AWS contribution to 3rd automotive CFD prediction workshop

Dr. Neil Ashton
Principal CFD Specialist Solution Architect, AWS
Visiting Fellow, University of Oxford

Special thanks to Vangelis Skaperdas from BETA-CAE Systems

Motivation

- After 1st and 2nd workshops, there are still unanswered questions.
- 2nd workshop showed there was a clear delta between CFD and exp. for C_D and most dramatically for C_{LF}
- The differences were largely consistent i.e majority of approaches could not capture C_{LF} and over-predicted C_D
- 3rd workshop is an opportunity to both explore this and whether we can capture the delta between two designs (common industry use-case)

Objectives

- Main objectives are:
  - Investigate grid resolution given the mandatory mesh was a fixed size (128M vs 585M cells)
  - Free-air vs wind-tunnel given exp. data was from a WT
  - Time-Step study (does it influence delta accuracy?) – won't show all today due to time constraints
  - Underlying RANS test of DDES – won't show all today due to time constraints
  - HPC performance

Simulation setup

- Committee ANSA high y + grids (free-air and wind-tunnel)
- Siemens Simcenter STAR-CCM+ v16.02 (Mixed Precision unless specified otherwise) to be consistent to previous workshop
- Segregated pressure-based
- 2nd order temporal (dual-time stepping)
- 2nd order hybrid upwind-central (STAR-CCM+ default ct changed from 1 to 0.1)
- Minmod gradient reconstruction
- Time step default $1e^{-4}$ (unless stated otherwise) CFL<~1 in critical regions
- Range of turbulence models (SA,SST DDES)
- All solutions started from a RANS pre-courser simulation
- All simulations run to at least 4s but some run to 10s to ensure enough time-averaging

HPC Setup

- Amazon Web Services (AWS)
- Amazon EC2 hpc6a.48xlarge (AMD Milan 96 core – 384Gb RAM – 100Gbits EFA) in Stockholm region (eu-north-1)
- Elastic Fabric Adapter (EFA) network interconnect & IntelMPI
- AWS ParallelCluster to create the HPC cluster dynamically
- 20TB Lustre (Amazon Fsx for Lustre) file system for parallel I/O
- NICE DCV to remotely visualize results on a GPU (Amazon EC2 g4dn.16xlarge) – Nvidia T4

Mesh resolution

AutoCFD2 reminder

Small overall difference

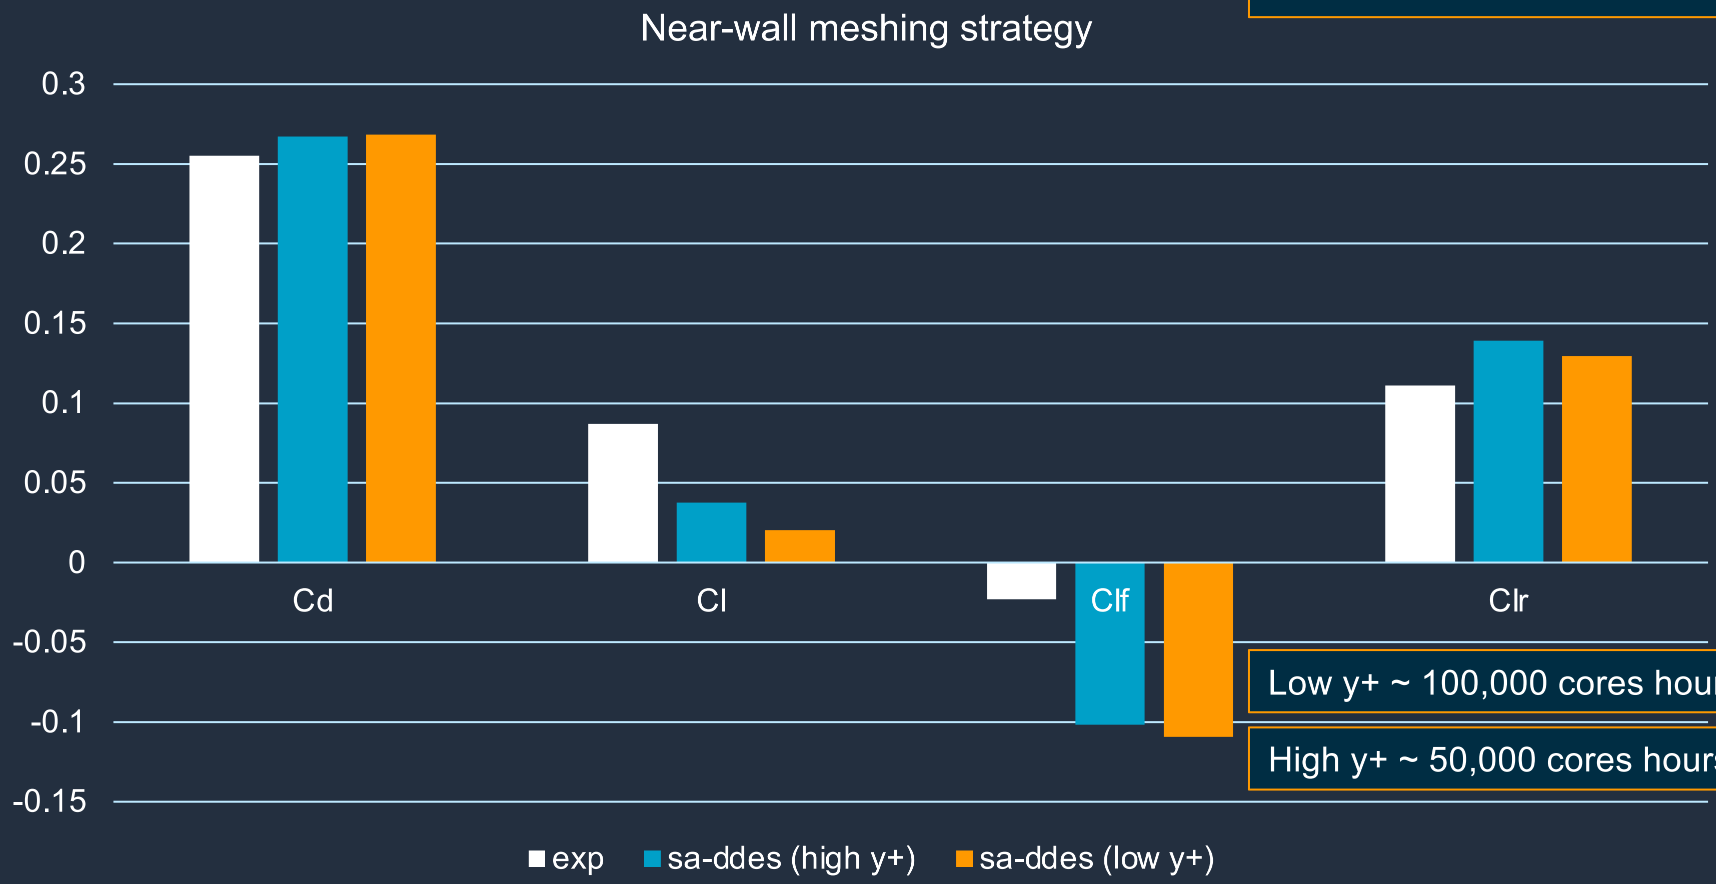

Low y+ ~ 100,000 cores hours

High y+ ~ 50,000 cores hours

Resolution requirements

I made a comment in AutoCFD2 that perhaps the grids were too coarse and they were not resolving the turbulence as much as possible – as indicated by the vorticity and the turbulent viscosity ratio.

The motivation being what quantities can we use to gauge appropriate HRLES resolution and set best-practices

Special thanks to Vangelis Skaperdas who built a 585M cell keeping exactly the same design as the committee 128M grid with just reducing the size of the cells in the refinement regions.

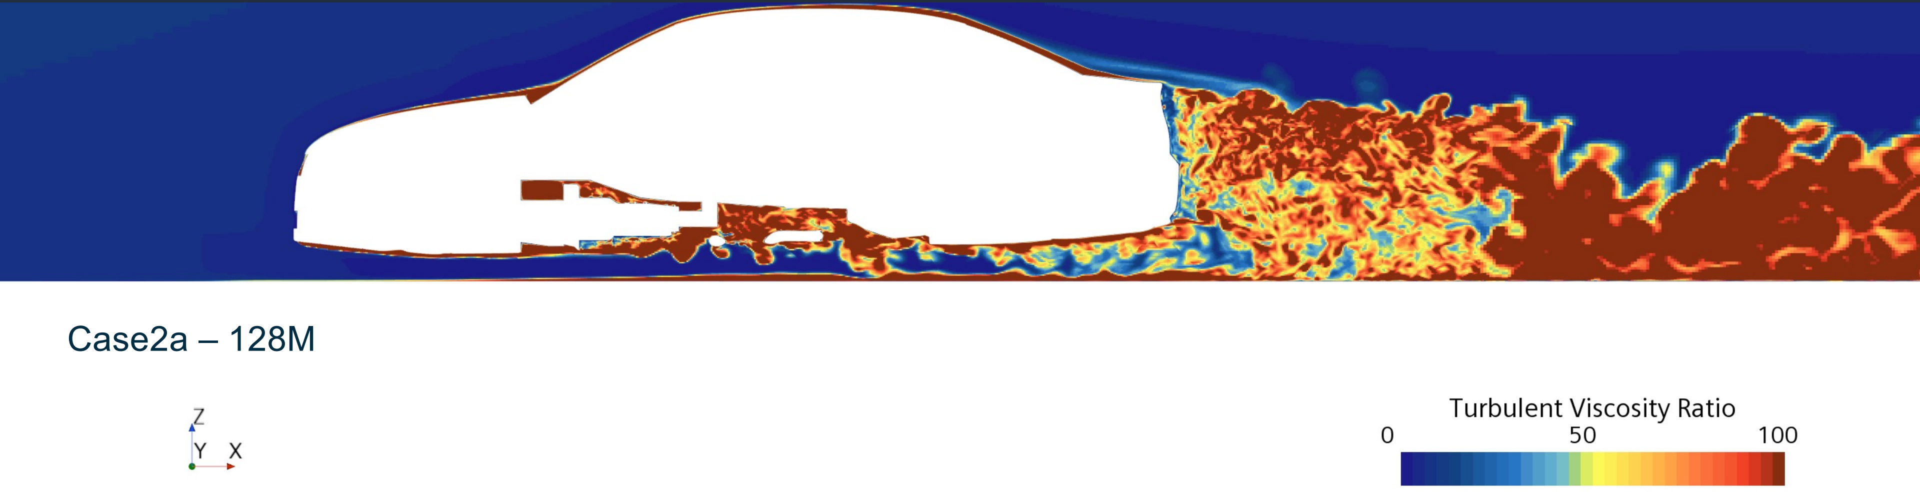

Case2a – 128M

Resolution requirements

The turbulent viscosity ratio is now less than 20 in key regions i.e shear layers

