SYNOPSYS®

Advanced Modeling & Simulation (AMS) Seminar Series
NASA Ames Research Center, March 19th, 2026

Advancements and Challenges in Aerodynamic Simulations in the Age of GPU Computing

Florian Menter

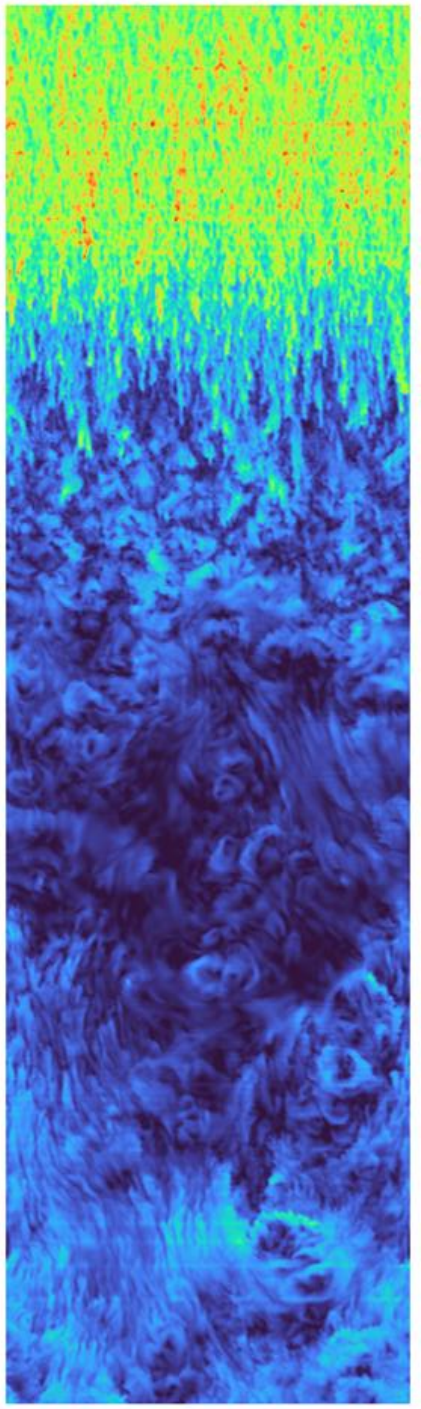
Fellow

Synopsys : Andreas Hüppe, David Flad, Katia Guseva, Rachid Malk, Ashwini Dalvi, Cristhian Aliaga...

ESSS Team in Brazil: Rodrigo Peralta, Alex J. Elias, Arthur M. Farias, Daniel Graebin

Ansys/Synopsys PBNS GPU Solver

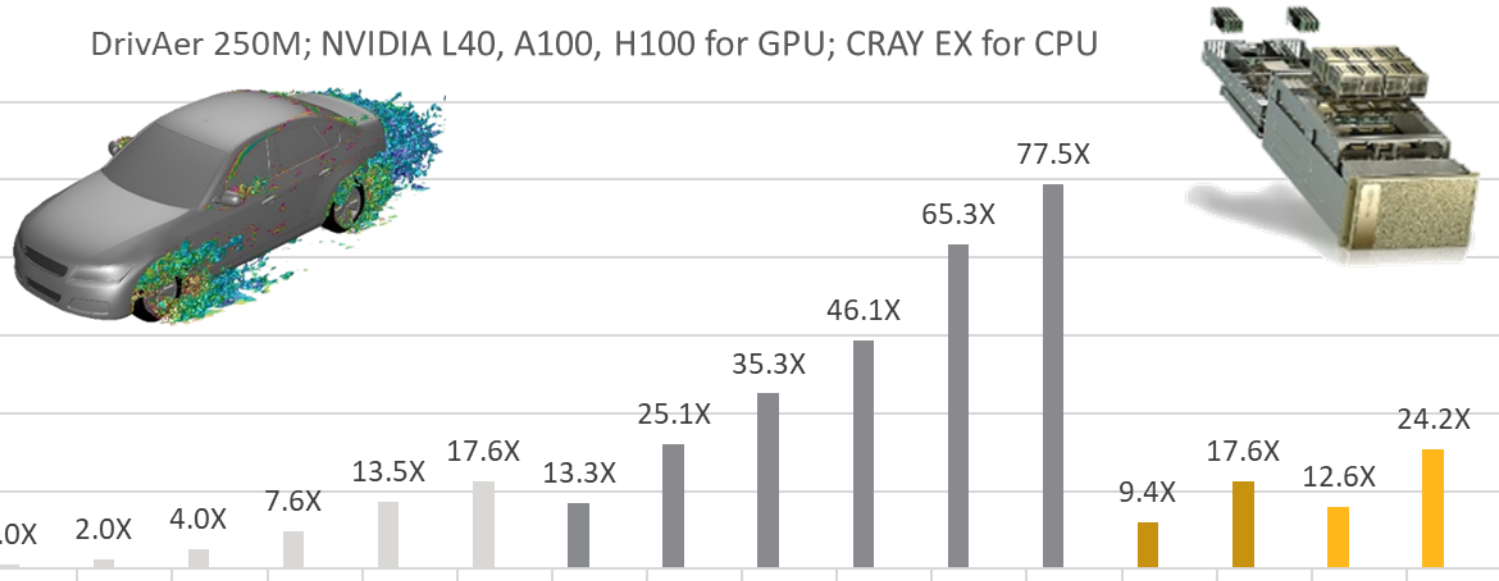
- General:
  - New solver (no off-loading of existing solver loops)
  - 2nd order cell centred FV
  - 2nd order implicit time stepping (Gear's method)
  - Pressure-base discretization using SIMPLEC
- LES runs:
  - Skew symmetric formulation (central reconstruction)
  - High under-relaxation (~1.0) and low iteration count
  - Hybrids RANS-LES (SBES), Wall function LES and WMLES
- Body Fitted:
  - Best representation of geometry and internal interfaces and thin gaps
  - Mass conservation
  - Possibility for high aspect ratio grids (laminar-turbulent transition etc.)
  - Optimal for surface interactions
    - FSI, Icing, Ablation, Reflections, ...
  - Modern meshing techniques reduce meshing cost drastically
- Implicit time integration:
  - Avoid very small time steps on 'small' grid cells or local high speed flow
  - Large time step for initial conditions



Native GPU implementation shows astounding performance gains

- CPU
  - CRAY EX
  - AMD EPYC 7763 64-Core
  - 128 cores per node
- GPU
  - Cluster 1
    - NVIDIA L40, 4 GPUs per node, Infiniband
  - Cluster 2
    - NVIDIA A100 80GB, 8 GPUs per node
  - Cluster 3
    - NVIDIA H100 80GB, 8 GPUs per node
- With 64 L40 GPUs
  - 0.05 seconds per iteration
  - 2 iterations per step
  - 250 million cells
  - 36,000 time steps in 60 minutes

Ansys / FLUENT NVIDIA



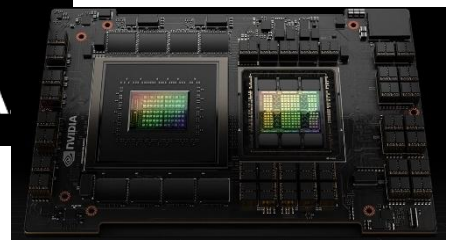
	512	1024	2048	4096	8192	12288	8 L40 GPUs	16 L40 GPUs	24 L40 GPUs	32 L40 GPUs	48 L40 GPUs	64 L40 GPUs	4 A100 GPUs	8 A100 GPUs	4 H100 GPUs	8 H100 GPUs
MIUPS	64	125	252	483	863	1124	848	1598	2254	2939	4164	4943	602	1124	802	1541
Parallel efficiency	100%	98%	99%	95%	85%	73%	100%	94%	89%	87%	82%	73%	100%	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

A single H100 GPU ≈ 1,500 CPU cores

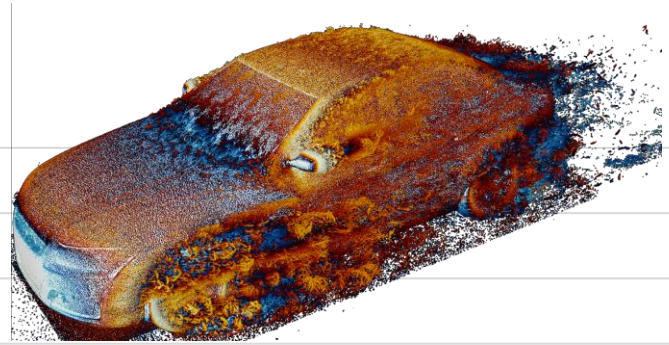
© 2026 Synopsys, Inc.

Compute Capacity @ TACC on Nvidia

- Simulation on Vista system at Texas Advanced Computing Center (TACC) with Nvidia GH200







	32 GPU	48 GPU	64 GPU	96 GPU	128 GPU	160 GPU	320 GPU
MIUPS	8678.56	12114.74	15983.05	21634.11	23784.45	27739.13	43638.61
Parallel efficiency	100%	93%	92%	83%	69%	64%	50%
Time per iteration	0.277	0.198	0.150	0.111	0.101	0.087	0.055

- Astounding performance
- Still scaling to at least 320 GPU
  - DrivAer case with **2.4B elements**
  - **180.000 Δt**
  - Becoming ‘too small’ → efficiency drops

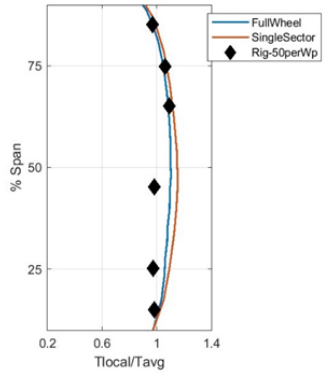
From 29 days (2048 cores) → 6 hrs (320 GPUs)

© 2026 Synopsys, Inc.

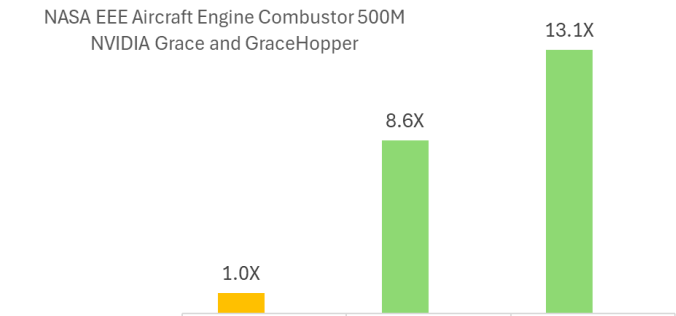
Full Annular Combustor Study with 500M Cells

- Energy Efficient Engine (EEE) Combustor
- 500M Rapid Octree Mesh
- Large Eddy Simulation with optimized numerics
- Flamelet generated manifold (FGM) combustion model
- Partition tracking for liquid fuel spray with SSD Break-up model





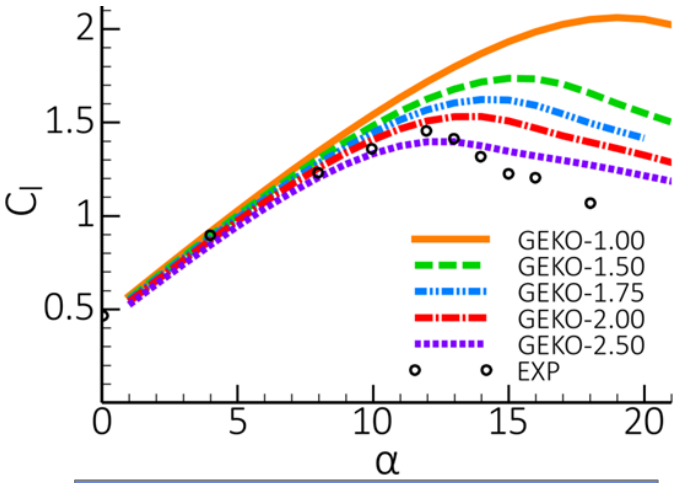
- Ansys Fluent ARM64 @ TACC
  - NVIDIA Grace CPU Superchip
  - NVIDIA GH200 Grace Hopper Superchip
- Performance
  - 92 hours to solve on 2304 CPU cores (16 nodes) ~ 1GPU=1240CPUs
  - **10.5 hours to solve on 16 Grace Hopper GPUsfor 50,000 time steps**

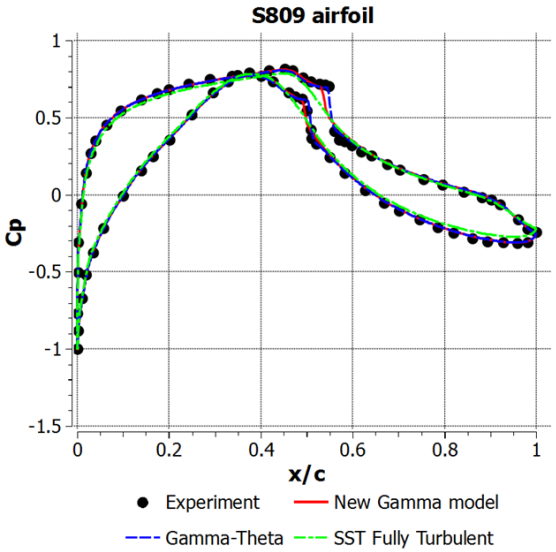


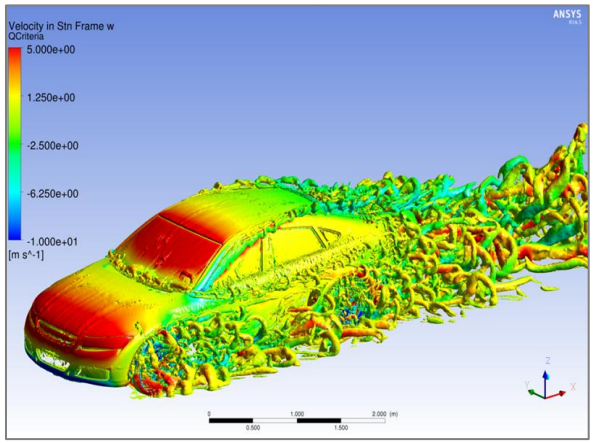
© 2026 Synopsys, Inc.

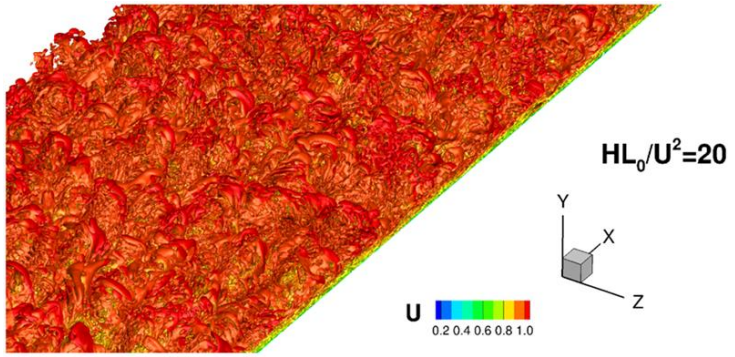
Turbulence Models

- RANS Modeling
  - Generalized k-ω (GEKO) model
  - Algebraic laminar turbulent transition model
- Hybrid RANS-LES model
  - Stress-Blended Eddy Simulation (SBES) model
- LES Models
  - Wall Function LES (WFLES)
  - Wall-Modeled LES (WMLES)







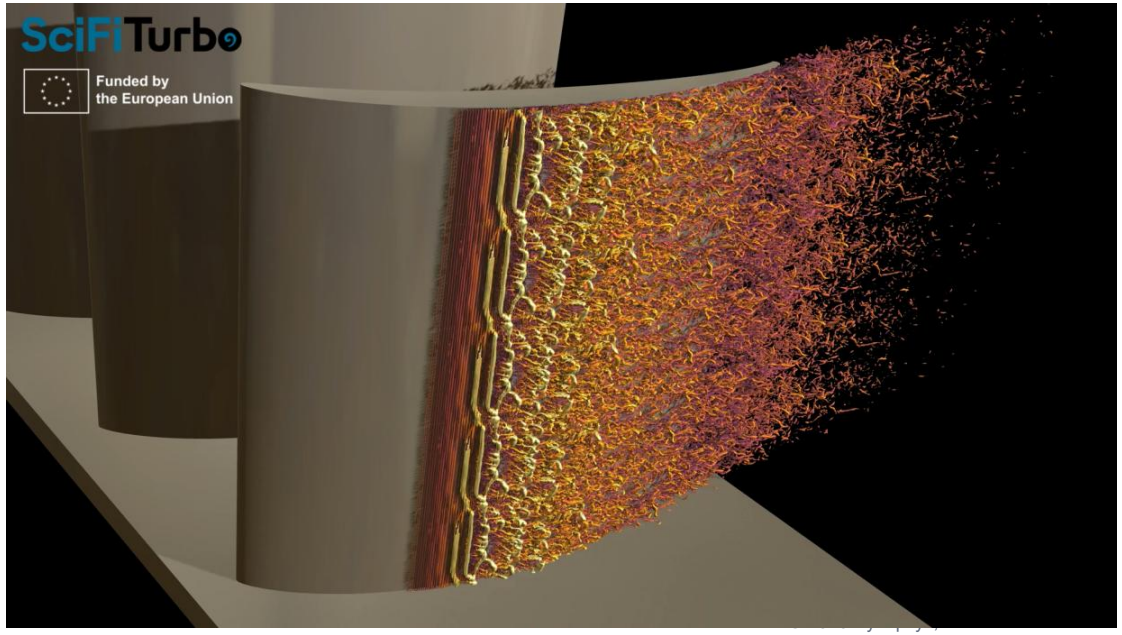


© 2026 Synopsys, Inc.

Scale Resolving Simulations (SRS) of Boundary Layer Flows

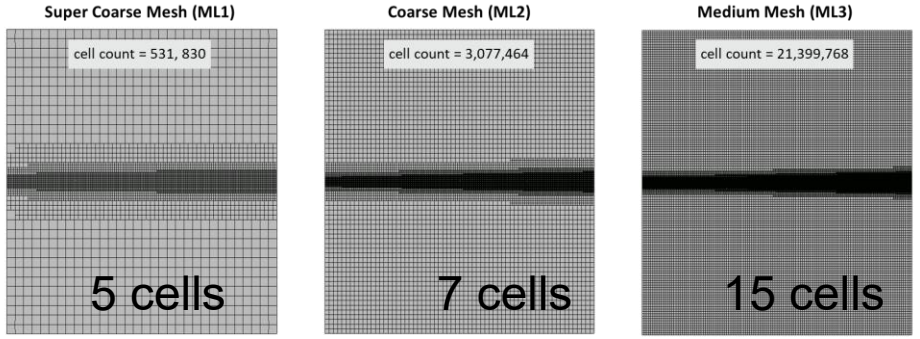
- LES in industrial flows was mostly used for free shear flows
  - Jets, Mixers, Combustion chambers...
- LES of boundary layers only industrially feasibly with GPU clusters
- However Wall-Resolved LES (WRLES) still too expensive
- Different approaches to bridge the resolution gap near the wall needed

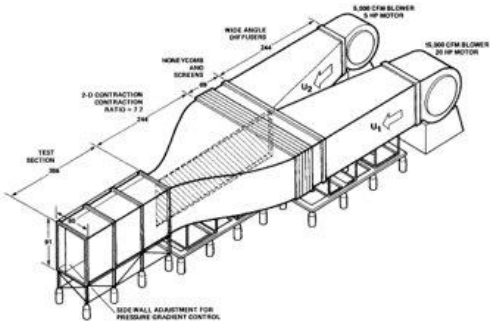


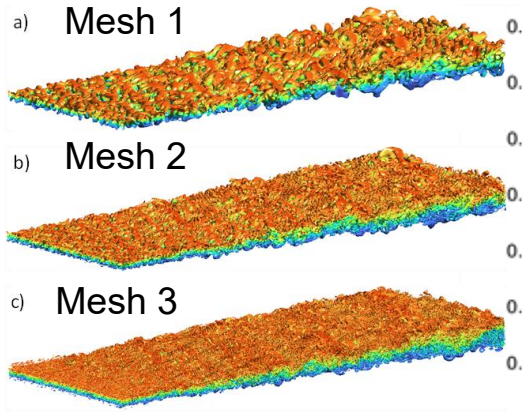


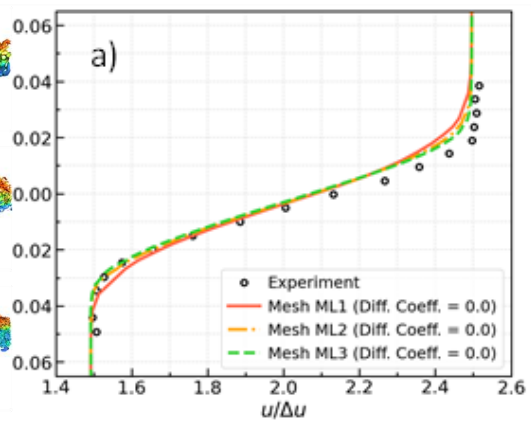
Example: Mixing Layer

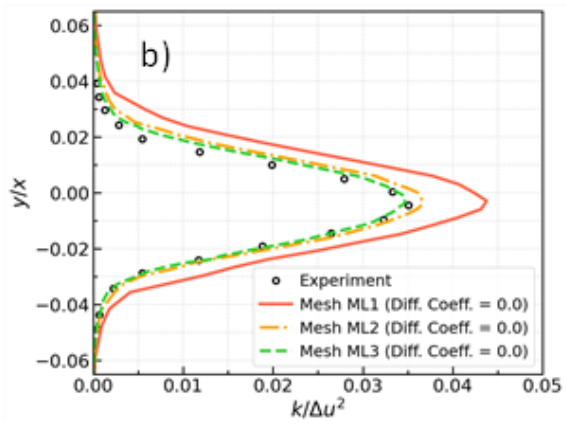
- 3 meshes
- Even with only 5 cells across the inlet mixing layer the velocity profile is captured well
- With increasing cells count the turbulent kinetic energy converges quickly
- The maximum CFL number can be increased
  - CFL=2 very small error increase
  - CFL=5 still reasonable accuracy in velocity profile

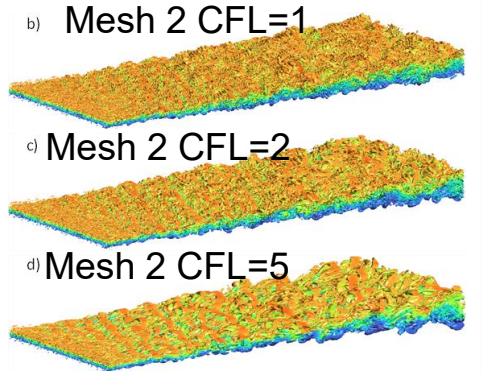


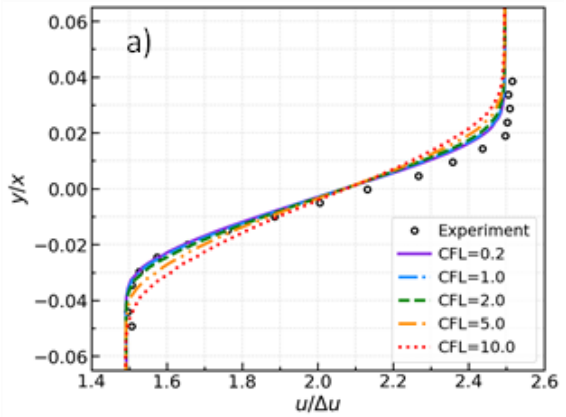


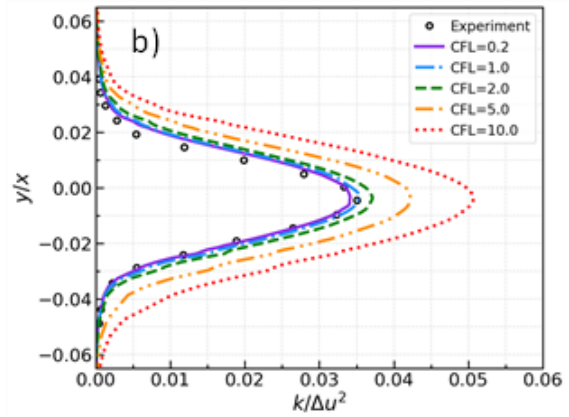












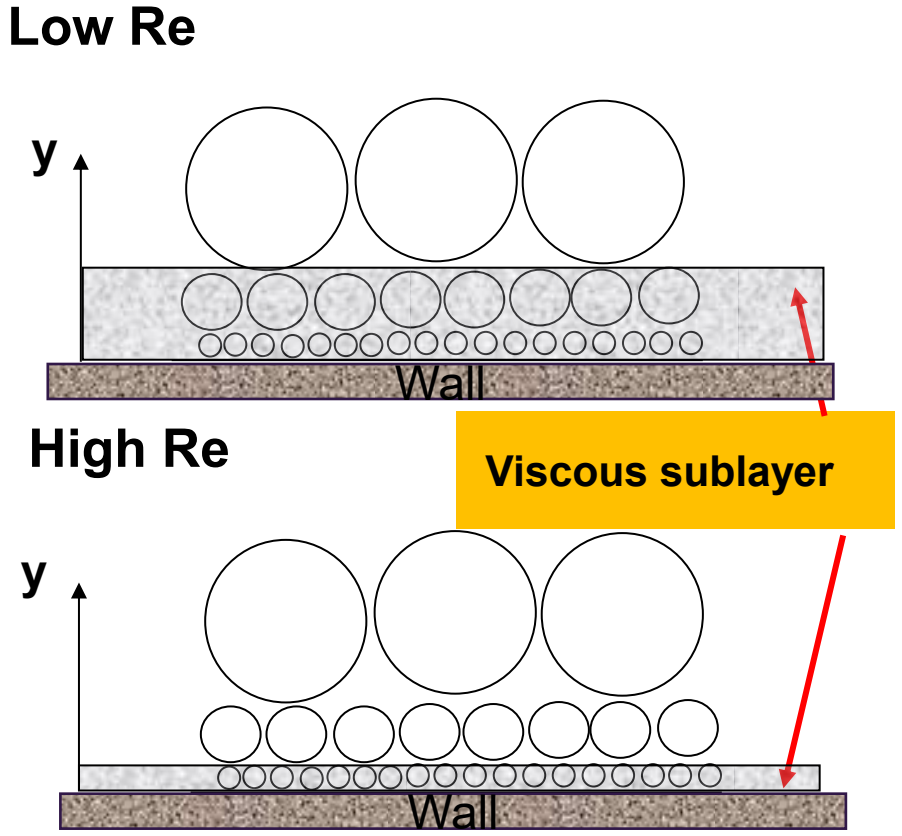
*Bell, J.H., Mehta, R.D., Development of a Two-Stream Mixing Layer from Tripped and Untripped Boundary Layers, AIAAJ (1990), Vol. 18, N. 12

© 2021 Synopsys, Inc.

Near Wall Scaling – What's the Problem

$$L_t = \kappa y$$

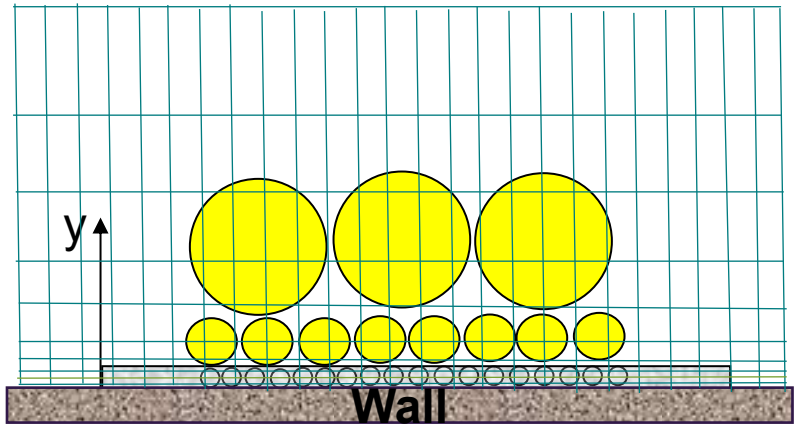
- Turbulent length, L_t, scale is independent of Re number
- Turbulent structures inside sublayer (grey box) are damped out and only structures above the sublayer need to be resolved
- However, the thickness of the viscous sublayer decreases with increasing Re number
- Smaller turbulence structures near the wall get "exposed" as Re increases – these structures need to be resolved in 3D space and time
- Strong Re number scaling for wall-resolved LES costs



© 2026 Synopsys, Inc.

