How Cloud Computing, Machine Learning and High Performance Computing (HPC) are Accelerating the Adoption of HRLES and WMLES Methods for Industry

Advanced Modeling & Simulation (AMS) Seminar Series
NASA Ames Research Center, March 28th, 2023

Dr. Neil Ashton
Principal Computational Engineering Specialist, AWS
neashton@amazon.com

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

The role of CFD in industry









Companies need a way to assess the performance and efficiency of new designs

Aerodynamics is key, especially in the age of electric vehicles

They want to go from conceptual design to full design in the shortest possible time with the least expense

Digital certification is the dream for many industries to reduce cost and time compared to physical tests; let's explore that further

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Use of CFD is almost always compared to cost & accuracy of physical tests





The central theme of this talk is what can be done to make CFD either much cheaper, more accurate or ideally both. What role does HPC, cloud & ML/AI have in this?

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

Cost of CFD: 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
- Regardless of whether the code is able to run more complex geometries or the HPC resources are available – the cost compared to a physical test
- Has value for understanding fundamental nature of flow physics & helps turbulence modelling

---

High-Lift aircraft example (M=0.2) (WMLES use ~300M cells with $t=2e^{-6}$s at 0.05s per time-step at 30k cells per core for RK3 explicit)

N_C=10B

$\Delta t=1e^{-7}$

N_T=40M (4s)

N_p=333k CPU cores (23 days)*

N_p=1100 GPUs (23 days)***

~$5M**

~$3M**

*0.05s per time-step (30k cells per core)

**$0.03 per core hour & $5 per GPU hour

*** assuming 300 CPU cores is x1 latest gen GPU

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

Cost of CFD: Direct Numerical Simulation

- Could DNS go from $3M now to $3000? i.e x1000 cheaper
- Could DNS go from 23 days now to 5hrs? i.e x100 faster

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

Cost of CFD: Direct Numerical Simulation

- Could DNS go from $3M now to $3000? i.e x1000 cheaper
- Could DNS go from 23 days now to 5hrs? i.e x100 faster

- Reasonable to assume x2 faster due to highly optimized numerical schemes, linear solvers, code optimization i.e 12 days or $1.5M.

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

Cost of CFD: Direct Numerical Simulation

- Could DNS go from $3M now to $3000? i.e x1000 cheaper
- Could DNS go from 23 days now to 5hrs? i.e x100 faster

- Reasonable to assume x2 faster due to highly optimized numerical schemes, linear solvers, code optimization i.e 12 days or $1.5M.
- Reasonable to assume x2 faster due to next-gen GPU-like technology in the next 5-10 years i.e 6 days and $750,000.

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Cost of CFD: Direct Numerical Simulation

- Could DNS go from $3M now to $3000? i.e x1000 cheaper
- Could DNS go from 23 days now to 5hrs? i.e x100 faster

- Reasonable to assume x2 faster due to highly optimized numerical schemes, linear solvers, code optimization i.e 12 days or $1.5M.
- Reasonable to assume x2 faster due to next-gen GPU-like technology in the next 5-10 years i.e 6 days and $750,000.
- But that is $750,000 per run. Even if we assume $300,000 per run due to even greater technology innovation – we often want to run 10 different scenarios i.e $3M for 10 runs. It's hard to see DNS in the next 5-10 days ever breaking into industrial design because the cost of a physical test would be cheaper for 95% of problems.
- Also, does your company have data-center space for 100,000 GPUs? i.e 1000 GPU per run, so x100 for your company? Logical to look at hyperscalers like AWS to worry about data-centers.
- DNS is becoming more affordable to run larger Re problems and deal with more complex cases. But…

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

Cost of CFD: Direct Numerical Simulation

- Could DNS go from $3M now to $3000? i.e x1000 cheaper
- Could DNS go from 23 days now to 5hrs? i.e x100 faster

- Reasonable to assume x2 faster due to highly optimized numerical schemes, linear solvers, code optimization i.e 12 days or $1.5M.
- Reasonable to assume x3 faster due to next-gen GPU-like technology in the next 5-10 years i.e 6 days and $750,000.
- But that is $750,000 per run. Even if we assume $300,000 per run due to even greater technology innovation – we often want to run 10 different scenarios i.e $3M for 10 runs. It's hard to see DNS in the next 5-10 days ever breaking into industrial design because the cost of a physical test would be cheaper for 95% of problems.
- Also, does your company have data-center space for 100,000 GPUs? i.e 1000 GPU per run, so x100 for your company? Logical to look at hyperscalers like AWS to worry about data-centers.
- DNS is becoming more affordable to run larger Re problems and deal with more complex cases. But…

But what if Re was 10 higher? Lets therefore restrict focus on LES-like & RANS methods

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

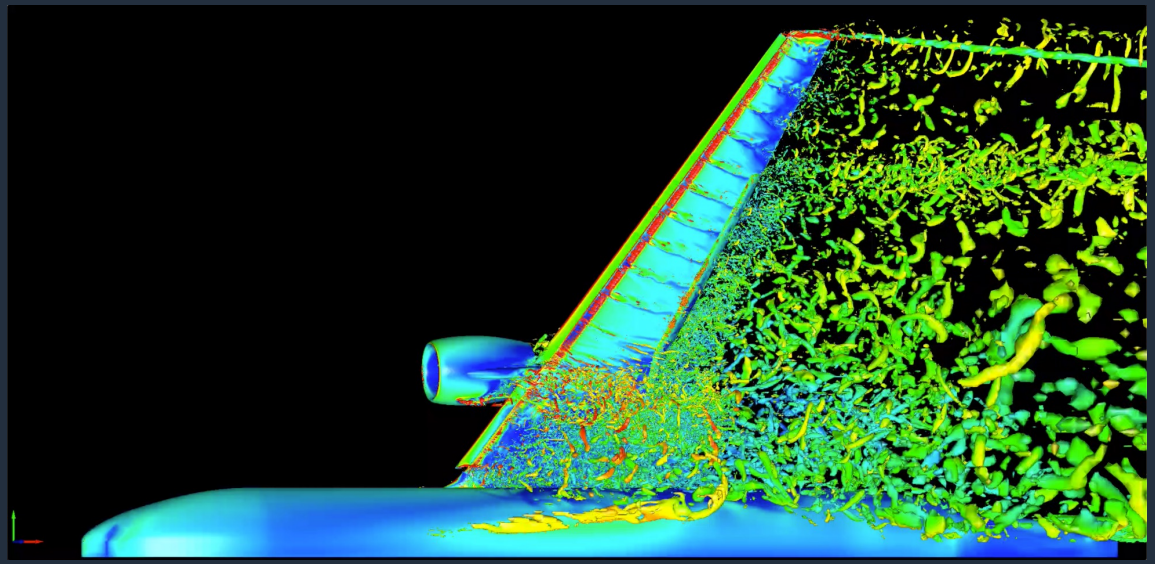
aws

High-Lift Prediction Workshop

Advance the state of the art for CFD with Boeing/NASA and many others.

Full details here: https://hiliftpw.larc.nasa.gov/

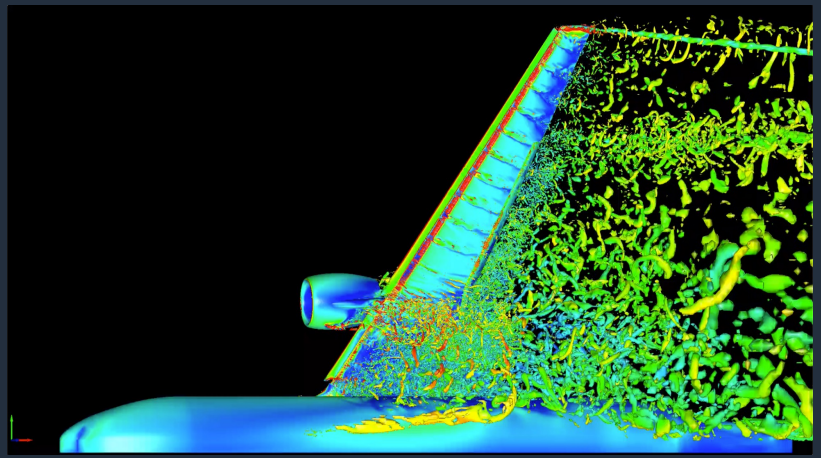
Summary papers are due to be out very soon in the AIAA Journal – for now keep an eye on the website and look through the summary presentations.



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

4th High-Lift Test-Case



Mach Number	0.2
Angles of attack	2.78, 7.05, 11.29, 17.05, 19.57, 20.55, 21.47° (wall corrected)
Reynolds Number based on MAC	5.49 million
Reference Static Temperature	521 °R
Reference Static Pressure	24.67 psi
Mean Aerodynamic Chord (MAC)	275.8 inches
Moment Reference Center (MRC)	x = 1325.9 inches, y = 0.0 inches, z = 177.95 inches
Semi-span model reference area	297,360.0 in^2
Flap Deflection	40°/37° inboard/outboard (nominal)
Important Details:	• Geometry is provided in full-scale inches. • For RANS, run simulations fully turbulent. • All simulations are “free air”; no wind tunnel walls or model support systems.
Data Delivery	Forces/moments, surface pressures, surface streamlines, skin friction contours, off-body velocity profiles, etc.

Work was conducted with Paul Batten (Metacomp Technologies, Inc.) & Vangelis Skaperdas (BETA-CAE Systems)

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Motivations

- Prior HLPW workshops showed a clear trend of the failure of RANS models to predict CL_Max, even at second-moment closure level.
- Where RANS did capture the lift/drag curves correctly it was typically for the wrong reasons
- Can we learn the lessons of the automotive sector – where RANS were failing and almost the entire industry has moved to hybrid RANS-LES methods even where the cost is higher.
- Hypothesis was that the same could be true for the aerospace sector – but we need to prove it and provide best-practices to those companies who wish to test HRLES vs RANS themselves. Also what are the HPC costs?

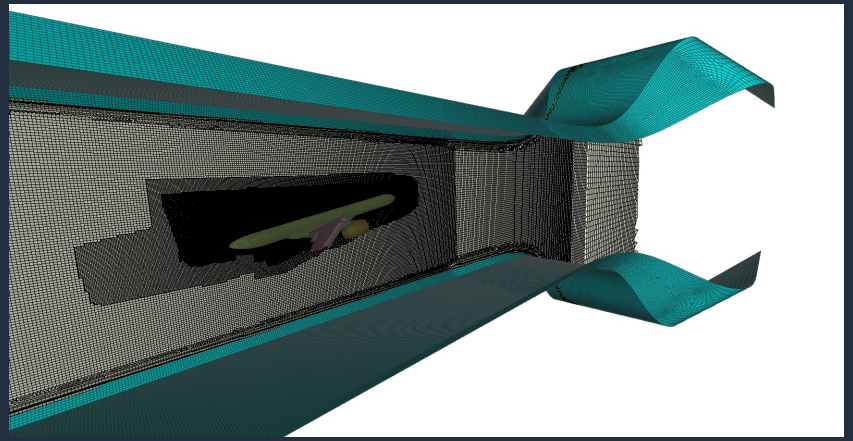
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

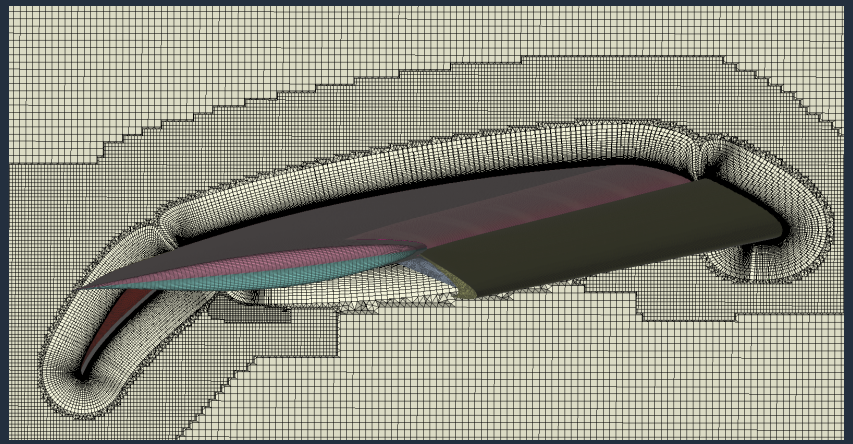
aws

Computational meshes

- Unstructured hexa/prismatic grids built using ANSA v21 by BETA-CAE Systems to give more control over committee grids for HRLES adaptation
- 74 prism layers with $y^+ < 1$
- Built for HRLES from the start
- All available on the HLPW4 website in various formats
- https://hiliftpw.larc.nasa.gov/Workshop4/grids_downloads.html

	Free-air	WT
Mesh A	91M	100M
Mesh B	173M	156M
Mesh C	276M	278M
Mesh C+	n/a	335M
Mesh D+	n/a	548M





© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Computational Setup

- Unstructured code; CFD++ 20.1 by MetaComp Technologies
- Smart SGS to remove numerical viscosity from SGS dissipation
- Implicit dual-time stepping
- SA-QCR RANS and SA-QCR DDES
- All simulations run for at least 40 CTU's to ensure sufficient time-averaging
- Summarized in the associated paper [3] – as well as journal paper with latest results in prep.

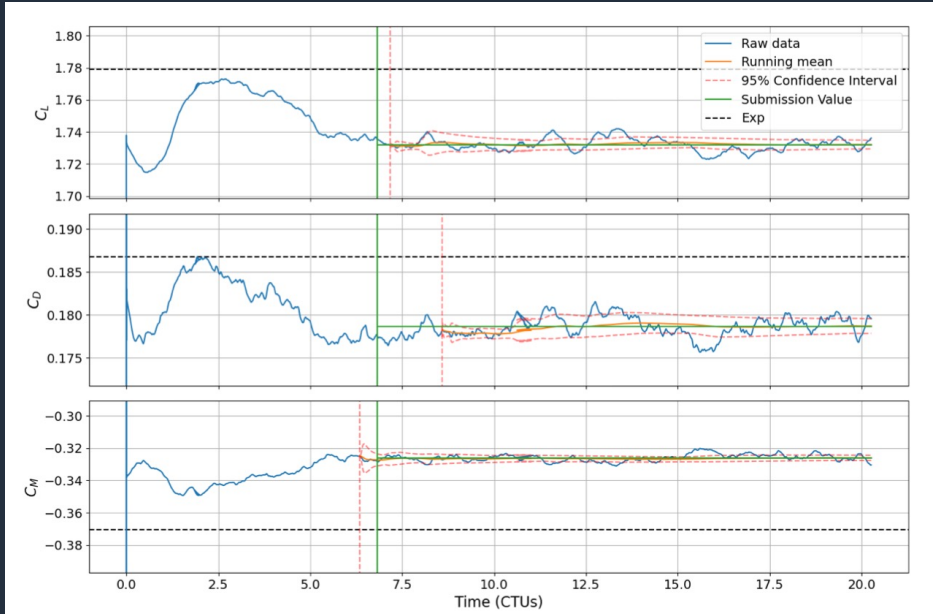


Figure 4: Example of meancalc tool output for L-001.2 case 2a free-air DDES at 7.05° angle of attack

[3] Ashton,N.,Eberhardt,S.,Batten,P.,andSkaperdas,V.,"Towards best-practicesforHybrid RANS-LES simulations of high-lift aircraft geometries," AIAA Aviation 2022, 2022.