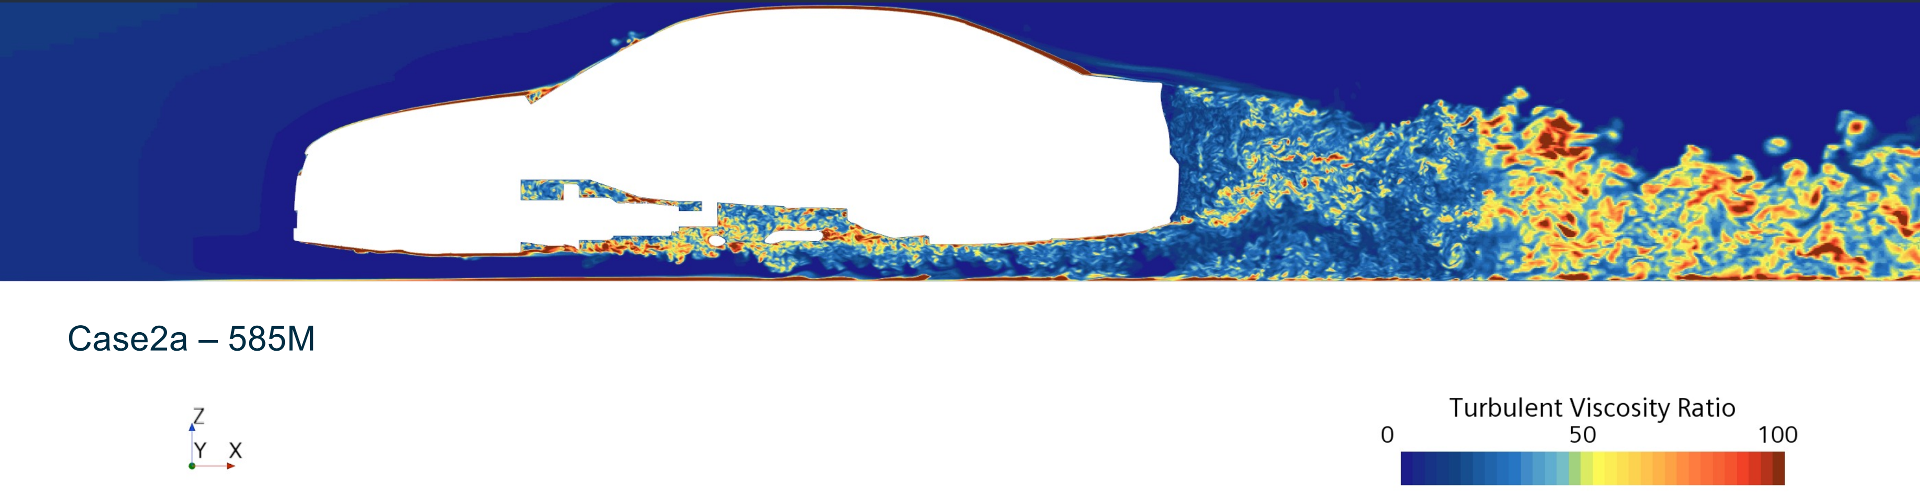

Case2a – 585M

Resolution requirements

However turbulent viscosity ratio is flow dependant and perhaps not the best general gauge of resolution.

Spanwise vorticity often gives a good feel of how well the structures are being resolved and typicallty correlations to low dissipation and low ‘modeled’ turbulence.

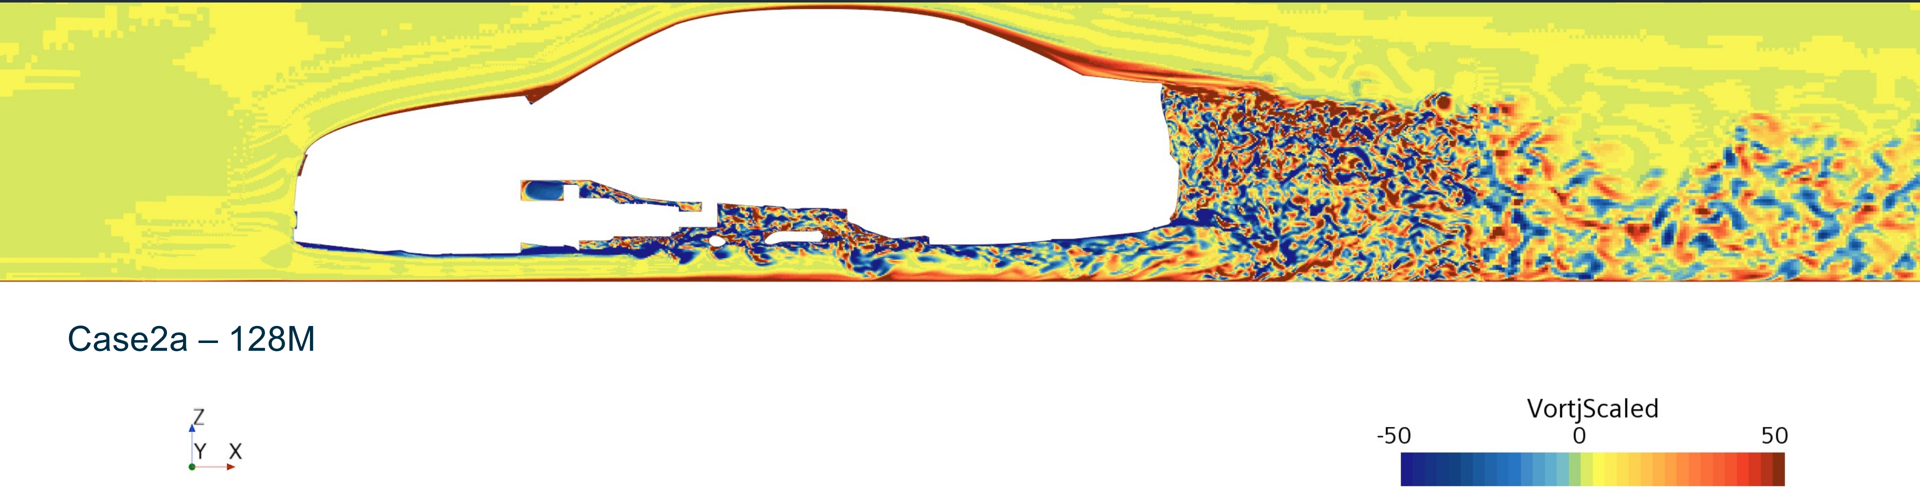

Resolution requirements

The finer grid has extremely fine-grain structures, suggesting the flow is being resolved in true LES mode.

However what is the difference in the flow field?