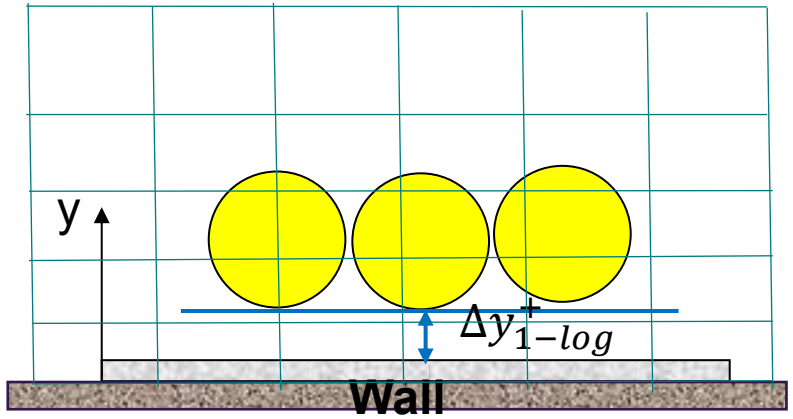
Near Wall Resolution – What are the Alternatives?

WRLES



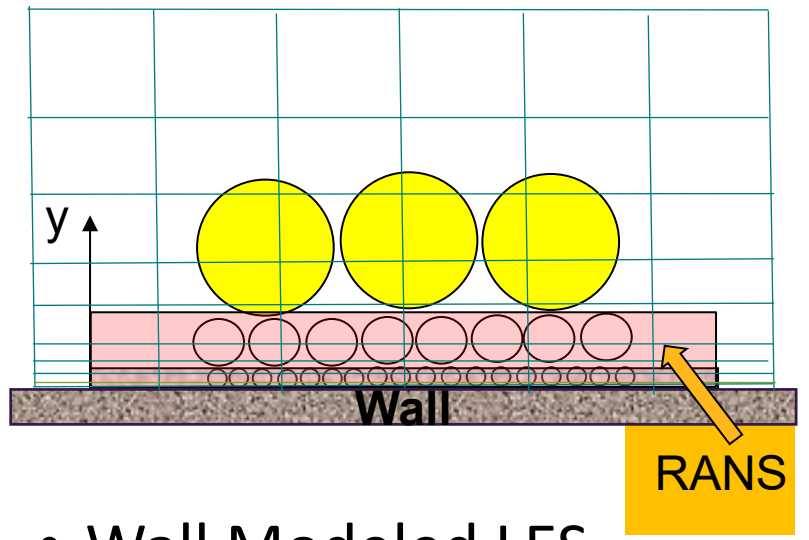
- Wall Resolved LES
  - Eddies are resolved in all three dimensions
  - Fine grid in dx, dy, dz
  - $dx^+ < 40, dz^+ < 20, dy^+ < 1$

WFLES



- Wall Function LES
  - Eddies are resolved only in logarithmic and out layer
  - Wall Function for inner layer
  - Coarse grid in dx, dy, dz
  - dx^+, dy^+, dz^+depend on Re number

WMLES

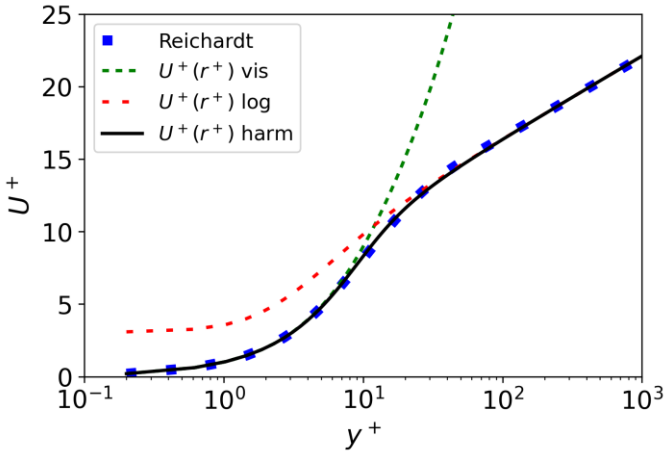


- Wall Modeled LES
  - Eddies are resolved only in logarithmic and out layer
  - Wall Function for inner layer
  - Coarse grid in dx, dz
  - dx^+, dz^+depend on Re number
  - $dy^+ = 1$

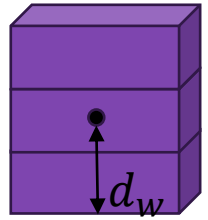
© 2026 Synopsys, Inc.

Non-Iterative Wall-Function (Wall-Stress Model)

- Based on r^+ not y^+
- Approximation of logarithmic layer
- Harmonic blending of sub- and log-layer
- Is ‘equilibrium’ good enough?



$$r^+ = y^+ u^+ = \frac{u d_w}{\nu}$$



$$u^+_{sub} = \sqrt{r^+}$$

$$u^+_{log} = C_1 ln(r^+ + C_2)$$

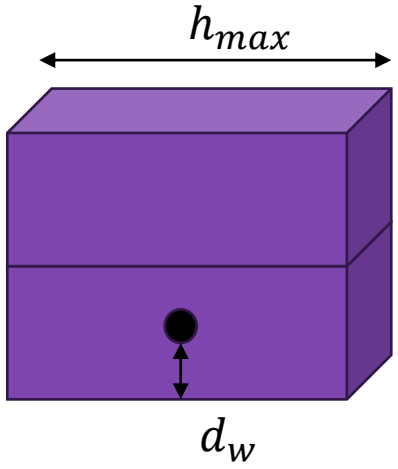
$$u^+_{eq}(r^+) = \left((u^+_{sub})^{-\alpha} + (u^+_{log})^{-\alpha} \right)^{-\frac{1}{\alpha}}$$

EVALUATION OF A NEW EXPLICIT WALL-FUNCTION FORMULATION FOR LARGE EDDY SIMULATION ETMM15
E.K. Guseva[1], D.G. Flad[1], A.L. Nogueira[2], D.P. Graebin[2] and F.R. Menter[1]
In Print

© 2026 Synopsys, Inc.

Algebraic WMLES Channel

- Blend of Prandtl and any alg. LES model
- Blending based on d_w/h_{max}
- Function f_{sw} has two parameters
  - C_{w1} switching location
  - C_{w2} steepness



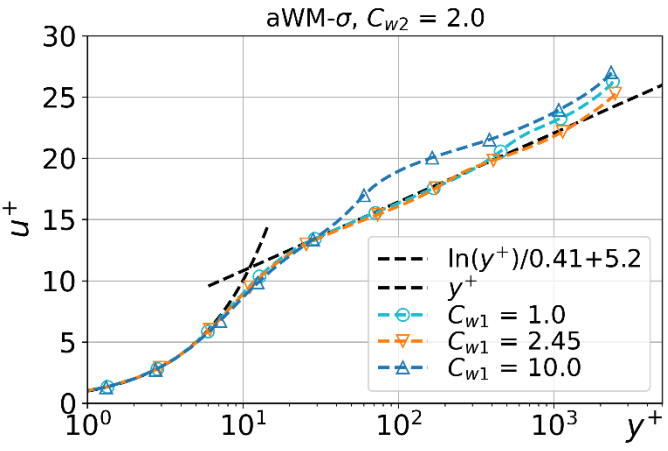
$$\nu_t = f_{sw}\nu_{t,RANS} + (1 - f_{sw})\nu_{t,LES}$$

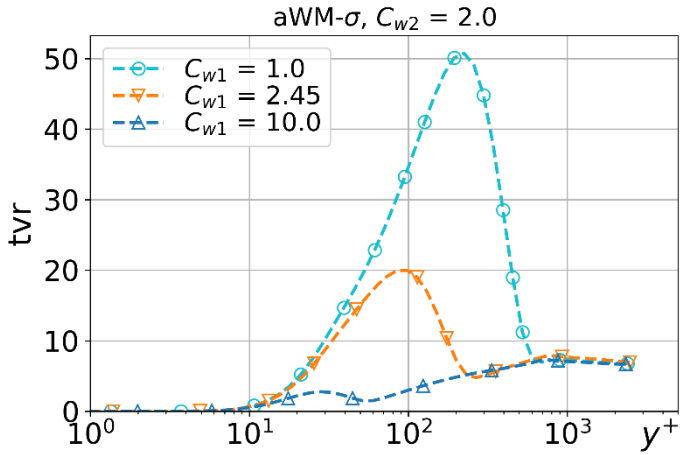
$$f_{sw} = e^{-\left(\frac{C_{w1}d_w}{h_{max}}\right)^{C_{w2}}}$$

$$\nu_{t,RANS} = l_{RANS}^2 S$$

$$l_{RANS} = 0.41 d_w \sqrt{f_{wd}}$$

$$f_{wd} = 1 - e^{-\left(\frac{l^+}{25}\right)^3}$$





ETMM14 Conference Barcelona 2023

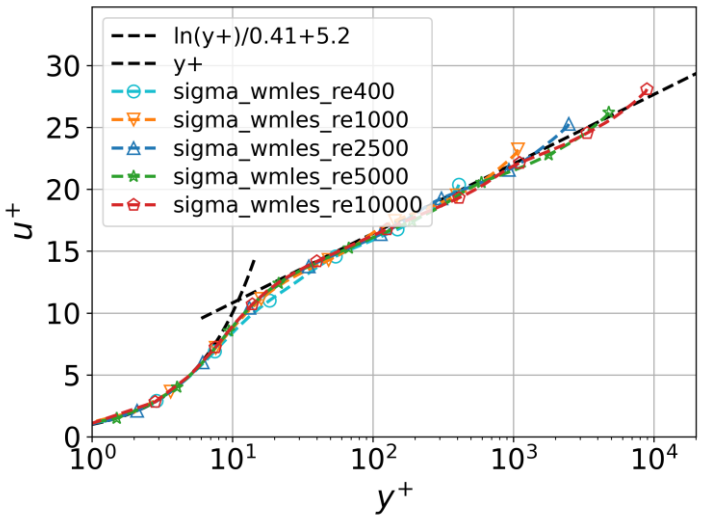
A Family of Wall-Modeled Large Eddy Simulation Formulations

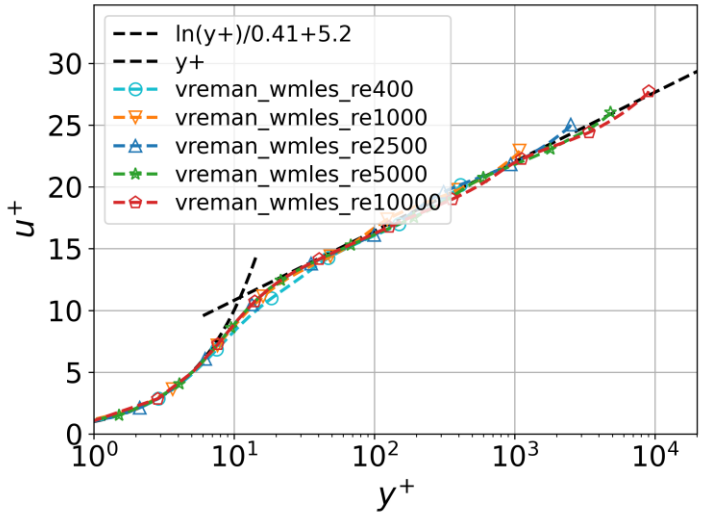
F. R. Menter[1], E. Guseva[1], A. Stabnikov[2], A. Matyushenko[2] and A. Garbaruk[2]

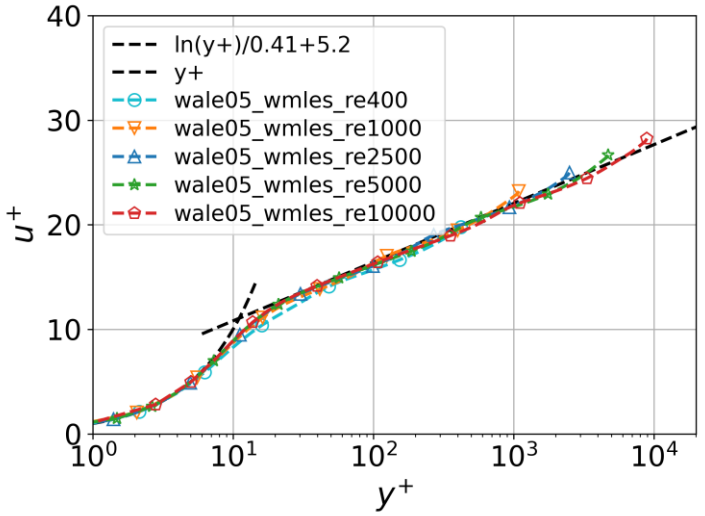
[1] Ansys Germany GmbH
[2] SPbPU St. Petersburg

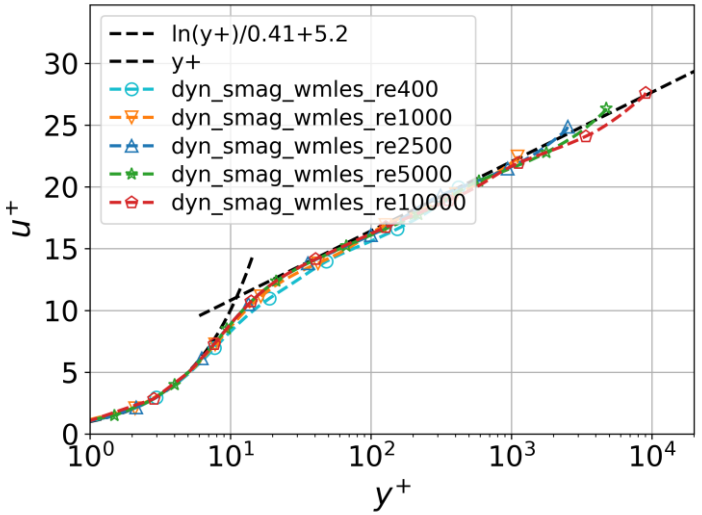
Variation in LES Models

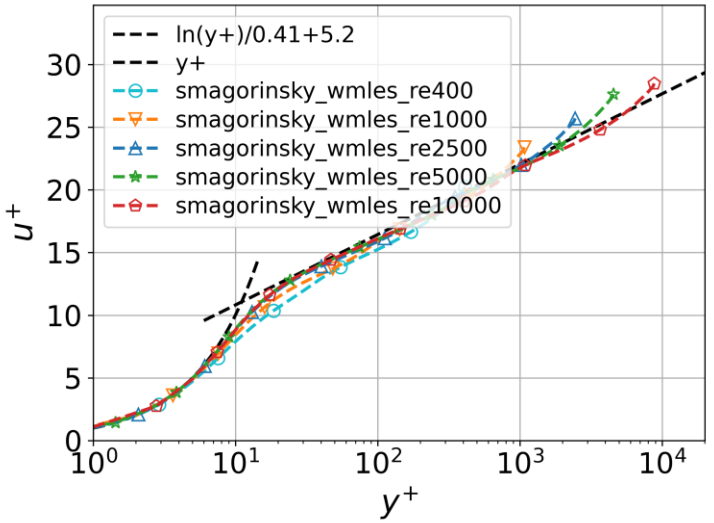
- σ, Vreman, WALE, Smagorinsky, dyn. Smagorinsky
- All LES models work
- No log-layer miss-match only for very low-dissipation numerics
- Grids ~600,000 cells







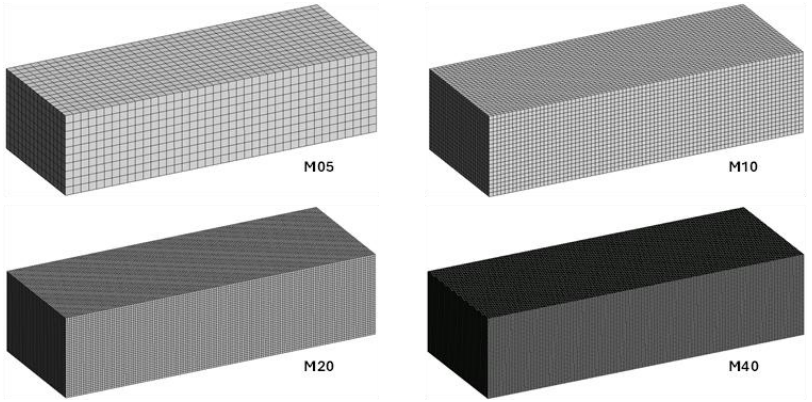




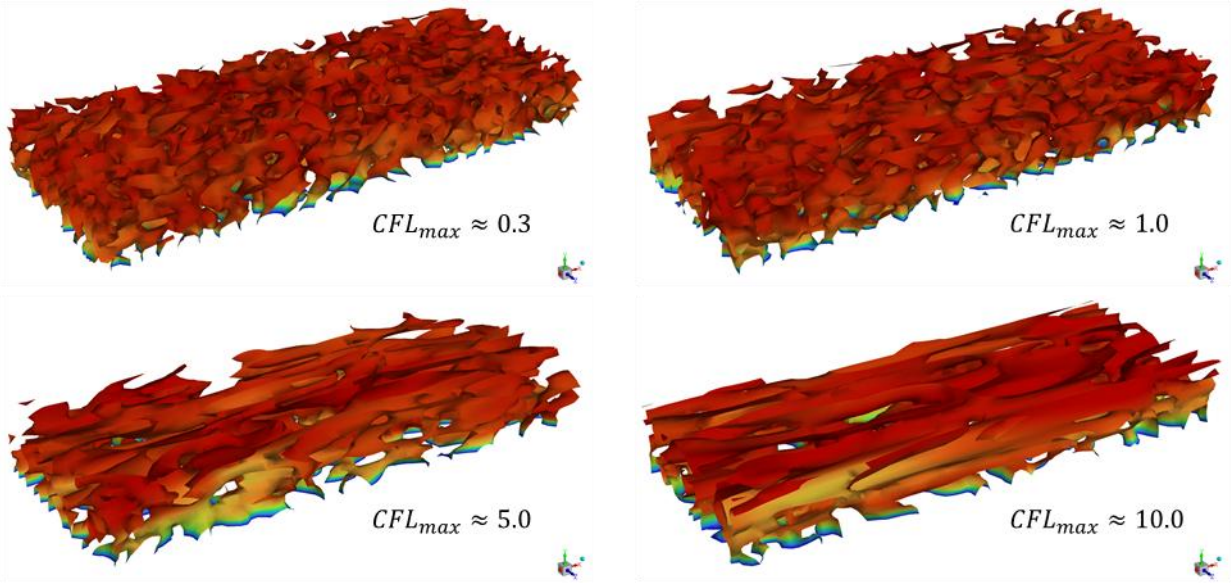
© 2026 Synopsys, Inc.

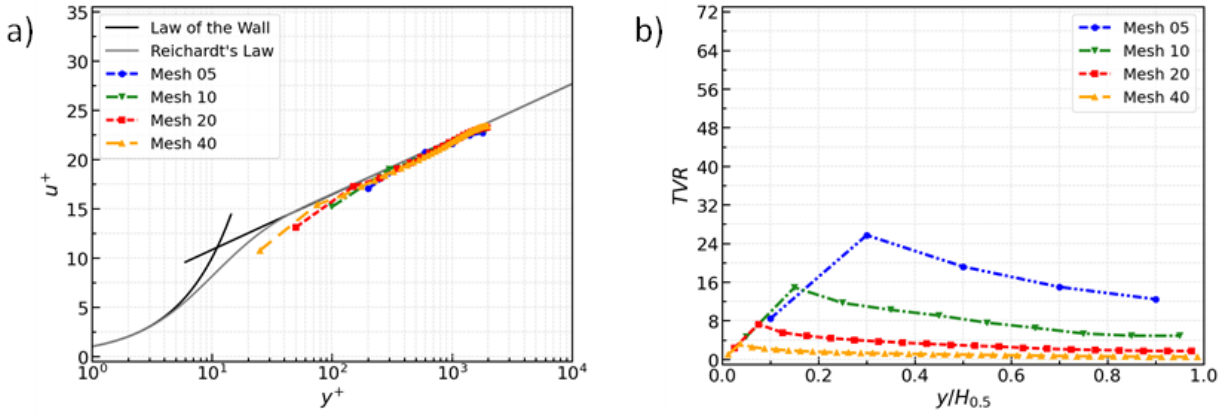
Example: WFLES Channel

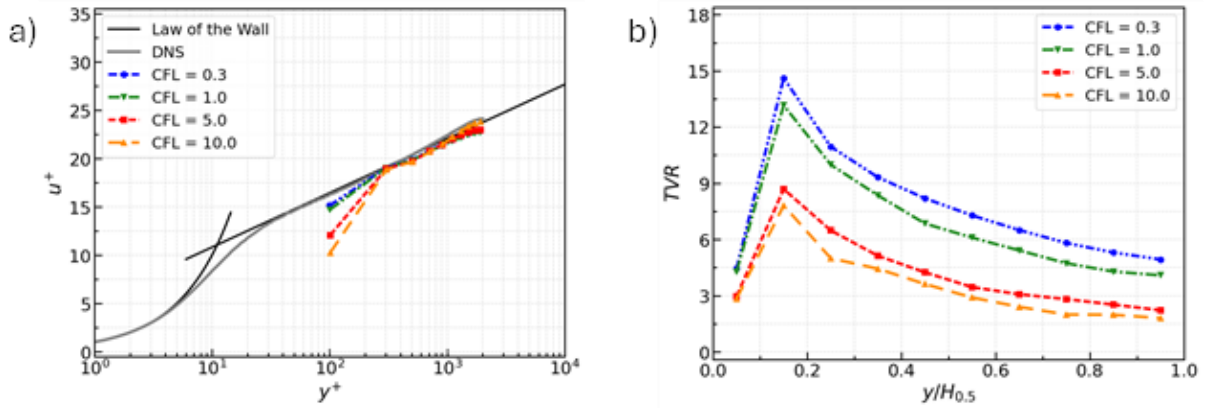
- Grid study
  - Even with only 5 cells across the half-channel – good match to log-layer
  - Sensitivity to CFL number moderate – even though turbulence structures affected



Mesh	Cells per $H/2$	Cell Count
M05	5	6,000
M10	10	48,000
M20	20	384,000
M40	40	3,072,000







© 2026 Synopsys, Inc.

Smooth ramp

- High-Fidelity CFD Verification Workshop 2024
- $Re_L = 6.67 \times 10^5$
- 4 Octree-type meshes with 2 prisms

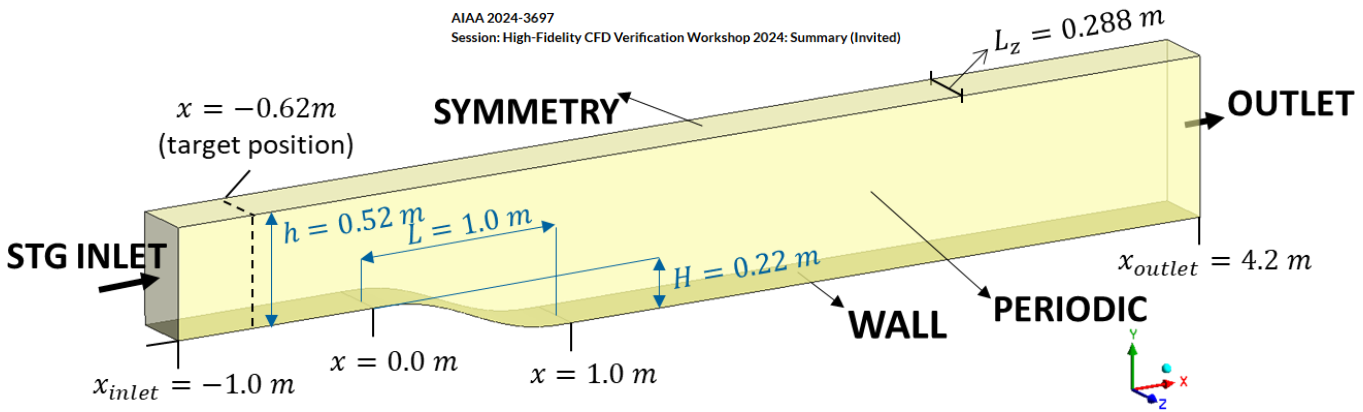
Mesh	Base cell size Δ, mm	$\delta_{target\ position}/\Delta$	Cells, M
M1	8	5	0.57
M2	4	10	3.7
M3	2	20	27
M4	1	40	200

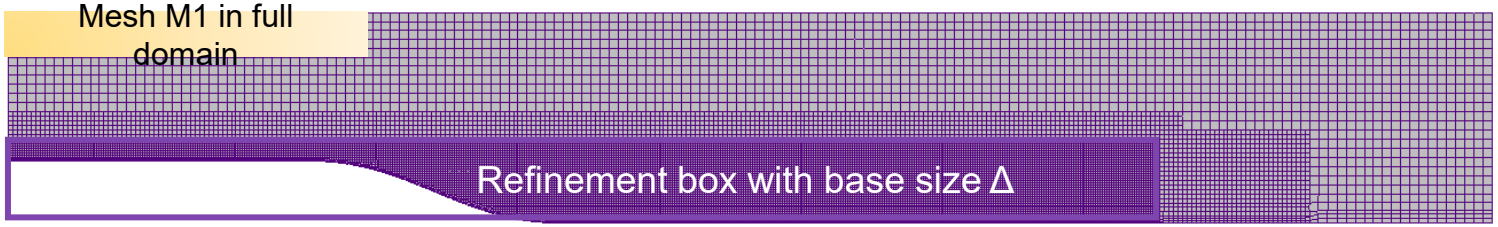
- RANS Velocity profile + unsteady fluctuations
  - ➢ Synthetic Turbulence Generator

Summary of the Wall-Modeled LES Test Suite From the 2024 AIAA High-Fidelity CFD Verification Workshop

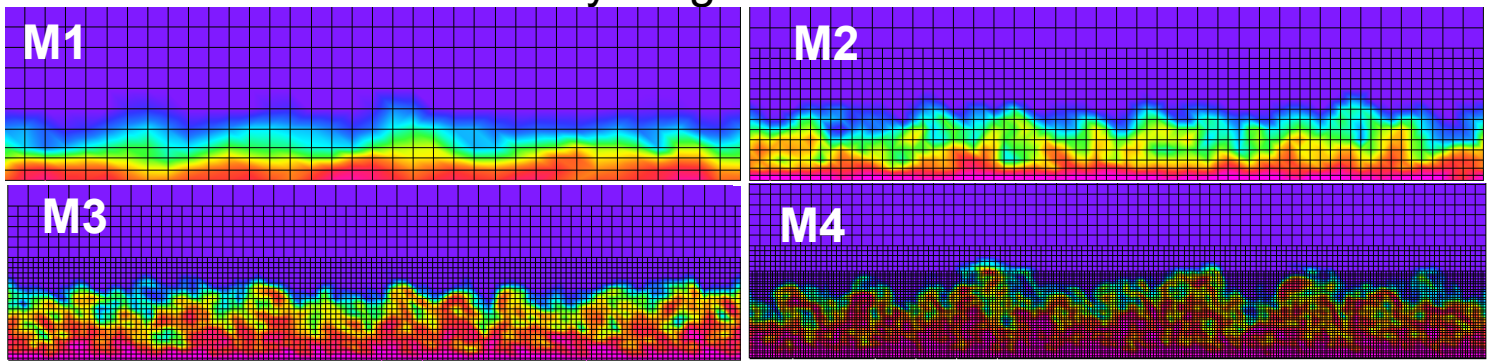
Daniel J. Garmann, Matthew J. Schwartz, Johan Larsson, Ivan Bermejo-Moreno, Ali Uzun, Sparsh Ganju, Christoph Brehm, Timofey Mukha, Matteo Parsani, David Flad and Jannik Schrempp

