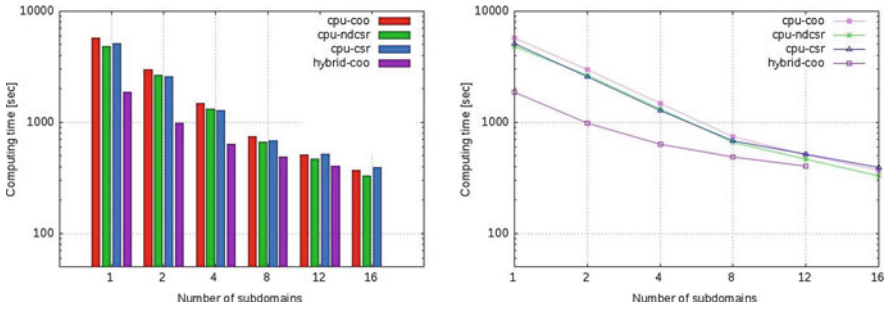


Fig. 8 NEC Nehalem cluster (CPU-GPU): Accumulated computational times for sparse matrix-vector operations of the PCG method using the CPU-only and hybrid CPU-GPU cluster, 8.9 million d.o.f.s

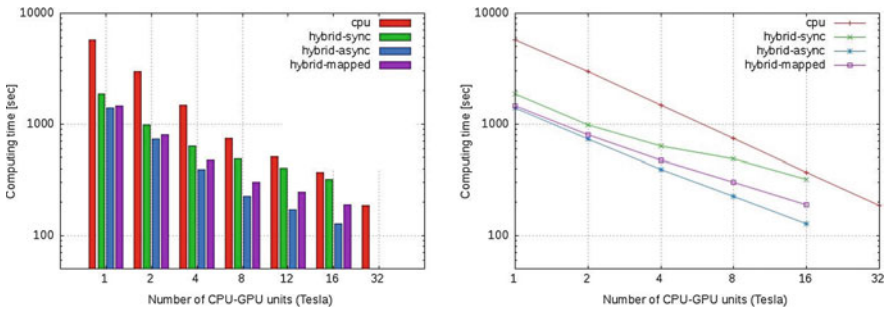


Fig. 9 NEC Nehalem cluster (CPU-GPU): Accumulated computational times for sparse matrix-vector operations of the PCG method using CPU-only and hybrid CPU-GPU cluster with synchronous (hybrid-sync), asynchronous (hybrid-async) and mapped memory (hybrid-mapped) CPU-GPU data transfer for the coo matrix storage format

device to CPU host, respectively, before and after the spmv computation. Hence, synchronous (hybrid-sync), asynchronous (hybrid-async) and mapped memory (hybrid-mapped) data transfer are compared with the computational times in respect to the cpu-only spmv computations considering the coo matrix storage format as shown in Fig. 9.

5.2 Cray XE6 Cluster: 3D Large-Scale Nickel-Alloy Specimen

The second high performance computing framework is based on the Cray XE6 cluster ‘Hermit’ at HLRS Stuttgart which is in production mode since the beginning of 2012. With the computational power of over (3,552) computing nodes, where each node consists one AMD Opteron Interlagos, it is possible to scale parallelized code up to several ten thousand cores with a peak performance in total of nearly one




Fig. 10 Multiphase Ni-alloy specimen with pore structure, voxel model based on computer-tomographic scans




Fig. 11 Cray XE6 cluster: Total computational time for the parallel assembling of global stiffness matrices (including the numerical integration) with increasing number of subdomains, 42.8 million d.o.f.s

petaflop. The AMD Interlagos processor, basically a 32-core x86-64 architecture, was introduced in 2011. The second example is based on a large-scale voxel discretization of a nickel-alloy specimen (Fig. 10). The numerical effort of this large-scale FE specimen are given with

- 14,093,177 hexaedral elements
- 14,292,274 FE nodes
- 42,876,822 nodal unknowns or d.o.f.s

The scaling starts with one single node and is increased to 256 physical cpu nodes. The following table shows the quantitative values for storing the FE data, memory demand as well as sequentially computing times which are necessary for the matrix assembly. Figures 11–14 illustrates the scaling behaviour according to FE assembly, PCG computation, matrix-vector computation, non-matrix-vector computation and MPI communication effort. Furthermore, a 4 byte data transfer for MPI communication was implemented, which improves moderately the scaling behaviour, especially if more than 64 subdomains are being taken into account (Table 2).




Fig. 12 Cray XE6 cluster: Total computational time based on the parallelized preconditioned conjugate gradient method, 42.8 million d.o.f.s

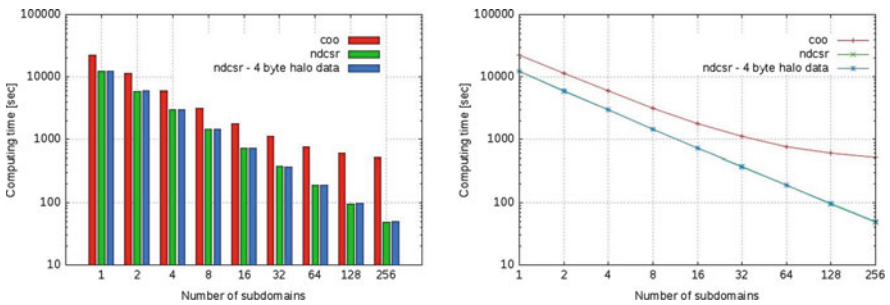


Fig. 13 Cray XE6 cluster: Accumulated time for sparse matrix-vector operations of the PPCG method, 42.8 million d.o.f.s

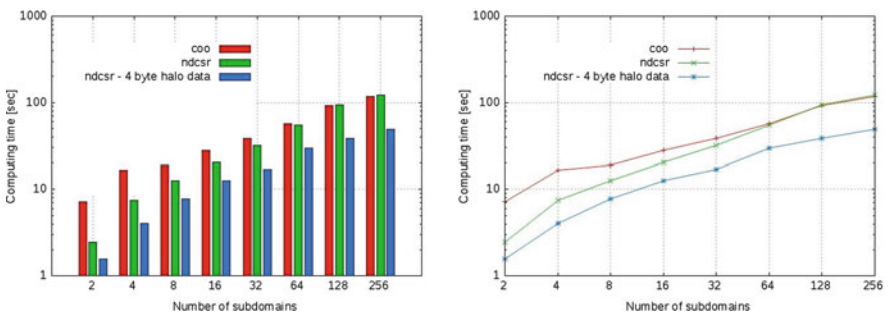


Fig. 14 Cray XE6 cluster: Scaling of accumulated computational times for the MPI based communication of the PPCG method, 42.8 million d.o.f.s

Table 2 Cray XE6 cluster: Quantitative values of total (3x3) FE blocks, the total number of block entries, memory demand and assembly time for the sequential case, 42.8 million d.o.f.s

Storage	FE blocks total	Block entries total	Memory (GB)	Seq. assembly (s)
coo	184,005,692	1,061,464,238	17.095	3,079
ndcsr	184,005,692	1,061,464,238	12.977	3,001

6 Outlook

In actual research work the computing model will be extended for the consideration of damage effects in multiphase composites. The resulting iterative analysis will be realized in a nonlinear simulation framework for high performance computers.

Acknowledgements This research is supported by the German Research Foundation (DFG), which is gratefully acknowledged by the authors.

References

1. Bathe KJ. *Finite Element Procedures*. Prentice Hall. 1995.
2. Raabe D. *Continuum scale simulation of engineering materials: fundamentals - microstructures - process applications*. Wiley-VCH. 2004.
3. Koenke C, Eckardt S, Haefner S, Luther T and Unger JF. *Multiscale simulation methods in damage prediction of brittle and ductile materials*. International Journal for Multiscale Computational Engineering vol 8(1), pp. 17–36. 2010.
4. Luther T, Koenke C. *Polycrystal models for the analysis of intergranular crack growth in metallic materials*. Engineering Fracture Mechanics vol. 76(15), pp. 2332–2343. 2009.
5. Kelley CT. *Iterative Methods for Linear and Nonlinear Equations*. Frontiers in Applied Mathematics, SIAM, Philadelphia, vol 16, 1995.
6. Gropp W, Lusk E, Skjellum A. *Using MPI: Portable Parallel Programming with the Message-Passing Interface*. MIT Press, 1999.
7. Cehavir A, Nukada A, Matsuoka A. *High performance conjugate gradient solver on multi-GPU clusters using hypergraph partitionings*. Computer Science - research and Development. vol 5, pp. 83–91. 2010.
8. Papadrakakis M, Stavroulakis G, Karatarakis A *A new era in scientific computing: Domain decomposition methods in hybrid CPU-GPU architectures*. Computational Methods Appl. Mech. Eng., accepted manuscript, 2011.
9. https://wickie.hlr.de/platforms/index.php/NEC_Nehalem_Cluster.
NEC Nehalem cluster technical description, HLR, Stuttgart, 2011.
10. Simo JC, Hughes TJR. *Computational Inelasticity*. Springer-Verlag. 1998.
11. Saad Y, van der Vorst HA. *Iterative solution of linear systems in the 20th century*. Journal of Computational and Applied Mathematics vol 123, issues 1–2, pp. 1–33. 2000.
12. Farhat C, Crivelli L, Roux F. *Extending substructure based iterative solvers to multiple load and repeated analysis*. Computer Methods in Applied Mechanics and Engineering vol 117, pp. 195–209, 1994.
13. Basermann A, Reichel B, Schelthoff C. *Preconditioned cg methods for sparse matrices on massively parallel machines*. Parallel Computing vol 23(3). 1997.
14. Toselli A, Widlund O. *Domain Decomposition Methods - Algorithms and Theory*. Springer series in computational mathematics, Springer-Verlag, Berlin, Germany, vol 34, 2005.
15. Mandel J. *Balancing domain decomposition*. Communications in Numerical Methods in Engineering vol 9, pp. 233–241. 1994.
16. Lohin D, Wathen AJ. *Schur complement preconditioning for elliptic systems of partial differential equations*. Numerical Linear Algebra with Applications vol 10(Issue 5–6), pp. 423–443. 2003.
17. Karypis G, Kumar V. *A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices*. Latest experimental release (5.0pre2), University of Minnesota, Department of Computer Science / Army HPC Research Center Minneapolis, MN 55455. 2005.

Reduction of Numerical Sensitivities in Crash Simulations on HPC-Computers (HPC-10)

Christiana Eck, Oliver Mangold, Raphael Prohl, and Anton Tkachuk

Abstract For practical application in engineering numerical simulations are required to be reliable and reproducible. Unfortunately crash simulations are highly complex and nonlinear and small changes in the initial state can produce big changes in the results. This is caused partially by physical instabilities and partially by numerical instabilities. Aim of the project is to identify the numerical sensitivities in crash simulations and suggest methods to reduce the scatter of the results. Work has been undertaken to identify sources of sensitivities through parameter studies, to improve existing mathematical formulations, e.g. of contact-impact and material models. Furthermore a tool is developed for generation of code from mathematical descriptions of finite elements with the aim to reduce effort required to make modifications of implementations of FEM models.

C. Eck (✉)

Automotive Simulation Center Stuttgart e.V., Nobelstraße 15, 70569, Stuttgart, Germany
e-mail: christiana.eck@asc-s.de

O. Mangold

High Performance Computing Center Stuttgart, Nobelstraße 19, 70569, Stuttgart, Germany
e-mail: mangold@hlrs.de

R. Prohl

Steinbeis Center of Innovation Simulation in Technology, Obere Steinbeisstr. 43/2,
75248 Ölbronn-Dürrn, Germany
e-mail: raphael@techsim.org

A. Tkachuk

Institute of Structural Mechanics, Pfaffenwaldring 7, 70550, Stuttgart, Germany
e-mail: tkachuk@ibb.uni-stuttgart.de

1 Identification of Sources of Sensitivity Through Parameter Studies

In simulation one obtain often different results for identical input values. To distinguish whether this scatter is physical or numerical nature and to eliminate or rather reduce numerical scatter is the aim of this research. To achieve this purpose parametric studies had been accomplished in the past with the commercial software packages Radioss and LS-Dyna.

Investigations with these general-purpose-programs have been expanded. In addition, further studies in Pamcrash have been conducted. As in the previous investigations the following parameters have been examined in Pamcrash: time step, mass-scaling, strain rate, rate filter, damping, contact-stiffness, stiffness, contact, element refinement, element formulation.

One can discover that in all general-purpose-programs scatter behavior is influenced by the same factors. In particular, this include damping, contact-stiffness, strain rate and rate filter.

Due to the fact that in Pamcrash a lot of parameters are predetermined, the approach of the examination has been changed. Moreover there is a given set of parameters which is based on experience. All calculations have been executed starting with this set of parameters. Following this practice the influence of the time step reduction is negligible. Moreover, the scatter is less therefor one need a measure to compare the results.

To appropriate a dimension for the scatter for each function $f_k \in f_M$, $M = \{1, \dots, n\}$, $n \in \mathbb{N}$ and $f_M : T \times F \rightarrow \mathbb{R}$, for T time and F force, in the time-force diagram the formula

$$\bar{f}_{k_l}(x_j) = \frac{1}{\max M} \sum_{i \in M} (f_{k_l}(x_j) - f_i(x_j)), \quad \forall k \in M, \forall j \in T, l \in M \quad (1)$$

has to be computed and subsequent the Euclidean norm has to be determined

$$\|\bar{f}_{k_l}\|_2 = \sqrt{\frac{1}{\#T} \sum_{j \in T} |\bar{f}_{k_l}|^2} := p_l, \quad p \in \mathbb{R}, l \in M. \quad (2)$$

These result has to be weighted in a suitable way and accordingly a median has to be calculated.

