National Aeronautics and Space Administration

Landing on Mars: Petascale Unstructured-Grid CFD Simulations on Summit

Ashley Korzun
Eric Nielsen
Aaron Walden
Bill Jones
Jan-Reneé Carlson
NASA Langley Research Center

Pat Moran, Chris Henze
NASA Ames Research Center

Tim Sandstrom
Inu Teq, LLC

Mohammad Zubair
Old Dominion University

Justin Luitjens
NVIDIA Corp.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

NVIDIA Theater, SC19 Denver, CO

FUN3D Overview

- Established as a research code in late 1980s; now supports numerous internal and external efforts across the speed range
- Solves 2D/3D steady and unsteady Euler and RANS equations on node-based mixed element grids for compressible and incompressible flows
- General dynamic mesh capability: any combination of rigid / overset / morphing grids, including 6-DOF effects
- Aeroelastic modeling using mode shapes, full FEM, CC, etc.
- Constrained / multipoint adjoint-based design, mesh adaptation
- Distributed development team using agile/extreme software practices including 24/7 regression, performance testing
- Capabilities fully integrated, online documentation, training videos, tutorials

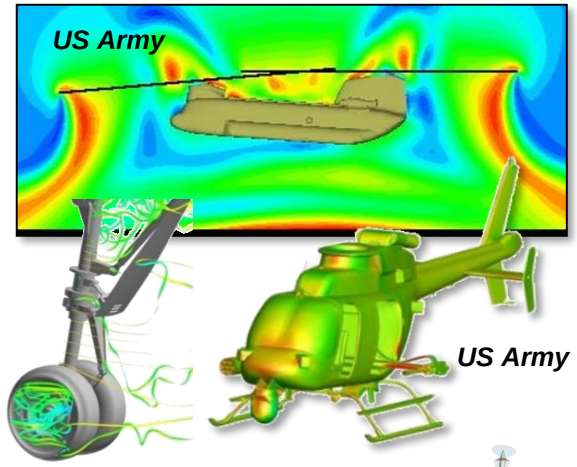

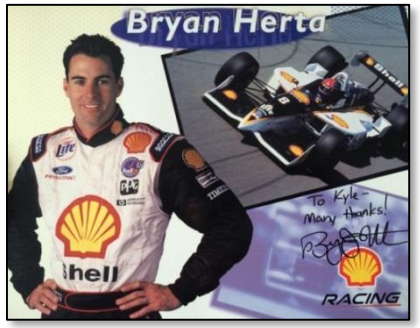

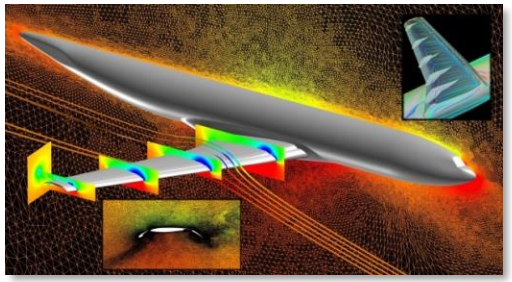

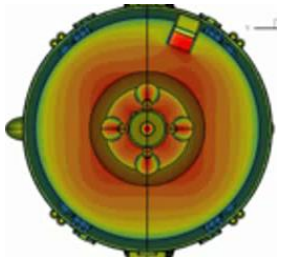

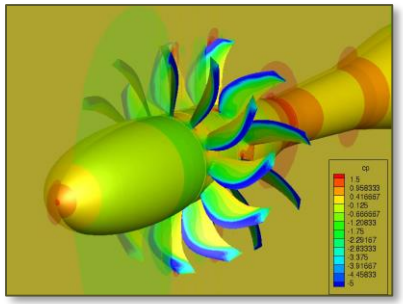

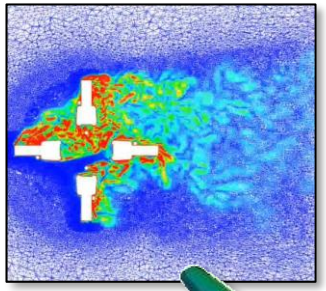

Georgia Tech

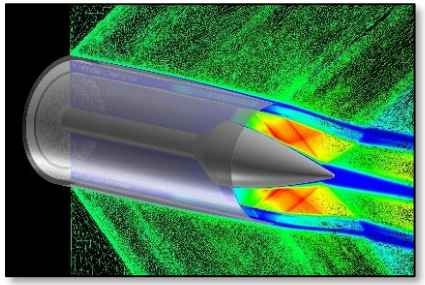

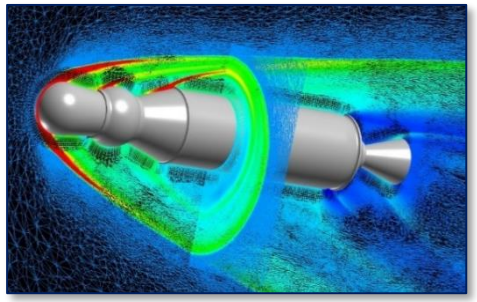

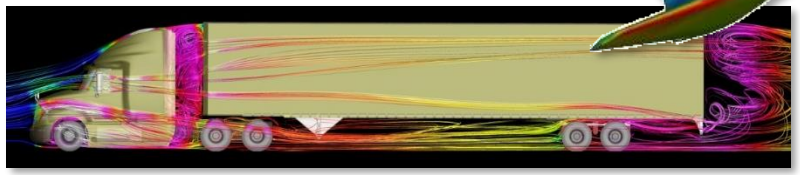

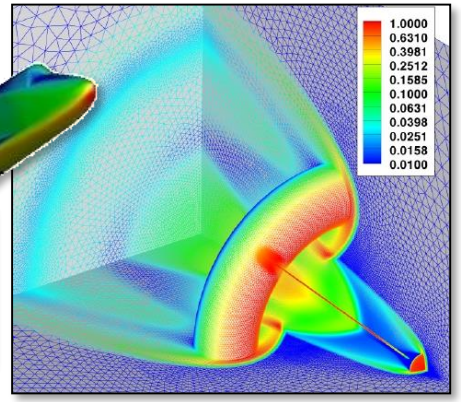

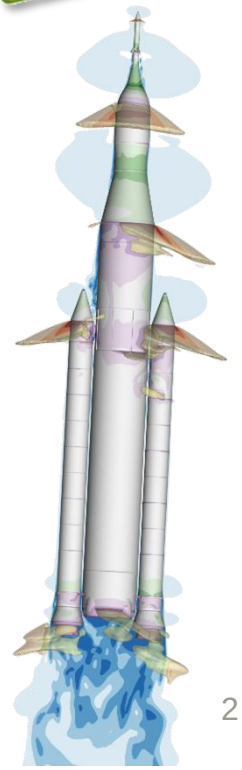

The Three Pillars of HPC

2018 — *Hardening and Deployment*

2016 — *Technology Demonstration*

2014

2012

2010 — *Basic Research*

Next-Generation Leadership Class Performance

Software

- Asynchronous, Overlapped Communication
- CUDA-Enabled MPI
- Language Interoperability
- Matrix Assembly
- Linear Algebra

Full App at Scale

Full App at Node Level

Mini-Apps

Workforce Skills

Programming Models

OpenACC CUDA Fortran
PTX CUDA C/C++ Kokkos

Partnerships & Hackathons

DOD HPC Modernization Program, NVIDIA, ParaTools, Oak Ridge National Laboratory, PGI, Argonne National Laboratory, Old Dominion University, ENGILITY

