NASA's LAVA CFD Solvers: HPC Perspective

Michael Barad, Emre Sozer, Jeffrey Housman, Jared Duensing

Launch Ascent and Vehicle Aerodynamics (LAVA)

Computational Aerosciences Branch (TNA)

NASA Ames Research Center

NASA/DOE-ECP Collaboration Meeting
April 11, 2023

POC: Jared.C.Duensing@nasa.gov

LAVA CFD Solvers: Introduction

- History
  - LAVA has been developed at NASA Ames since about 2009
  - Based on a long history of CFD at NASA and DOE
  - Focus on Navier-Stokes equations
  - We develop our tools to directly address mission critical flow physics
- LAVA has provided critical support to NASA missions, which require:
  - High accuracy, performance, parallel scalability
  - Support for high geometric complexity
  - Multiphysics: multispecies, multiphase, fluid-structure, heat transfer, acoustics etc
  - Rapid turn around, from CAD model to solution
- Current Implementation Details:
  - Mixed C++/Fortran
  - MPI/OpenMP parallelism for CPU
    - Cache management and explicit vectorization for performance
    - Overlapped MPI comm/compute
  - MPI/CUDA parallelism for GPU
    - Single kernel launch for each of the main operators: : apply boundary conditions, evaluate residual.
    - Streams for comm/no-comm MPI overlap
    - Careful shared memory usage for caching

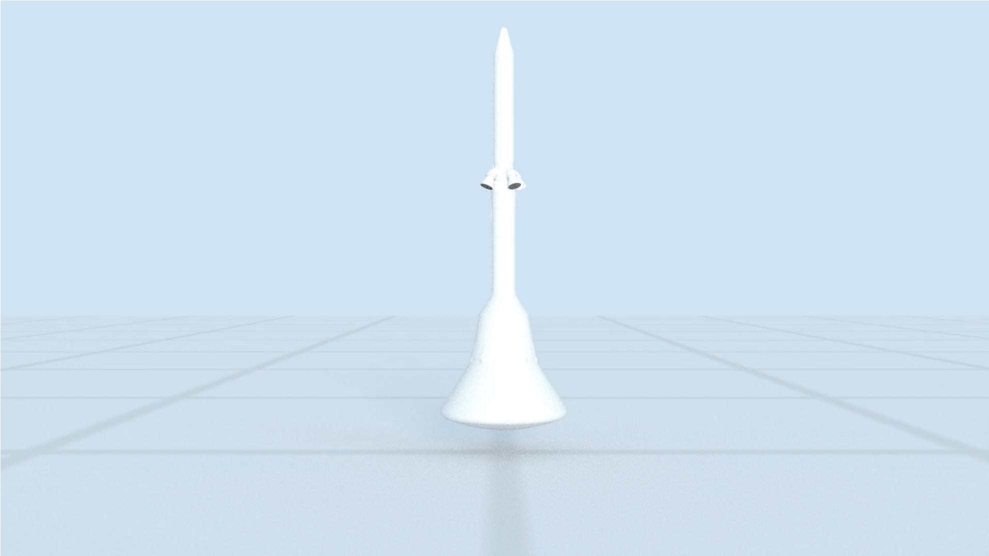

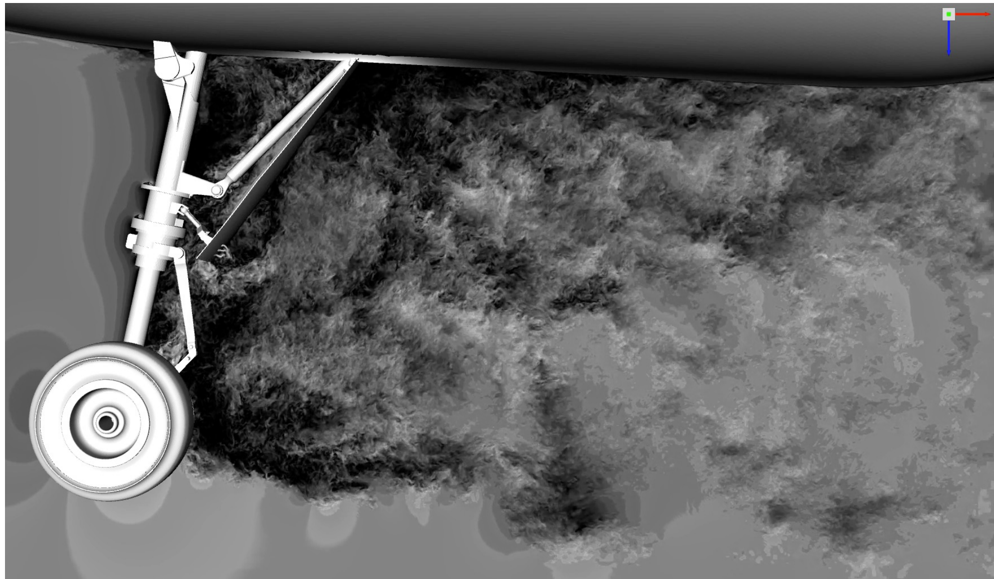

LAVA CFD Solvers: Capability Comparison

	Cartesian	Curvilinear	Unstructured
Meshing	Automated/AMR/Parallel	Manual	Automated
Wall Modelling	Immersed boundaries	Body-fitted	Body-fitted
Spatial Discretizations	2nd, 4th, 5th, 6th order	2nd, 4th order	2nd order
Time Discretizations	Explicit, 2nd, 3rd, 4th order	Explicit 3rd order or Implicit 2nd order	Explicit 3rd order or Implicit 2nd order
Body/Grid Motion	Prescribed, coupled rigid-body dynamics, or deformable (FSI)	Prescribed	None
Physics	Multispecies, multiphase, fluid-structure interaction (FSI), large-eddy simulation (LES)	Reynolds-Averaged Navier-Stokes (RANS), Hybrid RANS-LES (HRLES), LES, Wall-Modeled LES (WMLES)	Multispecies, RANS, LES, WMLES, Conjugate Heating
Mesh transitions	Octree interlevel operators	Overset interpolation	Fully conformal

LAVA CFD Solvers: HPC Performance Comparison

- Meshing is fully automated for Cartesian and Unstructured, but our curvilinear solver relies on a mostly manual approach.

- Bottlenecks across all codes are due load balance due to irregular workloads and implicit solving.

- All LAVA codes achieve similar scalability and performance due to a common multi-level parallelism approach using MPI/OpenMP/SIMD for CPU.

Application Example: Curvilinear WMLES for High-Lift Prediction Workshop 5 (HLPW5)

- 5.2% Scale Common Research Model High-Lift Configuration (CRM-HL)
- Structured Overset/Multi-Block Grid System
- Total of 915 zones with 805.6 million grid points
- Grid Generation Cost: **4 people 1.5 months**
- Resources: **100 Rome CPU nodes, 4 days, 16 cases**