The sensitivity analysis has been accomplished on a very basic model so far. The model consists of a tube and a plate. For the sensitivity analysis several calculations with little change of initial velocities have been accomplished for various parameter settings.

Figure 1 shows a force-strain diagram. It illustrates the impact of strain rate behavior. The left diagram displays a computation with strain rate dependency and the right diagram one without strain rate dependency. Both calculations have been conducted without rate filter.




Fig. 1 Force-time diagram. Computations with strain rate- (*left*) without strain rate-dependency (*right*) and in both cases without rate filter


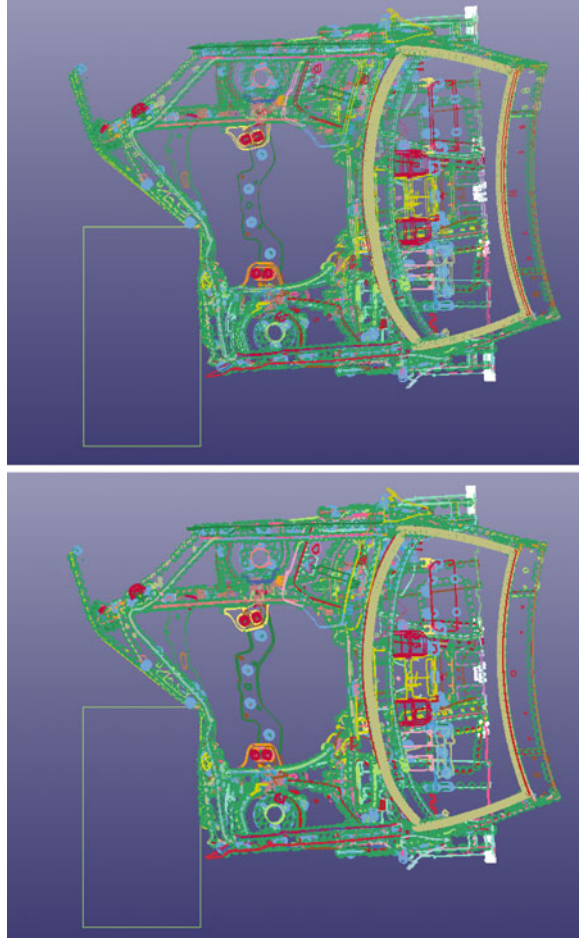


Fig. 2 Basic model computed frequently with the same parameters and the same initial velocities. The difference of the computations in detail is the distance between the plate and the tube

Moreover it was investigated whether the scatter is physical or numerical. One possibility to find out, is to compute frequently exactly the same problem. Numerical scatter would appear in each calculation, physical scatter probably not. Figure 2 shows results of several calculations performed with the same parameters and the same initial velocities. The only difference is the distance between the plate and the tube (the distance increase in each further calculation). This assures that the same problem is computed, merely time shifted. The result implies that the scatter has to be physical.

The next step will be to expand the findings of the simple model on whole- and part-vehicle model, respectively. Figure 3 shows first results of a part-vehicle model computed with under- and fully integrated elements.


Fig. 3 Part-vehicle-model computed with under- (above) and fully integrated elements



2 Improving Robustness of Contact Discretization Using Hybrid Singular Mass Matrices

Nowadays the majority of crash simulation are done using explicit finite element codes. In presence of contact-impact, buckling and non-elastic material response, explicit codes rarely fail computations. In contrast implicit codes, which rely on predictor-corrector methods, face problems with convergence of corrector step and may fail. However, further refinement of meshes in crash simulations requires very small time-steps for explicit codes. As an alternative, robust implicit space-time discretization schemes must be developed.

A novel spacial discretization of elastodynamic contact is proposed. It can be used both for bulk and thin-walled structures. The main idea of this discretization is to split dynamics and contact such that contact is collocated for a set of

Fig. 4 3-node line element

massless nodes whereas all inertia properties are condensed at the other nodes. This substantially improves conditioning of the problem reducing the differential index from 3 to 1 [2]. Usage of such mass redistribution improves robustness of time integration. It also leads to substantial reduction of oscillation in contact pressures [1, 3].

Variational framework that allows such mass redistribution relies on 3-field Hamiltonian's principle with independent variables for displacement, velocity and momentum [5, Appendix I]. Their spatial discretization together with a local elimination procedure delivers Hybrid-mixed mass matrices. If discretization spaces are specially built entries of mass matrix for certain nodes vanish. Such mass matrices are called Hybrid Singular Mass Matrices and were proposed first for triangular elements in a paper by Renard [4].

For example a 3-node line element uses corner node for contact conditions and the middle, inner node carries entire mass of the element, see Fig. 4. A simple setup for a benchmark and its results are given on Fig. 5a. A 2D Timoshenko beam hinged at both end bounces from a rigid obstacle. The quality of contact forces improve, see Fig. 5b.

Such split can be also found for 9-node quadrilateral and 27-node hexahedron elements if tensor product formula for the shape function is used, see Fig. 6. This elements are under testing.

This method can be proposed for general contact-impact problems. Deep drawing and crash of individual components may be seen as a benchmarks for study of compatibility of the approach with explicit codes. This will be covered in future work.

3 Development of Robust Algorithms in Finite Elastoplasticity

The simulations of crash tests are very sensitive numerical procedures, which require very accurate and robust algorithms in order to handle numerical instabilities. In addition to these, also physical instabilities, which may be due, for example, to the bifurcation behavior of the material under asymmetric loading, have to be detected and correctly treated by the numerical procedures in order to obtain reliable

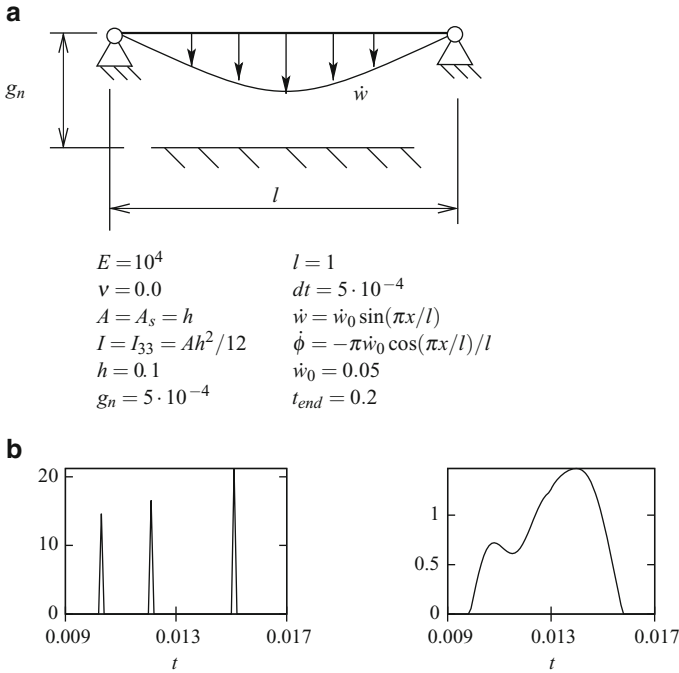


Fig. 5 Bounce benchmark for a Timoshenko beam. (a) Benchmark setup and (b) contact force at middle node computed with standard mass matrix (left) and singular mass matrix (right)




Fig. 6 9-node quadrilateral element

results. In this context, our task is to identify and eliminate numerical instabilities and improve the mathematical analysis of the physical ones.

For our purposes, we firstly formulated simplified versions of some full-vehicle-models used by automotive industries. By “simplified versions” we mean that our models are purely mechanical, take into account only a selected number of physical effects and are implemented for relatively simple (unrealistic) geometries. These preliminary steps are necessary for providing computationally cheap benchmark problems to be used for testing numerical codes and improving their functionalities, when needed. With this attitude, we focused on the stability of some algorithms

used in crash test simulations. In particular, we closely investigated the return-mapping-algorithm [7]. By reviewing this computational procedure in the context of nonlinear programming [9], we developed a new class of solution methods with better mathematical properties.

Our next step is to adopt the methods exposed in [6] in order to include large-deformation contact in our problem. After the stability of our numerical algorithms has been proven, the goal is to formulate a reliable criterion for identifying physical instabilities within a crash-simulation.

3.1 *Reduced Model for Finite Elastoplasticity*

In our mathematical model of crash tests, we consider large elastic and plastic strains as well as nonlinear material behavior. Initially we concentrate on geometric and kinematic nonlinearities as well as on nonlinear material properties and elastoplastic material behavior.

3.1.1 Balance Laws

We consider the balance of linear momentum in local, material form:

$$\begin{aligned} -\text{DIV}(FS) &= \rho_0 B & \text{in } \Omega \\ u &= d & \text{on } \Gamma_d \end{aligned} \tag{3}$$

In (3), F denotes the deformation gradient, S is the second Piola-Kirchhoff stress tensor, ρ_0 is the reference density, B is the body force per unit mass, and d represents the prescribed displacement on the Dirichlet boundary of the computational domain. For the first part of our study, inertial terms are not accounted for the force $\rho_0 B$. Apparently, dropping inertial terms in a mathematical model that should be the basis for crash test simulations may sound as a strong contradiction. However, our task here is to analyze and compare algorithms that should be applied in the numerical simulations of crash tests. In this respect, inertial terms are, at this stage, only “temporarily” switched off. This saves computational resources when different algorithms are compared, and allows for focusing on the possible numerical instabilities that, hidden behind a given algorithm, may exist independently on the consideration of inertial terms.

We restrict our investigations to a purely mechanical framework. Consequently, thermal phenomena are excluded from the outset and dissipation is expressed in terms of mechanical quantities only.

3.1.2 Constitutive Equations, Associative Flow Rule and Karush–Kuhn–Tucker Conditions

We consider the formulation of the J2 flow theory at finite strains for hyperelastic-plastic materials [7]. The constitutive equations, the flow rule and the KKT-conditions read:

$$S = 2 \frac{\partial W}{\partial C} \quad (4)$$

$$\frac{\partial \overline{C}_p^{-1}}{\partial t} = -\frac{2}{3} \gamma \cdot \text{tr}[b^e] F^{-1} \frac{\text{dev}[\tau]}{\|\text{dev}[\tau]\|} F^{-T} \quad (5)$$

$$\gamma \geq 0, f(\tau) \leq 0, \gamma f(\tau) = 0 \quad (6)$$

In (4)–(6), W is a stored energy-function, C is the Cauchy-Green strain tensor, b_e is the elastic part of the finger tensor, τ is the Kirchhoff stress tensor, f is the von-Mises flow-condition and γ is the (Karush–Kuhn–Tucker) plastic multiplier.

3.2 Numerical Methods

The numerical computations of the quasi-static case are performed by an incremental procedure, which contains a nonlinear sub-problem in every single incremental step.

3.2.1 Discretization

We obtain the discrete material model by an implicit Euler method in time for the evolution of the internal parameters, e.g. plastic strains or hardening variables. Inserting this update in the constitutive equation (4), we get a nonlinear incremental stress response, which has to satisfy the balance of momentum (3). Thus, we have to solve a nonlinear variational problem for the displacements. In order to discretize our equations in space we use a standard finite-element-method with trilinear hexahedron elements for the displacements.

3.2.2 Plasticity Algorithm

To start with the treatment of the governing equations of plasticity, we adopt the return-mapping-algorithm (RMA) [7]. As remarked in [9], this well-established procedure, which requires low computational effort, may be turn unstable because it computes stresses that do not necessarily satisfy the global equilibrium equations. This drawback can be improved by having recourse to an algorithm based on

the linearization of (4)–(6). In [9], it is shown that such a criterion leads to a computational method with higher robustness in the case of small strains. The additional iteration, introduced due to the linearization of (4)–(6), defines an algorithm, which now iterates along stresses satisfying equilibrium, until (4)–(6) are “sufficiently” fulfilled. We extended this method to the case of finite deformations.

3.2.3 Solution Methods

In our code the nonlinear variational problem is solved by a Newton method. The consistent tangent operator, introduced by Simo and Taylor in 1985, provides the basis for the linearization therein. We use a parallel multigrid solver for solving the linear sub-problems featuring in the Newton scheme [10].

3.3 Numerical Tests

The implementation of the RMA in our software framework UG (“Unstructured Grids” [8]) has been tested by some benchmark problems.

3.3.1 Benchmark Problems

As a reference for our numerical tests, we used a shear/compression-test of the unit-cube with perfect plastic behavior (Fig. 7).

Additionally, we simulated the well-documented necking of a circular bar as an example for exponential hardening behavior [7].

3.3.2 Software Framework UG

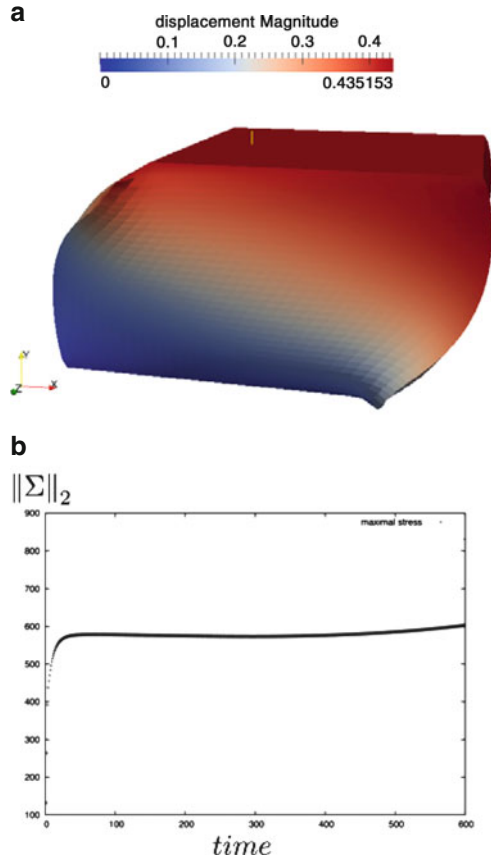
UG ([8]) is a general-purpose library for the solution of partial differential equations, which supports parallel adaptive multigrid-methods on high-performance computers. A novel implementation ensures the complete independence of grid and algebra. Cache aware storage for algebra structures and a parallel communication layer make UG4 well suited for current and next hardware-architectures.

