

CFD/CAA Simulations on HPC Systems

Michael Schlottke-Lakemper, Fabian Klemp, Hsun-Jen Cheng,
Andreas Lintermann, Matthias Meinke and Wolfgang Schröder

Abstract In this paper, a highly scalable numerical method is presented that allows to compute the aerodynamic sound from a turbulent flow field on HPC systems. A hybrid CFD-CAA method is used to compute the flow and the acoustic field, in which the two solvers are running in parallel to avoid expensive I/O operations for the acoustic source terms. Herein, the acoustic perturbation equations are solved by a high-order discontinuous Galerkin scheme using the acoustic source terms obtained from an approximate solution of the Navier-Stokes equations. Both solvers run simultaneously and operate on differently refined hierarchical Cartesian grids. This direct-hybrid method is validated by monopole and pressure pulse simulations and is used for performance measurements on current HPC systems. The results highlight the limitations of classic hybrid methods and show that the new approach is suitable for highly parallel simulations.

M. Schlottke-Lakemper (✉) · A. Lintermann
Jülich Aachen Research Alliance - High Performance Computing, RWTH Aachen University,
Aachen, Germany
e-mail: m.schlottke-lakemper@aia.rwth-aachen.de

A. Lintermann
e-mail: lintermann@jara.rwth-aachen.de

F. Klemp · H.-J. Cheng · M. Meinke · W. Schröder
Institute of Aerodynamics, RWTH Aachen University, Aachen, Germany
e-mail: f.klemp@aia.rwth-aachen.de

H.-J. Cheng
e-mail: h-j.cheng@aia.rwth-aachen.de

M. Meinke
e-mail: m.meinke@aia.rwth-aachen.de

W. Schröder
e-mail: office@aia.rwth-aachen.de

1 Introduction

One of the major challenges of today's aircraft development is noise reduction, which is also one of the central aims in European aircraft policy. The perceived noise levels of flying aircraft are to be reduced until 2050 by 65 % compared to the year 2000 [25]. Many sound-generating components of aircraft need to be assessed in sufficient detail to be able to improve their design, such as the optimization of the jet nozzle geometry to lower noise emissions at take-off without sacrificing the thrust efficiency. To achieve such optimizations, efficient, fully parallelized algorithms are needed to predict the flow field and the far-field noise of jet engines.

A hybrid method combining large-eddy simulation (LES) with computational aeroacoustics (CAA) for large-scale aeroacoustics simulations has been successfully applied in [7, 18]. It uses LES to determine the turbulent flow field for external flow configurations. From this solution, noise-generating source terms are extracted and used in a CAA simulation, where the acoustic field is predicted using the acoustic perturbation equations (APE) [6]. This scheme has been applied successfully to different problems in computational aeroacoustics, such as trailing edge noise [7], jet noise [12], or combustion noise [4, 11]. However, it suffers from the exchange of large data volumes for the acoustic source terms via I/O operations, which limits the efficiency of such a two-step approach especially on high-performance computing (HPC) systems.

To circumvent this bottleneck, the direct-hybrid method presented in this work combines the LES and CAA solvers in a single framework such that both solvers can run in parallel. The LES solver used for the prediction of the flow field is based on a finite-volume method, while the CAA approach makes use of a high-order discontinuous Galerkin (DG) method to solve the APE for the acoustic field. DG methods were first described by Reed and Hill [24] and were subsequently applied to various physical problems, such as incompressible and compressible flow [2, 22], magnetohydrodynamics [29], and aeroacoustics [1, 3].

The LES and CAA computations are performed on a joint Cartesian mesh. Based on a coloring scheme, cells are associated with different weights for the LES and CAA solution and a space-filling curve is used for the domain decomposition. The coupling mechanism between both simulations only requires memory transfer operations. That is, no additional communication between the subdomains is necessary, leading to an efficient algorithm to be used on massively parallel systems. Furthermore, this direct-hybrid approach allows a more fine-grained control over the coupling process itself, since the LES results are not obtained separately from the acoustic field anymore. This means that, e.g., the time step size or the grid size can be adapted during the simulation to account for time-dependent changes in the resolution requirements of both solvers, enabling in situ optimizations of the simulation process.

In this paper, the coupling approach for the direct-hybrid LES-CAA simulation is presented and results for performance measurements are shown. A CAA code is developed and integrated with an existing LES solver. After the governing equations are introduced in Sect. 2, the numerical methods are described in Sect. 3. In Sect. 4,

the coupling strategy is discussed in detail. The CAA solver is validated in Sect. 5, before it is used for strong scaling experiments on two state-of-the-art HPC systems. In Sect. 6, the presented methods and the obtained results are summarized.

2 Governing Equations

In this hybrid CFD-CAA method, two sets of governing equations are utilized. One solely describes the generation and propagation of acoustic waves, while the other set of equations predicts the physics of the underlying flow field. Here, the acoustic perturbation equations are used for the acoustic field and the Navier-Stokes equations for the flow field. Both are briefly summarized in the following.

