High-Fidelity CFD on the Cloud Thoughts and Perspectives

NASA Ames Research Center - Advanced Modeling & Simulation (AMS) Seminar Series– Dec 17th, 2020

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

Motivation for CFD – Big Picture changes

Electric

Efficiency

Range

Option 1 – Road/Flight tests

Build a car

Expensive

Accurate

Option 1 – Road/Flight tests

Build a plane

Expensive

Accurate

Option 2 – Wind Tunnel tests

Mock up

Expensive

Limited

Option 2 – Wind Tunnel tests

Re number

Expensive

Limited

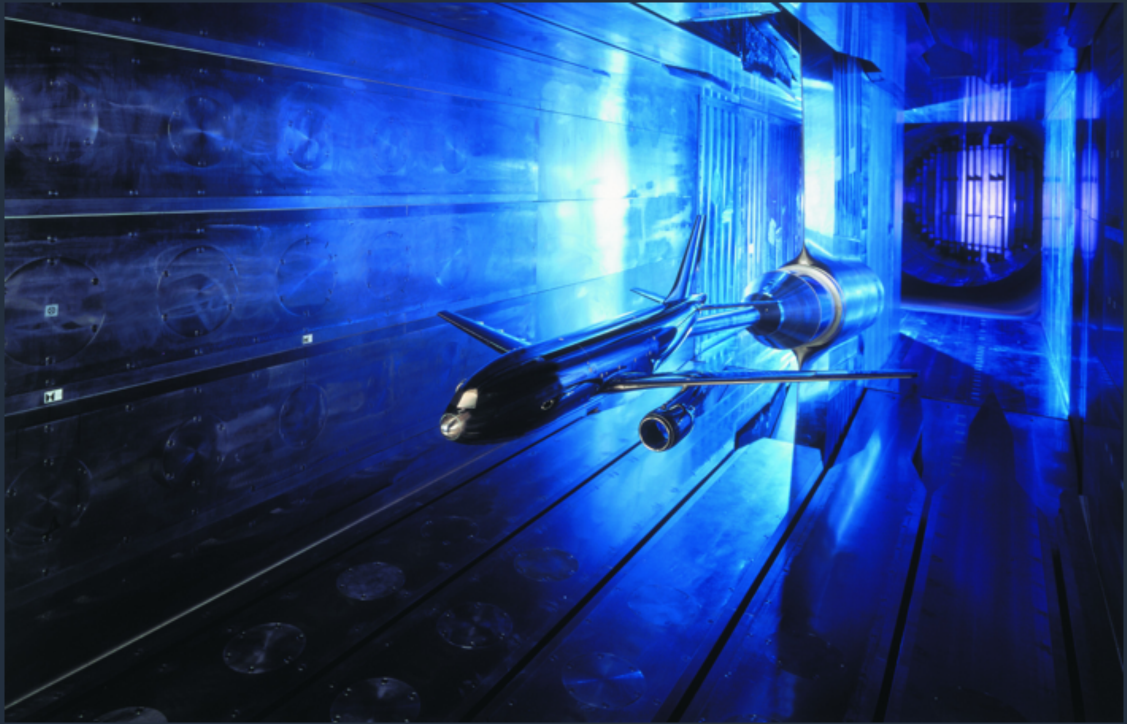

Option 2 – Wind Tunnel tests

Mock up

Expensive

Re number

Option 2 – Wind Tunnel tests

Re number

Expensive

Limited

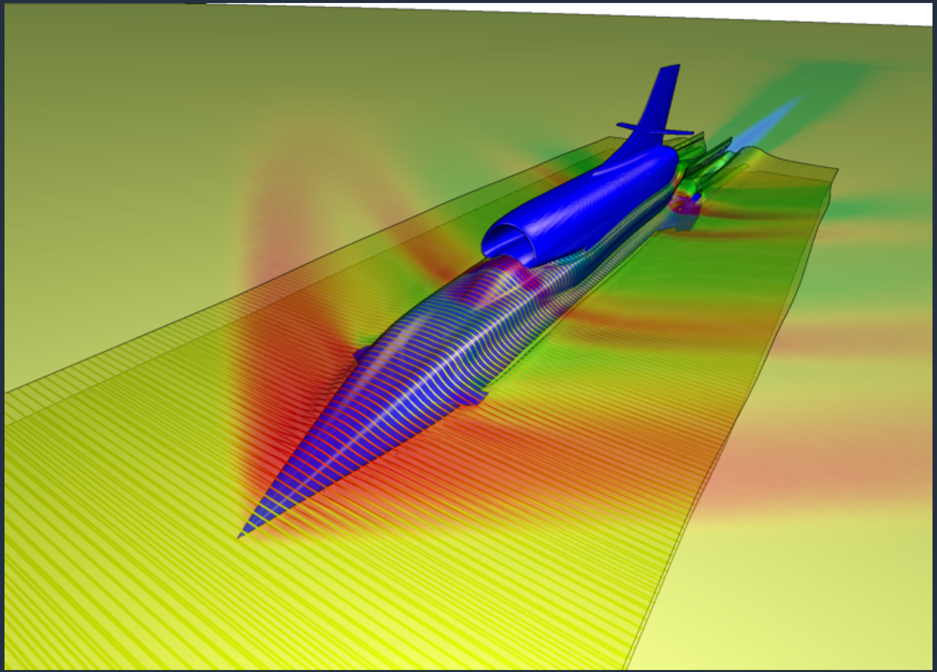

Prof. Ben Evans, Swansea University

Option 3 – Computational Fluid Dynamics

No Manufacture

Cheap & Fast*

Accurate?

$$\frac{\partial u_i}{\partial x_i} = 0$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho}\frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ij}}{\partial x_j}$$

Turbulence Modelling

DNS

Direct Numerical Simulation

- Best representation of the full Navier-Stokes; no modelling required
- All scales of motion are captured
- Extremely expensive; currently just a research tool

N_C=10B

Δt=1e^{-7}

N_T=40M (4s)

N_p=200k (92 days)*

~$13M**

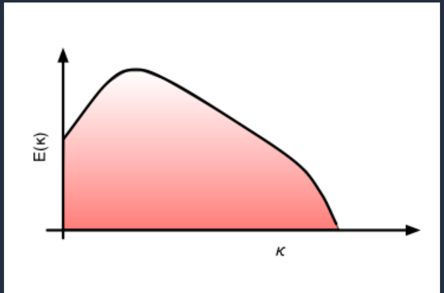

**0.2s per time-step (50k cells per core)

**$0.03 per core hour

Turbulence Modelling

LES

Large Eddy Simulation

- Resolve large scale but model small scales
- Explicit sub-grid scales or implicit through numerical schemes
- Remains expensive + mesh generation

N_C=2B

Δt=1e^{-6}

N_T=4M (4s)

N_p=40k (12 days)*

~$345k**

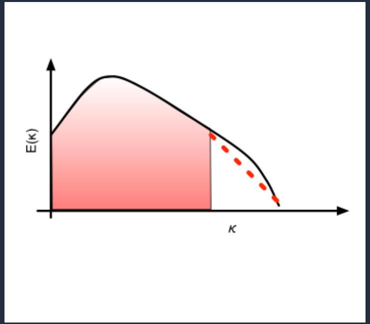

**0.25s per time-step
**$0.03 per core hour

Turbulence Modelling

RANS

Reynolds Averaged Navier-Stokes

- Required extensive closure modelling
- No information about instantaneous flow
- Cheap/fast for engineering applications

N_C=250M

Δt=n/a

N_T=4k

N_p=2.5k (2 hours)*

~$150**

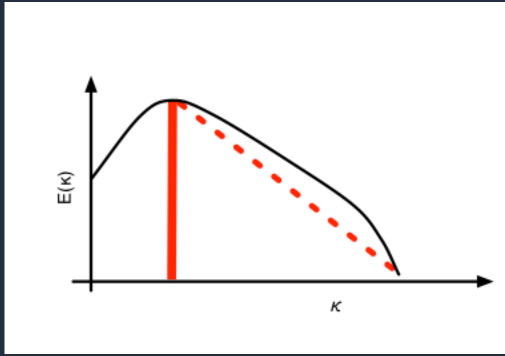

**2s per iteration (100k cells per core)

**$0.03 per core hour

Turbulence Modelling

Hybrid

Hybrid RANS-LES / WMLES

- RANS/WM in the near-wall regions, LES further away
- Can capture instantaneous flow
- 'Compromise' for industry

N_C=1B

Δt=1e^{-5}

N_T=400k

N_p=20k (108 hours)

~$64k*

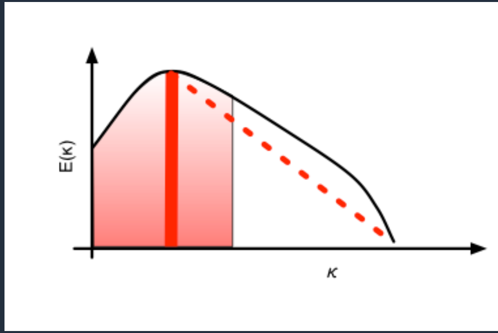

**1s per time-step (50k cells per core) - implicit
**$0.03 per core hour

Industry costs

Time-step

Mesh

HPC

$64k is still too expensive for industry compared to $150 for RANS (think about wind-tunnel costs)

- Larger time-step (1e^{-5} to 1e^{-4} = $6.4k*)
- Mesh from 1B to 250M = $1600*
- However we're assuming 50k cells per core for 1s e.g 5000 cores for 11 hours*
- 500 cores gives 4 days *.. (assume linear..)
- Commercial applications typically scales to 100k cells per core so assume 2,500 cores to give 24 hrs
- HPC bottleneck begins… cloud starts to help.

*depends on scaling/test-case

RANS Modelling

$$\frac{\partial U_i}{\partial t} + U_j \frac{\partial U_i}{\partial x_j} = -\frac{1}{\rho}\frac{\partial P}{\partial x_i} + \frac{1}{\rho}\frac{\tau_{ij}}{\partial x_j} - \frac{\partial \overline{u_i^{'} u_j^{'}}}{\partial x_j}$$

RANS Modelling

$$\frac{\partial U_i}{\partial t} + U_j \frac{\partial U_i}{\partial x_j} = -\frac{1}{\rho}\frac{\partial P}{\partial x_i} + \frac{1}{\rho}\frac{\tau_{ij}}{\partial x_j} - \frac{\partial \overline{u_i^{'} u_j^{'}}}{\partial x_j}$$

$$\frac{D\overline{u_i^{'} u_j^{'}}}{Dt} = P_{ij} + \phi_{ij} - \varepsilon_{ij} + D_{ij}$$

$$\overline{u_i^{'} u_j^{'}} = -2\nu_t S_{ij} + \frac{2}{3} k \delta_{ij}$$

$$\nu_t = L^2 T^{-1} = U^2 . T = U.L$$

2D wall-mounted hump

- Challenging flow separation case
- Re=936,000
- 3.2 million cells
- Range of RANS models
- No RANS model can predict correctly. Under-predicts turbulent shear-stress. (RSM's too)
- Please try it yourself! (*https://turbmodels.larc.nasa.gov/nasahump_val.html)*

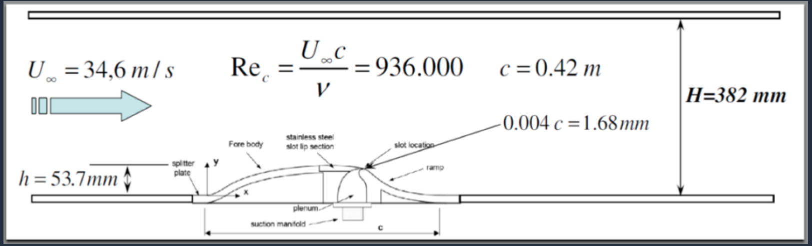

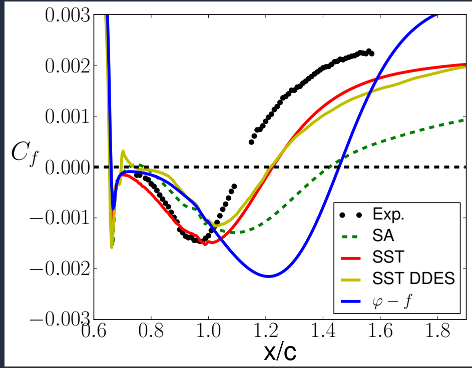

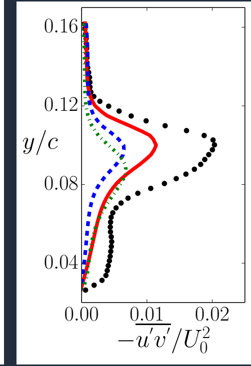

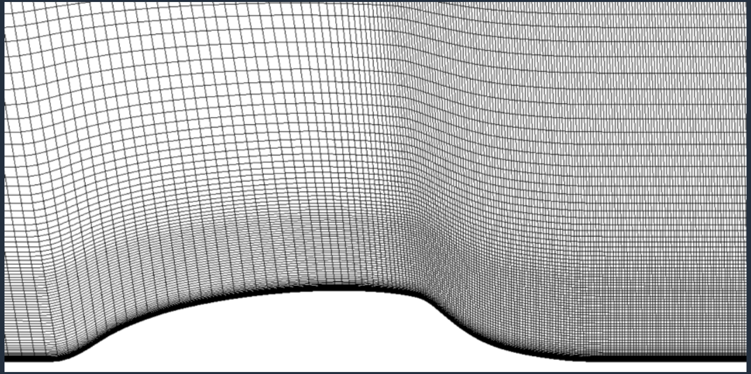

Ashton, N. (2012). *Implementation and testing of an alternative DDES formulation based on elliptic relaxation*. University of Manchester.