4 Automatic Generation of Efficient Finite Element Codes

4.1 Background


Implementing finite element codes is a time-consuming task. Much work unrelated to the mathematical model to realize which is just details of implementation is

Fig. 7 (continued)



involved. An actual implementation is a long, often confusing code where the underlying mathematical model is difficult to recognize, which is also often a source of errors. Because of this it is desirable to be able to describe the FEM model in a more mathematical abstract way, which is automatically translated to code. The nature of the error-prone, typically repetitive code with many similar cases to handle (e.g. for boundary conditions) make the code for element matrix computation ideal for automatic code generation. The same applies to the sparse matrix operations involved as typically the matrix format it adapted to the problem and the target hardware. The possibility to easily make modifications to the FEM model is desirable for the development of more robust models, a code generator is developed which allows abstract definition of FEM models and tries to generate Fortran or C code from it which executes efficiently on several different hardware platforms.

Fig. 7 (a) Unit cube in a shear/compression test, (b) max. stress at the midpoint of the cube, (c) 1/8 th of a circular bar in a tensile test, and (d) change of necking area in time



4.2 Abstract Model Description

The code generator is controlled via an input file which contains the necessary information of the mathematical model. The most important components of a FEM model were identified and an input format was specified. The following information is required to generate code for the computation of element matrices

- Element geometries, node coordinates
e.g. Quad, (0,0), (0,1), (1,0), (1,1)
- Base for space of shape functions
e.g. 1, x, y, xy
- Differential operator to discretize
e.g. $\frac{1}{2}\rho v^2$ (mass matrix) or $\frac{\partial x}{\partial X}$ (deformation gradient)
- Boundary conditions
e.g. Dirichlet, Neumann or combination
- Integration rules
Gaussian quadrature, or similar weighted points rule

The description of elements and integration rules is treated in a straightforward way. For example the description of a first-order 2D quad element would be done as follows:

```

Element Quad {
  Type = Square;
  Nodes myNodeSet {
    Coordinates {
      {0,0},
      {1,0},
      {0,1},
      {1,1}
    }
    ShapeFunctions {
      Polynomial { 1 },
      Polynomial { x },
      Polynomial { y },
      Polynomial { x y },
    }
  }
  Coordinates = myNodeSet;
  Integration myRule {
    0.25 -> { 0.211325, 0.211325 },
    0.25 -> { 0.789675, 0.211325 },
    0.25 -> { 0.211325, 0.789675 },
    0.25 -> { 0.789675, 0.789675 },
  }
}

```

It is possible to define multiple node sets and integration rules to treat different degrees of freedom differently.

The description of the operator to discretize is done in weak form, which is necessary to allow for the specification of partially integrated forms with derivatives of the test functions appearing. For example the Poisson equation in 2 dimensions

$$(\partial_x^2 + \partial_y^2)\psi = \rho$$

would be treated in the partially integrated weak form

$$\int_{\Omega} \partial_x \phi \partial_x \psi + \partial_y \phi \partial_y \psi \, d^2x = - \int \phi \rho \, d^2x$$

which can be expressed in the input file as

```

Matrix myExample {
  Elements { myQuad }
}

```

```

Coefficients { myNodeset:rho }
Variables    { myNodeset:psi }
TestFunctions { myNodeset:phi }
Operator leftSide {
  diff(phi,x)*diff(psi,x)+diff(phi,y)*diff(psi,y)
}
Vector Operator rightSide {
  -phi*rho
}
}

```

4.3 *Generated Code*

At the moment the generator can produce code for the languages

- C/C++
- Fortran 90
- CUDA C
- CUDA Fortran
- OpenCL

The code generator chooses for each target platform an implementation variant which takes into account hardware details to produce efficient code for this platform. For example the efficiency of sparse matrix storage formats is highly dependent of CPU and memory architecture and different patterns of vectorization should be used, depending on the SIMD length of the CPU.

4.4 *Current State*

At the moment the complete cycle of generation of element matrices is implemented and the generated code seems to produce correct result. Computation of nodal shape function from the basis, shape function derivatives, Jacobians, transformation of derivatives in the global coordinate system and numerical integration of operators are done as intermediate steps. This already relieves the developer of a great deal of work. The mathematical description can be stated in compact form and a complex computation code is automatically produced from the description. For example the complete input file of Poisson equation example as shown above is 45 lines long, which results in 1,047 lines of Fortran code.

At the moment work is in progress related to the generation of sparse matrix assembly code. Aim is to produce efficient parallel code for both multicore CPUs using OpenMP and accelerators using CUDA or OpenCL. Another step which is planned to be done in the near future is handling of boundary conditions, which is of course important but currently not available.

4.5 Results

Benchmarks of the generated code have been done for the Poisson equation example on a Intel Sandy Bridge CPU, a NVIDIA Tesla C2050 and an AMD Radeon 6970. The performance on all three platforms is promising as can be seen below but likely there is room for further optimization:

	Hardware peak performance (GFLOPS)	Performance of generated code (2D Poisson) (GFLOPS)
Sandy Bridge	105	≈30
Tesla C2050	500	≈200
Radeon 6970	500	≈200

The goal to create a tool which can be used to create real-world FEM applications has at least partially been achieved, but work will continue on adding more useful features and tuning performance for the different target platforms.

References

1. Laursen TA. *Computational contact and impact mechanics: fundamentals of modeling interfacial phenomena in nonlinear finite element analysis*. Springer, 2002.
2. Ascher UM and Petzold LR. *Computer Methods for Ordinary Differential equations and Differential-Algebraic equations*. SIAM: Philadelphia, 1998.
3. Hager C. Robust numerical algorithms for dynamic frictional contact problems with different time and space scales. *Ph.D. Thesis*, Universität Stuttgart, 2010. Available at: <http://elib.uni-stuttgart.de/opus/volltexte/2010/5868/>.
4. Renard Y. The singular dynamic method for constrained second order hyperbolic equations: Application to dynamic contact problems *Journal of Computational and Applied Mathematics* 2010; **234**:906–923.
5. Washizu K. *Variational Methods in Elasticity and Plasticity*. (3rd edition), Pergamon Press, 1982.
6. Krause, R. A Non-Smooth Multiscale Method for Solving Frictional Two-Body Contact Problems in 2D and 3D with Multigrid Efficiency, *SIAM Journal on Scientific Computing*, 31, no. 2, pp. 1399–1423. 2009.
7. Simo, J.C. Hughes, T.J.R., *Computational Inelasticity*, Springer, 1998
8. UG: <http://atlas.gcsc.uni-frankfurt.de/~ug/>
9. Wiens, C. Nonlinear solution methods for infinitesimal perfect plasticity, *Journal of Applied Mathematics and Mechanics*, 87, no. 8–9, pp. 643–660, 2007.
10. Wiens, C. Lang, S. Wittum, G. The application of adaptive parallel multigrid methods to problems in nonlinear solid mechanics. *Adaptive Methods in Solid Mechanics*. J. Wiley, 2002.

Characterization of Carrier Sense Multiple Access in Vehicular Propagation Channels

J. Mittag and H. Hartenstein

Abstract Wireless communications between vehicles is considered to be one of the building blocks in order to increase the safety level offered by future intelligent transportation systems. While it sounds intuitively convincing that a periodic exchange of status information, e.g. the current position, speed and driving direction, may help to avoid dangerous traffic situations or driving maneuvers, it is not clear whether the envisioned communications system, i.e. IEEE 802.11p, is sufficiently reliable and robust. In particular, it is not clear whether the employed Carrier Sense Multiple Access (CSMA) mechanism employed at the medium access control (MAC) layer is able to coordinate concurrent access by multiple network nodes in a highly dynamic environment as intended. In this paper, we evaluate the performance of the CSMA-based coordination mechanism employed by IEEE 802.11p. The evaluation is based on a network simulation framework that emulates the signal processing steps of a transceiver and accurately models the multi-path propagation effects of the wireless vehicular radio channel. Due to this accuracy, the execution of such high fidelity simulations is computationally highly expensive and represents a prominent example of the discipline called Computational Science and Engineering (CSE). Based on the results of our evaluation, we come to the conclusion that CSMA is able to coordinate concurrent access in vehicular environments, even if fading radio propagation characteristics are present.

J. Mittag (✉) · H. Hartenstein

Decentralized Systems and Network Services, Institute of Telematics, Steinbuch Centre for Computing (SCC), Karlsruhe Institute of Technology (KIT), D-76128 Karlsruhe, Germany
e-mail: jens.mittag@kit.edu; hannes.hartenstein@kit.edu

1 Introduction

Vehicle-to-vehicle (V2V) communications is required for numerous applications that aim to improve traffic safety. By a periodic wireless exchange of status related messages among neighboring vehicles, a mutual awareness should be established. Through this awareness vehicles are expected to detect potentially dangerous traffic situations in advance and are therefore able to avoid fatal driving maneuvers. In order to ensure a reliable operation of such active safety systems, the deployed wireless communications system has to be reliable and robust. In particular, the physical layer has to deal sufficiently well with fading radio propagation characteristics. Likewise, medium access control mechanisms have to coordinate concurrent access properly within a wide range of radio propagation conditions as well. Whether this is actually achieved is an open question that can not be answered easily.

To obtain a solid answer, we opt for a simulation based evaluation that considers primary and relevant influencing factors, in particular a valid representation of the channel effects. In [2, 5] we therefore presented an extension of the open source NS-3 network simulator [6], in particular the addition of a signal-level implementation of the IEEE 802.11p physical layer (PHY) and medium access control (MAC) specifications – which are envisioned to be used in a first generation of vehicle-to-vehicle communication systems. The proposed modeling approach replaces simplified models that have been used prior to our proposal and is accurate enough to evaluate whether the IEEE 802.11p based PHY and MAC are sufficiently robust against the fading exhibited in vehicular radio propagation channels. Since the highly data intensive signal processing leads to significant runtime performance slowdowns, we evaluated the benefit of parallel processing capabilities in [1], with the conclusion that noticeable speedups are only achievable in many-core system architectures with more than a few hundred compute units – for instance when using a GPU-based simulation architecture.

In this report, we present the results of a simulation campaign that addresses the question whether the IEEE 802.11p based MAC layer is robust with respect to the fading characteristics typically present in vehicular radio channels. In particular, we characterize the coordination performance of Carrier Sense Multiple Access (CSMA), which is the fundamental MAC mechanism employed by IEEE 802.11p, over a wide range of scenario configurations and radio propagation conditions. With CSMA, each network node listens to the channel prior to an own transmission, and starts a transmission only if the channel is not considered busy. If the channel is busy, the own transmission is deferred until the channel becomes idle again. Since we have no access to a GPU-based compute cluster (on which we can exploit our findings presented in [1] and on which the runtime performance of a single simulation experiment would be “optimal”) we use the most simple parallelization method: batch processing based on the parallel job scheduling capabilities offered by the HP XC4000 at the Steinbuch Centre for Computing.

The rest of this paper is structured as follows: Sect. 2 presents the evaluation methodology and introduces the performance metrics that were used to characterize

the coordination efficiency of CSMA. Section 3 then presents a selected subset of the obtained results to illustrate the outcome of the evaluation and provide an answer to the question whether CSMA is able to coordinate concurrent access sufficiently well in fading environments. Section 4 eventually concludes this paper with a brief summary.

2 Evaluation Methodology

The characterization of CSMA is performed considering a highway environment with vehicles being placed uniformly on a 5 km long road with 2 lanes per direction. The highway environment is chosen since considerably high velocities, and hence pronounced time- and frequency-selective propagation characteristics, can be expected in this setting. A simple broadcast application that is running on each vehicle generates periodic awareness messages at an average rate r (in Hz). In order to introduce a small randomness, a small jitter is applied to the interval between two subsequent awareness messages.

To evaluate different network saturation levels, application specific impact factors are varied, i.e. the packet generation rate r is set to either 2, 5 or 10 Hz, the transmission power is set to either 5, 10, 15 or 20 dBm, and the size of an awareness message is set to 200 or 400 bytes. Furthermore, three different average vehicle densities are considered to vary the number of transceivers for which concurrent access has to be coordinated: 40, 80, and 120 vehicles/km. Although mobility is considered in order to simulate fast-fading channel conditions, vehicles are configured to keep their (initial) positions. Since CSMA does not employ any slot reservation technique, and vehicles would not alter their positions significantly during a few milliseconds (with respect to the dimension of the network in terms of communication range), the topology of the network can be considered stationary during the channel contention period. This configuration should therefore not affect the relevance of the obtained results. Nevertheless, in order to enable an application of channel fading models, a (fake) average mobility of 100 km/h is considered by radio propagation models.

With respect to IEEE 802.11p medium access control, a basic Distributed Coordination Function (DCF) with a Clear Channel Assessment (CCA) threshold of -91 dBm, a fixed contention window size of 15 slots, and a slot time of $13 \mu\text{s}$ is used. Further, each vehicle is configured to use a data rate of 6 Mbps in a 10 MHz channel at a carrier frequency of 5.9 GHz. An overview of additional relevant configuration parameters is given by Table 1.

