14th International Symposium on Turbulence and Shear Flow Phenomena (TSFP14)
Heidelberg, Germany, July 28-31, 2026

Turbulence Modelling in the Age of GPU Power

Florian R. Menter
Ansys Germany GmbH
Staudenfeldweg 20/ 83624 Otterfing
Germany
florian.menter@synopsys.com

Katia Guseva
Ansys Germany GmbH
Staudenfeldweg 20/ 83624 Otterfing
Germany
katia.guseva@synopsys.com

David Flad
Ansys Germany GmbH
Staudenfeldweg 20/ 83624 Otterfing
Germany
david.flad@synopsys.com

Patrick Sharkey
Ansys UK Ltd.
United Kingdom
patrick.sharkey@synopsys.com

André Nogueira
ESSS
Florianópolis - SC, 88032-700 Brazil
anogueira@synopsys.com

Rodrigo Peralta
ESSS
Florianópolis - SC, 88032-700 Brazil
peralta@esss.com

ABSTRACT

The impact of newly available GPU compute power on industrial CFD simulations will be analysed. The computational elements required for industrial LES will be presented. The method will be applied to numerous test cases. Special consideration will be given to the minimum resolution requirements for the prediction of boundary layer within engineering accuracy expectations.

INTRODUCTION

Computational Fluids Dynamics (CFD) is one of the most computationally intensive disciplines in science and engineering. This is a direct consequence of the underlying multiscale dynamics of turbulence, featuring eddies which are many orders of magnitude smaller than the overall dimensions of the bodies investigated. The most extreme examples are aero/hydro-dynamic flows, like commercial aircrafts or marine vessels with body dimensions in the range of $L_{body}{\sim}100m$ featuring turbulence eddies as small as $L_{eddy}{\sim}10^{-6}m$ and with correspondingly high frequencies. In Direct Numerical Simulations (DNS), all such eddies need to be resolved in space and time, which is entirely out of reach even for more benign technical applications. A detailed discussion of meshing requirements for DNS/LES can be found in Spalart (2000), Yang and Griffin (2021), Agrawal et al. 2024; Piomelli (2014).

It is generally accepted that the smallest Kolmogorov scales do not significantly contribute to technical mixing processes and can therefore safely be ignored, leading to the concept of Large Eddy Simulation (LES). However, this assumption breaks down in the inner region of boundary layers, where the smallest eddies act as a transfer mechanism between the outer flow and the wall shear stress. LES in industrial CFD was therefore limited for the most part to free shear flows and mixing processes, the most prominent examples being applications like gas turbine combustors or mixing vessels etc.

Even for flows with dominating free shear characteristics, there are typically walls with corresponding boundary layers involved. This led to the introduction of hybrid RANS-LES models like Detached Eddy Simulation (DES) and its many variants (Spalart 2000; Shur et al. 2015; Menter 2016; Deck and Renard 2020). Their initial intent was to cover all boundary layers entirely in RANS and all free shear flows in LES mode. RANS and Hybrid models can be solved even on moderate CPU clusters and are routinely used for industrial CFD applications. However, in such simulations, the accuracy of the boundary layer results is still entirely reliant on the underlying RANS model.

The advent of GPU power in the last decade is opening new options for turbulence strategies even for boundary layer flows. While DNS and Wall-Resolved LES (WRLES) are still out of reach, simplified LES approaches which bridge the eddies in the innermost near wall layer are being developed and deployed in industrial settings. Such approaches can roughly be divided into Wall-Function LES (WFLES) and Wall-Modelled LES (WMLES) (note that different groups used different naming conventions). In WFLES (Guseva et al. 2025), a coarse grid (relative to the scale of the near wall eddies) is utilized in all three space directions and a wall-stress model is employed to bridge the gap. Many proposals for such models are available, but the most widely used models are still classical equilibrium formulations. In the WMLES concept (e.g. (F R Menter et al. 2023), the numerical grid is of similar resolution as for WFLES but with an option for wall-normal refinement. The lack of wall-parallel resolution is balanced by RANS models which are only active in the inner layer. In the current study only WFLES will be presented.

Aside of turbulence modelling details, two other CFD technologies become increasingly important for LES. The first is numerical accuracy. It is well understood that LES requires optimized numerical schemes to avoid excessive numerical damping of resolved eddies. This is especially true for boundary layer LES where the available resolution is still at the lower end of established resolution criteria. The second is High-Performance Computing (HPC) which need to be reformulated for efficient use on GPU clusters.

In the following, a brief description of the numerical and HPC aspects of the newly developed Fluent®-GPU solver will be given. The main part of the paper discusses turbulence model formulations and resolution requirements. Resolution requirements will be investigated somewhat different from more academic studies, where the question is typically 'which resolution is required to achieve mesh independent targets?'. At this early stage of adoption of LES to wall-bounded flows, an equally relevant question for an industrial CFD practitioner is 'how coarse can I go in resolution and still get acceptable target values?'. Such low resolution meshes can then serve as a starting point for local or global mesh refinement strategies.

THE GPU SOLVER

There are essentially two strategies for utilizing GPU power in industrial CFD codes. The first is to preserve an existing CPU code and to simply off-load computationally intensive loops onto the GPU. The second is to develop an entirely new code which is fully optimized for GPU architectures and GPU memory access pattern. An initial comparison of the two strategies carried out by the current groups HPC team concluded that only the second option was able to utilize the full potential of the new hardware. The new solver architecture allows for hardware agnostic C++ solver programming using a hardware specific programming layer which enables execution on a wide range of hardware (including CPUs).

