

High Fidelity Scale-Resolving Computational Fluid Dynamics Using the High Order Discontinuous Galerkin Spectral Element Method

Muhammed Atak, Andrea Beck, Thomas Bolemann, David Flad, Hannes Frank, and Claus-Dieter Munz

Abstract In this report we give an overview of our high-order simulations of turbulent flows carried out on the HLRS systems. The simulation framework is built around a highly scalable solver based on the discontinuous Galerkin spectral element method (DGSEM). It has been designed to support large scale simulations on massively parallel architectures and at the same time enabling the use of complex geometries with unstructured, nonconforming meshes. We are thus capable of fully exploiting the performance of HLRS Cray XE6 (Hermit) and XC40 (Hornet) systems not just for academic benchmark problems but also industrial applications. We exemplify the capabilities of our framework at three recent simulations, where we have performed direct numerical and large eddy simulations of turbulent compressible flows. The test cases include a high-speed turbulent boundary layer flow utilizing close to 94,000 physical cores, a DNS of a NACA 0012 airfoil at $Re = 100,000$ and direct aeroacoustic simulations of a close-to-production car mirror at $Re_c = 100,000$.

1 Introduction

The focus of our activities within this project is the development of a numerical framework, which combines the high accuracy of academical codes with the flexibility necessary to tackle complex industrial applications. Our target applications are large eddy simulation (LES) and direct numerical simulation (DNS) and require large amounts of computational resources, which can only be provided by optimally utilizing massively parallel systems. Therefore, each component of the framework has been designed to contribute to a high parallel efficiency. To achieve these

M. Atak (✉) • A. Beck • T. Bolemann • D. Flad • H. Frank • C.-D. Munz
Institute of Aerodynamics and Gas Dynamics, Universität Stuttgart, Pfaffenwaldring 21, 70569 Stuttgart, Germany
e-mail: atak@iag.uni-stuttgart.de; beck@iag.uni-stuttgart.de; bolemann@iag.uni-stuttgart.de; flad@iag.uni-stuttgart.de; frank@iag.uni-stuttgart.de; munz@iag.uni-stuttgart.de

goals we have chosen a variant of the discontinuous Galerkin (DG) methods called the discontinuous Galerkin spectral element method (DGSEM) as the frameworks foundation. As in Finite Elements, DG schemes use a polynomial approximation of the solution within a grid cell, but do not require continuity at the cell boundaries. Instead they rely on Finite Volume like flux functions for the coupling between grid cells, making them more robust in the presence of strong solution gradients than classical Finite Element schemes. High-order DG methods in general are known for their superior scale resolving capabilities, making them well suited for LES and DNS [6] as these require the resolution of a broad range of scales. Furthermore, previous investigations have shown that DG methods can be implemented efficiently on parallel architectures and scale up to more than 100,000 physical cores. This can be attributed to the local nature of DG schemes, as only surface data needs to be communicated.

The DGSEM-based CFD solver FLEXI, which is used for the simulations, is the most recent development of IAG's DG-based codes. Stemming from the Spectral Element Method, DGSEM is considered as one of the most efficient DG formulations. It uses a tensor-product nodal basis to represent the solution and collocates integration and interpolation points. The operator can thus be implemented using a favorable dimension-by-dimension structure and supports very high polynomial degrees. The methods' restriction to hexahedral (but unstructured) meshes may introduce difficulties for complex geometries, but are easily circumvented by element splitting strategies and the use of non-conforming meshes.

We demonstrate the accuracy, versatility and performance of our framework at the example of three test cases, an airfoil, a car mirror and a flat plate boundary layer. The first case, addresses a highly accurate DNS of a NACA 0012 airfoil and serves as preparation step for the car mirror simulations. Here, a very high polynomial degree of $N = 11$ has been employed. In the second test case, a LES for direct aeroacoustics of a close-to-production car side-view mirror has been conducted, showing that our framework is well applicable in an industrial setting with complex geometries. The final case, a DNS of a supersonic turbulent boundary layer, proves that FLEXI is able to scale seamlessly from a few thousand cores up to the whole HLRS Cray XC40 Hornet system and that our postprocessing toolchain can efficiently handle this amount of data. This simulation involved 1.5 billion degrees of freedom (DOF) and represents the biggest DG simulation described in literature. We note that all simulations have been performed in an off-design setting concerning the optimal number of degrees of freedom per core.

2 An Overview of the FLEXI Framework

The CFD framework of our research group has been developed for the simulation of unsteady 3D turbulent flows, characterized by a broad range of turbulent scales. Earlier simulations of turbulent flows at low Reynolds numbers using the FLEXI framework showed very promising results [4]. Due to the framework's maturity after

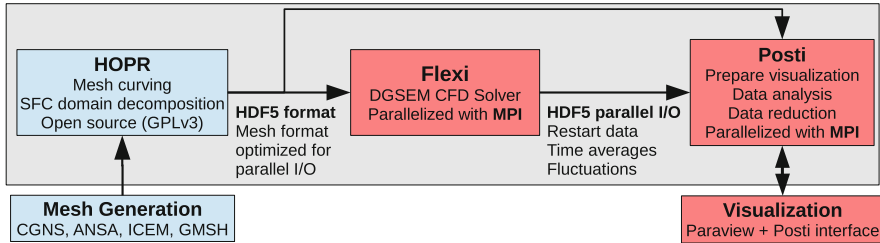


Fig. 1 Framework and workflow for a DGSEM simulation. *Blue*: serial, *red*: parallel

ongoing development and the encouraging results, we raised the bar and applied it to more complex flow phenomena in terms of higher Reynolds numbers and geometric complexity as well as even larger scale simulations.

The single components of the framework and the workflow are depicted in Fig. 1. The frameworks' three main components are the mesh preprocessor HOPR, generating high-order meshes, which are required for DG methods, the massively parallel DGSEM-based CFD solver FLEXI and the parallel postprocessing toolkit POSTI.

2.1 Pre- and Postprocessing

By design, high-order DG methods have a very high resolution within each cell, especially when operating at high polynomial degrees. It is therefore possible to use much coarser meshes compared to lower-order schemes, which is however in conflict with a sufficient geometry representation. Bassi and Rebay [3] have shown that in the case of curved boundaries a high-order DG discretization yields low-order accurate results when using straight-sided meshes. DG methods thus require a high-order boundary approximation to retain their accuracy in the presence of curved boundaries.