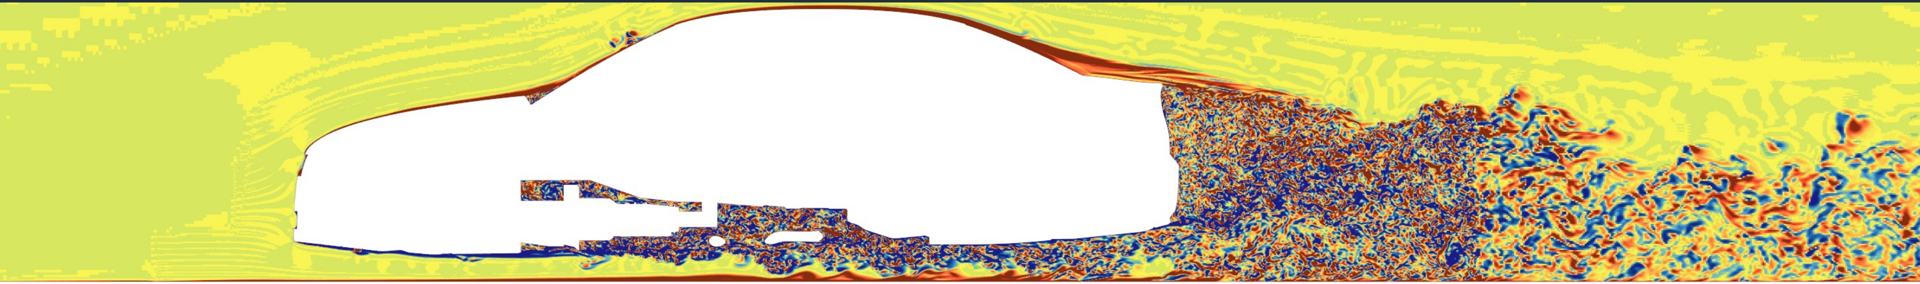

Case2a – 585M

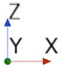

Case 2a – 128M – SA-DDES

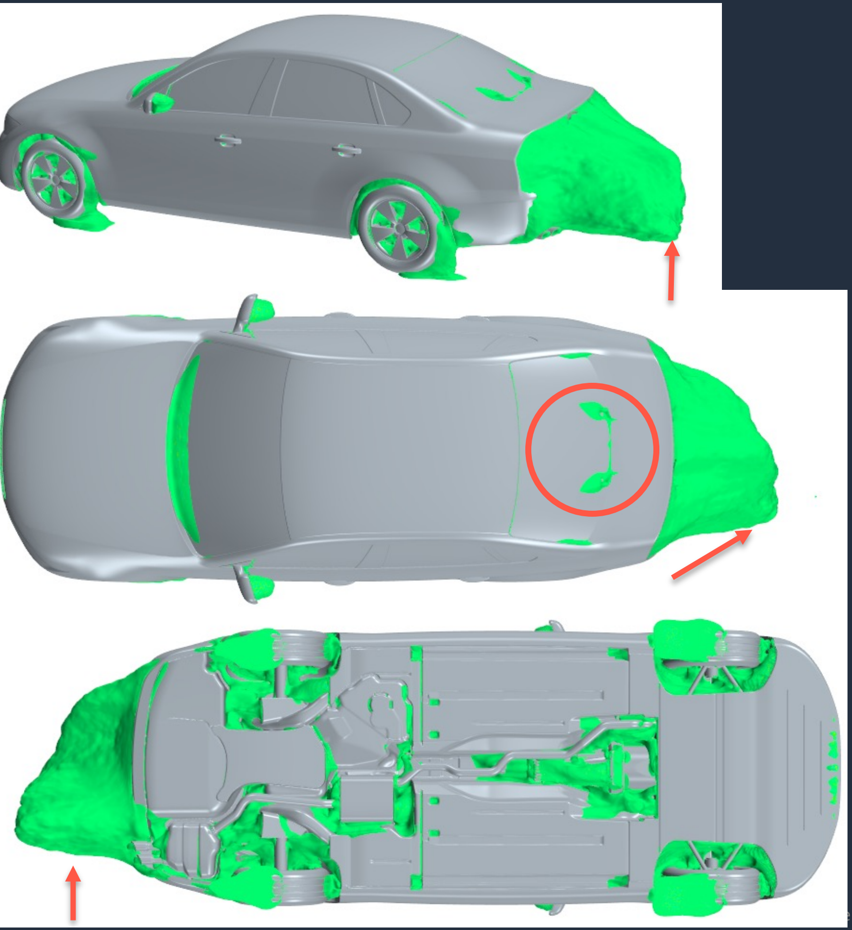

There are changes in the flow-field i.e

1) slight change in separation on the rear-window
2) the rear wake shifts more towards the center to form a tighter, thinner wake.
3) Greater separation from the wheelhouse

How does this change the forces?

Case 2a – 585M – SA-DDES

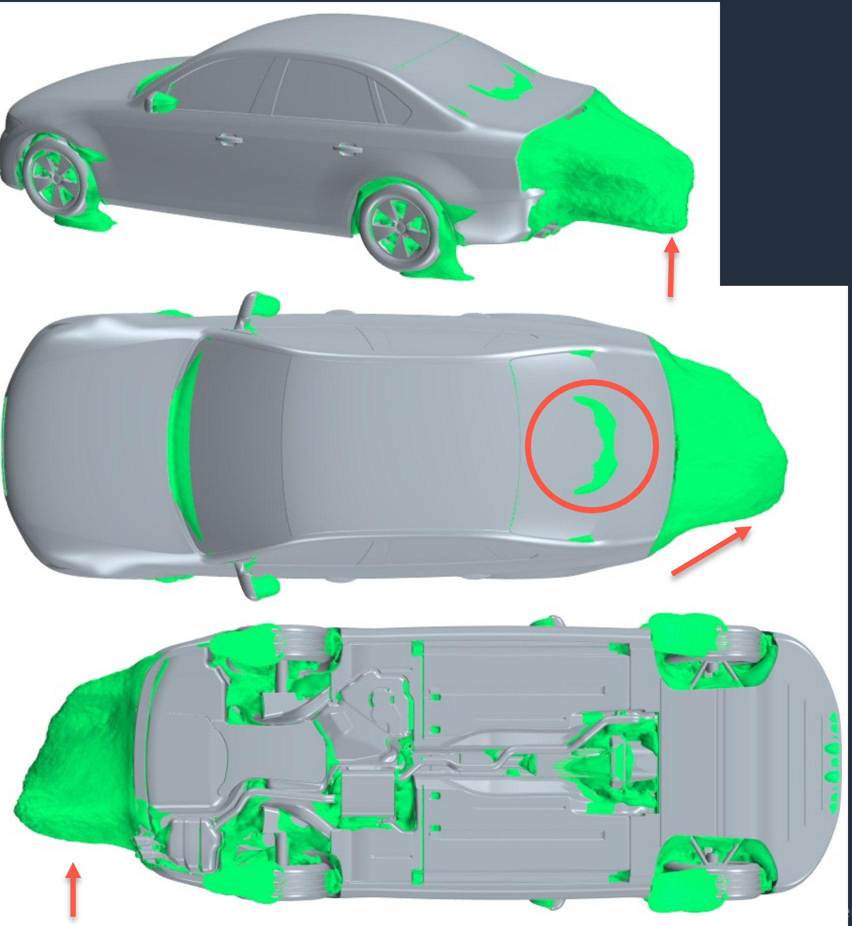

There are changes in the flow-field i.e

1) slight change in separation on the rear-window
2) the rear wake shifts more towards the center to form a tighter, thinner wake.
3) Greater separation from the wheelhouse

How does this change the forces?

585M cell vs 128M

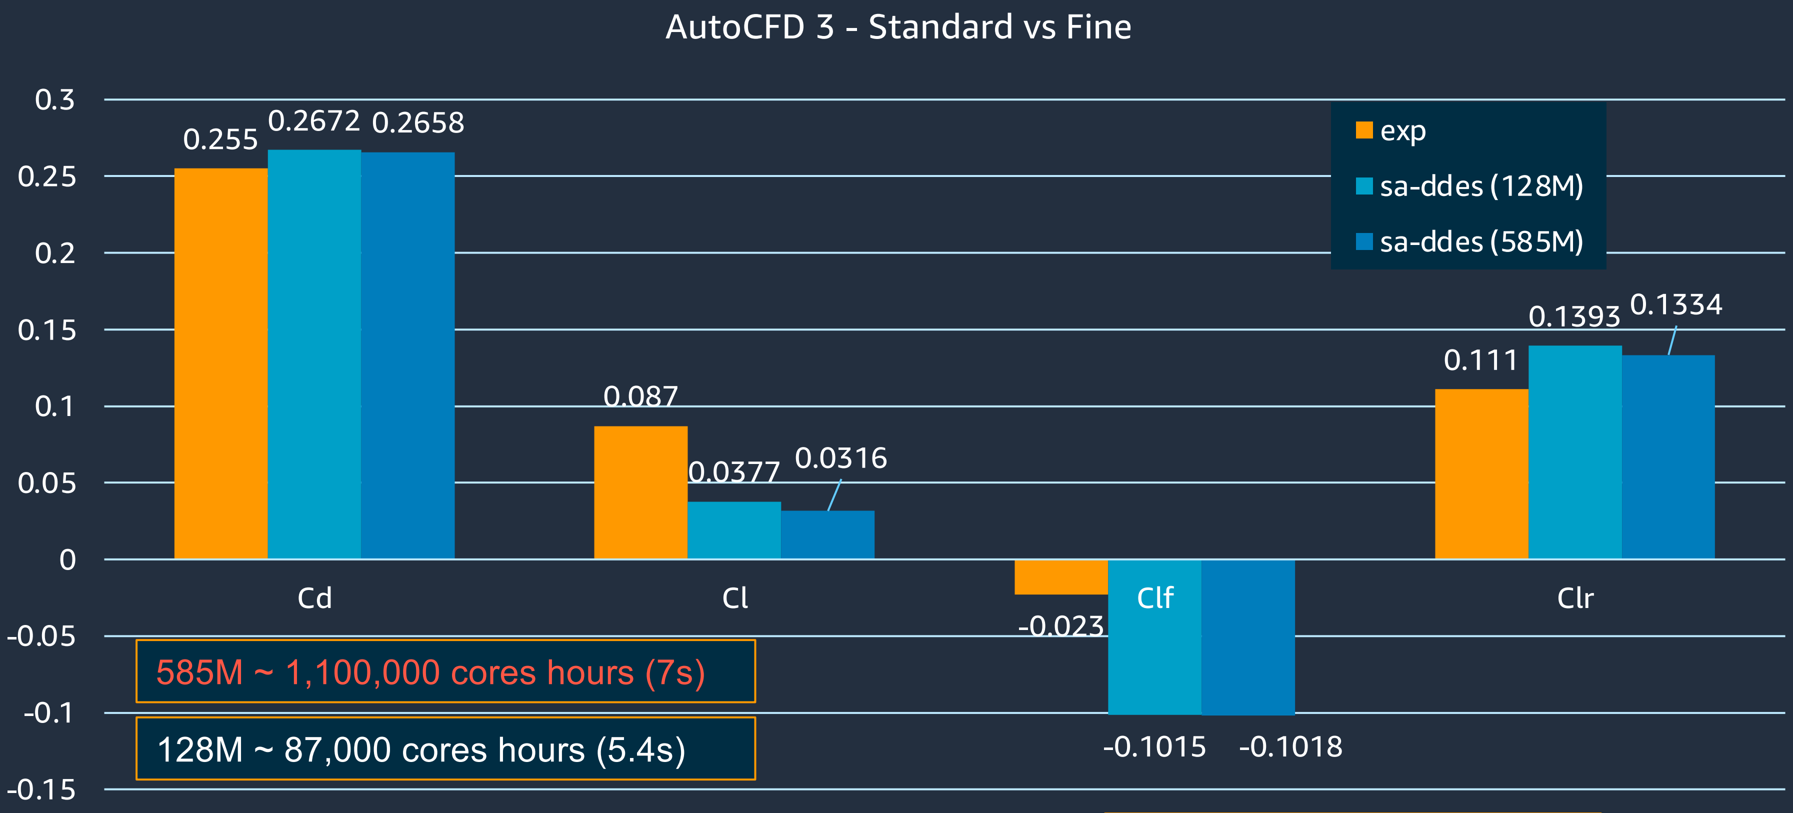

585M ~ 1,100,000 cores hours (7s)