The new Fluent®-GPU solver is based on the same underlying technology as the Fluent®-CPU solver:

- Cell centered finite volume scheme
- Polyhedral cell data structures
- Body fitted wall treatment
- 2nd order time and space discretization
- Implicit time integration (with potential explicit options)
- Pressure- and Density based solvers

All simulations shown here us the pressure-based solver with a SIMPLEC update strategy.

LES OPTIMIZE NUMERICS

In addition to the conventional numerics elements listed above, significant effort was invested into a low-dissipation space discretization. The space discretization is skew-symmetric with an optional 4th order dissipation term. All simulations in the present study are computed with the fully skew-symmetric scheme.

Time integration relies on a 2nd order Gears scheme (2nd order backward Euler).

HPC PERFORMANCE

The main motivation for the new solver lies in the optimal utilization of the compute power of modern GPUs. Figure 1 shows the scaling for the DrivAer test case with 250 million cells on different hardware platforms. From a hardware standpoint, it shows that one NVIDIA H100 GPU corresponds to roughly 1500 CPU cores (AMD EPYC 7763 64).

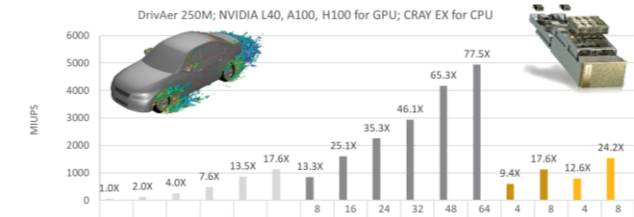

	512	1024	2048	4096	8192	12288	8 L40 GPUs	16 L40 GPUs	24 L40 GPUs	32 L40 GPUs	48 L40 GPUs	64 L40 GPUs	4 A100 GPUs	8 A100 GPUs	4 H100 GPUs	8 H100 GPUs
MIUPS	64	125	252	483	863	1124	848	1598	2254	2939	4164	4943	602	1124	802	1541
Parallel efficiency	100%	98%	99%	95%	85%	73%	100%	94%	89%	87%	82%	73%	100%	93%	100%	96%
Time per iteration	3.92	2.00	0.99	0.52	0.29	0.22	0.29	0.16	0.11	0.09	0.06	0.05	0.41	0.22	0.31	0.16

Figure 1. Performance in Million Updates per second (MIUPS) for an automotive application with 250 million cells.

Alternatively, using a more a mid-level GPU cluster with 64 L40 NVIDIA cores, around 36,000 time steps can be performed in 60 minutes (using 2 iterations per time step). It needs to be emphasized that the benchmark is a 'clean' aerodynamics WFLES test without additional complexity like rotating wheels etc.

TURBULENCE MODELING

Wall Function LES (WFLES)

The goal of a wall function is to provide the wall shear stress at the wall face of the near wall cell, based on information available at the cell center of the donor cell (typically the first or second cell from the wall). Applying the wall shear stress at the wall face closes the discretized momentum equations and allows the solution of the discretized system of equations. Historically, wall functions of the form $u^+ = f(y^+)$ have been based on the non-dimensional wall distance y^+defined as $y^+ = \frac{u_\tau y}{\nu}$ (u is the velocity at the matching location, $u^+ = \frac{u}{u_\tau}$, $u_\tau = \sqrt{\tau_{wall}/\rho}$, y the wall distance and ν the kinematic viscosity). Such y^+ based wall models require iterations as u_τ appears implicitly.

The current simulations use a non-iterative formulation (Guseva et al. 2025) based on $r^+ = y^+u^+ = \frac{uy}{\nu}$. In the sublayer this gives $u^+_{sub} = \sqrt{r^+}$ and a logarithmic layer can be matched using $u^+_{log} = C_1 ln(r^+ + C_2)$ with $C_1 = 2.21, C_2 = 4$. The two formulations are blended with a harmonic function $u^+_{eq}(r^+) = \left((u^+_{sub})^{-\alpha} + (u^+_{log})^{-\alpha}\right)^{-\frac{1}{\alpha}}$ with $\alpha = 6$. The function can be seen Figure 2 (left) which shows also the matching location (right) which is the second cell center from the wall for better near wall accuracy.

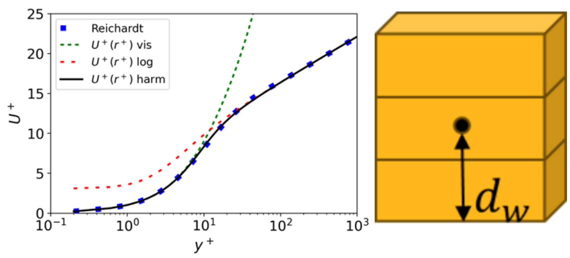

Figure 2: Blended non-iterative wall stress model

This wall stress model can be applied to any LES resolution, including WRLES and DBS, as in that limit the exact sublayer formulation is recovered. WFLES is defined here to mean that the matching location is in the logarithmic layer.

Unless stated otherwise, the LES model used in the studies is the WALE model (Nicoud and Ducros 1999) with ($C_W = 0.5$).

TEST CASES

Channel Flow

The turbulent channel flow is a widely used test case to evaluate LES models and numerics. In the current study, the emphasis is on the lower levels of resolution which can deliver 'engineering' accuracy. For this purpose, four meshes with uniform cubic cells have been created. The number behind the mesh indicator defines the number of cells across half of the channel (which can be viewed as an approximation of a boundary layer). M05 means 5 cubic cells across half the channel height, h, etc. The domain size is H=2h height, 4H length and 1.5H width. Different Reynolds numbers were achieved by adjusting the fluid viscosity.