We have developed the standalone high-order preprocessor HOPR¹ to generate these high-order meshes by using input data from standard, linear mesh generators. HOPR has been recently made open source under the GPLv3 license. Several strategies for the element curving are implemented, e.g. using additional data obtained by supersampling the curved boundaries or a wall-normal approach, see [7] for a complete overview. Though the current DG solver FLEXI uses only hexahedra, our mesh preprocessor is also capable of generating hybrid curved meshes. In addition, HOPR also provides a domain decomposition prior to the simulation by sorting all elements along a unique space-filling curve. The elements are then in a non-overlapping fashion stored in a contiguous array, together with

¹Available as open source under GPLv3 license at <http://www.hopr-project.org>.

connectivity information in a parallel readable file using the HDF5 format. The domain decomposition being already contained inside the mesh has some major advantages:

- The domain decomposition is directly available to all tools without requiring any additional libraries.
- The element list can be simply split into equally sized parts and each process reads its own set of contiguous non-overlapping data from the file. This massively reduces the file system load, which is especially beneficial for a large number of processes performing simultaneous I/O.
- As the solution is also stored according to the sorting in a single file, the simulations can be simply run, restarted and postprocessed on any number of processes without any need to convert or gather files.

We employ a Hilbert curve for the space-filling curve element sorting due to its superior clustering properties compared to other curves [9]. The main advantage of this approach embraces that proximate three-dimensional points are also proximate on the one-dimensional curve, which is an important property for domain decomposition. While specialized tools like METIS may provide a slightly better domain decomposition, we benefit from the fact, that for high-order DG methods a process has typically far less than 100, sometimes even as few as 1–2 single elements.

The domain decomposition obtained from the space filling curve approach of HOPR is depicted in Fig. 2 at the example of an unstructured hexahedral mesh. It shows the distribution of the grid cells on 128 and 4096 processes, each subdomain is shrunk towards its barycenter for visual purposes. The communication graph depicts for each domain its communication partners.

Postprocessing the simulation data for large datasets has become a major task in the last years. Like in mesh generation, DG, as a relatively young numerical technique, lacks the support of specialized visualization tools. DG solution data cannot be processed directly using standard tools like Tecplot or Paraview as those require linear solutions inside a cell, the DG solution, however, is a piecewise polynomial. We thus use a supersampling strategy and decompose a DG cell into multiple linear cells, which can be viewed using standard tools. However, the visualization is only a small part of the actual postprocessing procedure. Spatial averaging and computing correlations from time-averaged flow data as well as evaluating time accurate data logged each time step are only some examples for the very distinct tasks. Therefore our postprocessing toolkit POSTI has been designed to be easily adaptable to changing requirements. We employ a stripped-down modularized version of our solver as a basis, providing building blocks for I/O support, readers for geometry or solution data and interpolation and integration features. On top of this, we build separate tools, each specialized on single tasks. The resource consuming parts of our postprocessing toolkit POSTI feature an efficient parallelization. A recent and significant improvement of the framework is the development of a POSTI interface for Paraview. This enables the user to directly visualize FLEXI state files in Paraview without intermediate steps and provides Paraview access to POSTI features.

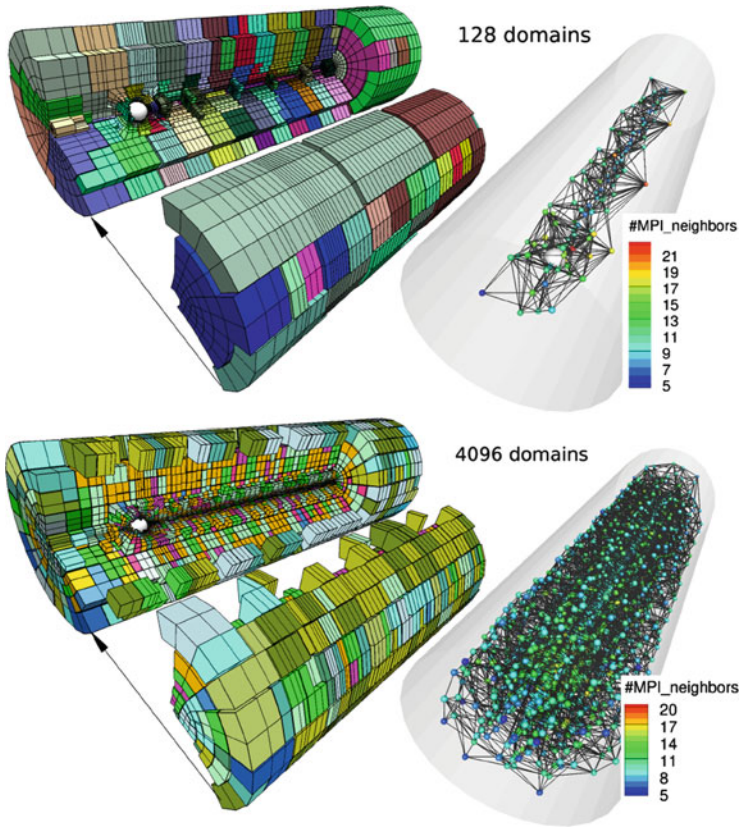


Fig. 2 Domain decomposition and communication graph for the unstructured mesh with 128 (*top*) and 4096 (*bottom*) domains

2.2 Parallel Performance

In this section, we present a parallel performance analysis of the DGSEM solver FLEXI. The strong scaling tests have been conducted on the new HLRS Cray XC40 Hornet cluster for a range of 48–12,288 physical cores, each test running for approximately 10 min. Figure 3 depicts two setups we have investigated: in the first setup, we employ a constant polynomial degree of $N = 5$ (216 DOFs/element) and use three different meshes from 768 to 12,288 elements. In the second setup we use a fixed mesh with 12,228 elements and benchmark 6 polynomial degrees from $N = 3$ –9 (64–1000 DOFs/element). In both cases we do not apply multithreading and start at less than 100 processes. We constantly double the number of processes until we reach the limit one single element per process. For all cases we can observe superlinear scaling over a wide range of processes and scaling only degrades at one

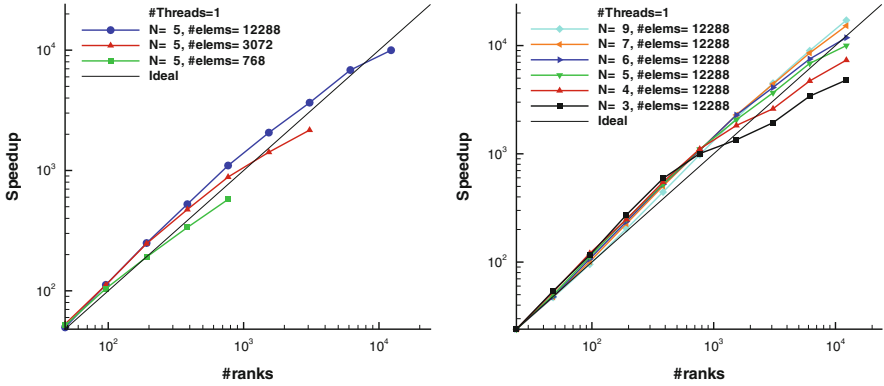


Fig. 3 Comparison of the strong scaling of the solver FLEXI without multithreading for varying number of elements and constant polynomial degree (*left*) and constant number of elements and varying polynomial degrees (*right*)