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Computational Resources

HPC is arguably one of the biggest bottlenecks to the adoption of HRLES – hoping that cloud computing can help address the issue of availability of resources (especially for enterprise who don't have access to DoE/DoD/NASA research supercomputers)

AWS ParallelCluster for dynamic SLURM-based cluster
Amazon EC2 HPC6a.48xlarge (AMD Milan 96 core node + 384GB RAM)
Elastic Fabric Adapter (EFA) for network interconnect
24-48 nodes per run (2304-4608 cores) – often up to 10 cases running at the same time i.e ~40,000 cores

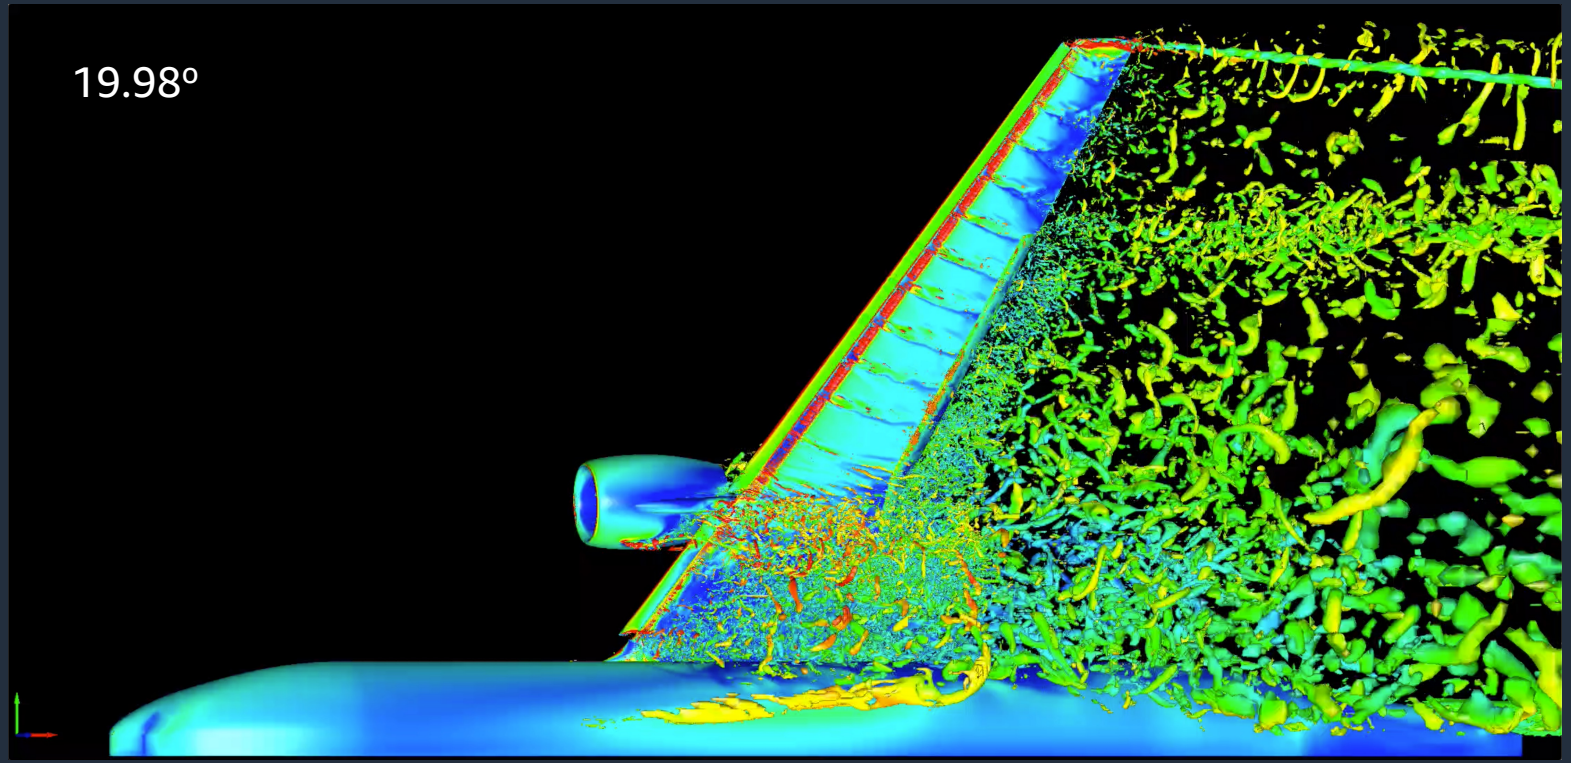
Amazon FSx Lustre parallel file-system (9.6TB + 70TB S3 bucket backend)

Amazon EC2 g3.16xlarge GPU node (488GB +32 core + Nvidia M60 GPU) + NICE DCV for remote visualization (shared sessions to more easily collaborate)

~2M core-hours per DDES simulation (x7 AoA) and ~200k core-hours per RANS simulation

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

General flow physics



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

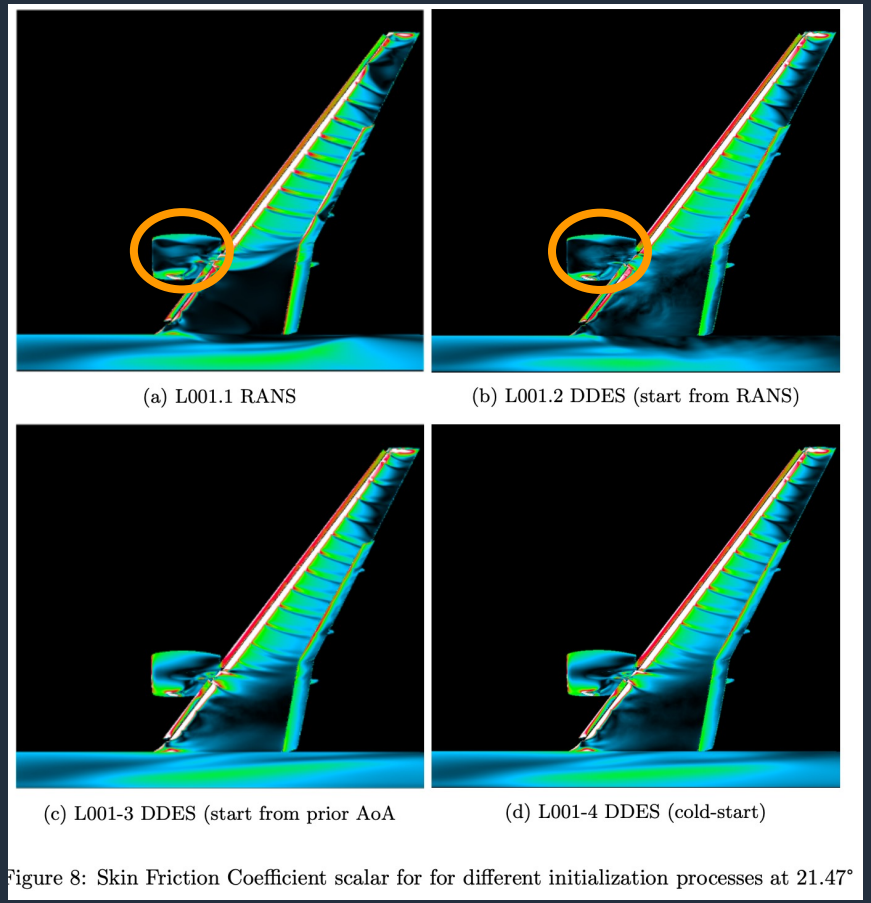
Solution Initialization and AoA Advancement (free-air) – KQ3

RANS vs HRLES - KQ1

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Solution Initialization and AoA Advancement (free-air)

- The choice of initial conditions for the HRLES was found to be critical.
- The common practice of starting an HRLES from an a priori steady-state RANS calculation proved inadequate in this context, with stalled RANS solutions resulting in a stalled HRLES that would take an indeterminately (possibly infinitely) long period of time to recover.
- More physically-realistic initial conditions helped to mitigate otherwise prohibitively-long run times.



(a) L001.1 RANS

(b) L001.2 DDES (start from RANS)

(c) L001-3 DDES (start from prior AoA

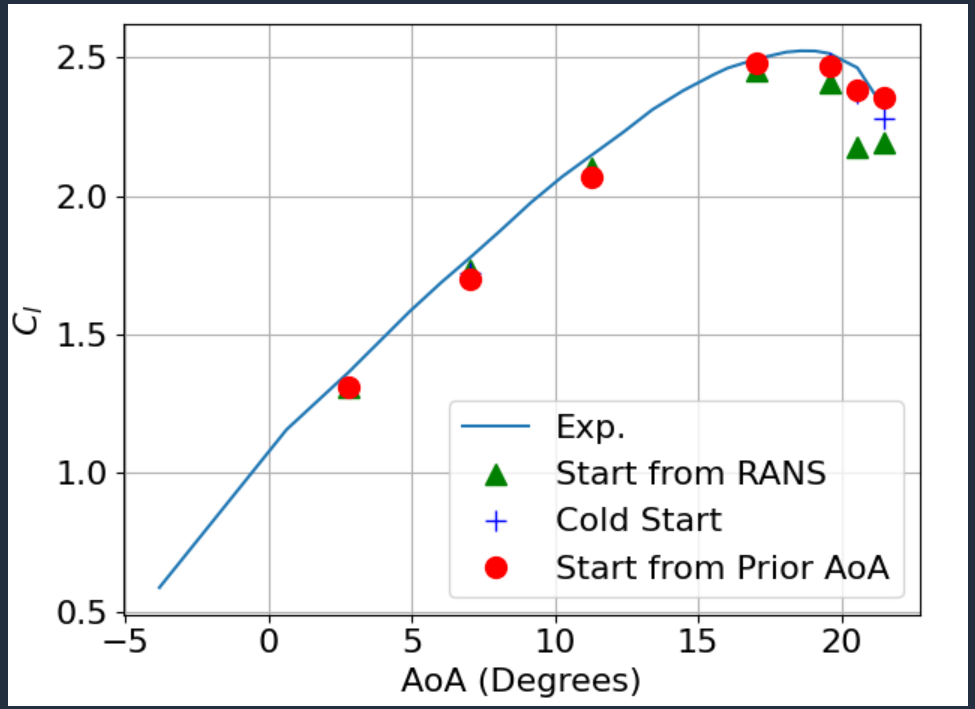
(d) L001-4 DDES (cold-start)

Figure 8: Skin Friction Coefficient scalar for for different initialization processes at 21.47°

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Solution Initialization and AoA Advancement (free-air)

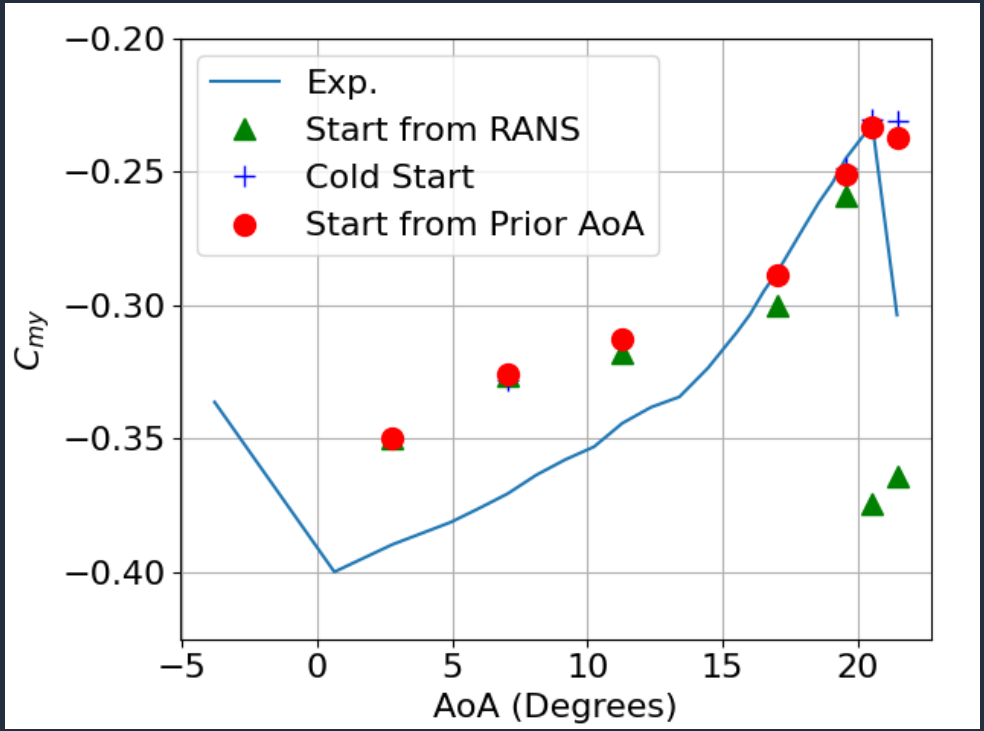
- The choice of initial conditions for the HRLES was found to be critical.
- The common practice of starting an HRLES from an a priori steady-state RANS calculation proved inadequate in this context, with stalled RANS solutions resulting in a stalled HRLES that would take an indeterminately (possibly infinitely) long period of time to recover.
- More physically-realistic initial conditions helped to mitigate otherwise prohibitively-long run times.



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

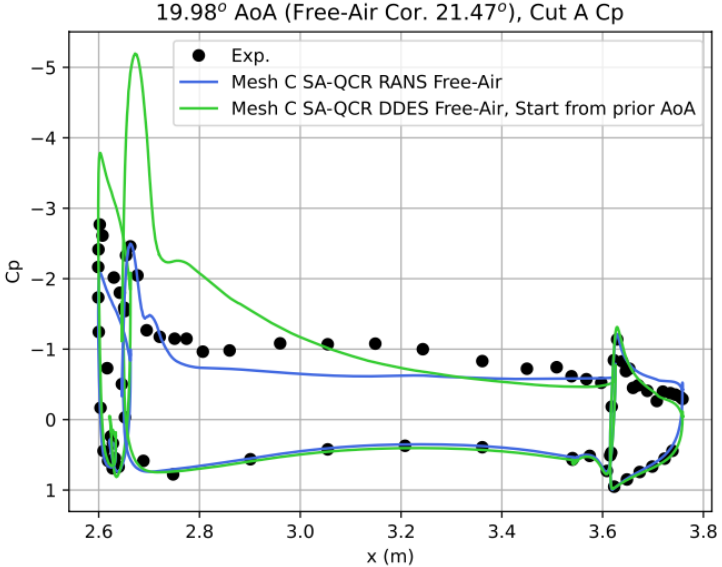
Solution Initialization and AoA Advancement (free-air)

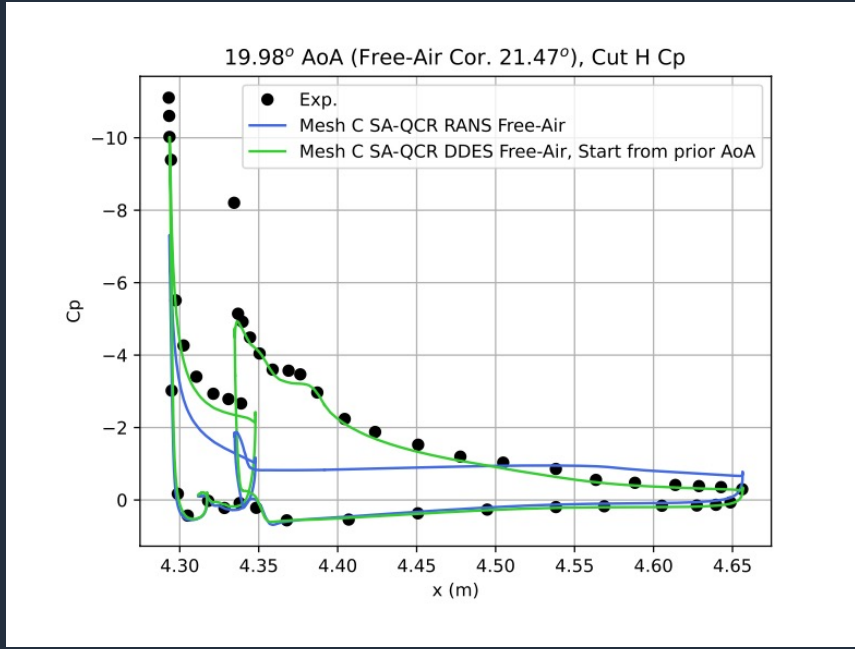
- The choice of initial conditions for the HRLES was found to be critical.
- The common practice of starting an HRLES from an a priori steady-state RANS calculation proved inadequate in this context, with stalled RANS solutions resulting in a stalled HRLES that would take an indeterminately (possibly infinitely) long period of time to recover.
- More physically-realistic initial conditions helped to mitigate otherwise prohibitively-long run times.