NASA / NVIDIA @ SC10

Hardware

Volta

- Memory Bandwidth
- Larger Memory

Pascal

- GPUDirect RDMA
- Power Efficiency
- Fast Atomics

Kepler

- Hyper-Q
- ECC Memory
- DP Support

Early GPUs

FUN3D Performance

Device Level

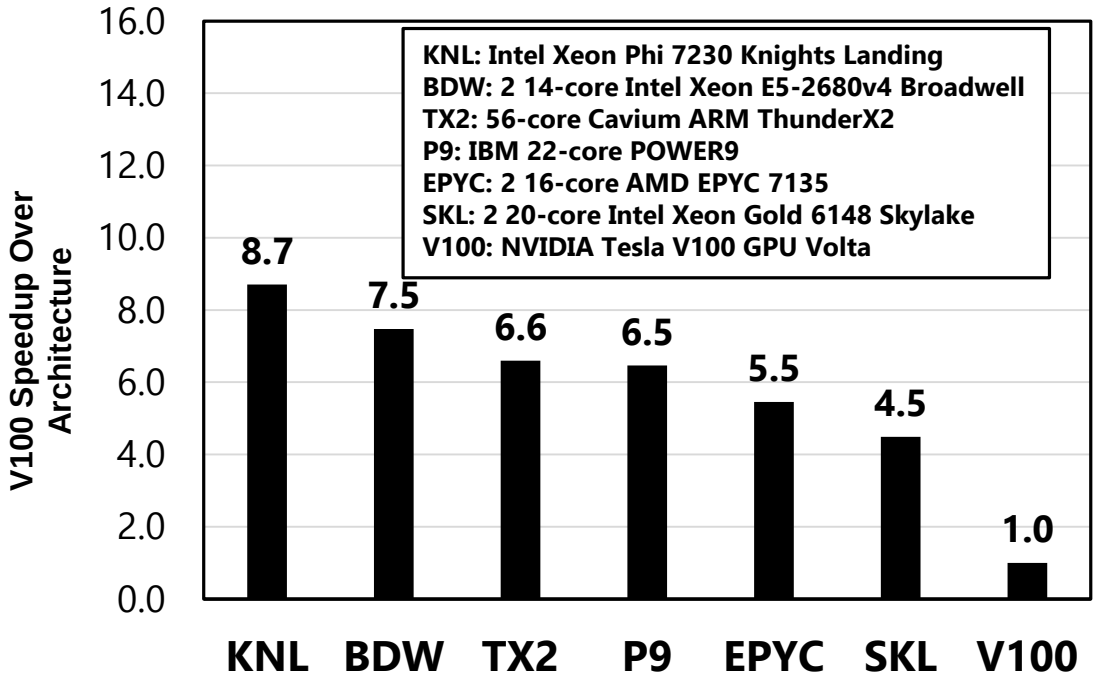

At Scale

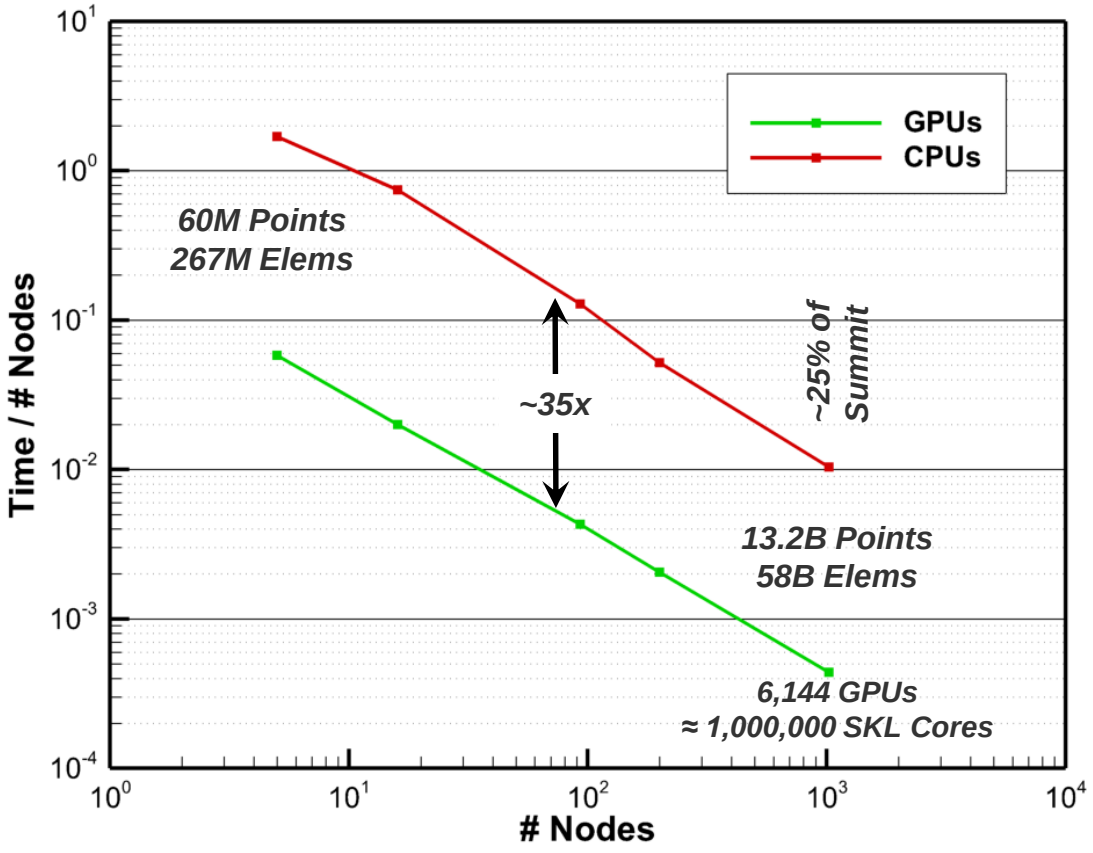

- Early access during Summit construction enabled early 2018 performance demonstrations shown here

2019 Summit Effort

"Enabling Human Exploration of the Red Planet"

- **CY19 allocations competitively awarded through Summit Early Science and INCITE programs**
  - Total award of 305,000 Summit node-hours
  - Equivalent of ~305,000,000 Xeon Skylake core-hours
- **Team includes NASA Langley, NASA Ames, NVIDIA, and Old Dominion University**
  - LaRC: Science and computational expertise