2D periodic hills

- Difficult case whilst looking easy
- Re = 37,000 based on hill height
- Periodic flow
- Unsteady separation and reattachment points
- Similar under-prediction of turbulent shear-stress by all models (RSM's don't help)

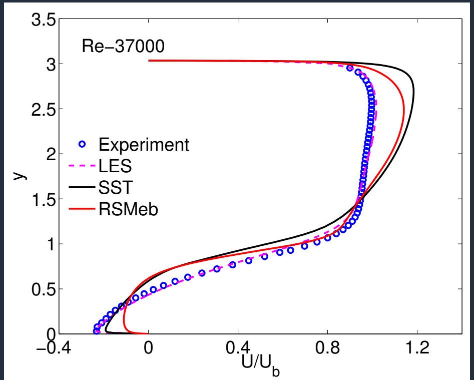

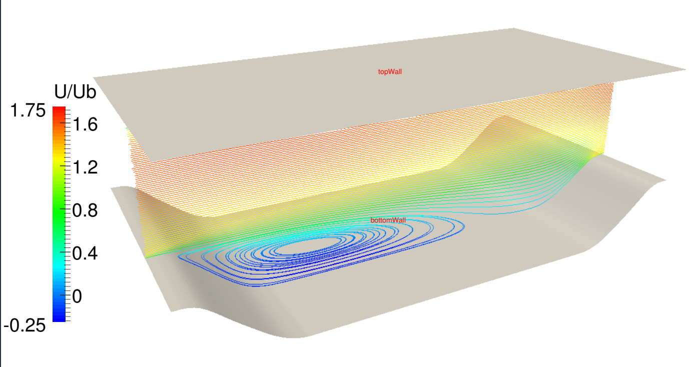

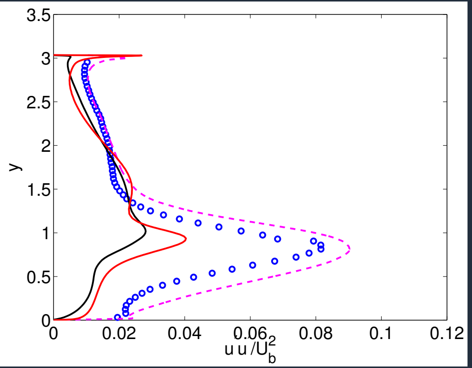

Stoellinger, M., Roy, R., & Ashton, N. (2015). Application of an elliptic blending Reynolds stress model in attached and separated flows. *22nd AIAA Computational Fluid Dynamics Conference.*

Ahmed car body (25 deg)

- Popular `simple’ car shape
- Re ~ 1 million based upon height
- 16 million cells
- 25 and 35 rear angle give different physics
- Unstructured/structured with range of RANS models

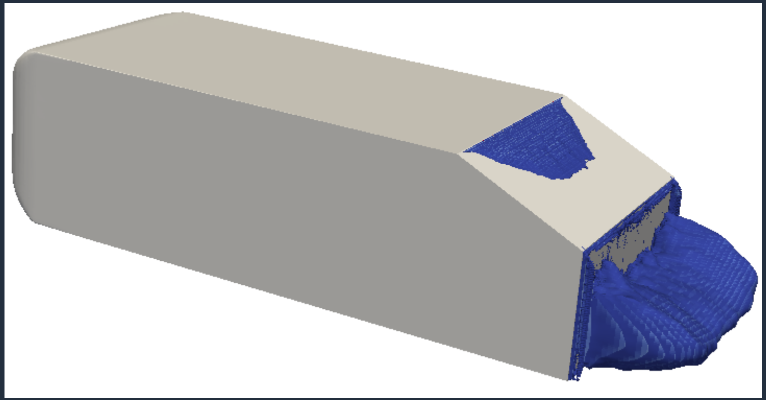

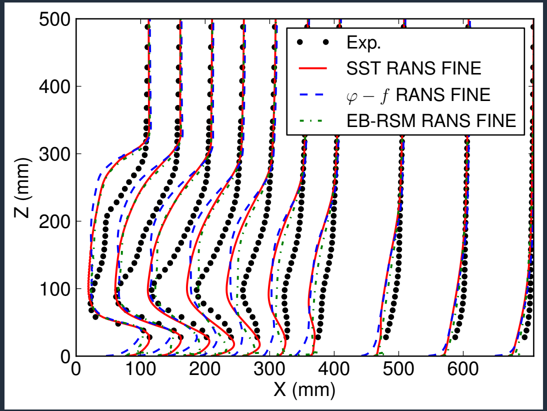

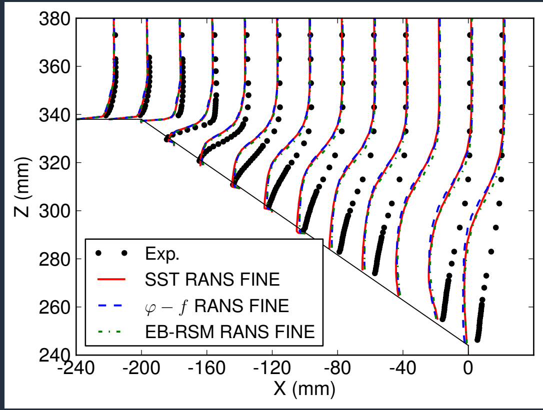

Ashton, N., & Revell, A. (2015). Key factors in the use of DDES for the flow around a simplified car. *International Journal of Heat and Fluid Flow*, *54*, 236–249. https://doi.org/10.1016/j.ijheatfluidflow.2015.06.002

Ahmed car body (25 deg)

At 25 degrees no RANS/URANS model (whether EVM/RSM, or ε/ω) can correctly prediction the velocity field.

All models under-predict the level of turbulence in the initial separated shear layer and as a result over-prediction the separation region. Same as 2D hump, 2D periodic hills

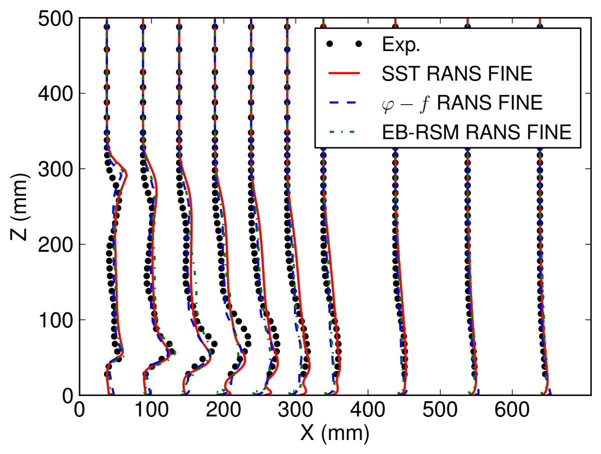

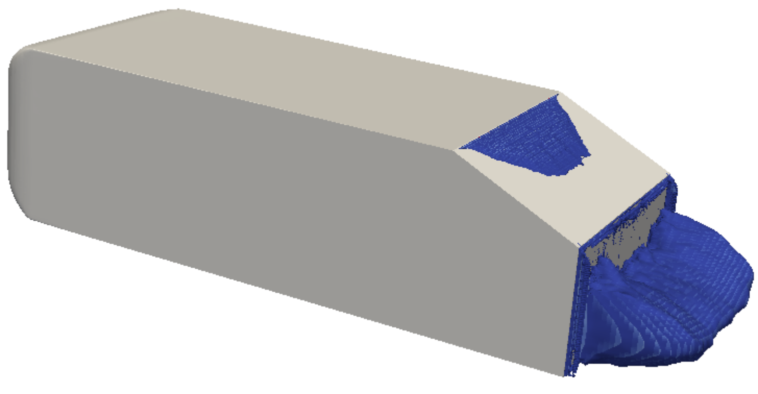

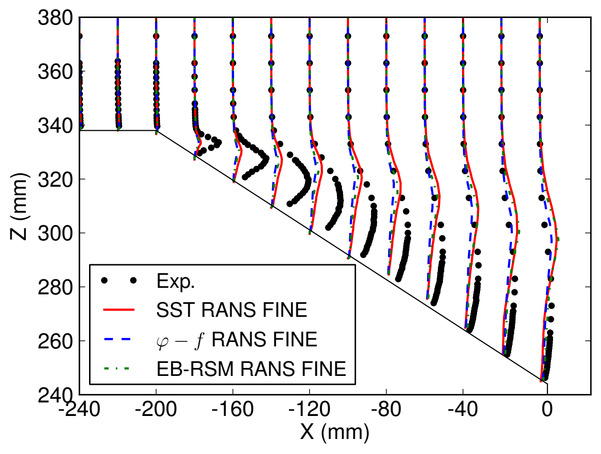

DrivAer model

- Popular open-source vehicle model
- Re=~1M, Grids up to 256M cells
- Fastback shows good RANS performance but estate with larger separation shows poor performance
- Recent Automotive CFD Prediction Workshop focused on this test-case (autocfd.eng.ox.ac.uk)

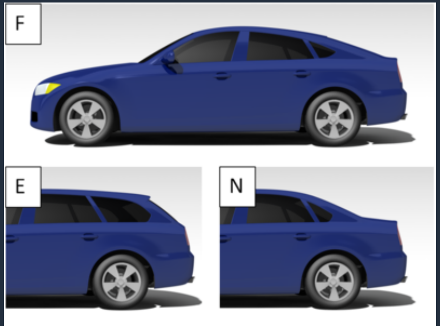

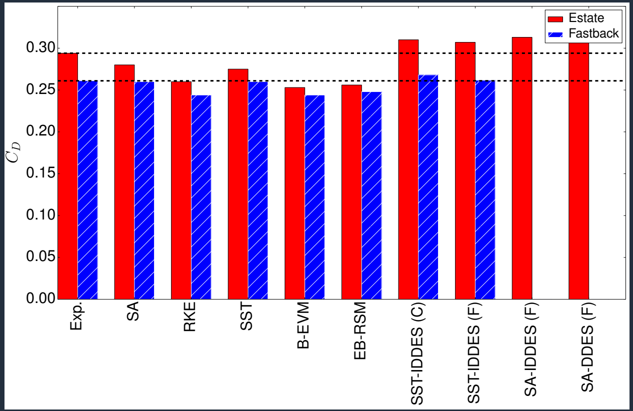

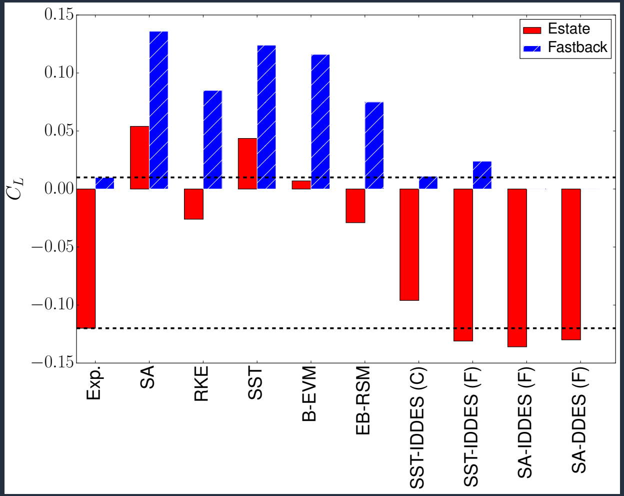

N. Ashton, A. West, S. Lardeau, A. Revell, *Assessment of RANS and DES methods for realistic automotive models*, Computers. Fluids. 128 (2016) 1–15. doi:10.1016/j.compfluid.2016.01.008.