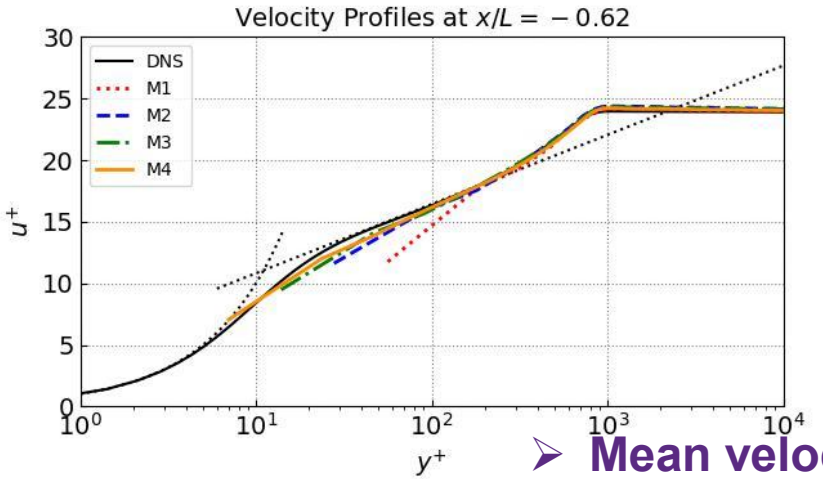
AIAA 2024-3697
Session: High-Fidelity CFD Verification Workshop 2024: Summary (Invited)





Mesh and vorticity magnitude contours at x/L = -0.62

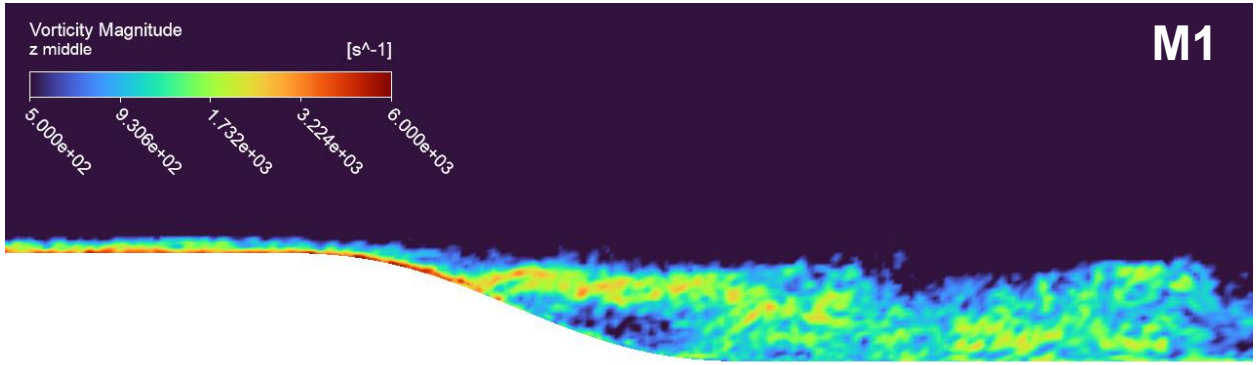


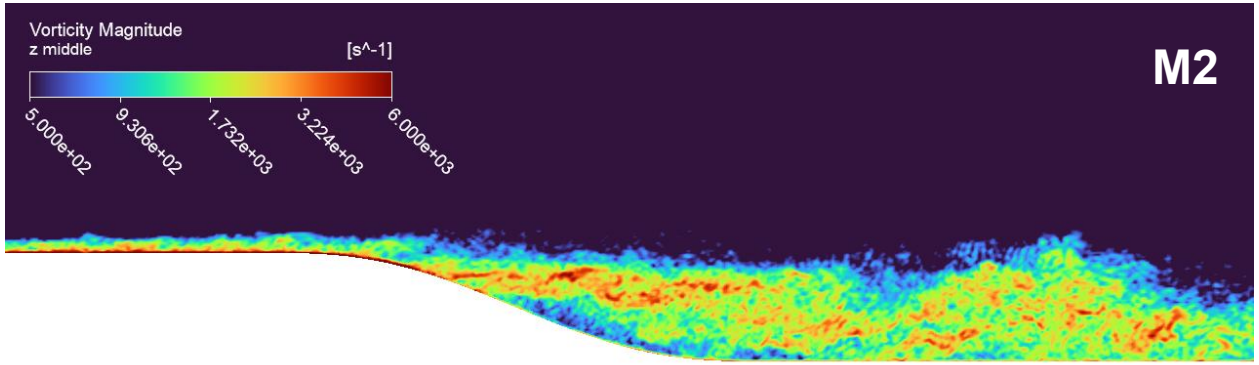


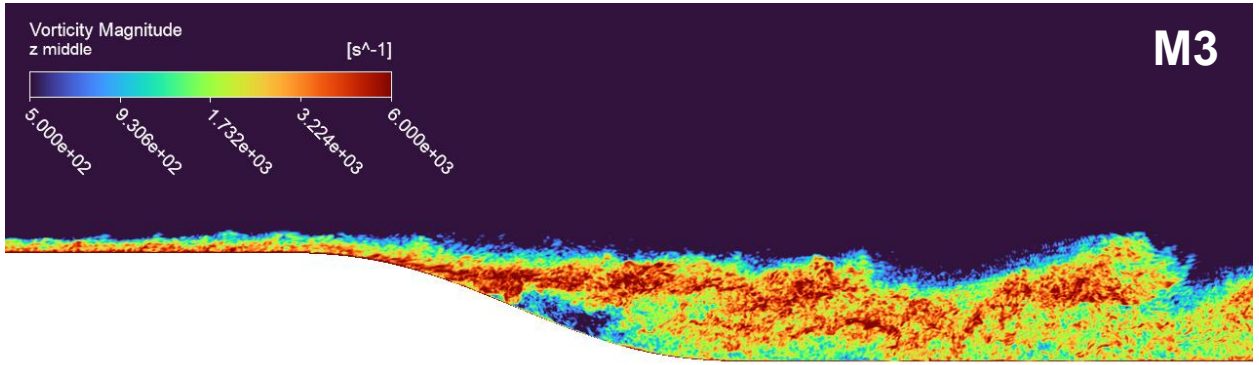
➢ **Mean velocity profile predicted well even on the extreme coarse mesh**

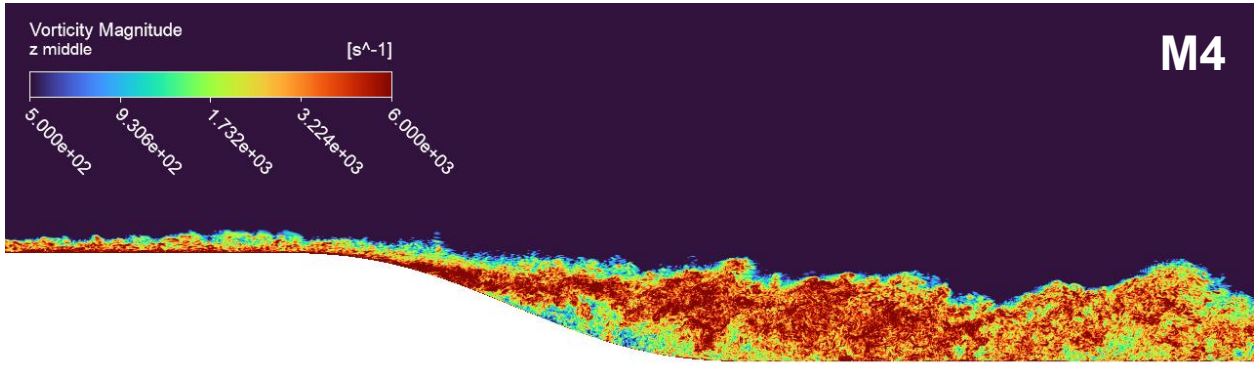
© 2026 Synopsys, Inc.

Flow visualization. Vorticity magnitude contours



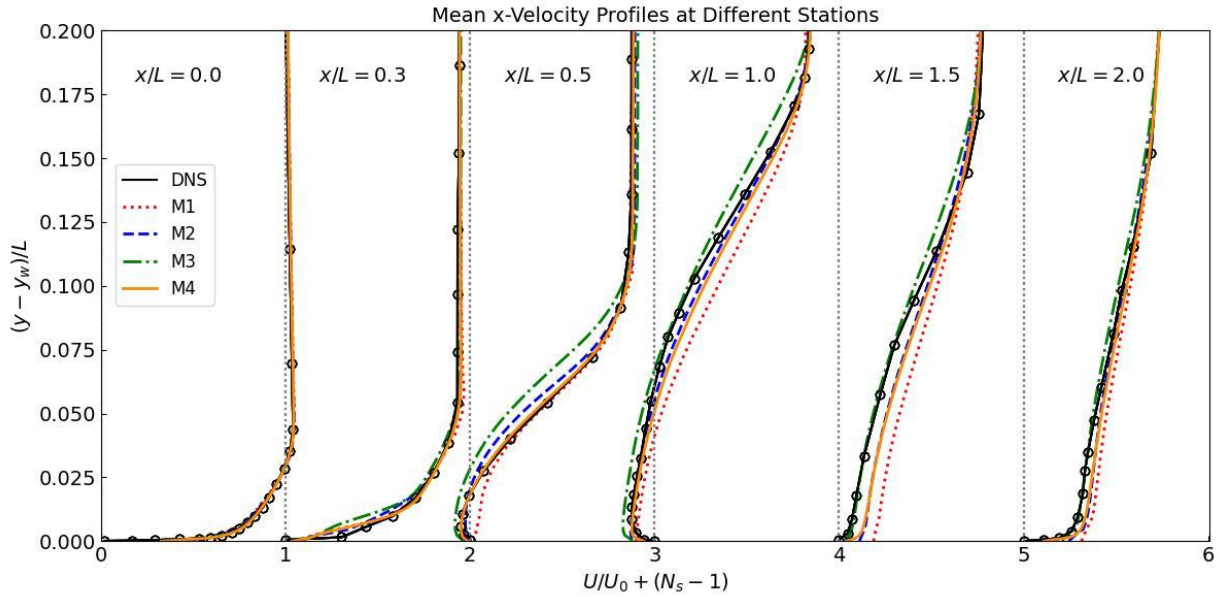


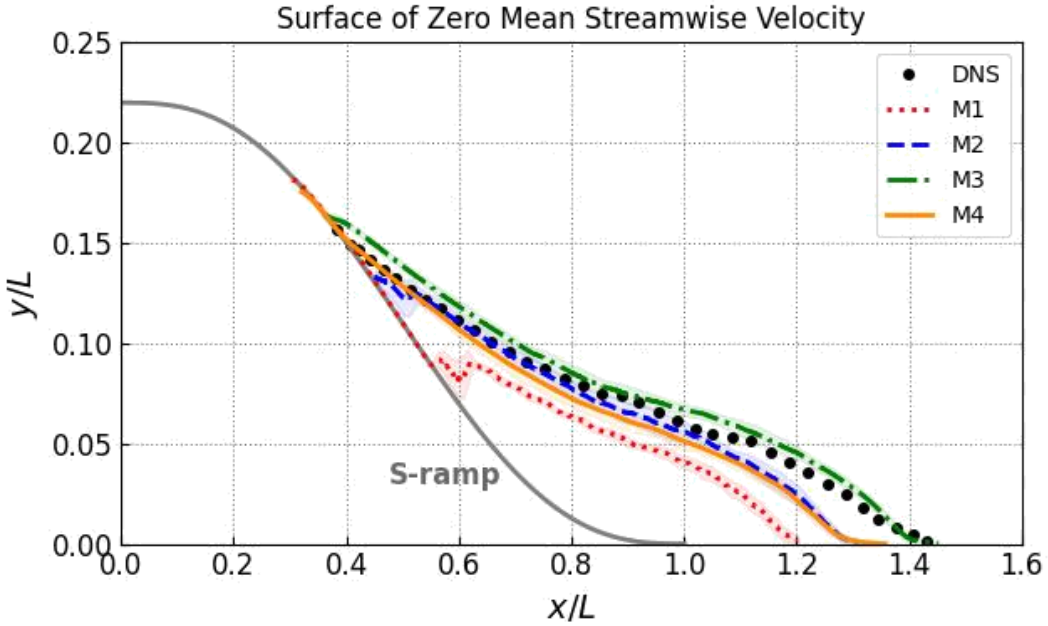




© 2026 Synopsys, Inc.

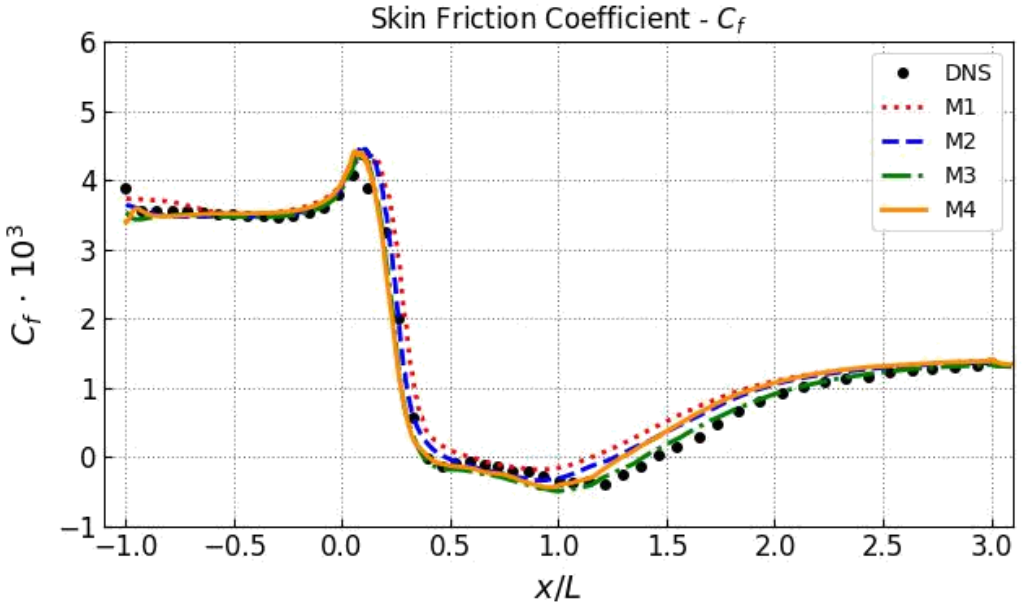
Smooth ramp





- On all the meshes, except on the coarsest mesh, the separation point is accurately captured
- The recirculation bubble size is somewhat underpredicted
- Some mesh dependency still observed between M2-M4
- Overall, good agreement with the DNS results was achieved

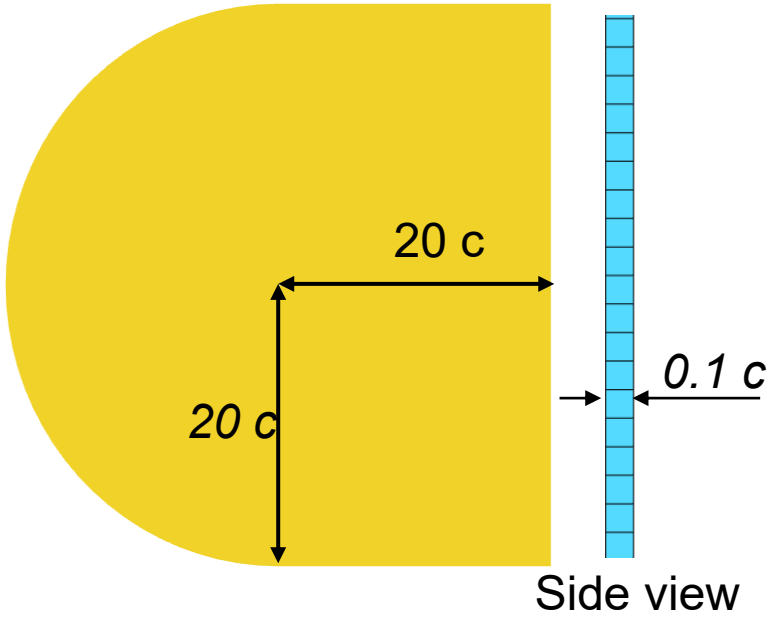
➢ **No systematic failure of equilibrium wall function in non-equilibrium conditions**



© 2026 Synopsys, Inc.

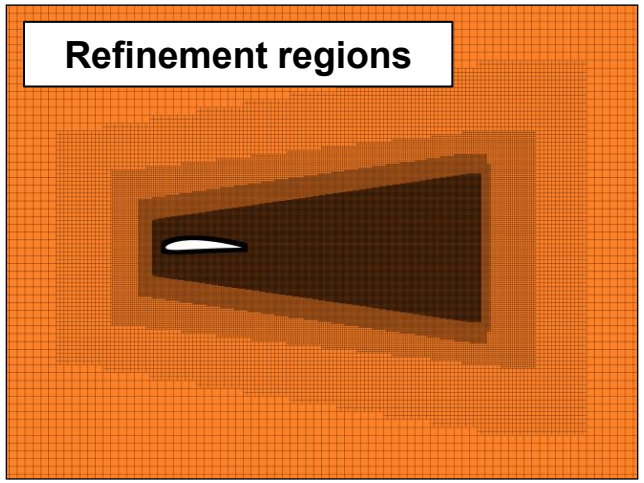
NACA 4412

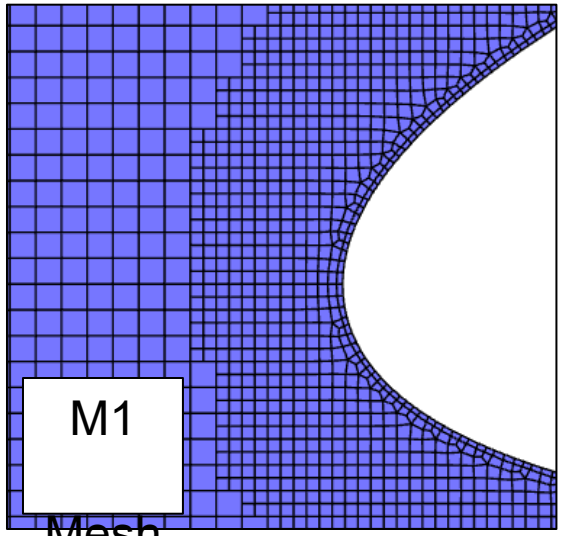
- NACA-4412 with $c = 1$ m
- Angle of attack: 13.87°
- Experiment - Wadcock (1979)
- $Re_c = 1.5 \cdot 10^6$
- 2 prisms layers near the wall

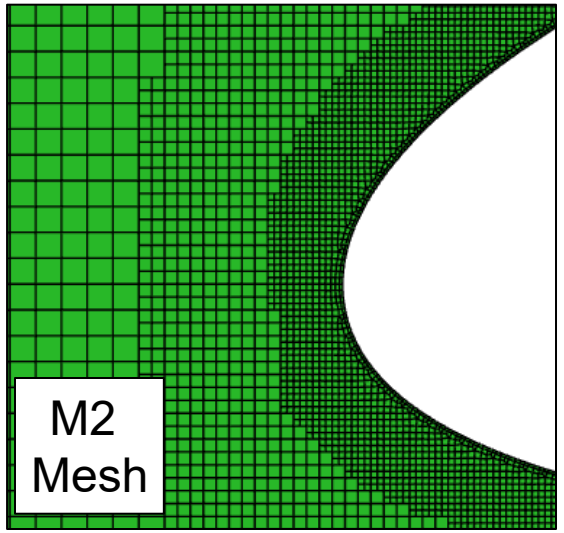


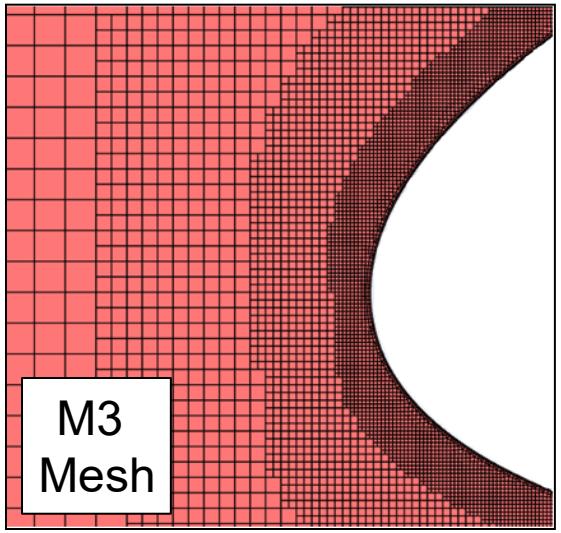
Side view

Mesh	Cell Size at Airfoil	Number of cells	y^+_{max}	Number of cells in BL*
M1	$1.6 \cdot 10^{-3}$ c	3,898,080	90	12
M2	$0.8 \cdot 10^{-3}$ c	8,117,776	48	19
M3	$0.4 \cdot 10^{-3}$ c	24,917,296	24	28









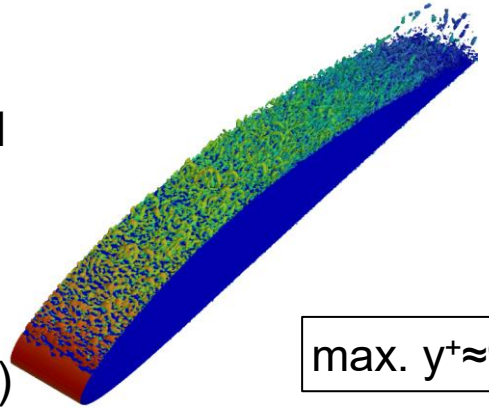
* Measured at *x/c = 0.5*

© 2026 Synopsys, Inc.

Results – iPBNS

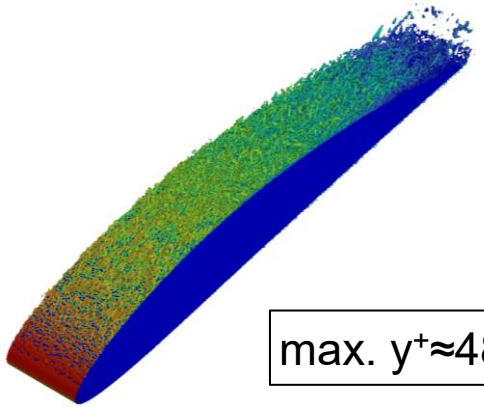
- Boundary Layer resolution
  - 12, 19 and 28 cells per BL at x/c = 0.5
  - 6, 9, and 16 cells per BL at x/c = 0.1
- Transition is noticeably delayed at meshes M1 and M2
  - (no trip was used in the simulations)
  - Local mesh refinement is an option

M1 mesh ($4 \cdot 10^6$)



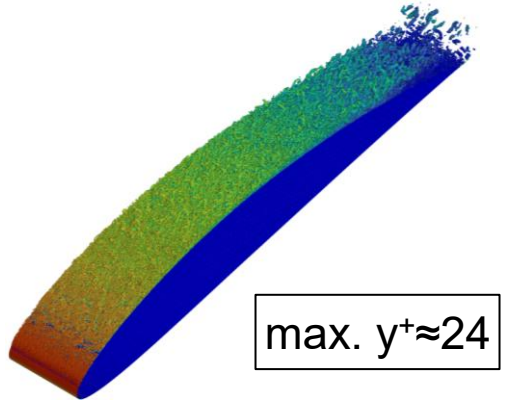
max. $y^+ \approx 90$

M2 mesh ($8 \cdot 10^6$)

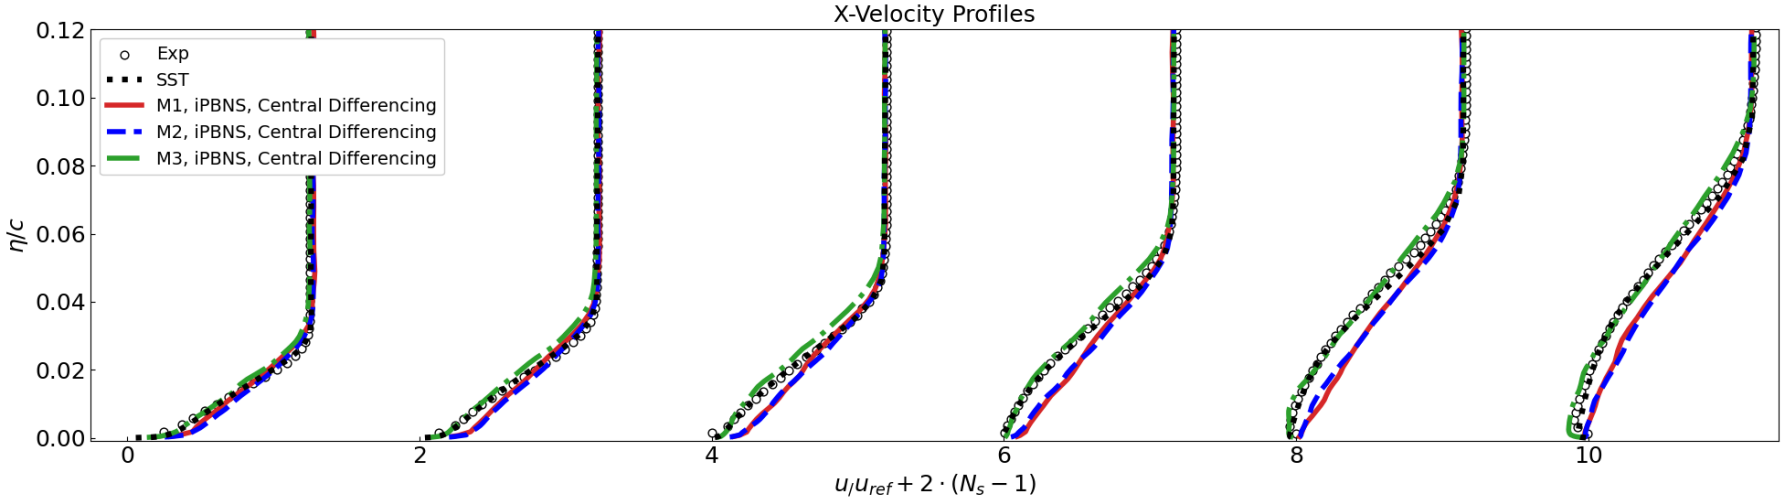


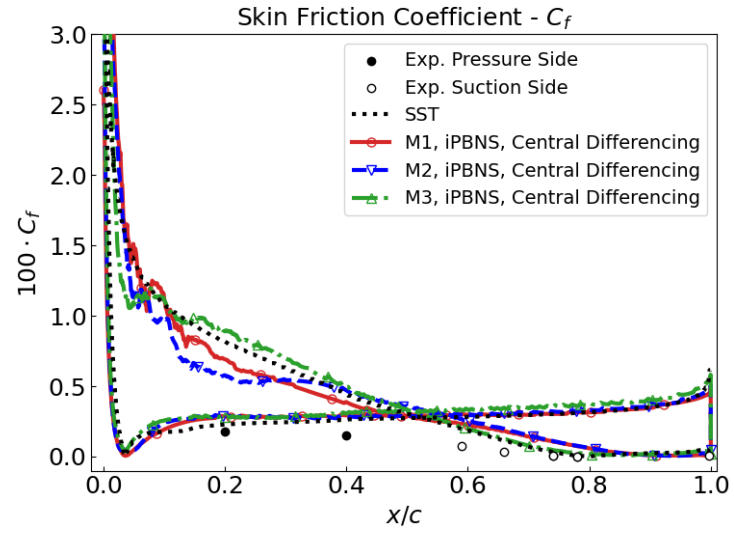
max. $y^+ \approx 48$

M3 mesh ($25 \cdot 10^6$)