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

RANS vs HRLES





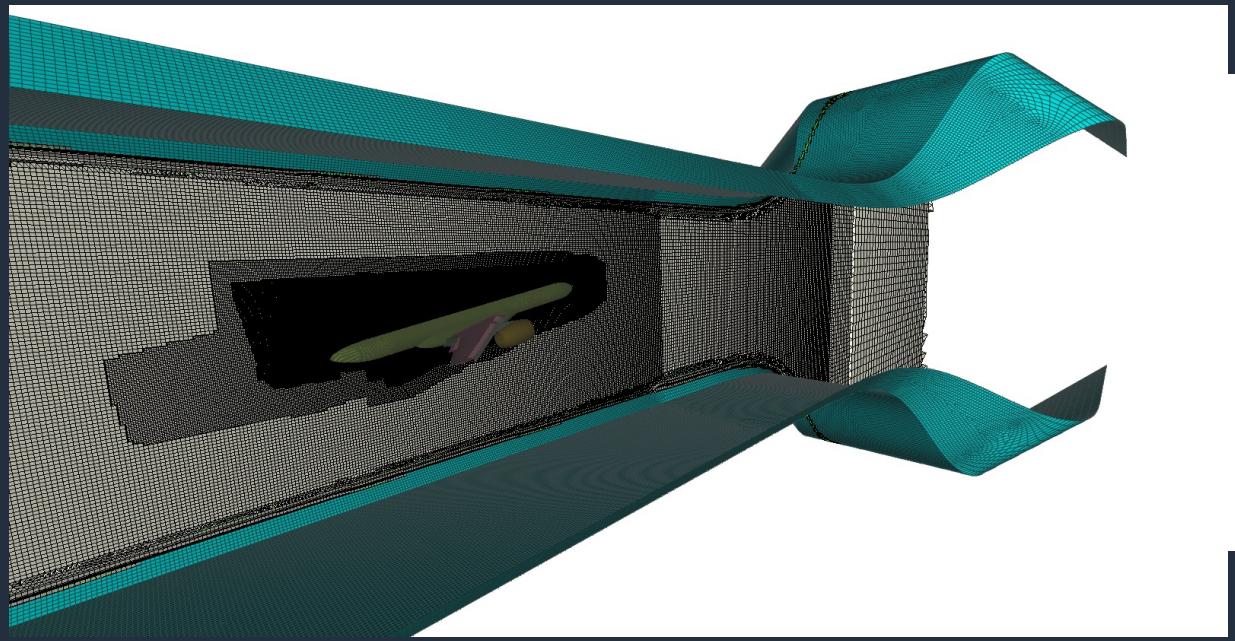
Interesting that RANS actually does better than HRLES at Cut A for free-air. Before going further we realized we need to investigate whether free-air is a good enough assumption or whether WT is needed

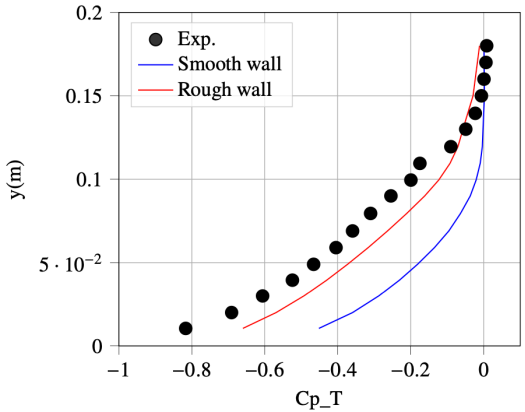
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind Tunnel Modeling – KQ6

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-Tunnel modelling

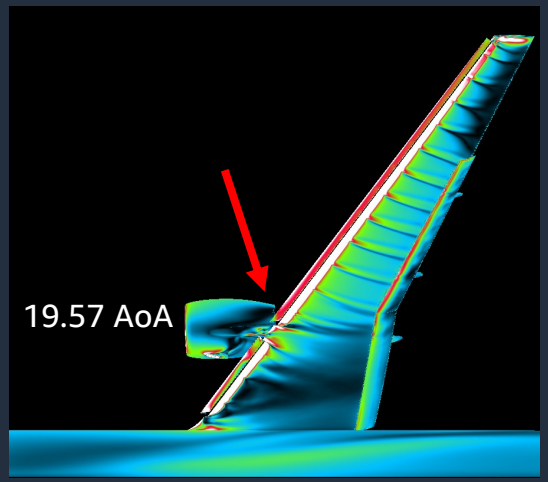


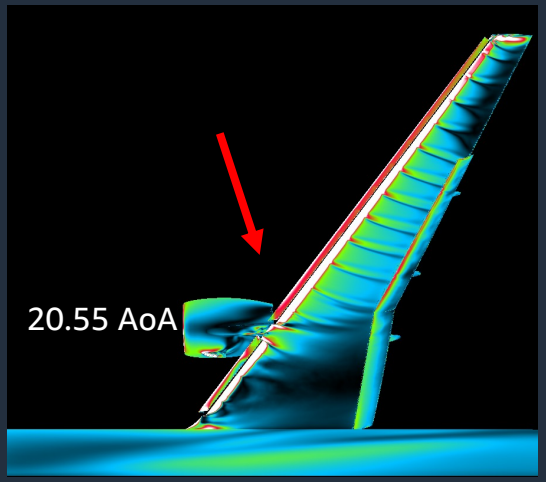


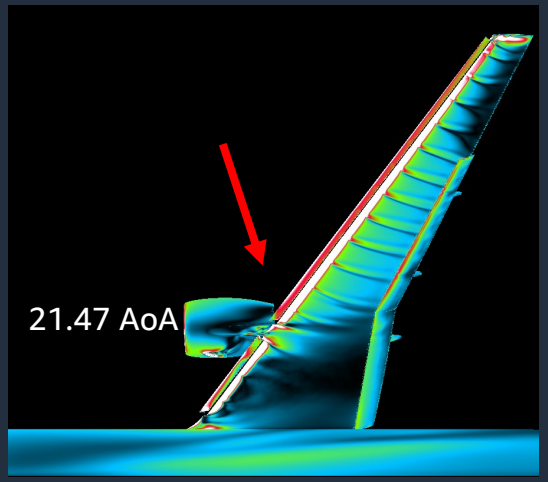
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

Case 2a – free-air



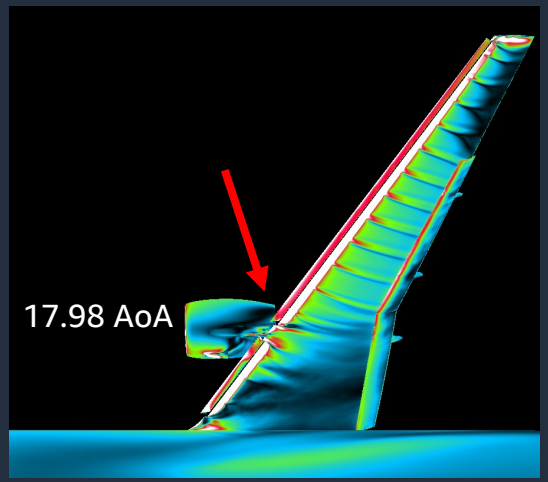


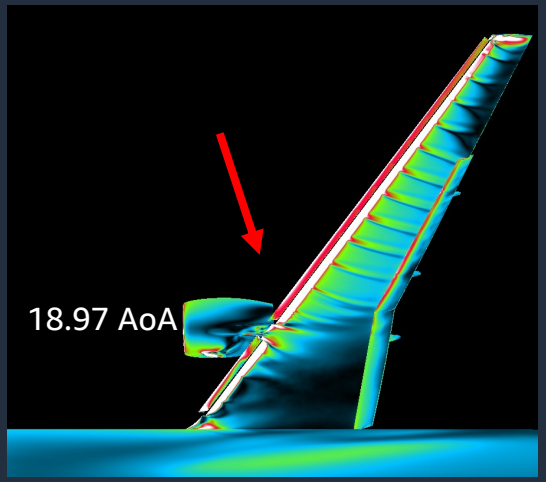


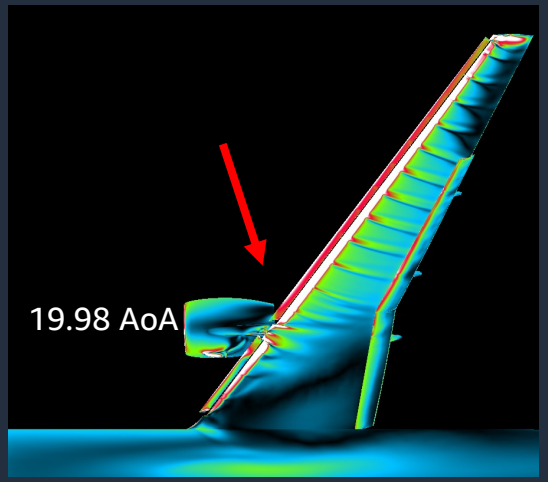
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

Case 2b – WT





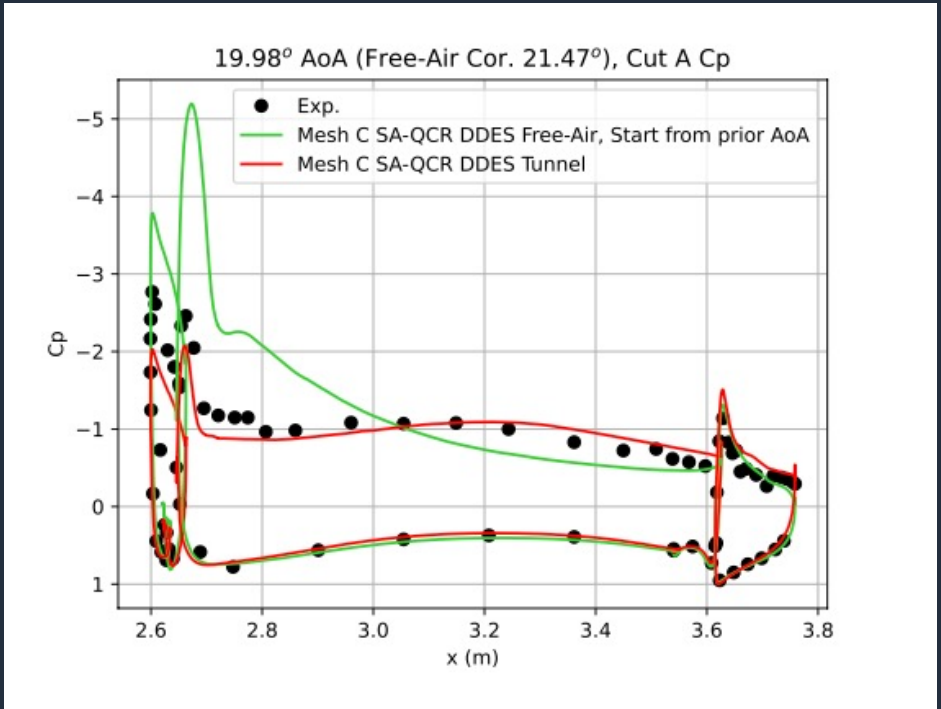


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

- Modeling of the wind-tunnel is important to correctly correlate against high-lift aircraft experiments that were performed using wall-mounted half-models.

- Strongest influence is on the inboard separation (cut A) , where in free-air RANS does better than HLRES – likely an influence of the tunnel-wall

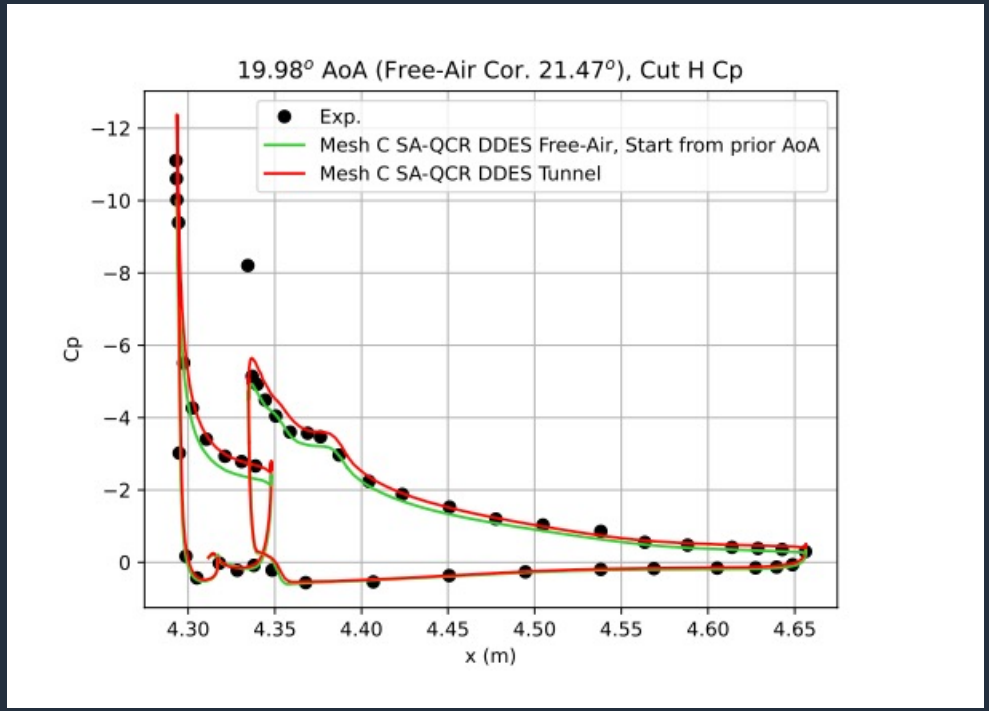


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

- Modeling of the wind-tunnel is important to correctly correlate against high-lift aircraft experiments that were performed using wall-mounted half-models.