Most importantly, the impact of three different radio propagation characteristics is evaluated. Initially, only a distance dependent deterministic path loss is considered to study the coordination performance in the absence of any channel fading characteristics. Such a consideration enables the identification of the fundamental CSMA weaknesses, and serves as a reference when analyzing the results of subsequent simulations in which fading is considered. As proposed by Kunisch et al. in [3], only

Table 1 Application layer, medium access control and physical layer parameters and the settings used for the packet collision characterization

Parameter	Value
Application layer	
Packet size	200, 400 bytes
Transmission rate	2, 5, 10 Hz
Transmission power	5, 10, 15, 20 dBm
Medium access control layer	
Slot time	13 μ s
Contention window size	15 slots
CCA busy threshold	-91 dBm
Physical layer	
Modulation scheme	QPSK
Coding rate	1/2
Channel bandwidth	10 MHz
Carrier frequency	5.9 GHz
Tx center frequency offset tolerance	20 ppm
Capture threshold	8 dB
Noise level	-99 dBm

a power law model with a reference loss of 59.7 dB and a path loss exponent of 1.85 is used. Then, large-scale fading characteristics that follow a Normal distribution with $\sigma = 3.2$ dB are introduced. In a last step, the effect of a small-scale fading is analyzed by the simulation of a Rayleigh fading channel using the Jakes Doppler spectrum (instead of the Normal shadowing).

During each simulation run, the “lifetime” of all packets transmitted by 20 selected reference nodes (which are located in the center of the scenario to avoid border effects) is monitored and evaluated with respect to different performance metrics. The applied metrics are classified with respect to receiver and transmitter perspective.

From the perspective of a transmitting (reference) node, the primary metric that describes the coordination efficiency of a MAC protocol is the probability that one or multiple other nodes are incoordinated to an own transmission.

Definition 1 (Packet Level Incoordination, PLI). The packet level incoordination, as observed from the perspective of a node r and one of its generated packets p , describes the probability that at least one node s , $s \neq r$, transmitted a packet q during the transmission period of p .

Apart from the quantification of the PLI, it is also important to resolve the type of incoordination. With respect to CSMA, the reason for an incoordination could be that either both nodes started their transmission at exactly the same points in time, e.g. due to simultaneous expiration of backoff timers, or that the incoordinated node did not sense the reference transmission, e.g. due to shadowing or fading channel characteristics. To distinguish between both cases, the time difference between both transmissions has to be evaluated, which leads to the definition of the *Incoordination Delay Profile*.

Definition 2 (Incoordination Delay Profile, IDP). The incoordination delay profile describes the probability distribution of the starting time differences between a set of packet transmissions $P = \{p_1, \dots, p_n\}$ and each packet's corresponding set of incoordinated transmissions $Q_i = \{q_{i1}, \dots, q_{ij}\}$, with $1 \leq i \leq n$, and j being the number of packets interfering with p_i .

In case of CSMA and distance decaying deterministic channel conditions, the IDP should indicate that all incoordinated nodes located within the carrier sense range – the range within which the received signal will exceed the CCA busy threshold – transmit more or less simultaneously. Only incoordinated nodes outside the carrier sense range should show significantly greater delays. To determine the effectivity of CSMA with respect to this controlled spatial reuse of the channel, the IDP is evaluated with respect to the distance between sender and incoordinated node in the following.

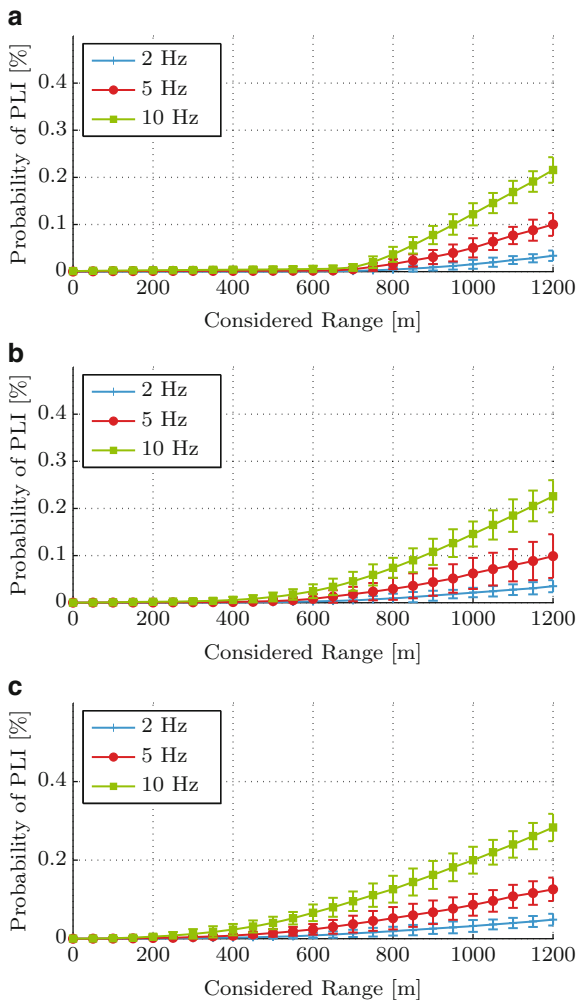
3 Results

The following section presents and discusses a representative subset of the obtained results. This subset is sufficient in order to illustrate the fundamental findings and to demonstrate the conclusions that are eventually drawn. In particular, due to space restrictions, only the 80 vehicles/km scenario with a 400 bytes packet size configuration is considered as a reference in the following. The reader who wishes to evaluate and screen the complete set of results is able to access the results as well as the full source code of all experiments online [4].

Figure 1 illustrates the observed packet level incoordination (with respect to the range in which incoordination is considered) using three different packet generation rates, the selected vehicle density of 80 vehicles/km, and three different channel configurations, i.e. a deterministic path loss only setup, a setup with additional Normal shadowing, and a setup with additional Rayleigh fading. As illustrated, almost no incoordination is observed within a range of approx. 700 m for most scenario configurations that exhibit only a deterministic path loss, cf. Fig. 1a. This is an expected result, since the range of 700 m reflects the area within which the received signal strength stays above the configured -91 dBm CCA busy threshold. CSMA can therefore be certified to achieve its design goal in such an environment and setup.

If a large-scale fading based on a Normal shadowing is assumed, the corresponding PLI curves are slightly increased in comparison to the deterministic path loss setup, cf. Fig. 1b. Particularly for ranges slightly below and above the deterministic carrier sense range of 700 m a noticeable increase of the PLI can be observed. This is also expected, since the large-scale fading leads to situations in which a vehicle located within the deterministic carrier sense range experiences received signal strengths below the -91 dBm CCA busy threshold, hence these vehicles will not block their MAC layer and might interfere with one of the reference transmissions.

Fig. 1 Probability of PLI with respect to the range within which incoordination is considered for a 20 dBm transmit power, different channel configurations, and a 2, 5 or 10 Hz setup. (a) Distance decaying pathloss. (b) Normal shadowing. (c) Rayleigh fading



This phenomenon is increased if a Rayleigh fading channel is considered, cf. Fig. 1c). Whether this small increase is still sufficiently low can not be answered at this point, since no statement about the resulting packet delivery ratios is made. In general terms, an additional evaluation of the MAC layer performance from a receiver perspective is required in order to answer the question whether a small increase of the PLI is significant or not.

As shown in Fig. 1a, a small amount of incoordinated transmissions is observed within the deterministic carrier sense range despite the absence of any signal fading. According to the design of CSMA, such an incoordination should only occur if the transmitting vehicle (i.e. the reference node) and the interfering vehicle start to transmit at (more or less) the same point in time. In such a situation, the interfering




Fig. 2 Observed incoordination delay profiles with respect to the distance between reference and incoordinated node for a 80 vehicles/km scenario, non-fading channel conditions, a 20 dBm transmit power and a 10 Hz packet generation rate

node is inherently not able to sense the signal of the reference node prior to its own transmission. In order to evaluate this aspect, the observed IDPs with respect to the distance between the reference node and a potentially interfering node are discussed in the following. Since the benefit of showing results for multiple scenarios is marginal, the setup of a 80 vehicles/km scenario with a 20 dBm transmit power, a 10 Hz packet generation rate and a packet size of 400 bytes is considered in the following.

Figure 2 illustrates the observed IDPs with respect to the distance between a reference node and interfering nodes for the deterministic channel configuration and a 10 Hz packet generation rate setup. Since the length of a 400 byte packet is equal to $576 \mu\text{s}$ in the time domain, the time difference between the transmission of a reference node and the transmission of an incoordinated node can be at most $576 \mu\text{s}$. This maximum value is however not observed for interferers located within a range of 700 m. Indeed, as expected in this setup, incoordination from vehicles located within the carrier sense range occurs only if interfering vehicles transmit exactly at the same point in time as the reference node. Hence, CSMA can again be certified to fulfill its objective. Please note that a more or less uniform distribution of the incoordination delay is observed for all distances greater than the carrier sense range.

Unsurprisingly, the situation changes if fading characteristics are introduced. As shown in Fig. 3a, which plots the IDPs obtained in a Normal shadowing configuration, the range within which incoordination occurs only due to identical transmission times decreases to approx. 300 m. For greater distances, the IDP tends towards a uniform distribution of the incoordination delays. A similar but more intense effect can be observed in the Rayleigh fading channel configuration, cf. Fig. 3b. In such an environment, the principle of CSMA has its difficulties to avoid incoordination that

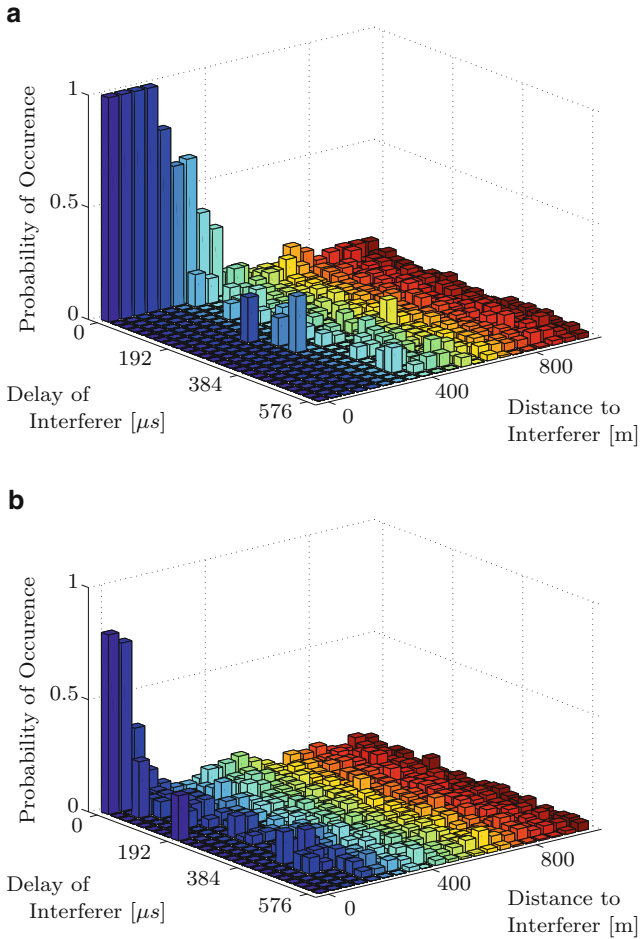


Fig. 3 Observed incoordination delay profiles with respect to the distance between reference and incoordinated node for a 80 vehicles/km scenario, fading channel conditions, a 20 dBm transmit power and a 10 Hz packet generation rate. **(a)** Normal shadowing. **(b)** Rayleigh fading

is not caused by identical transmission times. Recall that incoordination is still less likely to occur at such close distances (in comparison to remote distances). Only the time domain characteristic of incoordinated transmissions at close distances has changed, and is not fundamentally different from incoordinated transmissions at remote distances anymore (as it was in the deterministic path loss only channel configuration).

4 Conclusions

In this paper, we studied the coordination performance of Carrier Sense Multiple Access (CSMA) in vehicular environments, in particular under the influence of fading radio propagation conditions. We defined two metrics which characterize CSMA's ability to suppress concurrent and interfering packet transmissions by nodes located within each other's close surrounding. We further employed a high fidelity network simulation framework to capture and accurately model the impact of a multi-path radio propagation channel and the influence of signal processing algorithms at the physical layer. Since such an accurate modeling approach is highly expensive with respect to the computational costs we made use of the HP XC4000 operated by the Steinbuch Centre for Computing (SCC) at the Karlsruhe Institute of Technology (KIT). According to the obtained results, CSMA can be certified to fulfill its design objectives, in the sense that it is able to effectively suppress concurrent packet transmissions by nodes located within each others carrier sensing range. If fading propagation conditions are present, the effectiveness of CSMA is slightly reduced. In order to answer whether this reduction has a significant impact on the packet reception performance or not, we plan to evaluate the performance of CSMA from the perspective of a receiver in a follow-up work.

Acknowledgements This work was supported by the Steinbuch Centre for Computing (SCC) that is part of the Karlsruhe Institute of Technology (KIT).