Industrial automotive example

- Audi A1/Q3/A4 in OpenFOAM (full details in SAE paper)
- Real-life production example
- pimpleFOAM + ~100M snappyHexMesh grids
- RANS not good enough at capturing rear lift – similar to DriVaer

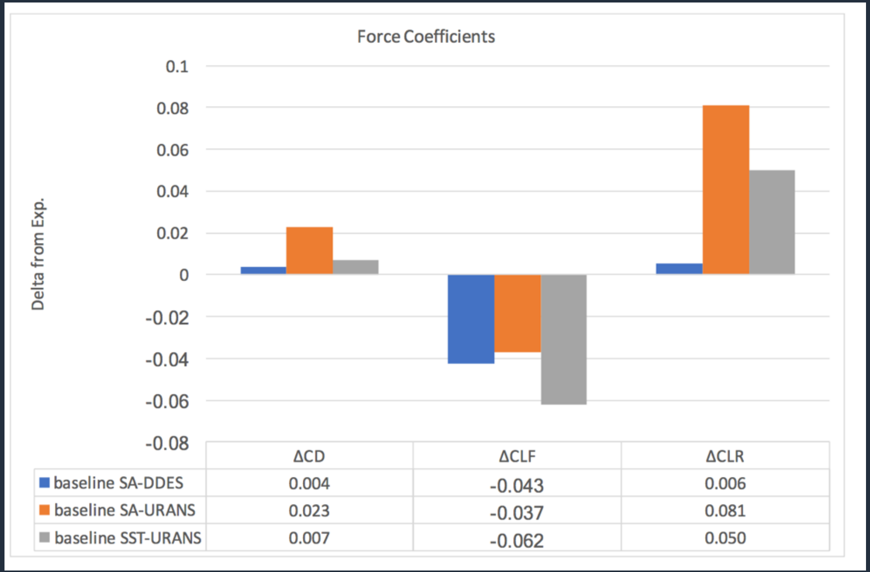

Ashton, N., Unterlechner, P., & Blacha, T. (2018). Assessing the Sensitivity of Hybrid RANS-LES Simulations to Mesh Resolution, Numerical Schemes and Turbulence Modelling within an Industrial CFD Process. *SAE Technical Papers*, *2018-April*. https://doi.org/10.4271/2018-01-0709

Aerospace

- 3rd AIAA High-Lift Prediction Workshop
- Aim to compared OpenFOAM to STAR-CCM+ and test RANS accuracy
- rhoPimpleFoam + local-time-stepping
- Up to 256M cells using ANSA generated grids
- Re=1.93M (M=0.172), AoA=0-21.57
- Spalart-Allmaras RANS model used

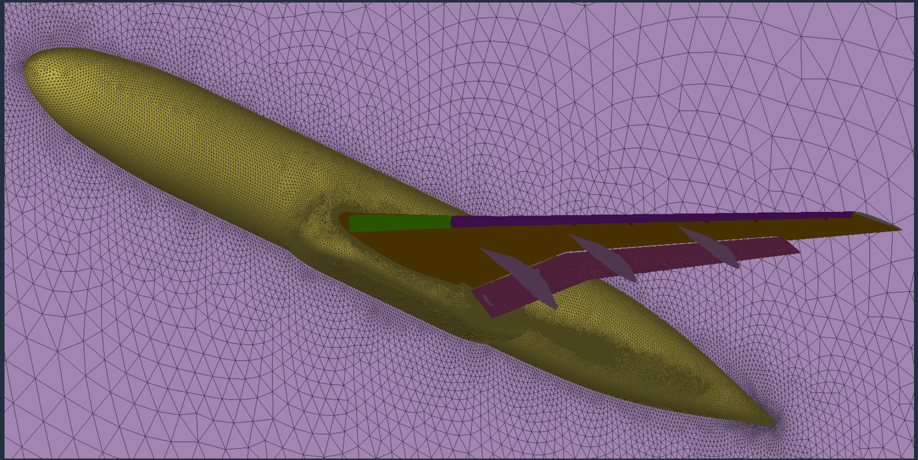

Ashton, N., & Skaperdas, V. (2017). Verification and Validation of OpenFOAM for High-Lift Aircraft Flows. *Journal of Aircraft*, 1–30. https://doi.org/10.2514/1.C034918

Aerospace

Excellent agreement between OpenFOAM and STAR-CCM+ - within 2% for all cases studied (see paper below)

Agreement looks good to experiment but this is a great case of error cancellation

Wrong separation inboard and outboard cancel out!

Most likely need high-fidelity methods for stall conditions

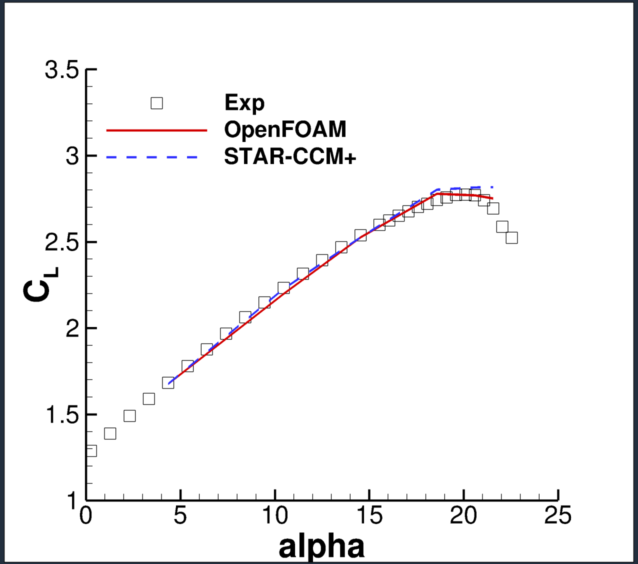

Ashton, N., & Skaperdas, V. (2017). Verification and Validation of OpenFOAM for High-Lift Aircraft Flows. *Journal of Aircraft*, 1–30. https://doi.org/10.2514/1.C034918

Aerospace

Excellent agreement between OpenFOAM and STAR-CCM+ - within 2% for all cases studied (see paper below)

Agreement looks good to experiment but this is a great case of error cancellation

Wrong separation inboard and outboard cancel out!

Most likely need high-fidelity methods for stall conditions (no other participant could predict correctly)

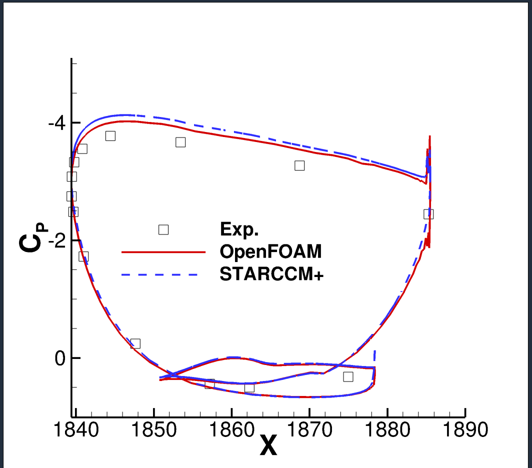

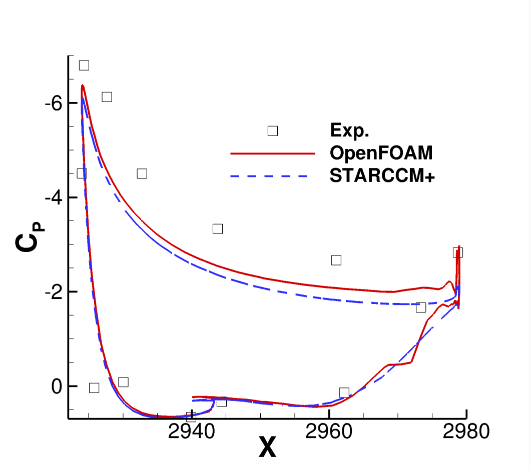

Ashton, N., & Skaperdas, V. (2017). Verification and Validation of OpenFOAM for High-Lift Aircraft Flows. *Journal of Aircraft*, 1–30. https://doi.org/10.2514/1.C034918

Aerospace

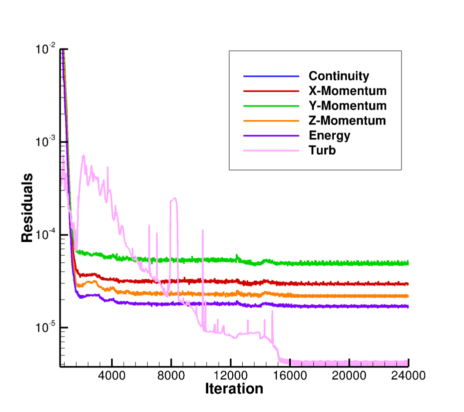

Stability of OpenFOAM was a real problem. Local-time stepping using rhoPimpleFOAM

We really need a high performing coupled density based solver to match the majority of aerospace codes.

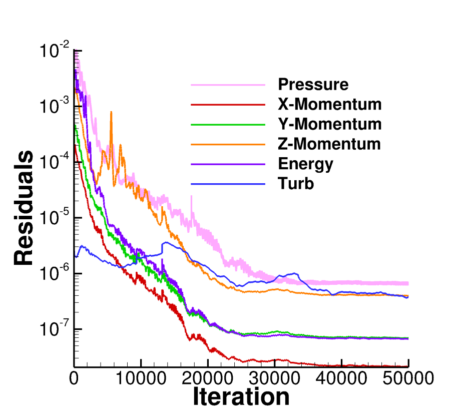

Ashton, N., & Skaperdas, V. (2017). Verification and Validation of OpenFOAM for High-Lift Aircraft Flows. *Journal of Aircraft*, 1–30. https://doi.org/10.2514/1.C034918

Motivation to go beyond RANS

RANS (at both eddy-viscosity and second moment closure level), fails in certain classes of flows.

- Commercial aircraft close to and beyond stall (e.g >18 AOA)
- Vehicles at certain rear window angles (estate cars)
- Noise reduction on aircraft (landing gear, high-lift systems)
- Noise reduction on cars (mirrors, tyre noise)
- Many other examples

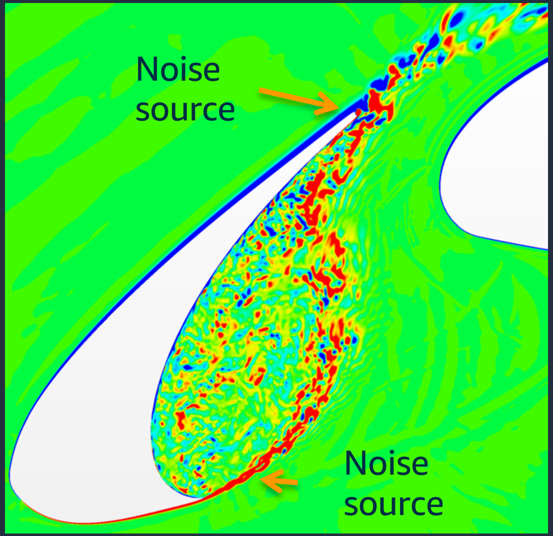

2D Periodic hills

DDES captures the shear-stress well and as a result matches the recirculation region

Only 2M cell grid compared to >20M for WRLES

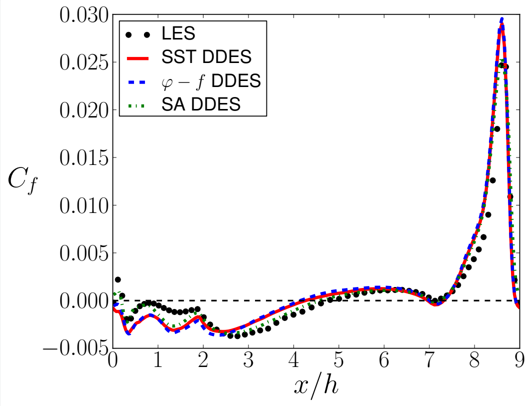

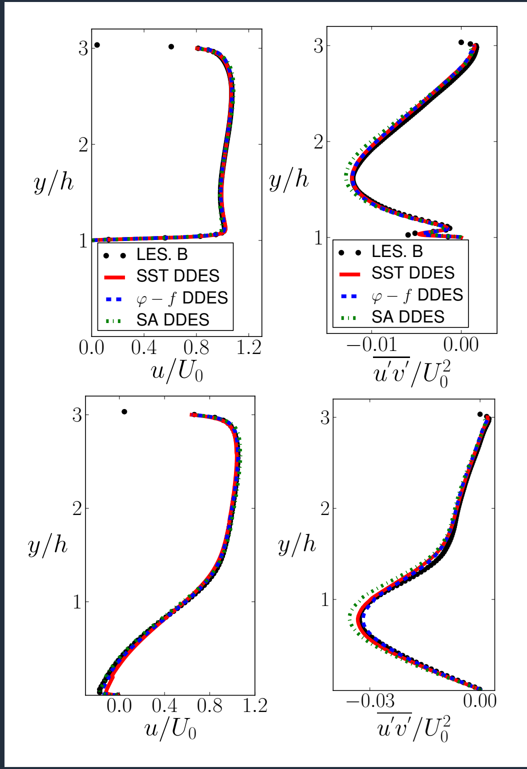

Ashton, N., Revell, A., Prosser, R., & Uribe, J. (2013). Development of an Alternative Delayed Detached-Eddy Simulation Formulation Based on Elliptic Relaxation. *AIAA Journal*, *51*(2), 513–519.