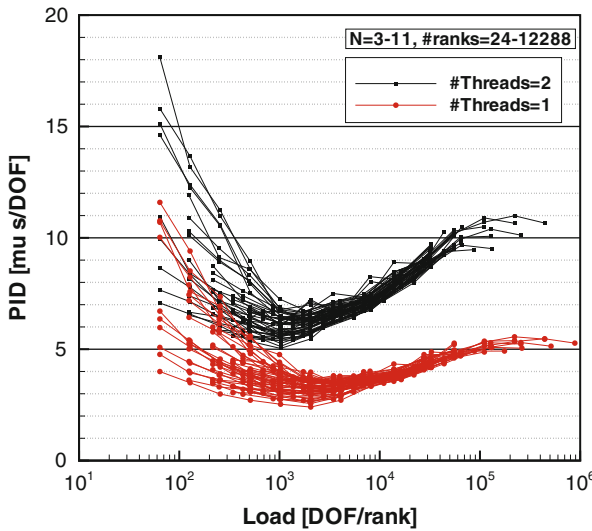


Fig. 4 Performance index (PID) with/without multithreading for polynomial degrees $N = 3-11$ and 24–12,228 processes

element per core. Note, that for the higher polynomial degrees $N = 7$ and $N = 9$, the scaling does not degrade even in the single element case.

In a third setup depicted in Fig.4 we investigate solvers’ performance index (PID), which is a convenient measure to compare different simulation setups: it represents the CPU-time required to update one degree of freedom within one time step, and is computed from the overall CPU time, the number of timesteps and

degrees of freedom of the simulation,

$$\text{PID} = \frac{\text{Wall-clock-time} \# \text{cores}}{\# \text{DOF} \# \text{Timesteps}}.$$

The PID boils down the contradicting effects of scaling, compiler optimization and CPU cache on the performance to a single efficiency measure, showing three regions. In the leftmost part in Fig. 4 is the latency-dominated region, where we observe a very high PID. Here, communication latency hiding by doing inner work has no effect due to the low load per core ratio. The rightmost region, however, which is characterized by a high PID and high loads, is dominated by the memory-bandwidth of the nodes. In the central “sweet-spot” region the load is small enough to fit into the CPU-cache while still ensuring sufficient latency hiding. Furthermore, we note that in our case multithreading had no beneficial effect on the solver performance.

3 DNS of a NACA 0012 Airfoil at $Re=100,000$

In this project, we have computed a direct numerical simulation of a NACA 0012 airfoil at $Re = 100,000$ and an angle of attack of 2° . This challenging flow case serves two purposes:

- The acoustic feedback mechanism described in Sect. 4 needs to be investigated further. To better understand the physical mechanisms, in particular in combination with turbulence, we have chosen a problem setup that generates a laminar vortex shedding and thus tonal noise. A comparison with literature [1, 8, 12] strongly suggest the existence of the feedback mechanism described for the car mirror case. Data analysis is currently underway to determine whether the feedback mechanism can also be identified for this case.
- The development of numerical schemes and modeling approaches for turbulent flows requires fully resolved DNS simulations as reference data. In particular for higher Reynolds numbers, complete spatial and temporal data sets are rarely available for comparison. This DNS computation provides a challenging test case for the LES models, as it incorporates fully laminar as well as turbulent flow regions and their interaction in the wake.

The NACA 0012 airfoil is located at the center of a C-shaped domain with radius $R = 100c$, where c denotes the airfoil chord. The chord is at an angle of attack of 2° to the free-stream flow. The spanwise extension is $0.1c$. A rectangular trip region of height $0.003c$ at position $x = 0.05c$ is located near the leading edge to force transition on the upper half. Periodic boundary conditions are applied on the spanwise faces, isothermal wall conditions on the geometry and free-stream condition on the outer boundaries. The Reynolds number was set to $Re = 100,000$, the Mach number to $M = 0.3$. The domain is discretized with 71,500 hexahedral

elements with a local tensor-product polynomial ansatz of degree 11, resulting in 123.5 million DOFs. The non-linear fluxes are represented by a tensor-product ansatz of degree 13, leading to 196.2 million flux integration points. Figure 5 shows a spanwise slice of the computational grid.

The simulation was started from a free-stream initial state, and required about 25 convection times $T^* = c/U_\infty$ to achieve stable vortex shedding. Figure 6 depicts the temporal development of the lift and drag coefficients after the initial transient. All computations were run on the HLRS Cray XC40 Hornet, using 14,300 processes.

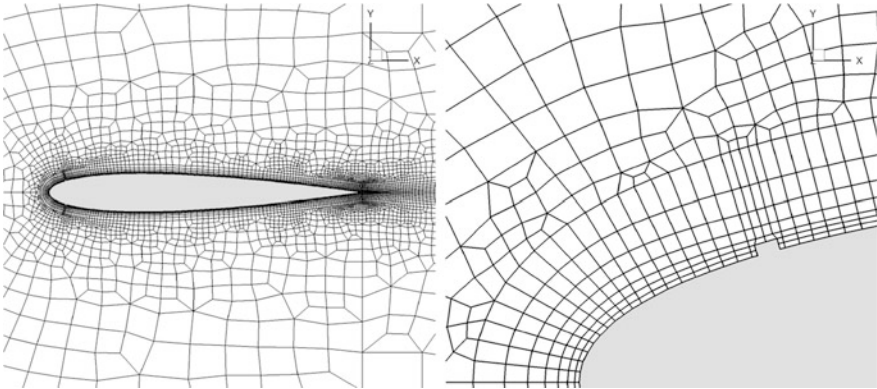


Fig. 5 Left: Computational grid for DNS of NACA 0012. Right: Zoom on upper trip region

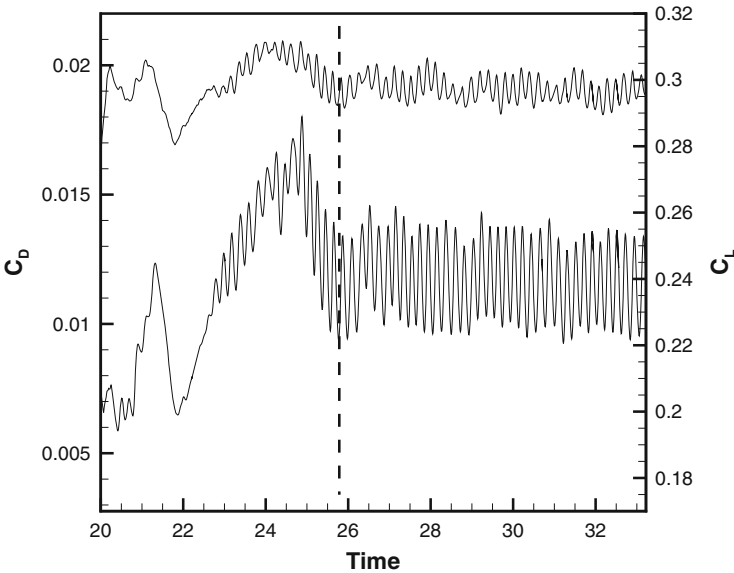


Fig. 6 Evolution of lift and drag coefficients over convective time units T^* . Dashed vertical line denotes onset of stable shedding

Table 1 summarizes the computational costs of this simulation. We have chosen the number of processes to achieve a balance between an optimal PID and queue waiting time, therefore, the core load is not optimal as for the case presented in Table 2.