References

1. P. Andelfinger, J. Mittag, and H. Hartenstein. GPU-Based Architectures and Their Benefit for Accurate and Efficient Wireless Network Simulations. In *Modeling, Analysis Simulation of Computer and Telecommunication Systems (MASCOTS), 2011 IEEE 19th International Symposium on*, pages 421–424, July 2011.
2. T. Gaugel, L. Reichardt, J. Mittag, T. Zwick, and H. Hartenstein. Accurate Simulation of Wireless Vehicular Networks Based on Ray Tracing and Physical Layer Simulation. In W. E. Nagel, D. B. Kröner, and M. M. Resch, editors, *High Performance Computing in Science and Engineering '11*, pages 619–630. Springer, January 2012.
3. J. Kunisch and J. Pamp. Wideband Car-to-Car Radio Channel Measurements and Model at 5.9 GHz. In *Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th*, pages 1–5. IEEE, September 2008.
4. J. Mittag. Characterization, Avoidance and Repair of Packet Collisions in Inter-Vehicle Communication Networks: The Complete Set of Results, February 2012. <http://dsn.tm.kit.edu/download/ns3-physim/results.html>.
5. J. Mittag, S. Papanastasiou, H. Hartenstein, and E. Ström. Enabling Accurate Cross-Layer PHY/MAC/NET Simulation Studies of Vehicular Communication Networks. *Proceedings of the IEEE*, 99(7):1311–1326, July 2011.
6. The NS-3 Network Simulator. <http://www.nsnam.org/>.

HMDB51: A Large Video Database for Human Motion Recognition

Hilde Kuehne, Hueihan Jhuang, Rainer Stiefelhagen, and Thomas Serre

Abstract With nearly one billion online videos viewed everyday, an emerging new frontier in computer vision research is recognition and search in video. While much effort has been devoted to the collection and annotation of large scalable static image datasets containing thousands of image categories, human action datasets lag far behind. Current action recognition databases contain on the order of ten different action categories collected under fairly controlled conditions. State-of-the-art performance on these datasets is now near ceiling and thus there is a need for the design and creation of new benchmarks. To address this issue we collected the largest action video database to-date with 51 action categories, which in total contain around 7,000 manually annotated clips extracted from a variety of sources ranging from digitized movies to YouTube. The goal of this effort is to provide a tool to evaluate the performance of computer vision systems for action recognition and explore the robustness of these methods under various conditions such as camera motion, viewpoint, video quality and occlusion.

H. Kuehne (✉) · R. Stiefelhagen
KIT, Institute for Anthropomatics, Karlsruhe, Germany
e-mail: kuehne@kit.edu; rainer.stiefelhagen@kit.edu

H. Jhuang
Perceiving Systems Department, Max Planck Institute for Intelligent Systems, Tübingen,
Germany
e-mail: hueihan.jhuang@tuebingen.mpg.de

T. Serre
Institute for Brain Sciences, Brown University, Providence, RI, USA
e-mail: Thomas_Serre@brown.edu

1 Introduction

We attempt to advance the field with the design and collection of a large video database dubbed Human Motion DataBase (HMDB) that tries to capture the richness and complexity of human actions. With 51 distinct action categories each containing at least 100 clips for just a little under 7,000 video clips extracted from a variety of sources, the proposed database is, to our knowledge, the largest and perhaps the most realistic available databases. Each clip of the database was validated by at least two human observers to ensure consistency. In addition to category labels, meta-data was added to facilitate further selection, pre-processing and training of recognition systems (i.e., number of actors involved in the action, view-point, approximate viewing-distance, presence/absence of camera motion, video quality, etc.).

First, we use this database to evaluate the performance of two representative computer vision systems: We considered the biologically-motivated action recognition system by Jhuang et al. [7], which is based on a model of the dorsal stream of the visual cortex and was recently shown to achieve human level of performance for the recognition of rodent behaviors in the constrained setting of the home environment [6]. We also considered the popular spatio-temporal bag-of-words system by Laptev and colleagues [9,10,15]. We here compare the performance of these state-of-the-art systems on the HMDB database, evaluate their robustness to various sources of image degradations (camera motion, occlusion and changes in view-point) and discuss on the relative role of shape vs. motion information for action recognition.

2 Related Work

With several billions of videos currently available on the internet and about 24 h of video uploaded to YouTube every minute, there is an immediate need for robust algorithms that could help organize, summarize and retrieve this massive amount of data. While much effort has been devoted to the collection of realistic internet-scale collection of static image databases [4,5,11,13,14,18], current action recognition datasets lack far behind. The three most popular benchmark databases (i.e., KTH [1], Weizmann [3] and the IXMAS [16] datasets contain around 6–11 actions each. These databases are not quite representative of the richness and complexity of real-world action videos as they are fairly well constrained in terms of illumination and camera position. A typical video clip contains a single (staged) actor with no occlusion and very limited clutter. Recognition rates on these datasets tend to be very high. For instance, a recent survey of action recognition system comparison [17] reported that 12 out of 21 systems tested perform better than 90% on the KTH dataset. For the Weizmann dataset, 14 out of 16 tested systems perform at 90% or better, 8 out of 16 better than 95% and 3 out of 16 scored a perfect 100% recognition rate.




Fig. 1 Comparison between current action recognition datasets

Recent action datasets try to build more realistic action recognition datasets (Hollywood and UCF50) by considering video clips taken from real movies and youtube. These datasets are more challenging due to large variations in camera motion, object appearance, pose, scale and viewpoint as well as cluttered background, etc. The *UCF50* database (available at <http://server.cs.ucf.edu/~vision/data/UCF50.rar>) extends the 11 original action categories from the *YouTube action dataset* and consists of 50 action categories with realistic videos taken from youtube. For all the 50 categories, the videos are grouped into 25 groups, where each group consists of more than 4 action clips.

The UCF50, its premier version UCF sports, and a recent Olympic sports dataset [19] contain mostly sport videos from YouTube. These types of action are relatively unambiguous (as a result of searching for specific titles on YouTube), and are highly distinguishable from shape cues alone (i.e., such as the raw positions of the joints or the silhouette extracted from single frames). To demonstrate this point, we conducted a simple experiment: We manually annotated stick-figures (i.e., 9 line segments corresponding to the two upper and lower arms, the two upper and lower legs and the body trunk) from 5 randomly selected clips from each of the 13 action categories on the UCF sport dataset. Using a leave-one-clip-out procedure, classifying the raw joint locations from single frames lead to a recognition rate above 98 % (chance level 8 %). We conducted a very similar experiment on the proposed HMDB database where we drew from 10 action categories similar to those used in the UCF (e.g., climb, climb-stairs, run, walk, jump, etc.) and manually annotated the joint locations for a set of over 1,100 random clips (see Sect. 4 for details). The accuracy reached by a classifier using the joint location as inputs reached only 35 % this time (chance level 10 %) and performed below the level of performance of the same classifier using motion features [7]. Such dataset may thus be a better indicator of the relative contributions of motion vs. shape cues for the recognition of real-world actions (see Sect. 4). Figure 1 shows a comparison between existing action recognition datasets.

3 Design of the Dataset

It has been estimated that there are over 1,000 human action categories. In order to isolate human actions that are representative of everyday actions, we first asked a group of students to watch videos using a subtitle annotation tool to annotate any segment of these videos that they deemed to represent a single non-ambiguous human action. Students were asked to consider a minimum quality standard (i.e., a single action per clip, a minimum of 60 pixels in height for the main actor, minimum contrast level, minimum action length of about 1 s and acceptable compression artifacts). Students considered videos from three sources: digitized films available on the internet, public databases such as the Prelinger archive as well as YouTube and Google videos. A first set of annotations was thus generated in this way with over 60 action categories. To further guarantee that we would be able to populate all action categories with at least 101 different video clips we considered the top 51 action categories and asked students to specifically look for these types of actions.

This database, that we call HMDB51, comprises 51 distinct human action categories, with overall 6,474 clips from 1,697 unique source videos. We also collected meta-data towards a precise evaluation of the limitation of current computer vision systems. Each clip annotation contains a field indicating the visible body parts (head, upper body, lower body, full body), whether the camera is moving or static, how many people are involved in performing the action (single action vs. two-people vs. multiple people actions), as well as the camera orientation relative to the actor (front, back, left or right). The clips were also annotated according to their video quality from (a) good (i.e., detailed visual elements such as the fingers and eyes of the main actor are identifiable through most of the clip with limited motion blur and compression artifacts), (b) medium (i.e., larger body parts like the upper and lower arms and legs are identifiable through most of the clip) and (c) fair (i.e., even larger body parts are not identifiable due in part to the presence of motion blur and compression artifacts). The distribution of labels in the HMDB50 is as follow:

- Clips with camera motion – 59.9 % vs. static camera – 40.1 %
- Camera position: front – 40.8 %, back – 18.2 %, left – 22.1 %, right – 19.0 %
- Clip quality: good – 17.1 %, medium – 62.1 %, fair – 20.8 %
- Visible body part: full body – 56.3 %, head – 12.3 %, lower body – 0.8 %, upper body – 30.5 %

3.1 Preprocessing Steps

Normalization

The original videos sources used to extract the action video clips varied in size (from 176×132 to $1,280 \times 720$) as well as frame rates (6–60 fps). To ensure consistency, we

thus resized all extracted clips to a height of 240 pixels (using bicubic interpolation over a 4×4 neighborhood). The width of the clips was scaled accordingly so as to maintain the original aspect ratio. We further normalized all video frame rates (by either dropping or duplicating frames) to ensure a fixed 30 fps frame rate over all clips. This was done using the *ffmpeg* video library. All clips were then compressed using the *DviX 5.0* codec with the *.avi* output format.

Video Stabilization Using Point Feature Matching

One significant challenge associated with the use of video clips extracted from real-world videos is the potential presence of significant camera/background motion (about $2/3$ of the clips in our database, see above). Such camera motion can interfere with the local motion computation (see Results) and should potentially be corrected. Stabilizing these videos is thus a key pre-processing step. A simple way to correct for camera motion is to apply a basic image stitching algorithm to align successive frames according to the camera motion. A background plane is estimated by first detecting and then matching salient features (using the Harris corner detector) between adjacent frames. Correspondences are computed using a distance measure that includes both the absolute pixel differences within a 15×15 image patch centered on the corner point and the Euler distance of the corner points. Corner points with the minimum distance are then matched and the RANSAC algorithm is used (50% inlier with 95% confidence) to estimate the geometric transformation between all neighboring frames from these noisy correspondences. This is done independently for every pair of frames and after smoothing a cumulative image transformation is then computed and movie frames are then warped using this estimate to achieve a stabilized video.

Training and Test Splits Generation

For evaluation purposes three distinct training and test splits were generated from the database. This splits were generated so as to ensure that (1) the same video source could not be used for both training and testing and that (2) the relative distribution of camera positions (view-point) as well as the proportion of clips with/without camera motion, the video quality and visible body parts would be balanced across the training and test sets. To do this, we implemented a very simple constraint satisfaction approach to select from a large number of completely random splits. We first picked the one best split according to these constrains (split_1) and then the second and third best splits that would be least correlated to this first split (see Table 1 for a measure of overlap between the splits as measured by a normalized Hamming distance)

Table 1 Normalized Hamming distance between the three training and test splits generated so as to minimize the number of video sources present in both the training and test set and to maintain similar statistics across view-points, amount of camera motion, etc.

(s_1, s_2)	(s_1, s_3)	(s_2, s_3)	Random splits
0.34	0.33	0.33	0.16 ± 0.10

4 Evaluation

4.1 Benchmark Systems

4.1.1 Biologically-Motivated Action Recognition System

Jhuang et al. have described a computational model of the dorsal stream for the recognition of actions [7]. The model starts with spatio-temporal filters modeled after motion-sensitive cells in the primary visual cortex [12]. Just like the V1-like simple units in the model of the ventral stream described above, these units are tuned to specific orientations. As opposed to those in the model of the ventral stream, which respond best to static stimuli, the V1-like simple units in the model of the dorsal stream respond best to a bar moving in a direction orthogonal to its preferred orientation. It has been suggested that motion-direction sensitive cells and static cells constitute two channels of processing, the former projecting to the dorsal stream and the latter to the ventral stream. During an unsupervised developmental-like learning stage, units in intermediate stages of the model become tuned to optic-flow patterns that appear with high probability in natural sequences of images. These optic-flow pattern units correspond to the combination of several complex cell receptive fields (tuned to different directions of motion instead of spatial orientations in the model of the ventral stream). Here we obtained code directly from the website of the authors.

4.1.2 Spatio-Temporal Bag-of-Features (ST-BoF)

Local space-time features have recently become a popular video representation for action recognition. Much like their static local spatial features counterpart for the recognition of objects and scenes, they have been shown to achieve state-of-the-art performance on several standard action recognition databases [10, 15]. An extensive comparison between existing methods (feature detectors and local descriptors) for the computation of space-time features in a common experimental setup was described in [15].

We implemented a system based one of the most commonly used system configuration using a combination of the *Harris3D* detector and the HOGHOF descriptors. For every clip we detected 3d Harris corners and computed the combination of histograms of oriented gradients(HOG) and oriented flow (HOF)

as local descriptor. To build the code book between 75000 and 250000 space-time-interest-point-descriptors were sampled from the dataset and a k-means clustering with $k = 2,000-10,000$ was applied on this sample set. The space-time-interest-point descriptors of every clips are matched to the nearest prototype as returned by k-means clustering and the histogram is build over the mapping. This leads to a k-dimensional feature vector where k is the number of clusters used for k-means. This feature vector is then used as input to an SVM classifier. For the classification a regularized support vector machine is used with an RBF kernel $K(u, v) = \exp(-\gamma * |u - v|^2)$. The parameters of the RBF kernel (cost term and γ) were optimized using a greedy search with a five-fold cross-validation on the training set over $c = 2^{-5}, 2^{-3}, \dots, 2^{25}$ and $\gamma = 2^{-15}, 2^{-13}, \dots, 2^{15}$.

4.2 Evaluation of the Two Systems Performance

We first evaluated the overall performance of both systems on the HMDB as well as on the ucf50 dataset. On the HKM, both systems show a similar mean recognition performance around 20 % whereas the performance of ST-BoF systems is 2 % lower than the system from Jhuang et al. [7] (Fig. 2). The recognition the ucf50 dataset, the recognition performance of both systems is again very close. Both have a performance about 45 % whereas in the case, the system of ST-BoF systems is 2.4 % better than the system of Jhuang et al. [7]. So there's no system showing a clear superiority above the other, but it becomes clear that the datasets are fairly different (Table 2).