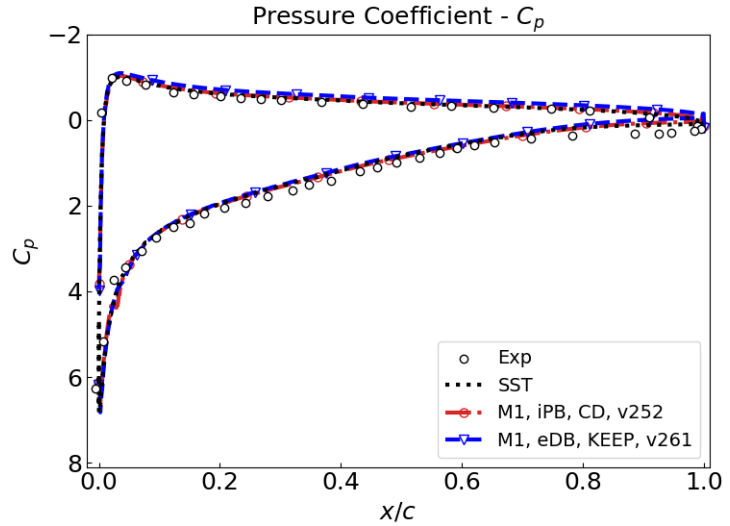
max. $y^+ \approx 24$

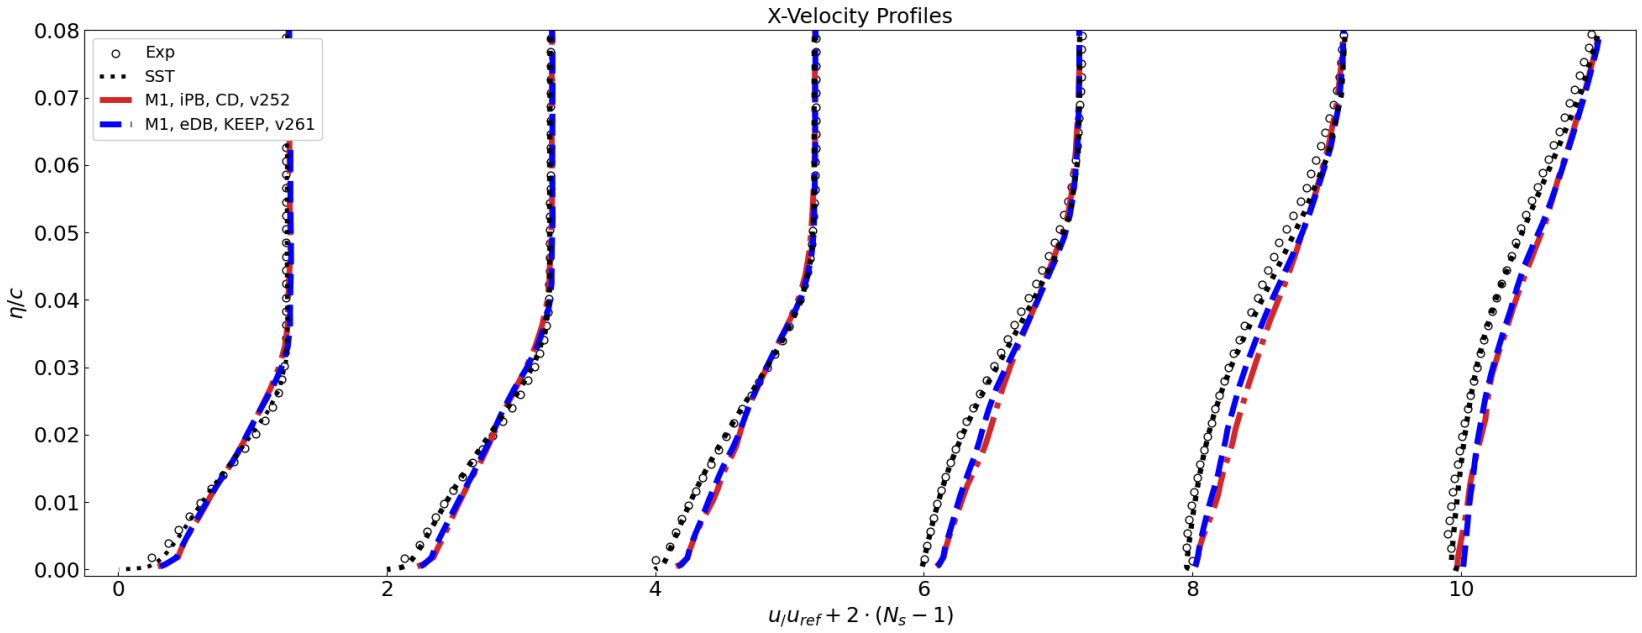


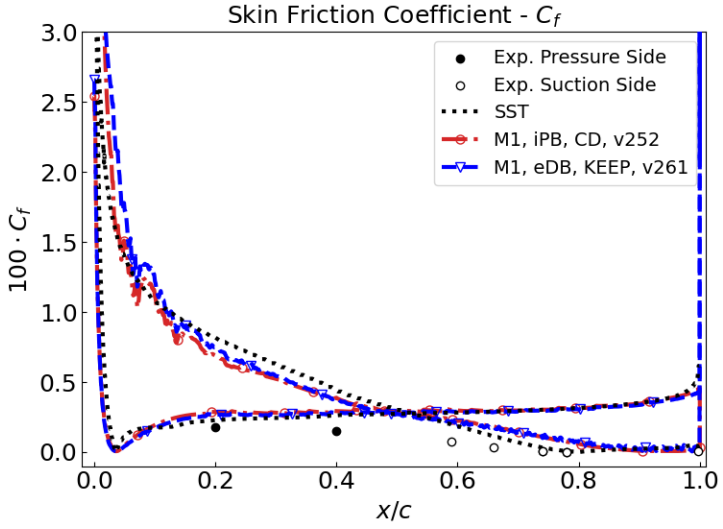


Results – M1 mesh (4 million cells)

- $\Delta t = 3.3 \cdot 10^{-6}$ s for iPBNS
- $\Delta t = 1 \cdot 10^{-6}$ s for eDBNS
  - Low dissipation KEEP scheme
- Both code central discretization





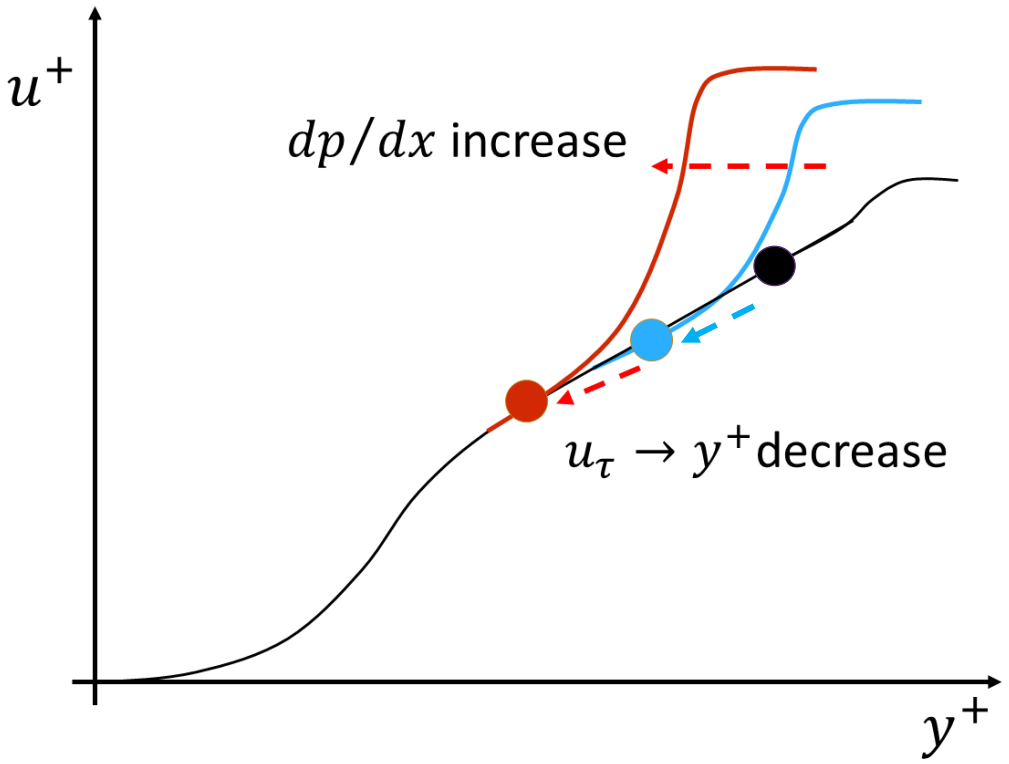


Using central scheme

© 2026 Synopsys, Inc.

Why Do Equilibrium Wall Functions Work under $dp/dx \neq 0$?

- Under strong positive dp/dx the logarithmic layer is lost from the outside
- At the same time u_τ and therefore $y^+ = \frac{u_\tau y}{\nu}$ decreases
- Near separation the wall-function point moves into the sublayer
- Equilibrium wall functions work well even under non-equilibrium conditions



"...combining the experimental uncertainties with the lack of proven grid-convergence in most studies, one must conclude that there has not yet been a truly convincing and conclusive demonstration of any superiority of non-equilibrium over equilibrium wall-stress-models in LES. Performing such a conclusive demonstration is a high priority in the field, regardless of the eventual conclusion from it."

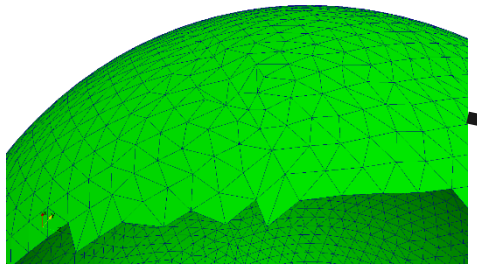
— Larsson et al. 2016

Top-Down Meshing Strategies

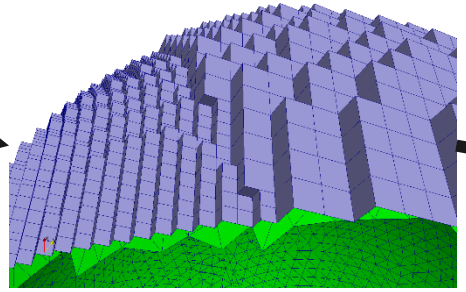
- For a body-fitted approach, meshing complex geometries can be bottle neck
- Classical bottom-up meshing do not scale well in highly parallel settings
  - From surface mesh to volume mesh
- Alternative approach is Top-Down meshing
  - Form volume to surface

© 2026 Synopsys, Inc.

Basic Steps in Normal Projection

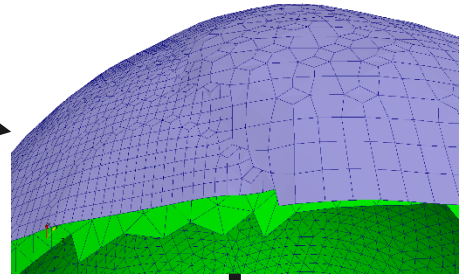


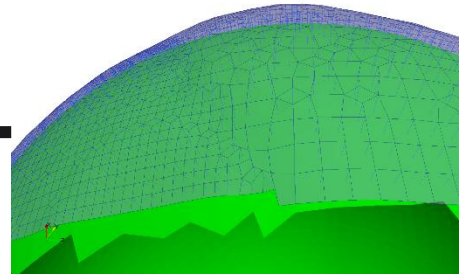
0. Triangulated Surface Data



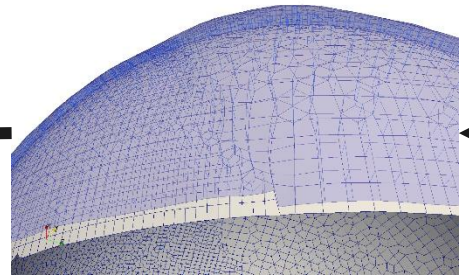
1. Cartesian Representation
(Cut-Out, Projection Front)

2. Smooth Projection Front

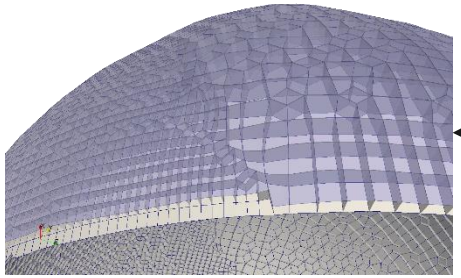




3. Compute Closest Points



4. Generate Boundary Faces

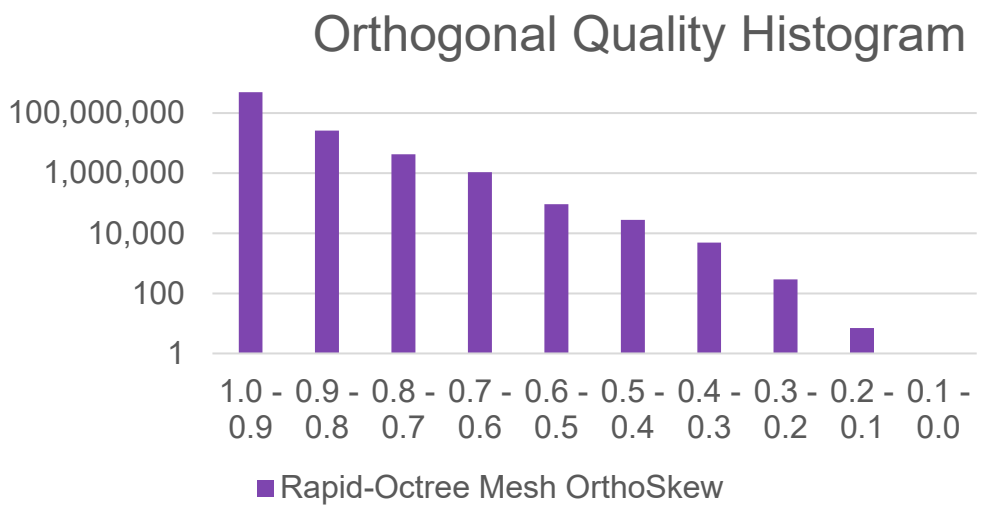


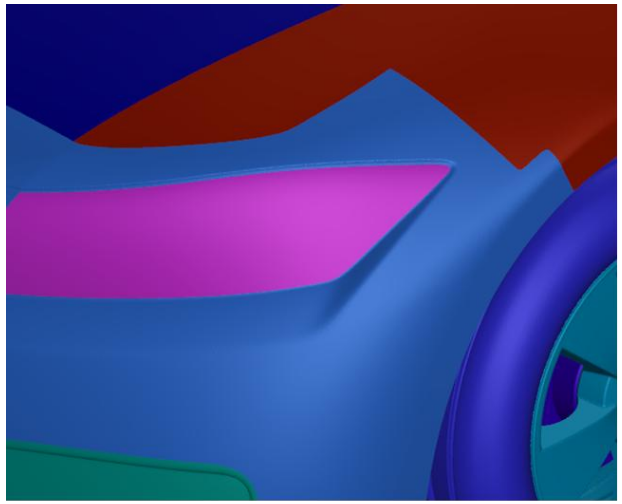
5. Create Projection Cell

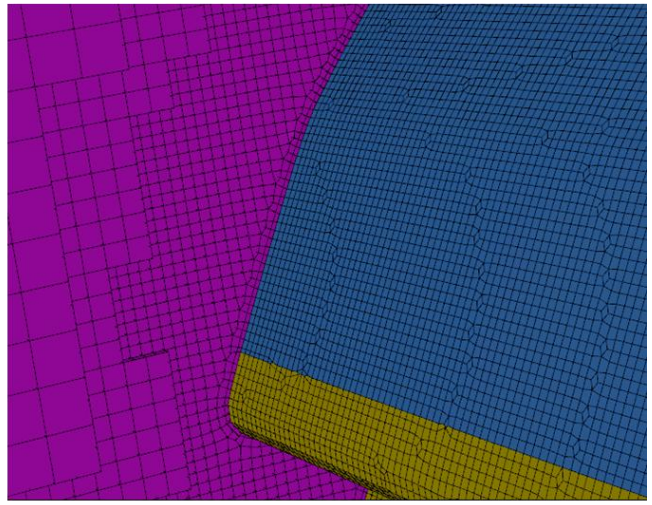
© 2026 Synopsys, Inc.

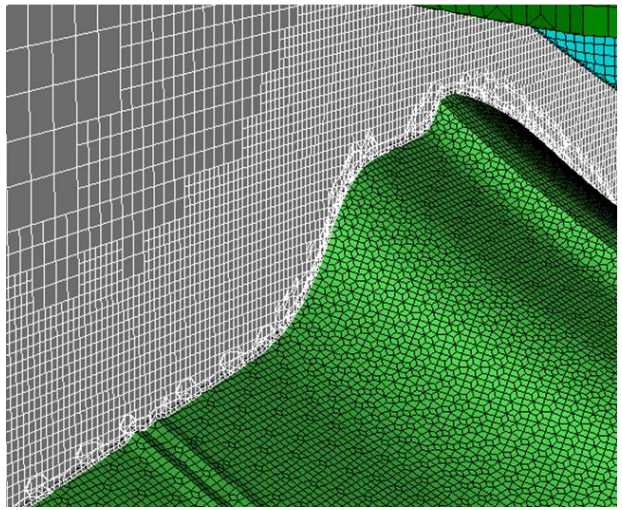
Base Mesh DrivAer Car

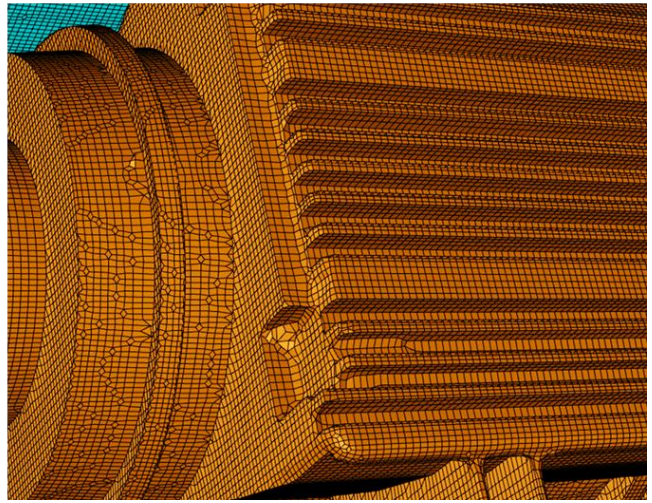
Number of CPU Cores for Meshing	480
Cell Count	525 252 268
Boundary Face Count	35 258 803
Meshing Time	1338.7s (~22min)
Worst Cell Orthogonal Skew	0.134
Maximum Aspect Ratio	24









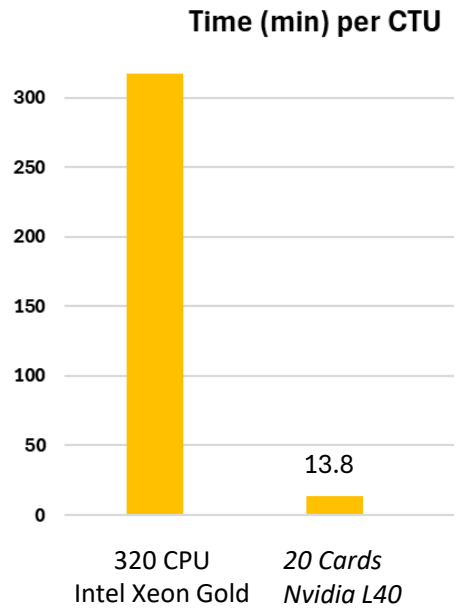


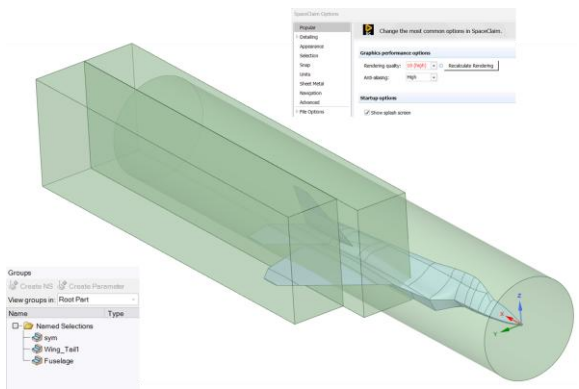
© 2026 Synopsys, Inc.

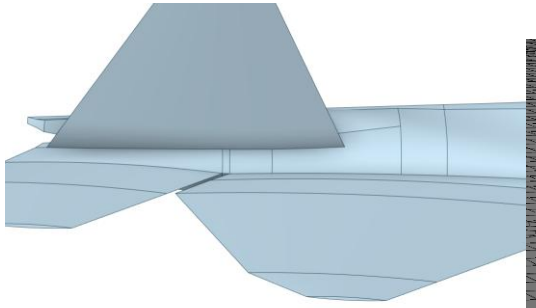
5th Gen Aircraft Workflow

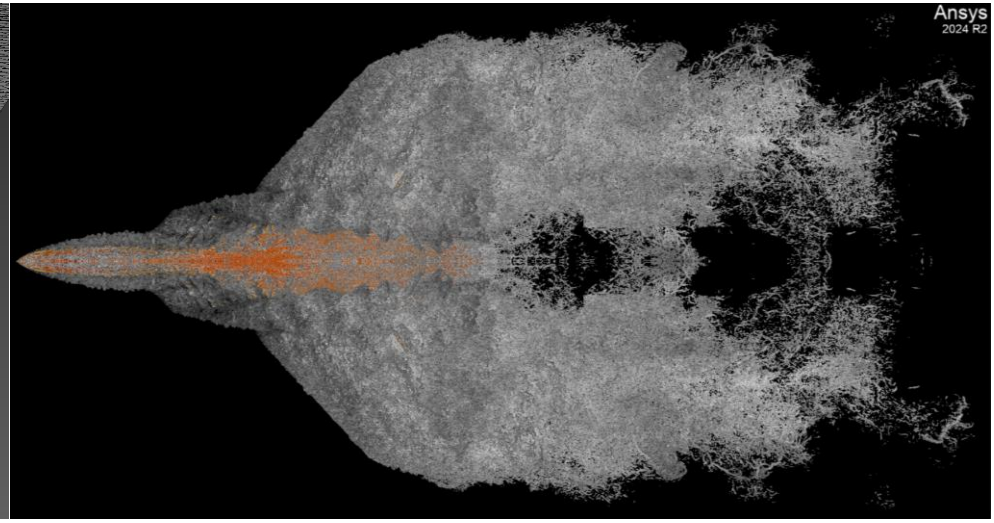
Giannelis, N.F.; Bykerk, T.; Vio, G.A. A Generic Model for Benchmark Aerodynamic Analysis of Fifth-Generation High-Performance Aircraft. Aerospace 2023, 10, 746.

GEOMETRY	***MESH***	***SOLVER***
DISCOVERY Modelling	FLUENT Meshing Rapid Octree	FLUENT GPU Single Precision
Rendering quality BOI creation Named Selections	RO settings BOI settings Prisms settings	WALE LES Optimized LES 2nd order implicit
Windows Laptop	CPU Cluster (480 Cores)	GPU cluster (20 x L40)
~ 20 min	600 M cells in 30min	14 hours (50 CTU)



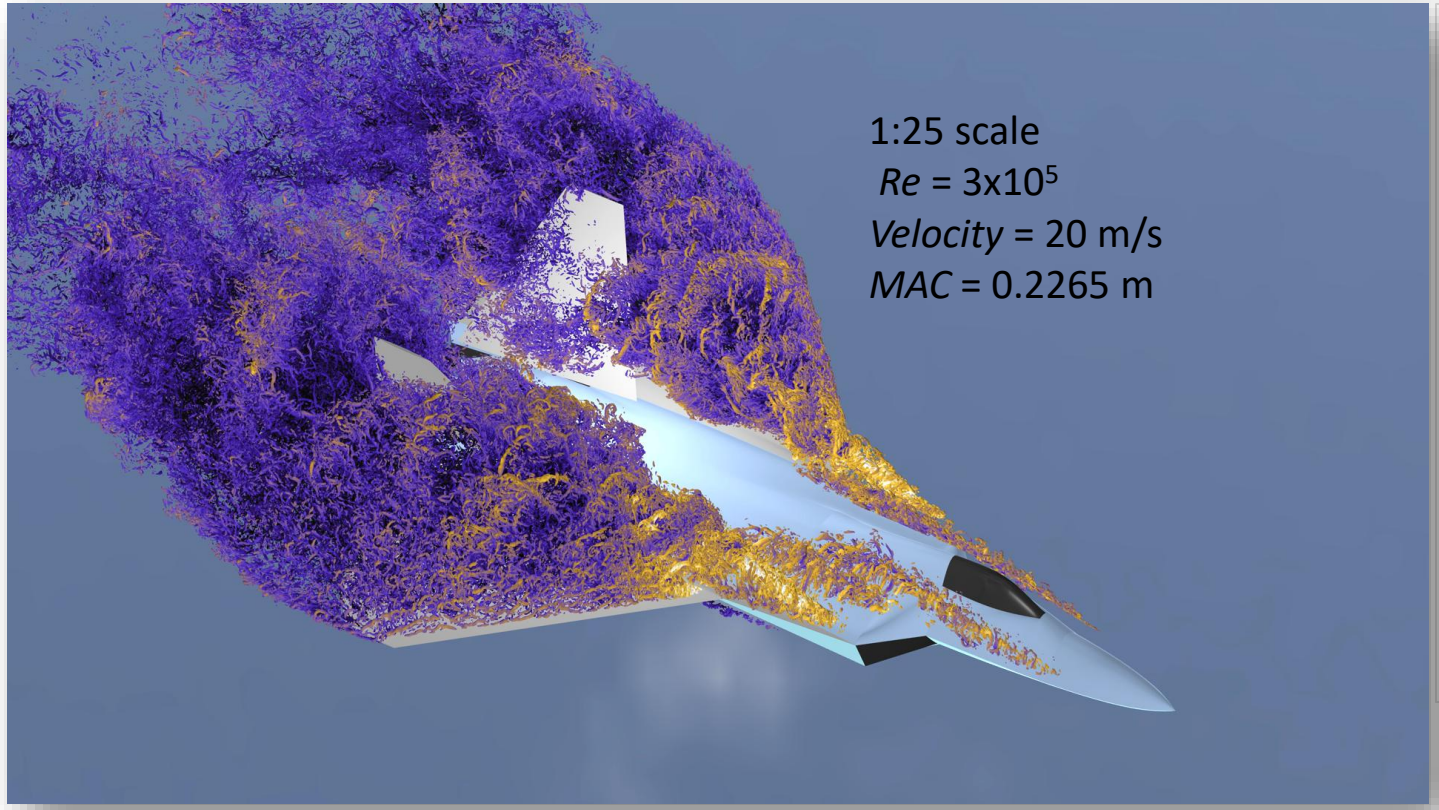




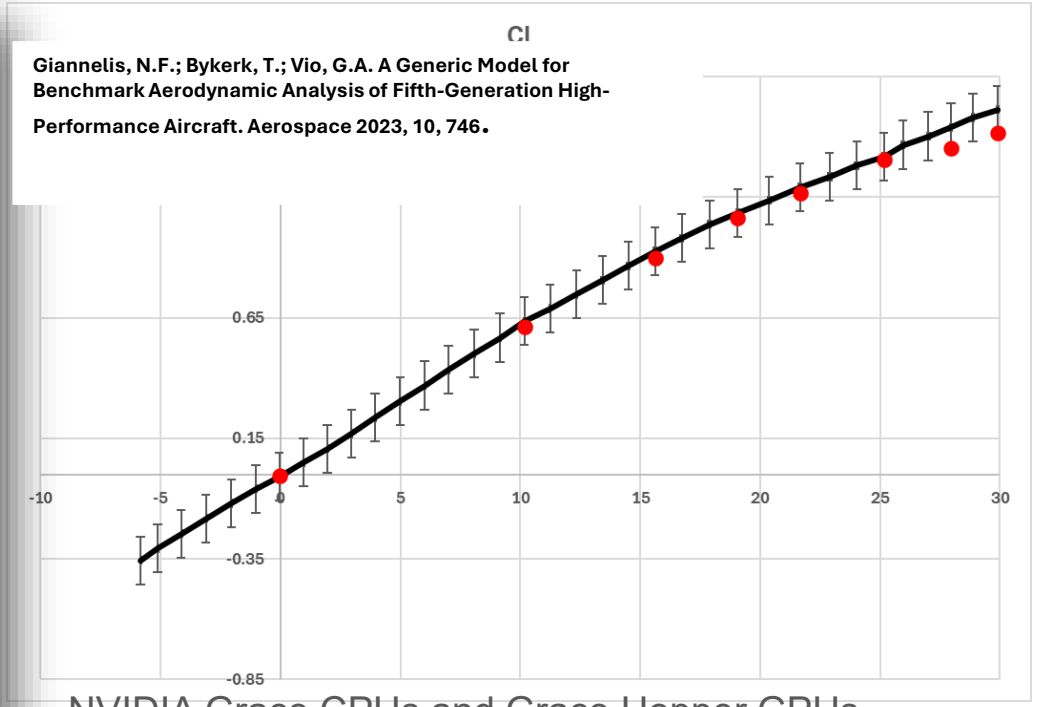


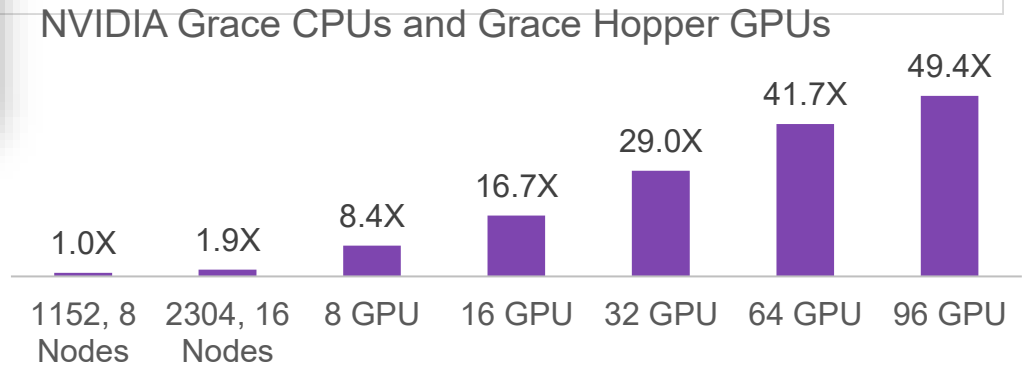
The Ansys Fluent **GPU Solver** in Practice

5th Gen High-Performance Aircraft



Lift coefficients at different attack angles

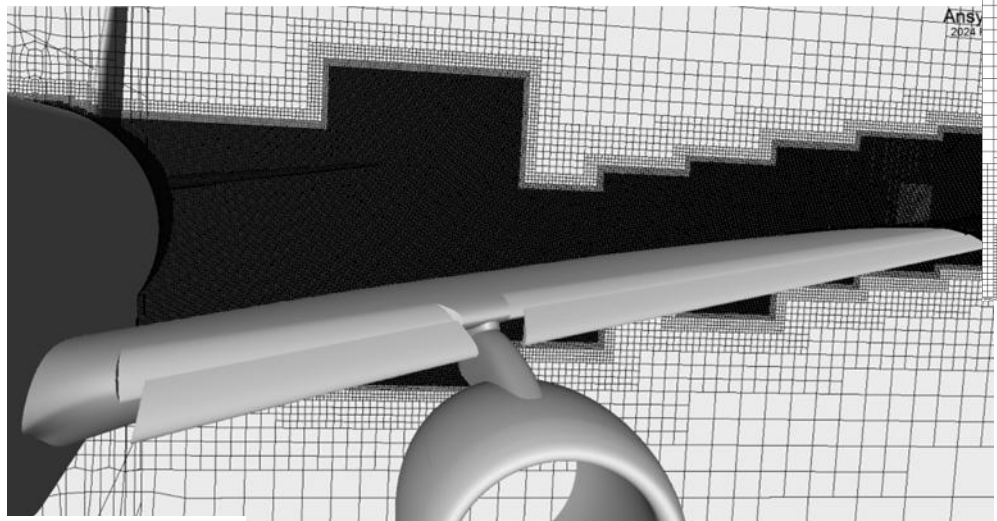


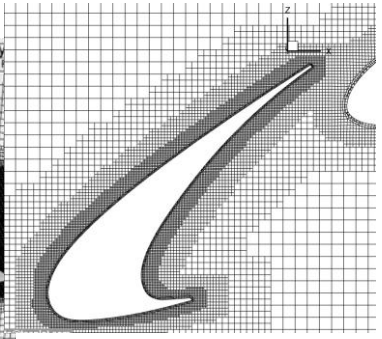


© 2026 Synopsys, Inc.