2D Hump

DDES (as well as other hybrid RANS-LES method) perform well for this case, although still not perfectly.
Grey-area problem exists (see Mockett new DDES variants)

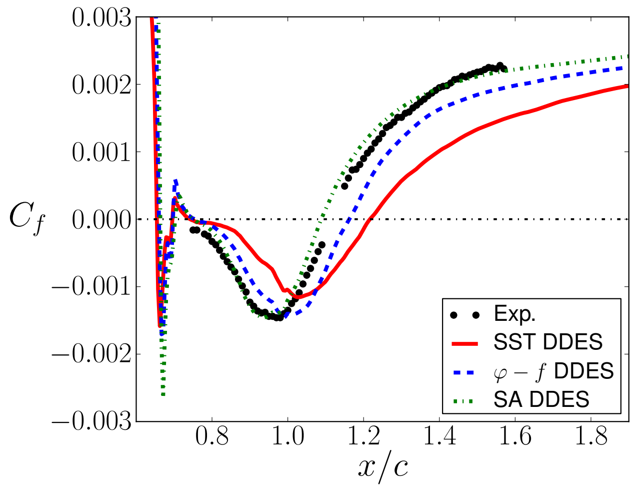

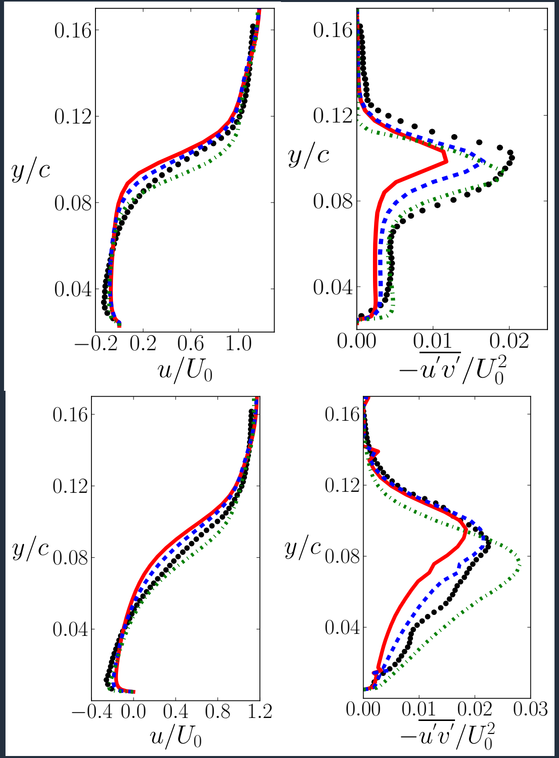

Ashton, N., Revell, A., Prosser, R., & Uribe, J. (2013). Development of an Alternative Delayed Detached-Eddy Simulation Formulation Based on Elliptic Relaxation. *AIAA Journal*, *51*(2), 513–519.

Ahmed car body

Considerable improvement over RANS but sensitive to the near-wall grid and filter width (See paper below)

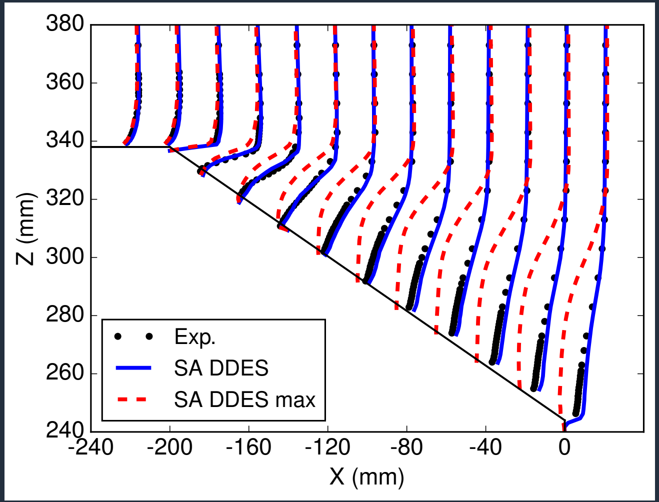

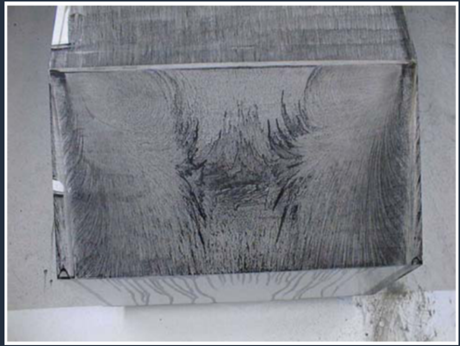

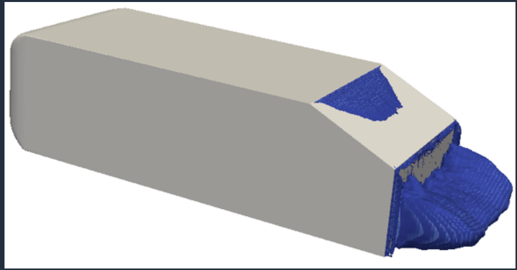

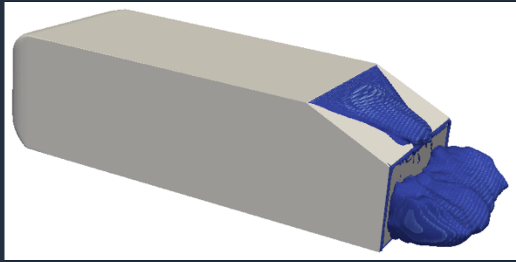

Ashton, N. (2017). Recalibrating Delayed Detached-Eddy Simulation to eliminate modelled-stress depletion. *23rd AIAA Computational Fluid Dynamics Conference*. https://doi.org/10.2514/6.2017-4281

DrivAer

Clear trend for DDES like methods to offer an improvement over RANS

Biggest difference for lift on estate i.e massive separation

Still sensitivity to time-step, mesh, underlying model etc

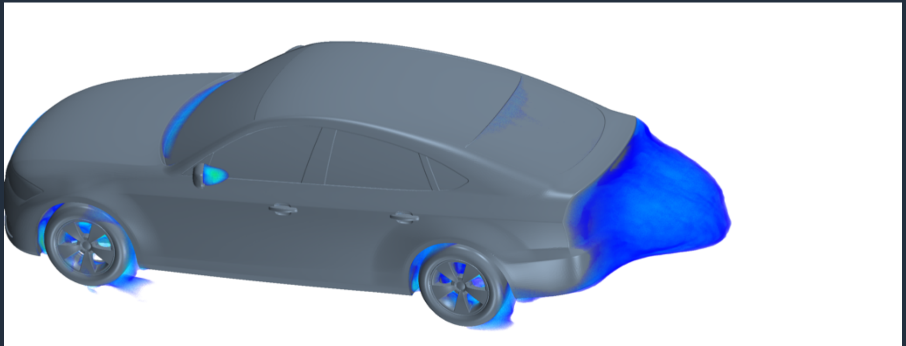

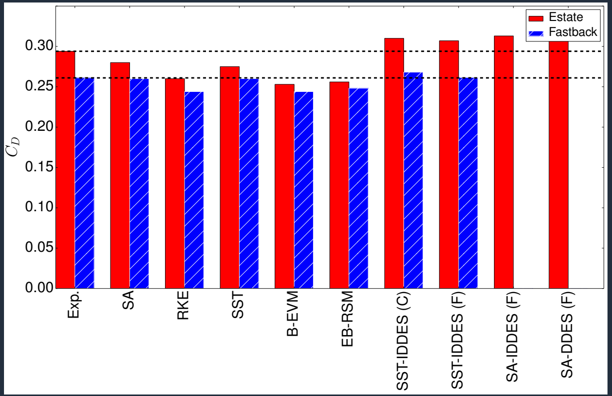

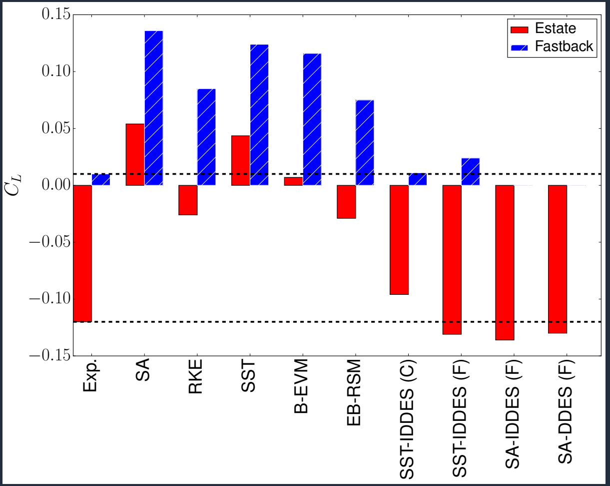

DrivAer

1st Automotive CFD Prediction Workshop (autocfd.eng.ox.ac.uk)

Range of inputs from Academic, Industry for automotive relevant test-cases

One key finding was the sensitivity of the time-step.

Tempting to increase time-step to reduce costs but clear trend away from experiment.

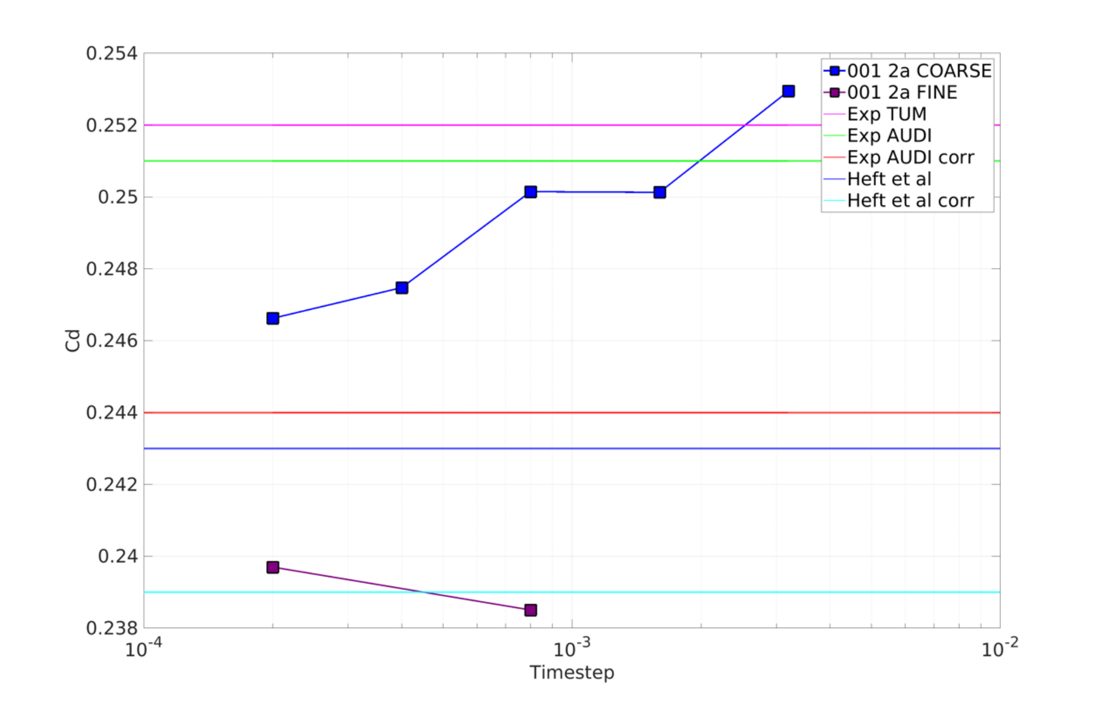

Additional considerations

"Grey-area"– Need methods to minimize this problem. Proper grid resolution and numerics help, but so does synthetic turbulence. Solved by WMLES? E.g Larson et al. approaches?

RANS/LES selection– Where is RANS and LES selected. The f_d function of DDES can be too strong, but can also fail is the boundary layer is too isotropic. True also for WMLES (y+ enough?)

Mesh resolution– Automatic mesh refinement to ensure cells are only there where they are needed. Needs more testing + maybe immersed boundary method?

Numerical schemes– Need to ensure as low dissipation as possible in LES, but need robust schemes for complex geometries and RANS. High-order?

Industry costs

Time-step

Mesh

HPC

- Clear need for high-fidelity methods but further improvements needed e.g WMLES.
- Compromises are still needed to reduce costs e.g time-step + mesh. Need faster solvers + hardware.
- To give <24hr turn-around time we need access to > 1000 cores per job.
- HPC bottleneck begins…
- But you need flexibility, sometimes you only do DES once a week..

Classic HPC issues

I'm stuck in the queue → Can you bump up my priority?

Why is UserX always hogging the queue?