2.1 Navier-Stokes Equations

The Navier-Stokes equations in non-dimensional, conservative form are given by

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla (\rho \mathbf{u}) &= 0, \\ \frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \left(\rho \mathbf{u} \mathbf{u} + p + \frac{\boldsymbol{\tau}}{\text{Re}_0} \right) &= 0, \\ \frac{\partial \rho e}{\partial t} + \nabla \left((\rho e + p) \mathbf{u} + \frac{1}{\text{Re}_0} (\boldsymbol{\tau} \mathbf{u} + \mathbf{q}) \right) &= 0. \end{aligned} \quad (1)$$

The quantity ρ represents the fluid density, \mathbf{u} the velocity vector, and e the total specific energy. The system in Eq. (1) is closed by the definition of the total specific energy for a perfect gas

$$\rho e = \frac{p}{\gamma - 1} + \frac{1}{2} \rho (\mathbf{u} \cdot \mathbf{u}), \quad (2)$$

where p is the pressure and γ is the specific heat ratio. For non-dimensionalization, the stagnation state is employed, which is denoted by the subscript 0. The Reynolds number based on the stagnation state is defined by

$$\text{Re}_0 = \frac{\rho_0 c_0 L}{\mu_0}, \quad (3)$$

where L is a reference length and ρ_0 , c_0 , and μ_0 are the stagnation density, the speed of sound, and the dynamic viscosity. A Newtonian fluid is assumed such that the components τ_{ij} of the stress tensor $\boldsymbol{\tau}$ can be written as

$$\tau_{ij} = -2\mu S_{ij} + \frac{2}{3}\mu S_{ij}\delta_{ij}, \quad (4)$$

where $S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$ is the rate of strain tensor. The dynamic viscosity μ is calculated by using Sutherland's law and the vector of heat conduction \mathbf{q} is determined by Fourier's law

$$\mathbf{q} = -\frac{k}{\text{Pr}(\gamma - 1)} \nabla T, \quad (5)$$

where T is the static temperature. The Prandtl number is defined with the specific heat at constant pressure c_p by $\text{Pr} = \frac{\mu_0 c_p}{k_0}$. For a constant Prandtl number, the relation $k(T) = \mu(T)$ holds for the thermal conductivity.

2.2 Acoustic Perturbation Equations

The acoustic perturbation equations (APE) were introduced in [6] and are used to predict the acoustic field for flow-induced noise. They are derived from the linearized Euler equations and modified to retain only acoustic modes without generating vorticity or entropy modes. Neglecting all viscous, non-linear and entropy-related contributions, the APE-4 system reads [6]

$$\frac{\partial \mathbf{u}'}{\partial t} + \nabla (\bar{\mathbf{u}} \cdot \mathbf{u}') + \nabla \left(\frac{p'}{\bar{\rho}} \right) = \mathbf{q}_m, \quad (6)$$

$$\frac{\partial p'}{\partial t} + \bar{c}^2 \nabla \cdot \left(\bar{\rho} \mathbf{u}' + \bar{\mathbf{u}} \frac{p'}{\bar{c}^2} \right) = 0, \quad (7)$$

where the source term \mathbf{q}_m is the linear Lamb vector

$$\mathbf{q}_m = -(\boldsymbol{\omega} \times \mathbf{u})' = -(\boldsymbol{\omega}' \times \bar{\mathbf{u}} + \bar{\boldsymbol{\omega}} \times \mathbf{u}'), \quad (8)$$

with $\boldsymbol{\omega}$ as the vorticity vector. The variables of the APE are perturbed quantities denoted by prime (\cdot)' and are defined by $\phi' := \phi - \bar{\phi}$, where the bar ($\bar{\cdot}$) denotes time-averaged quantities.

In the present work, the non-dimensional form of Eqs. (6) and (7) is used. As for the Navier-Stokes equations, the stagnation state is used for the definition of reference values. Furthermore, it is assumed here that the time-averaged values for the speed of sound and density are constant and equal to the stagnation state, i.e., $\bar{c} = c_0$ and $\bar{\rho} = \rho_0$, which is only valid in the low-Mach number regime. By using the following non-dimensional variables,

$$\tilde{t} = \frac{t c_0}{L}, \quad \tilde{x} = \frac{x}{L}, \quad \tilde{\mathbf{u}} = \frac{\mathbf{u}}{c_0}, \quad \tilde{p} = \frac{p}{\rho_0 c_0^2}, \quad (9)$$

the APE can be written as

$$\frac{\partial \tilde{\mathbf{u}}'}{\partial \tilde{t}} + \tilde{\nabla} \cdot (\tilde{\mathbf{u}} \cdot \tilde{\mathbf{u}}' + \tilde{p}') = \tilde{\mathbf{q}}_m, \tag{10}$$

$$\frac{\partial \tilde{p}'}{\partial \tilde{t}} + \tilde{\nabla} \cdot (\tilde{\mathbf{u}}' + \tilde{\mathbf{u}} \tilde{p}') = 0. \tag{11}$$

The non-dimensional source term is given by $\tilde{\mathbf{q}}_m = \frac{q_m}{c_0^2/L}$. For convenience, in the following discussion the tilde is dropped from the non-dimensional quantities.

3 Numerical Methods

In this section, the meshing process and the domain decomposition are outlined. Furthermore, the numerical methods for the acoustic perturbation equations and the Navier-Stokes equations are briefly described.