  - ARC: Large-scale visualization, network transfers
  - NVIDIA, ODU: Kernel optimizations

Campaign Goals

- ***Science***: Advance the understanding of retropropulsion flow physics during Mars EDL of a human-scale lander
- ***Computational***: Demonstrate production readiness and efficiency advantages of GPU implementation of the FUN3D CFD code at scale

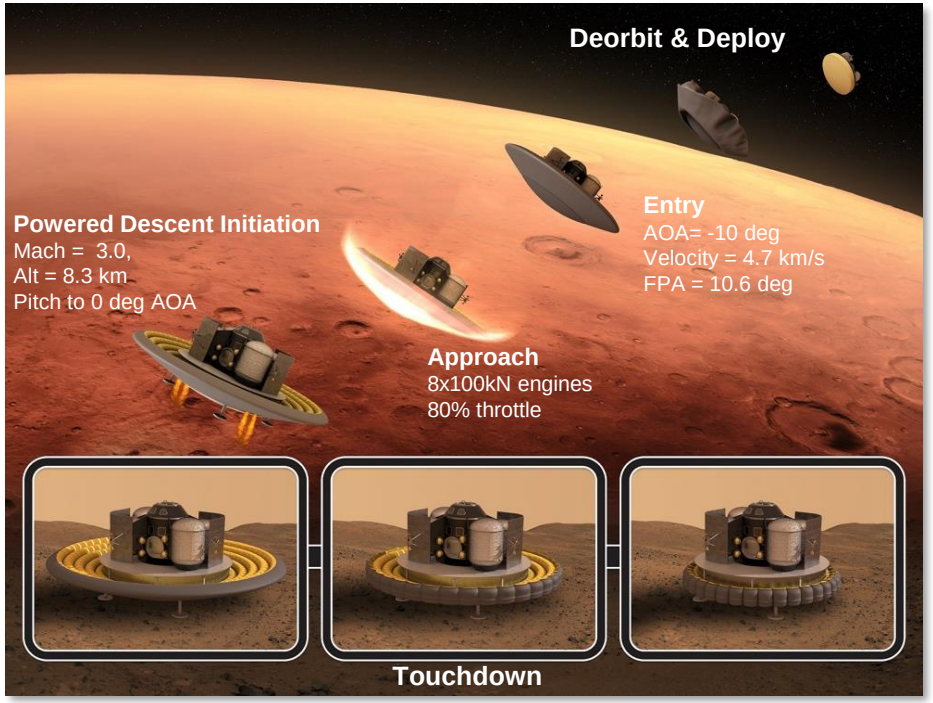

Retropropulsion for Human Mars Exploration

Human-scale Mars landers require new approaches to all phases of Entry, Descent, and Landing

- Cannot use heritage, low-L/D rigid capsules → deployable hypersonic decelerators or mid-L/D rigid aeroshells
- Cannot use parachutes → retropropulsion, from supersonic conditions to touchdown
- No viable alternative to an extended, retropropulsive phase of flight