Figure 3 shows velocity profiles for the meshes M05, M10, M20 and M40 for two Reynolds numbers. As seen, even with only 5 cells across half of the channel, the method can still reproduce the logarithmic layer faithfully.

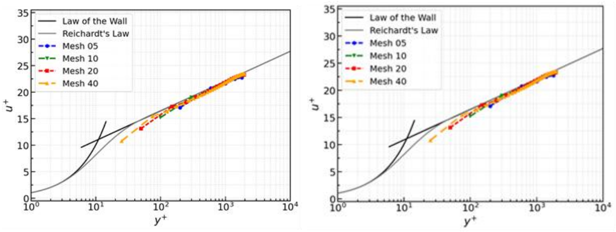

Figure 3: Velocity profiles for channel flow for meshes M05, M10, M20 and M40. Right $Re_\tau = 2000$, Left $Re_\tau = 4000$

Another stress test for the method was performed using a variation in CFL numbers from $CFL = 0.3, 1.0, 5.0$ and 10. As seen in Figure 4, the turbulence structures become highly distorted for the two higher CFL numbers. However, as seen in Figure 5, even under such drastic CFL changes, the logarithmic layer is still preserved (note that the 1st point value deteriorates – it is however passive as the matching location is at the 2nd point)

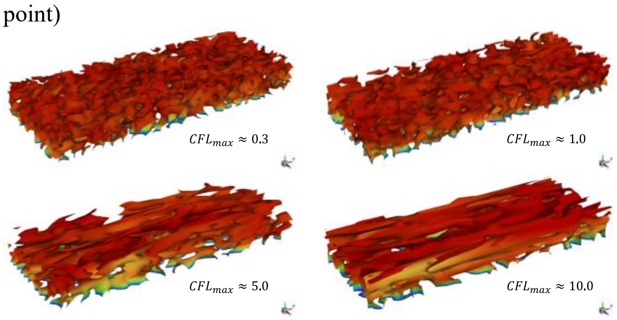

Figure 4: Channel flow using M10 with variation of CFL number ($CFL = 0.3, 1.0, 5.0, 10.0$).

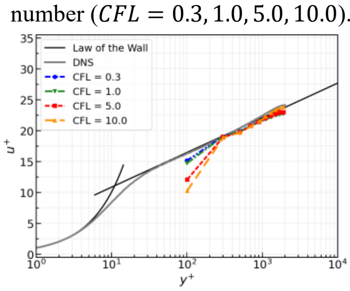

Figure 5: Velocity profiles for channel flow and M10 mesh under variation of CFL number ($CFL = 0.3, 1.0, 5.0, 10.0$).

It is to be stressed that such high CFL numbers are not recommended, but in complex industrial flows, local increases in CFL number beyond recommended values cannot always be excluded.

Further studies (not shown) have been performed which show that the LES turbulence model has only a marginal effect on this flow.

S-Shaped Diffuser

This test-case is based on the Smooth Ramp Test Case of the High-Fidelity CFD Verification Workshop 2024 which was held together with the AIAA SciTech 2024 conference (Garmann et al. 2024). The purpose of using this test case is to evaluate the accuracy of the WFLES model to predict the wall quantities (e.g., C_f and C_P) and the separation bubble over a S-Shaped Diffuser. DNS data for the skin-friction and pressure coefficients, velocity and Reynolds stresses profiles, as well as the bubble size (with the separation and reattachment positions of the flow), were used to evaluate WFLES modeling accuracy. The domain and mesh M1 are shown in Figure 6. Synthetic turbulence using the STG method of (Shur et al. 2014) was applied at the inlet.

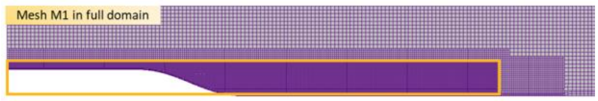

Figure 6: Sketch of the computational domain Mesh M1

Four meshes as shown in Figure 6 and given in Table 1 are used. The value of δ_t/Δ in the table refers to the number of cells across the boundary layer just upstream of the diffuser, staring with a very coarse mesh M1 with only around 5 cells.

Mesh	Δ (mm)	δ_t/Δ	Cells (millions)
M1	8	5	0.57
M2	4	10	3.7
M3	2	15	27
M4	1	30	200

Table 2: Mesh size of mesh with systematic refinement (Δ refers to the edge length of on the body surface).

The velocity profiles just upstream of the diffuser can be seen in Figure 7. Even on the coarsest mesh with only 5 cells across the boundary layer, the DNS profile is well reproduced.

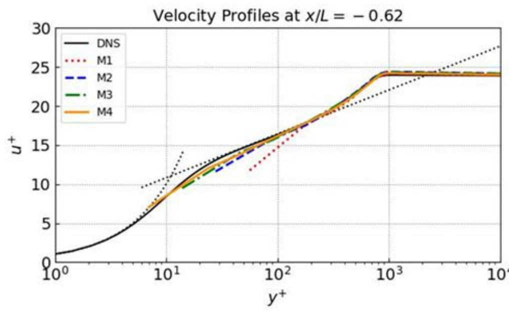

Figure 7: Velocity profiles just upstream of diffuser for the different meshes.

Figure 8 shows the corresponding wall shear stress coefficient for the same meshes. It needs to be stressed that the wall shear stress is essentially constant just downstream of the inlet, indicating a high-quality synthetic turbulence insertion.

The separation bubble is overall well represented, albeit somewhat smaller than for the DNS reference.