- Strongest influence is on the inboard separation (cut A) , where in free-air RANS does better than HLRES – likely an influence of the tunnel-wall

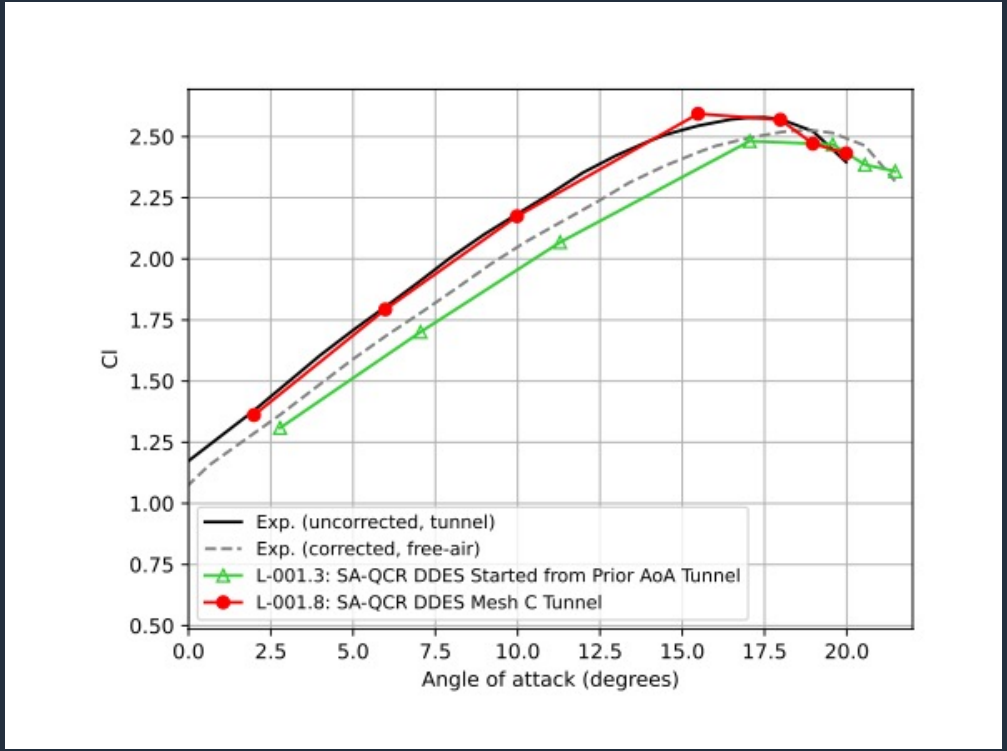


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

- Modeling of the wind-tunnel is important to correctly correlate against high-lift aircraft experiments that were performed using wall-mounted half-models.

- Strongest influence is on the inboard separation (cut A) , where in free-air RANS does better than HLRES – likely an influence of the tunnel-wall

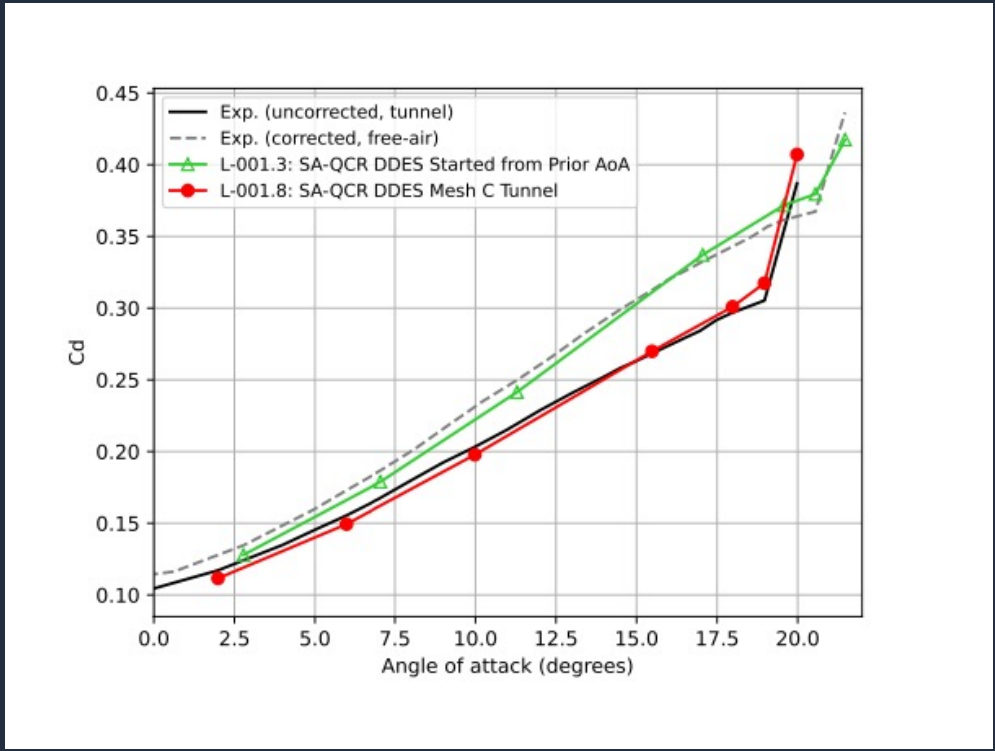


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

- Modeling of the wind-tunnel is important to correctly correlate against high-lift aircraft experiments that were performed using wall-mounted half-models.

- Strongest influence is on the inboard separation (cut A) , where in free-air RANS does better than HLRES – likely an influence of the tunnel-wall

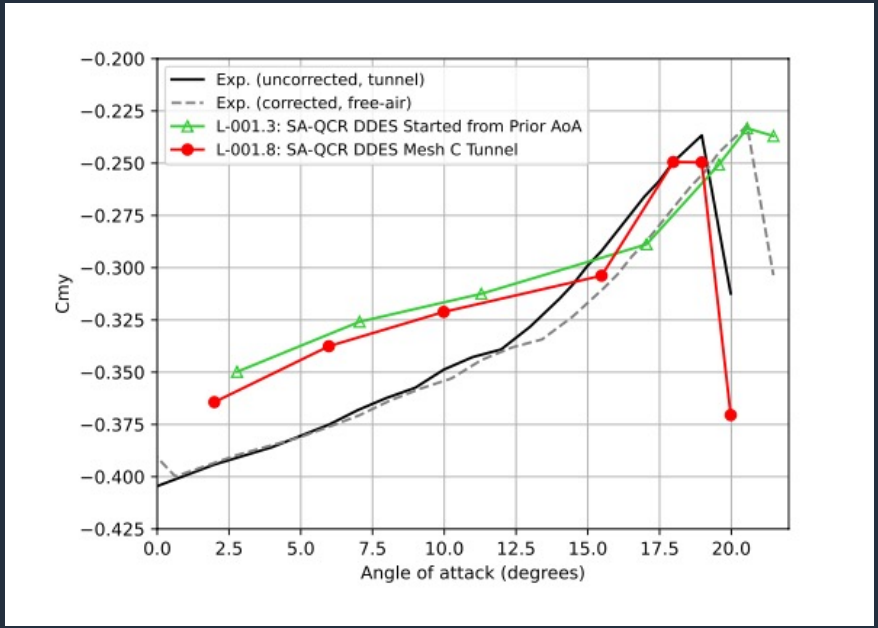


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Wind-tunnel modelling

- Modeling of the wind-tunnel is important to correctly correlate against high-lift aircraft experiments that were performed using wall-mounted half-models.

- Strongest influence is on the inboard separation (cut A) , where in free-air RANS does better than HLRES – likely an influence of the tunnel-wall



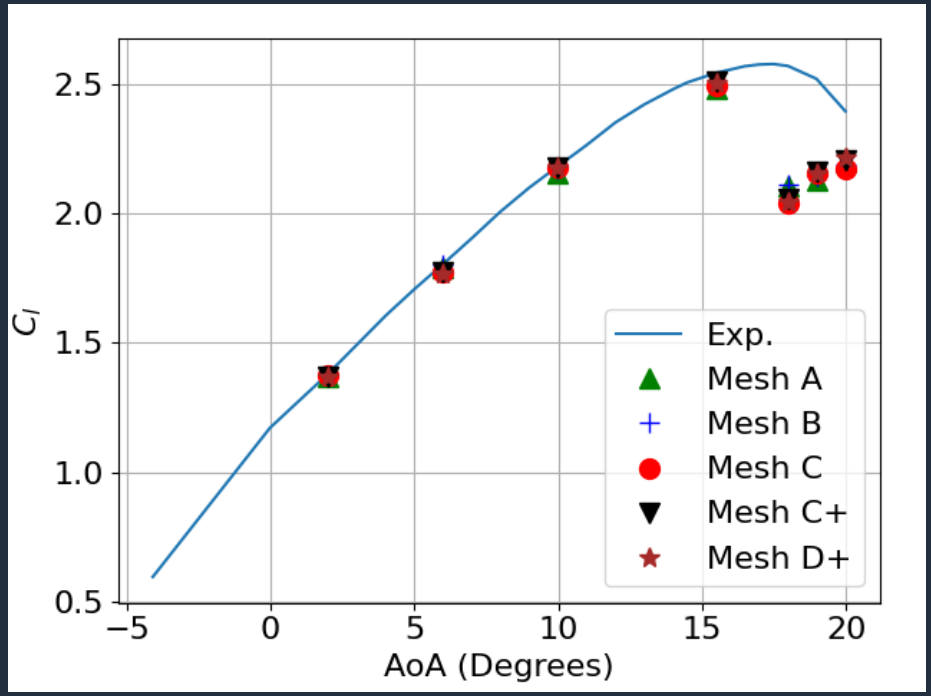
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution – KQ2

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (RANS)

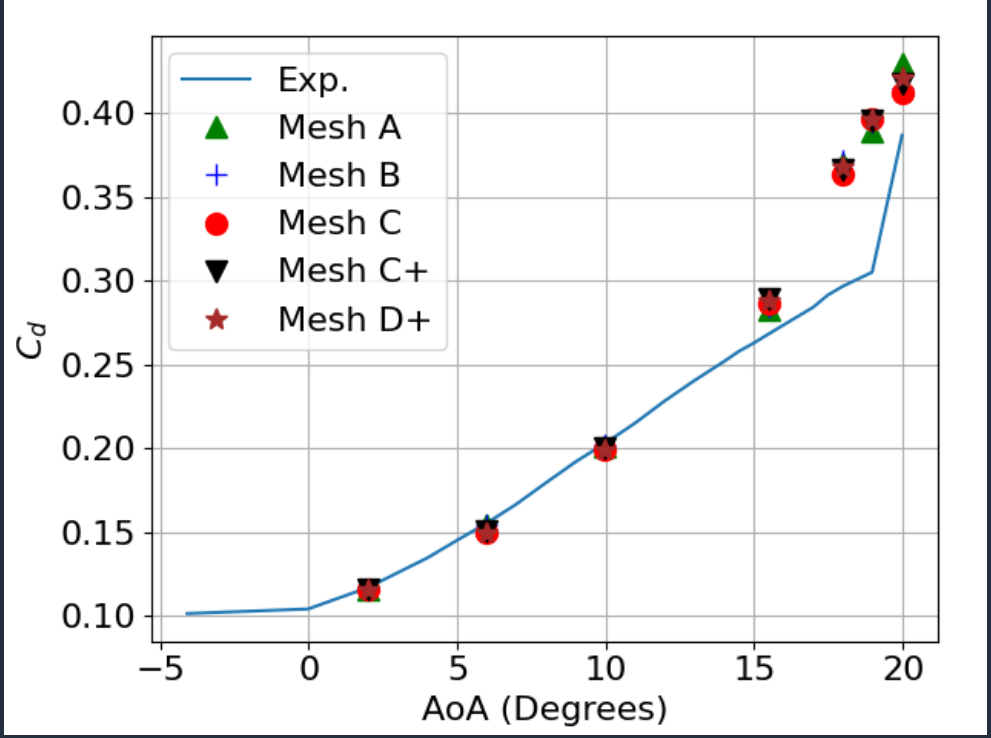
- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- Little difference between Mesh A and Mesh D+ for lift and drag but moments show mesh A & B to be under-resolved. Hence Mesh C being the optimum for RANS from an accuracy/HPC compromise



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (RANS)

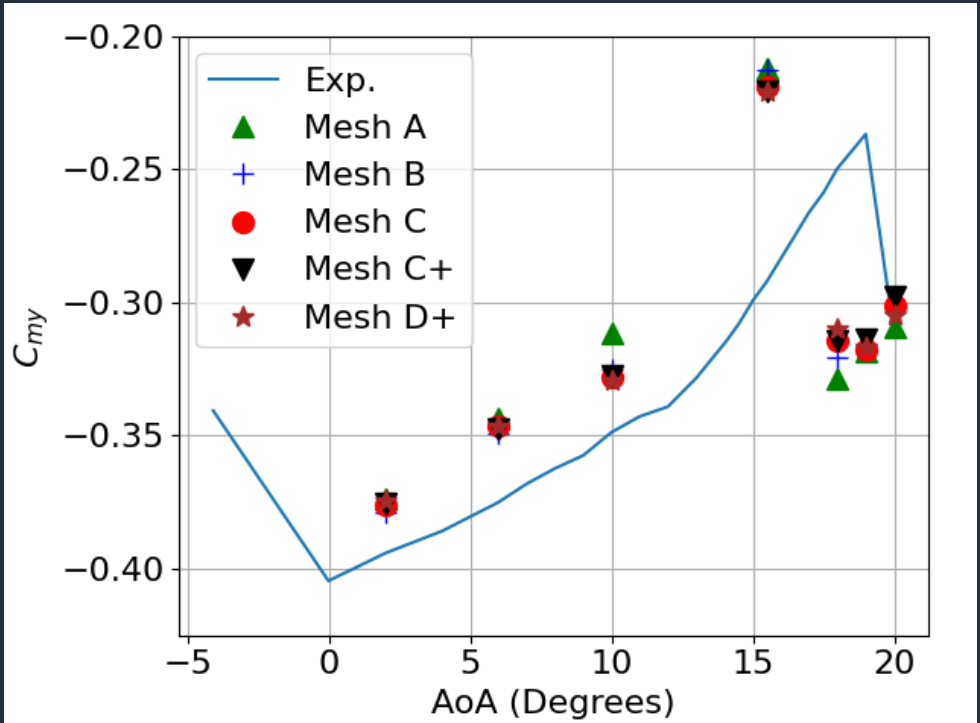
- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- Little difference between Mesh A and Mesh D+ for lift and drag but moments show mesh A & B to be under-resolved. Hence Mesh C being the optimum for RANS from an accuracy/HPC compromise



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (RANS)