	Viking	Pathfinder	MERs	Phoenix	MSL	InSight	M2020	Human-Scale Lander (Projected)
Entry Capsule (to scale)								
Diameter (m)	3.505	2.65	2.65	2.65	4.52	2.65	4.5	16 - 19
Entry Mass (t)	0.930	0.584	0.832	0.573	3.153	0.608	3.440	40 - 65
Parachute Diameter (m)	16.0	12.5	14.0	11.8	19.7	11.8	21.5	N/A
Parachute Deploy (Mach)	1.1	1.57	1.77	1.65	2.2	1.66	1.75	N/A
Landed Mass (t)	0.603	0.360	0.539	0.364	0.899	0.375	1.050	26 - 36
Landing Altitude (km)	-3.5	-2.5	-1.4	-4.1	-4.4	-2.6	-2.5	+/- 2.0
Landing Technology	Retro-propulsion	Airbags	Airbags	Retro-propulsion	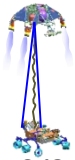 Skycrane	Retro-propulsion	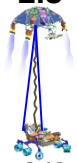 Skycrane	Retro-propulsion

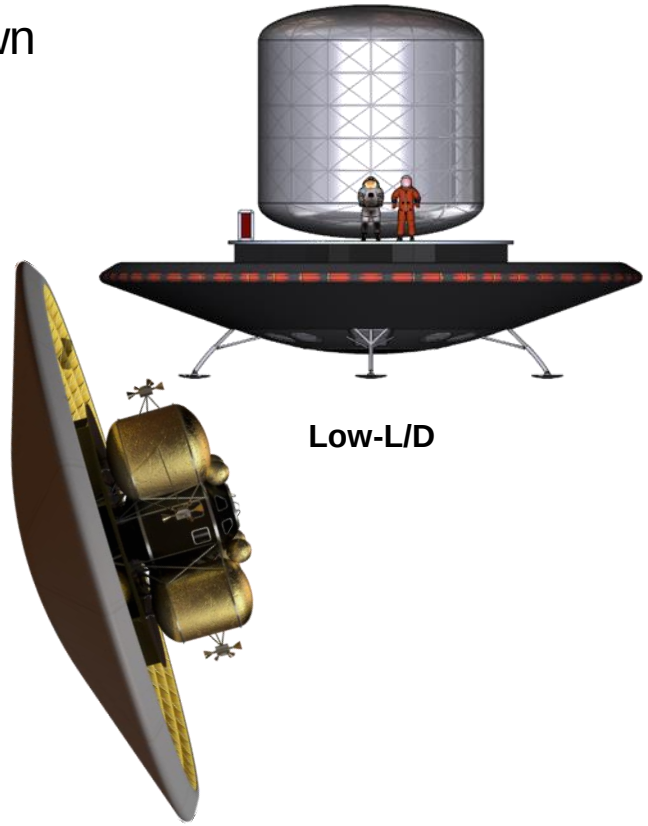

Low-L/D

Steady progression of "in family" EDL → **New EDL Paradigm**

Retropropulsion in an Atmosphere

- Retropropulsion environments impact vehicle performance
- Maturation requires balance between ground testing and computational analysis
- Infeasible to continue with conventional resources, given the computational expense of single solutions (several weeks-to-months each)
- Incremental performance gains in computing will not solve this issue