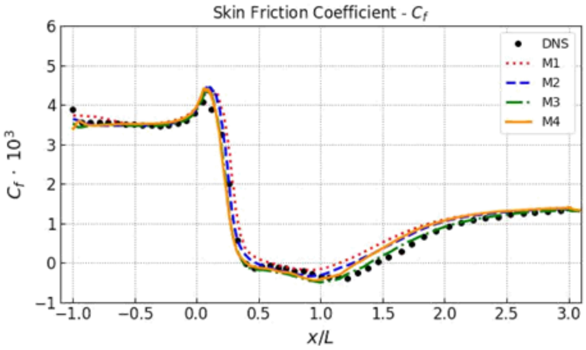

Figure 8. Wall shear stress coefficient for diffuser for different meshes.

The velocity profiles shown in Figure 9 depict again an overall good agreement with the DNS reference. The coarser grids tend to slightly delay separation, whereas mesh M3 shows a significant earlier separation onset, which is corrected on the finest mesh. The reason for the non-monotonic grid convergence is at current point not well understood.

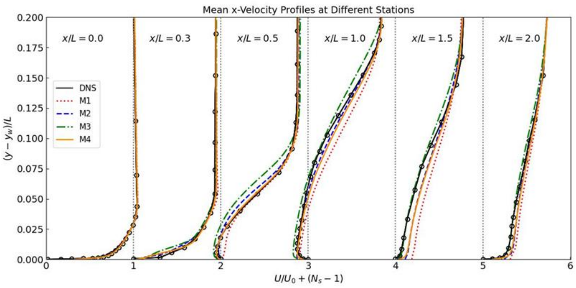

Figure 9. Velocity profiles for diffuser in separation and reattachment region for four meshes.

Overall, the accuracy of the WFLES is of good engineering quality, especially when considering the severe under-resolution of the coarser grids shown in Figure 10.

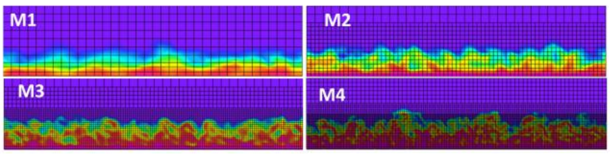

Figure 10: Contours of vorticity upstream of diffuser with grid lines.

NACA 4412 Airfoil

The NACA 4412 is a classic benchmark case widely used in RANS and LES validation studies (for a discussion see (Mallor et al. 2024). Experimental investigations of the flow past the NACA 4412 under near-stall conditions have been carried out by Wadcock (1987) and Hastings and Williams (1984). In the experiments, the flow was tripped near the suction peak using a transition strip. Since such a strip cannot be resolved by WFLES, the flow in the simulations is allowed to transition based on available mesh resolution. Since the upstream laminar boundary layer is thin, predicting laminar-turbulent transition with WFLES is a challenge as such grids do not have sufficient stretching in wall-normal direction to resolve the laminar layer.

The airfoil geometry and the CFD domain are shown in Figure 11. The Reynolds number based on the chord length, $c = 1m$, is $Re_c = 1.5 \cdot 10^6$ and the angle of attack is $\alpha = 13.87^o$. The spanwise domain size is 10% of chord. Standard inflow and no-slip conditions are applied with periodicity in spanwise direction.

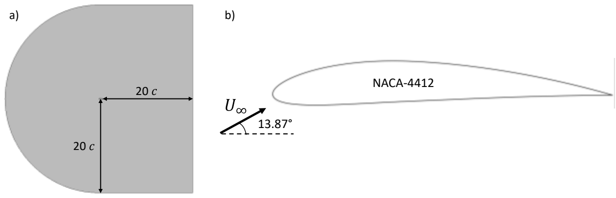

Figure 11: Geometrical details: a) computational domain characteristics; b) shape of the NACA 4412 airfoil.

Table 3 give the information on the Octree meshes employed and Figure 12 shows mesh details near the leading edge. The meshes feature a prism layer with one split which results in a cell aspect ratio of around two for the near wall cell, which is optimal for WFLES. No provisions are made to resolve the thin laminar boundary layer emanating from the stagnation point.

Table 3: Meshes for NACA 4412 Airfoil. $\boldsymbol{N_y}$ gives the number of cells across the boundary layer at mid-chord. Cell Size is the size is the edge length on the wall.

Mesh	Cell Count	N_y	Cell Size	Δt
M1	3,989,080	12	≈ 0.00144	3.33μs
M2	8,117,776	19	≈ 0.00072	2.50μs
M3	24,917,296	28	≈ 0.00036	1.25μs

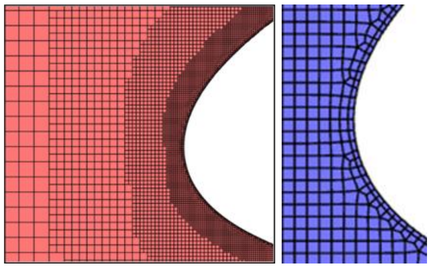

Figure 12. Mesh details near leading edge of airfoil

Figure 13 shows the non-dimensional mesh resolution on the airfoil. The wall-normal resolution lies mostly in the buffer layer ($\Delta y^+ \sim 15$). Note that the figure shows the Δy^+location of the first cell center, whereas the wall stress model uses the cell center of the 2nd cell).