128M ~ 87,000 cores hours (5.4s)

585M with $4e^{-5}$ time-step
128M with $1e^{-4}$ time-step

585M cell vs 128M

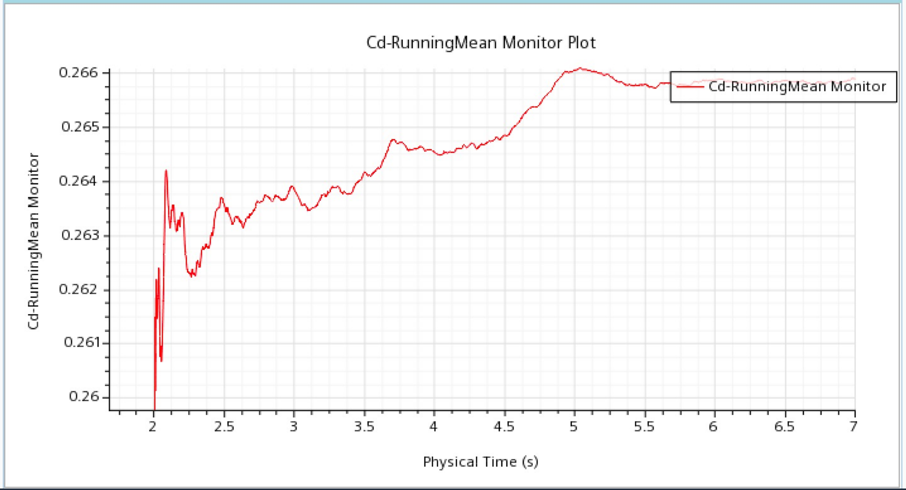

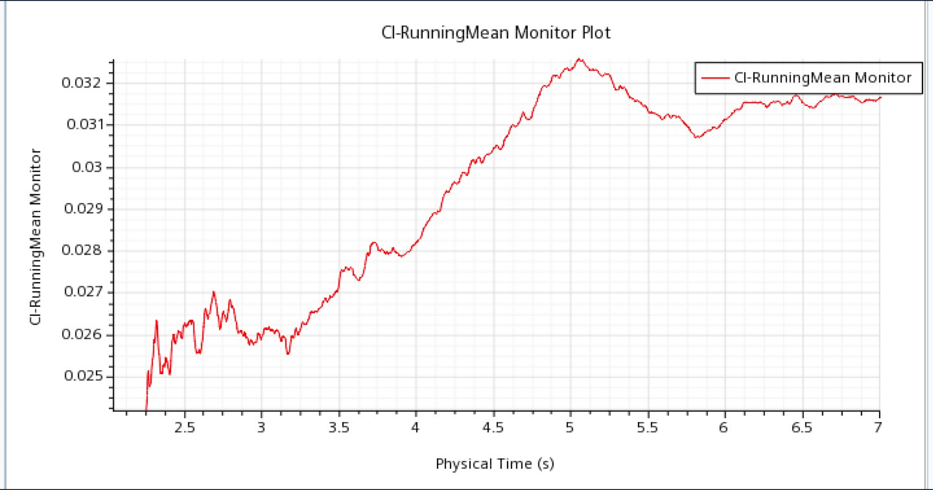

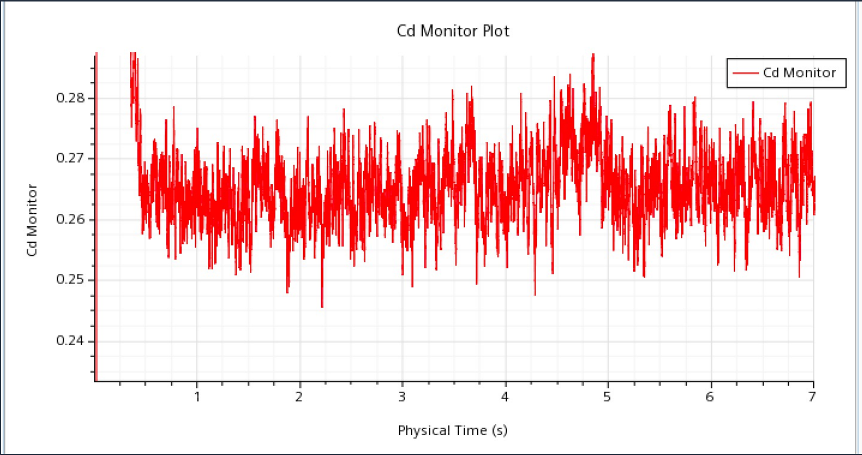

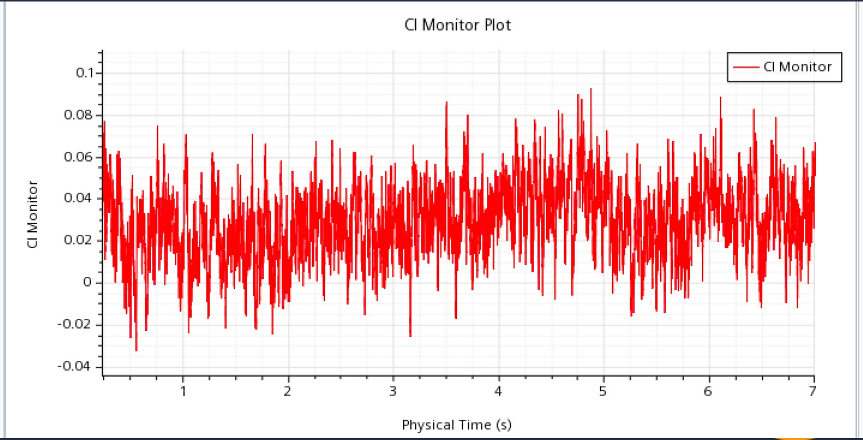

585M cell scaling

This shows that modern commercial codes investment in HPC has worked. More than 90% efficiency on more than 35k cores down to 16k cells per core. All run on AWS with HPC6a.48xlarge (AMD Milan 96 core)

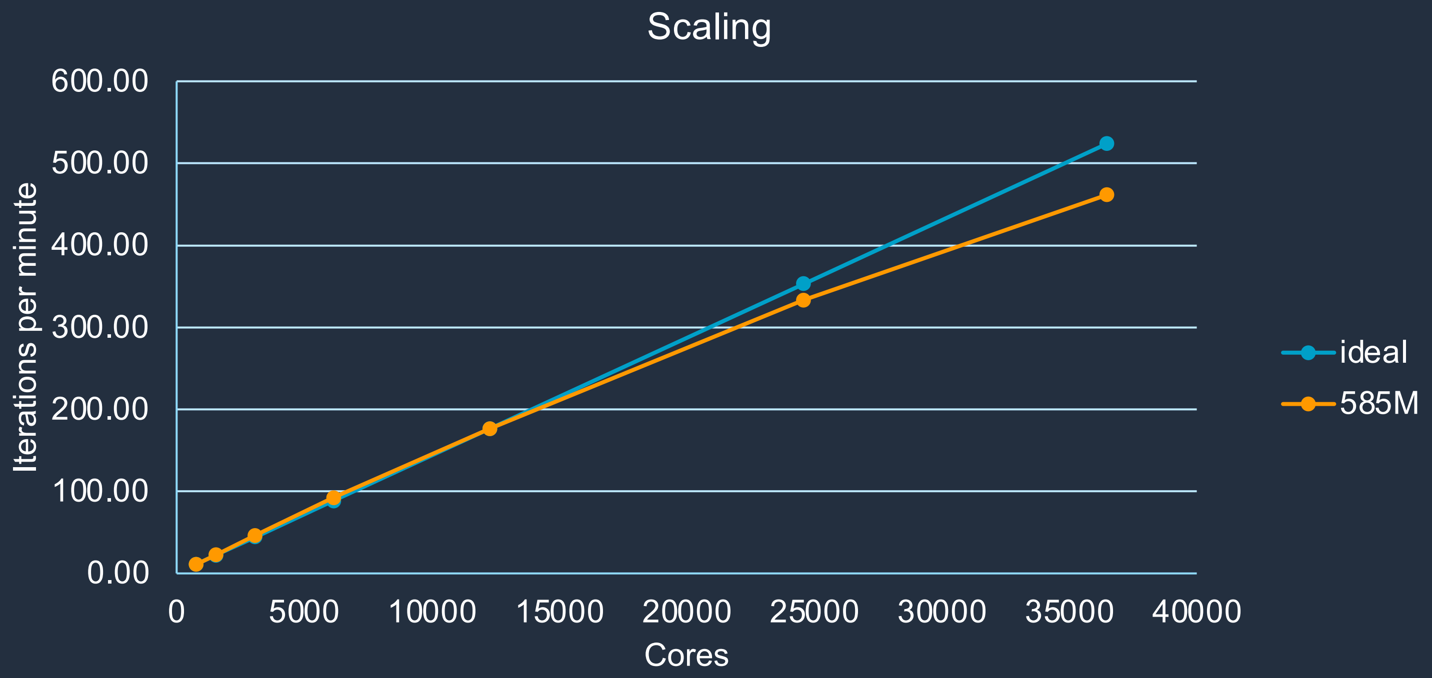

128M cell scaling

On 6144 cores for 200,000 iterations (i.e 4s at 1e^{-4} with 5 inner) = 11hrs

Have started testing on our p4de.24xlarge (A100 80GB system) with Siemens, hope to share results soon