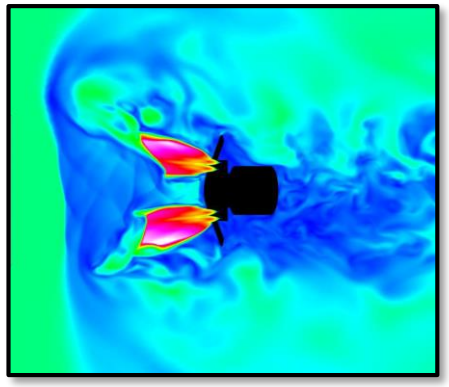

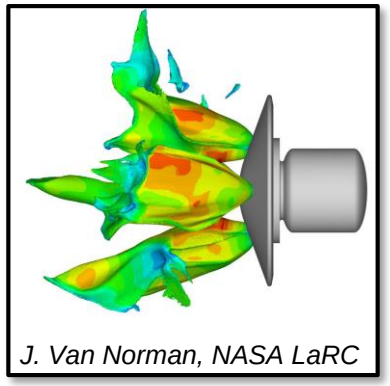
J. Van Norman, NASA LaRC

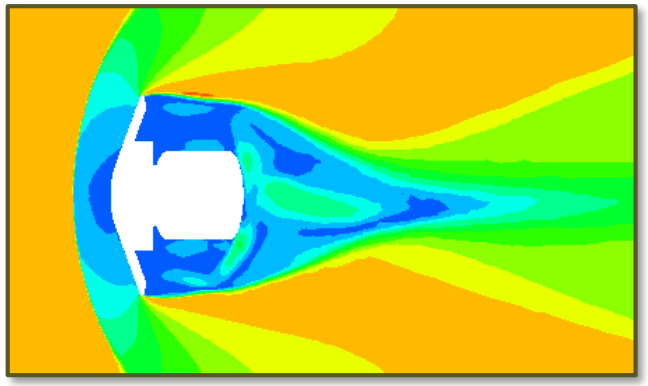

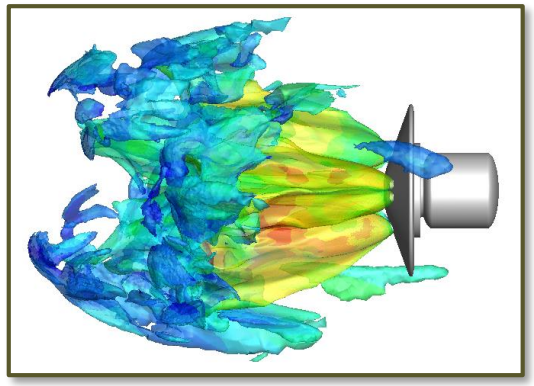

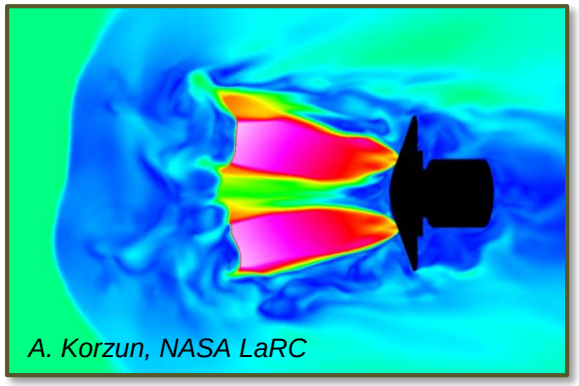
A. Korzun, NASA LaRC

Examples of solutions with insufficient spatial and temporal resolution

Infeasible to develop models and databases within current conventional computational paradigm

Summit Campaigns

- Rather than pursue small number of “hero” simulations, exploring large ensemble of asymmetric throttle conditions across freestream Mach numbers from 0.8 to 2.4
- Spatial mesh sizes ranging from ~1-10 billion elements
- Long temporal duration (~1.6 sec real time) to capture diverse transients and statistics
- Individual runs can reach 200 TB of output; average ~30 TB / day from ORNL to NASA Ames
- Total of ~2 PB of data generated