Figures 7 and 8 show some preliminary results from the DNS computation. As expected, the trip generates rapid transition and results in a fully turbulent boundary layer on the upper airfoil surface. Due to the negative pressure gradient, the lower boundary layer remains laminar, even across the trip, and shows characteristic two-dimensional rollers near the trailing edge, which suggests the occurrence of tonal noise.

Table 1 Computational details for the DNS of NACA 0012 airfoil case

N	n_{DOF}	Procs	$PID[\mu s]$	Runtime/ T^*
11	$1.235 \cdot 10^8$	14,300	8.2	23 min

Table 2 Computational setup for the NACA0012 airfoil case

N	n_{DOF}	Procs	$PID[\mu s]$	Runtime/ T^*
7	$43.6 \cdot 10^6$	3,288	6.0	266.7 min

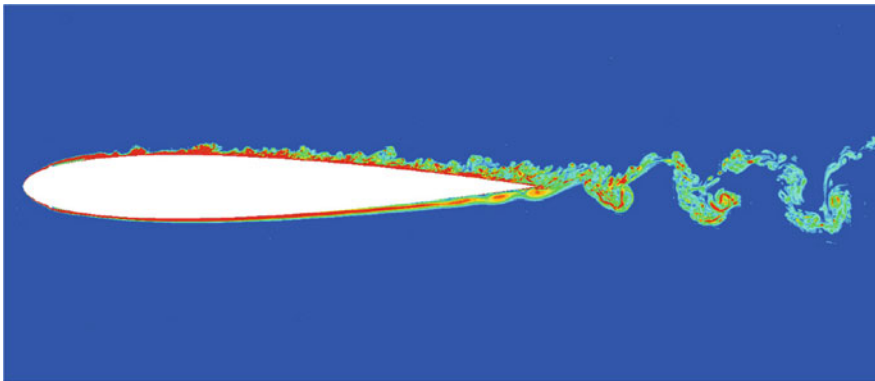


Fig. 7 Instantaneous vorticity magnitude at $t = 29$ and $z = 0.05c$

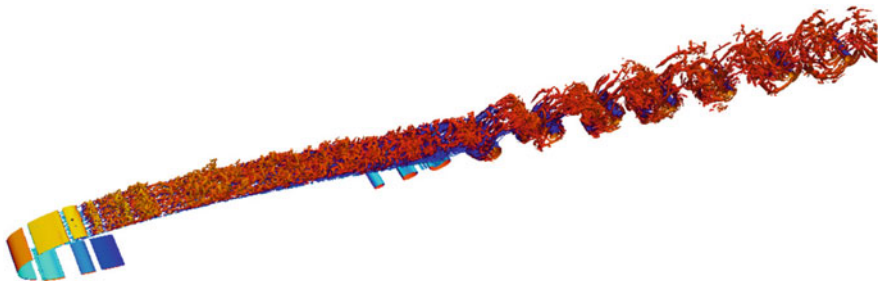


Fig. 8 Visualization of the vortical structures by the Q-criterion ($Q=100$). Tripped boundary layer on *top* surface becomes fully turbulent, while two-dimensional structures on lower side indicate laminar flow

As a next step, the DNS computation will be continued to gather sufficient statistical data for the analysis of the turbulent and acoustic quantities. It will then serve as a benchmark for the validation and assessment of LES modeling approaches.

4 Direct Aeroacoustic Simulation for the Analysis of Acoustic Feedback Mechanisms at a Side-View Mirror

For the acoustic optimization in automotive engineering the reduction of acoustic sources leading to narrow band noise is essential. As broad band noise sources such as engine noise or noise on road were significantly reduced, aeroacoustic noise sources leading to tonal noise are perceived as especially annoying by occupants. Particularly tonal noise at the side-view mirror is often found during development processes and needs to be removed in costly wind tunnel testing. To this end, the underlying mechanisms are not well understood. This work contributes to the understanding of the mechanism by comparing experimental and numerical data, both obtained at IAG. For the numerical analysis the fully compressible Navier-Stokes equations are solved time accurately using our in-house discontinuous Galerkin spectral element solver FLEXI. Using high order schemes allows for the direct simulation of acoustic within a hydrodynamic simulation, due to their excellent dissipation and dispersion properties. The simulation is scale-resolving in the sense of a large eddy simulation (LES). In comparison to state-of-the-art hybrid acoustic methods, this allows for the simulation of acoustic feedback mechanisms.

The idea of acoustic feedback leading to tonal noise is well established for airfoils at medium Reynolds number flows ($\mathcal{O}(10^5) - \mathcal{O}(10^6)$) and was investigated experimentally [1, 12] and numerically [8]. When there is a laminar flow almost up to the trailing edge, laminar separation is likely to occur. Within the separated shear layer instabilities are excited, rolling up as vortices downstream. The interaction of vortices with the trailing edge leads to sound emission and the upstream traveling acoustic wave may excite boundary layer instabilities of the same frequency. This mechanism is only self-maintaining for such frequencies with negligible overall phase difference. This criterion leads to the typical ladder-like structure of the acoustic spectrum when plotted over inflow velocity. The acoustic spectra of the side view mirror under consideration show a similar behavior. In addition, the absence of tonal peaks in the spectrum, when turbulators are applied to the surface further hint at the presence of acoustic feedback.

Prior to the study of the side view mirror, the simulation method was validated by an airfoil simulation described in [8]. The flow around a NACA 0012 airfoil is simulated fully scale-resolving (DNS) in two dimensions. The Reynolds number with respect to the chord length is $Re_c = 100,000$, the Mach number $M = 0.4$ and the angle of attack is 0° . Vortex shedding downstream of the separation lead to several discrete spectral components. The perturbation ansatz also used

in [8], examining only a single isolated wave package of vanishing amplitude, definitely demonstrates the existence of a feedback mechanism in the simulation. Figure 9 shows the non-dimensional power spectral density of DNS and perturbation simulation evaluated two chords above the trailing edge. In both cases several narrow band peaks are observed for which the frequencies collapse.