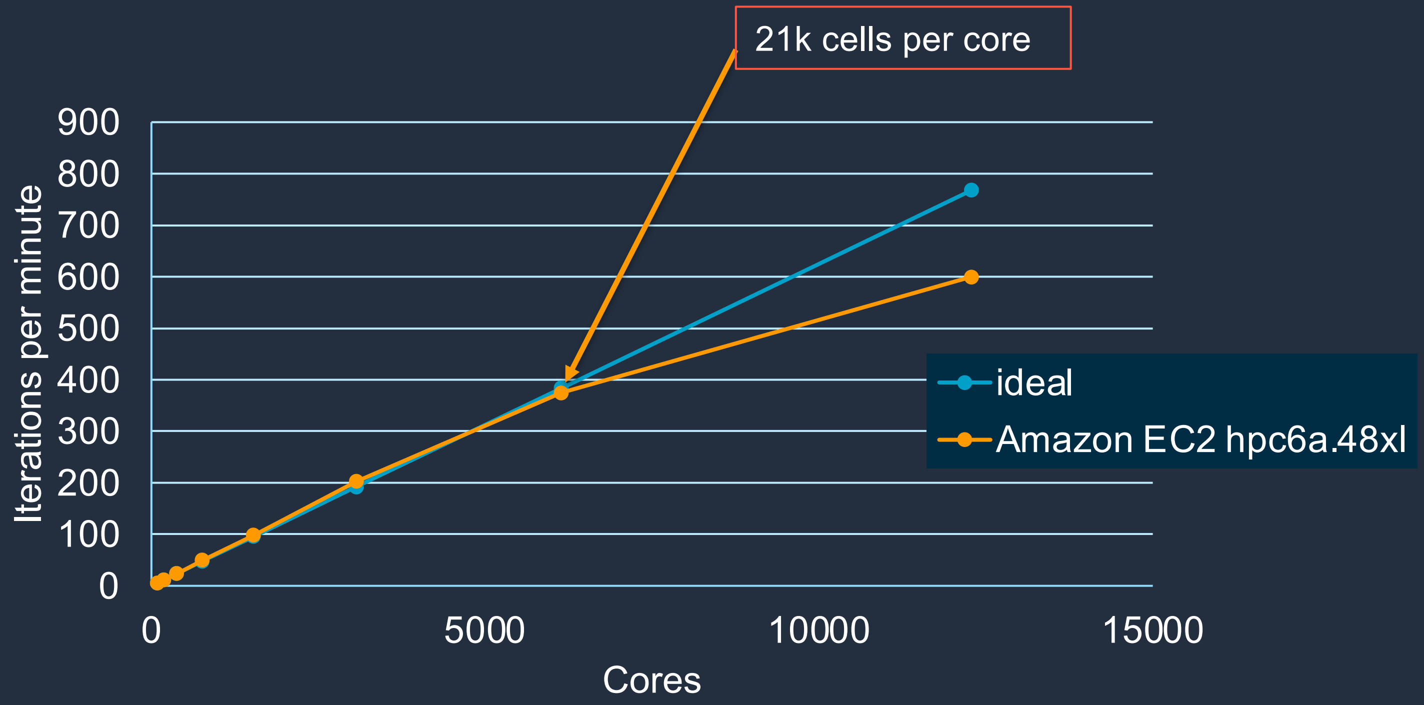

Free-air vs Wind-Tunnel

Case 2a – Free-air vs WT

Again massive thanks to Vangelis for building a model of the WT that was used in the ford experiments. We fully realize it's not perfect but it was to start a process of better understanding WT effects.

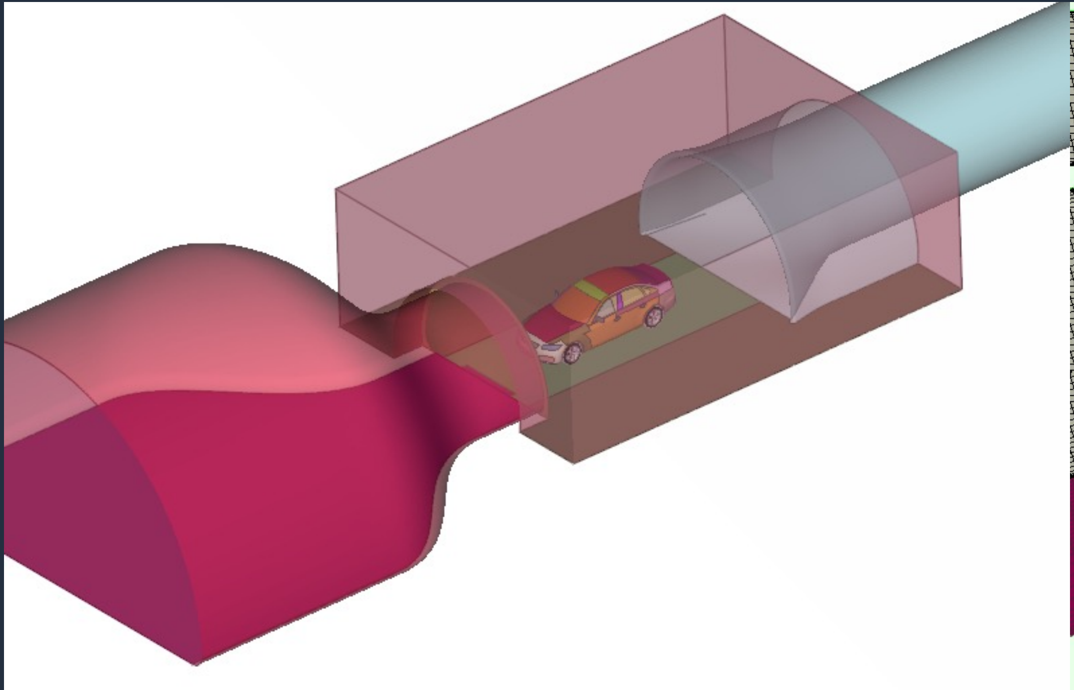

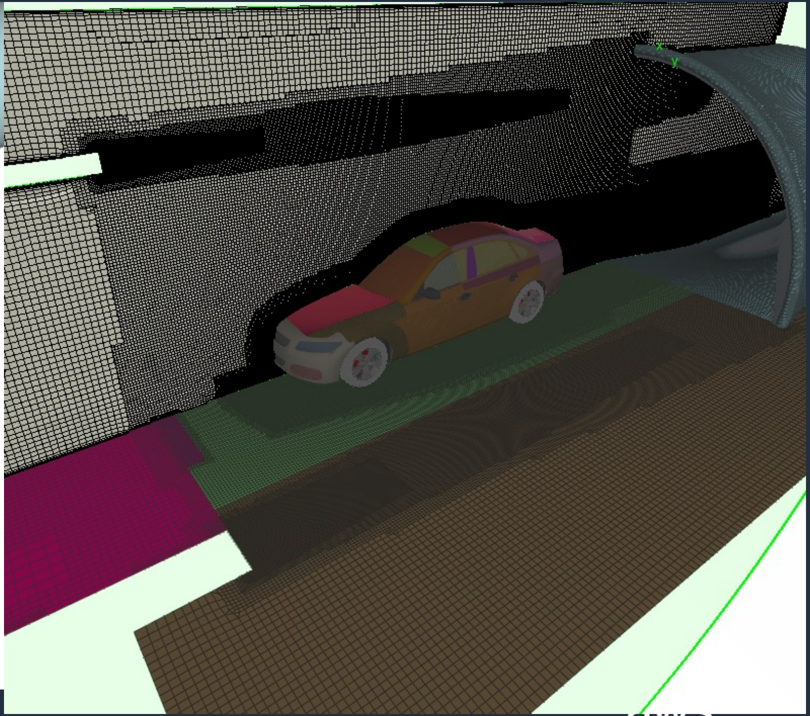

Case 2a – Free-air vs WT

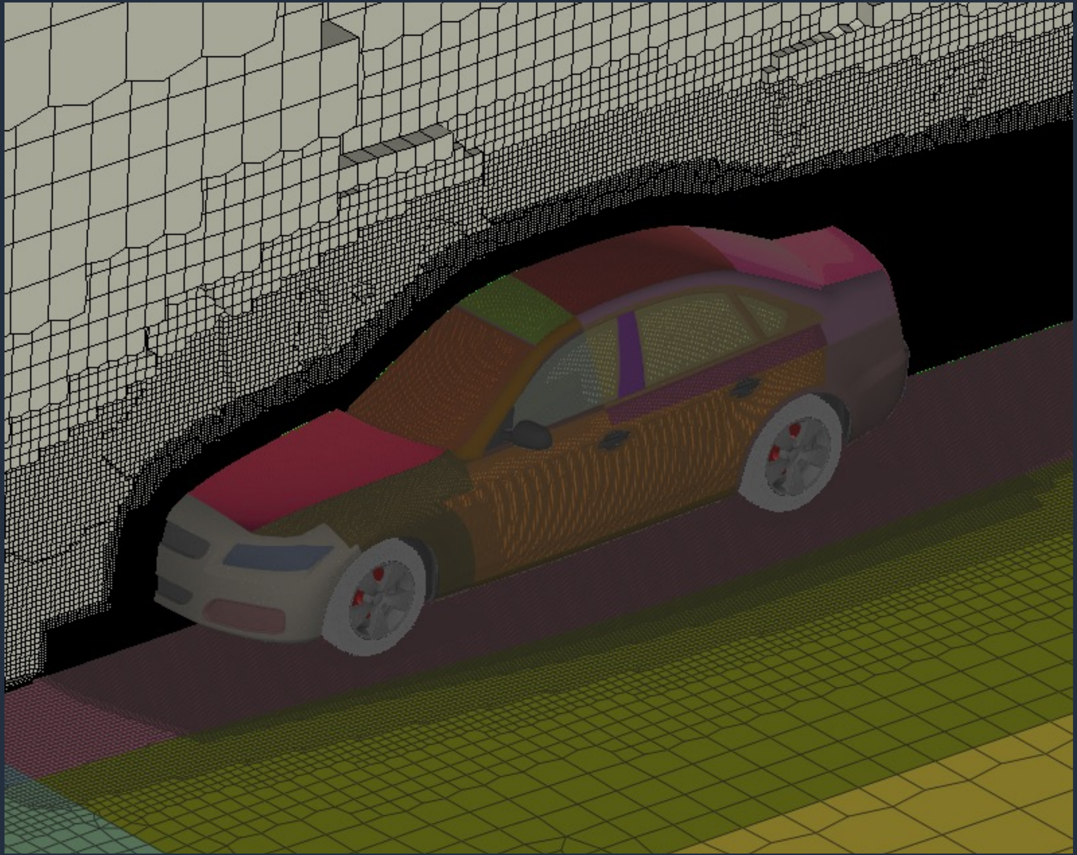

The grid has exactly the same resolution in the car region as the free-air case. The only difference being the tunnel itself.