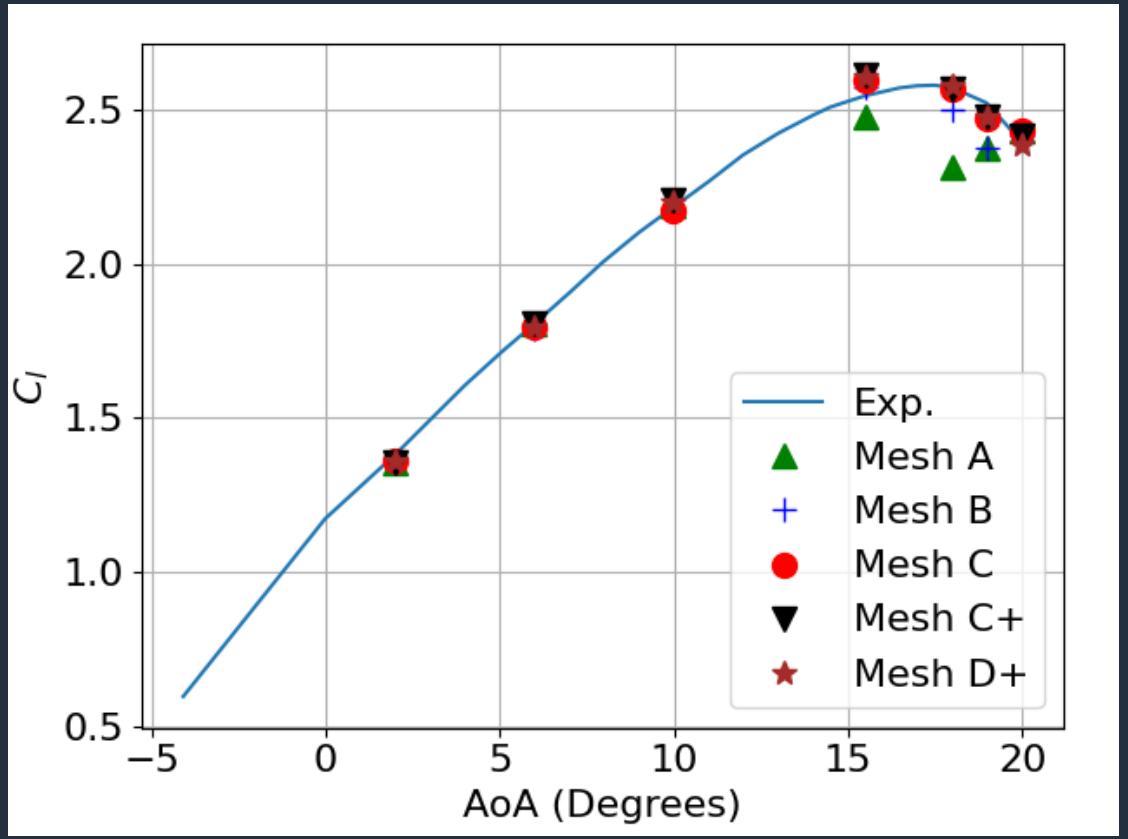
- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- Little difference between Mesh A and Mesh D+ for lift and drag but moments show mesh A & B to be under-resolved. Hence Mesh C being the optimum for RANS from an accuracy/HPC compromise



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (HRLES)

- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- HRLES shows close to mesh convergence by Mesh C for lift and drag, with Mesh C+ and D+ being very similar.
- The moment highlights greater differences but by Mesh C, the differences are relatively small, which again points to Mesh C being optimum for accuracy/HPC cost.

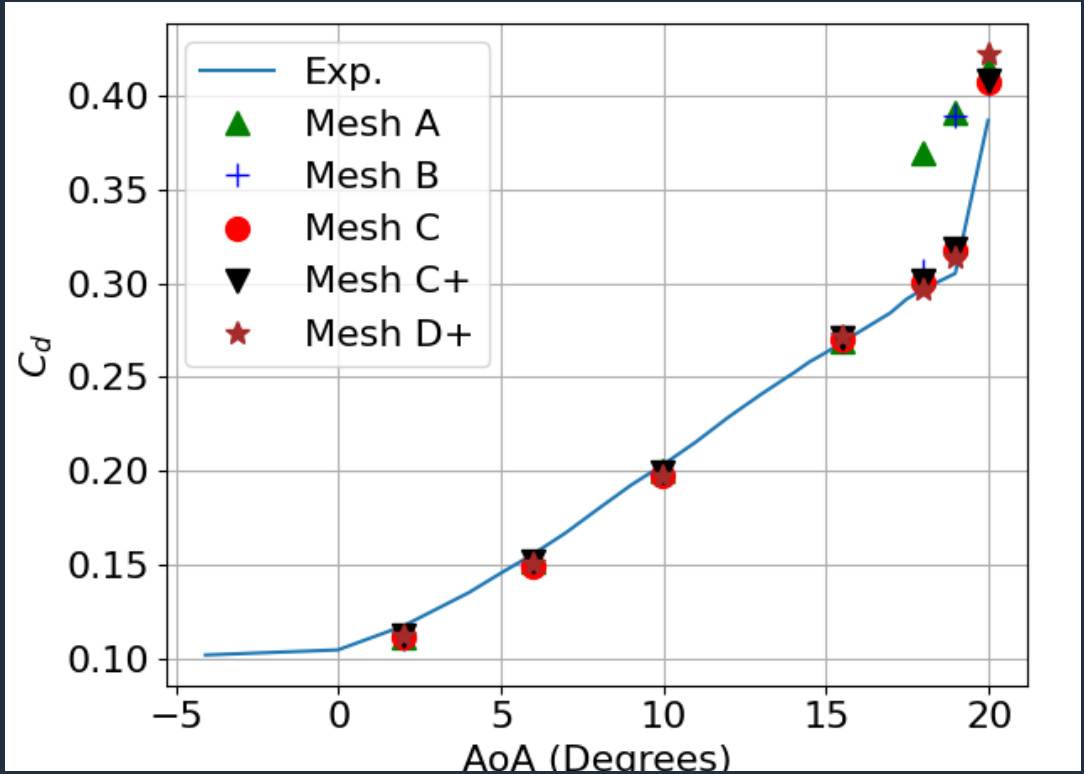


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

Mesh Resolution (HRLES)

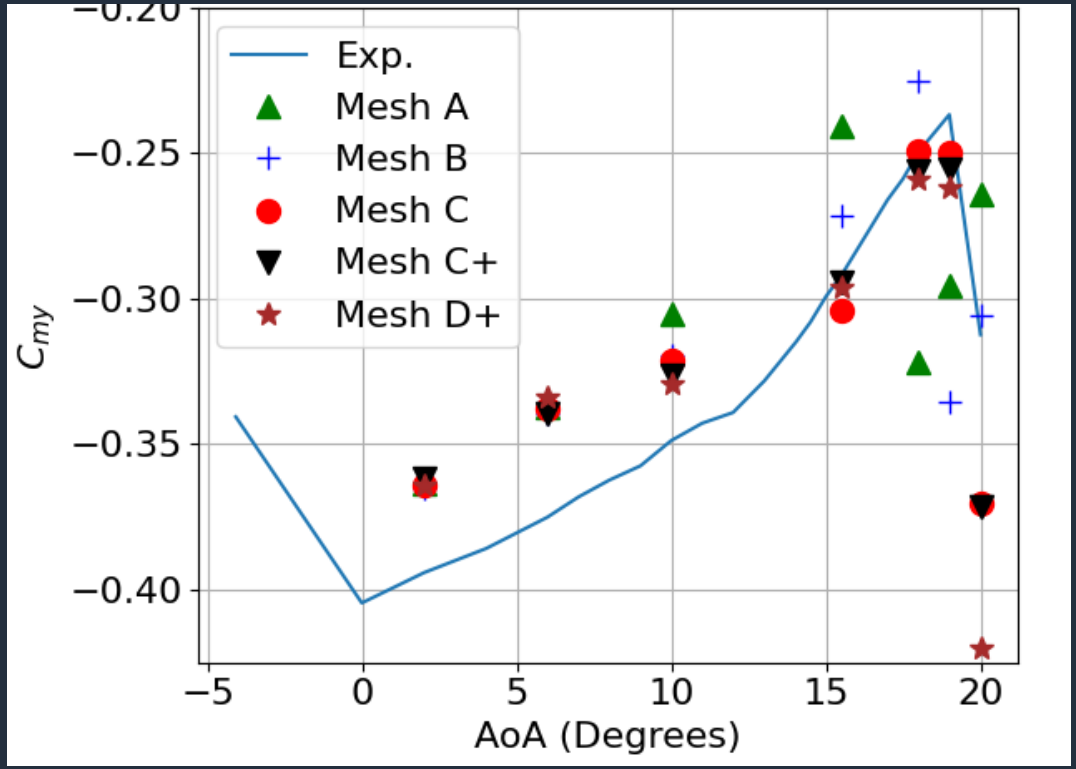
- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- HRLES shows close to mesh convergence by Mesh C for lift and drag, with Mesh C+ and D+ being very similar.
- The moment highlights greater differences but by Mesh C, the differences are relatively small, which again points to Mesh C being optimum for accuracy/HPC cost.



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (HRLES)

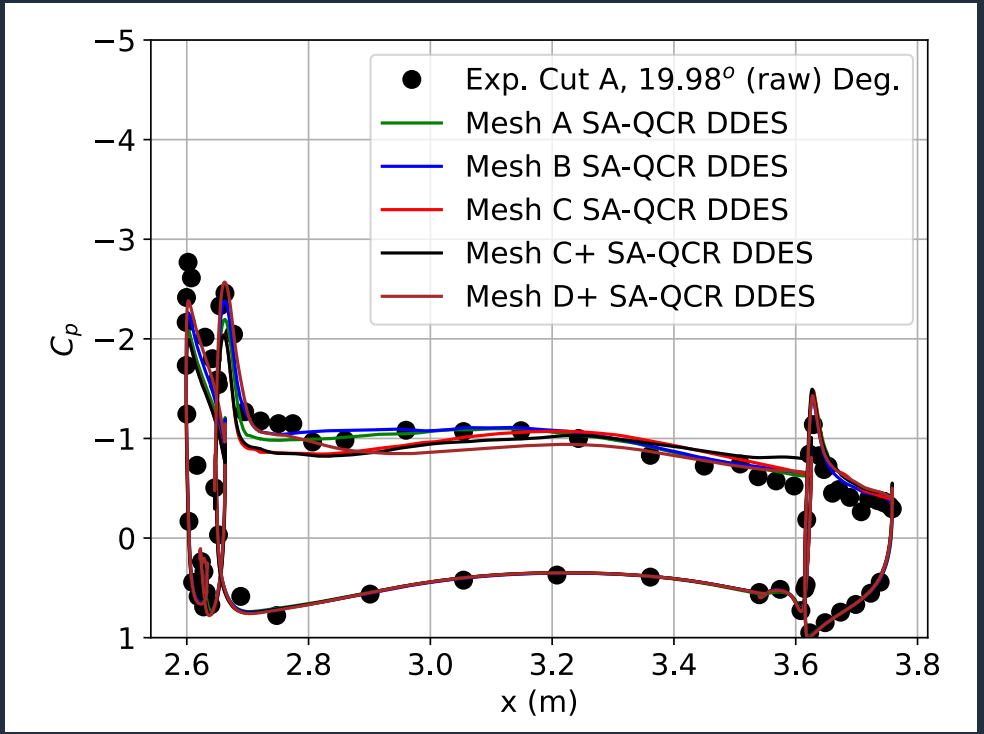
- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- HRLES shows close to mesh convergence by Mesh C for lift and drag, with Mesh C+ and D+ being very similar.
- The moment highlights greater differences but by Mesh C, the differences are relatively small, which again points to Mesh C being optimum for accuracy/HPC cost.



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (HRLES)

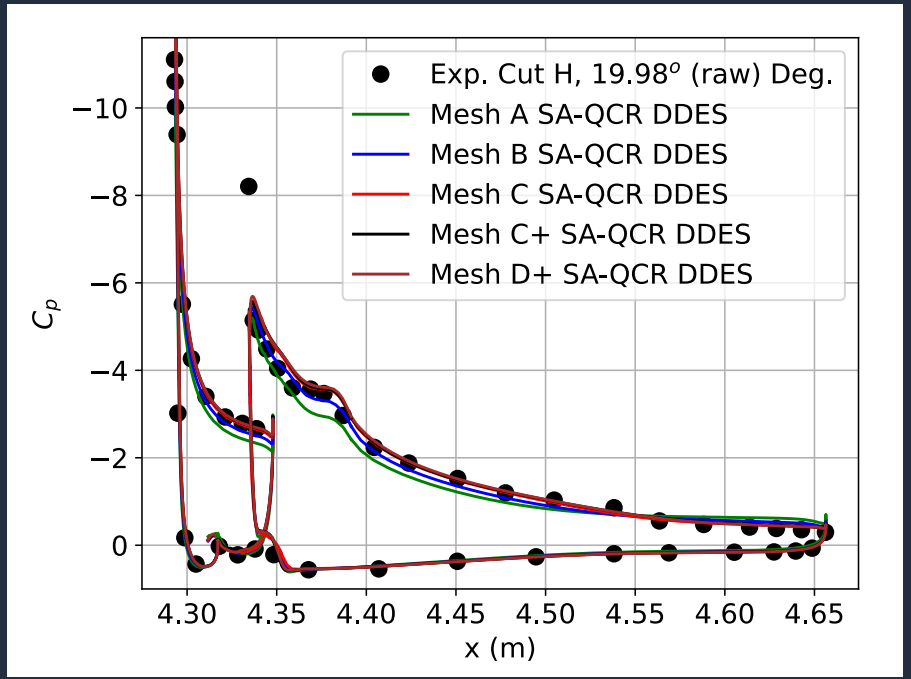
- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- RANS looks largely converged by Mesh C, whereas HRLES still requires even finer meshes to confirm mesh convergence



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Mesh Resolution (HRLES)

- Adequate mesh resolution must be supplied in the unsteady/separated-flow regions. HRLES requirements are different than RANS
- HRLES shows close to mesh convergence by Mesh C for lift and drag, with Mesh C+ and D+ being very similar.
- The moment highlights greater differences but by Mesh C, the differences are relatively small, which again points to Mesh C being optimum for accuracy/HPC cost.



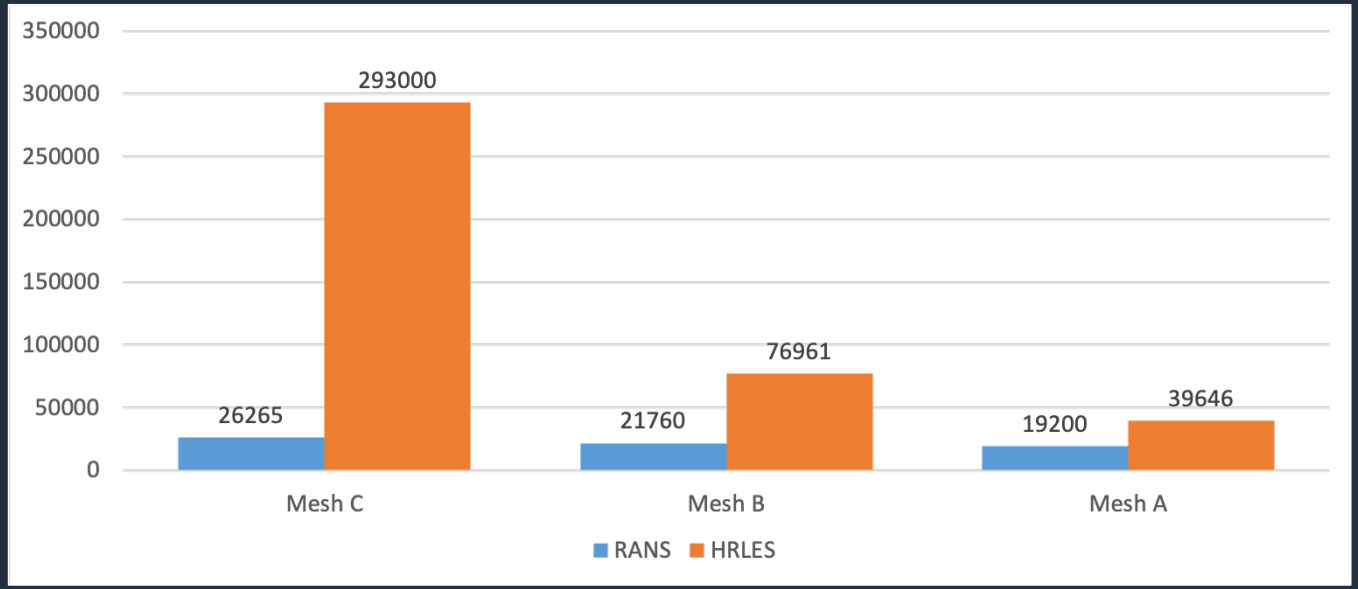
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

HPC Cost – KQ3

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

HPC cost

- Approximately 293k core-hours required for 40CTU's (need to do this x7 times) – x10 more than RANS for Mesh C
- Advances in HPC hardware and more efficient software will naturally reduce this cost over the next 3-5 years but how does this compare against other methods and participants?