I've run out of disk space → Can I have 1TB more, I promise to delete some files…

Cloud HPC Motivation

- Most companies will phase in the use of high-fidelity CFD alongside their standard RANS processes
- They need to burst to much higher cores for HRLES/WMLES runs which may not be possible on-prem
- Two main camps – all in on cloud or hybrid bursting mode for spikes

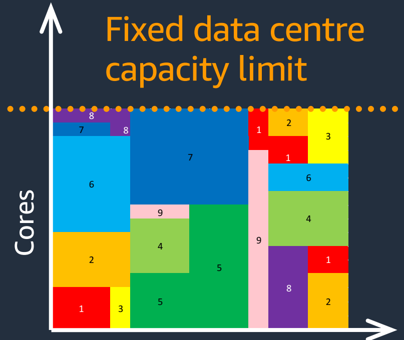

Finite capacity, usually with long queues to wait in

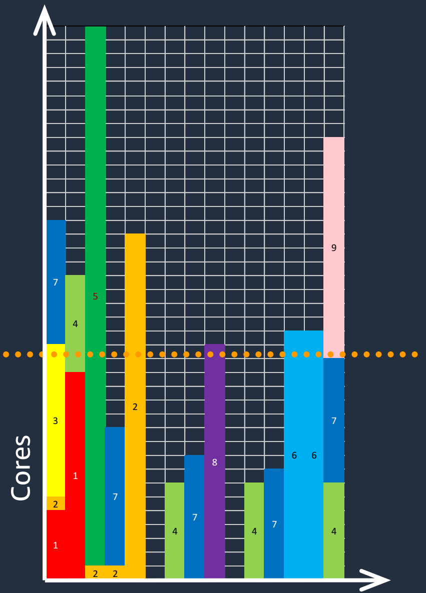

Massive capacity when needed to speed up time to results

AWS Global Infrastructure

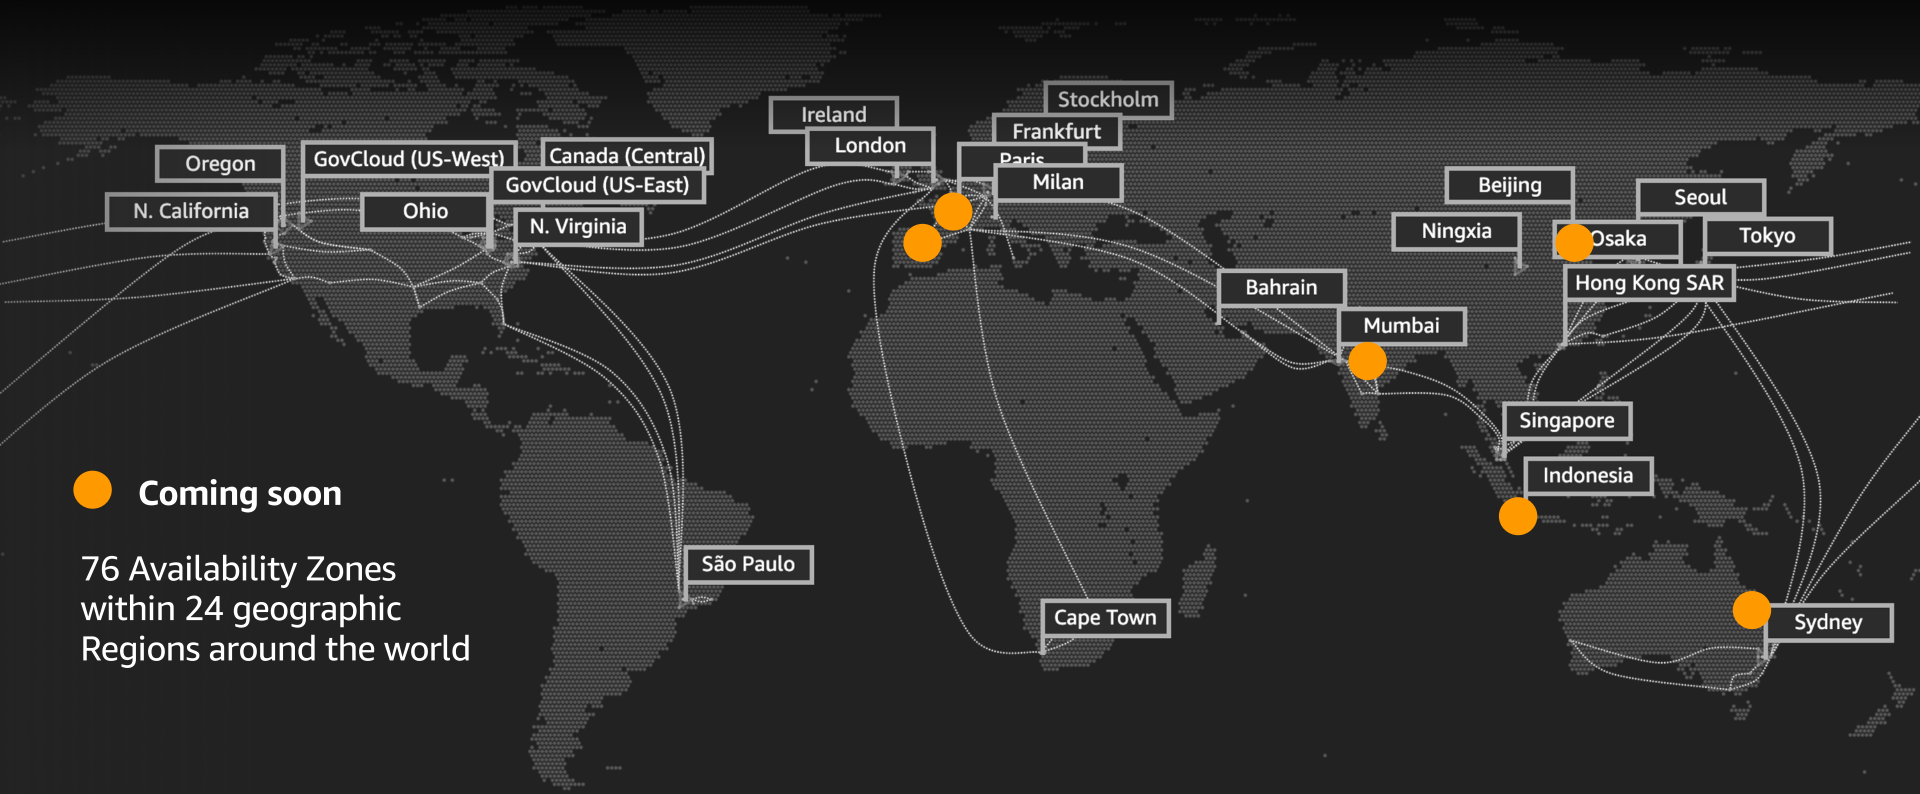

76 Availability Zones within 24 geographic Regions around the world

Security/ITAR : AWS GovCloud (US)

Isolated AWS infrastructure and services for customers with strict regulatory and compliance requirements and sensitive data. In addition AWS Secret Region for Classified, Secret, Top Secret.

OVERFLOW + FUN3D running on thousands of cores every day in AWS GovCloud

Addresses the most stringent US Government regulations, policies and security requirements

Example AWS Customers

amazon Prime Air

BOOM

INEOS | TEAM UK

JOBY AVIATION

™

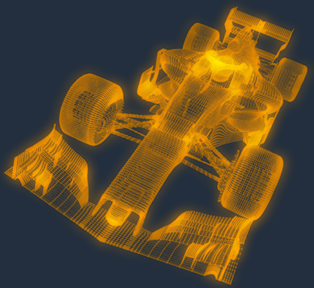

aws

Performance

- Coupled Steady-State RANS
- Multi-grid

- C5n.18xlarge = Intel Skylake cores, 192 GB RAM + 100 Gbi Elastic Fabric Adapter (EFA)

- C5.24xlarge = Intel Cascade L (48 cores, 192GB RAM) + 25GBit/s

- Elastic Fabric Adapter (EFA)* makes a meaningful differen CFD Scaling

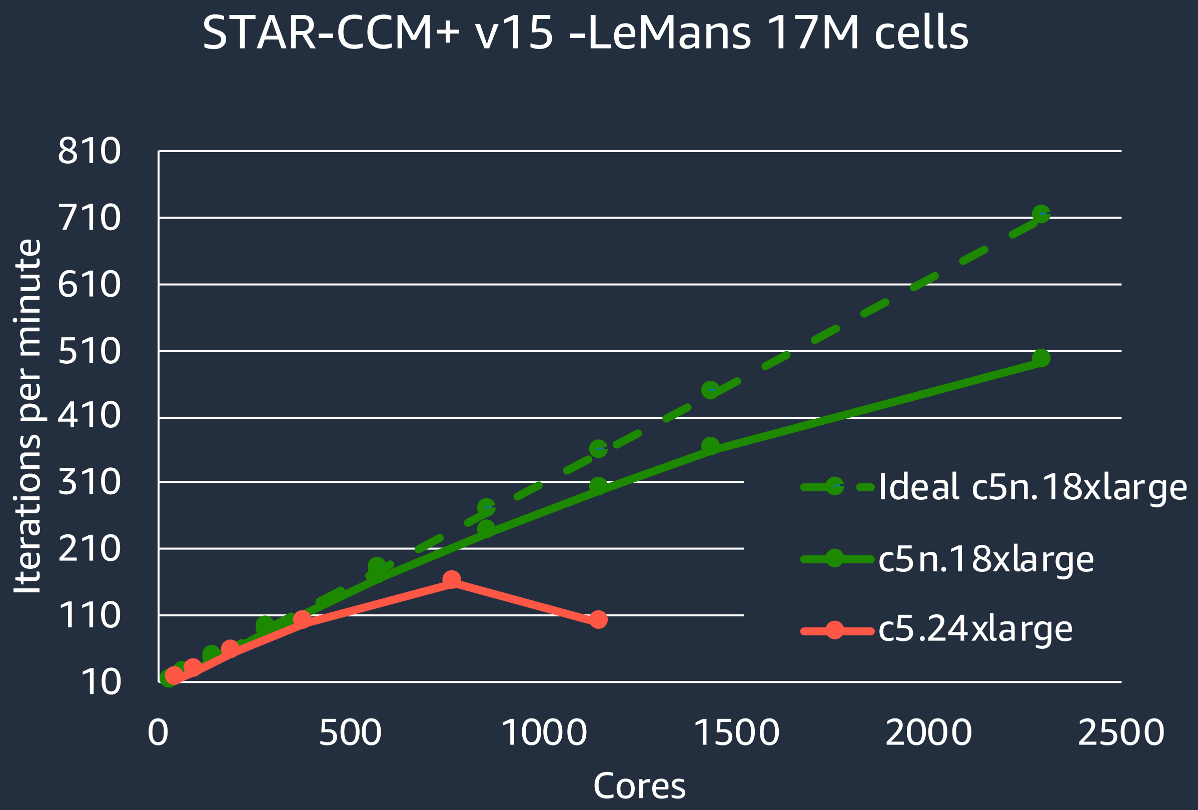

*Shalev, L., Ayoub, H., Bshara, N., & Sabbag, E. (2020). A Cloud-Optimized Transport Protocol for Elastic and Scalable HPC. *IEEE Micro*, *40*(6), 67–73. https://doi.org/10.1109/MM.2020.3016891

Performance

- Coupled Steady-State RANS
- Multi-grid

- C5n.18xlarge = Intel Skylake 36 cores, 192 GB RAM + 100 Gbit/s Elastic Fabric Adapter (EFA)

- C5.24xlarge = Intel Cascade Lake (48 cores, 192GB RAM) + 25GBit/s

- Elastic Fabric Adapter (EFA)* makes a meaningful difference to CFD Scaling

STAR-CCM+ v15 LeMans 17M cells

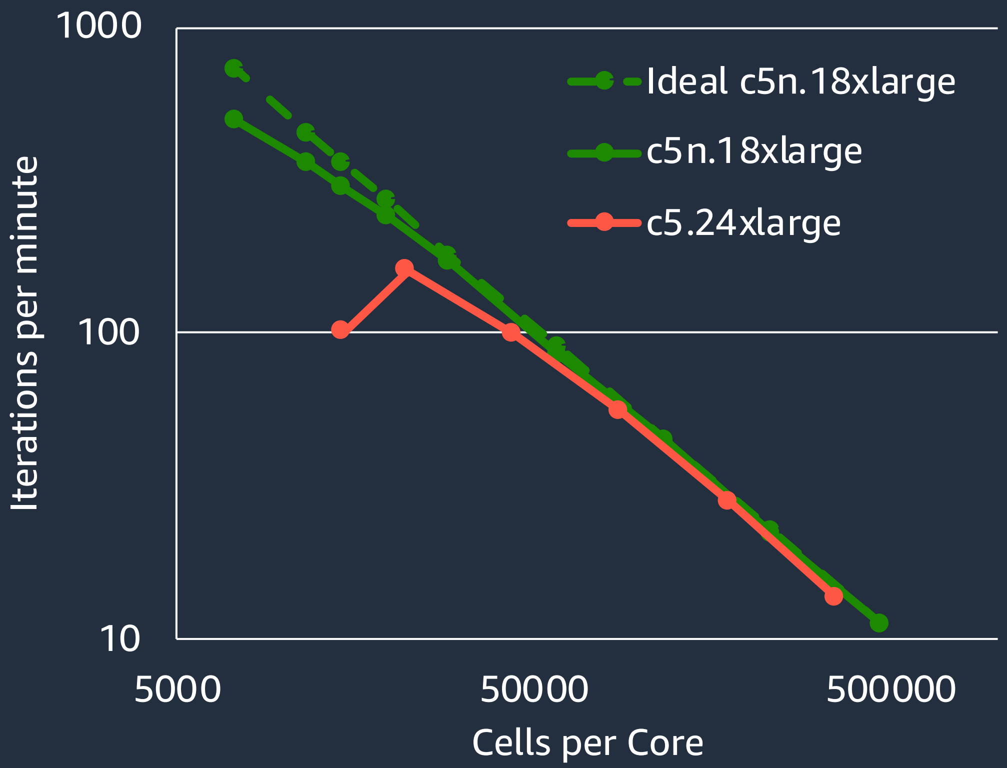

*Shalev, L., Ayoub, H., Bshara, N., & Sabbag, E. (2020). A Cloud-Optimized Transport Protocol for Elastic and Scalable HPC. *IEEE Micro*, *40*(6), 67–73. https://doi.org/10.1109/MM.2020.3016891