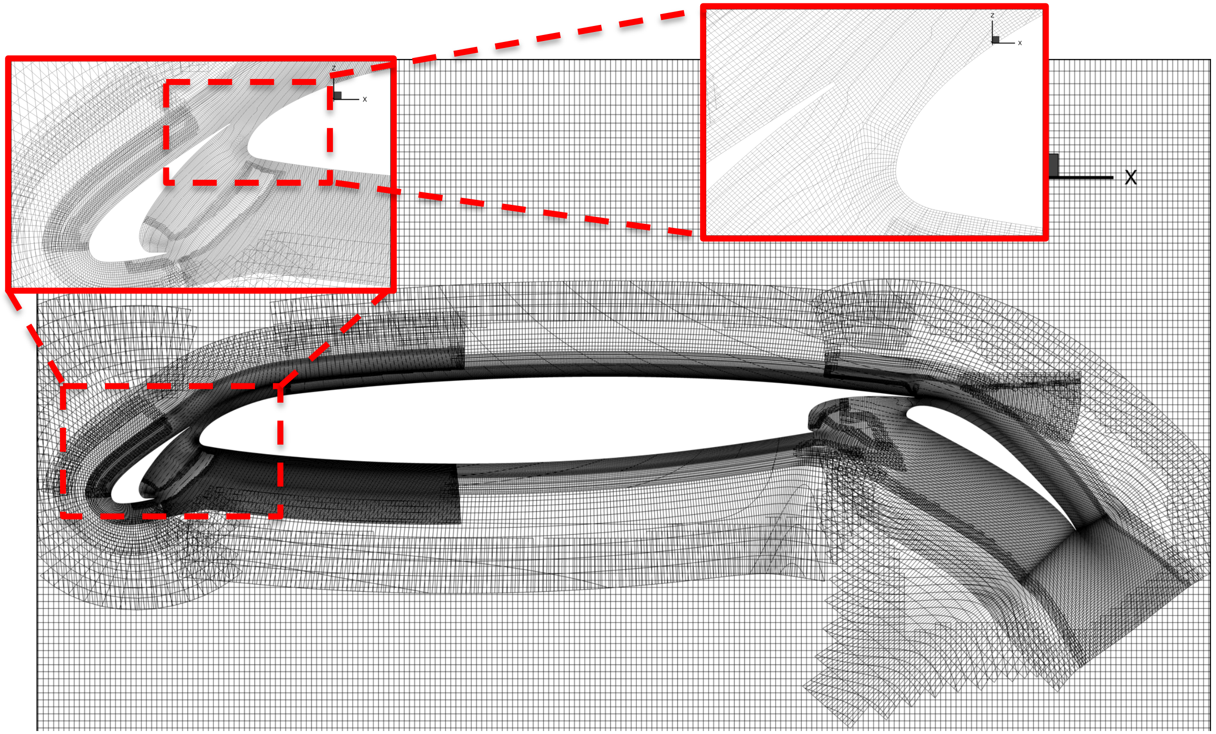

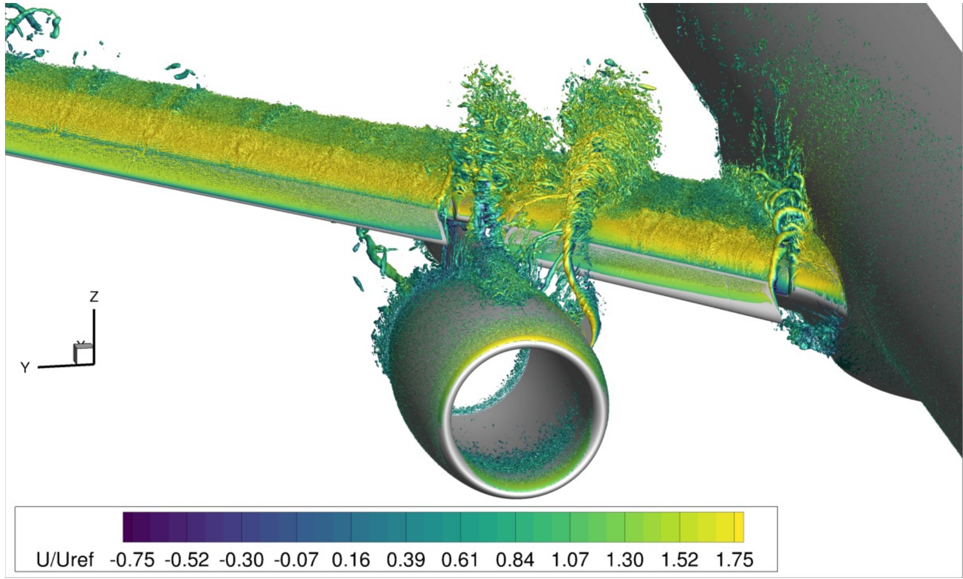

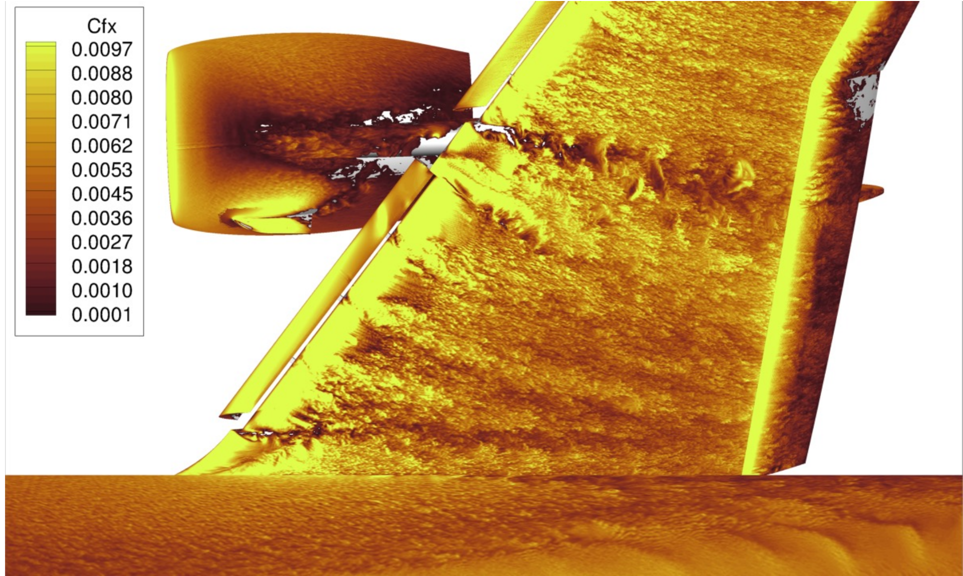

Major Contributors

Gerrit-Daniel Stich, Daniel Maldonado, Leonardo Machado, Luis Miguel dos Santos Fernandes, James Koch, Oliver Browne, Jeffrey Housman, Jared Duensing

LAVA Voronoi Mesher for Complex Geometry: Automated and Body-Fitted

Development effort has been ongoing to implement the LAVA Voronoi mesher:

- Inherent Voronoi properties results in a high-quality mesh, ideally suited to LES applications
- Completely automated (user inputs: surface triangulation, simple text file for sizing regions)
- Coarse/fine interface refinement ratios and blending smoothness can be easily controlled
- Body-aligned layers, where near-body cells are clipped using the geometry with automatic de-featuring of sub-cell details, i.e. no manual work needed to simplify geometries
- High performance, scalable implementation with MPI/OpenMP parallelism (in progress)

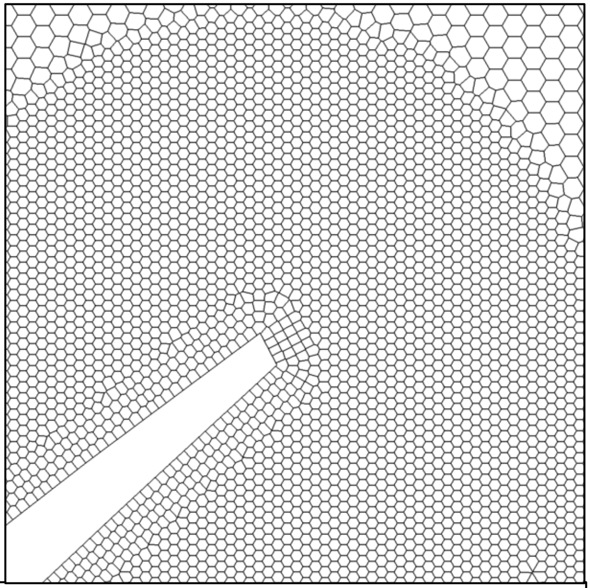

Various cell types are supported

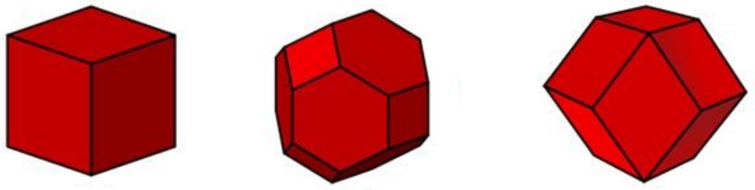

Cube (Cartesian)

Truncated octahedral (body-centered cubic)

Rhombic dodecahedral (face-centered cubic)