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

What about other participants?

L001.1 - SA-QCR RANS: ~25,000 core hours = $750
Steady-state, unstructured 300M cells, 2nd order, implicit, AMD Milan 96c per node (hpc6a), $0.03 per core/hour

L001.3 - SA-QCR DDES: ~250,000 core hours = $7,500
40CTU, unstructured 300M cells, 2nd order, implicit AMD Milan 96c per node (hpc6a) - $0.03 per core/hour

W-020 - WMLES : ~220,000 core hours= $6,000
40CTU, structured overset 360M cells, 2nd order, explicit RK3, assumed AMD Milan hpc6a (10% faster than Rome), $0.03 per core/hour

W-021 - WMLES : ~6hr*96 Nvidia V100= $2,300
40CTU, unstructured 362M cells, 2nd order, explicit, assumed p3dn.24xlarge pricing (x8 V100), $31 per node hour

How much do GPUs help with cost and speed?

https://hiliftpw.larc.nasa.gov/Workshop4/WorkshopPresentations/07_GMGW3_HLPW4_WMLES-LB.pdf
https://hiliftpw.larc.nasa.gov/Workshop4/ParticipantTalks/Aviation2022/S2_02_summary_hrles_ashton.pdf

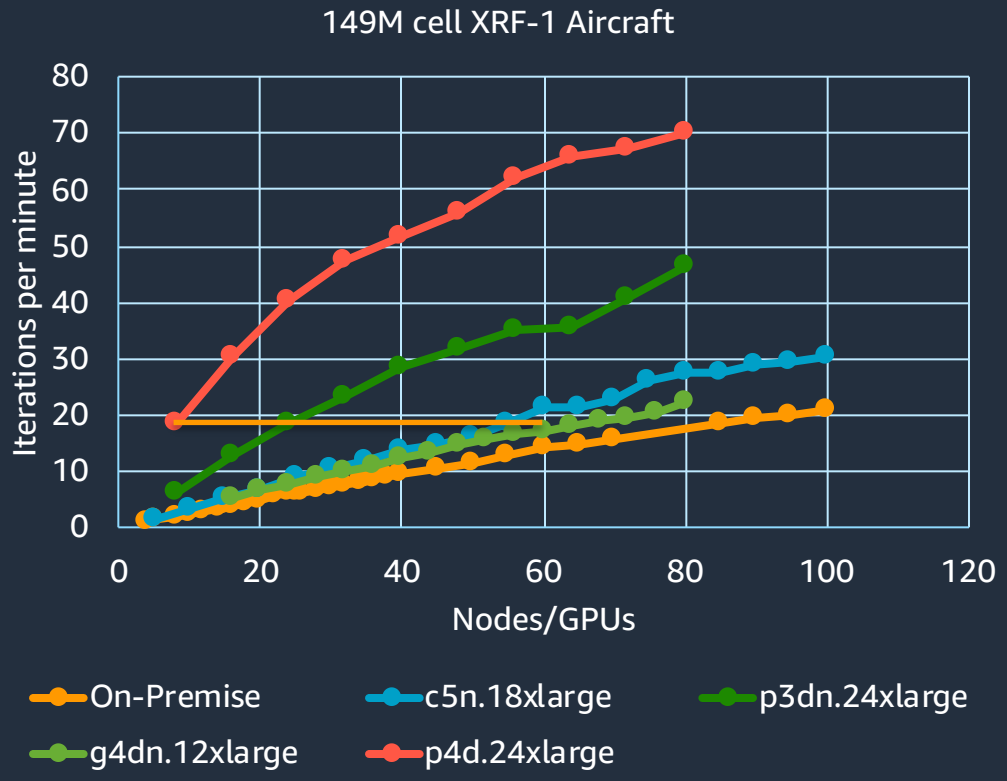
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

GPUs?

ZENOTECH
SIMULATION UNLIMITED

- Similar conclusions on previous work with Zenotech. 1 GPU node (i.e x8 A100) = ~2000 CPU cores or 1 GPU =~250-300 CPU cores
- GPU's (p4d.24xlarge, p3.24xlarge and g4dn.16xlarge Amazon EC2 instances i.e Nvidia A100/v100 and T4's) deliver faster and cheaper results
- *Appa, J, Turner, M, **Ashton N**, 'Performance of CPU and GPU HPC Architectures for off-design aircraft simulations ', AIAA Scitech 2021, DOI: 10.2514/6.2021-0141*

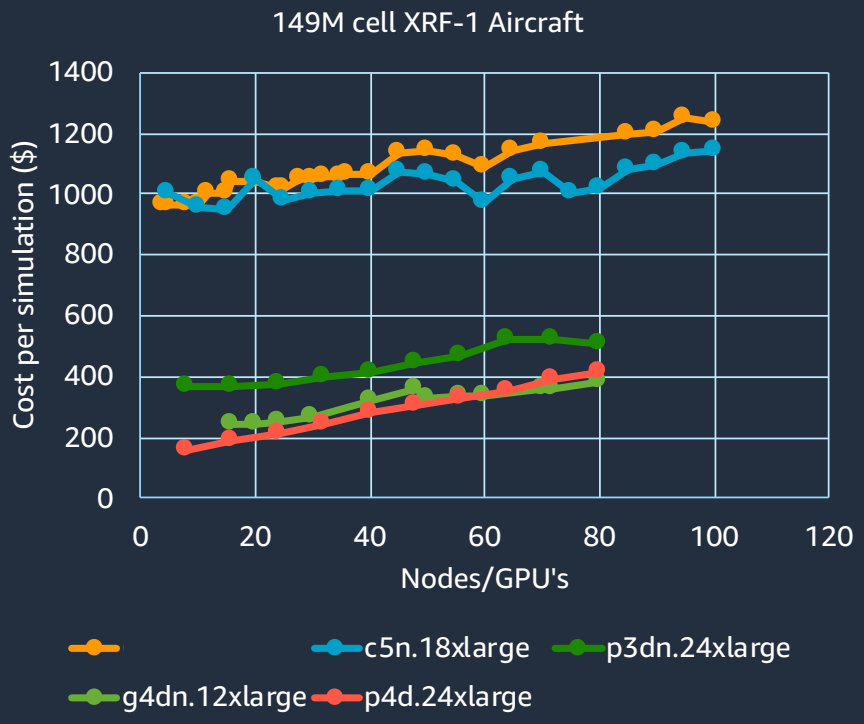


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

aws

GPUs?

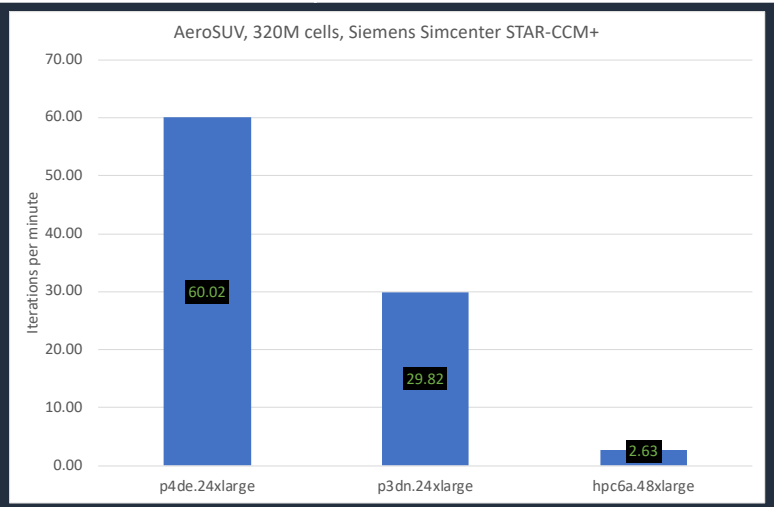
- Similar conclusions on previous work with Zenotech. 1 GPU node (i.e x8 A100) = ~2000 CPU cores or 1 GPU =~250-300 CPU cores
- GPU's (p4d.24xlarge, p3.24xlarge and g4dn.16xlarge Amazon EC2 instances i.e Nvidia A100/v100 and T4's) deliver faster and cheaper results
- *Appa, J, Turner, M, **Ashton N**, 'Performance of CPU and GPU HPC Architectures for off-design aircraft simulations ', AIAA Scitech 2021, DOI: 10.2514/6.2021-0141*

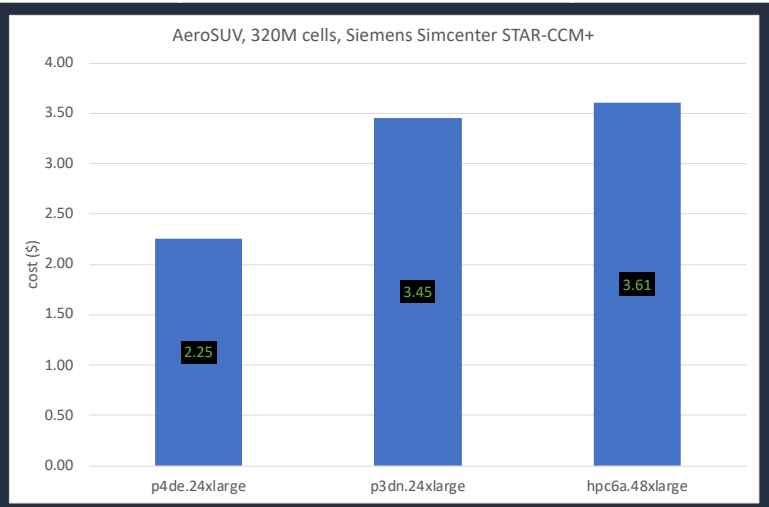


© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Example of GPUs cost/performance

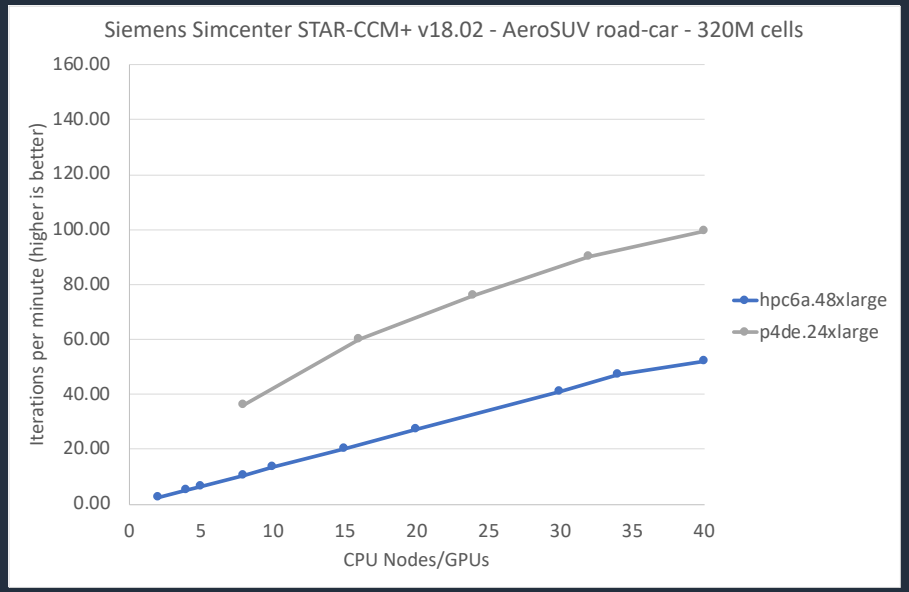
	Memory Bandwidth (MB)	Price-per-node ($)	Its/per/**min for 2 nodes**	Sim cost (5000 it) - $
Hpc6a.48xlarge (AMD Milan 96c)	~300GB/s	2.88	2.63	3.61
P3dn.24xlarge - V100 (32GB GPU memory)	~1000 GB/s (typically 8 per node) i.e 8,000 GB/s	31	29.82	3.45
P4de.24xlarge - A100 (80GB GPU memory)	~2000 GB/s (typically 8 per node) i.e 16,000 GB/s	41	60.02	2.25

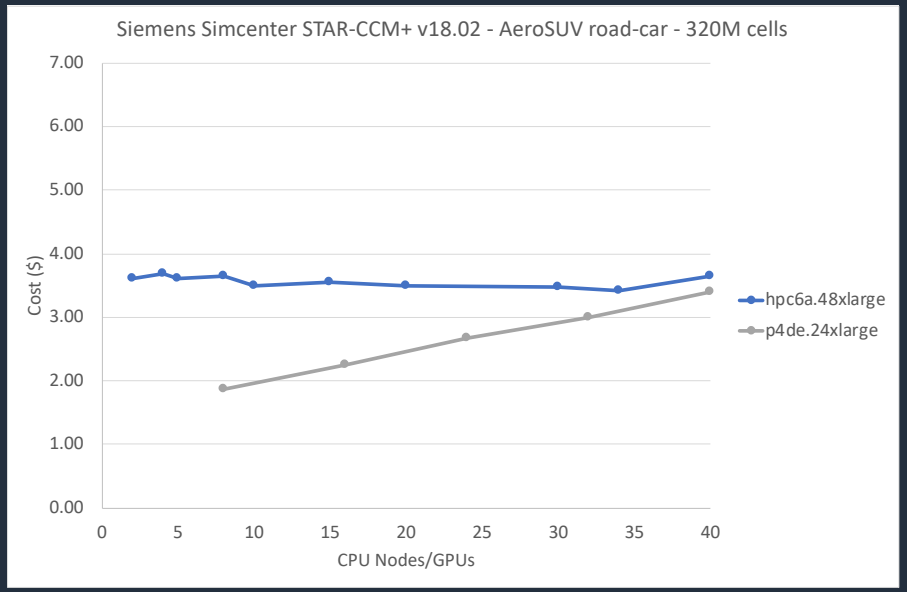




© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Example of GPU cost/performance





Keeping the GPUs busy means you can't often scale as well for a relatively small problem so cost benefit can reduce as you scale. However, different codes will show better or worse performance depending on the porting and the underlying numerical methods and the way they can use the GPU performance.

So far mainly showed Nvidia GPUs but next few years will see Intel & AMD grow their offerings in the market

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

But what about meshing and the need for wall-resolved?