Case 2a – Free-air vs WT

The grid has exactly the same resolution in the car region as the free-air case. The only difference being the tunnel itself.

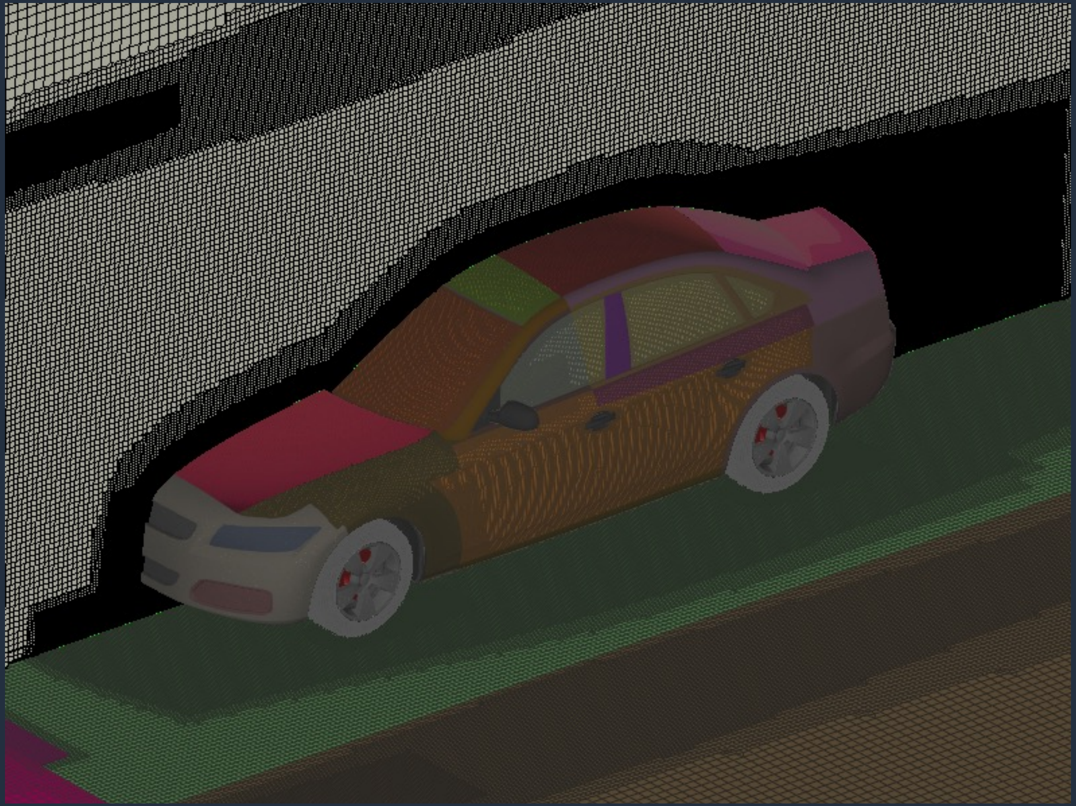

Case 2a – Free-air vs WT

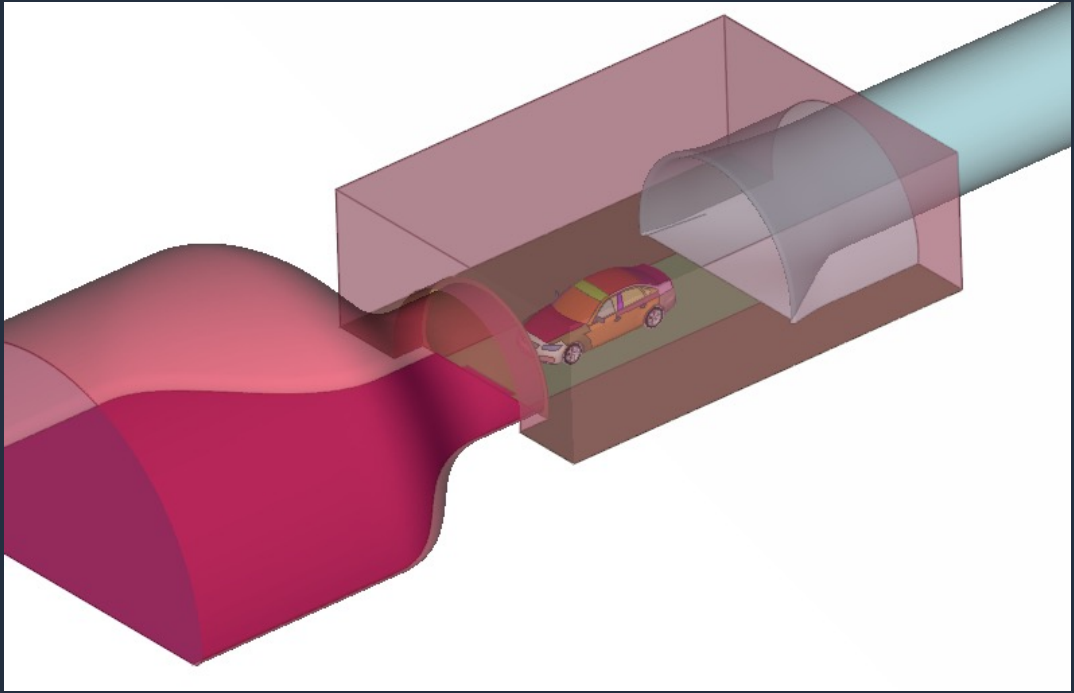

1) The floor ahead of the car is set to slip, as per the free-air case (even though we realize in the tunnel there are suction/blowing to control BL layer)
2) We first ran an empty WT with no car to set the inlet condition to achieve the desired flow speed in the test-section (38.88ms-1)
3) We use the plenum method to calculate a correction factor when the car is then placed into the tunnel.
4) We took 16 points around the plenum to measure the pressure but there is uncertainty in this approach so results should be treated as preliminary

Case 2a – Free-air vs WT

Clearly there is an influence of the WT and nozzle position. We are not modelling correctly but at least it's more realistic than free-air

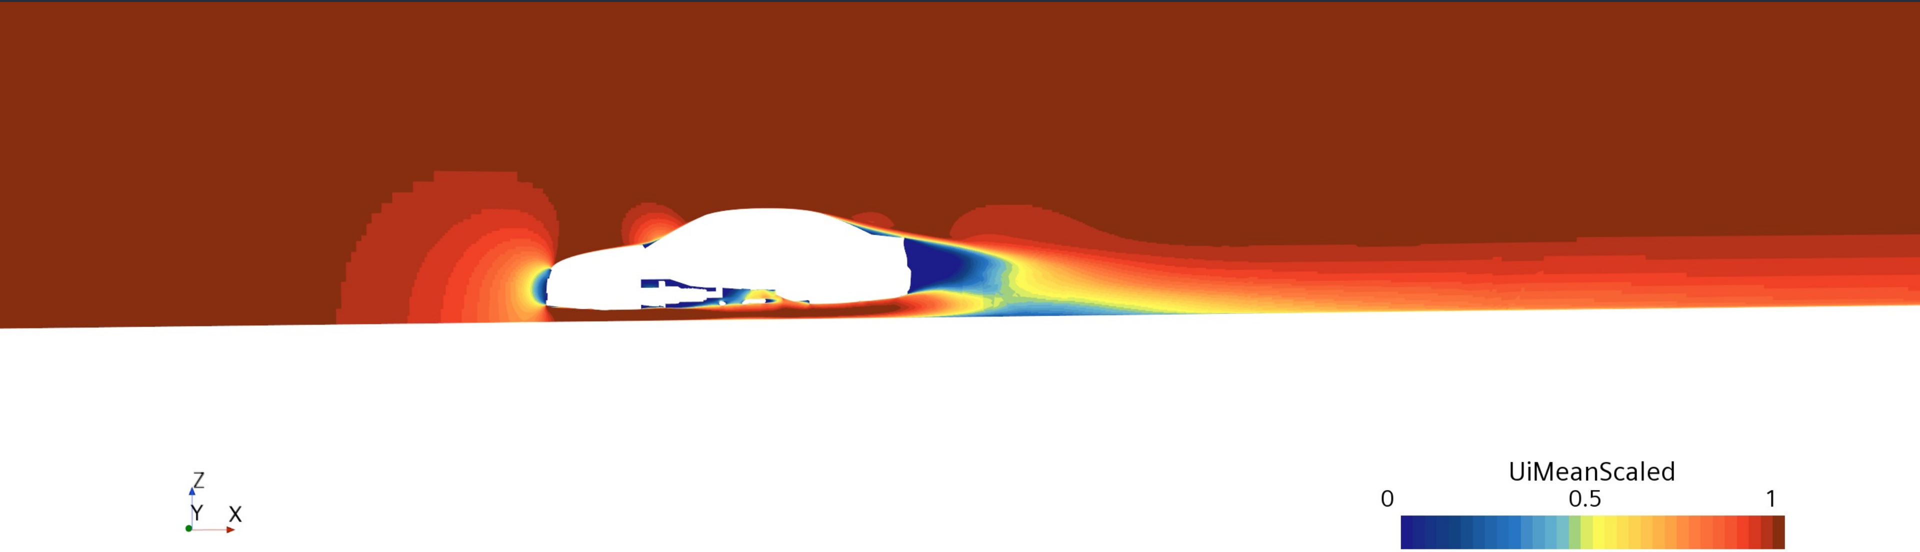

Case 2a – Free-air vs WT

Clearly there is an influence of the WT and nozzle position. We are not modelling correctly but at least it's more realistic than free-air