3.1 Hierarchical Mesh Topology

Both the LES solver and the CAA solver operate on a joint hierarchical Cartesian mesh. The cells of the grid are organized in a tree structure (2D: quadtree, 3D: octree), with parent-child relationships between different levels and neighbor relationships within a level. The discretization process follows the method described in [21] and starts with a single square/cube cell which encloses the whole computational domain. This zero-level cell is then refined uniformly until the desired refinement level is reached (see Fig. 1a). A cell to be refined is isotropically subdivided into 2^d square/cube cells, with d being the number of spatial dimensions and with the original cell becoming the parent cell of the new child cells. Individual regions of

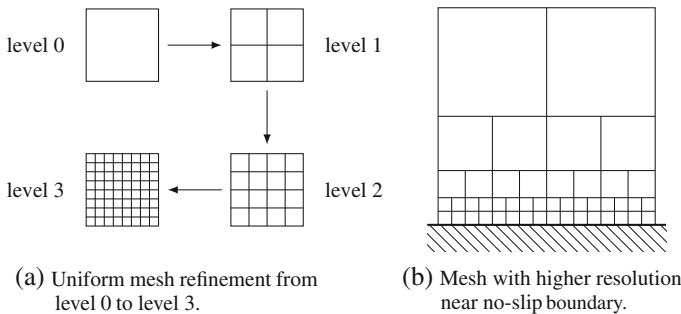



Fig. 1 Cell refinement for a hierarchical Cartesian grid in 2D

Fig. 2 Domain partitioning on two domains with four subtrees starting at level l_α



the mesh can be further refined to meet resolution requirements, e.g., in areas with small-scale physical features such as wall-bounded shear layers or to accurately resolve boundaries (see Fig. 1b). A smoothing algorithm ensures that the level difference between neighboring cells does not exceed one, i.e., each cell has at most 2^{d-1} neighbor cells in each spatial direction. Special treatment is necessary for cells that are intersected by the body geometry. In this paper, only non-intersected cells are considered. During grid generation, the zero-level cell is homogeneously refined to a minimum level l_α and all coarser cells are discarded [21]. These cells at level l_α become the roots of their subtrees and are further subdivided until the required refinement level is reached.

For the domain decomposition, a Hilbert space-filling curve [26] is used to map the grid at level l_α to the interval $[0, 1]$. Each cell at level l_α is assigned a load that depends on the number of cells in its subtree and on the type of the cells, i.e., whether they are LES or CAA cells. Load balancing is achieved by taking into account these load values when distributing the cells among the processes and for each l_α cell the entire subtree is placed on the same rank (see Fig. 2). By consecutively placing l_α cells and their subtrees on the MPI ranks according to their position on the Hilbert curve, spatial compactness is ensured, reducing the overall communication cost.

3.2 Discontinuous Galerkin Approximation of the APE

A discontinuous Galerkin spectral element method (DGSEM) is used to determine the acoustic field. In Kopriva et al. [19], the DGSEM was proposed and has been used extensively [9, 17]. Since it was derived for quadrilateral/hexahedral mesh elements, it is well-suited for the use on hierarchical Cartesian grids. Furthermore, its compact formulation allows a very efficient parallelization, when explicit time stepping is used, and the parallel efficiency is independent of the chosen order of the scheme.

Since the DGSEM elements correspond to cells in a finite-volume context, the words cell or element will be used interchangeably. In the following, the main components of the DGSEM are outlined. First, the system of equations is mapped to a reference element for efficiency reasons. The derivation of the DG formulation

then starts with the weak formulation, choosing Lagrange polynomials to represent the solution within each element. This gives rise to an integral equation, which is approximately solved using Gauss quadrature. Finally, the discrete DG operator is integrated in time using a Runge-Kutta scheme.

A general system of hyperbolic conservation equations in three dimensions reads

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{f}(\mathbf{U}) = 0, \quad (12)$$

where $\mathbf{U} = \mathbf{U}(\mathbf{x}, t)$ is the vector of conservative variables $\{u_i\}_{i=1}^{n_v}$ and \mathbf{f} is the flux vector. For efficiency reasons, the differential equation is mapped to a reference element E , which is in three dimensions given by a cube of size $[-1, 1] \times [-1, 1] \times [-1, 1]$. Introducing the reference coordinate vector $\boldsymbol{\xi} = (\xi^1, \xi^2, \xi^3)^\top$, the final transformed equation reads [17]

$$\hat{J} \mathbf{U}_t + \nabla_{\boldsymbol{\xi}} \cdot \mathbf{f} = 0, \quad (13)$$

where \hat{J} is the Jacobian, which for cube-to-cube transformations is just $\frac{h}{2}$, h being the side length of the cube, and \mathbf{U}_t is the time derivative of the vector of conservative variables.

The derivation of the DG method starts with the weak form of the equation. Therefore, Eq. (13) is multiplied by a test function $\phi = \phi(\boldsymbol{\xi})$ and integrated over the reference element E

$$\int_E (\hat{J} \mathbf{U}_t + \nabla_{\boldsymbol{\xi}} \cdot \mathbf{f}) \phi \, d\boldsymbol{\xi} = 0. \quad (14)$$