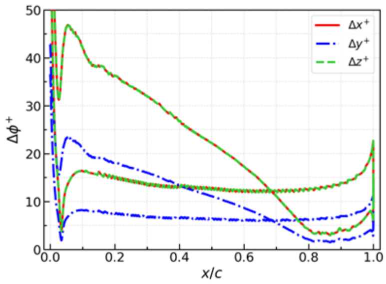

Figure 13. Non-dimensional mesh resolution $(\Delta x^+, \Delta y^+, \Delta z^+)$ (streamwise, normal, spanwise) on airfoil

Figure 14 shows q-criterion plots of the turbulence structures on the suction side of the airfoil. The results on the coarsest grid (left) show a delayed transition with a correspondingly slow formation of three-dimensional turbulence eddies. The other two mesh results appear visually as acceptable.

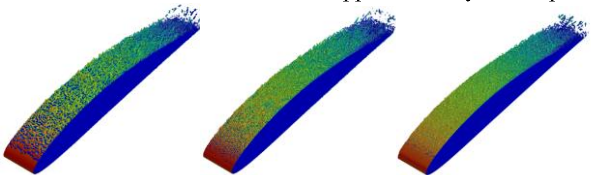

Figure 14: Q-criterion plots for NACA 4412 airfoil and meshes M1-M2-M3 (from left to right) colored by velocity

Figure 15 shows the wall shear stress for the NACA 4412 airfoil for the different meshes. Included in the plots are experimental data and a solution of the SST turbulence model which is known to capture this flow well. It is obvious that the wall shear stress for meshes M1 and M2 show delayed transition and only catches up with the reference data at mid-chord, while mesh M3 follows the SST model already much further upstream ($x/c \approx 0.1$). This is a direct result of the under-resolution of the laminar boundary layer emanating from the leading edge for M1 and M2.

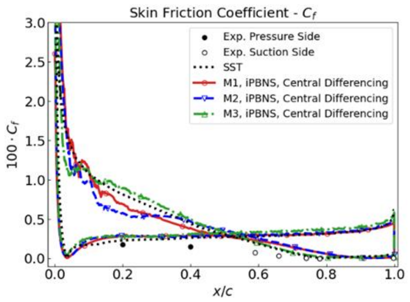

Figure 15: Wall shear stress distributions for NACA 4412 airfoil on different meshes.

Figure 16 shows the corresponding velocity profiles, again in comparison with experimental data and the SST model. The results from Meshes M1 and M2 are very close to each other, and both predict a delayed separation onset, whereas mesh M3 is very close to the data. The delayed separation is most likely strongly related to the delayed transition observed in Figure 15.

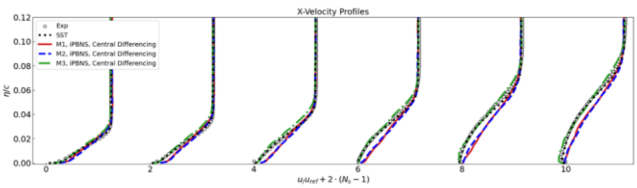

Figure 16: Velocity profiles for NACA 4412 airfoil on different meshes.

The test case can also be used to evaluate numerical performance. The solver is optimized to operate at high under-relaxation coefficients (factors of 1.0 for pressure and 0.95 for momentum). Such high values are designed to allow low iteration counts within a time step. Figure 17 shows velocity profiles when using two, five and ten iterations. Even with only two iterations, the solution is already established.

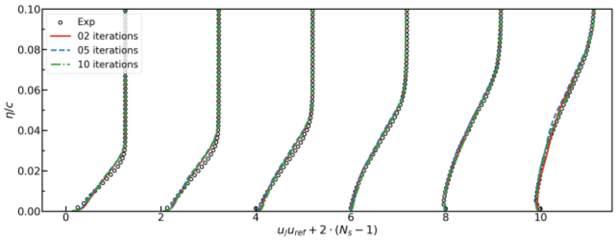

Figure 17: Velocity profiles for mesh M3 using different number of iterations for the SIMPLEC method

To test numerical accuracy, the pressure-based skew-symmetric scheme is compared to the well-established LES/DNS KEEP method (Asada and Kawai 2018). The comparison can be seen in Figure 18 for mesh M2. Both methods produce almost identical results, confirming the viability of the current discretization for LES applications.

.

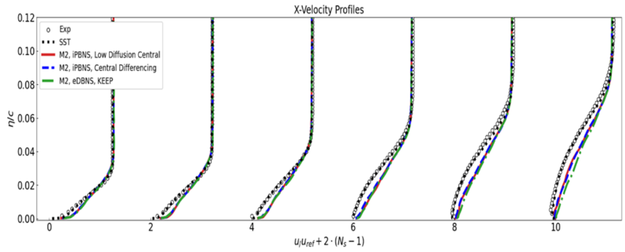

Figure 18: Velocity profiles for NACA 4412 airfoil for SIMPEC and KEEP numerical schemes.

DrivAer Car

The DrivAER (Heft et al. 2012) notchback geometry is shown in **Error! Reference source not found.**. The geometry is based on a generic design with specific characteristics to represent features akin to a common passenger vehicle. These include wheels, wheel hubs and wheel housing; side mirrors; and detailed underbody components. The notchback configuration features an inclined rear passenger section and a horizontal trunk. A configuration with a closed-grill and closed-engine bay was used in this investigation. The wheels of the configuration are stationary as specified by for the Auto-CFD workshop website for this test case.

Experimental results are available for two versions of the notchback configuration: one without and one with a front wheel deflector. Results for both cases can be found on the Auto-CFD workshop website.

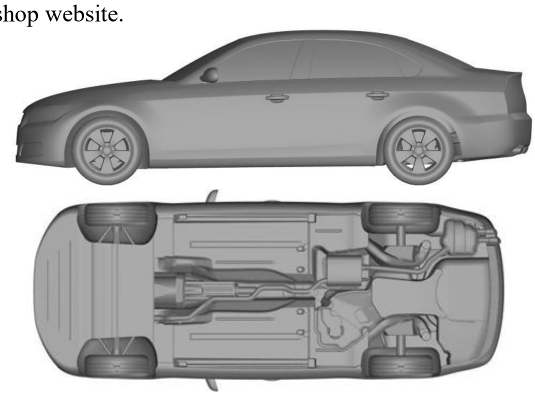

Figure 19: Geometry of the DrivAer notchback car

When designing WFLES meshes it is best to start with an estimate for the desired boundary layer resolution. Assuming a $4mm$ edge length on the body and the minimum resolution requirement of 5 cells across the boundary layer, a $2cm$ boundary layer could be resolved on such mesh, assuming that 5 cells with $4mm$ are also provided in wall normal direction. Estimating a projected surface area of a car with approximate dimensions of $5m$ long and $2m$ wide results on $A_p \approx 10m^2$. Accounting for the top, the floor and the side surfaces by using a factor $4x$, one gets an estimate of the aerodynamic surface of $A_{AS} \approx 40m^2$. Note that this is on the optimistic side, and additional complexities like under-floor, under-hood (if present) and wheels etc. need to be considered (often using a different resolution there). A $4mm$ surface mesh would then result in approximately $2.5 \cdot 10^6$ faces on the surface. With a wall-normal extent of 6 cells of this size (5 cubes and one prism split) one gets $\sim 15 \cdot 10^6$ cells to cover the boundary layer region of the aerodynamic surfaces.

This is reasonable starting point for a baseline mesh, considering that automotive companies currently tend to strive for overall meshes of around $100 - 200 \cdot 10^6$ cells. The baseline mesh leaves sufficient space for adding local refinement boxes in regions where they are required. It is to be emphasized however, that the assumption of a $2cm$ boundary layer thickness is not satisfied for the front aerodynamic surfaces of the car, and some decrease of accuracy needs to be expected there. Starting with such a base wall cover, numerous local refinement boxes are applied:

Refinement boxes are added (see Figure 20 and Figure 21):

- Front wheel box
- Underfloor between the car and the tunnel wall
- Wake box
- Front lip refinement box
- Front wheel deflector area
- Wind tunnel wall

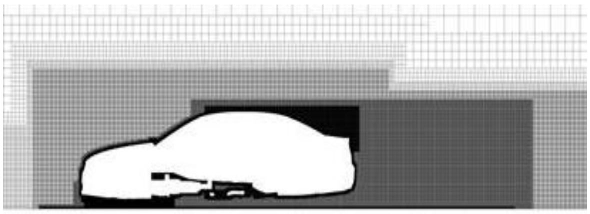

Figure 20: Refinement boxes in center plane

The final grid has a cell count of $115 \cdot 10^6$ cells for case 2a. Pictures of the refinement boxes used are shown in Figure 21. This mesh is also used as a starting point for further optimization in the future.

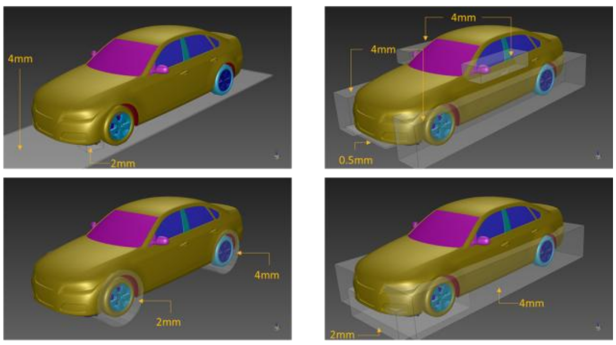

Figure 21: Refinement boxes in three-dimensional view

The time step size for the simulations was $4.5 \cdot 10^{-5}s$ for both meshes. The time step was selected to achieve CFL numbers below CFL<1. The convergence criterion was 0.5 which was achieved typically in 2-3 iterations.

The SIMPLEC P-V coupling algorithm, using the arithmetic Rhie-Chow method, linear discretization for pressure and under-relaxation factors of 1.0 and 0.95 for pressure and momentum, respectively, were used in the computations.

The flow was initiated by the GEKO RANS model. The WFLES was then run for $2s$ physical time to establish the solution and another $2s$ for averaging. The $2s$ correspond to $\sim$17 convective passings of the flow over the car). Figure 22 shows the turbulence structures resolved on the baseline mesh.

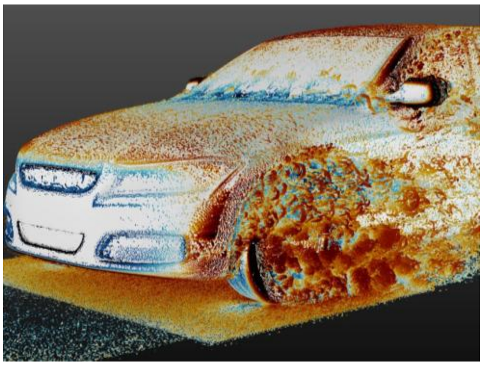

Figure 22. Q-criterion colored with mean velocity

Figure 23 shows the velocity profile below the car in the center plane at $z = -0.2376m$ halfway between the car bottom surface and the tunnel wall. The agreement with the experiment is satisfactory with only a moderate deviation around the mid-section of the car.

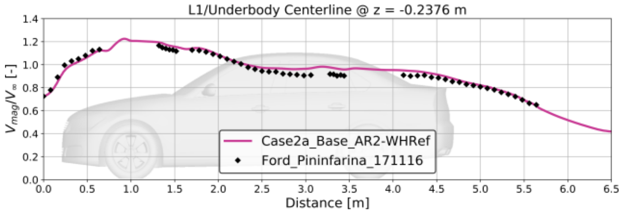

Figure 23: Velocity profile along centerline below the car for WFLES against exp. data.

Figure 24 shows the velocity profiles below the car in transverse direction also located at a height of $z = -0.2376m$. The location of the profiles is given in the picture. It is important to note that these stations are in highly complex and sensitive regions of the car. The agreement between CFD and experiments is overall very good with some more noticeable differences observed in the region around the back wheel.

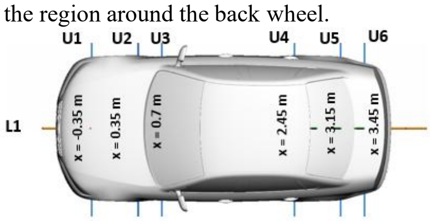

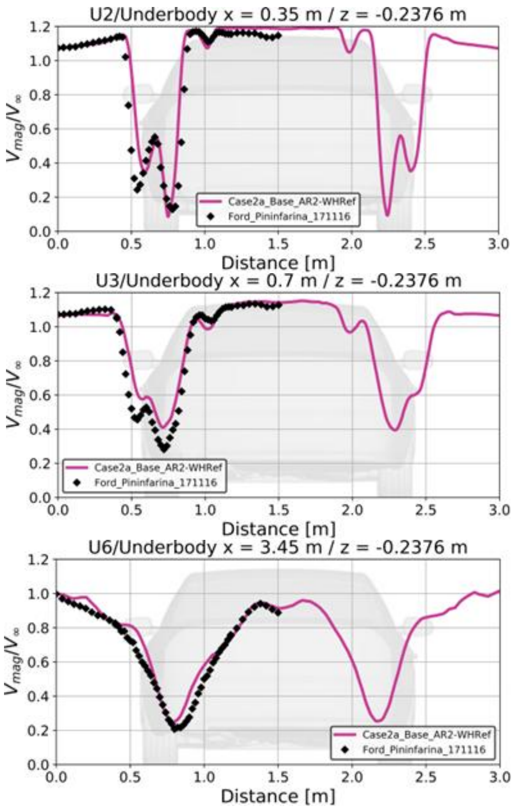

Figure 24: Velocity profiles at various critical transfers lines below the car for WFLES against exp. data.

The pressure distributions in the center of the car are shown in Figure 25. The agreement with the experimental data is again very close. A slightly flatter distribution between the back window and the trunk section on the upper centerline indicates a slightly stronger separation then in the experiments.

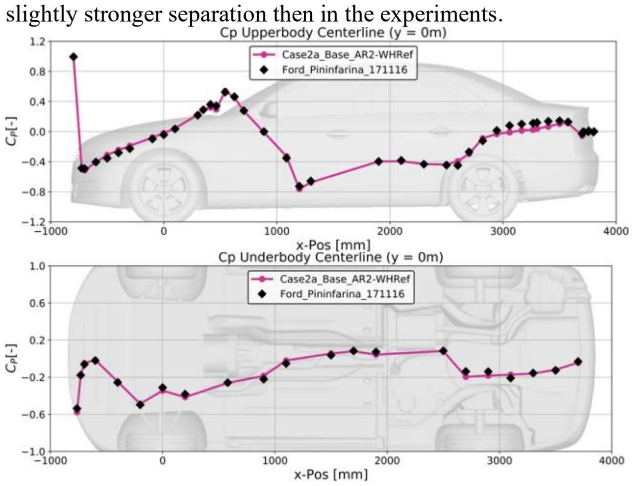

Figure 25: Pressure distribution in center plane of the car. Upper picture. Top centerline. Lower picture: Bottom centerline

Finally, a comparison for the wall shear stress is shown in Figure 26. Since no experimental data are available, the results of the WFLES are compared against a hybrid RANS LES simulation for which the boundary layer was resolved on a $y^+{\sim}1$ mesh using the SST model. The SST model is considered a reliable reference to judge overall model consistency. As expected, there is a downward shift on the hood and on the windshield, where boundary layers cannot be resolved properly on the current WFLES grid. Note however, that some turbulence structures can be seen in Figure 22 indicating at least some turbulence content in these regions. The wall shears stress is very well captured on the roof of the car and shows a slightly stronger separation at the back window-trunk intersection. The overall agreement in the wall shear stress is surprisingly good, considering the relatively low boundary layer resolution.

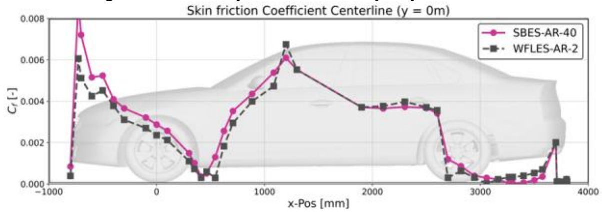

Figure 26: Wall shear stress distribution on upper center line. Comparison of WFLES on against SST-SBES on $y^+{\sim}$ grid.

The drag coefficient from the WFLES simulation is somewhat higher than the experiment $C_D^{WFLES} = 0.264, C_D^{Exp} = 0.2546$ which follows a trend observed during the AutoCFD4 workshop, where simulations generally tend to overpredict the drag forces. Possible reasons for this discrepancy will be the subject of investigations for AutoCFD5.

The above simulations demonstrate the considerations applied to generating a suitable baseline grid and set-up for the DrivAer test case based on previous building-block test cases. The comparison with experimental data demonstrates that an overall good agreement with the data was achieved. This qualifies the 115 million cell grid as a good starting point for further detailed investigations.

SUMMARY

The elements employed by the author's group for industrial LES simulation were described. The central aspects are:

- High computational speed due to native optimizing of the CFD code for GPUs.
- Low dissipation numerical method.
- Careful linearization to allow for high under-relaxation factors and low iteration counts within a time step.
- Carefully formulated and optimized near wall treatment.
- Octree based meshing strategy with well defined near wall prism split.
- Step-by-step validation of the method

The combination of these elements with the power of modern GPU clusters allows for the first time the application of LES to industrial flows with boundary layers within reasonable turn-around times. Considering the high rate of advancements in GPU technology and the corresponding increase in computational speed, memory and memory throughput, it is only a question of time, until the current coarse grid strategies will evolve into multi-billion cell simulations in the foreseeable future.

REFERENCES

Agrawal, Rahul, Sanjeeb Bose, and Parviz Moin. 2024. "Reynolds Number Dependence of Length Scales Governing Turbulent Flow Separation with Application to Wall-Modeled Large-Eddy Simulations." AIAA J. 62 (10). https://doi.org/https://doi.org/10.2514/1.J063909.

Asada, Kengo, and Soshi Kawai. 2018. "Large-Eddy Simulation of Airfoil Flow near Stall Condition at Reynolds Number 2.1 × 10 6." Physics of Fluids 30 (8): 8. https://doi.org/10.1063/1.5037278.

Deck, Sébastien, and Nicolas Renard. 2020. "Towards an Enhanced Protection of Attached Boundary Layers in Hybrid RANS/LES Methods." Journal of Computational Physics 400 (January): 108970. https://doi.org/10.1016/j.jcp.2019.108970.

Garmann, Daniel J., Matthew J. Schwartz, Johan Larsson, et al. 2024. "Summary of the Wall-Modeled LES Test Suite From the 2024 AIAA High-Fidelity CFD Verification Workshop." AIAA Aviation and Ascend (Las Vegas, Nevada), July 29. https://doi.org/10.2514/6.2024-3697.

Guseva, E. K., D. G. Flad, A. L. Nogueira, D. P. Graebin, and F. R. Menter. 2025. "Evaluation of a New Explicit Wall-Function Formulation for Large Eddy Simulation." ETMM15.

Hastings, R. C., and B. R. Williams. 1984. Studies of the Flow Field near a NACA 4412 Aerofoil at Nearly Maximum Lift. RAE AD-A157 750.

Heft, A., T. Indinger, and N. Adams. 2012. "Introduction of a New Realistic Generic Car Model for Aerodynamic Investigations." Paper presented at SAE 2012 World Congress & Exhibition. https://www.sae.org/publications/technical-papers/content/2012-01-0168/.

Mallor, Fermin, Ricardo Vinuesa, Ramis Örlü, and Philipp Schlatter. 2024. "High-Fidelity Simulations of the Flow around a NACA 4412 Wing Section at High Angles of Attack." International Journal of Heat and Fluid Flow 110 (December): 109590. https://doi.org/10.1016/j.ijheatfluidflow.2024.109590.

Menter, F R, E. Guseva, A. Stabnikov, A. Matyushenko, and A. Garbaruk. 2023. "A Family of Wall-Modeled Large Eddy Simulation Formulations." Paper presented at ETMM14.

Menter, Florian R. 2016. "Stress-Blended Eddy Simulation (SBES)—A New Paradigm in Hybrid RANS-LES Modeling." Progress in Hybrid RANS-LES Modelling. 137. https://doi.org/10.1007/978-3-319-70031-1_3.

Nicoud, F., and F. Ducros. 1999. "Subgrid-Scale Stress Modelling Based on the Square of the Velocity Gradient Tensor." Flow, Turbulence and Combustion 62 (3): 3. https://doi.org/10.1023/A:1009995426001.

Piomelli, U. 2014. "Large Eddy Simulations in 2030 and Beyond." Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 372 (2022): 2022. https://doi.org/10.1098/rsta.2013.0320.

Shur, Michael L., Philippe R. Spalart, Michael K. Strelets, and Andrey K. Travin. 2014. "Synthetic Turbulence Generators for RANS-LES Interfaces in Zonal Simulations of Aerodynamic and Aeroacoustic Problems." Flow, Turbulence and Combustion 93 (1): 1. https://doi.org/10.1007/s10494-014-9534-8.

Shur, Mikhail L., Philippe R. Spalart, Mikhail Kh. Strelets, and Andrey K. Travin. 2015. "An Enhanced Version of DES with Rapid Transition from RANS to LES in Separated Flows." Flow, Turbulence and Combustion 95 (4): 709–37. https://doi.org/10.1007/s10494-015-9618-0.

Spalart, P. R. 2000. "Strategies for Turbulence Modelling and Simulations." International Journal of Heat and Fluid Flow, International Journal of Heat and Fluid Flow, vol. 21 (3): 3. https://doi.org/10.1016/S0142-727X(00)00007-2.

Wadcock, A. J. 1987. Investigation of Low-Speed Turbulent Separated Flow Around Airfoils. NASA-CR-177450.

Web-Site. n.d. "AutoCFD4 Workshop." https://autocfd.org/.

Yang, Xiang I. A., and Kevin Patrick Griffin. 2021. "Grid-Point and Time-Step Requirements for Direct Numerical Simulation and Large-Eddy Simulation." Physics of Fluids 33 (1): 015108. https://doi.org/10.1063/5.0036515.