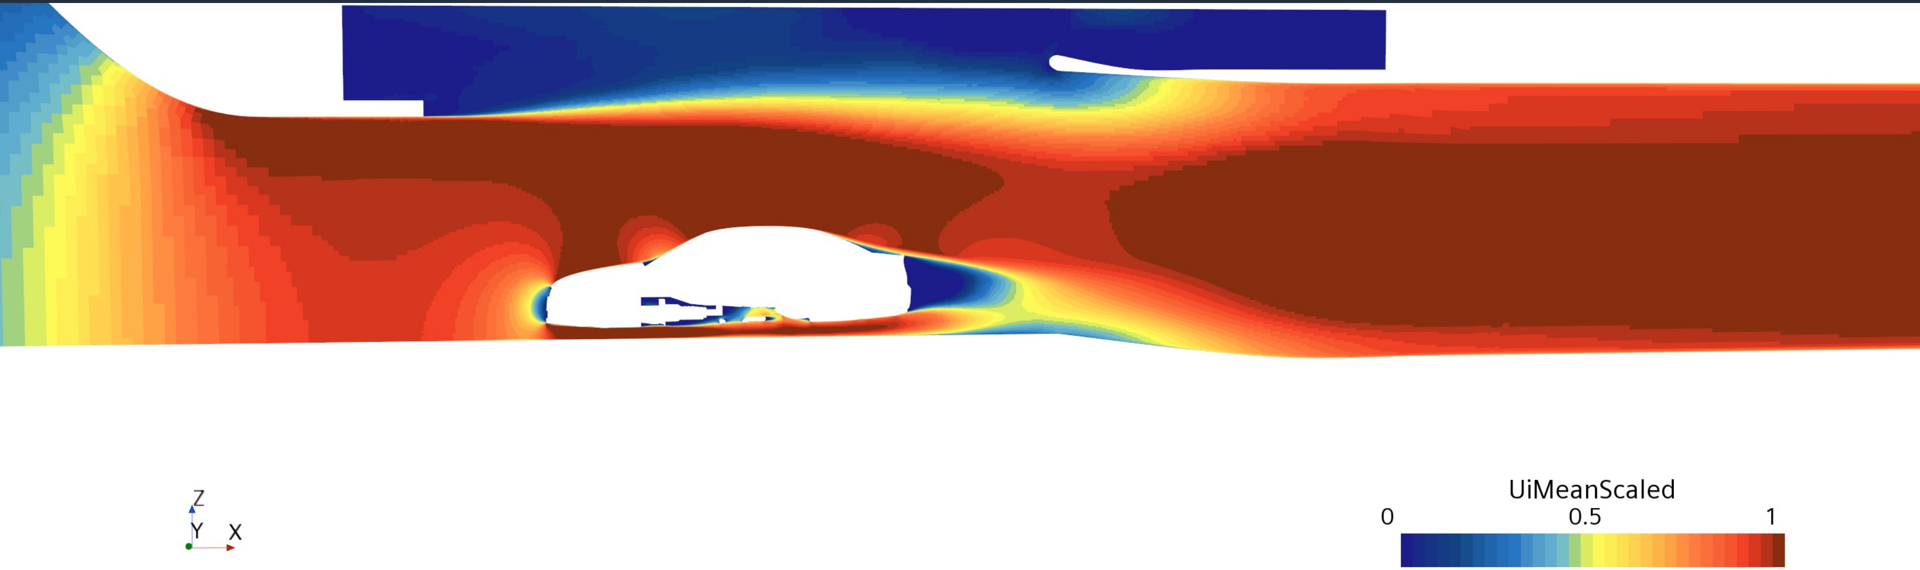

Case 2a – Free-air vs WT

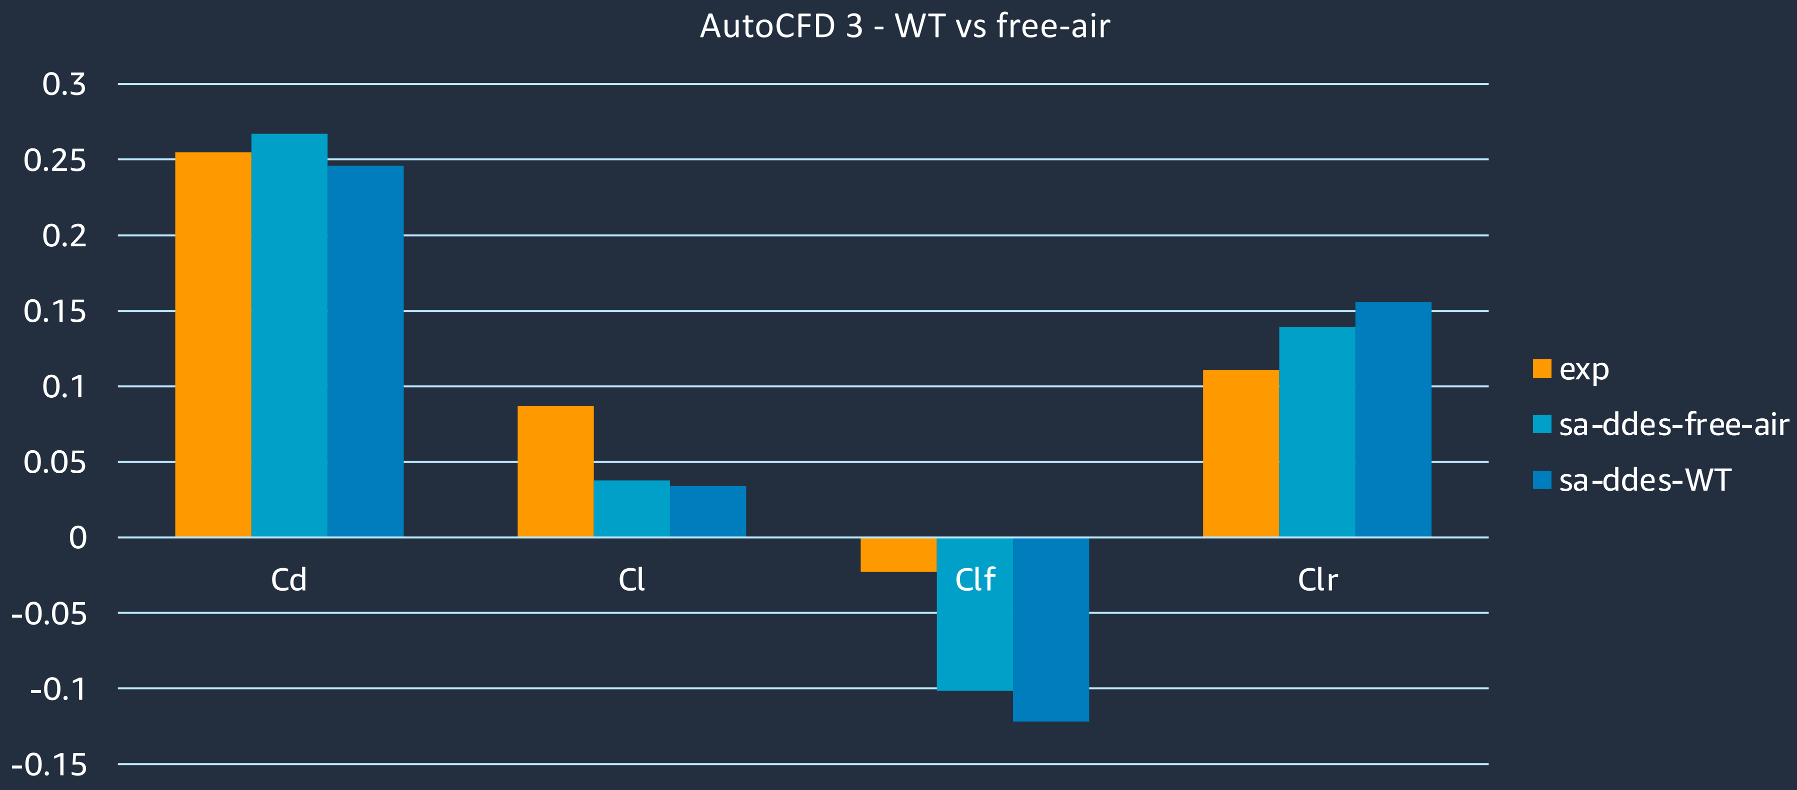

Case 2a – 128M – Free-air - SA-DDES

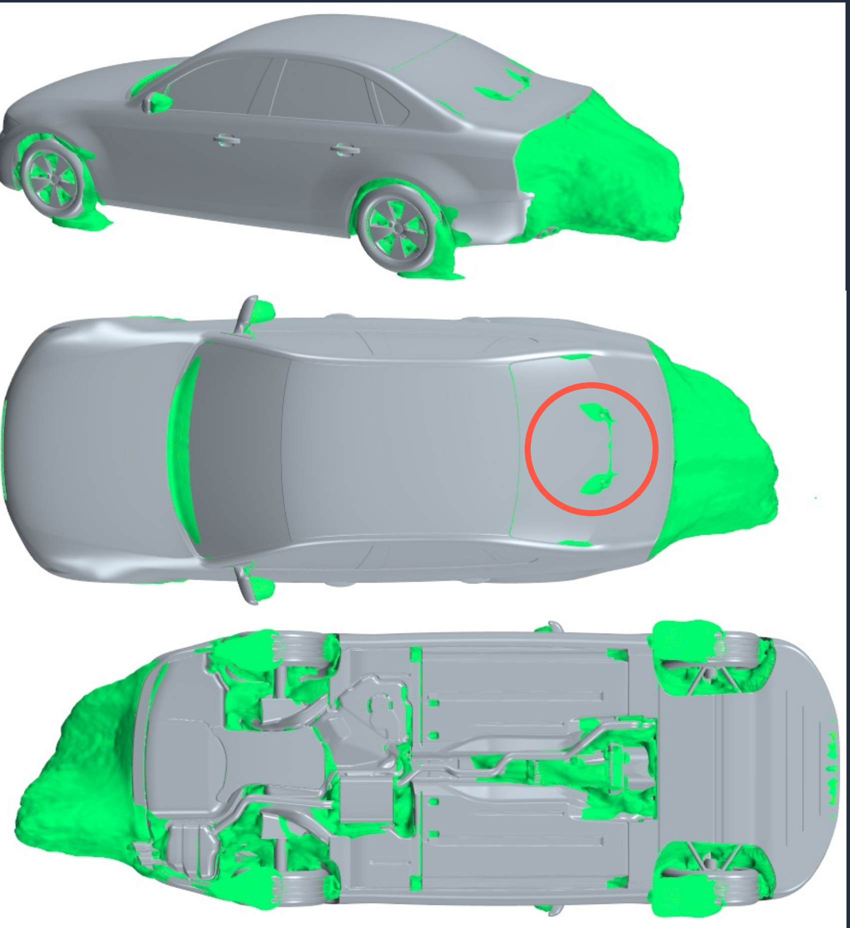

Front of the car sees actually very little difference. Main influence is on the rear separation