Analogous to the experiments, described in detail in [13], we consider an isolated side view mirror mounted on the floor. The setup and the computational mesh on the surface can be seen in Fig. 10. The inflow velocity is $U_\infty = 100$ km/h corresponding to a Reynolds number of $Re_S \approx 184,000$, with respect to the mirror inner surface length $S = 0.1$ m, and Mach number $M = 0.082$. The computational domain has 17 million degrees of freedom, using a polynomial degree $N = 7$.

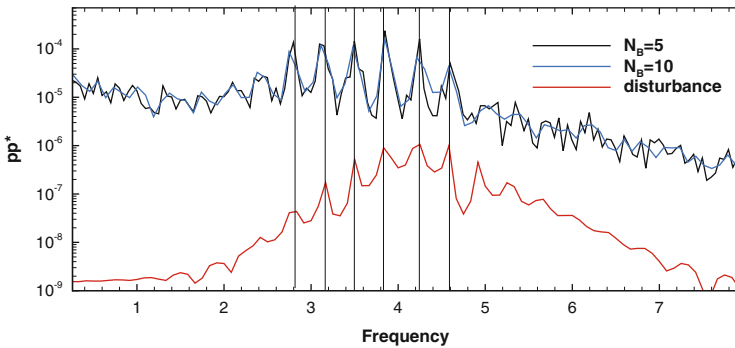


Fig. 9 Power spectral density of emitted acoustic field for DNS (*upper curves*) and perturbation calculation (*lower curve*)

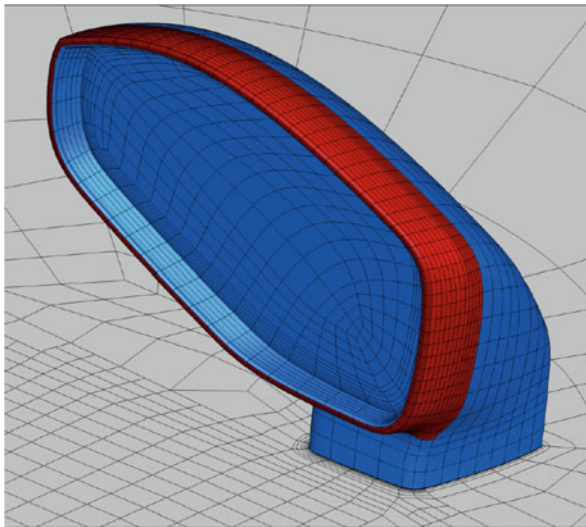


Fig. 10 Isolated mirror on windtunnel floor. Surface mesh with refinement near the trailing edge

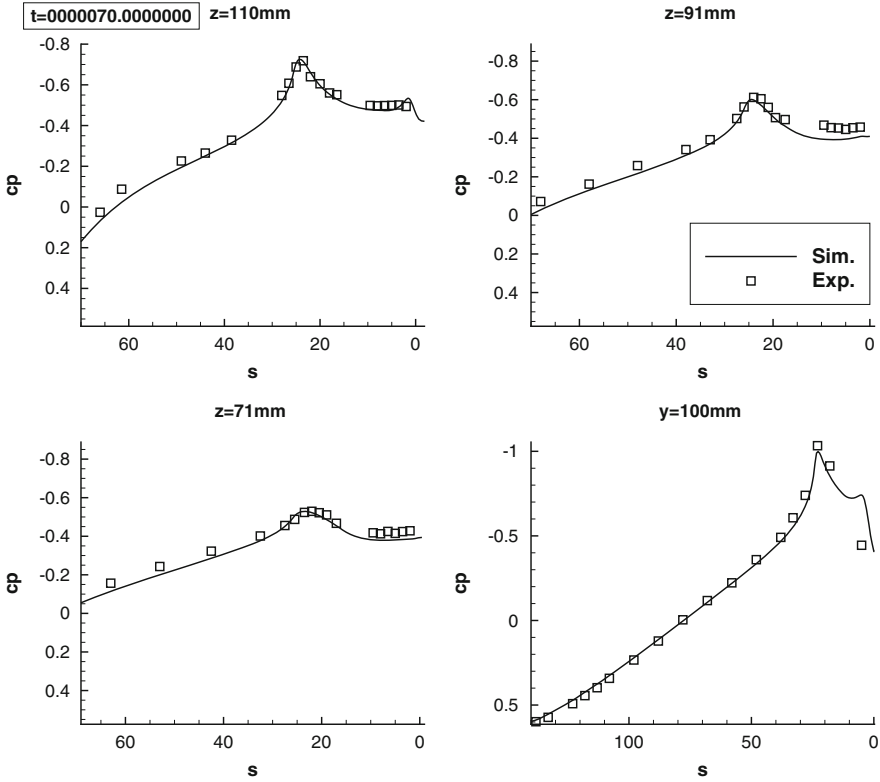


Fig. 11 Pressure distribution on the inner ($z = 110$, $z = 91$, $z = 71$) and top ($y = 100$ mm) mirror surface

The computation was carried out on HLRS Cray XC40 using 3288 processes. The simulation was run for 70 convective time units ($T^* = S/U_\infty$), where the last 40 time units were used to analyze the results.

Figure 11 shows the surface pressure distribution for three positions on the inner surface and one on the top surface of the side view mirror, where s is the wall-tangential coordinate up to the trailing edge. Excellent agreement of simulation and experimental data is observed at all positions. The pressure distribution shows that the flow is strongly accelerated up to an edge upstream of the trailing edge. At the edge the pressure peaks and the flow downstream is unable to follow the strong pressure rise. Consequently the boundary layer separates, characterized by the flat pressure distribution observed in Fig. 11 after the peak. The boundary layer separation is visualized in Fig. 12 by means of the time averaged velocity field obtained from Particle Image Velocimetry (PIV) and LES results. Again, there is an excellent agreement between experiments and simulation for the hydrodynamic quantities.

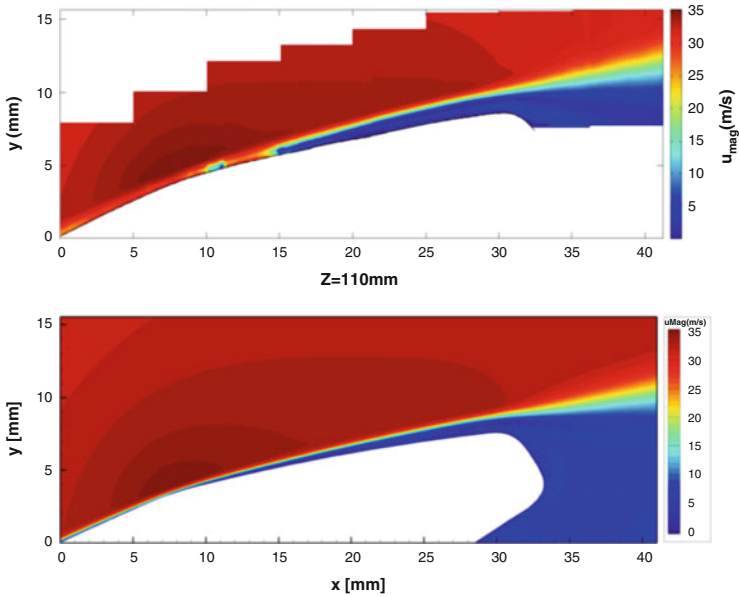


Fig. 12 Time averaged velocity magnitude at $z = 110$. *Top:* PIV, *bottom:* simulation