Scaling on AWS - OpenFOAM

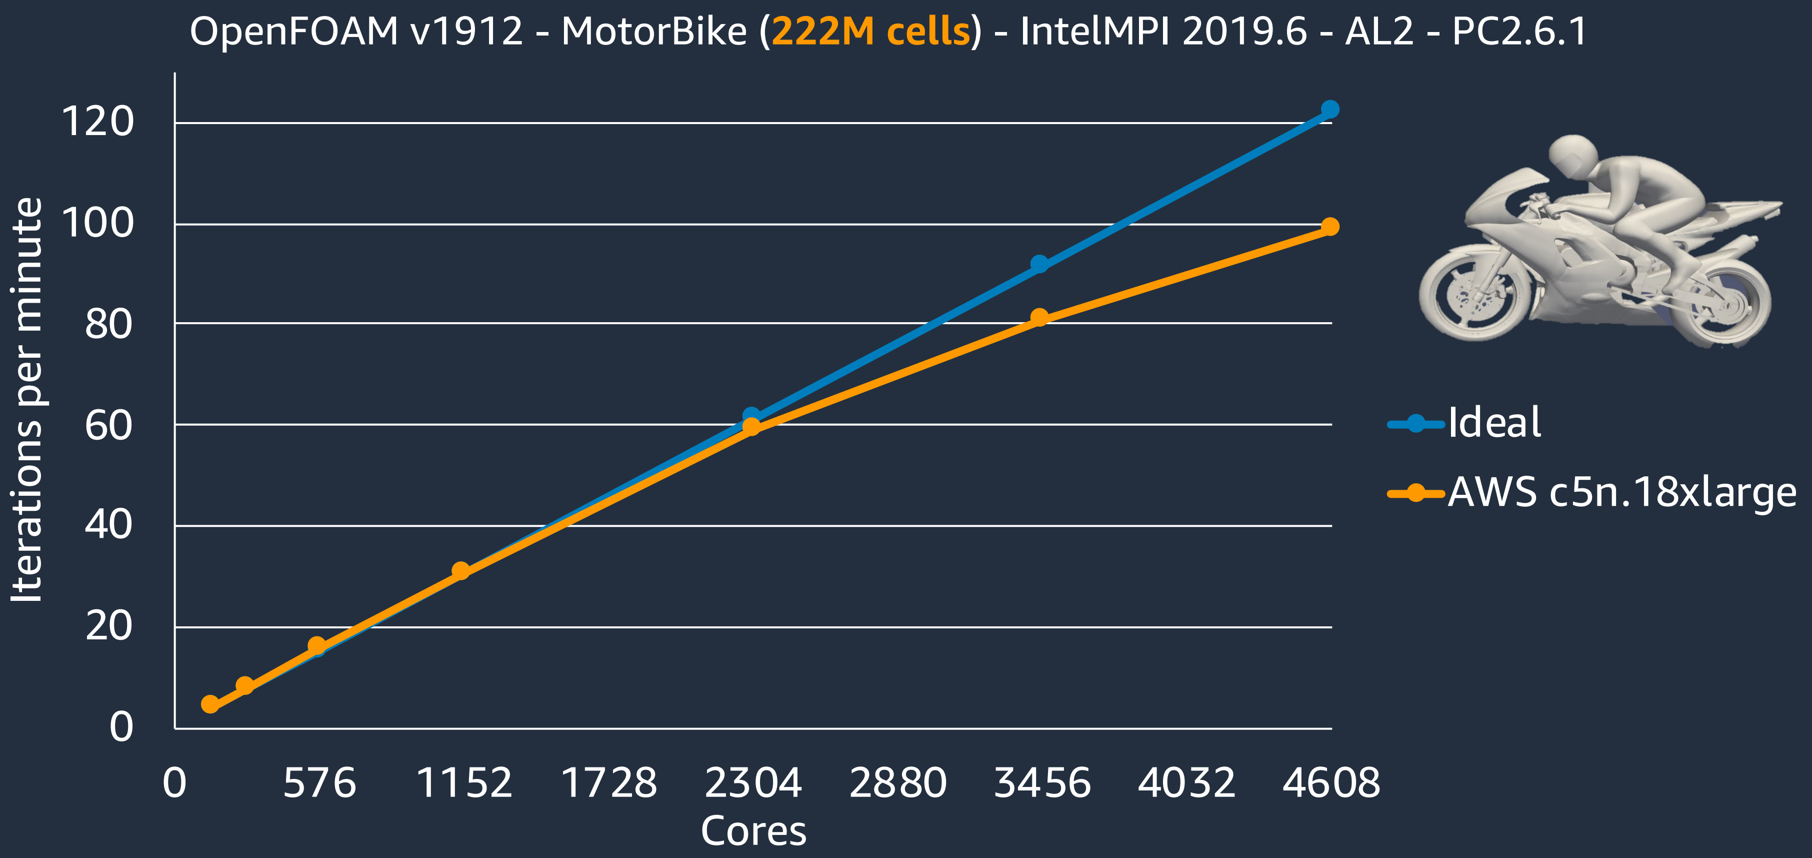

Scaling on AWS - OpenFOAM

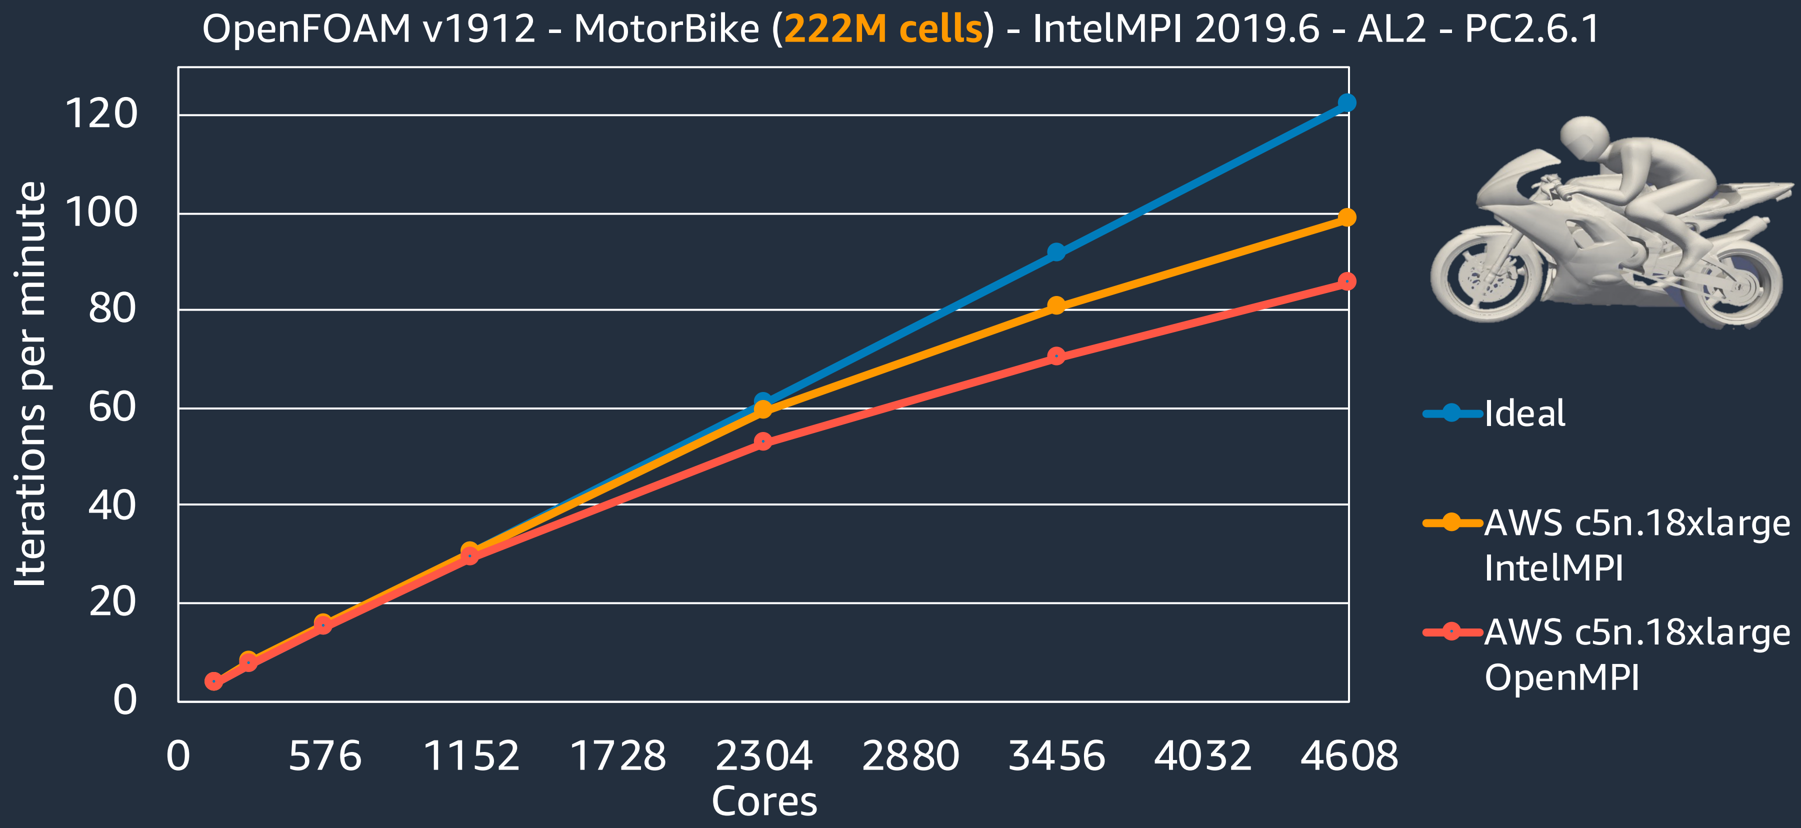

Scaling on AWS - OpenFOAM

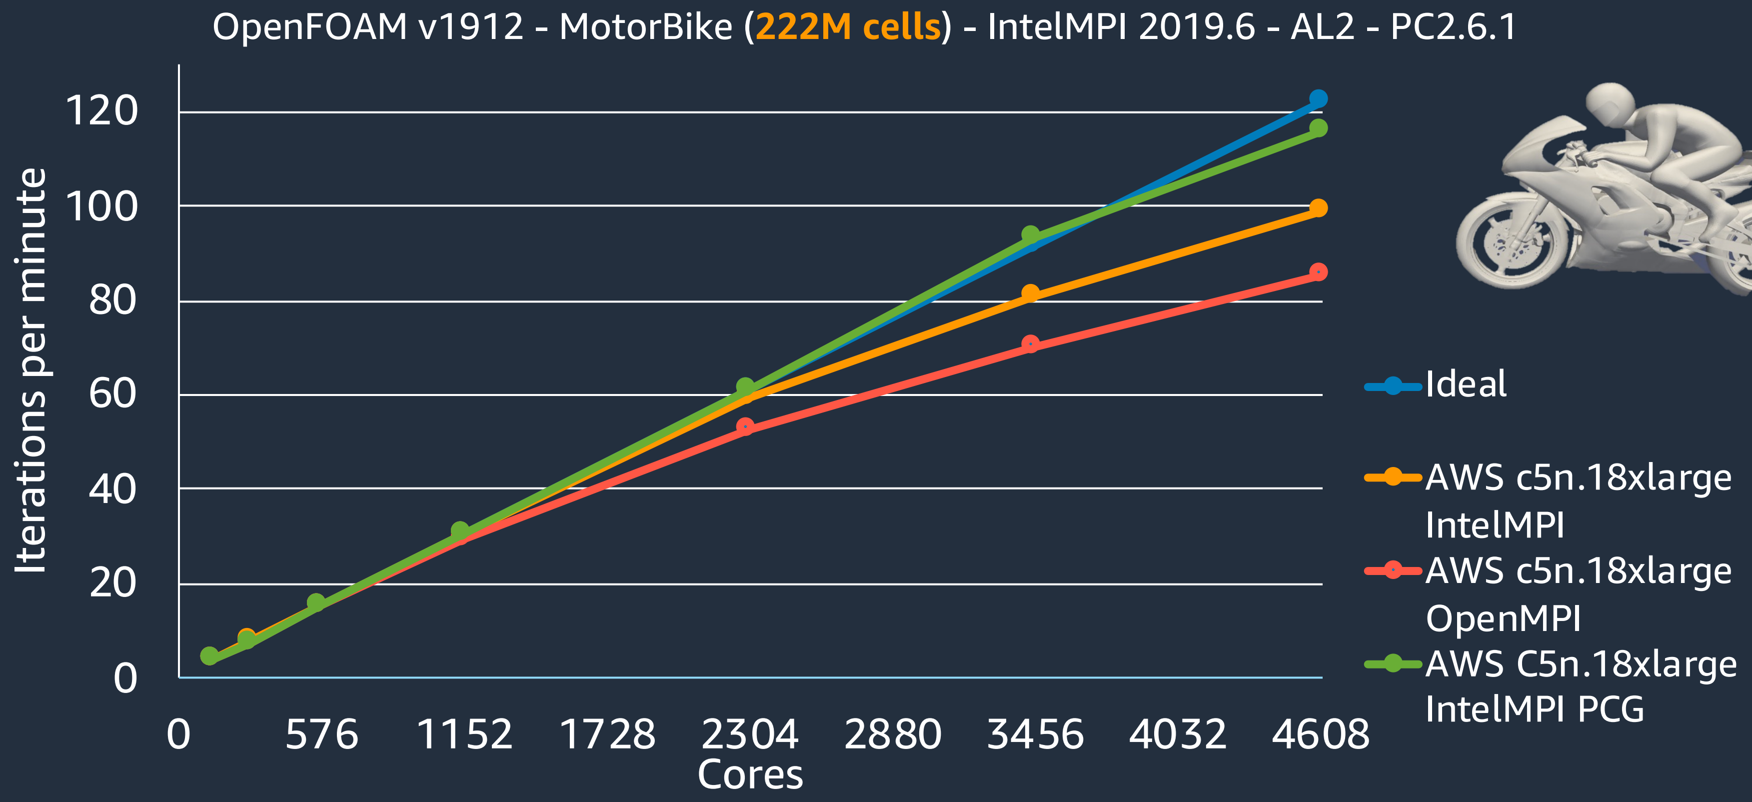

Scaling on AWS – Fluent

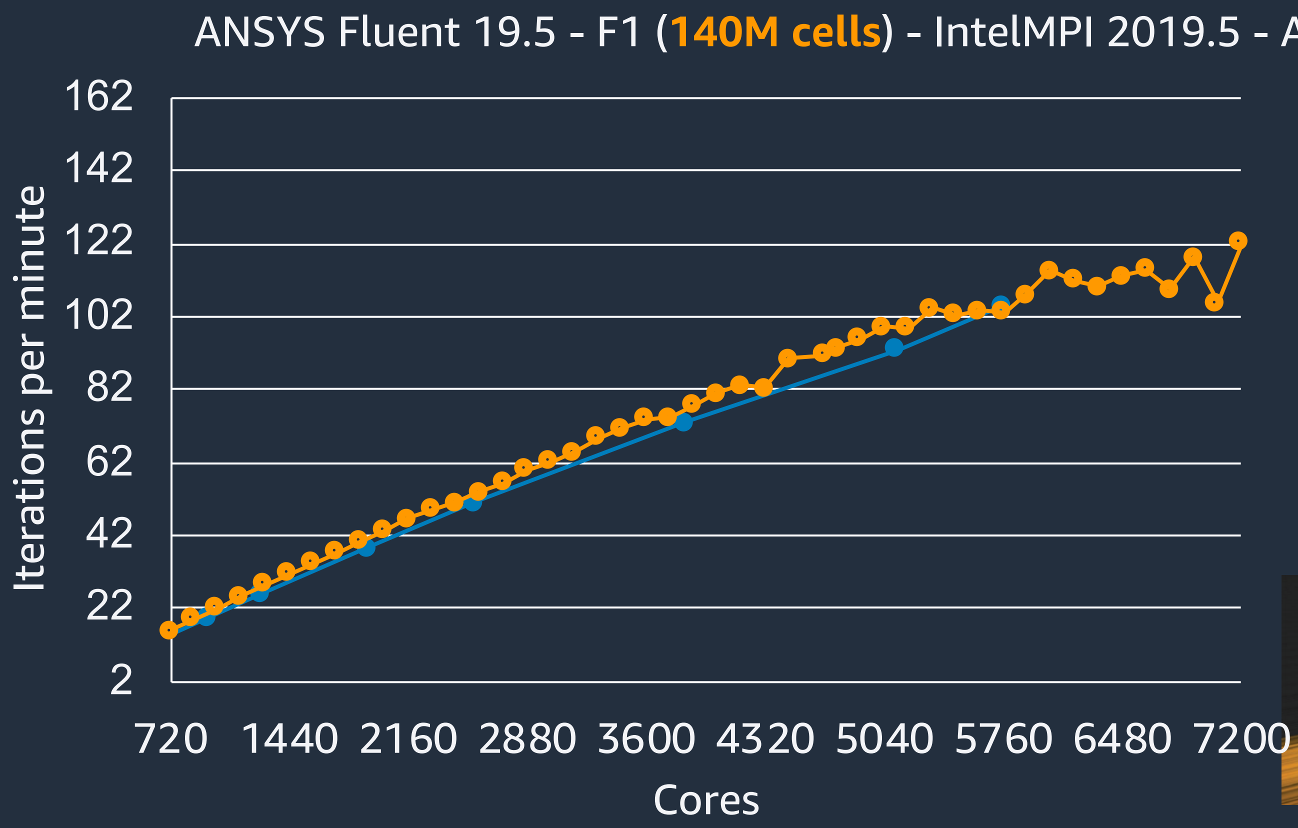

Scaling on AWS – STAR-CCM+

Simcenter STAR-CCM+ 2020.1 - F1 (**403M cells**) - IntelMPI 2019.6 - AL2 - PC2.6.1

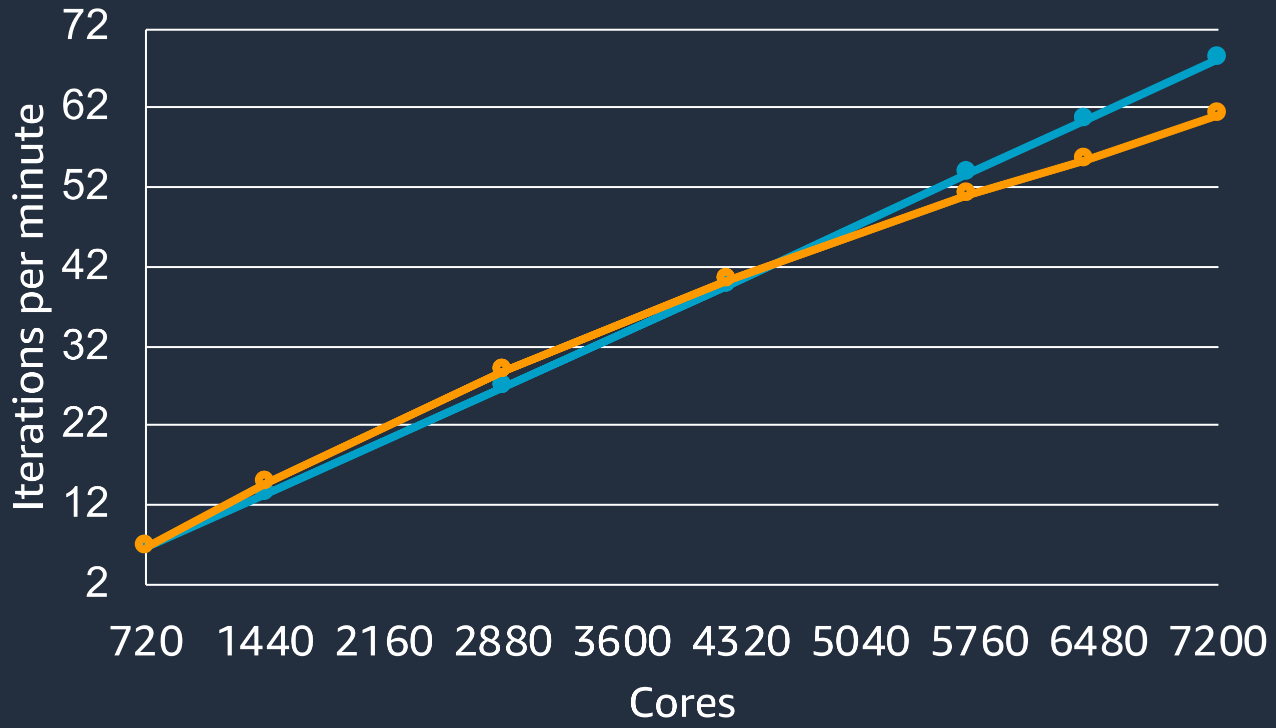