High-Lift AIAA Workshop Case

- High Lift Workshop 5 config 2.4
- Octree grid 20 min on 500 CPUs
- WFLES on GPU solver is tested on Rapid Octree mesh

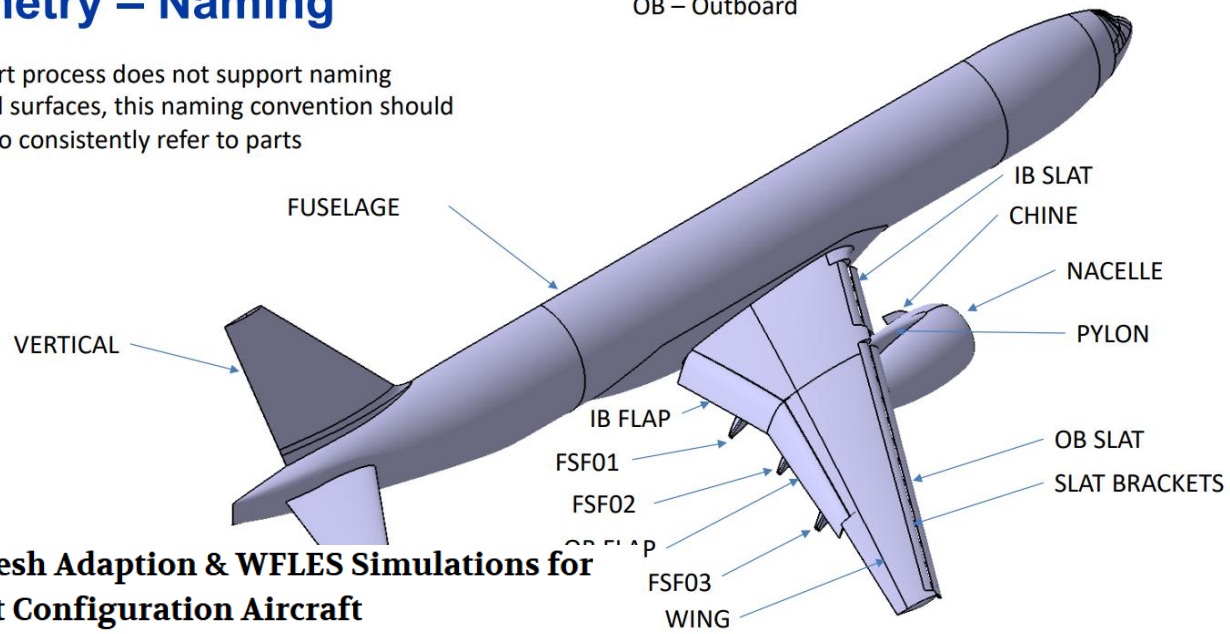


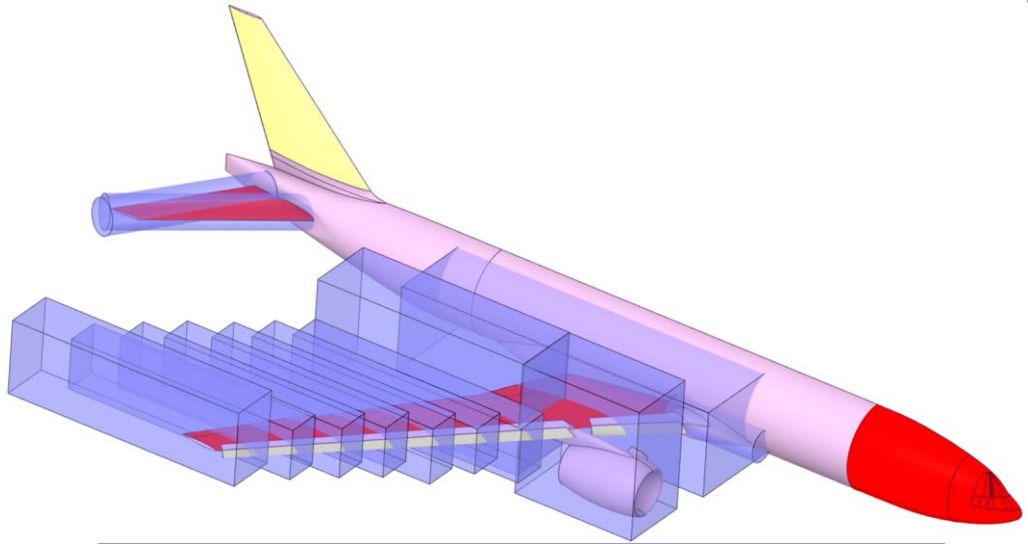


Geometry – Naming

IB – Inboard
OB – Outboard

The export process does not support naming individual surfaces, this naming convention should be used to consistently refer to parts





1 billion cells in 19 min on 512 Cores

RANS Mesh Adaption & WFLES Simulations for High Lift Configuration Aircraft

Corrections for this article

Cristhian Aliaga, Krishna Zore, Rachid Malk, Tobias Berg, Achilleas Krikas, Jeya Selva, Laith Zori and Boris Makarov

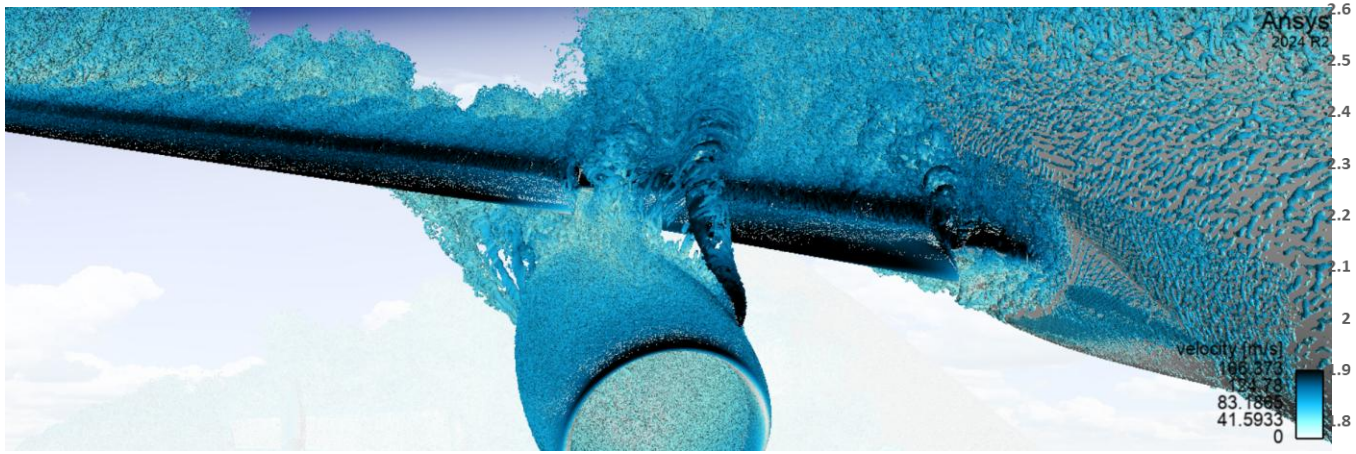
AIAA 2025-0498
Session: Special Session: High-Lift-Prediction-Workshop (HLPW-5): RANS and Adaptive Simulations (Invited)
Published Online: 3 Jan 2025 • https://doi.org/10.2514/6.2025-0498

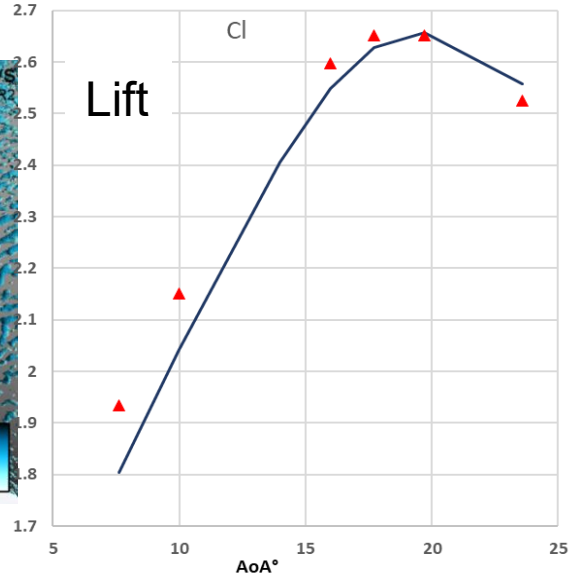
© 2026 Synopsys, Inc.

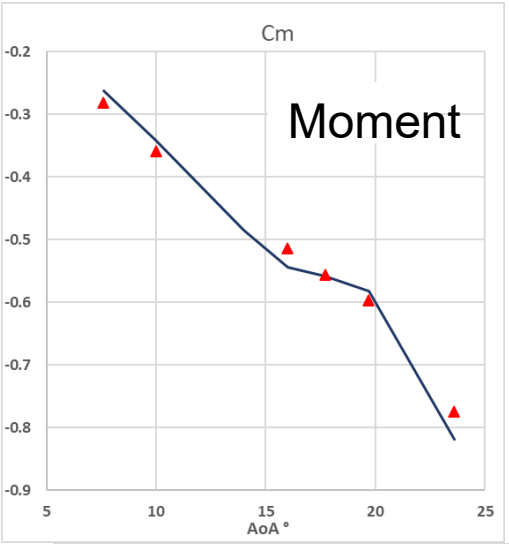
New Era for Fast Transient

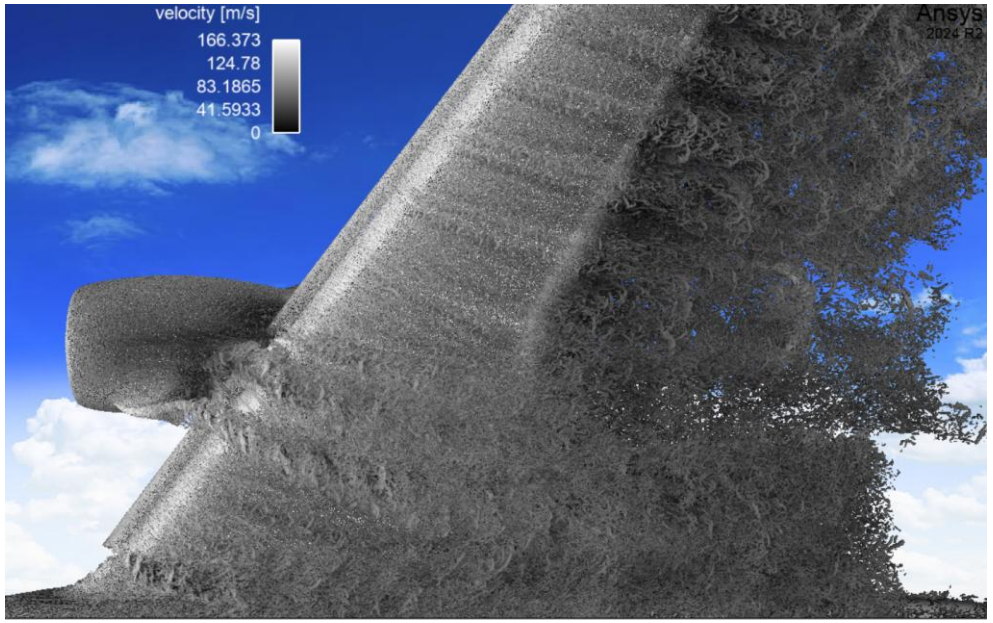


Solutions

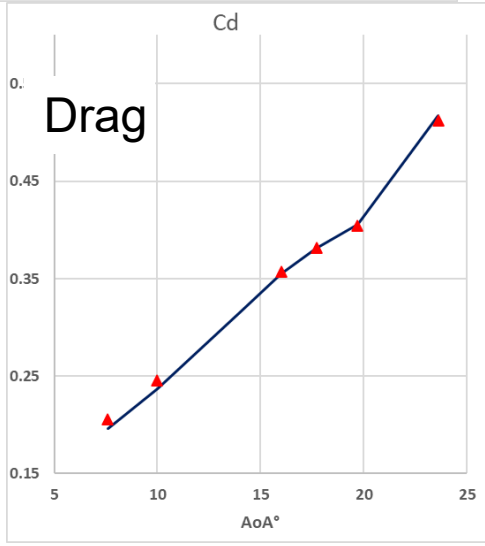






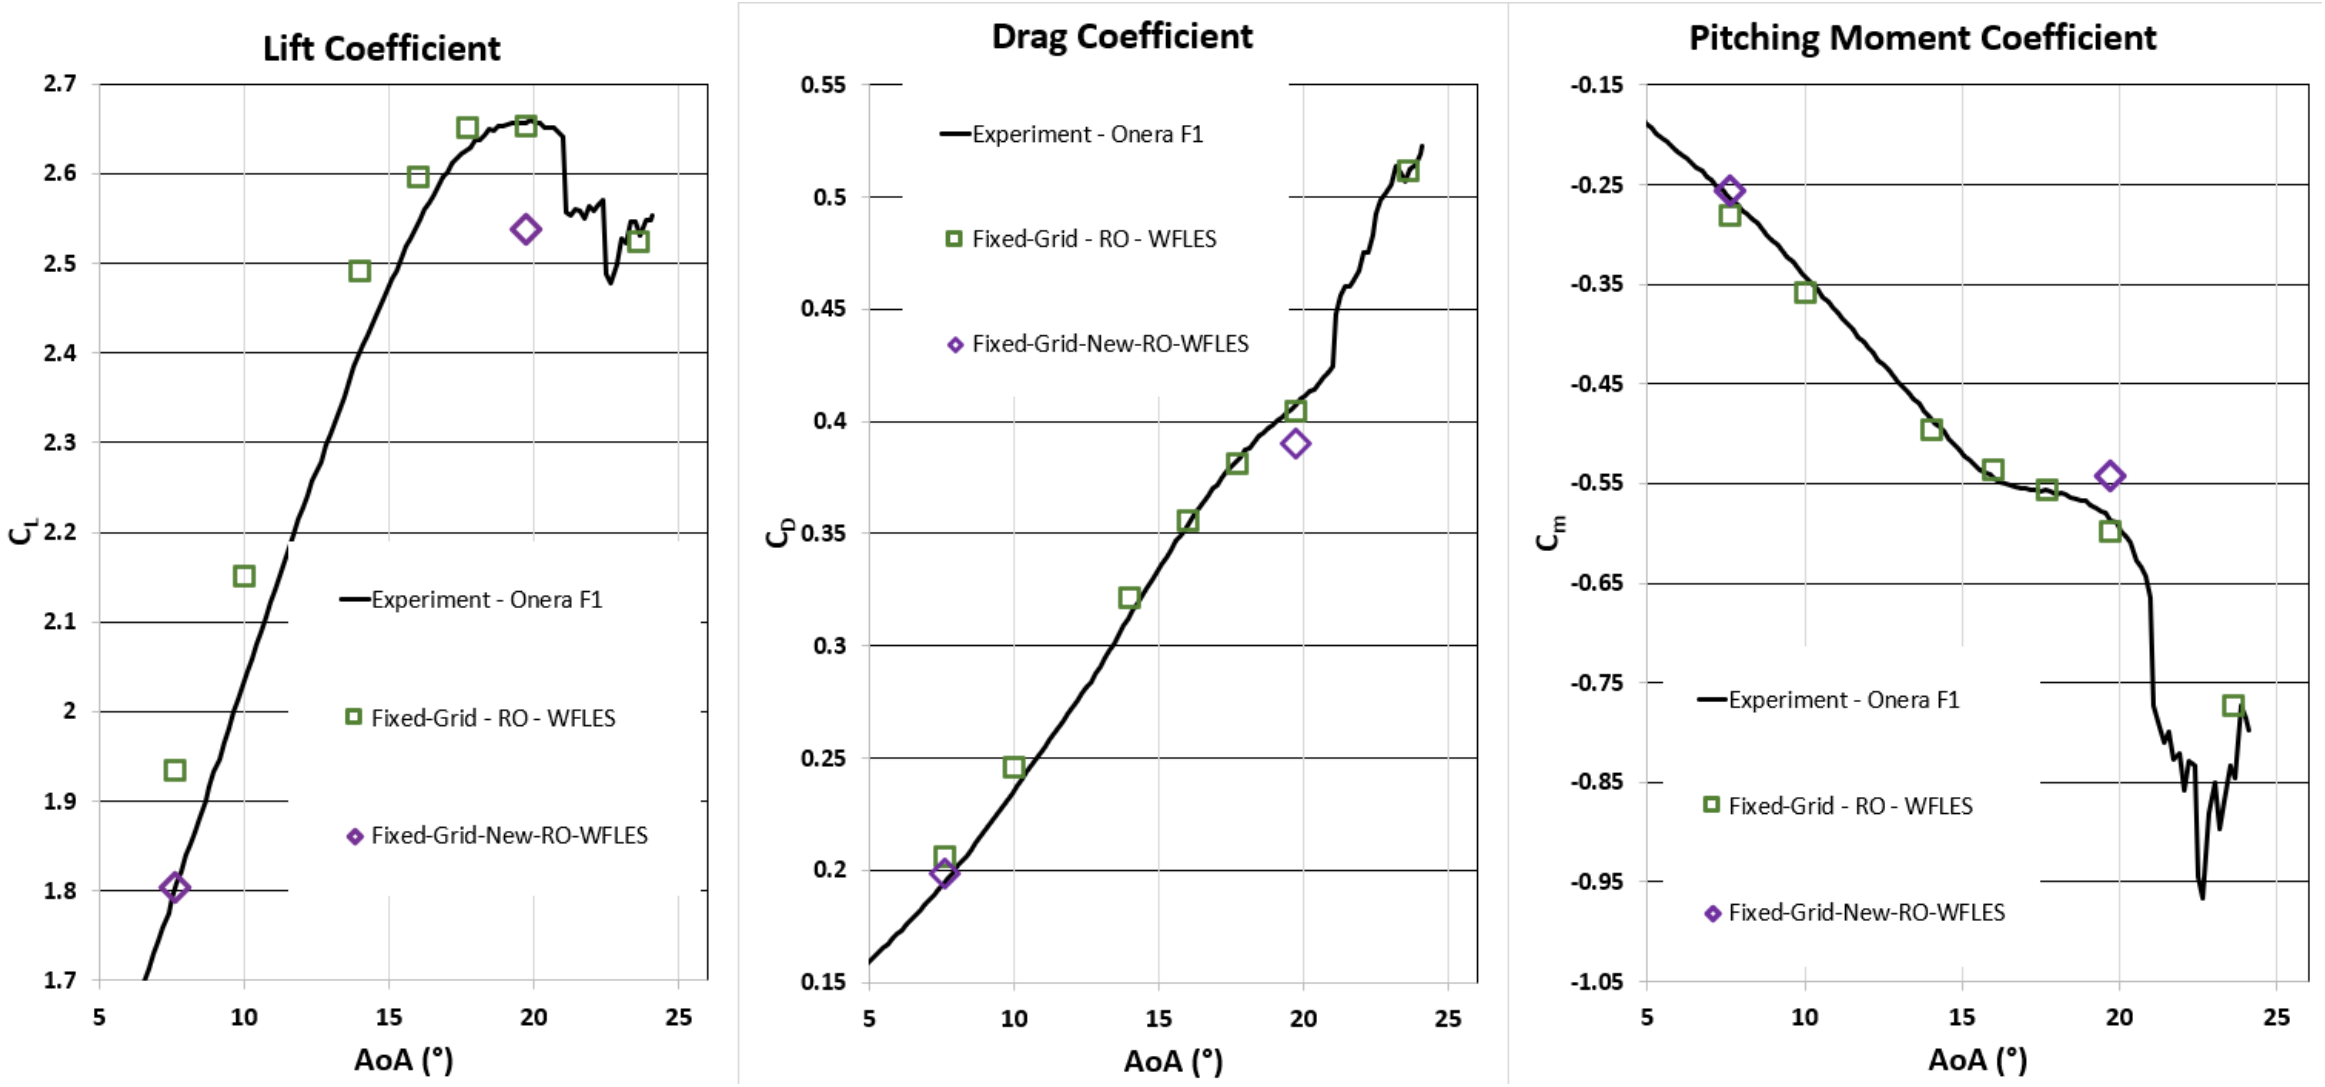


- Optimized LES numerics
- 2nd order Implicit time discretization
- Time step 4e-5 s / 2it per timestep
- 32 QuadL40 Nvidia cards (44 Gb)
- 1 hour / CTU (0.1s)



© 2026 Synopsys, Inc.

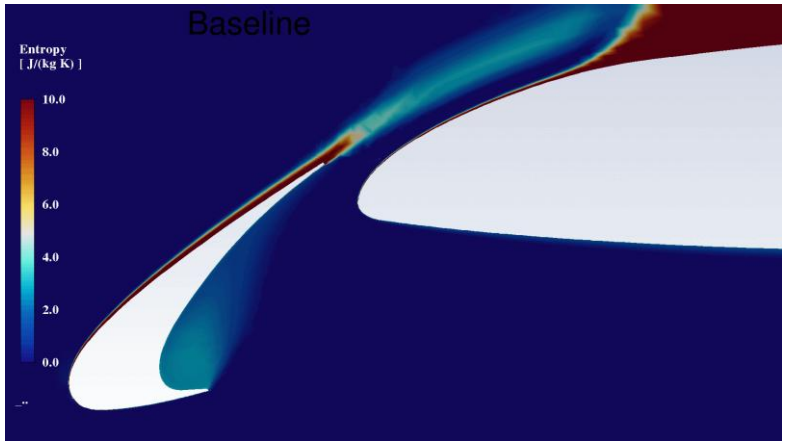
Aerodynamic Coefficient

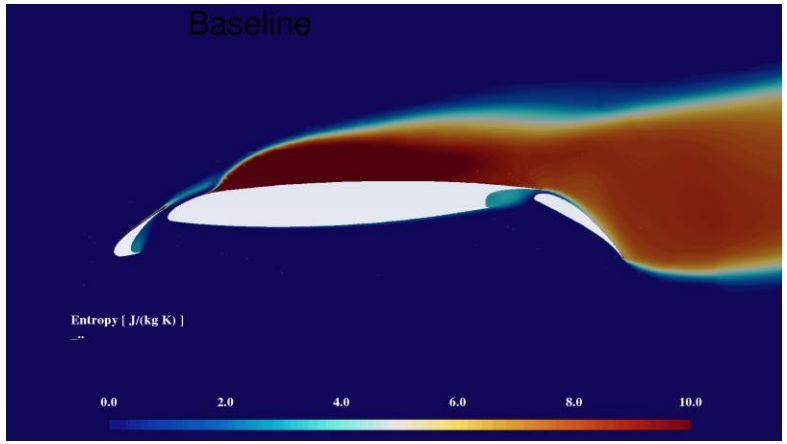


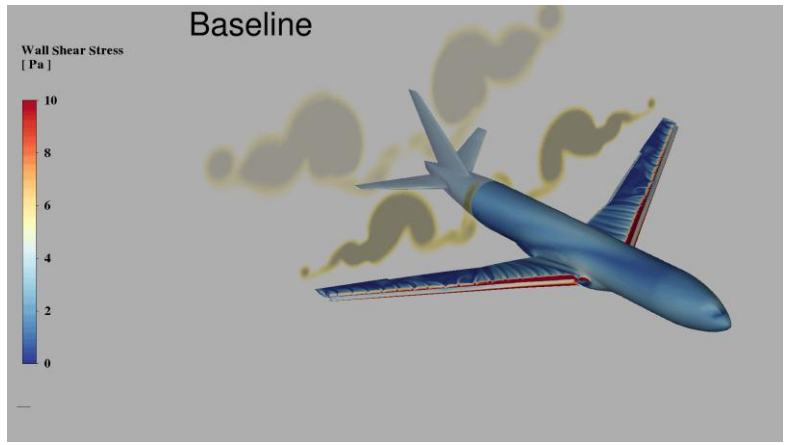
© 2026 Synopsys, Inc.