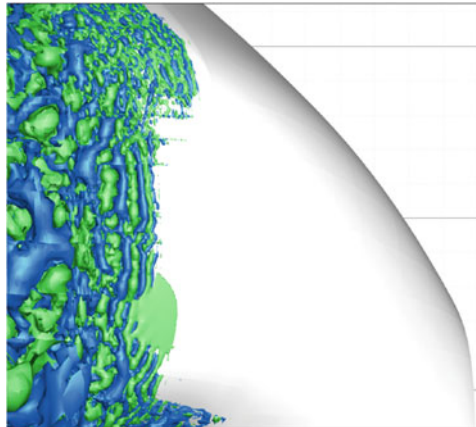
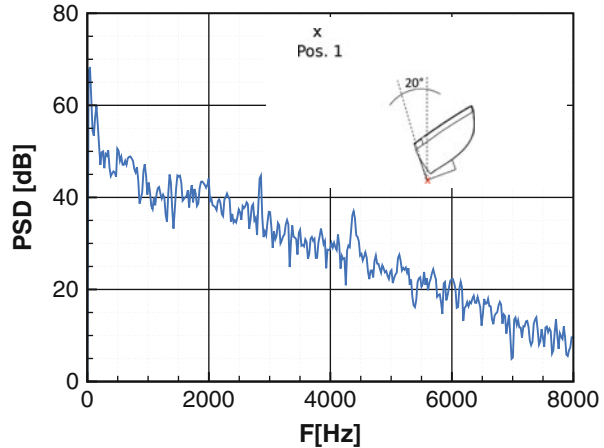


Fig. 13 Visualization of vortex structures by means of isosurfaces of instantaneous y -velocity fluctuations with $v' = \pm 0.01U_\infty$. In proximity to the side surface, the y -velocity is aligned with the wall-normal direction

From the velocity field, turbulent re-attachment of the boundary layer can be ruled out, therefore all prerequisites for the feedback loop are given. The evaluation of the corresponding RMS-fluctuations (not plotted) shows the growth of perturbations along the separated shear layer. A qualitative impression of the acoustic source region at the mirror inner side is given in Fig. 13. Coherent vortices are identified by

Fig. 14 Akustisches Spektrum bei Position 1, $z = 300$ mm



velocity fluctuation isosurfaces. The interaction of those vortices (or rollers) with the trailing edge at an almost constant frequency results in distinct tonal noise. Similar structures are also found at the mirror top surface. The power spectral density (PSD) of the acoustic pressure in Fig. 14 shows two dominant tonal peaks. The frequencies are at 2860 and 4370 Hz. A comparison to the hydrodynamic pressure spectra reveals that the source of the low frequency is located at the mirror inner surface, the higher frequency source at the mirror top surface. This is in good agreement with the experiments. Despite the excellent agreement of hydrodynamic quantities and the location of the sources, the measured frequencies deviate significantly from those predicted by the simulation. The measured frequency at the top surface is 5100 Hz and at the inner surface 3600 Hz. Up to now, there is no conclusive identification for the deviation.

In summary, we showed that the tonal noise generation mechanism observed at an automotive side view mirror is closely related to acoustic feedback loops known from airfoils or wing sections. The noise generation was successfully reproduced within a direct acoustic simulation, resolving all relevant scales in time and space. Such simulations are only possible and meaningful when using high order, highly efficient methods, such as our DGSEM solver FLEXI. The calculation also shows the advantage compared to state-of-the-art hybrid acoustic-hydrodynamic methods, which cannot account for acoustic feedback in general. Table 2 summarizes the computational details of the side view mirror simulation, run on HLRS Cray XC40 Hornet.

5 The Multicore Challenge: DNS of a Spatially-Developing Supersonic Turbulent Boundary Layer up to High Reynolds Numbers

The payload ratio and the economic effectiveness of space transportation systems are important characteristic factors for today's aerospace projects. These factors can even decide whether a project will be realized or discarded. In this context, air-breathing supersonic and hypersonic aircrafts may become a key technology for the future of space flight as they could significantly reduce the payload costs. Particularly scramjets—a ramjet engine where the combustion occurs at supersonic velocities—are considered as a favorable alternative for tomorrow's space missions. In contrast to conventional transportation technologies, which need to carry tons of oxygen supplies, the scramjet inhales the atmospheric air to gain the oxygen required for combustion. Thus, the aircraft becomes lighter, faster and eventually more economic since the payload can be drastically increased. The incoming air will not be compressed by moving parts like compressors, but through a series of shock waves generated by the specific shape of the intake and the high flight velocity, hence, the intake plays a decisive role for air-breathing vehicles. Furthermore, as the supply of compressed air is of fundamental importance for the subsequent efficient combustion of the fuel-air mixture to produce thrust, the intake also determines the operability limits of the whole system. The intake flow itself is characterized by laminar and turbulent boundary layers and their interaction with shock waves, yielding a three-dimensional unsteady complex flow pattern. Beyond that, due to interactions with shocks the boundary layer may experience intense heat loads leading to serious aircraft damages [2]. Hence, the investigation of the turbulent boundary layers are crucial not only for new cooling concepts, but also to ensure the structural integrity of such supersonic/hypersonic air-breathing propulsion systems.

As mentioned before, the flow within the intake (particularly at the upper part) is dominated by laminar and turbulent boundary layers emerging along the walls. Thereby, the initially laminar boundary layer becomes turbulent further downstream of the intake. The phenomenon describing the changeover from laminar state to turbulent is referred to as transition and depicts one of the main flow features encountered in the intake. During the operation state of the scramjet, however, the occurring shock waves traverse through the spatially-developing boundary layer until a quasi-equilibrium state between the shock strength and the boundary layer will be reached. Thus, in order to generate a reliable database for further complex studies, e.g. shock wave/boundary layer interactions (SWBLI) at different impingement locations, and also to demonstrate the high potential of DG schemes for compressible turbulent wall-bounded flows, we will show results of a DNS of a compressible spatially-developing flat-plate turbulent boundary layer up to a momentum thickness based Reynolds number of $Re_\theta = 3878$. We note that, the simulation of a turbulent boundary layer along the flat-plate represents the base study for canonical SWBLIs as it can be transferred into a SWBLI study by imposing shock waves via appropriate boundary conditions.