Case 2a – 157M – WT - SA-DDES

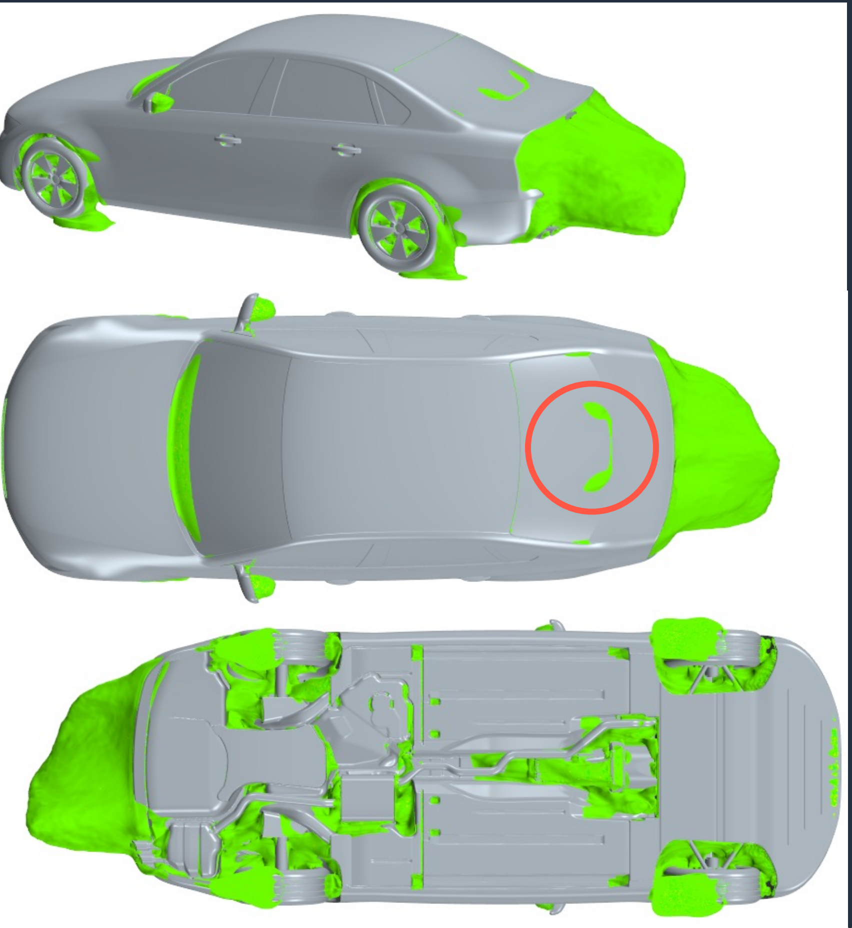

Front of the car sees actually very little difference. Main influence is on the rear separation

Case 2a – 157M – WT - SA-DDES

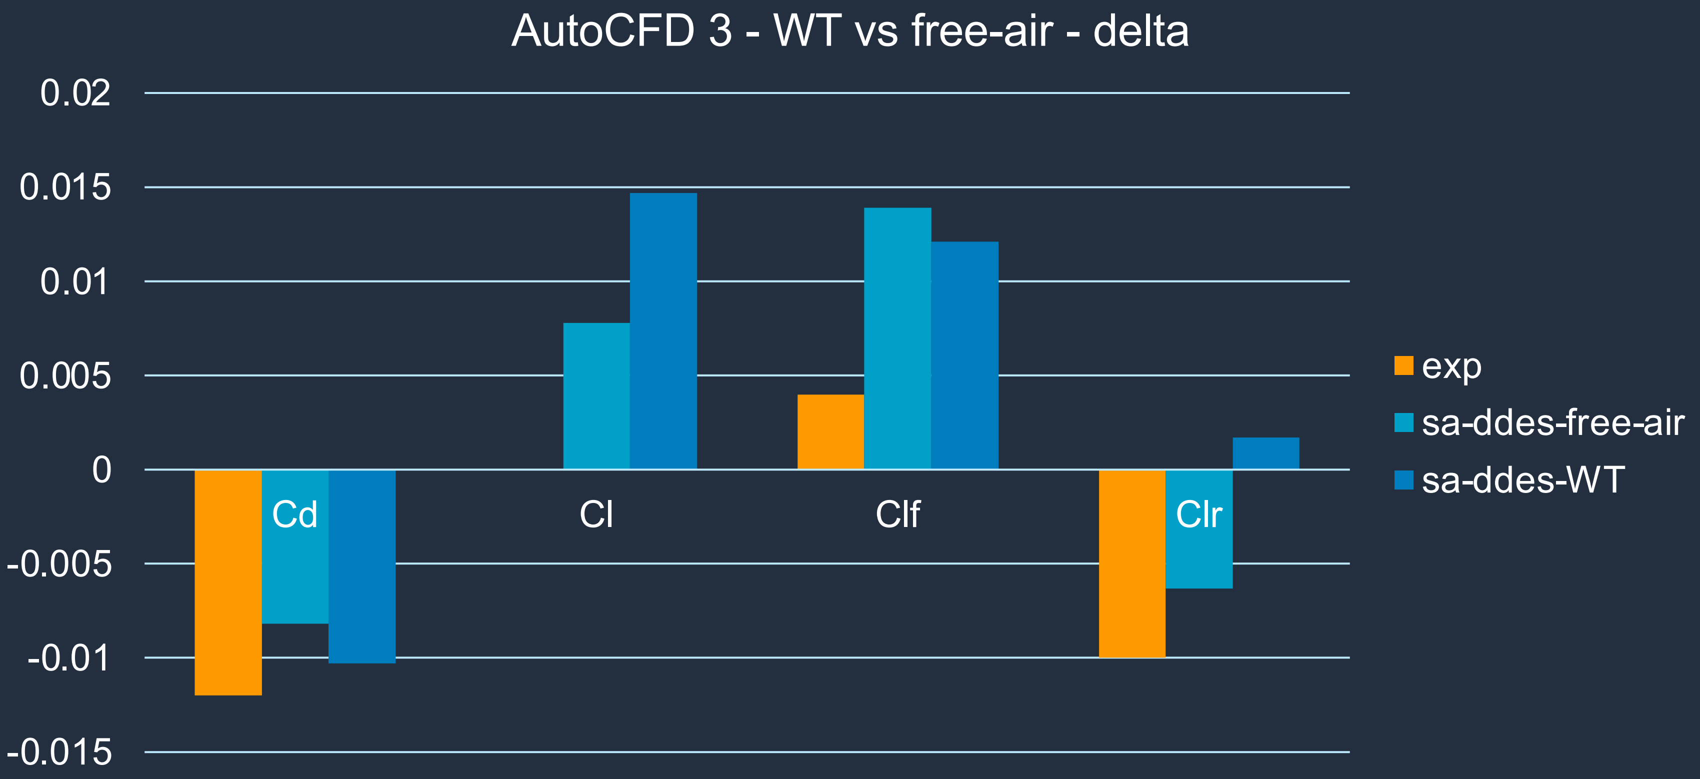

Case 2a – 157M – WT - SA-DDES

Even running for 12.5s we are still not time-averaged... Needs more investigation, so please treat this work as preliminary but keen to share to get feedback.

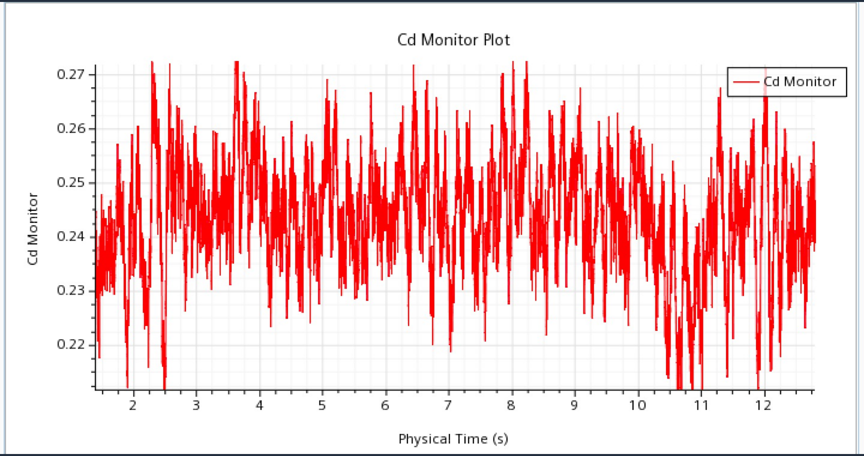

Summary

- Correlation is still a major challenge – lowering time-step and increasing mesh resolution to > 580M cell with still a CFL of 1 in key regions only has minor changes.

- Including the WT environment raises many questions that still need answering. The flow-field has changed but now the drag force is too low! The drag delta is however similar so not a magic bullet but needs much more investigation.

- All of these methods require significant HPC investment and we should continue to factor in the runtime and cost of the simulation. Further work will continue on testing next generation CPU and GPUs

- Keen to put more effort into WMLES methods to compare against DDES