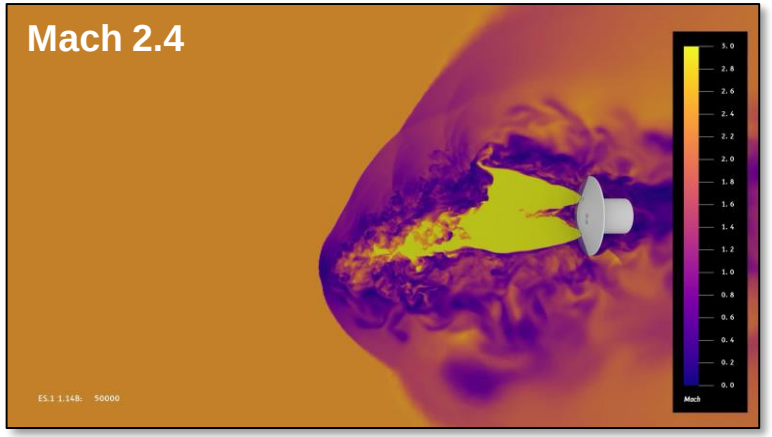

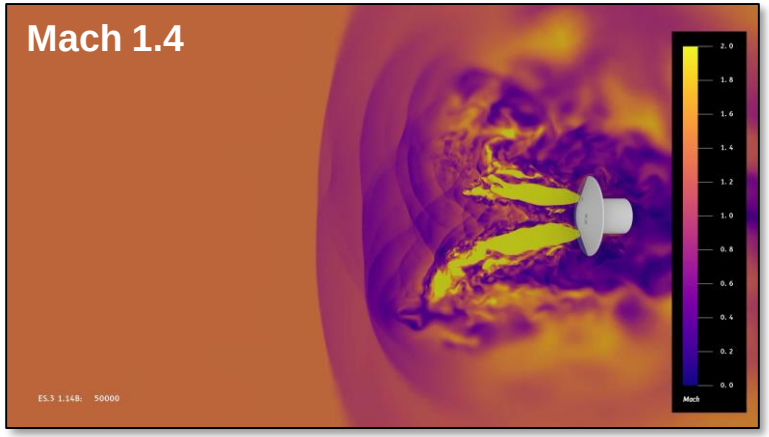

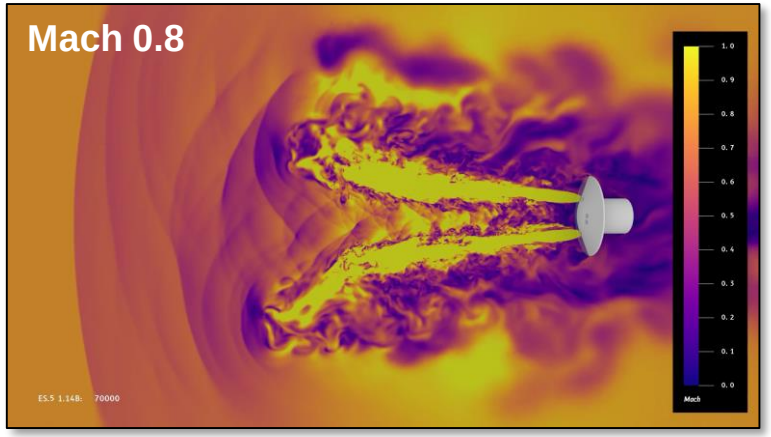

Results: Mach 2.4

All engines at 80% throttle (individual engines are under-expanded)
Vorticity magnitude contours