Using integration by parts on the flux term, the weak formulation of the differential equation is obtained

$$\int_E \hat{J} \mathbf{U}_t \phi \, d\boldsymbol{\xi} + \int_{\partial E} (\mathbf{f} \cdot \mathbf{n})^* \phi \, ds - \int_E \mathbf{f} \cdot \nabla_{\boldsymbol{\xi}} \phi \, d\boldsymbol{\xi} = 0, \quad (15)$$


where \mathbf{n} is the surface normal vector in the reference system. Similar to the finite-volume approach, the value for the normal flux $\mathbf{f} \cdot \mathbf{n}$ is not uniquely defined on the element boundaries ∂E , since the solutions in the left \mathbf{U}^- and right \mathbf{U}^+ elements are discontinuous. Therefore, a numerical flux $(\mathbf{f} \cdot \mathbf{n})^* = \mathbf{g}(\mathbf{U}^+, \mathbf{U}^-, \mathbf{n})$ is chosen that combines values from both sides to a single flux. In this work, the local Lax-Friedrichs flux formulation is used,

$$\mathbf{g}(\mathbf{U}^+, \mathbf{U}^-, \mathbf{n}) = \frac{1}{2} (\mathbf{f}(\mathbf{U}^+) + \mathbf{f}(\mathbf{U}^-)) \cdot \mathbf{n} + \frac{1}{2} \left(\max_{U \in [\mathbf{U}^+, \mathbf{U}^-]} |\mathbf{a}(U) \cdot \mathbf{n}| (\mathbf{U}^+ - \mathbf{U}^-) \right), \quad (16)$$

where \mathbf{a} is the vector of eigenvalues of the flux Jacobian. The solution \mathbf{U} is approximated by a polynomial basis

$$\mathbf{U}(\boldsymbol{\xi}, t) \approx \sum_{i,j,k=0}^N \bar{\mathbf{u}}_{ijk}(t) \psi_{ijk}(\boldsymbol{\xi}), \quad \psi_{ijk}(\boldsymbol{\xi}) = l_i(\xi^1) l_j(\xi^2) l_k(\xi^3), \quad (17)$$

Fig. 3 Legendre-Gauss nodes in a 2D reference element for $N = 3$



where the basis functions ψ_{ijk} are the product of one-dimensional Lagrange polynomials l of degree N in each spatial direction and $\bar{u}_{ijk}(t)$ are the coefficients to be determined. The nodal basis is defined on a set of interpolation points $\{\xi\}_{i=0}^N$ on the interval $\xi \in [-1, 1]$, which in this work are the Legendre-Gauss nodes (Fig. 3). For the fluxes, the same approach is used for the approximation.

The three integrals in Eq. (15) are approximated by Gauss quadrature. Generally, the Gauss quadrature of an arbitrary function $f(x)$ on the interval $[a, b]$ with $N + 1$ nodes can be written as

$$\int_a^b f(x) dx \approx \sum_{i=1}^N \omega_i f(x_i), \quad (18)$$

where the weights ω_i and the integration nodes x_i are specific to the chosen quadrature type. These weights are pre-calculated and stored to make the algorithm efficient. With the interpolation points $\{\xi_i\}$ collocated at the Gauss nodes, all sums collapse into single values, yielding the discrete DG operator $\mathcal{L}(\mathbf{U}, t) = \mathbf{U}_t$ [17]. In the next step, the semi-discrete formulation is integrated in time to obtain the solution at the next time step, for which a low-storage fourth-order Runge-Kutta scheme is used [5].

3.3 Finite-Volume Method for the Flow Simulation

A second-order finite-volume method is used to solve the unsteady Navier-Stokes equations for compressible flow as given in Sect. 2.1. The solver has been extensively validated and used for various flow problems previously [15, 16]. A detailed description of the method can be found in [13, 15, 16, 28].

4 Coupling Strategy

To solve the acoustic perturbation equations, the averaged quantities $\bar{\mathbf{u}}$ and \bar{c} and the source term \mathbf{q}_m have to be determined first. The flow solution is advanced without coupling until the averaged quantities are statistically converged. The coupling process for each time step of the LES reads:

1. Advance the LES solution.
2. Calculate the source terms from instantaneous and averaged quantities.
3. Advance the CAA solution.

The actual coupling takes place via the source terms computed from the LES solution, which are then used to solve the APE. This means that there is a one-way coupling from the flow solution to the acoustic field, while the flow solution is not influenced by the acoustic field.

In the direct-hybrid method described here, the LES and the CAA simulation are both performed within a single simulation framework and by using the same grid topology. This makes certain aspects of the coupling process more efficient and allows a more fine-grained control over the interface between the two solvers. In the following, some details of the method are presented.

4.1 Spatial Coupling

The instantaneous variables of the source term q_m are available after each time step from the flow simulation. They have to be transferred, however, from the LES to the acoustic grid. Since both simulations typically operate on different levels of the same grid, identification of corresponding cells is possible by traversing the octree constituting the hierarchical Cartesian mesh. While LES and CAA leaf cells can generally be of different size, the coupling always happens within a single subtree. Since the domain decomposition algorithm distributes entire subtrees on different processes (see also Sect. 3.1), no additional inter-rank communication is required for the exchange of data between CFD and CAA cells.

This type of mesh also guarantees that there are no partially overlapping cells, i.e., a smaller cell is always fully contained inside a larger cell. Note that the DG elements are generally of higher order than the finite-volume cells. Depending on the resolution of the fluid and acoustics problems, four types of transformations are possible.