The free-stream Mach number, temperature and pressure of the supersonic boundary layer are given by $M_\infty = 2.67$, $T_\infty = 568$ K and $p_\infty = 14,890$ Pa, respectively. Air was treated as a non-reacting, calorically perfect gas with constant Prandtl number $Pr = 0.71$ and with a constant specific heat ratio $\kappa = c_p/c_v = 1.4$. Sutherland's law was used to take the temperature dependency of the dynamic viscosity μ into account. The Reynolds number at the leading edge of the plate was set to $Re = 1.156 \cdot 10^5$ and isothermal no-slip conditions were applied at the wall. The wall temperature was equal to the adiabatic wall temperature $T_w = T_{ad} = 1242$ K. Furthermore, sponge zones were used at the outflow regions to suppress any reflections from the outflow boundaries. To trigger the laminar-turbulent transition, periodic disturbances—given by amplitude and phase distributions along the inflow boundary—were added to the initial Blasius solution. In this work, we superimposed 5 discrete disturbances (determined by the eigenfunctions from linear stability theory) with a disturbance amplitude of $A = 0.02$ which yield a maximum RMS velocity of 0.43 % of the free-stream velocity. As for compressible flows the most unstable modes are obliquely traveling disturbance waves, we introduced two single oblique waves with a fundamental spanwise wave number of $\gamma_0 = 21$ to rapidly reach a turbulent state. The DNS was performed with the DGSEM code FLEXI using a polynomial degree of $N = 5$. The mesh consisted of $1500 \times 100 \times 45$ elements in x-, y- and z-direction resulting in 1.458 billion DOFs. To the best of the authors' knowledge, this simulation is the biggest computation within the DG community.

Due to the excellent scaling attributes of the solver, we were able to cope with the demanding costs by exploiting the whole computational resources on the HLRS Cray XC40 Hornet in an efficient way and run the DNS with up to 93,840 physical cores. An overview of the performance index (PID) for different numbers of processors used for this study is given in Table 3. The overview particularly reveals that the PID stays constant up to the limit of available physical cores, which proves that the solver achieves an ideal speedup without any performance losses even at running simulations on the whole machine.

The non-dimensional simulation time was $T = 4.25T^*$ whereas $T^* = 2L/U_\infty$ labels the characteristic flowthrough time with plate length L and free-stream velocity U_∞ . Figure 15 shows the instantaneous distribution of the vortices along the plate at the final simulation time visualized by the λ_2 -criterion and colored by the streamwise velocity component. Here, we can clearly see how the oblique disturbance waves introduced in the inflow generate hairpin vortices which in turn breakdown into turbulence further downstream.

The streamwise development of the time- and spanwise-averaged skin friction coefficient C_f is shown in Fig. 16. Whilst the leading edge of the plate always maintains its laminar state, the flow experiences transition where a sudden increase of C_f is observable around $x_{tr} \approx 3.5$ ($Re_{x_{tr}} = 3.5 \cdot 10^5$). After the overshoot in skin-friction, which is typical for transition, the C_f -profile in the turbulent regime, however, agrees very well with the skin-friction correlation according to the reference temperature approach [14].

Table 3 Performance index for varying numbers of processes on the HLRS Cray XC40 Hornet

Procs	12,000	24,000	93,840
Grid	$1500 \times 100 \times 45$	$1500 \times 100 \times 45$	$1500 \times 100 \times 45$
n_{DOF}	$1.498 \cdot 10^9$	$1.498 \cdot 10^9$	$1.498 \cdot 10^9$
$n_{DOF}/Procs$	121,500	60,750	15,537
PID [μs]	6	6	6

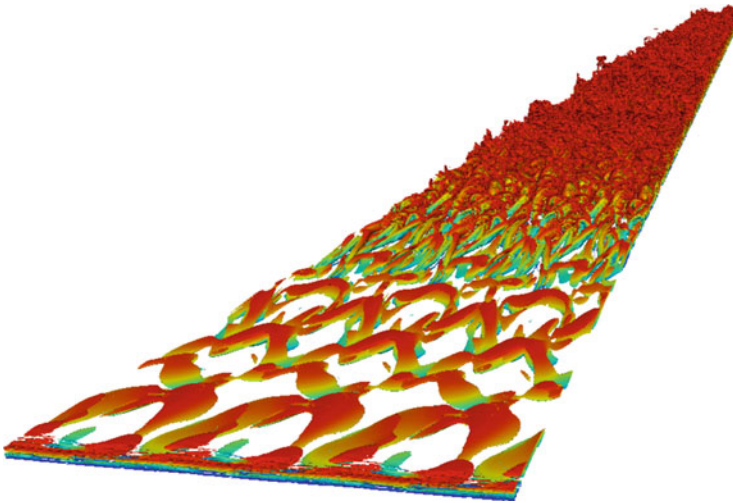


Fig. 15 λ_2 -visualization of the turbulent structures along the flat-plate colored by the streamwise velocity component u

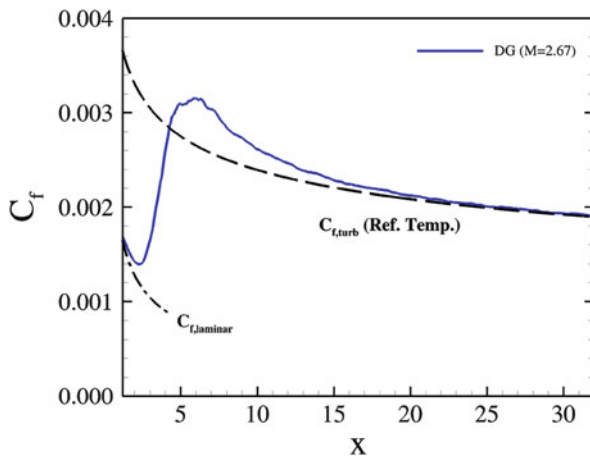


Fig. 16 Development of the skin friction coefficient C_f along the plate length

The Van Driest transformation of the streamwise velocity u accounts for compressibility effects and enables the comparison with incompressible scaling laws [14]. Figure 17 compiles the transformed velocity profiles u_{vD}^+ at different downstream positions along the plate. Here, we can clearly observe that the results perfectly match both the theoretical incompressible law of the wall ($u^+ = y^+$ and $u^+ = 2.44 \cdot \ln(y^+) + 5.2$) and well-established DNS studies of spatially-developing turbulent boundary layers by Bernardini and Pirozzoli [5] and Schlatter and Örlü [11]. In Fig. 17 we also notice that the higher the momentum thickness based Reynolds number, the longer the velocity profile follows the incompressible log-law.