Smoothing of coarse/fine interfaces

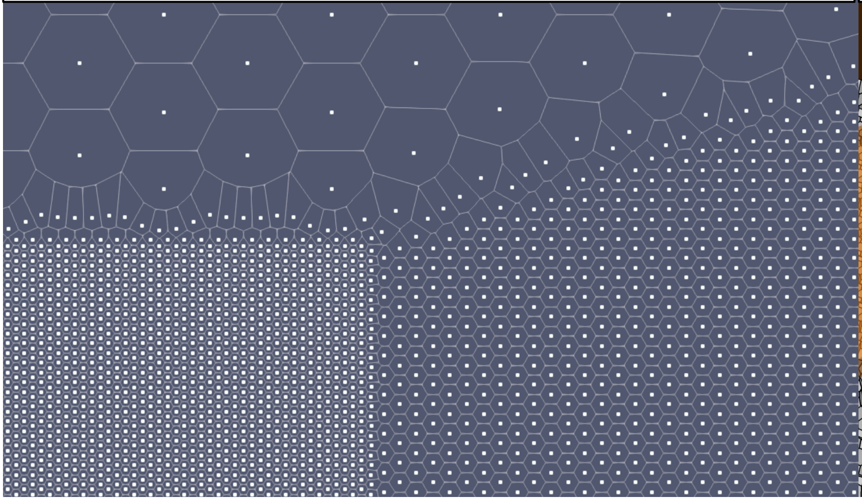

Distance-based refinement and wall-aligned layers

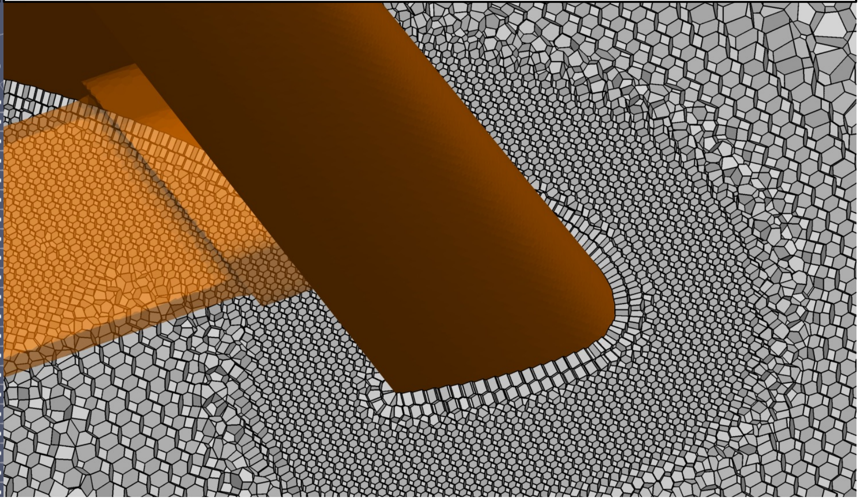

Application Example: Unstructured Voronoi WMLES for High-Lift Prediction Workshop 5 (HLPW5)

- 5.2% Scale Common Research Model High-Lift Configuration (CRM-HL)
- Body fitted Voronoi mesh, body-centered cubic (BCC) interior cell types
- Total of 125 million cells
- Grid Generation Cost: **1 person, 1 day**
- Resources: **120 Skylake CPU nodes, 1 day**

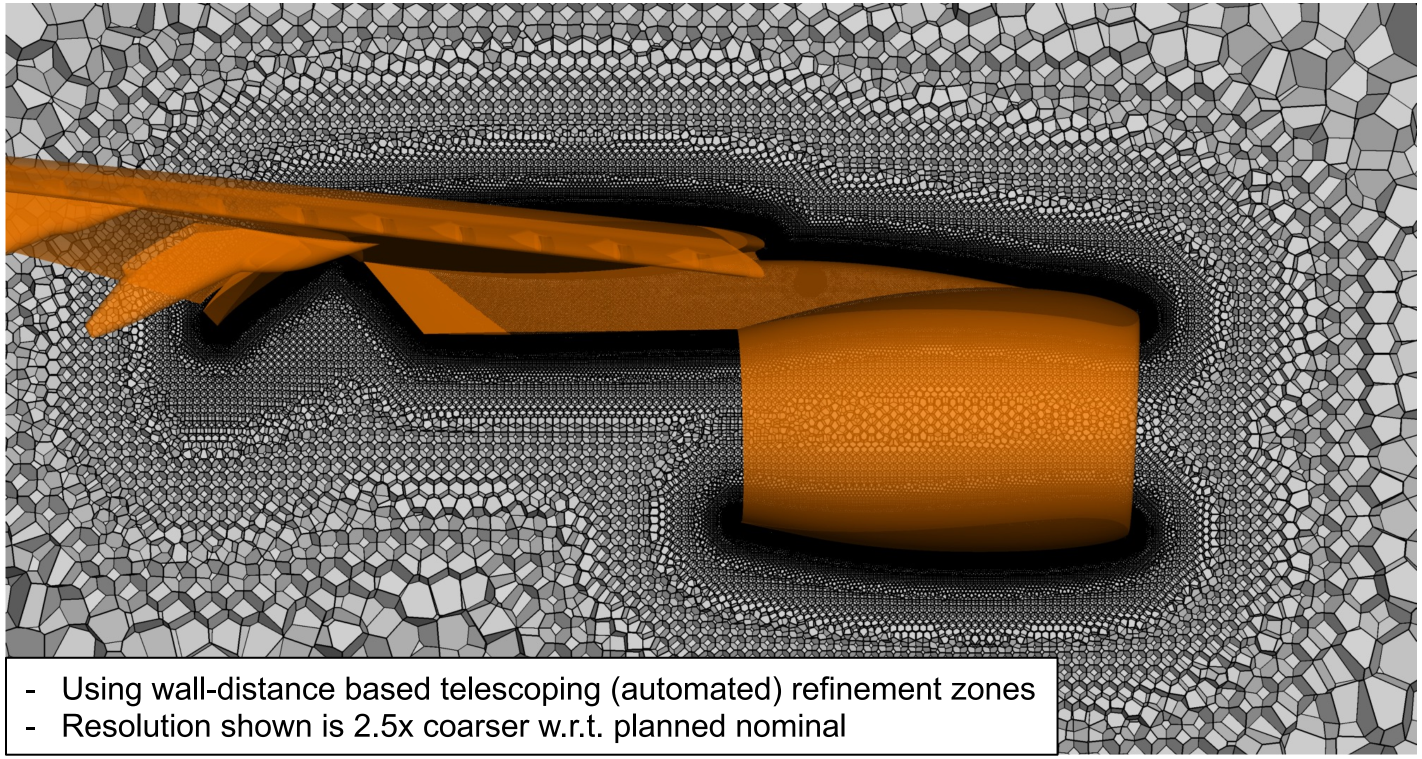

- Using wall-distance based telescoping (automated) refinement zones
- Resolution shown is 2.5x coarser w.r.t. planned nominal

Major Contributors: Emre Sozer, Michael Barad, Abram Rodgers, Victor Sousa, Keshav Sriram