In the simplest case, one fluid cell corresponds exactly to one acoustics cell (Fig. 4a). That is, the source term is calculated once in the finite-volume part and the same value is used at all Gauss nodes of the DG element. This approach is used exclusively in the present work, i.e., no spatial interpolation is performed. Similar to the one-to-one mapping, the source term is calculated once and then used at all Gauss nodes of all elements if one fluid cell is mapped to multiple acoustics cells (Fig. 4b).

Having multiple finite-volume cells mapped onto one DG element (Fig. 4c) requires the values at the Gauss nodes to be interpolated from several flow cells. A natural choice would be to interpret the finite-volume cells as equidistant nodes of a polynomial and to obtain the values at the Gauss nodes through projection. This, however, can lead to spurious oscillations if the number of finite-volume cells and thus the polynomial degree is high, especially in regions with large flow gradients.




Fig. 4 Possible spatial mappings for coupled simulations. Aeroacoustics cells (*top*) are *white*, fluid cells (*bottom*) are *grey*

Other possibilities are weighted least squares methods, nearest neighbor interpolation, or inverse distance weighting. Which approach is best depends on a number of factors. A practical consideration is the computational cost of the chosen method, e.g., whether the effort scales linearly with the number of degrees of freedom or worse, since the interpolation has to take place at each flow simulation time step. The smoothness of the interpolated function is also important, especially in high-gradient zones. Furthermore, it is desirable to have a conservative interpolation scheme such as proposed by Farrell and Maddison [8], to avoid distorting the source terms.

If there are regions without either a flow or acoustics grid, no coupling is performed. If only acoustic cells exist, far-field values for the averaged quantities \bar{c} and \bar{u} have to be specified for the APE, e.g., the freestream values from the flow field. The source term \mathbf{q}_m is set to zero with a smooth transition from non-zero to zero values.

4.2 Temporal Coupling

The coupling between the flow and the acoustics simulations has to be realized at each time step. Due to the explicit global time stepping it is possible that the time step size differs between the two solvers. In this case, at each time step the source term from the LES solution needs to be interpolated to the simulation time of the CAA solver.

Depending on the features of the geometry, the time step for the aeroacoustics simulation may be smaller than that for the flow simulation or vice versa and thus the source terms have to be interpolated between two flow time steps. As for the spatial coupling, there are many different interpolation methods to choose from. Linear interpolation is the most straightforward approach, with sometimes inferior results. Several temporal interpolation methods suitable for hybrid aeroacoustics simulations are compared and evaluated by Geiser et al. [10] and least-squares optimized interpolators were found to have the best properties when it comes to broadband error reduction.

The simplest approach is using the same time step for both simulations, which requires no interpolation between the two datasets. In this case, the next time step based on the CFL condition is determined for the CFD and the CAA method and the minimum of both methods is used, which is also the procedure that is used in this work.

4.3 Data Transfer

There exist two options for transferring data between the flow solution and the acoustics solution: via data files written to disk, i.e., offline coupling as used in standard hybrid approaches, or through in-memory data access, i.e., online coupling as done in the new direct-hybrid approach. Both methods are discussed in the following.

In offline coupling, the processes of obtaining the flow solution and running the aeroacoustics simulation are completely separated. At first, the flow solution is obtained and the source term q_m is written to a file at certain time intervals. During the acoustics simulation, the source terms are determined from the files by interpolation in time. Conceptually, this is the simplest approach, since except for the I/O routines nothing has to be changed inside the two simulations. However, the high amount of data that has to be transferred to and from the disk makes this method expensive in terms of computational cost, especially for large-scale simulations on thousands of cores. However, it is also the first step towards a simulation which makes use of online coupling as outlined next.

In online coupling, the flow and the acoustics simulations are fully integrated and run synchronously at the same time. Typically, the flow solution will be advanced by one time step and the acoustics solution has to be updated until they are both synchronized. Since no files have to be written to disk, this approach is more efficient than offline coupling. If the acoustics cells are kept on the same computational core as the corresponding flow cells, the acoustics simulation can directly access the relevant information by simple memory transfer operations. This locality of data is achieved by the specific subdomain decomposition, which operates on the joint LES-CAA grid. On the other hand, the increased memory consumption makes it necessary to use more computational cores. Furthermore, due to the different number of operations for the finite-volume and the DG operator, paired with different numbers of flow cells per acoustics cell, load balancing between the cores becomes mandatory to achieve reasonable parallel efficiency. This is accomplished by assigning appropriate loads to the fluid and acoustics cells.

5 Results

The CFD solver has already been extensively tested and used in the past, e.g., in [13, 15, 16, 23, 28]. Thus in Sect. 5.1, only the new CAA solver is validated. Additionally, parallel performance results for the CAA solver are presented in Sect. 5.2.

5.1 Validation of the Aeroacoustics Solver

The DG method described in Sect. 3.2 is validated by solving the acoustic perturbation equations for several generic problems. It is demonstrated that the solver is able to correctly predict the acoustic pressure field for sheared mean flow, for acoustic reflection at a solid wall, and for sound waves emanating from a boundary layer.

5.1.1 Monopole in Sheared Mean Flow

Figure 5 shows the results for wave propagation in a sheared mean flow. The example was chosen since mixing layer-type flow configurations with sheared mean flow are typical for noise generation, e.g., for turbulent jets. An S-shaped velocity profile is prescribed for the mean velocity,

$$\bar{u} = \frac{1}{2} \tanh\left(\frac{2y}{\delta_w}\right), \quad (19)$$

where the shear-layer thickness is set to $\delta_w = 50$ and an analytical source term is used to generate an acoustic monopole [6]. The domain was discretized using 200×200 elements with a polynomial degree $N = 3$. Figure 5 shows the result in comparison to the perturbed pressure field obtained in [6] from the linearized Euler equations (LEE). It can be seen that the DG results agree well with the reference solution.