Furthermore, in Fig. 18 we compared the Reynolds-stresses, again, at different streamwise locations with previous DNS data and also with experimental PIV measurements of Piponniau et al. [10]. The Reynolds-stresses were rescaled according to Markovin’s density scaling [14] to account for the varying density. Figure 18 exhibits that the results are in an excellent agreement with other DNS studies as

Fig. 17 Van Driest transformed velocity profiles u_{vD}^+ at different streamwise locations compared with DNS data by Bernardini and Pirozzoli (BP) [5] and Schlatter and Örlü (SÖ) [11]

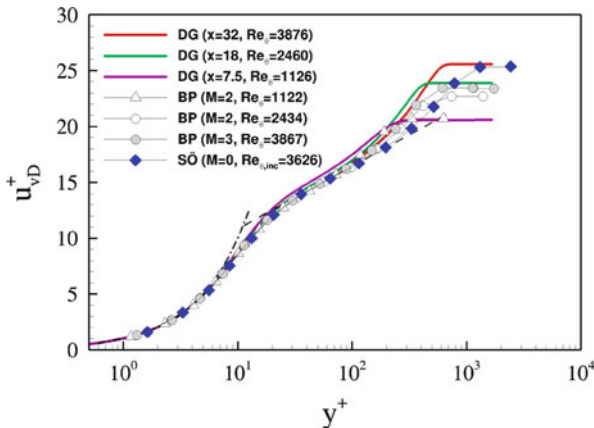


Fig. 18 Density weighted Reynolds-stresses at different streamwise locations compared with DNS data by Bernardini and Pirozzoli (BP) [5], Schlatter and Örlü (SÖ) [11] and experimental data by Piponniau et al. [10]

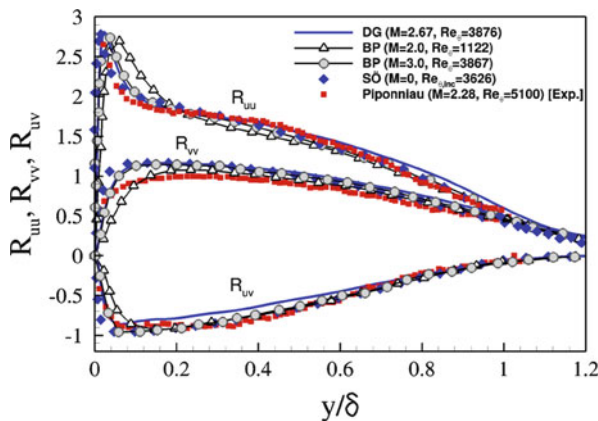


Table 4 Computational setup for the compressible supersonic turbulent boundary layer case

N	n_{DOF}	Processors	$PID[\mu s]$	Runtime/ T^*
5	$1.458 \cdot 10^9$	93,840	6	23.41 <i>h</i>

well as with experimental data. A summary of the computational details is given in Table 4.

With this study, we laid the foundation of subsequent SWBLI studies, in which we are interested in assessing the impact of different shock impingement locations on the flow field and thus to gain a deeper insight into the complex physics involved. In this context, we will apply the FV subcell shock capturing which is already incorporated in our DG framework and enables the robust approximation of shock waves.

6 Summary and Outlook

In this report, we summarize the ongoing development in numerical schemes, parallelization efforts and toolchain improvements for our HPC simulation framework and provide details for large scale application cases.

We demonstrate that the excellent scaling results transfer from the HLRS Cray XE6 Hermit to the new XC40 Hornet, and our postprocessing toolchain is able to handle massive amounts of data. The efficiency of our simulation framework allows us to conduct large scale simulations with 10^7 to 10^9 spatial degrees of freedom, using up to 93,840 processes on the XC40 and 14,000 on the XE6. The presented test cases show that our framework is versatile and mature enough to conduct industrially relevant simulations of complex engineering flows, as well as very large scale computations of canonical test cases for basic research.

In the future, the capabilities of the framework will be extended to retain efficiency beyond 10^5 processes. We will apply our framework to direct numerical simulations of previously unfeasible and expensive computations of multiscale problems, both in basic research as well as in applied engineering.

Acknowledgements The research presented in this paper was supported in parts by the Deutsche Forschungsgemeinschaft (DFG), the Boysen Stiftung and the Audi AG. We truly appreciate the ongoing kind support by HLRS and Cray in Stuttgart.

References

1. Arbey, H., Bataille, J.: Noise generated by airfoil profiles placed in a uniform laminar flow. *J. Fluid Mech.* **134**, 33–47 (1983)

2. Babinsky, H., Harvey, J.K.: *Shock Wave-Boundary-Layer Interactions*. Cambridge University Press, Cambridge (2011)
3. Bassi, F., Rebay, S.: High-order accurate discontinuous finite element solution of the 2D Euler equations. *J. Comput. Phys.* **138**(2), 251–285 (1997)
4. Beck, A.D., Bolemann, T., Flad, D., Frank, H., Gassner, G.J., Hindenlang, F., Munz, C.-D.: High order discontinuous Galerkin spectral element methods for transitional and turbulent flow simulations. *Int. J. Numer. Methods Fluids* **76**(8), 522–548 (2014)
5. Bernardini, M., Pirozzoli, S.: Wall pressure fluctuations beneath supersonic turbulent boundary layers. *Phys. Fluids* **23**(8), 085102 (2011)
6. Gassner, G.J., Beck, A.D.: On the accuracy of high-order discretizations for underresolved turbulence simulations. *Theor. Comput. Fluid Dyn.* **27**(3–4), 221–237 (2013)
7. Hindenlang, F., Bolemann, T., Munz, C.-D.: Mesh curving techniques for high order discontinuous galerkin simulations. In: Kroll, N., Hirsch, C., Bassi, F., Johnston, C., Hillewaert, K. (eds.) *IDIHOM: Industrialization of High-Order Methods-A Top-Down Approach*, vol. 128, pp. 133–152. Springer International Publishing AG, Cham (2015)
8. Jones, L.E., Sandberg, R.D.: Numerical analysis of tonal airfoil self-noise and acoustic feedback-loops. *J. Sound Vib.* **330**(25), 6137–6152 (2011)
9. Moon, B., Jagadish, H., Faloutsos, C., Saltz, J.: Analysis of the clustering properties of the Hilbert space-filling curve. *IEEE Trans. Knowl. Data Eng.* **13**(1), 124–141 (2001)
10. Pionniau, S., Dussauge, J.P., Debieve, F.F., Dupont, P.: A simple model for low-frequency unsteadiness in shock-induced separation. *J. Fluid Mech.* **629**, 87–108 (2009)
11. Schlatter, P., "Orl"u, R.: Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.* **659**, 116–126 (2010)
12. Tam, C.: Discrete tones of isolated airfoils. *J. Acoust. Soc. Am.* **55**, 1173–1177 (1974)
13. Werner, M., Würz, W., Krämer, E.: Experimental investigations of tonal noise on a vehicle side mirror. In: *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, Munich, Germany. Springer 2014, 2014, submitted
14. White, F.M.: *Viscous Fluid Flow*. McGraw-Hill, New York (1991)