Skin friction in x-direction

Preliminary coarse simulation

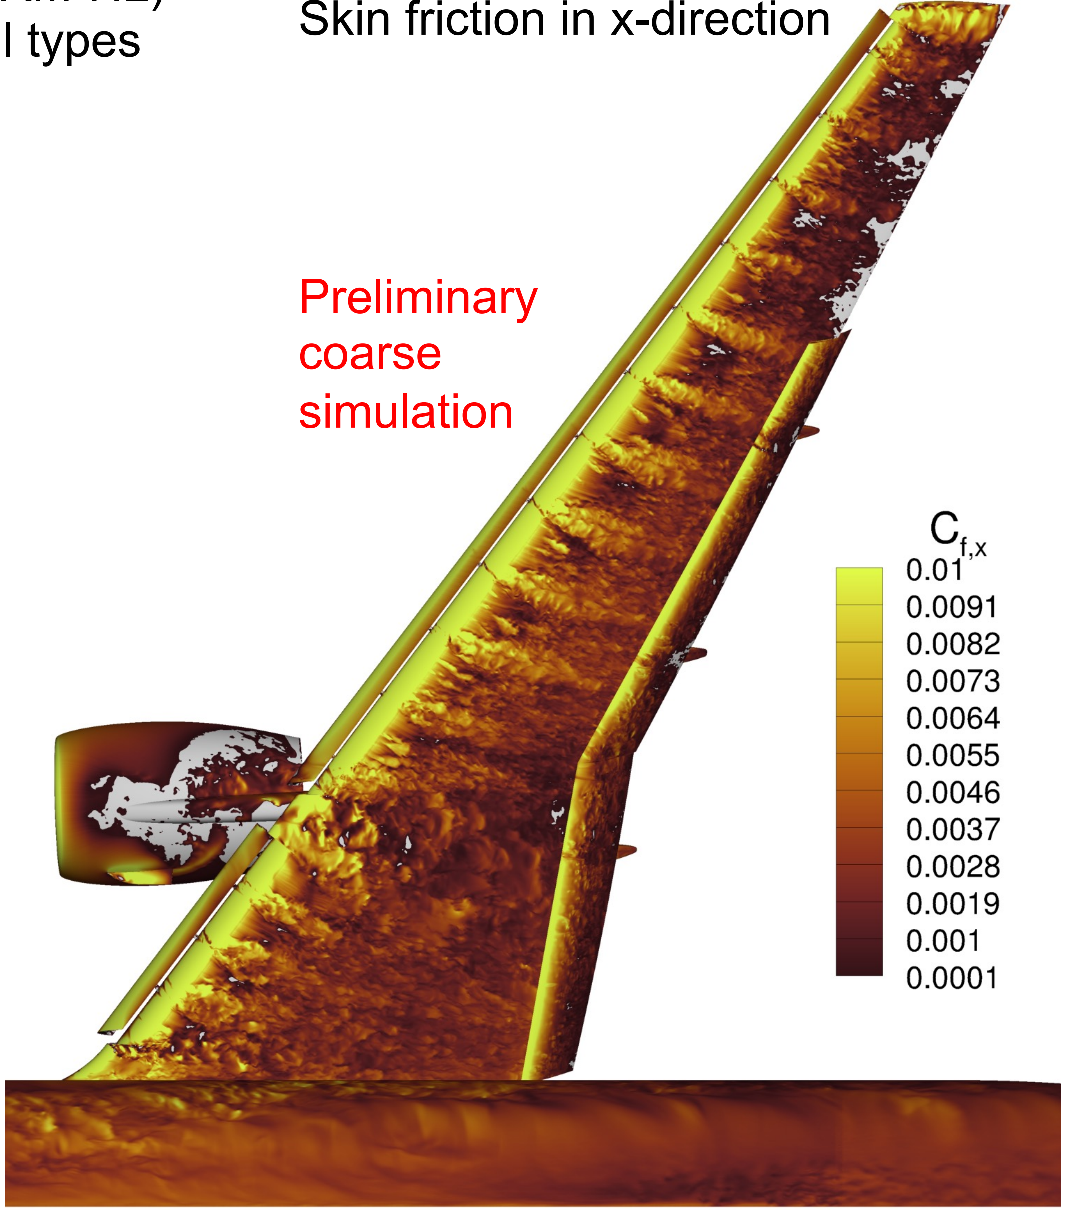

Application Example: Curvilinear WMLES for Jet Acoustics Predictions for Commercial Supersonic Technologies (CST)

- Body fitted curvilinear mesh
- Total of 400 million cells
- Grid Generation Cost: **1 person, 4-14 days**
- Resources: **100 Rome CPU nodes, 1 day**

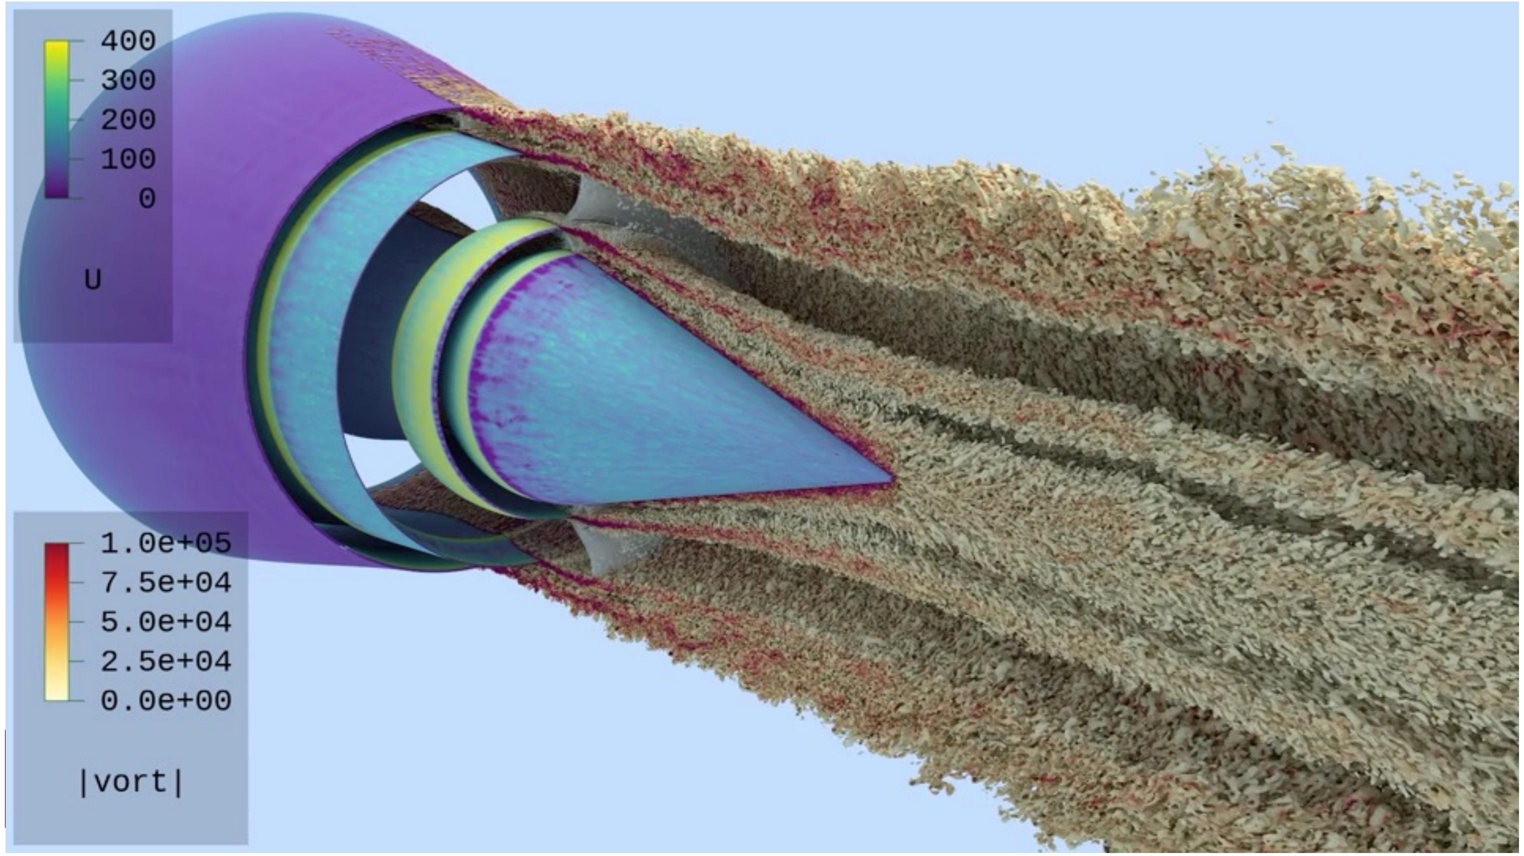

Major Contributors
Gerrit Stich and Jeffrey Housman

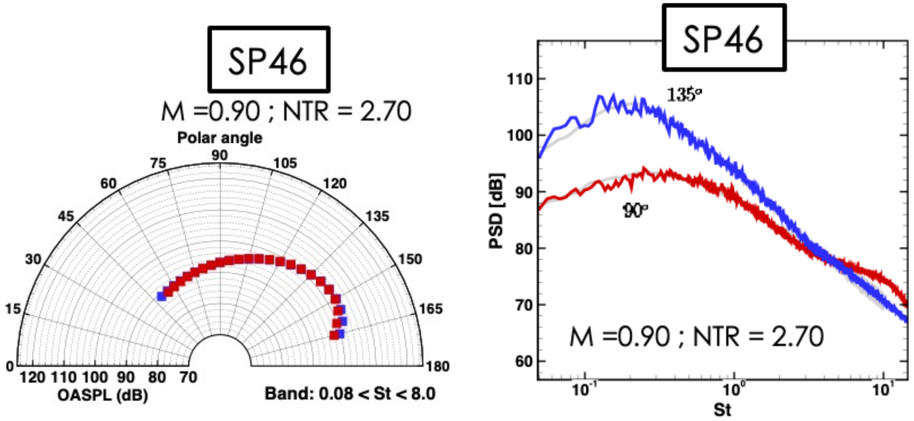

Overall Sound Pressure Level (OASPL) and Power Spectral Density (PSD) for SP46 at 100D from nozzle exit