4.3 Comparison of the Systems Performance on Ten Common Action Categories from UCF Versus HMDB

In order to achieve a baseline for the performance of HMDB according to other databases as UCF50, we identified ten categories that were similar or equal in both datasets (UCF50/HMDB50): basketball/shoot_ball, biking/ride_bike, diving/dive, fencing/stab, golf swing/golf, horse riding/ride_horse, pull ups/pullup, push ups/pushup, rock climbing indoor/climb and walking with dog/walk. We evaluated the performance of ST-BoF recognition system with the original clips as well as with the stabilized ones on both datasets. On the HMDB the mean classification rate for the ten categories is at 54.3 % (57.3 % for motion stabilized clips) whereas the classification rate for the UCF50 for the same categories is at 66.3 % (68.7 % for stabilized clips). These results suggest that the HMDB is a more challenging and perhaps richer dataset.




Fig. 2 Confusion matrix for the ST-BoF system on the stabilized clips

Table 2 Average performance for the two benchmark systems used here (Laptev et al. [10] with $k = 5,000$ clusters and Jhuang et al. [7]) on the HMDB and UCF50 datasets

Databases (orig./stab.)	Laptev et al. (%)	Jhuang et al. (%)
HMDB51	20.4/21.9	22.8/23.1
UCF50	45.4/49.3	43/-

4.4 Robustness of the Two Systems

In order to assess the relative strengths and weaknesses of the two benchmark systems, we broke down their performance in terms of: (1) the quality of the video clips, (2) occlusions, (3) camera position and (4) camera motion. The results are summarized in Table 3.

Camera Motion

For the ST-BoF approach, 23.1 % of the clips with camera motion were classified correctly while the classification rate dropped down to 17.6 % for clips with camera motion. When tested on stabilized clips, the recognition rate for clips without motion remained stable at 23.8 % whereas the classification rate for the clips with camera motion increased by 7.7–25.3 %. Here it can be seen, that camera motion does effect the results of the ST-BoF system. Surprisingly the performance of the C2 features by Jhuang et al. was higher on the clips with camera motion than on the clips without camera motion (17.4 vs. 25.8 %). As evident from the decrease in performance when clips were stabilized for motion, it seems that the system is somehow able to pick on the camera motion, which, on this dataset seems to be correlated with the category label. This will require further investigation.

Table 3 Average performance for the two benchmark systems used here (Laptev et al. [10] with $k = 5,000$ clusters and Jhuang et al. [7] with $n = 2,000$ C2 features) analyzed by their robustness to various sources of video degradations

		Camera motion			
HMDB(orig./stab.)	Without camera motion (%)		With camera motion (%)		
Laptev et al.	23.1/23.8		17.6/25.3		
Jhuang et al.	17.4/22.9		25.8/21.6		
		Clip quality			
HMDB(orig./stab.)	Good (%)	Medium (%)	Fair (%)		
Laptev et al.	32.0/26.3	16.4/24.7	19.3/22.8		
Jhuang et al.	21.9/23.8	22.9/23.0	21.6/16.4		
		Camera position			
HMDB(orig./stab.)	Front (%)	Back (%)	Left (%)	Right (%)	
Laptev et al.	20.6/20.9	16.8/20.7	16.5/30.1	24.9/30.3	
Jhuang et al.	19.2/20.4	20.1/23.3	26.8/23.1	27.6/23.4	
		Visible body part			
HMDB(orig./stab.)	Full body (%)	Upper body (%)	Lower body (%)	Head (%)	
Laptev et al.	17.8/24.8	22.7/27.2	27.3/27.3	21.1/19.1	
Jhuang et al.	24.6/20.3	20.7/25.8	46.7/21.4	16.7/21.0	

Clip Quality

Here again we found a surprising result. While the performance of the ST-BoF approach seems to be influenced by the quality of the clips, the system by Jhuang et al. seems almost completely invariant to degradations in the quality of the videos. While the two systems achieve comparable rates on the medium and fair quality clips, the ST-BoF seems to perform significantly better on the good quality clips (on the original clips the performance of the ST-BoF vs. Jhuang et al. approach is 32.0 vs. 21.9%). A further evaluation of the dataset shows that in the subset of good quality clips, 1/3 of the clips contain camera motion, against 2/3 for the medium-quality and 3/4 for the bad quality clips. This suggest that the evaluation of the clip quality on systems performance is most likely contaminated by the presence of camera motion. We feel that this confound will likely start to decrease as the database grows in size.

Camera Position

With the exception of the lower body subset (with only 0.9% of the total clips results on this subset are non-significant), it seems that the position of the camera in the video does not influence the performance of the two systems significantly.

Table 4 Average performance for the two benchmark systems used here (Laptev et al. [10] and Jhuang et al. [7]) on the HMDB and UCF50 datasets

HMDB50	HOGHOF (%)	HOF (%)	HOG (%)
orig	19.7	16.3	15.9
stab	24.7	24.6	17.1

Visibility of Different Body Parts

For the evaluation of the different body parts, it can be said that the recognition is robust for clips with camera motion as well as for those without. Only the clips showing the full body of the character can be improved. This holds as well for the C2 features. The increase of the recognition of clips where only the lower body is visible shows that the algorithm can also deal with uncommon types of clips as this property holds only for a small fraction of clips and is not as much present in the training and testing set as the variations.

4.5 Shape Versus Motion Information

Table 4 shows a comparison between the performance of the original STBoF system (using the HOGHOF descriptor) as well as the contribution of shape alone (HOG) and motion alone (HOF) cues.

5 Conclusion

We have described an effort to advance the field of action recognition with the design of what is to our knowledge currently the largest action recognition database. With currently 50 action categories and a little under 7,000 video clips, the proposed database is still far from capturing the richness and the full complexity of video clips commonly found on the internet. However given the level of performance of representative state-of-the-art computer vision algorithms (i.e., about 25 % correct classification with chance level at 2 %), this initial database is arguably a good place to start (performance on the CalTech-101 database for object recognition started around 16 % [8]). Furthermore our exhaustive evaluation of these two systems suggest that performance is not significantly affected over a range of factors such as camera position and motion as well as occlusions. This suggests that current methods are fairly robust with respect to these low-level video degradations but remain limited in their representative power in order to capture the complexity of human actions.

Acknowledgements This report describes research done in part at the Center for Biological & Computational Learning, affiliated with MIBR, BCS, CSAIL at MIT. This research was sponsored by grants from DARPA (IPTO and DSO), NSF (NSF-0640097, NSF-0827427), AFSOR-THRL (FA8650-05- C-7262). Additional support was provided by: Adobe, King Abdullah University Science and Technology grant to B. DeVore, NEC, Sony and by the Eugene McDermott Foundation. This work is also done and supported by Brown University, Center for Computation and Visualization, and the Robert J. and Nancy D. Carney Fund for Scientific Innovation, by DARPA (DARPA-BAA-09-31), and ONR (ONR-BAA-11-001). H.K. was supported by a grant from the Ministry of Science, Research and the Arts of Baden Württemberg, Germany.

References

1. C. Schudt, I. Laptev, B. Caputo: Recognizing human actions: A Local SVM Approach. In: Proc. ICPR 2004, Cambridge, UK.
2. R. Blake and M. Shiffrar: Perception of human motion. In: *Annu. Rev. Psychol.*, 58: 47–73, 2007.
3. M. Blank, L. Gorelick, E. Shechtman, M. Irani, and R. Basri: Actions as space-time shapes. In: *IEEE International Conference on Computer Vision*, 2005. ICCV 2005, 2, 2005.
4. J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei: Imagenet: A large-scale hierarchical image database. In: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009 2009.
5. M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman: The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results. <http://www.pascalnetwork.org/challenges/VOC/voc2010/workshop/index.html>.
6. H. Jhuang, E. Garrote, X. Yu, V. Khilnani, T. Poggio, A. D. Steele, and T. Serre. Automated home-cage behavioural phenotyping of mice. *Nature Communications*, 1(5): 1–9, 2010.
7. H. Jhuang, T. Serre, L. Wolf, and T. Poggio. A biologically inspired system for action recognition. *Proceedings of the Eleventh IEEE International Conference on Computer Vision (ICCV)*, 2007.
8. R. F. L. Fei-Fei and P. Perona. Learning generative visual models from few training examples: an incremental bayesian approach tested on 101 object categories. *IEEE CVPR Workshop on Generative-Model Based Vision*, 2004.
9. I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2–3): 107–123, 2005
10. I. Laptev, M. Marszaek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Jan 2008.
11. B. Russell, A. Torralba, K. Murphy, and W. Freeman. Labelme: a database and web-based tool for image annotation. *International Journal of Computer Vision*, 77(1): 157–173, 2008
12. E. Simoncelli and D. Heeger. A model of neuronal responses in visual area mt. *Vision Research*, 38(5): 743–61, 1998.
13. A. Torralba, R. Fergus, and W. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE T. Pattern Anal.*, pages 1958–1970, 2008.
14. A. Torralba, B. Russell, and J. Yuen. Labelme: online image annotation and applications. *Proceedings of the IEEE*, 98(8): 1467–1484, 2010.
15. H. Wang, M. Ullah, A. Klaser, I. Laptev, and C. Schmid. Evaluation of local spatio-temporal features for action recognition. *British Machine Vision Conference*, London, UK, pages 1–11, 2009
16. D. Weinland, E. Boyer, R. Ronfard, and G. LJK-INRIA. Action recognition from arbitrary views using 3d exemplars. *Computer Vision*, Jan 2007.

17. D. Weinland, R. Ronfard, and E. Boyer. A survey of visionbased methods for action representation, segmentation and recognition. *Computer Vision and Image Understanding*, pages 1–51, Oct 2010.
18. J. Xiao, J. Hays, K. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. *Computer Vision and Pattern Recognition (CVPR)*, 2010 IEEE Conference on, pages 3485–3492, 2010.
19. J. C. Niebles, C.-W. Chen and L. Fei-Fei. Modeling Temporal Structure of Decomposable Motion Segments for Activity Classification. *Proceedings of the 12th European Conference of Computer Vision (ECCV)*. 2010.

Efficient Parallelization of a Three-Dimensional High-Order Particle-in-Cell Method Applied to Gyrotron Resonator Simulations

J. Neudorfer, A. Stock, T. Stindl, R. Schneider, S. Roller, C.-D. Munz,
and M. Auweter-Kurtz

Abstract Growing computational capabilities and simulation tools based on high-order methods allow the simulation of complex shaped plasma devices including the entire nonlinear dynamics of the Maxwell-Vlasov system. Such simulations model the particle-field-interactions of a non-neutral plasma without significant simplifications. Thereby, new insights into physics on a level of detail that has never been available before provide new design implications and a better understanding of the overall physics. We present a high-order discontinuous Galerkin method based Particle-In-Cell code for unstructured grids in a parallelization framework allowing for large scale applications on HPC clusters. We simulate the geometrically complex resonant cavity of the 170 GHz gyrotron aimed for plasma resonance heating of the fusion reactor ITER and we demonstrate that a highly efficient parallelization is a crucial requirement to simulate such a complex large-scale device.

J. Neudorfer · A. Stock · C.-D. Munz (✉)
Institut für Aerodynamik und Gasdynamik, Universität Stuttgart, Stuttgart, Germany
e-mail: munz@iag.uni-stuttgart.de

T. Stindl
Institut für Raumfahrtssysteme, Abt. Raumtransporttechnologie, Universität Stuttgart,
Stuttgart, Germany
e-mail: fasoulas@irs.uni-stuttgart.de

S. Roller
Applied Supercomputing in Engineering, RWTH Aachen University
German Research School for Simulation Sciences GmbH, 52062 Aachen, Germany
e-mail: s.roller@grs-sim.de

R. Schneider
Karlsruher Institut für Technologie, Institut für Hochleistungsimpuls- und Mikrowellentechnik,
PF3640, 76021 Karlsruhe, Germany
e-mail: rudolf.schneider@kit.edu

M. Auweter-Kurtz
German Aerospace Academy ASA, 71034 Böblingen, Germany
e-mail: m.auweter-kurtz@german-asa.de

1 Introduction

For the numerical simulation of highly rarefied plasma flows, a fully kinetic modeling of the Boltzmann equation completed by the Maxwell equations is necessary. In the present report, we neglect collisional effects and focus our attention on the Maxwell-Vlasov equations allowing the self-consistent investigation of collective plasma phenomena. The numerical methods to tackle this non-linear problem in six-dimensional phase space are very briefly reviewed in Sect. 2. In Sect. 3 we present the highly scalable parallelization algorithm followed by Sect. 4 showing the results obtained from the scaling experiments and the gyrotron resonator simulations. Finally, a short outlook on further activities is given in Sect. 5.

2 Numerical Framework

A powerful method to numerically treat the non-linear Maxwell-Vlasov problem in six-dimensional phase space is the PIC approach [1, 5], which has a long history of more than five decades. The peculiarity of the PIC method is the ingenious particle-mesh technique for the coupling of an Eulerian grid-based model for the Maxwell equations with a Lagrangian approach for the Vlasov equation. To get an overview of the numerical methods applied within the PIC framework, a single PIC cycle, schematically depicted in Fig. 1, is discussed. The rarefied non-neutral plasma inside a device is represented by a sample of charged simulation particles of, in general, different species. In each time step, the electromagnetic fields are obtained by the numerical solution of the Maxwell equations. In the context of the present PIC solver, a discontinuous Galerkin (DG) method for these equations is applied where, in addition, a hyperbolic divergence cleaning technique [9] is considered. Especially the use of a powerful mixed nodal and modal DG approach [3, 4] allows a fast and high-order space discretization. Afterwards, the electromagnetic fields are interpolated to the actual spatial positions of the simulation particles [11]. These charged particles experience a force and thus an acceleration due to the electromagnetic fields. According to the Lorentz force, the charges are advanced and the new phase space coordinates are determined by numerically solving the usual laws of Newtonian dynamics. For this purpose we use an explicit low-storage fourth-order Runge-Kutta approach (LSERK4) [7]. To close the chain of self-consistent interaction, the simulation particles have to be located with respect to the computational grid in order to assign the contribution of each charged particle to the changed charge and current densities [11], which are then the sources for the Maxwell equations in the subsequent time step.