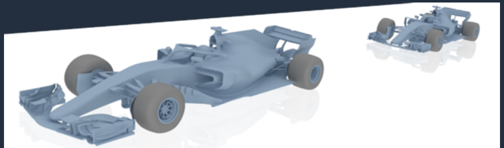

Scaling on AWS – STAR-CCM+

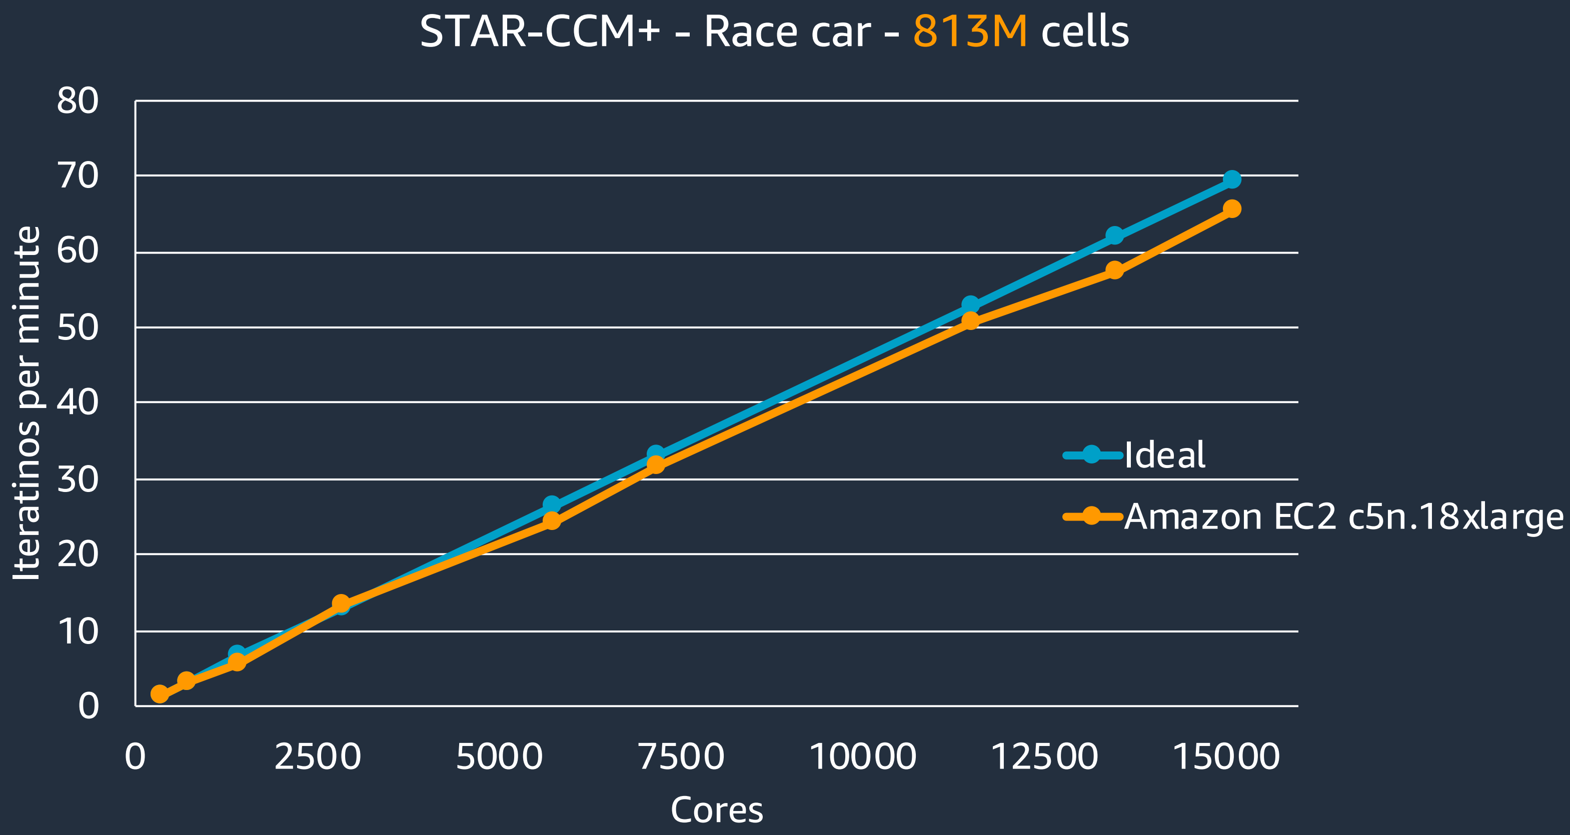

GPUs?

- Recent work with Zenotech (zCFD) is very interesting
- GPU's (p3.24xlarge and g4dn.16xlarge Amazon EC2 instances i.e Nvidia v100 and T4's) deliver faster and cheaper results
- P4d.24xlarge (NVIDIA A100) gave even higher boost for a lower cost.
- Full results presented at AIAA Scitech 2021

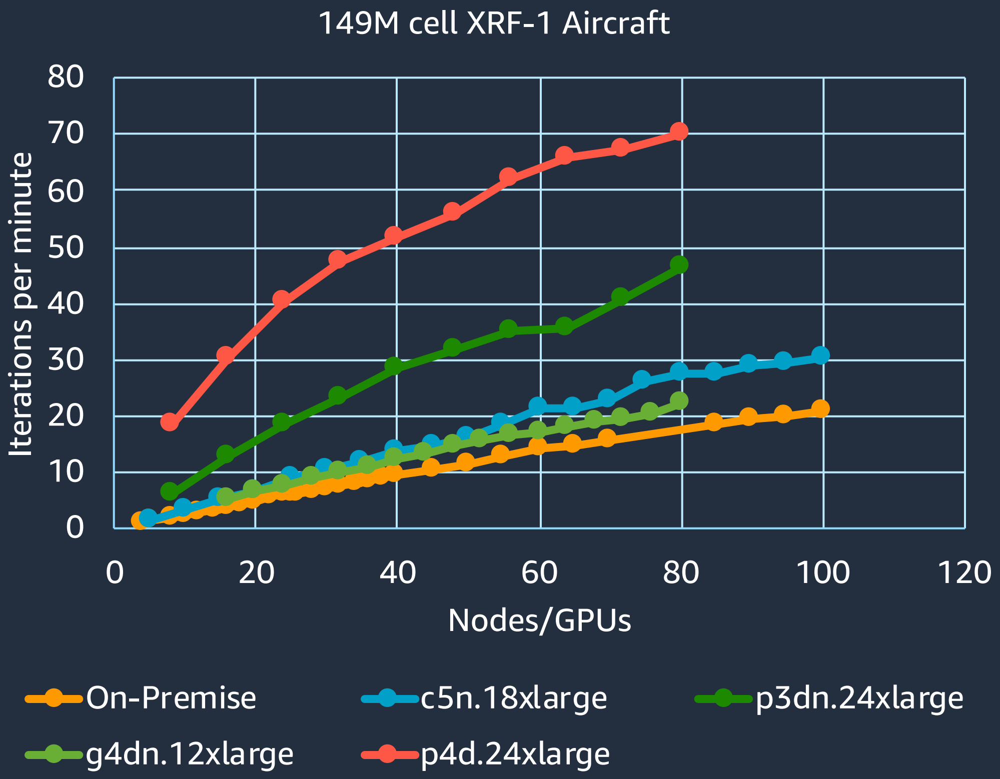

GPUs?