Stich et. al. AIAA2022-3002 & AIAA2022-0684

Path towards a robust, reliable and fast WMLES solver for jet noise database generation

November 2019

Hybrid RANS/LES	6.5 days [100 Broadwell]

63 x speedup

Today

WMLES + Optimized Code & Best Practices & Scaling	150 minutes [100 Rome]

Application Example: Unstructured Voronoi WMLES for Jet Acoustics Predictions for CST

- Body fitted Voronoi mesh
- Total of 125 million cells
- Grid Generation Cost: **1 person, 1 day**
- Resources: **120 Skylake CPU nodes, 16 hours**

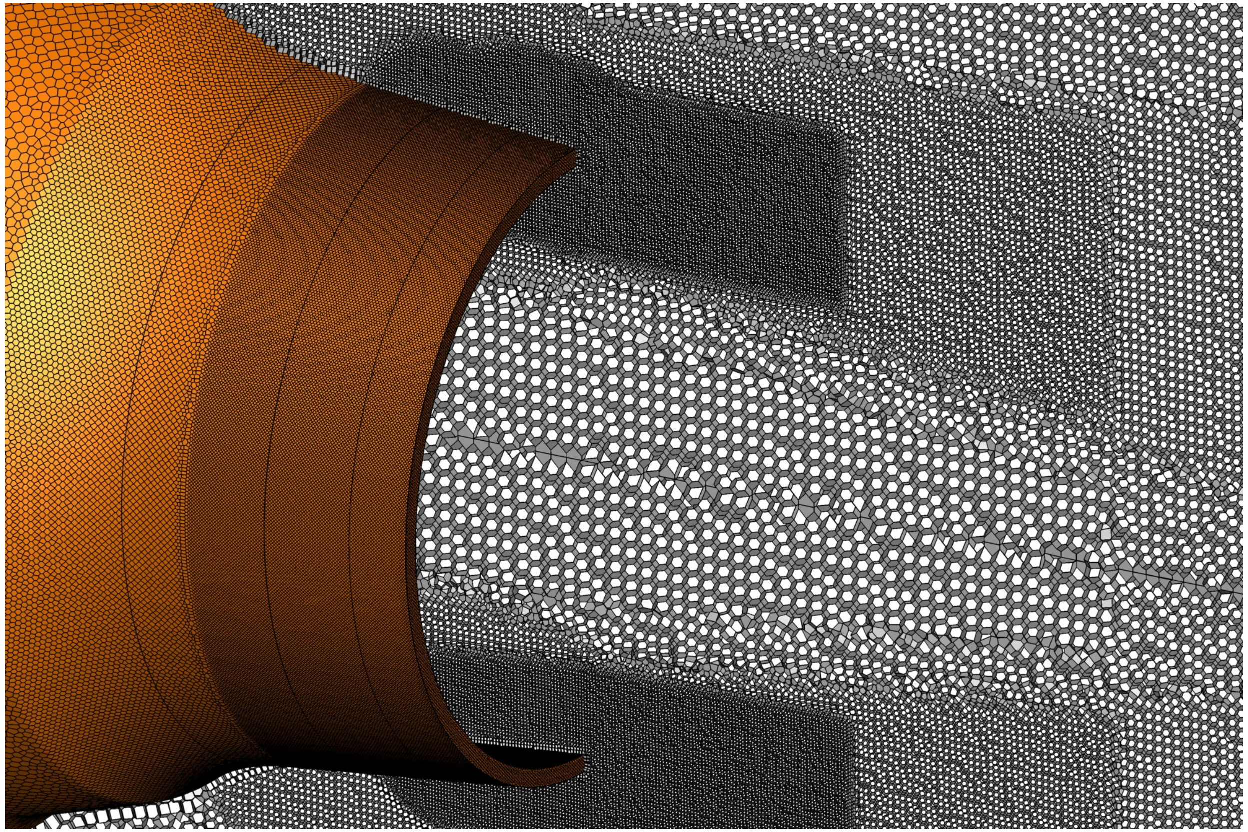

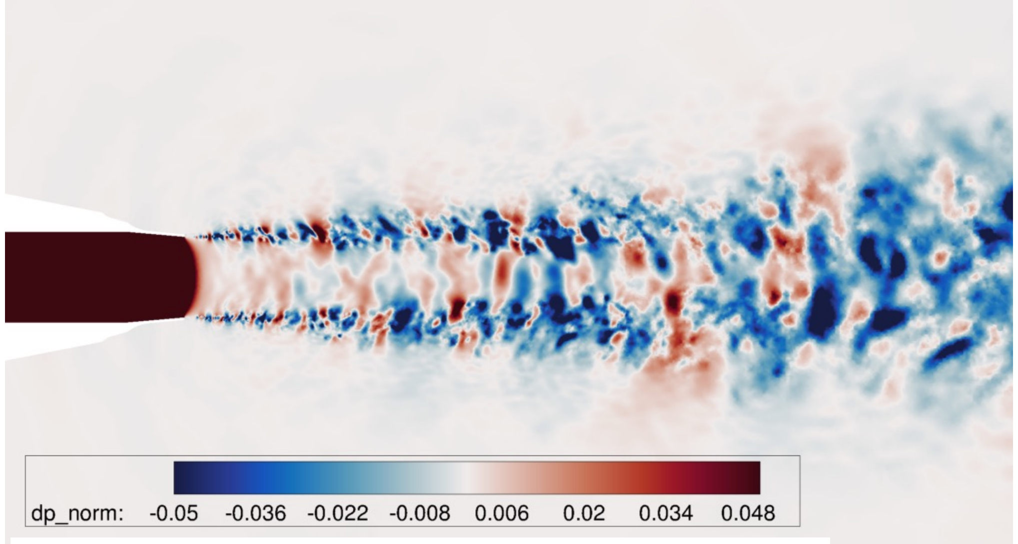

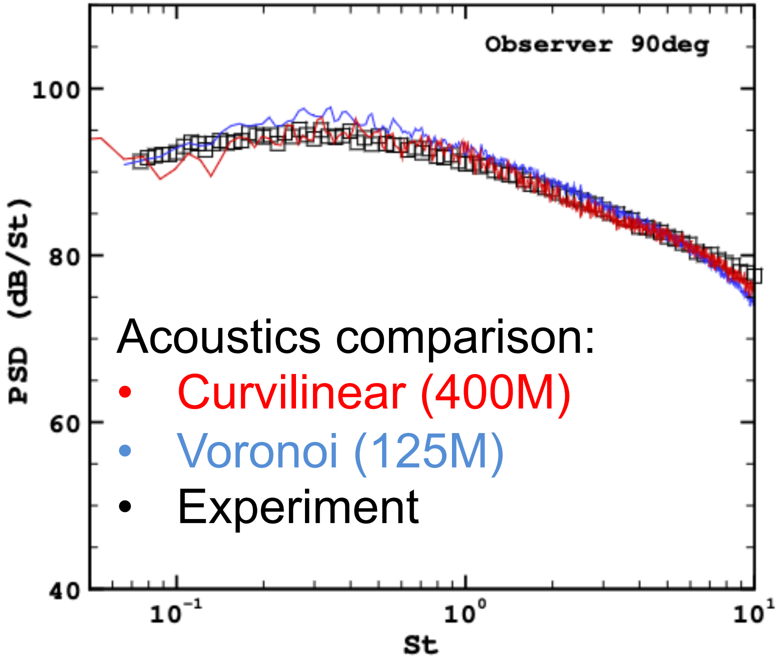

Major Contributors
Emre Sozer, Michael Barad, Victor Sousa, Abram Rodgers, Keshav Sriram, Gerrit Stich