Case 2.3: RANS Mesh Adaption Solution - AoA = 20.7° - Stall AoA

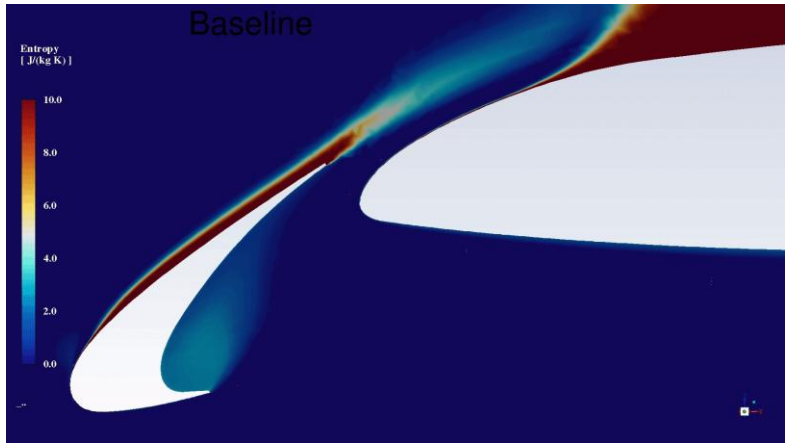
SA

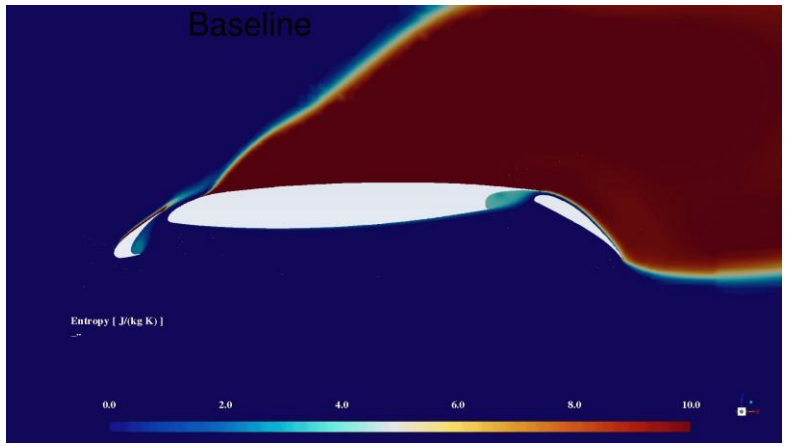






K-ω SST + γ-Re_θ







RANS Mesh Adaption & WFLES Simulations for High Lift Configuration Aircraft

Corrections for this article

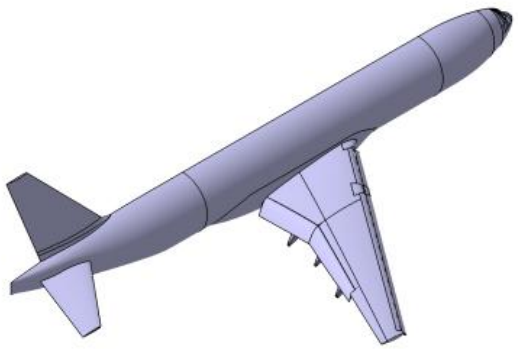
Cristhian Aliaga, Krishna Zore, Rachid Malk, Tobias Berg, Achilleas Krikas, Jeya Selva, Laith Zori and Boris Makarc

AIAA 2025-0498
Session: Special Session: High-Lift-Prediction-Workshop (HLPW-5): RANS and Adaptive Simulations (Invited)

Published Online: 3 Jan 2025 • https://doi.org/10.2514/6.2025-0498

Case 2.3: RANS Fixed Grid vs RANS Mesh Adaption

- Experimental results in solid lines
- Fluent Fixed Grid RANS SA results in dash lines
  - ANSA Level C (Fine) mesh
- Fluent PUMA Combined Hessian Indicator (RANS Primitive + Entropy) with RANS results in no fill markers
  - ANSA Level A (71 million cells) is used to start the adaption
  - No surface mesh adaption
  - Adapt ANSA Level A with SA in yellow squares
  - Adapt ANSA Level A with K-ω SST + γ-Re_θ in red triangles

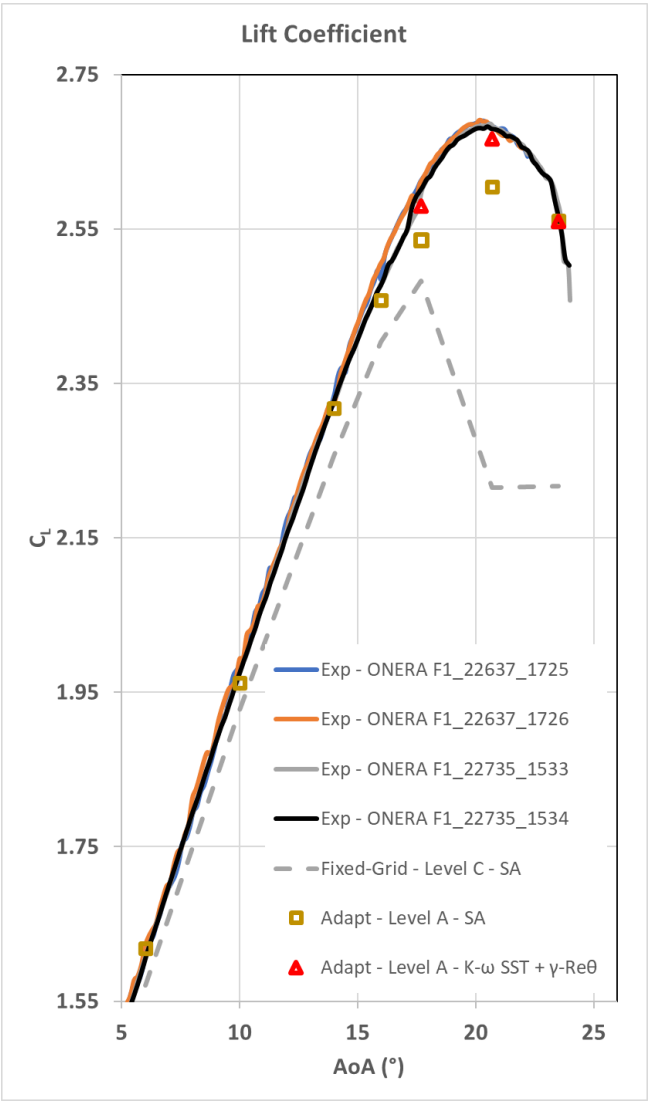


Case 2.3: ONERA_LRM-WBSFHV

Slats & Flaps Deployed

Mesh Size in Million Cells

	Coarse	Medium	Fine
Fluent SA	71	107	226
Fluent PUMA SA	~300		
Fluent PUMA K-ω SST + γ-Re_θ	~400		



© 2026 Synopsys, Inc.

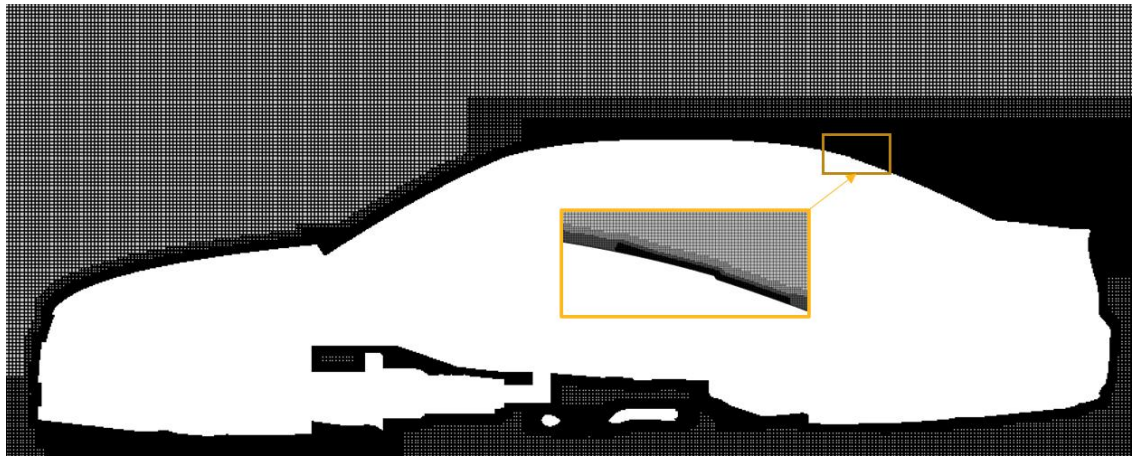
Auto CFD Workshop

5th Automotive CFD Prediction Workshop

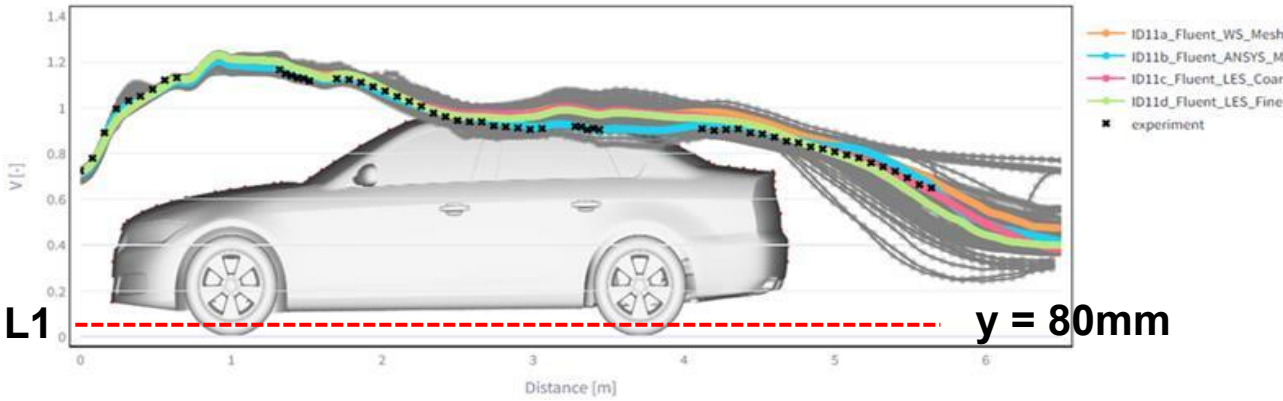
October 8th-9th 2026, Dama Technopolo, Bologna, Italy

- DrivAer Hatchback
- Showcasing our solution for auto external aero
  - Application of new, automated meshing
  - Emphasis on speed





Ansys solutions (colour) for WFLES and SBES against all workshop results and data

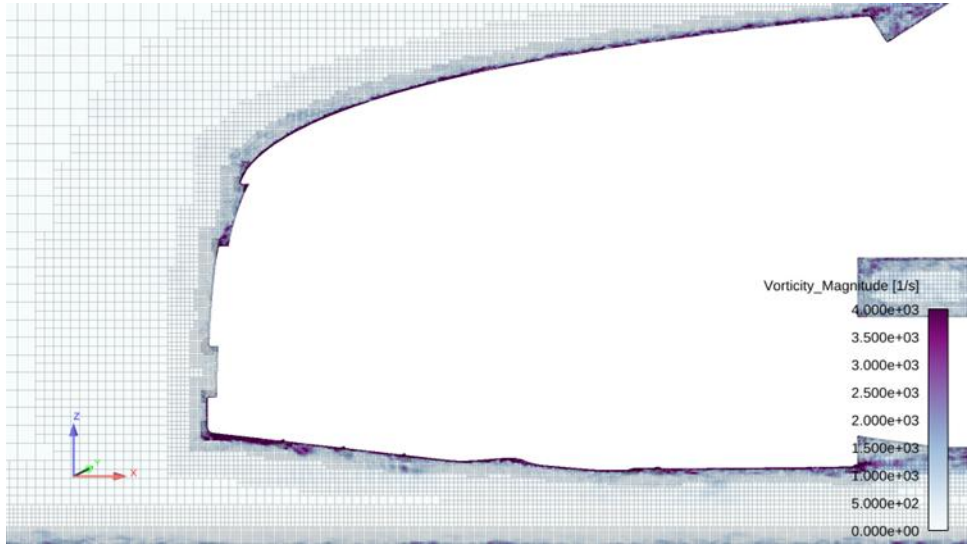


© 2026 Synopsys, Inc.

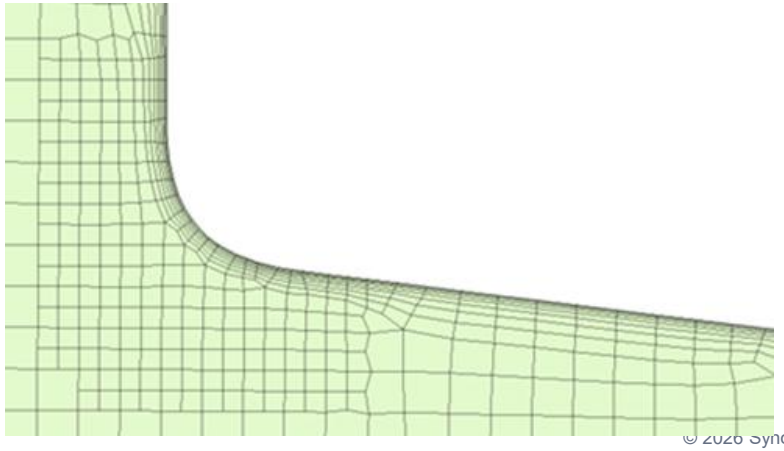
Work in progress for AutoCFD5 Workshop

- Run WFLES and SBES on ‘minimum meshes’
  - WFLES has 128 million cells
  - SBES has 135 million cells (same mesh + 8 prisms)
- Explore relative accuracy of both methods
- Use best method and explore potential for improved predictions
  - Overall forces
  - Delta between case 2a/2b variants

Base mesh for WFLES

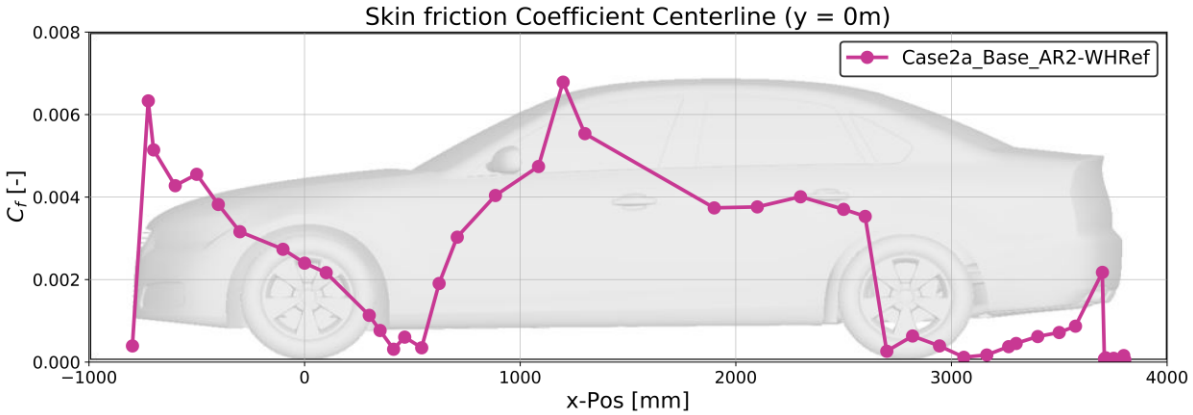


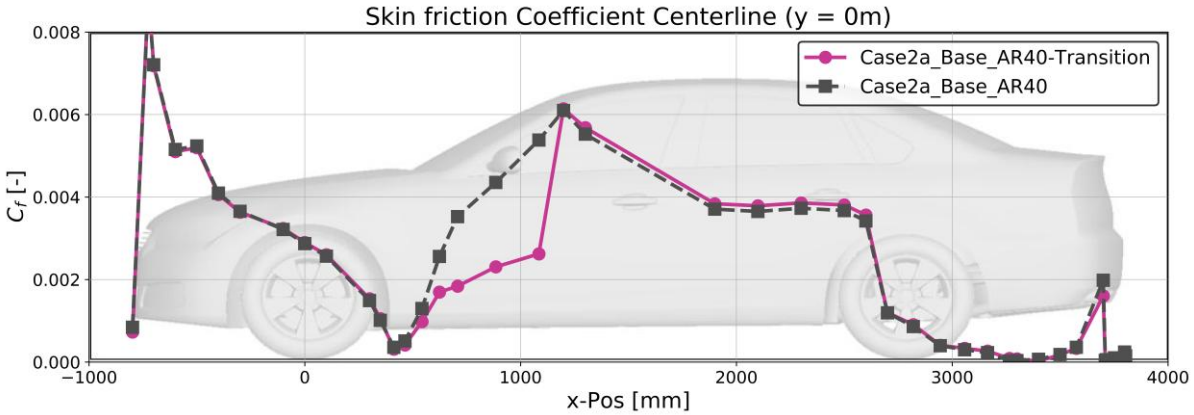
Base mesh + prisms for SBES



© 2026 Synopsys, Inc.

Skin Friction Coefficient





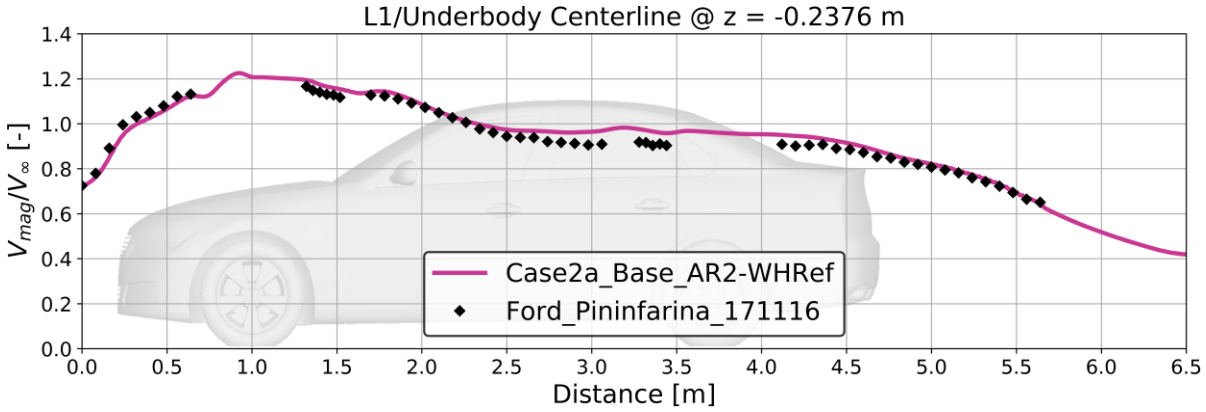
Consistency check for WFLES

- WFLES gives Cf very similar to SBES/RANS

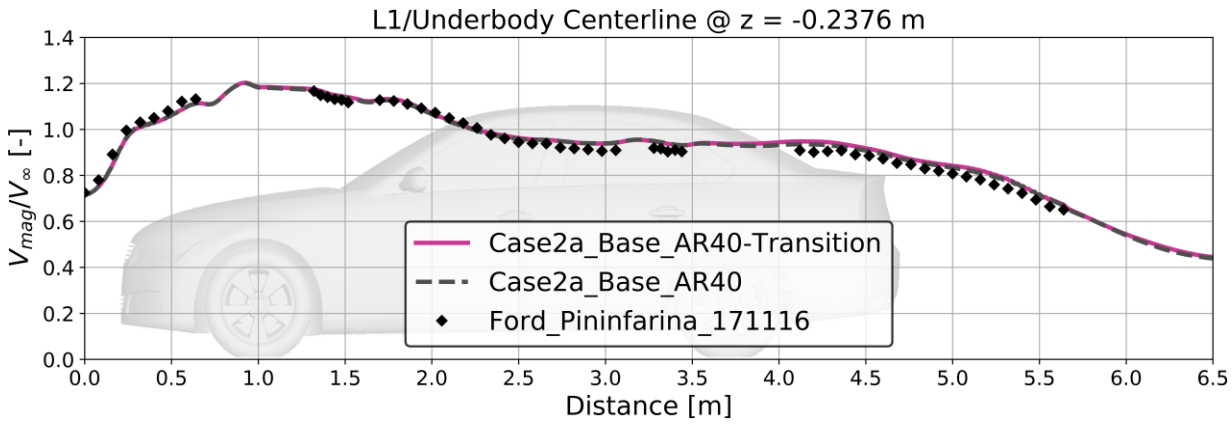
© 2026 Synopsys, Inc.

L1-Underbody Centreline case 2a

WFLES



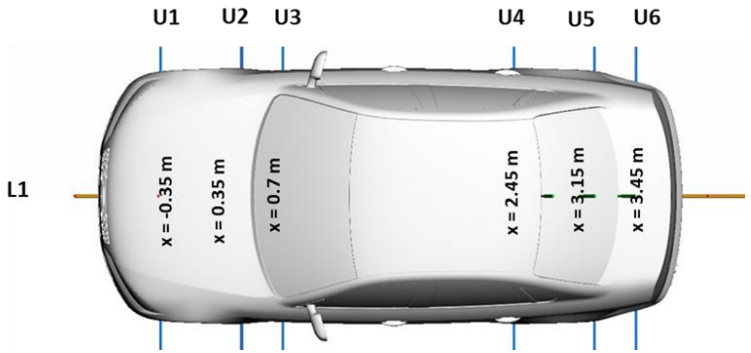
SBES





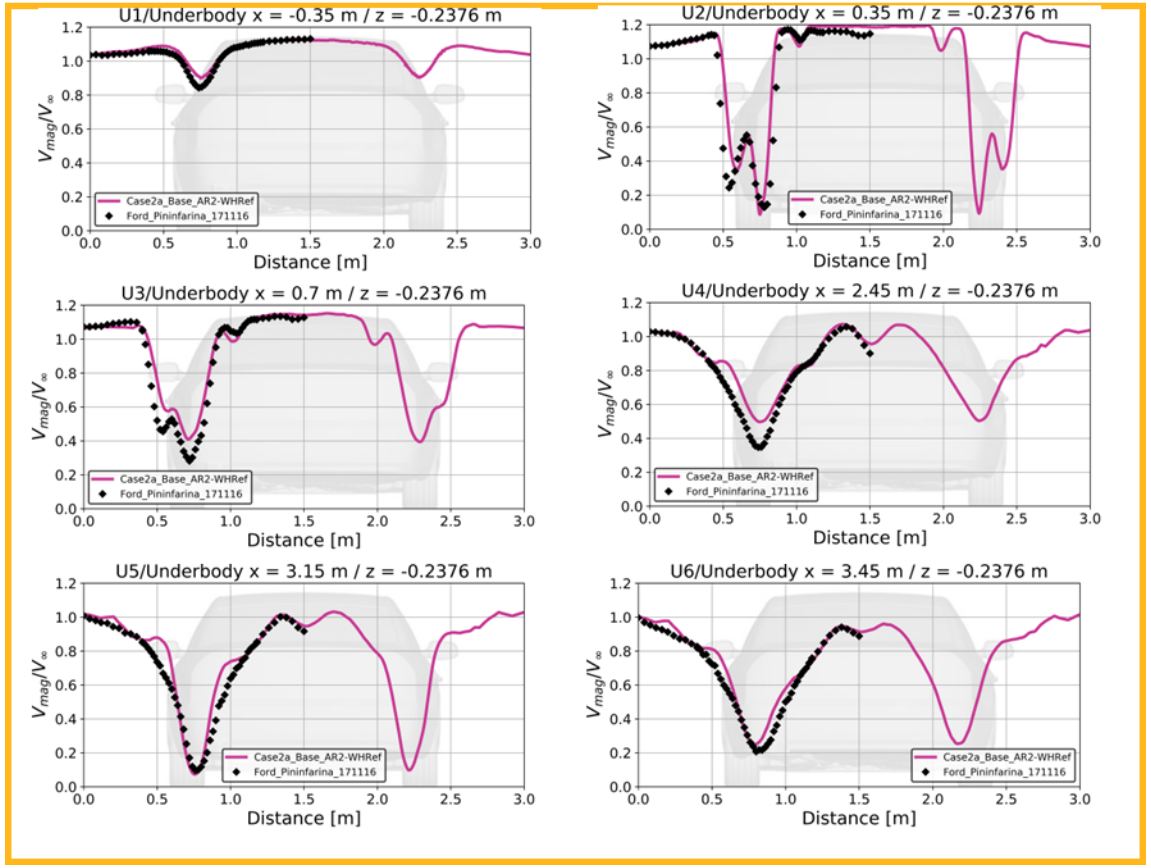
© 2026 Synopsys, Inc.