Fig. 1 Standard PIC concept

3 Highly Scalable Parallelization of the PIC Method

Highly parallel PIC simulations are often handicapped by the fact that in many simulated problems, the particles are locally confined. This causes load balancing problems if a simple domain decomposition approach is used for the parallelization. The presented load balancing algorithm is based upon the concept that on each MPI domain, the two different solvers (i.e. the particle and the grid solver) each require different computation times and that these computation times can be traded to some extent. On a domain with lots of particles, the particle solver might require a much longer time than the grid solver. On a different domain, the work load might be equal while on yet another domain there might be no particles at all, so the particle solver does not need any computation time. And yet, if the load balancing is chosen in such a way that computation times of Maxwell and particle methods on each domain sum up to an average that is about the same on all domains, the computational load is balanced across MPI domains and the parallel computation is still efficient.

However, the efficient load balancing relies on two assumptions: First, that an ideal average computational load can be identified, and second, that the computation can be distributed in such a way that all MPI domains receive a computational load close to the average. The following paragraphs briefly investigate whether these two assumptions are justified.

Identifying an optimal average computational load. The domain decomposition assigns elements and the particles that are located in these elements to the available MPI processes. Pre-determined weights predict the computational load for each element and for each particle. The decomposition is performed with the goal that the sum of the weights of all elements and particles on each MPI domain equals the average of the sum of the weights of all particles and elements in the simulation.

A DG element is weighted by its number of interpolation points. In order to find a suitable prediction for the computational load of a particle, a parameter study was conducted for the weighting factor of the particles versus a DG element. On 1,024

Fig. 2 A parameter study with a variation of particle weights shows that a ratio of about $w \approx 0.003$ yields the most efficient load balancing for 1,024 MPI processes





Fig. 3 For different grids, the maximum deviation from the medium (assumed-to-be ideal) weight is plotted over the number of MPI processes. The upper curves show simulations without load balancing while the lower curves show results from simulations with load balancing. The deviations are normalized by the average weight for each process



MPI processes, the weighting factor of a single particle vs. a single element was varied in order to find the best particle-element ratio. The results of the study as depicted in Fig. 2 show that a ratio $w = \frac{w_{part}}{w_{elem}}$ of about $w \approx 0.003$ yields the most efficient load balancing for the fourth-order test simulation of the gyrotron resonator. It should be noted that this ratio of $w \approx 0.003$ is an optimum for that specific simulation on 1,024 MPI processes. For simulations with different numbers of MPI processes, the ratio can vary slightly. For different computational grids, the ratio will probably vary a bit more. It will vary even more for different types of plasma devices.

Distribution of the computational weight: Deviation from the average. In order to characterize the quality of the load balancing, several simulations were inspected with regard to their distribution of computational load. The total weight w_{total} of all elements and particles in the computational domain was calculated using the ratio w mentioned above. An ideal weight for an MPI domain was then identified as $\frac{w_{total}}{n_{procs}}$ with the number of MPI domains n_{procs} . The total weight $w_{i_{procs}}$ of all particles and elements on each MPI domain i_{proc} was then compared to w_{total} , and




Fig. 4 Domain decomposition for varying MPI process numbers. (a) Load balanced decomposition into 32 MPI domains. (b) Load balanced decomposition into 2,048 MPI domains. (c) Decomposition into 2,048 MPI domains without load balancing. All domains contain between 98 and 99 elements. (d) Load balanced decomposition into 2,048 MPI domains. Domains with more than 140 elements are blanked

the maximum deviation is plotted in Fig. 3 for each simulation. As is to be expected, the deviation increases for increasing n_{procs} because the smallest delta of computational load that can be shifted from one domain to another is the load of a single element and the number of elements on each MPI domain decreases as n_{procs} increases. Likewise, the deviation increases with a decreasing number of grid cells in a computational mesh. From Fig. 3 it is also evident that the deviations of all load balanced simulations remain lower than the deviations of any non-load balanced simulation.


It is thus to be concluded that both assumptions are justified. In the following, the effects of the balancing of the computational load on the number of grid cells in each MPI domain are inspected. Visualizations of the domains resulting from the load balancing algorithm are shown in Fig. 4 for the resonant cavity simulation presented in Sect. 4. The color of each domain reflects the number of elements it contains.

It is evident that the outer domains generally contain more elements than the inner domains. This is because the particles are confined to the inner region close to the coaxial insert (a detailed description is given in Sect. 4.1, especially Fig. 8). The more particles a domain contains, the smaller it will be in terms of the number of grid cells in order to keep the total computational load approximately uniform across all MPI domains. In order to contrast the load balanced decomposition, Fig. 4c shows a domain decomposition without load balancing. Here, the number of grid cells is approximately uniform across all MPI domains: All domains contain 98 or 99 elements. It is clear, however, that in this decomposition, the Maxwell-only domains of the outer region will have to wait for one or more inner domains which have an additional computational load due to the particles of the electron beam. For the load balanced decomposition, these inner domains are highlighted in Fig. 4d where domains with more than 140 elements are blanked. The significant effect of the load balancing is evident because here some domains include only about 20 elements. Of course, this low number of elements results from the large number of particles contained in these domains.

4 Gyrotron Simulation

Gyrotrons are high-power milli- and micrometer wave generators. The gyrotron presented here is used in the context of fusion plasma heating, i.e. electron cyclotron resonance heating. A detailed overview to the state-of-the-art gyrotron research can be found in Ref. [13]. This section deals with the simulation of the 170 GHz resonant cavity which excites the $TE_{34,19}$ mode (as shown for an x-y-slice in Fig. 5). The $TE_{34,19}$ wave mode is generated by gyrating electrons and their interaction with their self-generated electromagnetic fields in what is known as the electron cyclotron maser instability. This interaction takes place in a cylindrical part of the gyrotron known as the resonant cavity or resonator. The radius of this cylindrical part has to be chosen depending on the desired wave mode. The gyrating electrons will then emit waves that are reflected from the walls and stimulate an azimuthal bunching of the electrons in a desired phase on the gyro-circle. The bunching amplifies the emitted electromagnetic waves at the desired frequency, due to a resonance effect between the gyro-frequency and the radio-frequency field. This process is described in detail by Kern [8] and Illy [6]. Simulations of a $TE_{2,3}$ resonant cavity using the simulation code presented here were introduced by Stock et al. [12]. Clearly, the excitation of the wave modes relies on the self-consistent particle-field interaction. Therefore, the resonant cavity was considered a suitable test case to demonstrate the capabilities of the coupled particle and Maxwell solver. Since the simulation deals with high frequency wave generation at small wavelengths, the computational demand as well as the requirements with respect to memory are enormous. Thus, it is also a good test case for the efficiency of the parallelization. This is the case especially because the particle distribution is not homogeneous but all particles are confined to the small fraction of the volume that is occupied by the electron beam.

Fig. 5 The E_{\perp} component of the $TE_{34,19}$ mode shows the 34 wavelengths in circumferential direction and the 19 wavelengths across the diagonal that form the indices of the TE wave mode



This accumulation of particles in a small part of the domain is a challenge to the balancing of the computational load. While one process might compute a region of the domain outside the beam which does not include particles at all, another region may include a large portion of the beam with a large number of particles inside.

The simulation of the resonant cavity is detailed in the following two sections, beginning with the description of the simulation setup and the computational grid in Sect. 4.1, followed by a presentation and discussion of the simulation results in Sect. 4.2, and concluded by the discussion of scalability aspects in Sect. 4.3.

4.1 Setup

The geometry for the presented simulation was taken from Piosczyk et al. [10]. During the startup of the experiments presented in Ref. [10], various wave modes are temporarily excited as the voltage is continually increased. However, we do not simulate the complete startup of the gyrotron since the establishment of the $TE_{34,19}$ mode can be reached much more quickly numerically than experimentally. Therefore, a different operating point was chosen for the numerical simulation than the working point presented in [10]. Preliminary simulations with the established cavity simulation tool SELFT [8] have shown that with the chosen parameter set, a $TE_{33,19}$ mode develops first, followed by the excitation and final establishment of the $TE_{34,19}$ mode.

The simulated coaxial cavity is a cylindrical tube with radius $r_o = 29.55$ mm and length $\Delta z_r = 16$ mm. Below the resonator ($z < 22$ mm), the outer wall is down-tapered at an angle of $\beta_1 = 3^\circ$ while above the resonator ($z > 38$ mm), the outer wall is up-tapered at an angle of $\beta_2 = 2.5^\circ$ as depicted in Fig. 6.




Fig. 6 Gyrotron resonator geometry

Inside the cavity, there is an insert, down-tapered at an angle of $\beta_i = 1^\circ$ with a radius of $r_i = 7.86$ mm at $z = 30$ mm. In the outer surface of the insert, there are 75 rectangular corrugations (see Fig. 6).

Boundary Conditions: Both the corrugated insert and the outer wall are modeled by a perfectly conducting boundary. Entry and exit at $z_0 = 0$ mm and $z_1 = 60$ mm, respectively, are modeled by non-reflecting boundaries. For the particles, entry and exit are open boundaries. Based on a current of $I = 65$ A, $n_{part,e} = 405,698$ particles with an MPF¹ of 10^6 are inserted per nanosecond at the entry at $z_0 = 0$ mm. These particles are distributed in space on a circular beam with radius $r_b = 10.0$ mm and length $\Delta z_e = v_{\parallel} \cdot \Delta t$. The particles gyrate around the externally applied magnetic field of $B_z = 6.88$ T with a Larmor radius of $r_g = 0.108$ mm (see Fig. 7). The angular positions of the inserted particles on the beam as well as the phase of their gyration around B_z is chosen at random with a uniform distribution. Their initial velocity of $v_{p,i} = 1.52 \times 10^8 \frac{\text{m}}{\text{s}}$ is determined from the acceleration due to an applied potential of $U_c = 82.2$ kV. The velocity is directed tangentially along a gyrating path around B_z . The relation of the circumferential velocity v_{\perp} around B_z versus the velocity in z -direction v_{\parallel} is given by the factor of $\alpha = \frac{v_{\perp}}{v_{\parallel}}$ which is fixed here to $\alpha = 1.1$.

Initial Conditions: The field solver is initialized with zero for the electromagnetic fields. The external magnetic field of $B_z = 6.88$ T is applied to the particles only and is not included in the Maxwell simulation. The particles are initialized with the same beam parameters used for the emission also. However, for the initialization, the beam is created throughout the whole domain with $\Delta z = 60$ mm. The initial number of particles is $n_{part,0} = n_{part,e} \frac{\Delta z}{v_{\parallel}}$.

Table 1 summarizes the relevant parameters for the resonator simulation. Note that for the scalability studies shown in Sect. 4.3, different computational parameters were used. This is due to the fact that the scalability analysis was done while a work-

¹Macro particle factor (MPF), i.e. numbers of real particles per simulated particle.

Fig. 7 Particle parameters for the electron beam




Table 1 Parameters for the 170 GHz resonant cavity simulation

Symbol	(Unit)	Value	Symbol	(Unit)	Value
Geometry of the boundaries			Computational parameters		
z_0	(mm)	0.0	MPF	(—)	10^6
z_1	(mm)	60.0	$n_{Part,0}$	(—)	237,668
Δz_r	(mm)	16.0	$n_{Part,e}$	($\frac{1}{ns}$)	405,698
r_o	(mm)	29.55	\mathcal{O}	(—)	5–6
r_i	(mm)	7.86	Computational parameters for the scalability study		
Technical parameters			MPF	(—)	10^4
U_c	(kV)	82.2	$n_{Part,0}$	(—)	29,191,600
I	(A)	55.0	$n_{Part,e}$	($\frac{1}{ns}$)	46,811,300
B_z	(T)	6.88	\mathcal{O}	(—)	4
Beam parameters					
r_b	(mm)	10.0			
r_g	(mm)	0.108			
α	(—)	1.1			
γ	(—)	1.176			

ing setup for the simulation had not yet been found. For all resonator simulations presented in this work, Cell Mean Value particle-grid coupling was used [11].

Computational Grids: Due to their small geometric extensions in radial and circumferential direction (0.44 and 0.35 mm, respectively), the corrugations in the insert dictate the minimum size of grid cells (and thus also the time step). The most natural way to discretize these corrugations is with hexahedral elements. The cavity itself was also meshed using hexahedra. Building upon the experiences from the TE_{34,19} waveguide and launcher simulations, an edge length of $\Delta x \approx 1$ mm was chosen for the cavity, yielding a total of 201,720 hexahedral grid cells.

The simulation was started with a spatial convergence order of 5. In the course of the simulation, however, it was realized that the polynomial degree of the basis functions in the fifth-order simulation was too small, resulting in a filtering of the higher wave modes. Therefore, the spatial order of the scheme was changed during the simulation from $\mathcal{O}5$ to $\mathcal{O}6$ at 11 ns which yielded the expected results.