Application Example: Cartesian Multiphase Simulations of Kennedy Space Center Flame Trench

- Models ignition over-pressure water suppression system (IOPSSS) impact on launch vibrations and acoustics
- Used for launch pad redesign
- Up to 550 million cells
- Grid Generation Cost: **none**
- Resources: **100 Sky CPU nodes, 5 weeks**

Dry

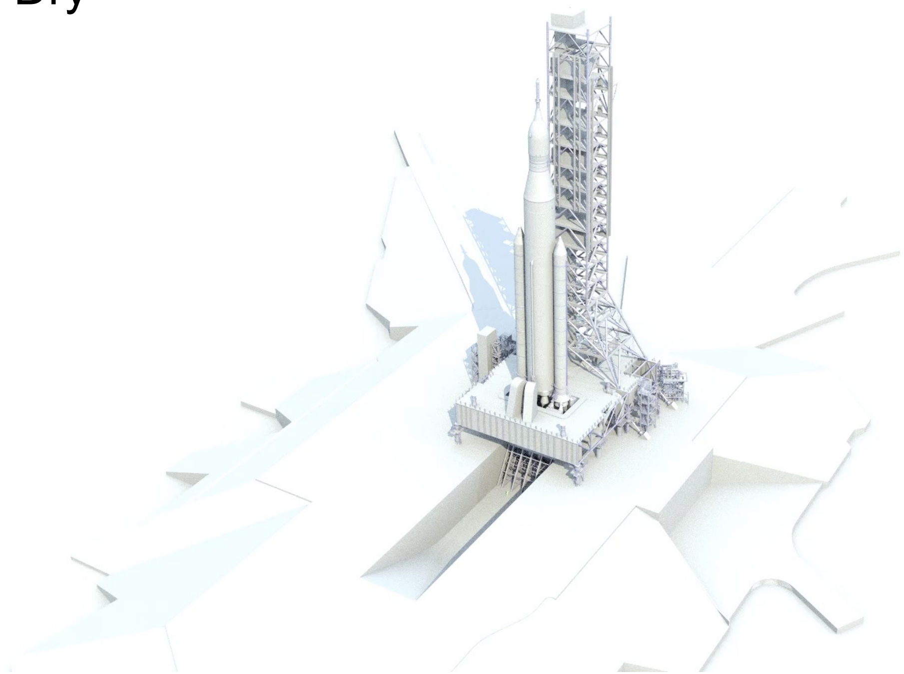

Wet

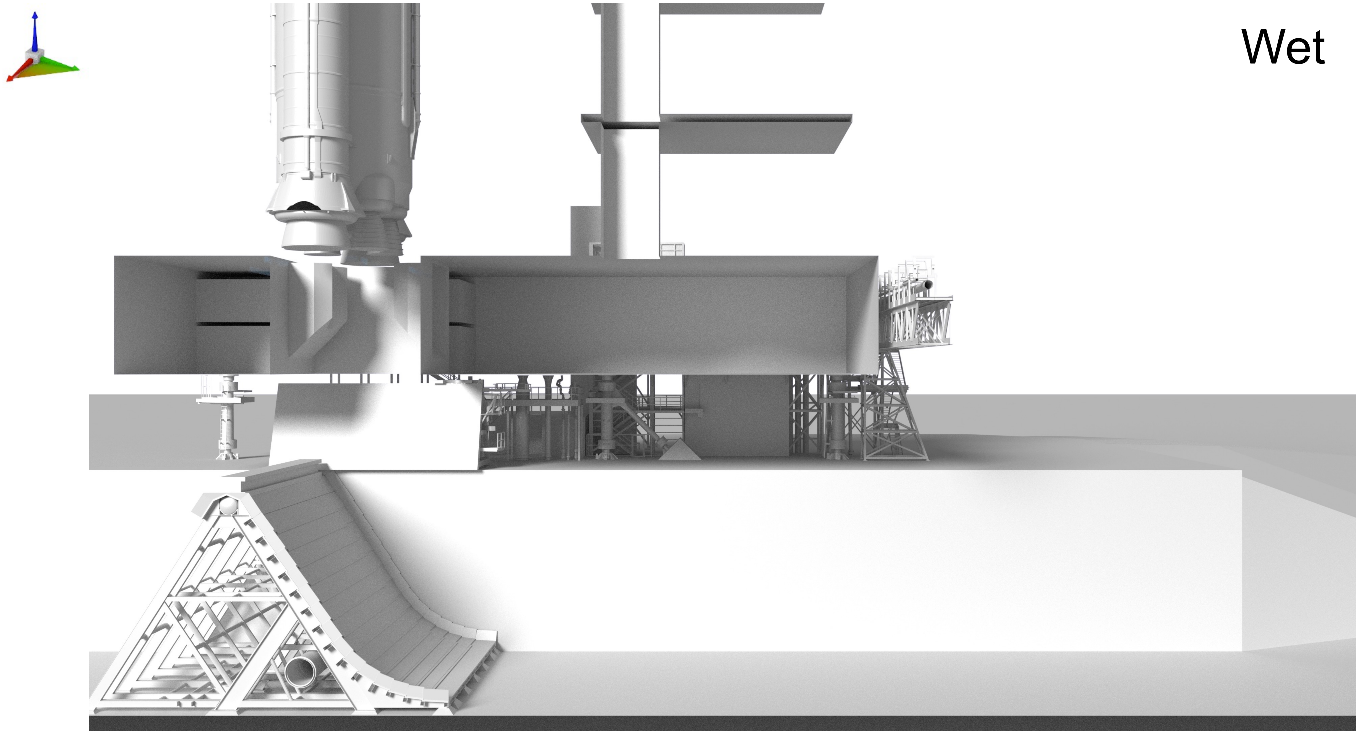

Major Contributors: Michael Barad, Scott Neuhoff, Francois Cadieux, Emre Sozer, Jeffrey Housman, Timothy Sandstrom

Application Example: Cartesian AMR Simulations of Supersonic Parachute Inflation

- Fluid-structure interaction: CFD loosely coupled with geometrically non-linear structural dynamics code to model parachute inflation
- Block-structured with octree based adaptive mesh refinement, up to 300 million cells
- Grid Generation Cost: **only triangulation/lines prep, otherwise automatic**
- Resources: **50 SkyLake CPU nodes, 2 days**