- WMLES and HRLES methods both showed great promise within the HLPW as well as more generally for external aerodynamics applications where there is significant separation.
- The current debate is whether WMLES is the natural successor to HRLES methods i.e faster, more generalizable, potentially more suited to GPUs or whether high y+ HRLES approaches would potentially provide the same runtime and with advanced DDES-type closures, the same accuracy. To date, most HRLES have used y+ grids so future tests are needed.
- Assuming WMLES (or high y+ HRLES) is the current best option for external aerodynamics, then meshing becomes a key question.
- Immersed boundary methods [4] have shown great promise, particularly when linked to cartesian approaches but the inability to resolve boundary layers were often cited as a blocker. However if high y+ approaches are becoming the norm then such approaches may help to remove the often major blocker of meshing.
- Open question however onto the broader suitability or whether it will remain a tool that has great promise but only for specific applications. However, lets park this debate for now.

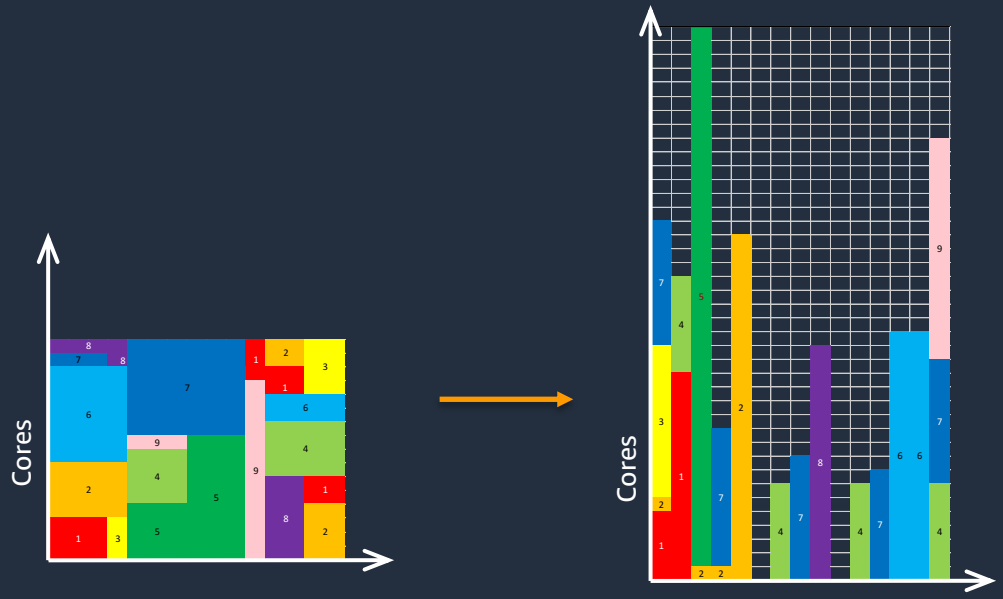
[4] C Brehm, O Browne, N Ashton " Towards a viscous wall model for immersed boundary methods" 2018 AIAA Aerospace sciences meeting, 2018

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Cloud Computing

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Main motivations for the cloud

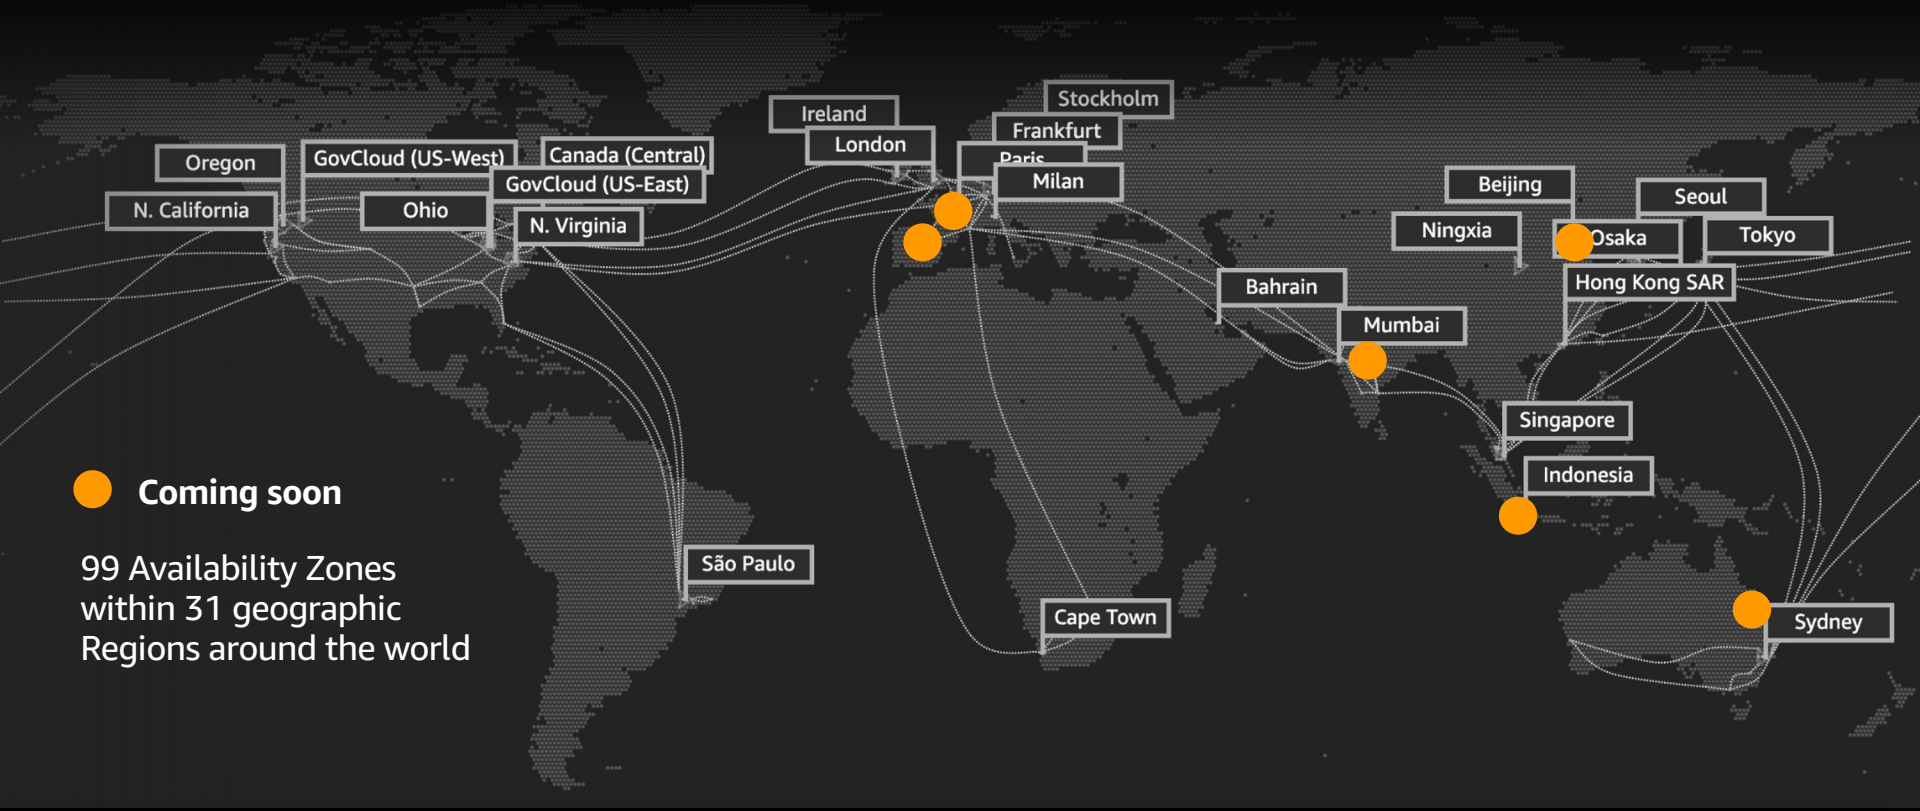


HPC

Software as a service (SaaS)

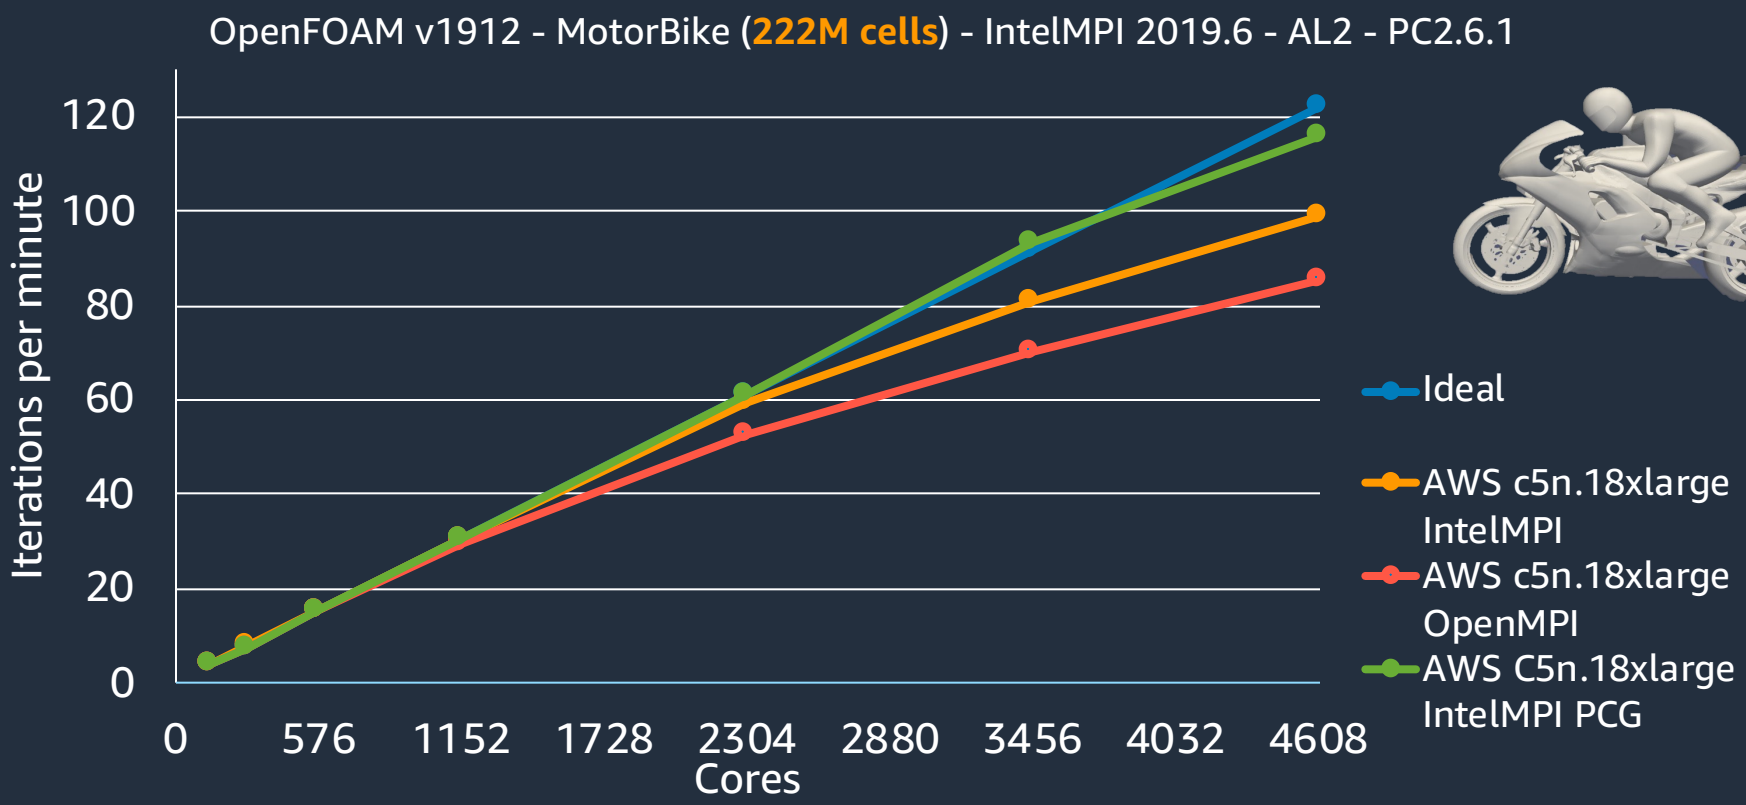
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Helps to work with an increasingly global workforce



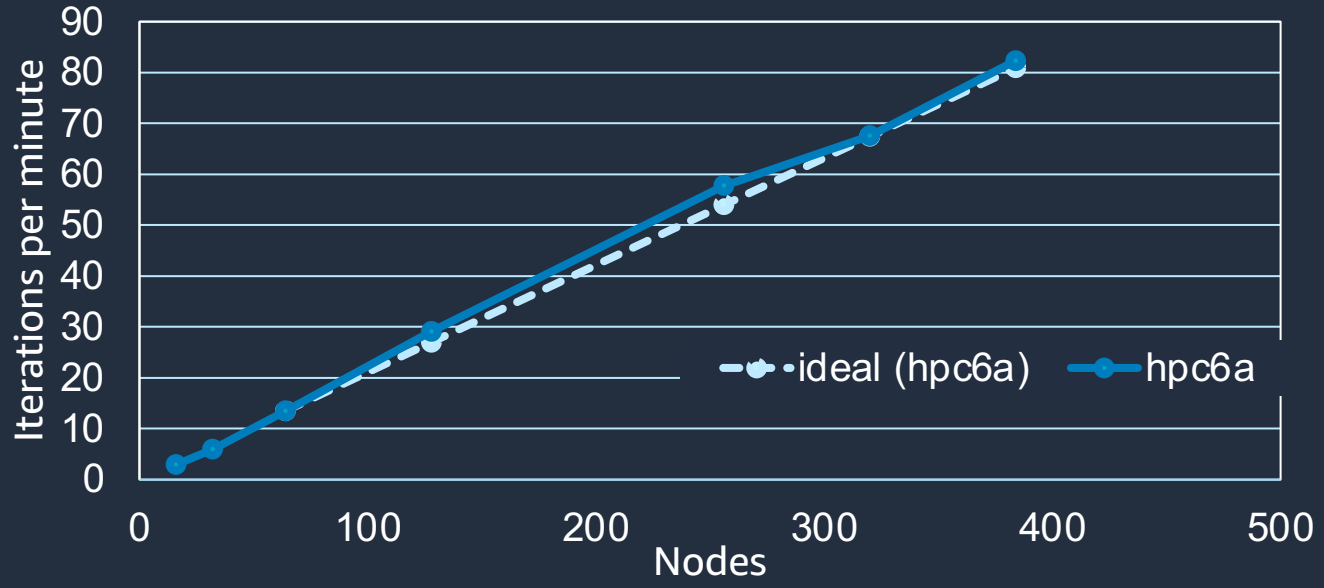
99 Availability Zones
within 31 geographic
Regions around the world

Importance of code optimization and linear solver



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Cloud is providing supercomputer level of scaling