U-Velocity



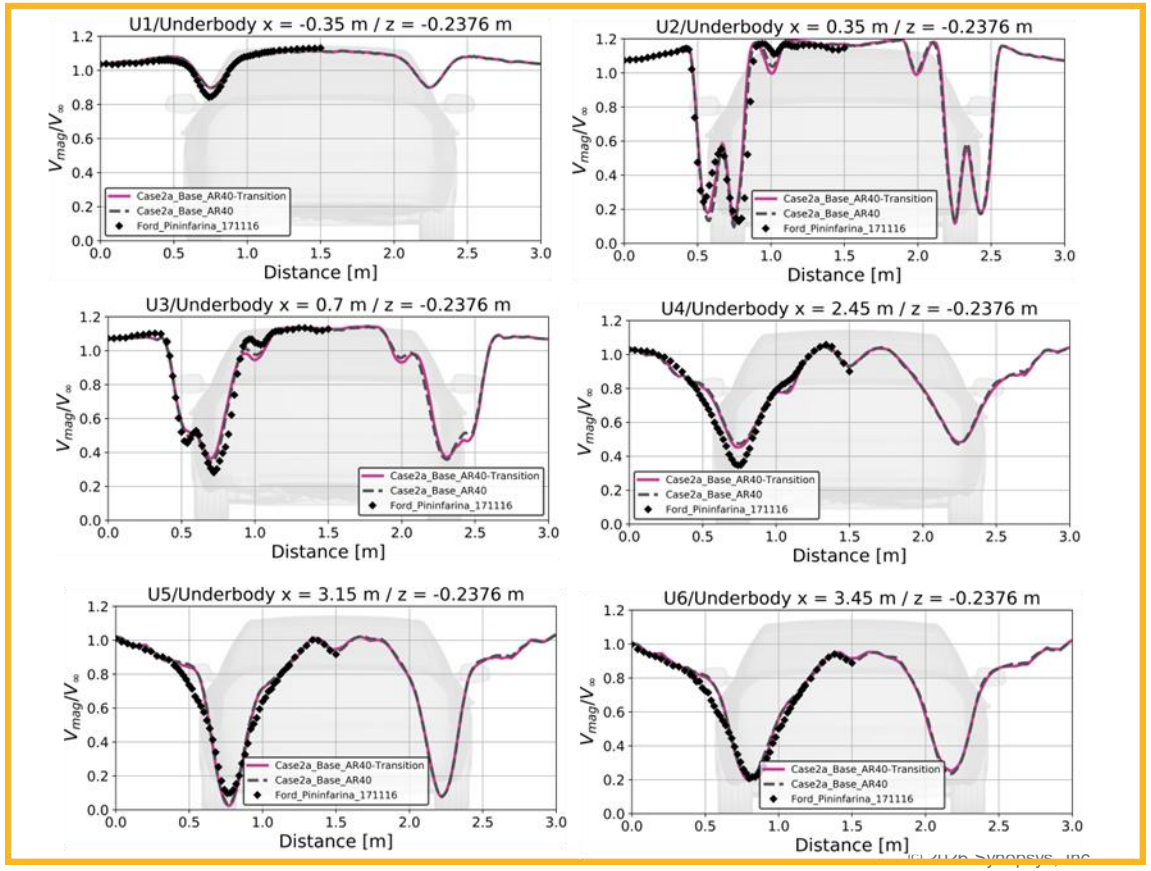
WFLES

Case2a



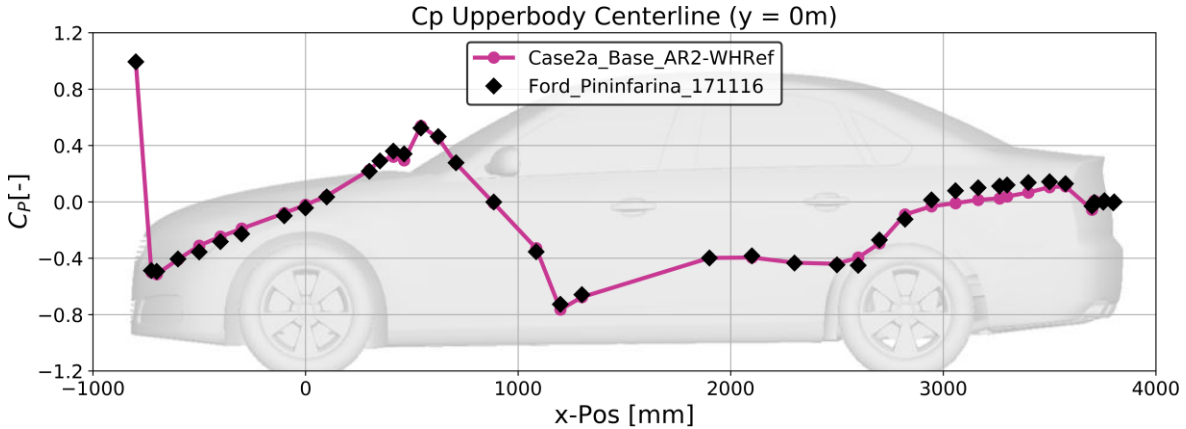
SBES

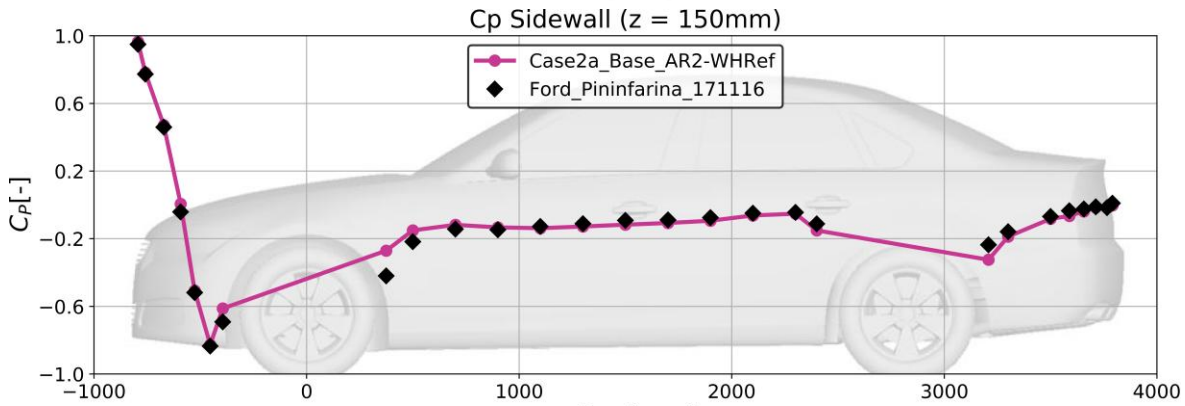
Case2a

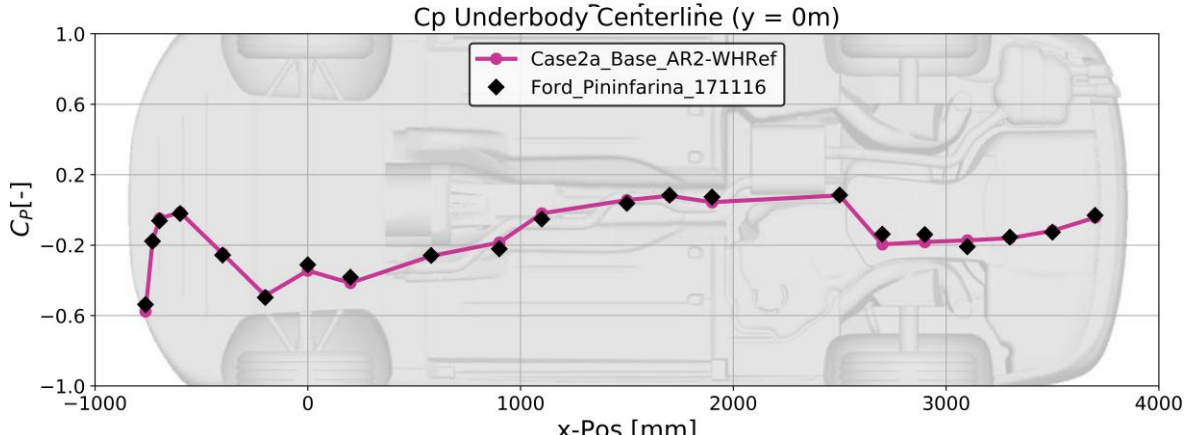


Surface Pressure

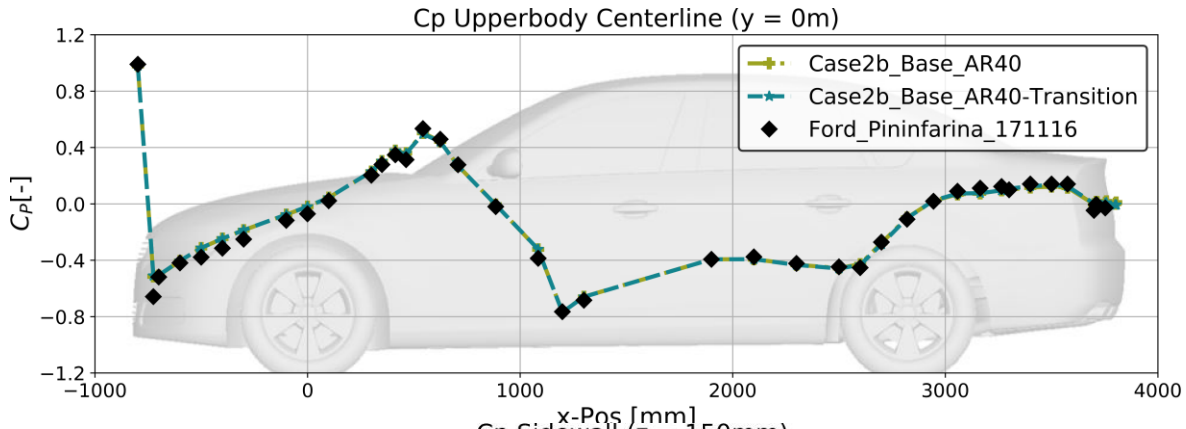
WFLES

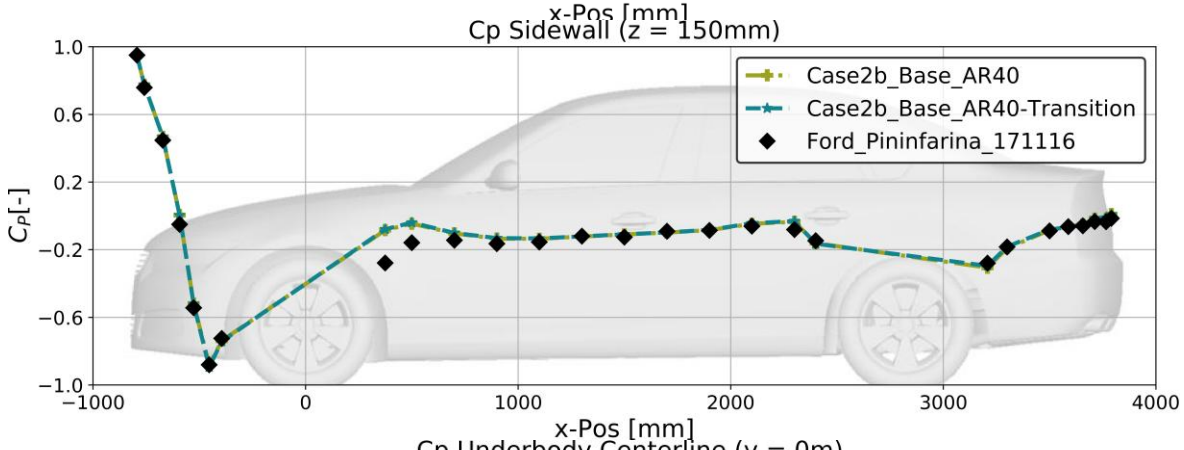


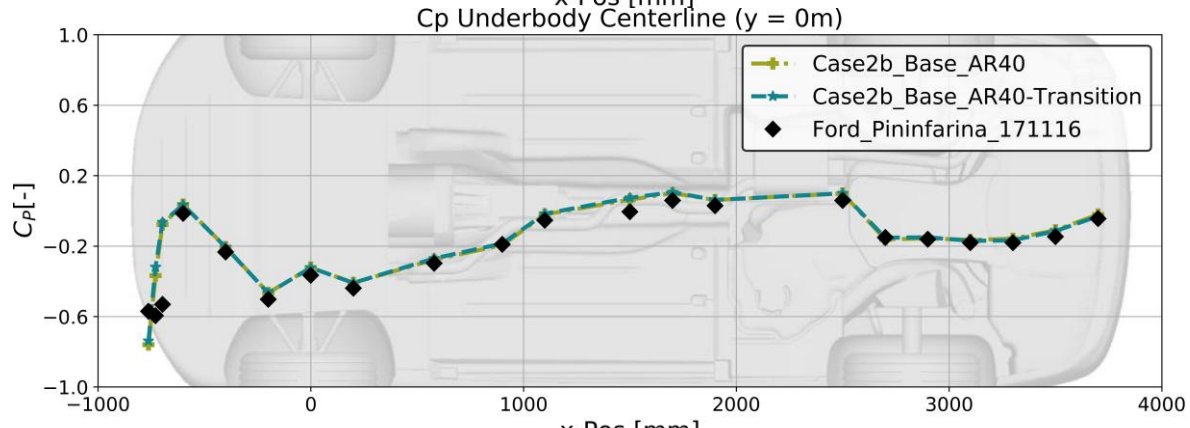




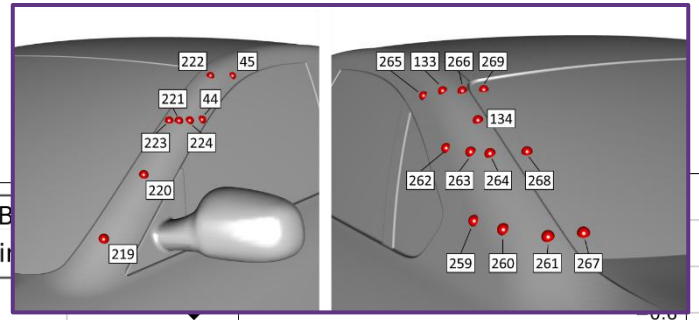
SBES



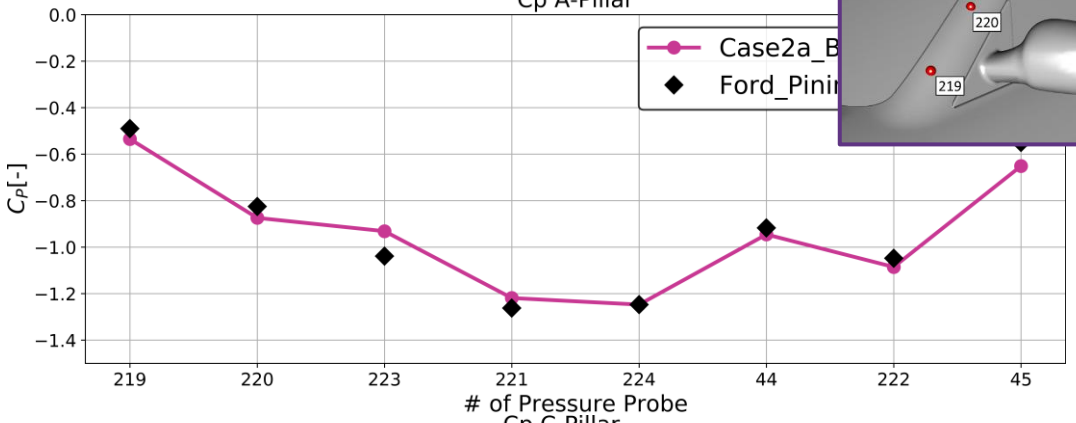




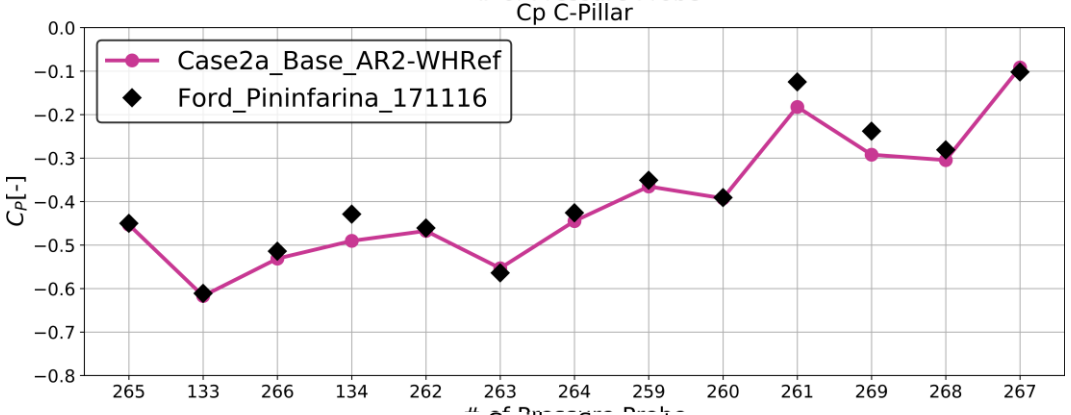
WFLES



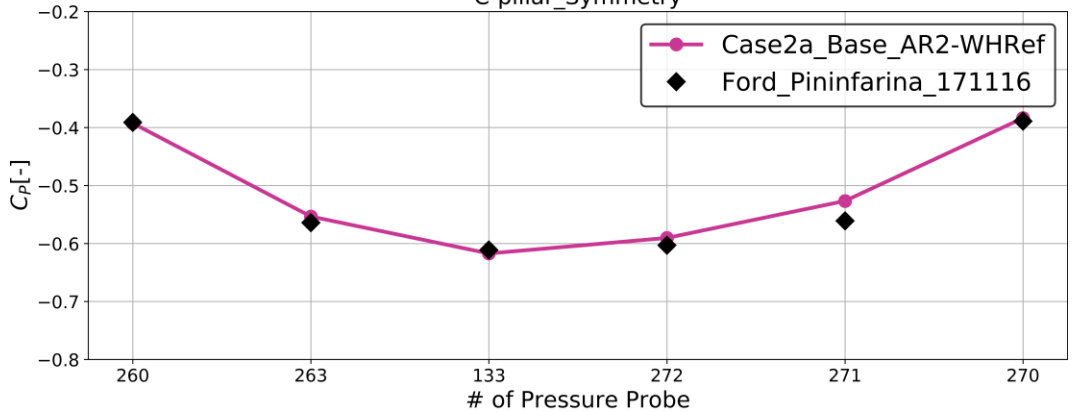
Cp A-Pillar



Cp C-Pillar

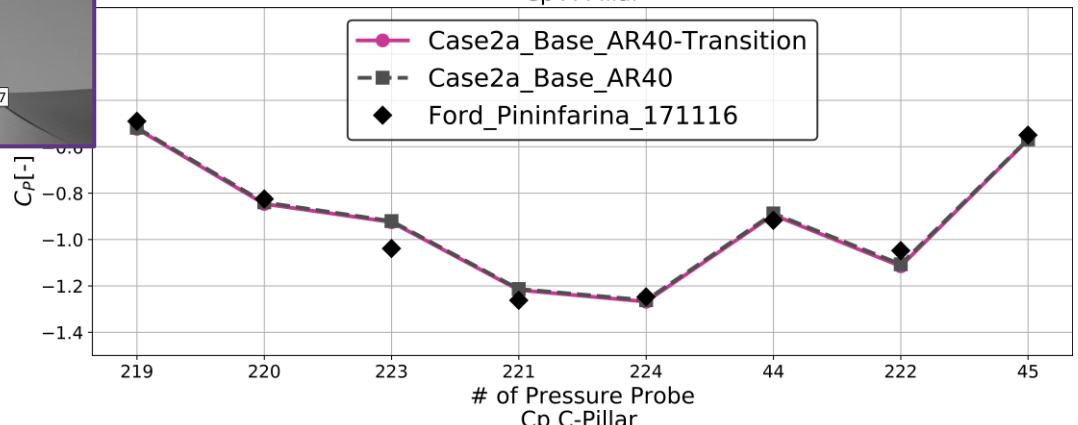


C-pillar_Symmetry

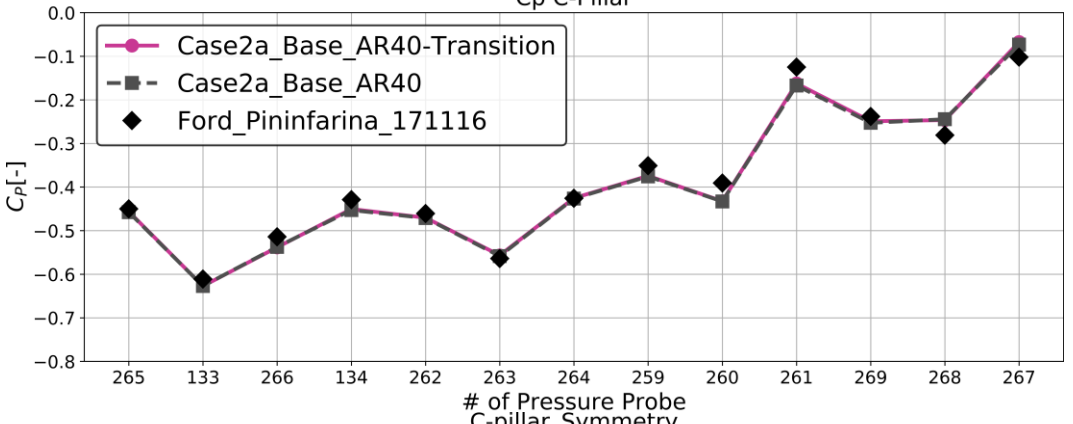


SBES

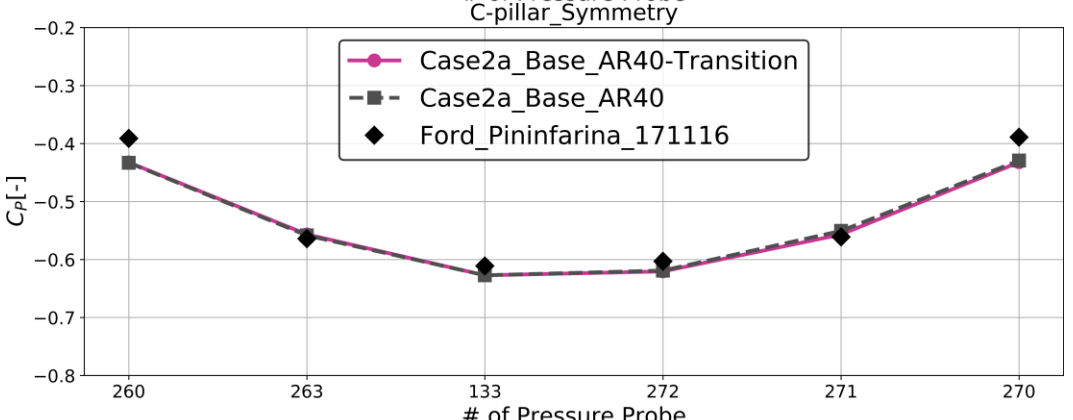
Cp A-Pillar



Cp C-Pillar

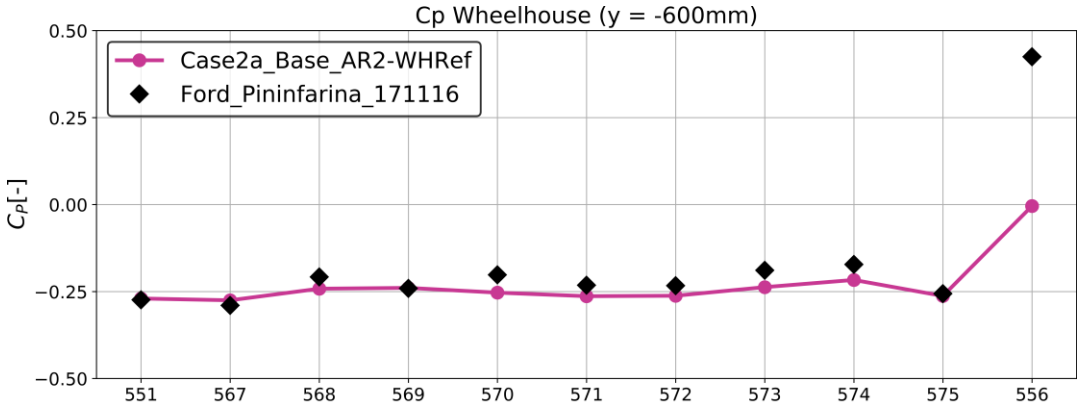


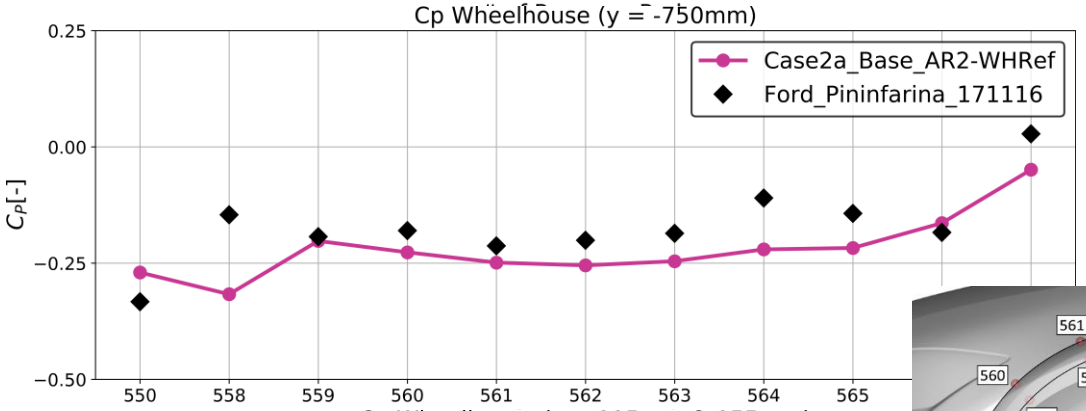
C-pillar_Symmetry

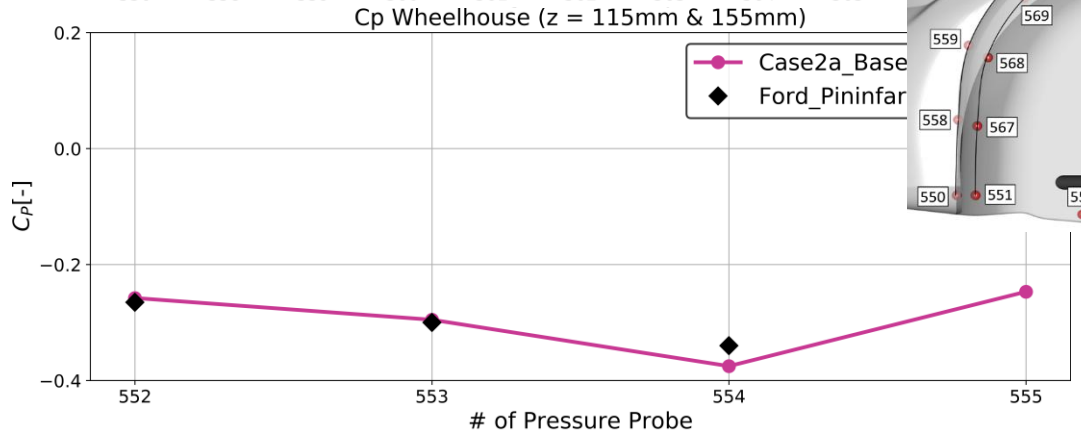


Wheelhouse

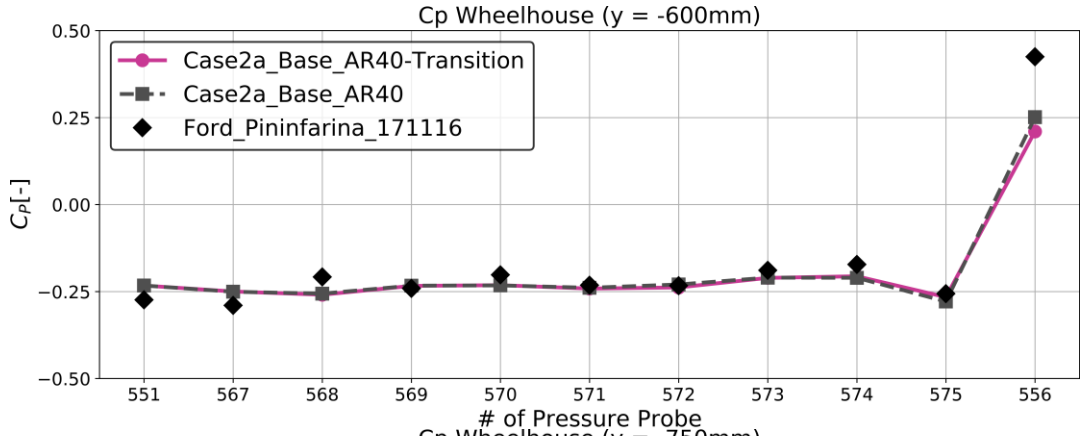
WFLES

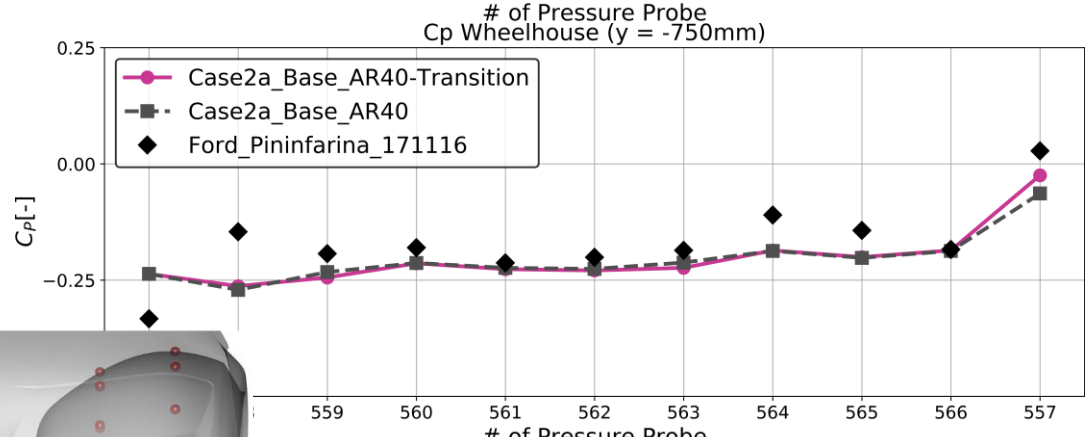


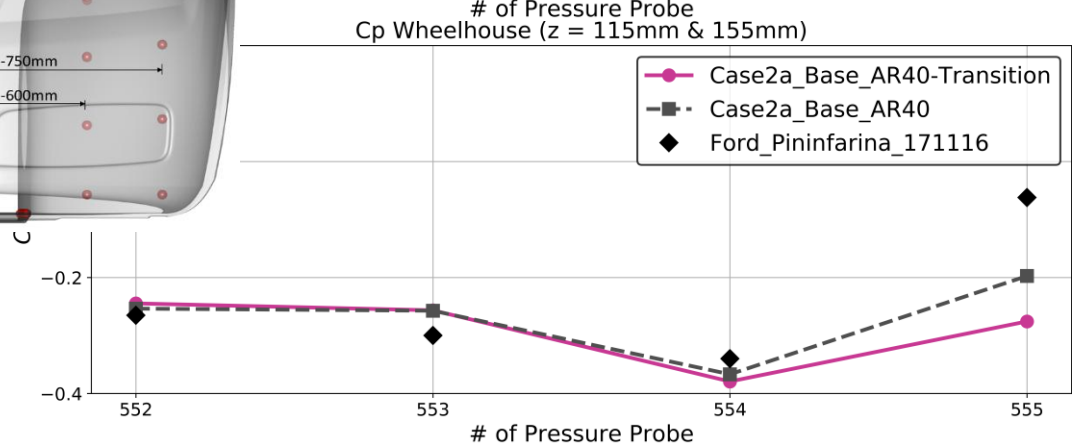


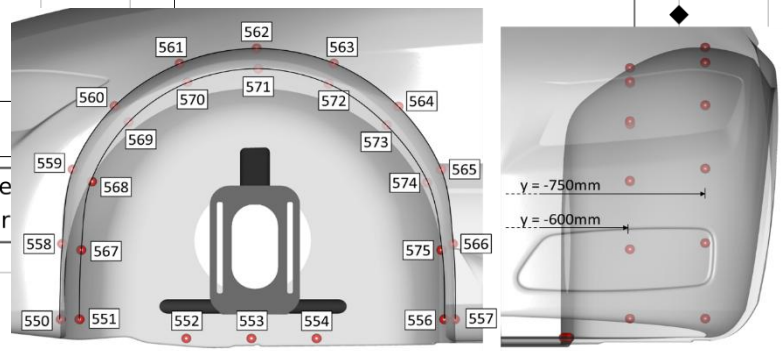


SBES





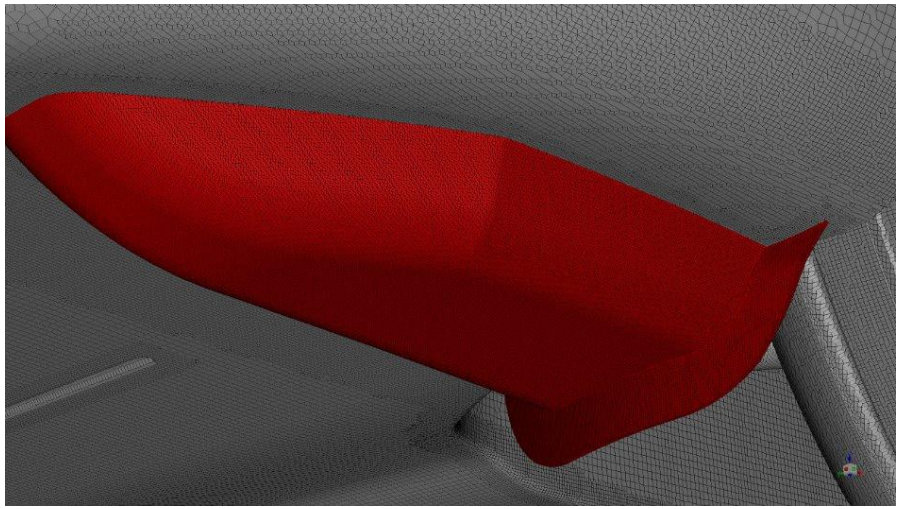




Aerodynamic Coefficient

- There was a tendency in the workshop to over-predict drag
- Differences in C_D by bulk of participants currently not understood
- Current results on minimum meshes
- Goals:
  - Get closer on overall drag
  - Get closer of ΔC_D