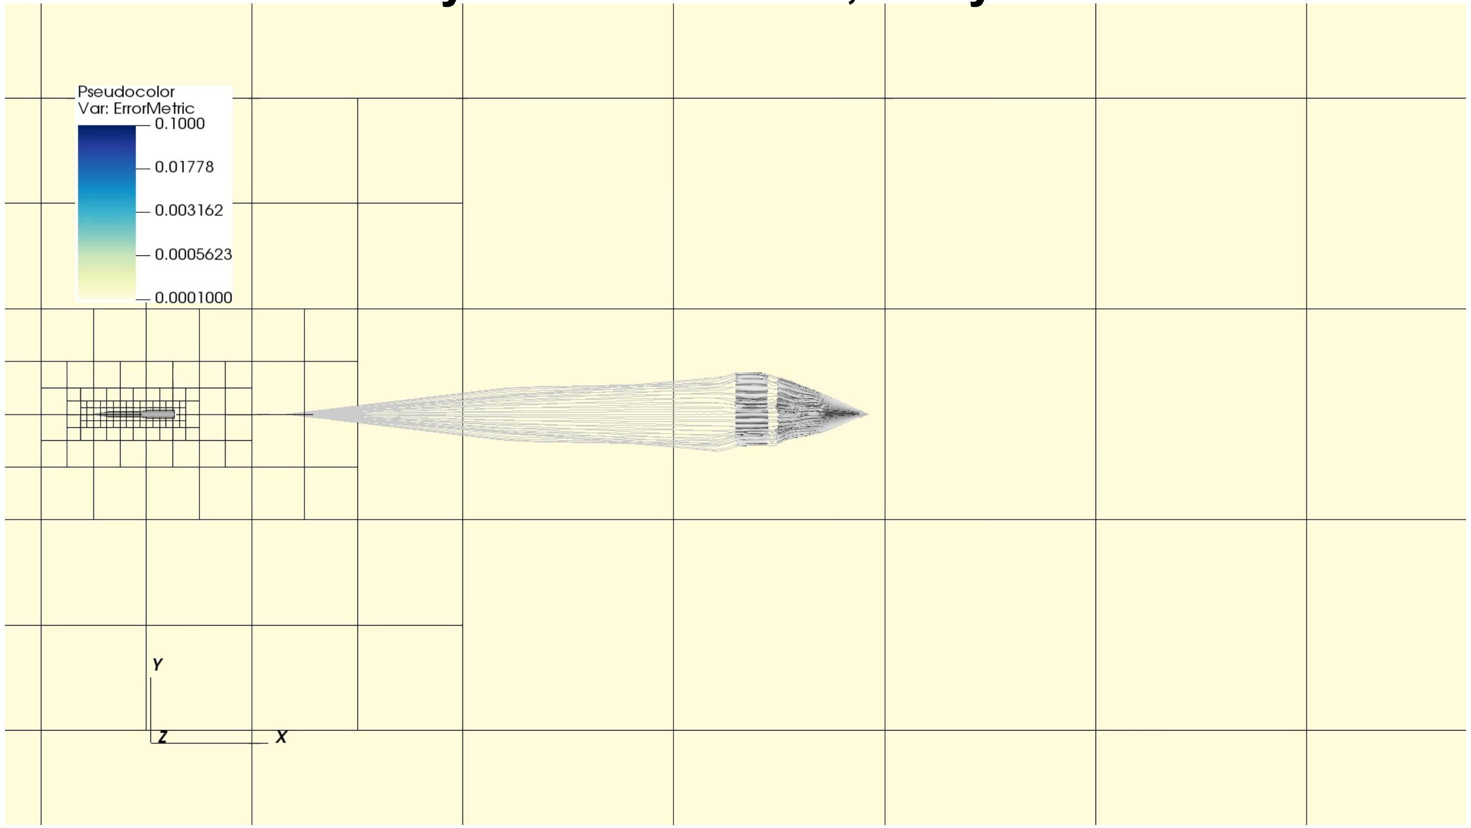

- Error metric on the cut plane and parachute surface for FSI
- Payload and suspension lines shown in light gray.
- Solution-based AMR (every 50 time steps).
- Mesh refinement using error indicator is limited to concentric frustrums defined by the user to avoid larger mesh that could blow past memory limit.

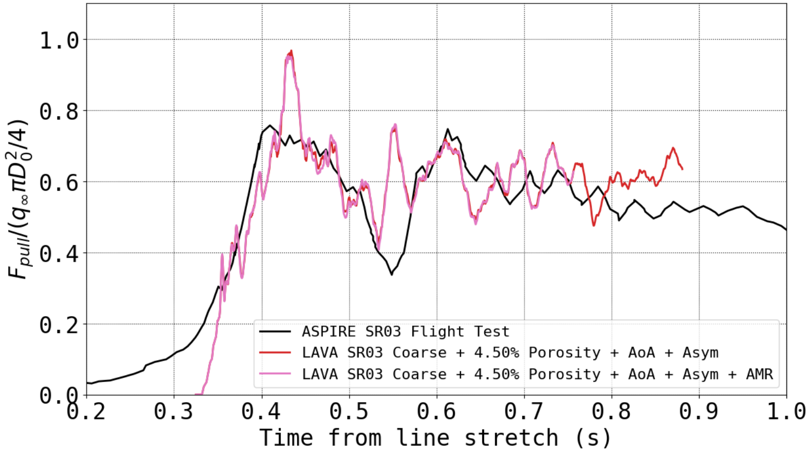

Effect of Solution-based AMR on Accuracy

Major Contributors
Francois Cadieux and Michael Barad

Application Example: Cartesian LES of Supersonic Retro-Propulsion Concept

- Simulate descent of human Mars lander concept
- Model Mars atmosphere and combustion exhaust mixture (no combustion physics)
- Up to 1.62 billion cells
- Grid Generation Cost: **none**
- Resources: **70-100 Sky/Rome CPU nodes, 0.5-7 days**