- Recent work with Zenotech (zCFD) is very interesting
- GPU’s (p3.24xlarge and g4dn.16xlarge Amazon EC2 instances i.e Nvidia v100 and T4’s) deliver faster and cheaper results
- P4d.24xlarge (NVIDIA A100) gave even higher boost for a lower cost.
- Full results presented at AIAA Scitech 2021

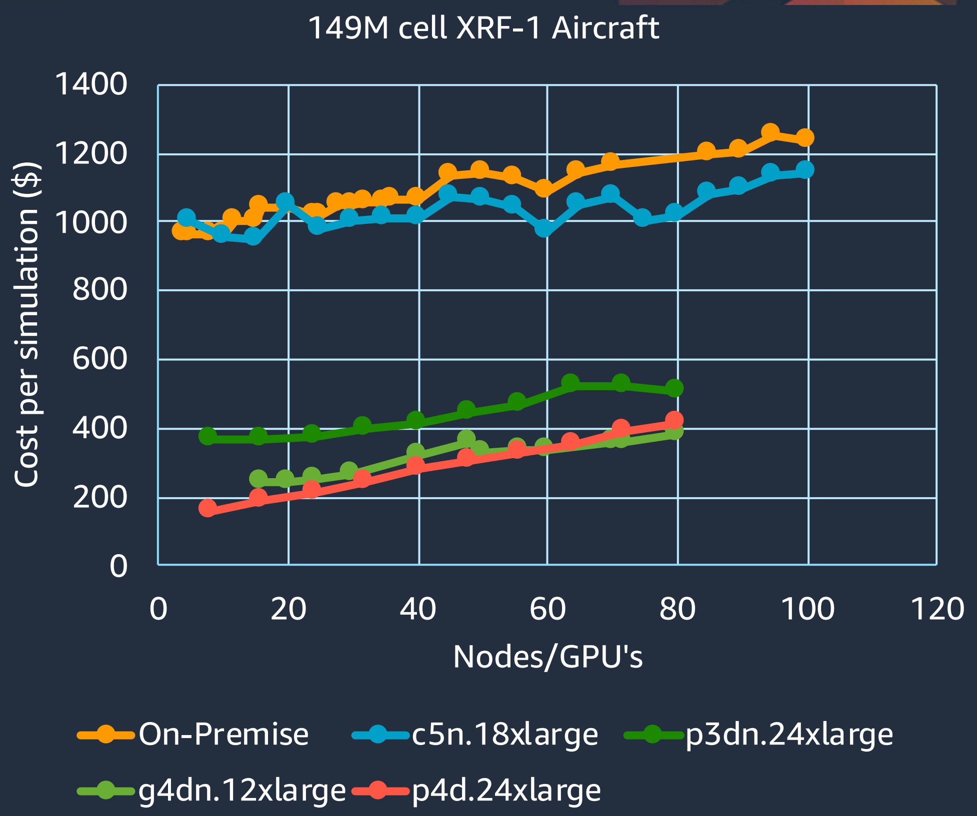

Graviton2 Arm option

C6gn.16xlarge (GA soon) - 64 core, 128GB RAM, 100 Gbit/s EFA offers > 20% saving over x86 Intel

Combined with AMD's + GPU's, allows CFD developments to test their code on a wide range of hardware

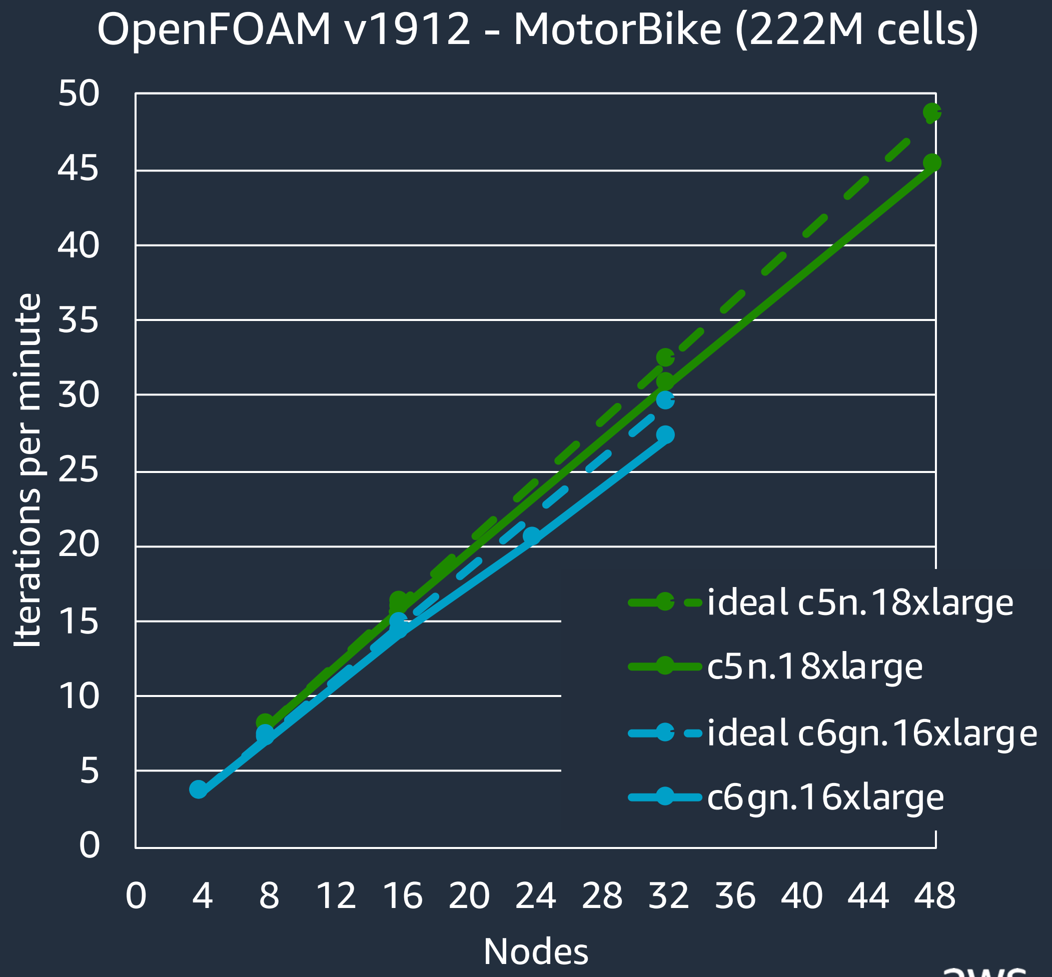

Wide choice of Architectures + CPU/Mem options

Categories

General purpose

Burstable

Compute intensive

Memory intensive

Storage (High I/O)

Dense storage

GPU compute

Graphics intensive

Capabilities

Choice of processor
(AWS, Intel, AMD)

Fast processors
(up to 4.5 GHz)

High memory footprint
(up to 24 TiB)

Instance storage
(HDD and NVMe)

Accelerated computing
(GPUs and FPGA)

Networking
(up to 400 Gbps)

Bare Metal

Size
(Nano to 32xlarge)

Options

Elastic Block Store

Elastic Graphics

Elastic Inference

+ +=

~400
instance types

Amazon FSx for Lustre - High and scalable performance

High and scalable performance

Parallel File System

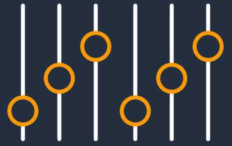

100+ GiB/s throughput
Millions of IOPS
Consistent sub-millisecond latencies

SSD or HDD-based

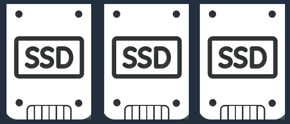

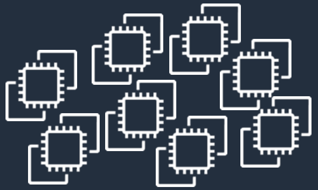

Supports concurrent access from hundreds of thousands of cores

Each terabyte (TB) of storage provides 200 MB/second of file system throughput and ~5,000 IOPS

Several cost models

Scale using Spot, On-Demand, or both

Use Reserved Instances/Saving plans for known/steady-state workloads

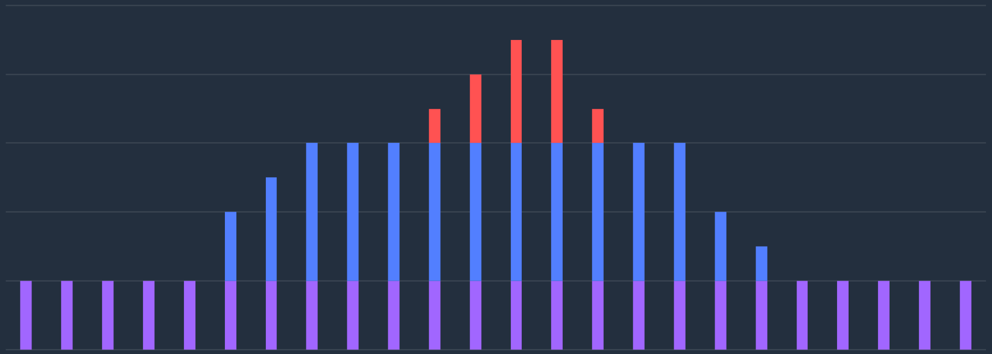

Evaluate the trade-off of time to solution vs. cost for scaling

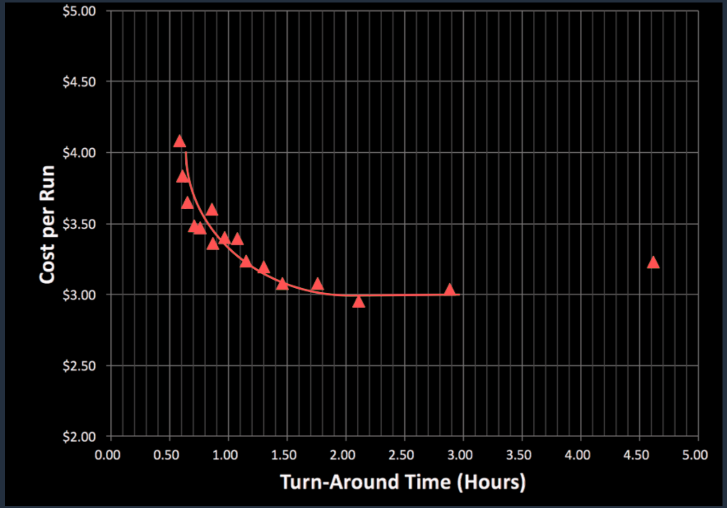

Ways to use AWS for collaboration

Collaborations

- Spin up an EC2 instance (or use HPC orchestration services like AWS ParallelCluster or AWS Batch), install your code and run your cases
- Save as a Amazon Machine Image (AMI) – share that AMI securely with your collaborators
- e.g Prof Peter Vincent (Imperial) used AWS for the PyFR Symposium. Created an AMI with PyFR installed that everyone can use on a Nvidia V100. Head to http://www.pyfr.org/events.php to use it.

Pre/Post Processing blockers for peta-exa scale

Meshing is still largely limited to a single-node for many packages e.g Pointwise, ANSA

Key challenge is to have access to a node with enough RAM to build and then visualize them. Majority of on-prem machines only have 256GB/512GB fat nodes and these are limited.

Cloud has very high RAM because of needs beyond just HPC e.g X instances up to 4TB per node

Remote visualization has been a key COVID-19 requirement – which helps look at huge meshes. Low latency because of global infrastructure

Enabling ML/AI for CFD

Very active research area, deserves a talk of it's own. Turbulence modelling, design optimization etc.

Key need is for data i.e never want to delete a simulation result.

GPU's for training models but also we want to train models using more than CFD data e.g sensor data from physical tests etc

Huge increase in ML services on the cloud which CFD can take advantage of

Summary

- Turbulence modelling is the key to accuracy for many problems.
- HPC is the main bottleneck to achieve RANS like turn-around time
- Cloud model can help to broaden HPC access to companies/researchers and allow testing on latest hardware
- Enables easy global collaboration and sharing large files
- Exciting times ahead for ML+CFD, hope cloud can help to accelerate that.

Thank you!

neashton@amazon.com