Fig. 5 Monopole in sheared mean flow (left perturbed pressure p' , right p' at $y = 70$ and $t = 180$)




Fig. 6 Reflection of a pressure pulse at a solid wall (*left* perturbed pressure p' , *right* p' at $x = y$ and $t = 30$)

5.1.2 Acoustic Reflection at a Solid Wall

A pressure pulse impinging on a plane wall in the presence of a uniform mean flow was simulated to validate the wall boundary conditions. The wall is located at $y = 0$ and the initial conditions at time $t = 0$ are

$$u' = v' = 0, \quad p' = \exp \left\{ -(\ln 2) \frac{x^2 + (y - 25)^2}{25} \right\}. \quad (20)$$

The mean flow is prescribed parallel to the wall by setting $\bar{u} = 0.5$, $\bar{v} = 0.0$. Both the setup and the analytical values are taken from [14]. The square-shaped computational domain with side length $l = 200$ was discretized using 256 elements in each spatial direction with a polynomial degree of $N = 5$. In Fig. 6, results for the acoustic pressure field of the reflected pulse are shown. They confirm that the CAA solver is able to correctly predict the reflection of acoustic waves from a solid wall.

5.1.3 Monopole in a Boundary Layer

In this case, a plane sound wave is assumed to travel through a small channel and to exit through a small orifice in a plane wall. Due to the small size of the channel, the emanating wave is an approximation for a singular monopole at the wall [3]. The domain is defined by $x \in [-25.6, 25.6]$ in the x -direction and $y \in [0.0, 20.0]$ in the y -direction, and it is discretized by 400,000 elements with a polynomial degree of $N = 3$. The monopole has a size of $\epsilon = 0.1$ and is located at the origin. It is created by enforcing a sinusoidal boundary state by setting

$$u' = 0, \quad v' = p' = \sin(2\pi t). \quad (21)$$


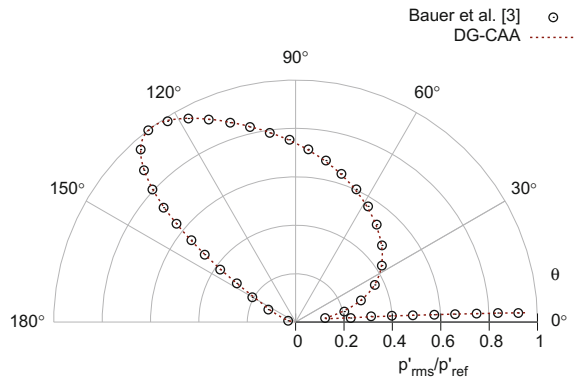


Fig. 7 Contour plot of perturbed pressure p' for monopole in boundary layer

Fig. 8 Directivities for rms pressure along $r = 15$



In addition to the monopole at the wall, a non-zero mean velocity is prescribed, which decreases to zero in the boundary layer region:

$$\bar{u} = \begin{cases} M_x(2y - 2y^2 + y^4), & \text{if } 0 \leq y \leq 1, \\ M_x, & \text{if } y > 1, \end{cases} \quad \bar{v} = 0, \quad (22)$$

where the Mach number is set to $M_x = 0.3$. Figure 7 shows a contour plot of the resulting pressure field. In Fig. 8 the results are compared to those in [3]. The DG-CAA solution is virtually indistinguishable from the reference solution, which demonstrates that the DG-APE method is able to adequately capture the refraction and reflection of sound waves in flow fields with velocity gradients, both in the channel region at $\theta < 10^\circ$ and in the shadow region at $140^\circ < \theta < 180^\circ$.

5.2 Parallel Performance Analysis

To assess the parallel performance of the newly developed aeroacoustics solver, a strong scaling experiment with two setups was performed on HPC systems. In each




Fig. 9 Strong scaling experiments for the CAA solver on a Cray XC 40 (*left*) and a BlueGene/Q (*right*)

setup, the three-dimensional domain is cube-shaped. To obtain meaningful error measures, a manufactured solution approach was used, i.e., an auxiliary source term was added to the system of equations such that the analytical initial conditions, which are based on trigonometric functions, fulfill the system of equations exactly. In the first setup, a grid with 16.8 million cells and a polynomial degree $N = 3$ was used (low-order case). For the second setup, the number of cells was reduced to 2.1 million and the polynomial degree was set to $N = 7$ (high-order case). This yields the same global number of degrees of freedom for both cases (1.1 billion). The setups were chosen to be representative of typical large-scale aeroacoustics simulations under realistic conditions.

Figure 9 shows the strong scaling results for both setups on two state-of-the-art supercomputers, i.e., the Cray XC 40 of the High-Performance Computing Center Stuttgart and the BlueGene/Q of the Forschungszentrum Jülich. On both machines, the simulations were executed with one MPI rank per core and two OpenMP threads per rank. For the Cray system, the low-order case has a parallel efficiency of 79 % on 93,600 cores, which improves to 98 % for the high-order case. Both values are very satisfactory. On the BlueGene/Q, the efficiency for the low-order case on the full machine is 80 %. From these results, it can be concluded that the CAA solver is highly scalable and that it is well-suited for large-scale aeroacoustics simulations. Furthermore, the comparison of the two setups on the Cray XC 40 shows that it is beneficial for the parallel efficiency to use a higher-order approximation in the DG scheme.