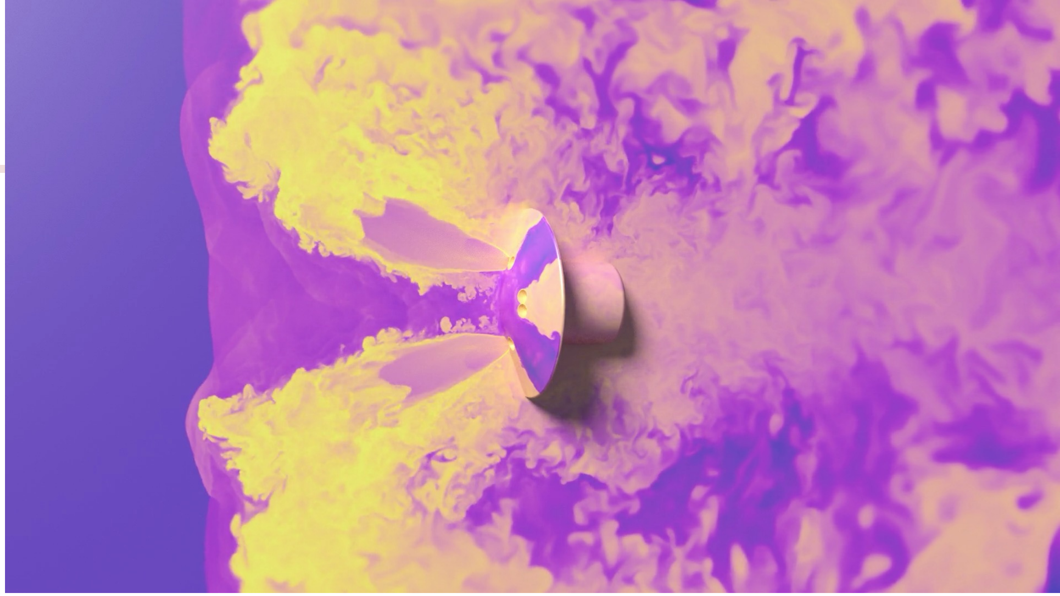

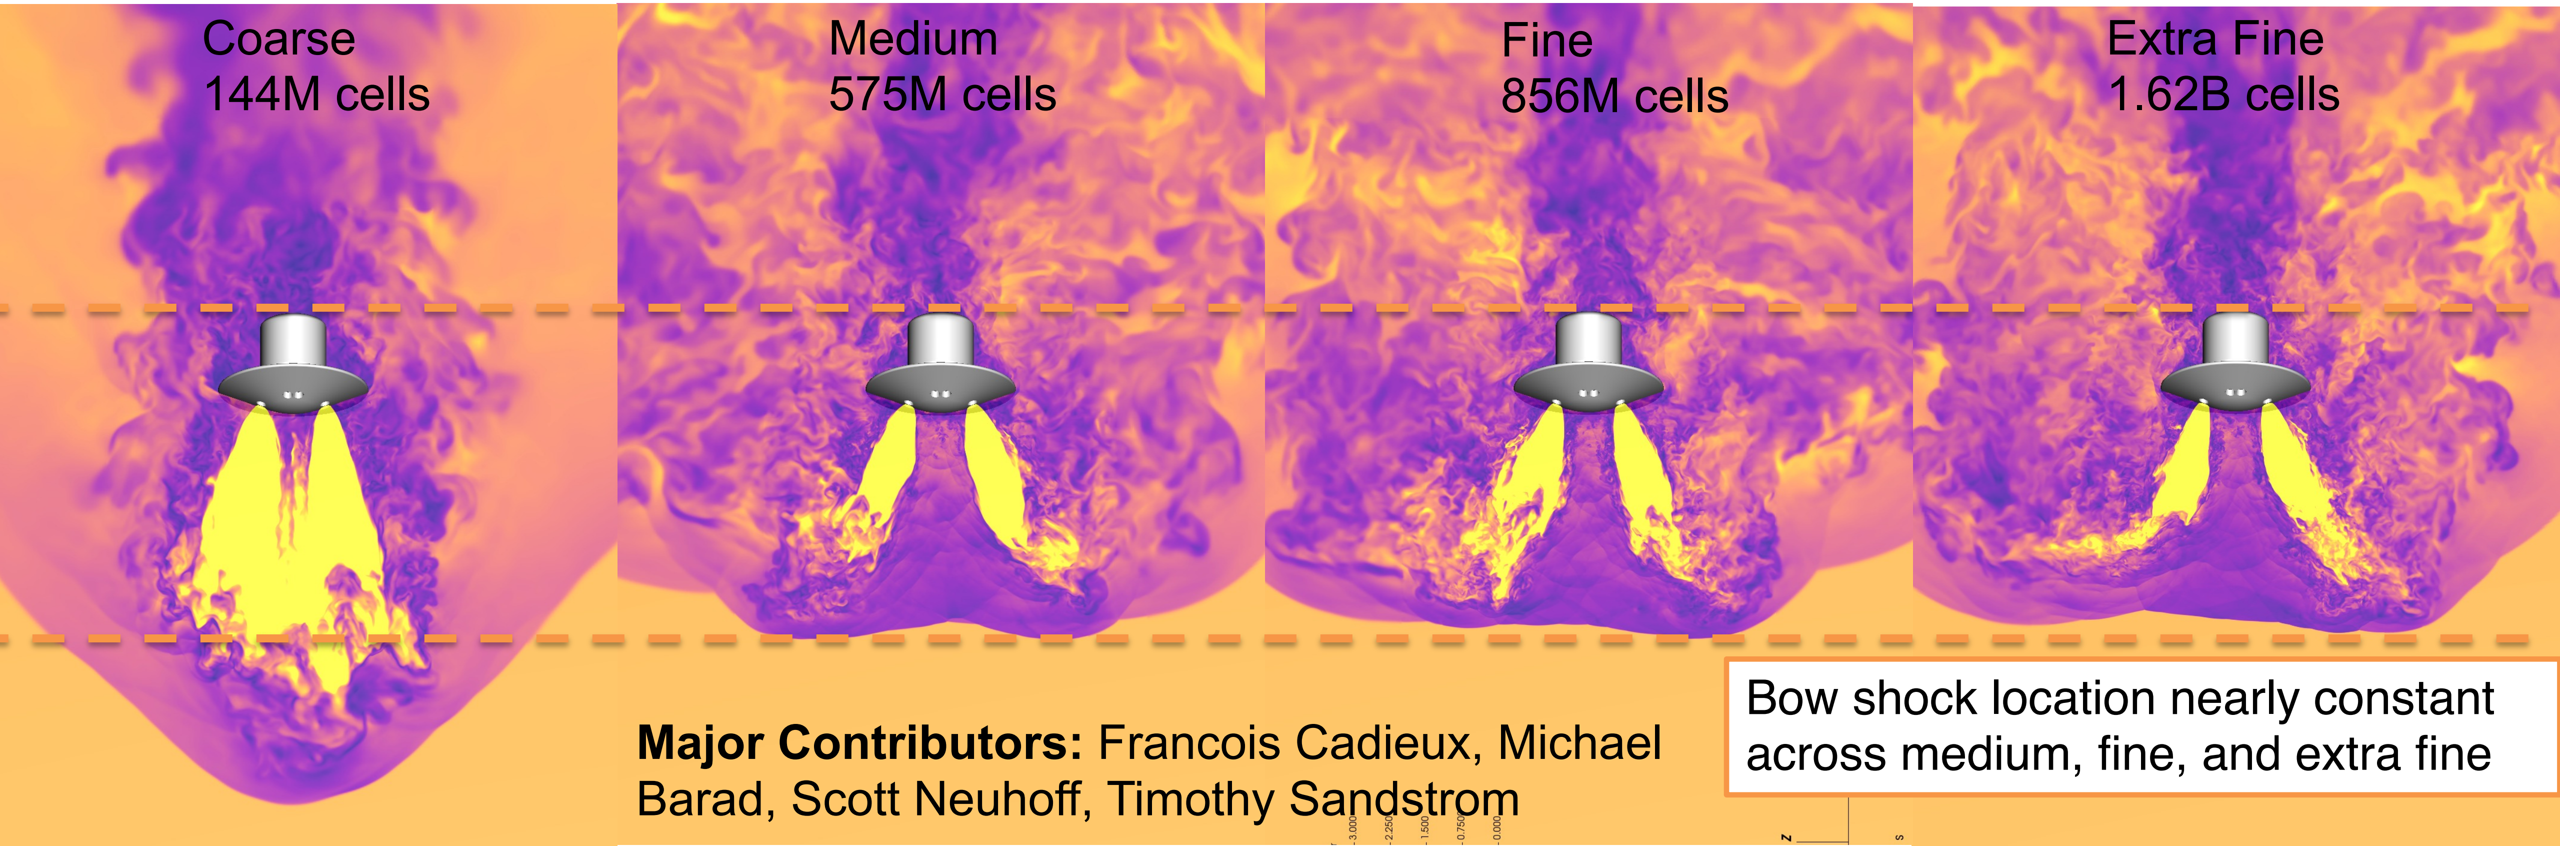

Coarse
144M cells

Medium
575M cells

Fine
856M cells

Extra Fine
1.62B cells

Major Contributors: Francois Cadieux, Michael Barad, Scott Neuhoff, Timothy Sandstrom

Bow shock location nearly constant across medium, fine, and extra fine

LAVA GPU Porting – Outlook

- **Out of the top 15 fastest systems worldwide, all but 2 rely on GPU accelerators**
- GPU porting of LAVA is in progress
- Unstructured benchmarks on hexahedral meshes were performed during 2022 NAS/NVIDIA GPU Hackathon
  - MUPS: Millions of UPdates (timesteps) per Second → Higher is better
  - **MUPS/W: MUPS per unit Power** → Higher is better

Architecture	Type	MUPS	Power, TDP (W)	MUPS/W	Theoretical Peak Flops (TF)	Theoretical Peak Bandwidth (GB/s)
2x Intel Skylake 6148 (AVX512)	CPU	35	300	0.12	3.1	230
1x Nvidia A100 (CUDA)	GPU	**298**	400	**0.75**	9.7	**1600**
GPU/CPU Ratio		**8.5**	**1.3**	**6.4**	**3.1**	**7.0**

Challenges and Future Plans

- Current simulations are typically run in the O(0.1-1) billion grid cell range and their turn-around time is limited by hardware availability & current solver strong scaling limits.

- Challenges for O(10-100) billion grid cell scale simulations/applications is predominantly parallel scalability of the full CAD to solution process (CAD cleanup, mesh generation, simulation, visualization and post-processing). Pipeline automation is critical.

- Plans for greater use of unsteady CFD will have increasing demands on hardware availability and constantly changing hardware architecture requires developers to write portable software while retaining performance.

Challenges and Future Plans

- Current simulations are typically run in the O(0.1-1) billion grid cell range
  - Typical run is 100-400 Sky/Cascade/Rome nodes on NASA’s Supercomputers
  - Turn-around time (1-30 days) is limited by:
    - hardware availability
    - current strong scaling limits
- Challenges for O(10-100) billion grid cell scale simulations/applications:
  - Mesh generation readiness:
    - Cartesian: tested to 100 billion, fully automated parallel/distributed volume meshing
    - Curvilinear: manual meshing, untested above 10 billion
    - Unstructured: planning on our in-house Voronoi mesher for this scale, parallel underway
  - Software readiness for this scale:
    - Cartesian: CPU ready, GPU underway, wall model maturing
    - Curvilinear: not planned
    - Unstructured: complete GPU WMLES implementation with strong scaling is close
    - Long term plan for coupled body-fitted Voronoi for near body, Cartesian AMR off-body
    - Post-processing (Paraview/VisIt), IO (HDF5, CGNS), in situ processing, etc
  - Hardware availability
    - Will require an increase in hardware resources
    - Performance portability is a concern (ECP guidance on SYCL etc would be helpful)

Acknowledgements

- This work was funded by many NASA projects under ARMD, SMD, and ESDMD.
- Current and past LAVA Team members for their technical contributions and valuable insight
- Computer time has been provided by the NASA Advanced Supercomputing (NAS) facility at NASA Ames Research Center.