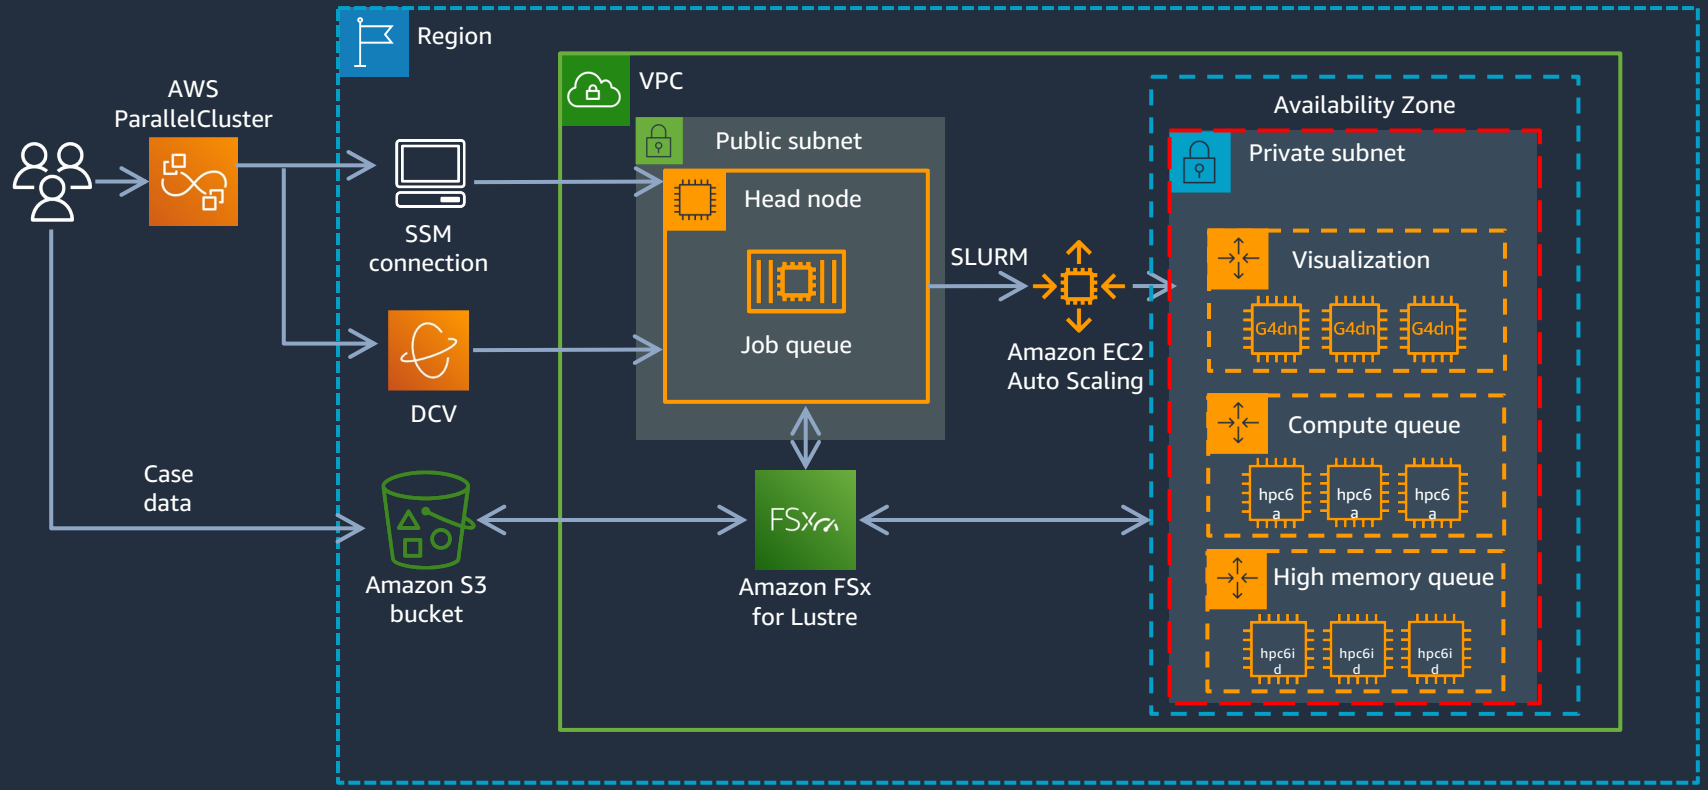


- 1.4B cell STAR-CCM+ automotive case
- IntelMPI
- Coupled implicit Solver
- Linear scaling up to 400 nodes (38,400 cores)
- Hpc6a, AMD Milan 96 core, 384GB – 100Gbit/s EFA



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Example architecture for cloud-centric CFD



© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Cloud is also becoming a way to ease access to your software

Growing range of companies offering software as a service. Started with SimScale, Dive Solutions, OnScale etc

Major ISVs such as Siemens, Cadence, ANSYS have released SaaS versions of their main packages.

New CFD startups such as FlexCompute & Luminary Cloud are actively promoting cloud as a way to access their software

ReScale, TotalCAE, UberCloud and others are providing a catalog of ISVs for people who just want to pick a code and hardware and run it.

Other customers want to build their own solution and just use services such as AWS ParalleCluster to orchestrate their jobs.

Main question for many companies is the location of the data i.e are they running in a SaaS providers account or are they running in their own account.

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Machine Learning for CFD



Extract relationships from simulation data

Accelerate HPC simulations with embedded ML models

Learn solutions to governing equations directly

Physics-enabled

Physics-informed

AI-driven Infrastructure Optimization

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

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.

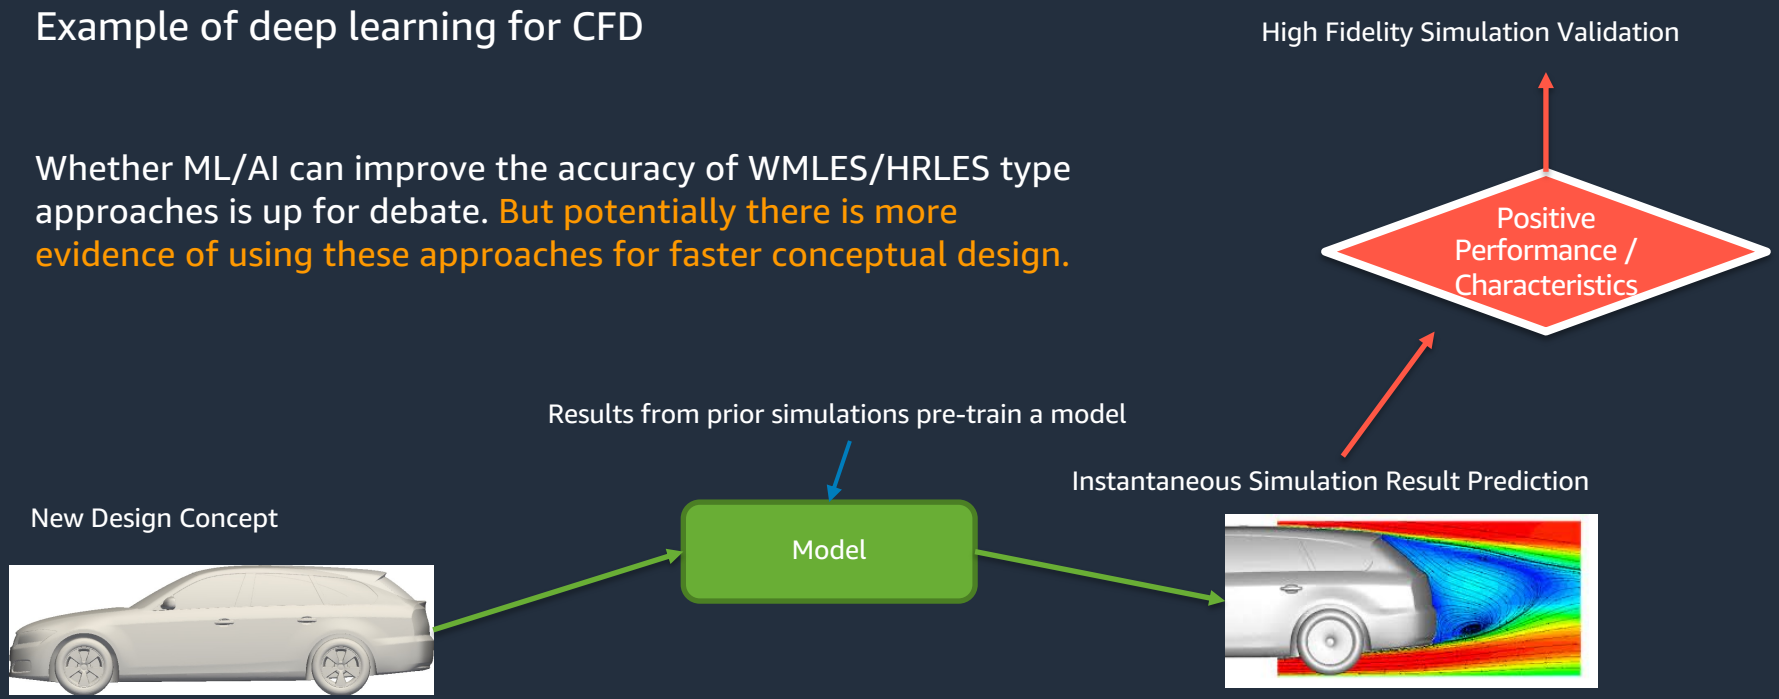
GPU's for training models but also we want to train models using more than CFD data e.g sensor data from physical tests etc

Huge increase in ML services on the cloud which CFD can take advantage of i.e AWS SageMaker and highly-tuned pyTorch, TensorFlow implementations.

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Example of deep learning for CFD

Whether ML/AI can improve the accuracy of WMLES/HRLES type approaches is up for debate. But potentially there is more evidence of using these approaches for faster conceptual design.



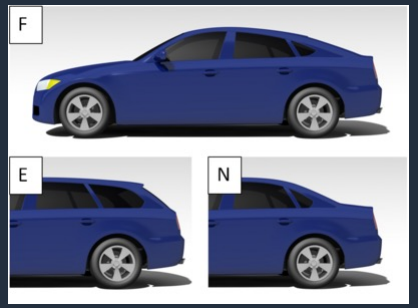
© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Example application

Goal: Predict Drag Coefficient C_d for unseen geometries.

Dataset: Open source automotive DrivAer Model dataset.

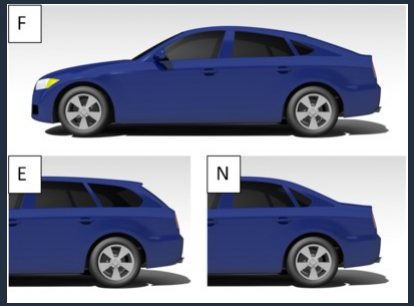
- Different tops (e.g. Fastback, estate, notchback), wheels, underbodies.
- 20 Geometries Training – 3 Geometries Validation.
- All run on AWS with high-fidelity transient DDES (~300M cells)



https://www.epc.ed.tum.de/en/aer/research-groups/automotive/drivaer/

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Example application



https://www.epc.ed.tum.de/en/aer/research-groups/automotive/drivaer/

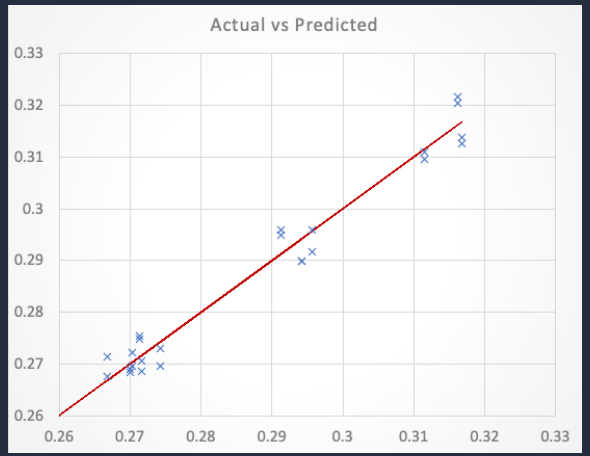
Goal: Predict Drag Coefficient **Cd** for unseen geometries.

Dataset: Open source automotive DrivAer Model dataset.

Model: CNN Regression algorithm based on PointNet.

Results:

- *Mean Absolute Error:* 0.00461
- *Mean Absolute Percentage Error :* **1.72%**



Training Time: 20 minutes

Inference: <5 seconds

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

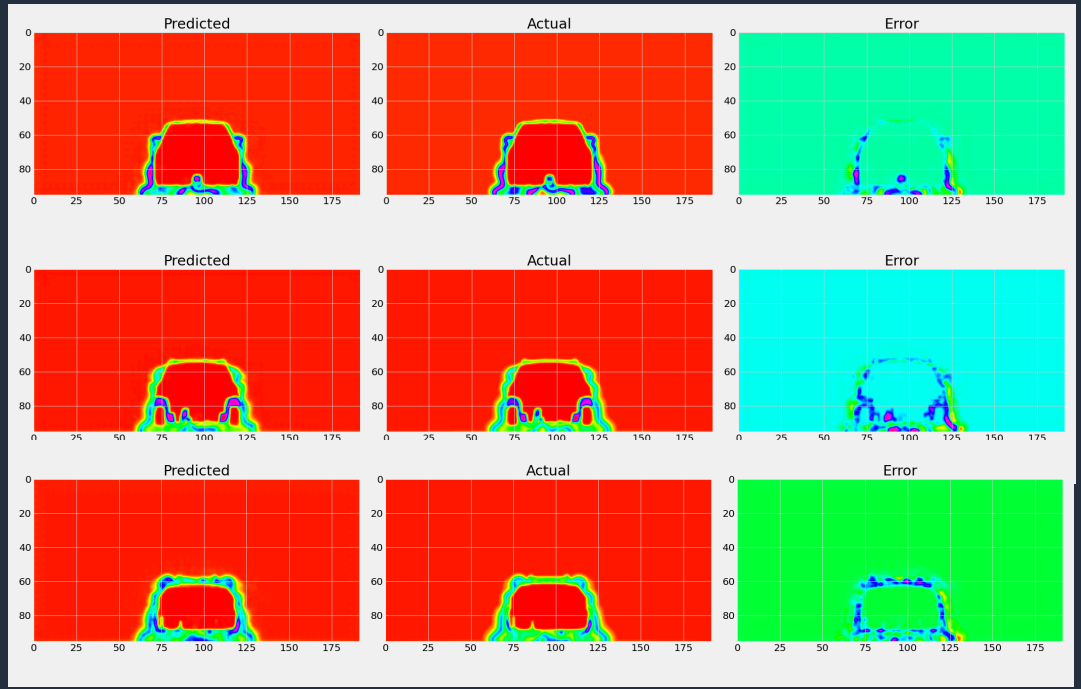
aws

Higher Dimensional Field Regression

Training Time: 20 minutes

Inference: <5 seconds

- Results: Model was trained to regress CpT Slices (192x96) longitudinally across vehicle.



Will this be accurate enough is the debate. This is only a quick academic exercise but shows there is some potential.

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Summary

- Accuracy, cost and runtime are keen drivers for CFD. High-Fidelity is a must but runtime and cost are important considerations compared to wind-tunnel and physical tests
- Success of HRLES/WMLES approaches are also based upon best-practice implementation i.e mesh resolution, initialization, wind-tunnel modelling for greater accuracy
- GPU-porting is a must to bring high-fidelity CFD cost down
- Cloud model can help to broaden HPC access to companies/researchers and allow testing on latest hardware.
- AI/ML may be able to fill in the gap of lower-fidelity methods by using high-fidelity to train a model but still early days

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Thank you!

neashton@amazon.com

© 2023, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark