Reducing turnaround time of high-fidelity automotive aerodynamics simulations using cloud HPC resources

H. Hetmann[1], N. Ashton[2], M. Fuchs[1], T. Knacke[1], F. Kramer[1],
C. Mockett[1] & N. Schönwald[1]

[1]Upstream CFD GmbH, Bismarckstr. 10-12 / 10625 Berlin / Germany, CFD Consultant,
+49 (0)30 9921 14 004, hendrik.hetmann@upstream-cfd.com

[2]Amazon Web Services, Principal CFD Specialist

To squeeze the last percentage points of aerodynamic performance out of vehicle designs, the increased accuracy of scale resolving high-fidelity simulations is becoming an integral part of design processes within the automotive industry. The gain in accuracy unfortunately comes at the price of significant computational effort, as meshes for full-scale car geometries readily exceed 100M cells for hybrid RANS-LES approaches such as DDES. Upstream CFD's ambition is to enable the adoption of high-fidelity CFD in industrial design processes and the presentation will cover two key factors for this: an advanced, validated and automated simulation process and the availability of high-performance computing (HPC) resources to minimise turnaround times. The former is demonstrated by comparison of DDES results for the DrivAer Notchback geometry to wind tunnel data (Chang et al., 2021) and RANS. The comparison will feature the performance of a new boundary layer shielding function (Deck & Renard, 2020) combined with a grey-area enhanced DES approach (Fuchs et al., 2020).

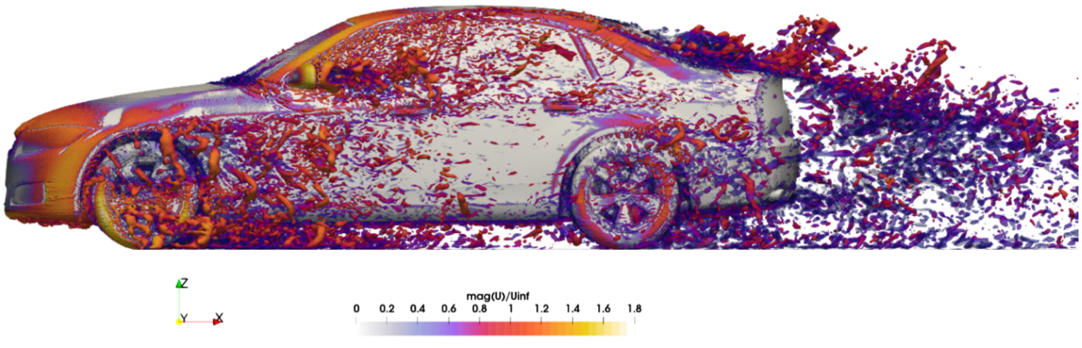

Figure 1: DrivAer case: Iso-contour of normalised Q-criterion $Q_{norm} = 400$ *coloured with normalised velocity magnitude. Simulation on a mesh of 121M cells generated automatically using bramble (https://bramblecfd.com).*

In the second part, the reduction of turnaround time using cloud HPC resources (Amazon EC2) will be discussed. Performance-tuned OpenFOAM settings (e.g. compiler flags, single precision) and the usage of the highly scalable AWS ParallelCluster infrastructure are addressed. The latter in particular promises to be a game changer to reduce simulation times for large cases, considering the scalability test conducted for a 30M cells SAE notchback geometry, see Figure 2. It is shown that the choice of PCG for the linear pressure solver has a much better scalability than GAMG. In this case, 93% scalability is achieved even at very high levels of parallelisation (only 12000 grid cells per CPU core).

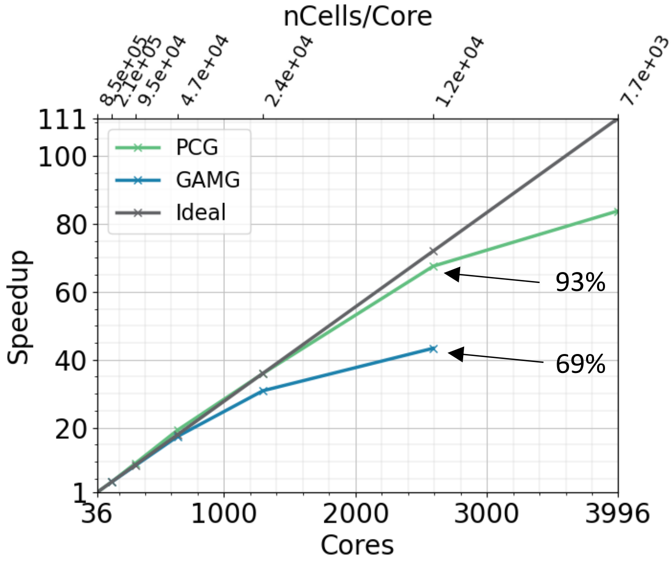

Figure 2: Scalability for SAE Notchback (30M grid cells) on Amazon EC2, showing the effect of the chosen linear pressure solver

References

Fuchs, M., Mockett, C., Sesterhenn, J., & Thiele, F. (2020). The Grey-Area Improved σ-DDES Approach: Formulation Review and Application to Complex Test Cases. *Progress in Hybrid RANS-LES Modelling* (pp. 119-130). Berlin: Springer. doi:10.1007/978-3-030-27607-2

Deck, S., & Renard, N. (2020). Towards an enhanced protection of attached boundary layers in hybrid RANS/LES methods. *Journal of Computational Physics, 400*, 108970. doi:https://doi.org/10.1016/j.jcp.2019.108970

Chang, J.-H., Shin, Y.-s., Hupertz, B., Chalupa, K., Krueger, L., Howard, K., . . . Lewington, N. (2021, 4). On the Aerodynamics of the Notchback Open Cooling DrivAer: A Detailed Investigation of Wind Tunnel Data for Improved Correlation and Reference. *SAE WCX Digital Summit.* SAE International. doi:https://doi.org/10.4271/2021-01-0958