To highlight the necessity of developing a new coupling approach for hybrid CFD-CAA simulations, another scaling experiment was conducted. In this case, a CAA simulation of a two-dimensional mixing layer was performed with offline coupling, i.e., the source term information was read from data files [27]. Figure 10 shows the speedup and the absolute wall-clock time for a single-threaded scaling from 32 to 4,096 cores. In the left figure, the speedup is shown once for the overall simulation, with an ultimate efficiency of 61 % at 4,096 cores. When excluding the I/O time, i.e., the time spent reading the source term data from disk, the efficiency improves to 92 %. The reason for this behavior can be understood when looking at the




Fig. 10 Speedup (*left*) and wall-clock time (*right*) for a offline coupling simulation with 31.4 million cells and $N = 1$ on a BlueGene/Q




Fig. 11 Strong scaling results for the I/O write performance of a 63 GiB file on a Cray XC 40 using two Lustre file systems with 96 and 168 object storage targets (OST) respectively (*left*) and a BlueGene/Q using a GPFS file system (*right*)

wall-clock time for computation and I/O separately (see right figure): while the time for computation continuously decreases when using higher core counts, the curve for I/O time flattens out when going from 2,048 to 4,096 cores. This means that the I/O component ceases to scale beyond a certain number of cores, effectively turning the I/O into a bottleneck for the overall simulation.

The degradation of the parallel efficiency for offline-coupled simulations due to I/O performance limits can be further substantiated by examining the I/O bandwidth on current HPC systems. In Fig. 11, the measured maximum write speed for a single 63 GiB file is shown at increasing numbers of cores. The numbers were obtained with the Parallel netCDF library [20] using collective I/O and one MPI rank per core. On both machines, i.e., a Cray XC 40 with a Lustre file system (left figure) and a BlueGene/Q with a GPFS file system (right figure), the I/O bandwidth peaks at a certain number of cores and actually decreases for higher core counts. These results strongly suggest the need for an online coupling approach, where the CFD and the CAA solvers do not have to rely on the file I/O system to exchange data.

6 Conclusions

A direct-hybrid method suitable for large-scale aeroacoustic simulations has been presented. The flow field is predicted using an LES solver based on the finite-volume method. For the CAA solution, a nodal DG method is used to solve the acoustic perturbation equations for the determination of the acoustic field. In the novel approach, both solvers use the same hierarchical Cartesian grid, enabling an efficient data exchange between the two solvers. Appropriate strategies for the spatial and temporal coupling are described.

The CAA method is shown to correctly predict the acoustic pressure field for a monopole in sheared mean flow, acoustic reflection at a solid wall, and a monopole in a boundary layer. In addition, the parallel performance of the new scheme is investigated in several strong scaling experiments. They show that the new DG-CAA solver is capable of efficiently running simulations on hundreds of thousands of cores. Furthermore, while the direct-hybrid method with offline coupling involving disk I/O scales well up to a 128-fold increase in MPI ranks, the I/O operations necessary for reading the source terms from disk are identified as a bottleneck towards extreme scaling. This observation is further corroborated by an analysis of the I/O bandwidth on two current HPC systems, which emphasizes the need for the online coupling approach.

Overall, the proposed direct-hybrid method has shown to be a good candidate for efficient, highly parallel CAA simulations. As a next step, spatial as well as temporal interpolation schemes need to be investigated to lessen the restriction on the resolution requirements in space and time. A dynamic load balancing scheme will be developed to further improve the parallel performance for moving geometries.

Acknowledgements This work has been performed with the support from the JARA-HPC SimLab Fluids & Solids Engineering of the RWTH Aachen University, Germany and the Forschungszentrum Jülich, Germany. The authors gratefully acknowledge the allocation of supercomputing time as well as the technical support by the High-Performance Computing Center Stuttgart of the University of Stuttgart, Germany and by the Jülich Supercomputing Centre of the Forschungszentrum Jülich, Germany.

References

1. Atkins, H.L.: Continued development of the discontinuous Galerkin method for computational aeroacoustic applications. AIAA Paper (97-1581) (1997)
2. Bassi, F., Rebay, S.: A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-stokes equations. *J. Comput. Phys.* **131**(2), 267–279 (1997)
3. Bauer, M., Dierke, J., Ewert, R.: Application of a discontinuous Galerkin method to discretize acoustic perturbation equations. *AIAA J.* **49**(5), 898–908 (2011)
4. Bui, T.Ph., Schröder, W., Meinke, M.: Numerical analysis of the acoustic field of reacting flows via acoustic perturbation equations. *Comput. Fluids* **37**(9), 1157–1169 (2008)