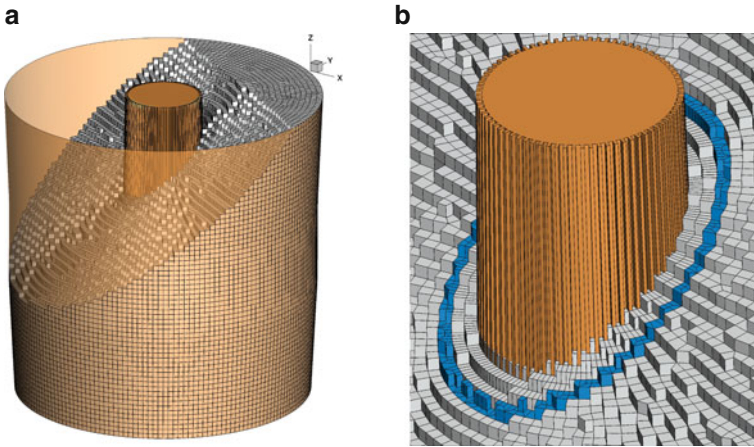


Fig. 8 Compute grid for coupled simulation of the $TE_{34,19}$ resonant cavity. (a) Complete view of the computational domain. (b) Detailed view on the computational domain with the zone containing particles marked *blue*

The computational grid is depicted in Fig. 8. Figure 8a shows a cut-view of the mesh while Fig. 8b provides a detailed view of the corrugated insert, showing also in blue the grid cells containing particles.

4.2 Resonator Simulation Results

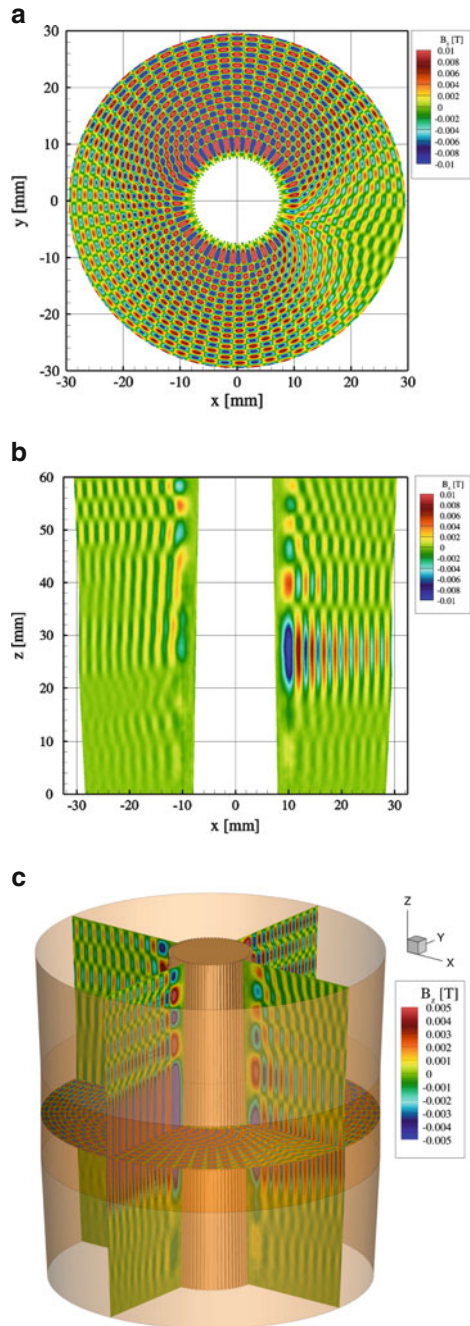
In the presented simulation, the $TE_{34,19}$ mode established after about 33 ns as shown in Fig. 9 for $t = 36$ ns. Figure 9 shows that the $TE_{34,19}$ wave mode really is in a stable state: The regular pattern of the $TE_{34,19}$ mode dominates the domain, most notably the resonator itself and the uptapered region above the resonator.

A result of more general interest in the resonator simulation is the computation time required per nanosecond. A sixth-order simulation on 512 CPU cores requires a wall-clock time of 16,100 s per nanosecond simulation time on an Intel® Xeon X5570 (Nehalem) cluster with Infiniband network.

4.3 Scaling Experiments

In the presented 170 GHz resonator simulation, particle and Maxwell computation each require a significant part of the total computation time. Moreover, particle computation is confined to the region of the electron beam (see Fig. 8b). Even on small numbers of MPI domains, there are domains with more than 50% particle

Fig. 9 The wave mode at $t = 36.0$ ns. **(a)** B_z on an x-y-slice at $z = 30$ mm. **(b)** B_z on an x-z-slice at $y = 0$ mm. **(c)** B_z on three slices



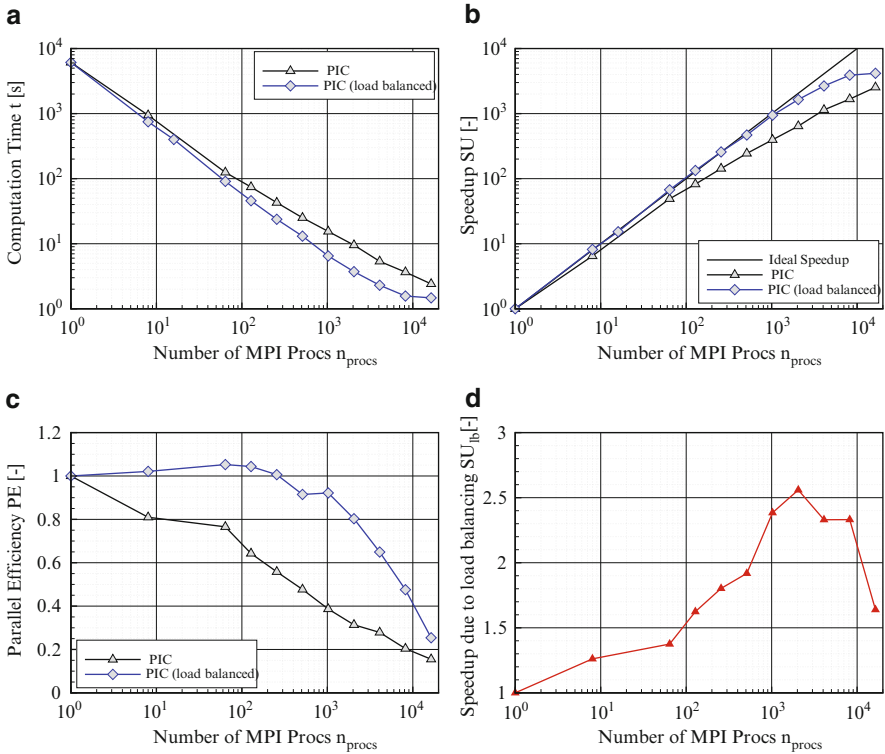


Fig. 10 Scalability curves resonator (strong scaling). (a) Simulation time resonator. (b) Speedup resonator. (c) Parallel efficiency resonator. (d) Speedup due to load balancing

computation and, at the same time, domains with 100 % Maxwell computation. Therefore, efficient load balancing is the key to any efficient parallel computation of the problem.

Measurements were taken as pure computation times, i.e. they represent the wall clock time, not including initialization or I/O. The reason for this is that the simulation times can then be normalized as simulation time per time step. This normalized time provides a more general quantity, applicable for simulation runs of different length. Moreover, in a typical 170 GHz gyrotron resonator simulation, initialization, I/O and other computations at analyze time levels require less than 1 % of the total wall clock time.

Furthermore, it should be noted that the scalability studies with the 170 GHz gyrotron resonator were performed when a working parameter set for the successful simulation had not been found yet. Apart from small changes in the physical parameters, the computationally most notable differences are the MPF of 10^4 and the spatial order of the Maxwell scheme of $\mathcal{O}4$ which were used for the scalability studies presented below. We note, therefore, that the computational load in the

scalability studies is shifted more towards the particle algorithms. In other words, in the simulation presented in Sect. 4.2 with the parameters summarized in Table 1, the Maxwell solver requires more computation time while the particle methods require less computation time compared to the simulations of the scalability studies. We emphasize that this does not affect the relevance of the results and their interpretations. Rather, the scalability simulations were performed with a more challenging setup with respect to the load balancing than actually required by the simulation presented in Sect. 4.2.

The performance of the parallelization is best shown in the strong scaling study depicted in Fig. 10. The study was conducted on the mesh described in Sect. 4.1 with spatial order 4. Serving as reference, one strong scaling study was done using a compilation without the particle algorithms, here referred to as “Maxwell only.” PIC simulations were done both with and without load balancing.

The Maxwell only simulation scales almost perfectly from 1 to 16,384 MPI processes. The PIC simulations scale up to 16,384 processes but the decline in the speedup shows that a significant scalability beyond 16,384 MPI processes is not to be expected. The reason is that here, some MPI domains include one single element only. Thus, the computational load cannot be further distributed or balanced without a hybrid parallelization. Further speedups can most likely be achieved if an additional shared memory parallelization is used. Over a wide range of MPI process numbers, the load balanced PIC scheme is significantly faster than the scheme without load balancing. In the range from 1,024 to 8,192 processes, the load balanced simulations are more than twice as fast as the non-load balanced computations (as shown in Fig. 10d).

5 Conclusions and Outlook

We demonstrated the ability of our PIC code to simulate a 170 GHz gyrotron resonator with the currently available HPC capabilities. Since the 170 GHz gyrotron plays a crucial role in the development of future fusion reactors like ITER [13] we expect a high relevance of our work for the research in the field of gyrotron design. An efficient parallelization approach is crucial for achieving any results for a problem of the size presented here. Without the presented load balancing approach the ability to waste HPC capabilities can be enormous. A parallelization of the PIC method when dealing with problems of this size is therefore an essential feature which has to be centrally focused on in development and research. We believe that computations like the presented resonator simulation are essentially needed if open problems like after cavity interactions or beam misalignment are to be investigated. In future work we expect to verify a recent theory which questions coaxial super power gyrotron feasibility [2], concluding that maximal values of the azimuthal index exist, beyond which stationary single mode operation of gyrotron is not possible, due to the onset of stochastic oscillations of the radio-frequency fields: This issue definitely needs deeper investigations with a full-wave time-domain

code like the presented method. Both tasks of this challenging project are of great importance in the field of gyrotron development.

Acknowledgements We gratefully acknowledge the Deutsche Forschungsgemeinschaft (DFG) for funding within the project “Numerical Modeling and Simulation of Highly Rarefied Plasma Flows”. T. Stindl wishes to thank the Landesgraduiertenförderung Baden-Württemberg and the Erich-Becker-Stiftung, Germany, for their financial support. Computational resources have been provided by the Bundes-Höchstleistungsrechenzentrum Stuttgart (HLRS).

References

1. C.K. Birdsall and A.B. Langdon. *Plasma Physics via Computer Simulation*. Adam Hilger, Bristol, Philadelphia, New York, 1991.
2. O. Dumbrajs and G. S. Nusinovich. Coaxial gyrotrons: past, present, and future (review). *IEEE Transactions on Plasma Science*, 32(3):934–946, June 2004.
3. G. Gassner, F. Lörcher, C.-D. Munz, and J. S. Hesthaven. Polymorphic nodal elements and their application in discontinuous Galerkin methods. *J. Comput. Phys.*, 228(5):1573–1590, 2009. doi:10.1016/j.jcp.2008.11.012.
4. J.S. Hesthaven and T. Warburton. *Nodal Discontinuous Galerkin Methods*. Springer, New York, 2008.
5. R. Hockney and J. Eastwood. *Computer Simulation using Particles*. McGraw-Hill, New York, 1981.
6. S. Illy. *Untersuchungen von Strahlinstabilitäten in der Kompressionszone von Gyrotron-Oszillatoren mit Hilfe der kinetischen Theorie und zeitabhängiger Particle-in-Cell-Simulationen*. PhD thesis, Universität Karlsruhe und Forschungszentrum Karlsruhe (FZKA 6037), December 1997.
7. Christopher A. Kennedy, Mark H. Carpenter, and R. Michael Lewis. Low-storage, explicit Runge-Kutta schemes for the compressible Navier-Stokes equations. *Applied Numerical Mathematics*, 35:177–219, 2000.
8. S. Kern. *Numerische Simulation der Gyrotron-Wechselwirkung in koaxialen Resonatoren*. PhD thesis, Forschungszentrum Karlsruhe GmbH, FZKA 5837, 1996.
9. C.-D. Munz, P. Omnes, R. Schneider, E. Sonnendrücker, and U. Voß. Divergence correction techniques for Maxwell solvers based on a hyperbolic model. *J. Comput. Phys.*, 161:484–511, 2000.
10. B. Piosczyk, G. Dammertz, O. Dumbrajs, O. Drumm, S. Illy, J. Jin, and M. Thumm. A 2-MW, 170-GHz coaxial cavity gyrotron. *IEEE Transactions on Plasma Science*, 32(2):413–417, April 2004.
11. T. Stindl, J. Neudorfer, A. Stock, M. Auweter-Kurtz, C.-D. Munz, S. Roller, and R. Schneider. Comparison of coupling techniques in a high-order discontinuous Galerkin based particle in cell solver. *J. Phys. D: Applied Physics*, 44:194004, 2011.
12. A. Stock, J. Neudorfer, M. Riedlinger, G. Pirrung, G. Gassner, R. Schneider, S. Roller, and C.-D. Munz. Three-dimensional numerical simulation of a 30 GHz gyrotron resonator with an explicit high-order discontinuous galerkin based particle-in-cell method. *accepted for publication in: IEEE Transaction on Plasma Science*, XX:yyy–yyy, 2012.
13. M. Thumm. Progress on gyrotrons for iter and future thermonuclear fusion reactors. *IEEE, Transaction on Plasma Science*, 39(4):971–979, April 2011.