Front Wheel Deflector in 2b

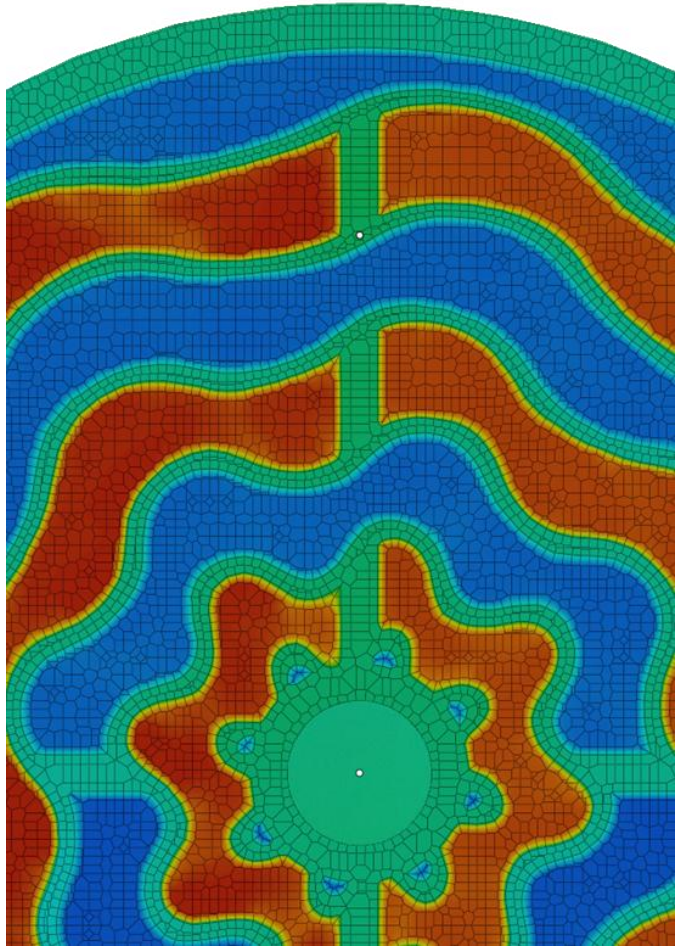


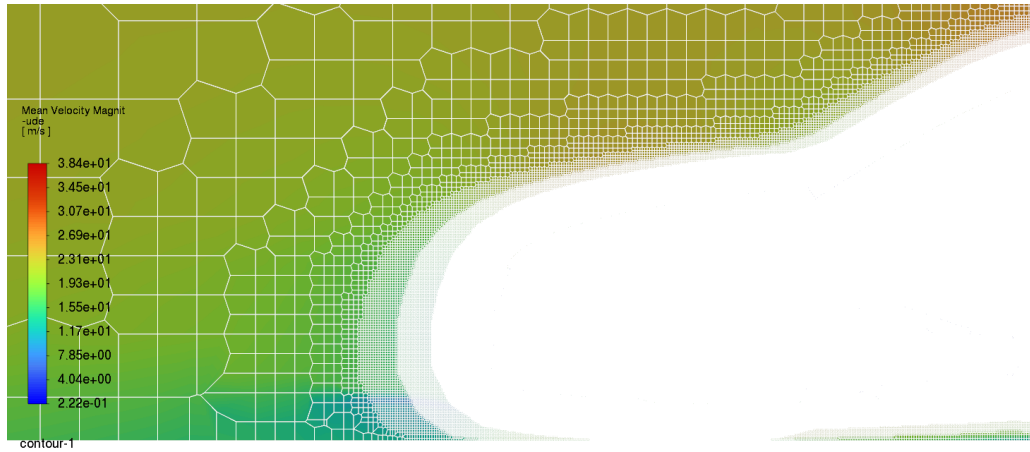
C_D0.267	2a	2b	ΔC_D
Exp	0.2546	0.2420	0.0126
WFLES	0.2670	0.2600	0.0070
SBS	0.2700	0.2627	0.0073

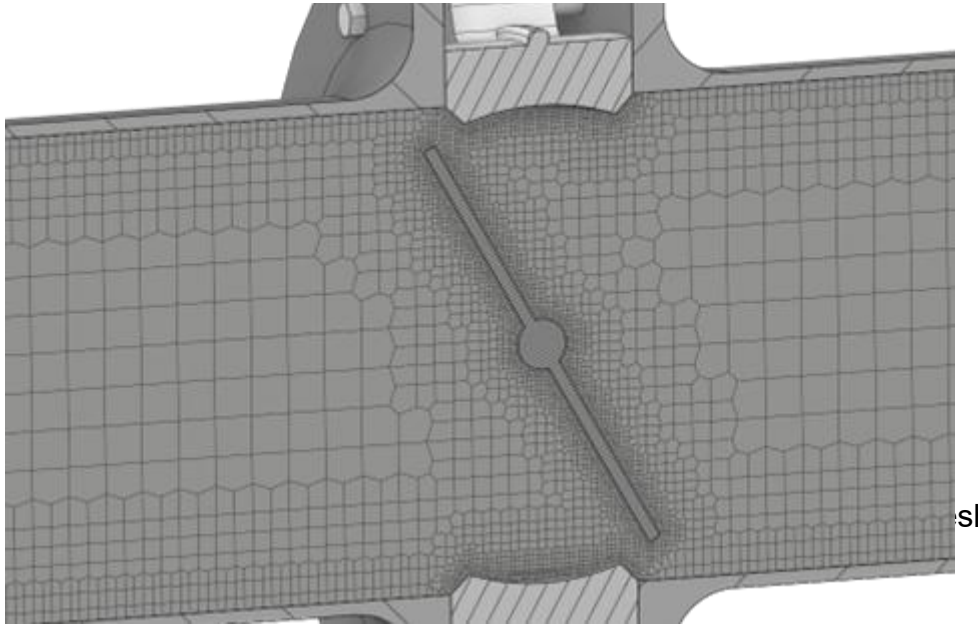
© 2026 Synopsys, Inc.

Groundwork for MultiGPU – Octree/Voronoi Meshing

- Top-down meshing
  - Both Octree and Voronoi technologies under development
  - Highly scalable (1000s of CPUs, GPUs)
  - Can handle ‘dirty’ geometries





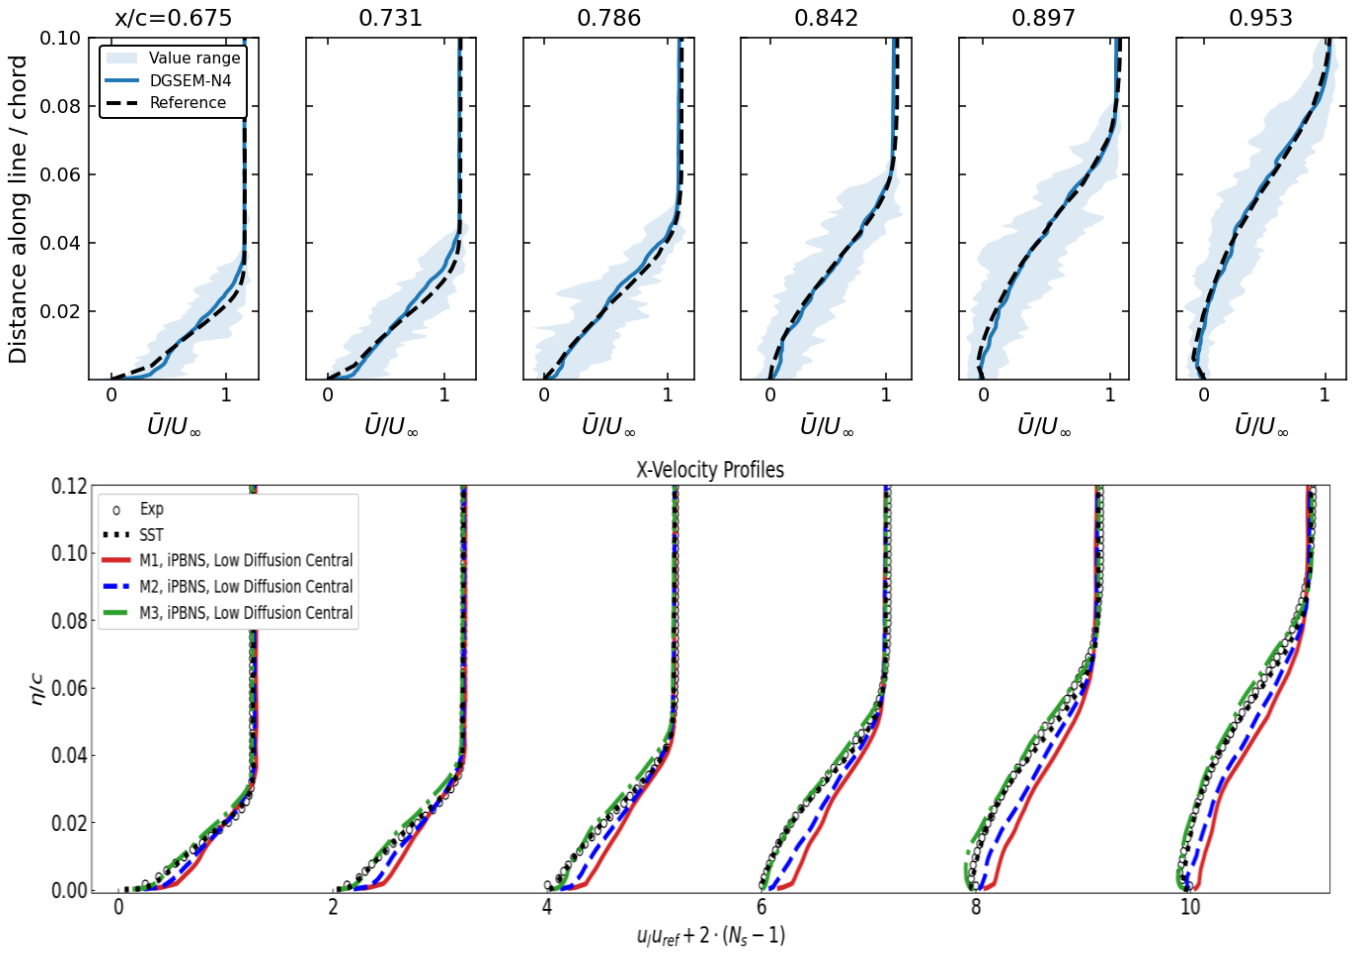


	1 GPU	2 GPUs	3 GPUs	4 GPUs
13M Cells	1.19 s	1.06 s	0.92 s	0.77 s
69M Cells	-	4.5 s	3.65 s	2.28 s

© 2026 Synopsys, Inc.

Higher Order Discontinuous Galerkin Prototype - NACA4412

- Prototype development of higher order DG method
- First results for NACA 4412 airfoil
- 3 mm mesh 4th order DGSEM results comparable to 0.4 mm mesh 2nd order results
- Compared to: 2nd order FV, central scheme, mesh convergence study
  - 1.6 mm ; 0.8 mm ; 0.4 mm meshes



A GPU-native high-order DG solver for Large-Eddy-Simulation

David G. Flad, Andreas Hueppe, Alex Main and Vito Pasquariello

AIAA 2026-2333
Session: High-Order Numerical Methods

Published Online: 8 Jan 2026 • https://doi.org/10.2514/6.2026-2333

© 2026 Synopsys, Inc.

Overview Synopsys GPU-based CFD Solvers

- Pressure-based Navier Stokes (PBNS) solver
  - Coupled
  - Segregated (SIMPLEC,...)
  - CPU and GPU
- Density-based Navier Stokes (DBNS)
  - Implicit coupled
  - Explicit Runge-Kutta (KEEP) scheme (prototype)
- Higher Order DG solver (R&D)
  - Density-based
  - Explicit
- SPH, DPM, ...

© 2026 Synopsys, Inc.

LES Summary

- GPU solver is ‘enabling’ technology
- GPU solver allows resolution of turbulence eddies in boundary layers for industrial flows
- Reynolds number scaling of LES can be reduced by WFLES/WMLES
- Highly optimized numerical methods essential
- Laminar-Turbulent transition is an open question both for WFLES and WMLES
- User training is non-trivial
  - Meshing strategy different from RANS/SBES
  - Averaging requires flow-dependent process
  - Mesh independence hard to achieve

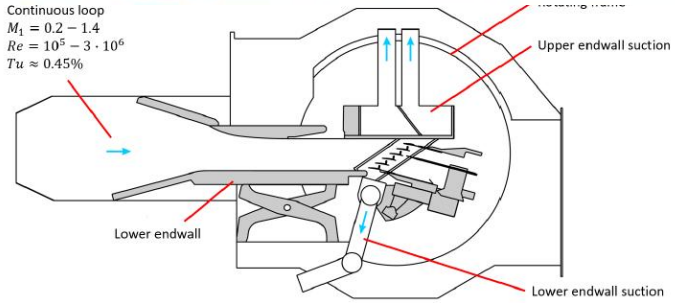
© 2026 Synopsys, Inc.

Overview Test-Cases

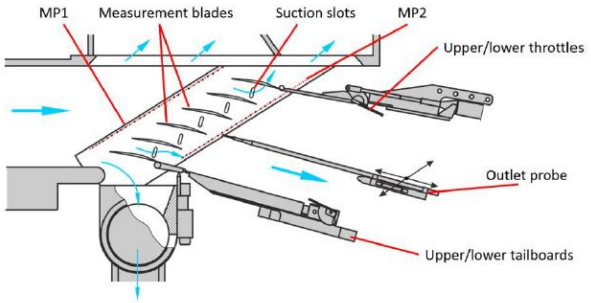
- Transonic Cascade Wind Tunnel, DLR



1. **L**ow **R**eynolds **N**umber **O**utlet **G**uide **V**ane (LRNOGV)
   - $Ma_{in} = 0.6$
   - $Re_{in} = 150.000$



2. **T**ransonic **C**ascade **T**eam**A**ERO (TCTA)
   - $Ma_{in} = 1.21$
   - $Re_{in} = 1.350.000$



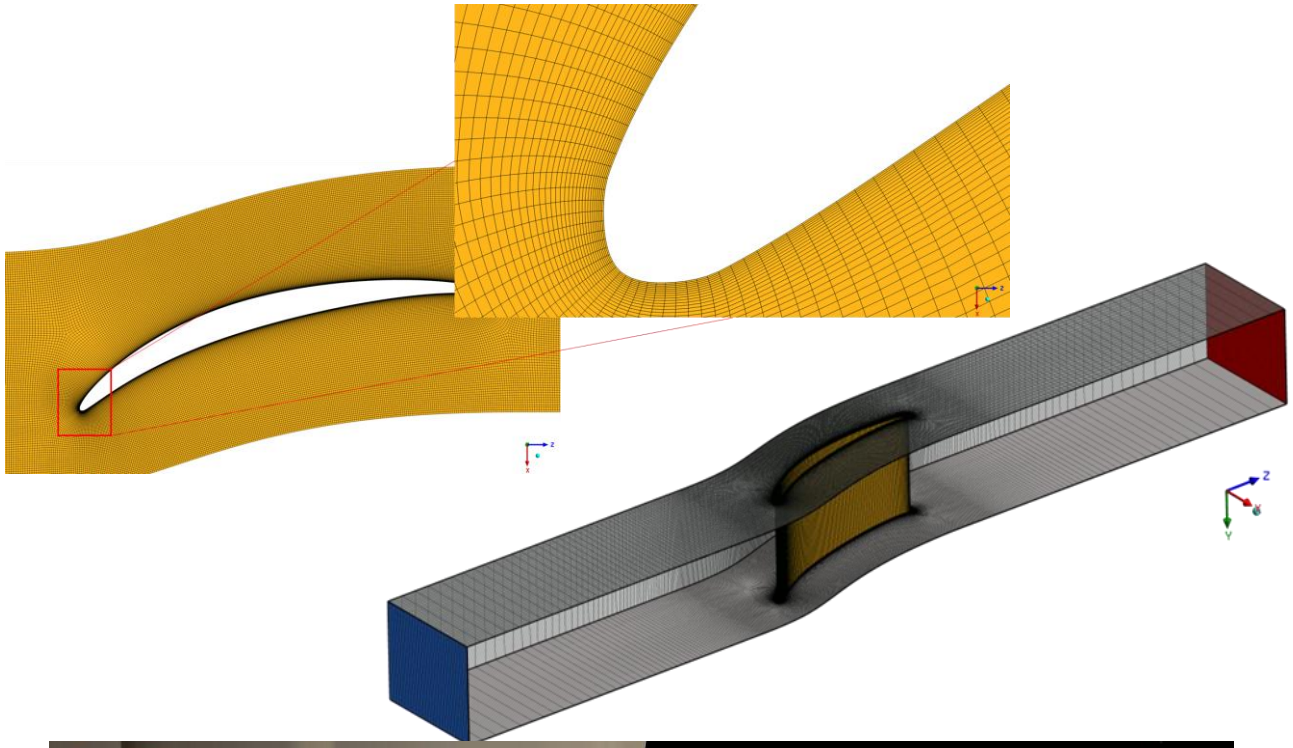
© 2026 Synopsys, Inc.

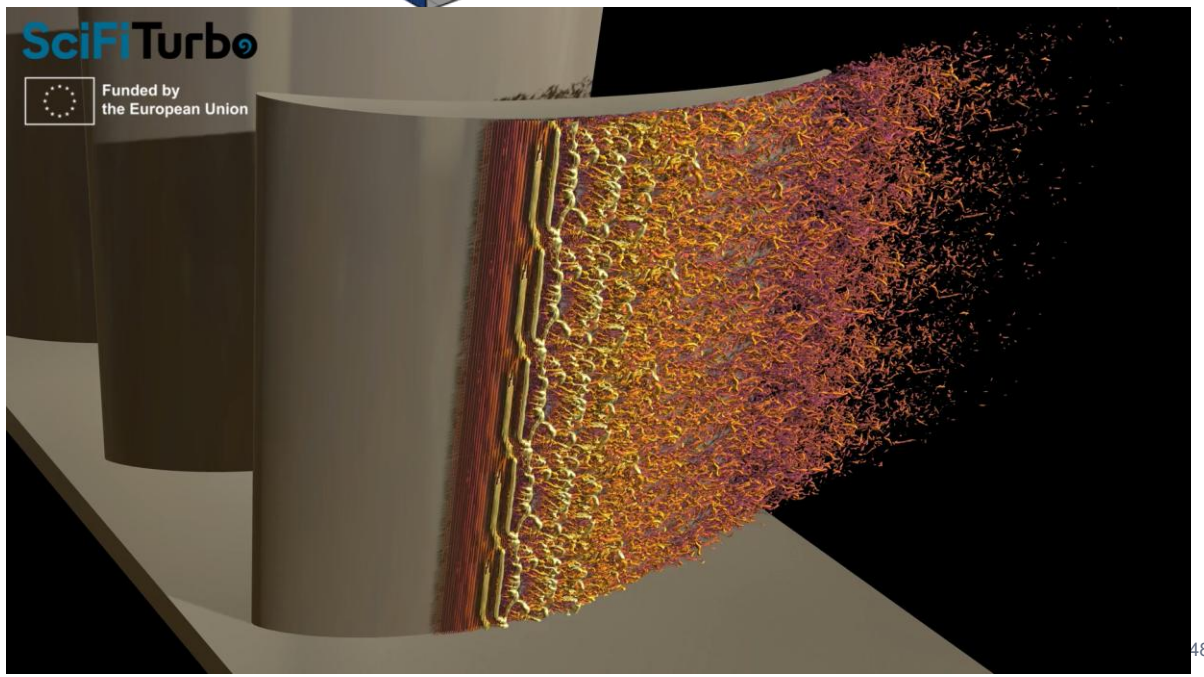
Baseline Mesh

Ansys TurboGrid: Structured Mesh

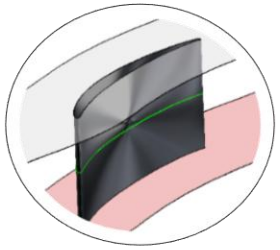
Mesh Properties

- Size: 9.4M cells
- Mesh Metrics
  - Min. Orthogonal Quality: 0.92
  - Max. Aspect Ratio: 56.3
- RANS: $y^+_{avg} = 0.39$, $y^+_{max} = 1.7$

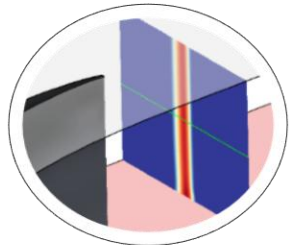


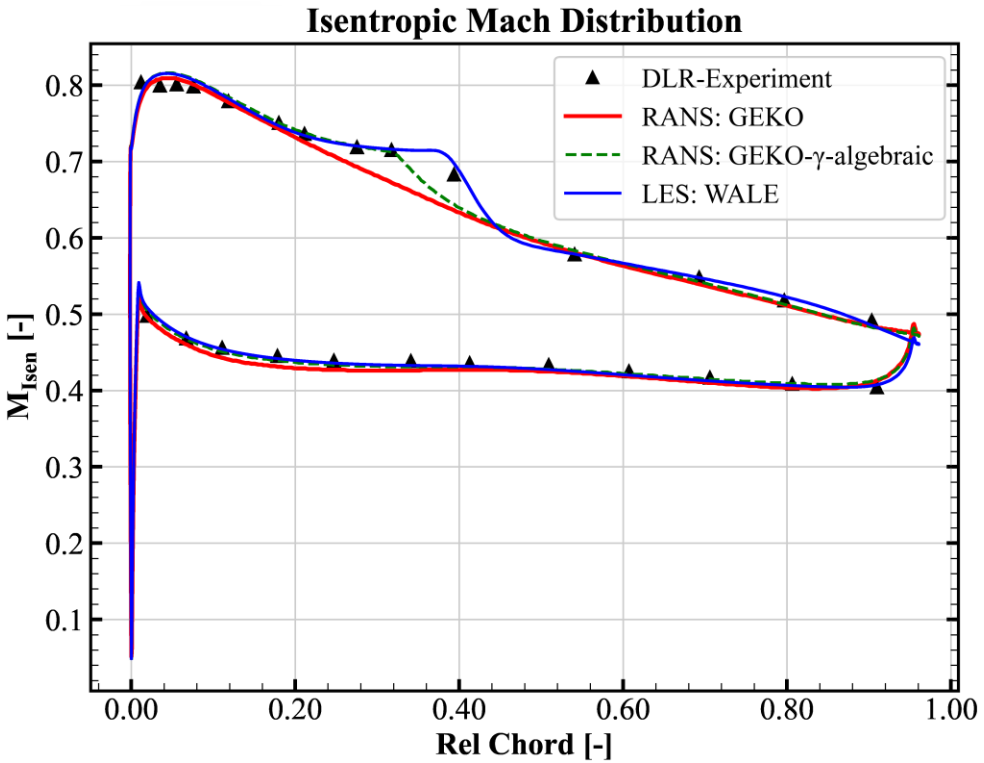


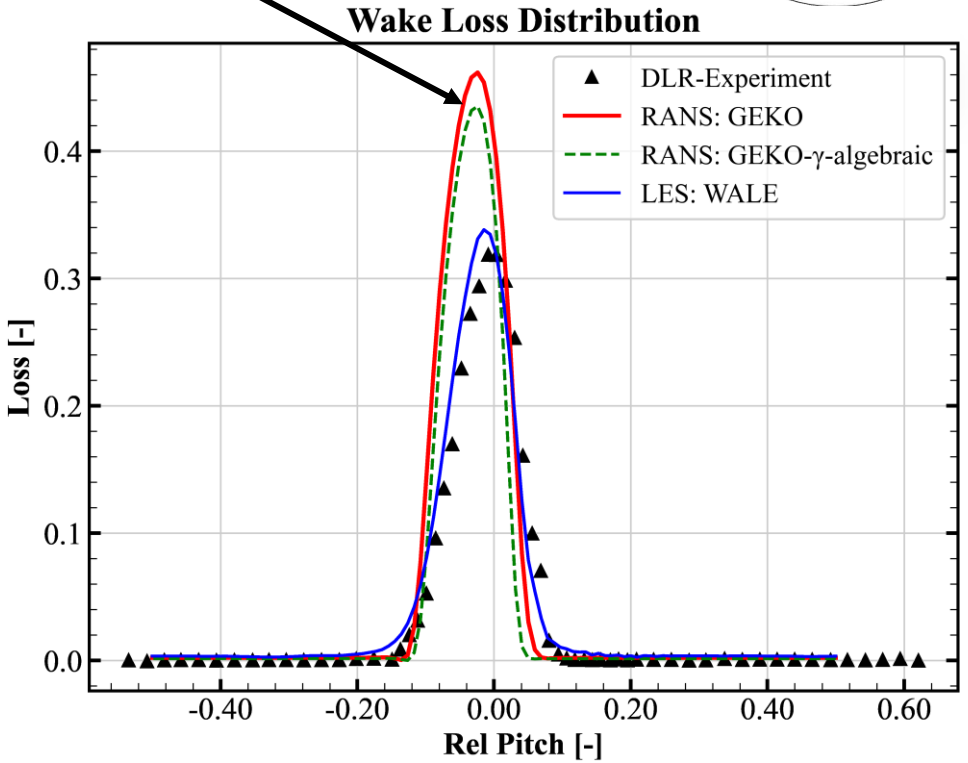
Impact Turbulence Model



RANS Models over-predict profile losses







$$M_{Isen} = \sqrt{\frac{2}{\gamma - 1}\left[\left(\frac{p_{t,1}}{p}\right)^{\frac{\gamma-1}{\gamma}} - 1\right]}$$

$$\omega = \frac{p_{t,1} - p_t}{p_{t,1} - p_1}$$

© 2026 Synopsys, Inc.