5. Carpenter, M.H., Kennedy, C.: Fourth-order 2N-storage Runge-Kutta schemes. NASA Report TM 109112, NASA Langley Research Center (1994)
6. Ewert, R., Schröder, W.: Acoustic perturbation equations based on flow decomposition via source filtering. *J. Comput. Phys.* **188**, 365–398 (2003)
7. Ewert, R., Schröder, W.: On the simulation of trailing edge noise with a hybrid LES/APE method. *J. Sound Vibr.* **270**(3), 509–524 (2004)
8. Farrell, P., Maddison, J.: Conservative interpolation between volume meshes by local Galerkin projection. *Comput. Meth. Appl. Mech. Eng.* **200**(1–4), 89–100 (2011)
9. Flad, D., Frank, H., Beck, A.D., Munz, C.D.: A Discontinuous Galerkin spectral element method for the direct numerical simulation of aeroacoustics. *AIAA Paper* (2014-2740) (2014)
10. Geiser, G., Marinc, D., Schröder, W.: Comparison of source reconstruction methods for hybrid aeroacoustic predictions. *International Journal of Aeroacoustics* **12**(7–8), 639–662 (2014)
11. Geiser, G., Schlimpert, S., Schröder, W.: Thermoacoustical noise induced by laminar flame annihilation at varying flame thicknesses. In: 18th AIAA/CEAS Aeroacoustics Conference (33rd AIAA Aeroacoustics Conference), 04–06 June 2012, Colorado Springs, CO, AIAA 2012–2093 (2012)
12. Gröschel, E., Schröder, W., Renze, P., Meinke, M., Comte, P.: Noise prediction for a turbulent jet using different hybrid methods. *Comput. Fluids* **37**(4), 414–426 (2008)
13. Günther, C., Meinke, M., Schröder, W.: A flexible level-set approach for tracking multiple interacting interfaces in embedded boundary methods. *Comput. & Fluids* **102**, 182–202 (2014)
14. Hardin, J., Ristorcelli, J.R., Tam, C.K.W. (eds.): *ISCASE/LaRC Workshop on Benchmark Problems in Computational Aeroacoustics (CAA)*, vol. NASA Conference Publication 3000. NASA (1995)
15. Hartmann, D., Meinke, M., Schröder, W.: An adaptive multilevel multigrid formulation for Cartesian hierarchical grid methods. *Comput. Fluids* **37**, 1103–1125 (2008)
16. Hartmann, D., Meinke, M., Schröder, W.: A strictly conservative Cartesian cut-cell method for compressible viscous flows on adaptive grids. *Comput. Meth. Appl. Mech. Eng.* **200**, 1038–1052 (2011)
17. Hindenlang, F., Gassner, G.J., Altmann, C., Beck, A., Staudenmaier, M., Munz, C.D.: Explicit discontinuous Galerkin methods for unsteady problems. *Comput. Fluids* **61**, 86–93 (2012)
18. Koh, S., Schröder, W., Meinke, M.: Turbulence and heat excited noise sources in single and coaxial jets. *J. Sound Vibr.* **329**, 786–803 (2010)
19. Kopriva, D., Woodruff, S., Hussaini, M.: Discontinuous spectral element approximation of Maxwell's equations. In: B. Cockburn, G. Kariadakis, C.W. Shu (eds.) *Proceedings of the International Symposium on Discontinuous Galerkin Methods*. Springer (2000)
20. Li, J., Zingale, M., Liao, W.k., Choudhary, A., Ross, R., Thakur, R., Gropp, W., Latham, R., Siegel, A., Gallagher, B.: Parallel netCDF: a high-performance scientific I/O interface. In: *Proceedings of the 2003 ACM/IEEE Conference on Supercomputing - SC '03*, p. 39. ACM Press, New York, USA (2003)
21. Lintermann, A., Schlimpert, S., Grimmen, J.H., Günther, C., Meinke, M., Schröder, W.: Massively parallel grid generation on HPC systems. *Comput. Meth. Appl. Mech. Eng.* **277**, 131–153 (2014)
22. Liu, J.G., Shu, C.W.: A high-order discontinuous Galerkin method for 2D incompressible flows. *J. Comput. Phys.* **160**(2), 577–596 (2000)
23. Pogorelov, A., Meinke, M., Schröder, W.: Cut-cell method based large-eddy simulation of tip-leakage flow. *Phys. Fluids* **27**(7), 075106 (2015)
24. Reed, W., Hill, T.: *Triangular mesh methods for the neutron transport equation*. Tech. Rep. LA-UR-73-479, Los Alamos Scientific Laboratory (1973)
25. Directorate-General for Research, Innovation European Union: *Flightpath 2050: Europe's Vision for Aviation: Maintaining Global Leadership and Serving Society's Needs*. Office for Official Publications of the European Communities (2011)
26. Sagan, H.: *Space-filling curves*, 1st edn. In: *Universitext*. Springer, New York (1994)
27. Schlottke, M., Cheng, H.J., Lintermann, A., Meinke, M., Schröder, W.: A direct-hybrid method for computational aeroacoustics. In: *AIAA Aviation*, 22–26 June 2015, Dallas, TX, 21st AIAA/CEAS Aeroacoustics Conference, AIAA-2015-3133 (2015)

28. Schneiders, L., Hartmann, D., Meinke, M., Schröder, W.: An accurate moving boundary formulation in cut-cell methods. *J. Comput. Phys.* **235**, 786–809 (2013)
29. Yakovlev, S., Xu, L., Li, F.: Locally divergence-free central discontinuous Galerkin methods for ideal MHD equations. *J. Comput. Sci.* **4**(1–2), 80–91 (2013)