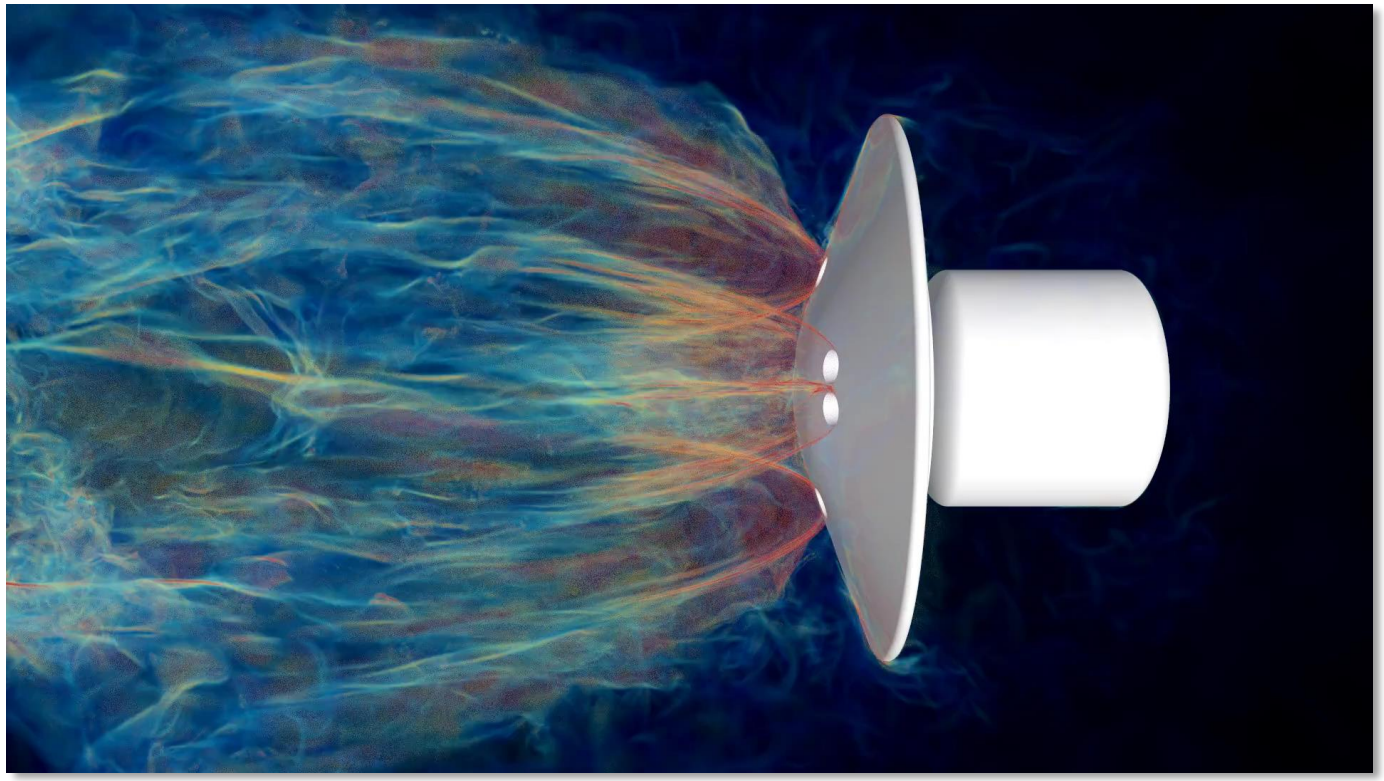

Side view

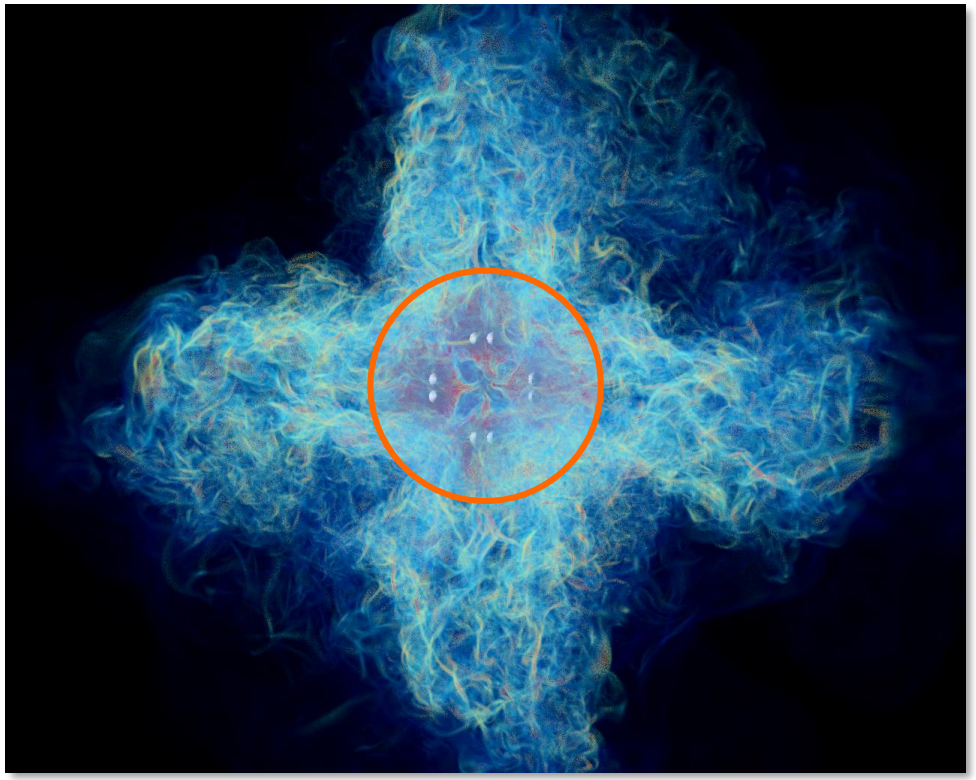

Facing the heatshield

Vehicle-level design decisions are directly impacted by the ability to characterize and bound aerodynamic-propulsive interference effects

Game-Changing Performance

Typical Job of 6.5B Elements, 200K Time Steps, 200TB Output

Conventional Computing Approach

- 9 months per run on 5,000 Xeon Skylake cores (3 months compute, 6 months queues)
- Multiple runs would take years

Current Summit Campaign

- 4 days per run on 552 V100s
- 6 simultaneous runs on 3,312 V100s
  - → 6 jobs done in a workweek
  - → Equivalent throughput of ~600,000 Xeon Skylake cores

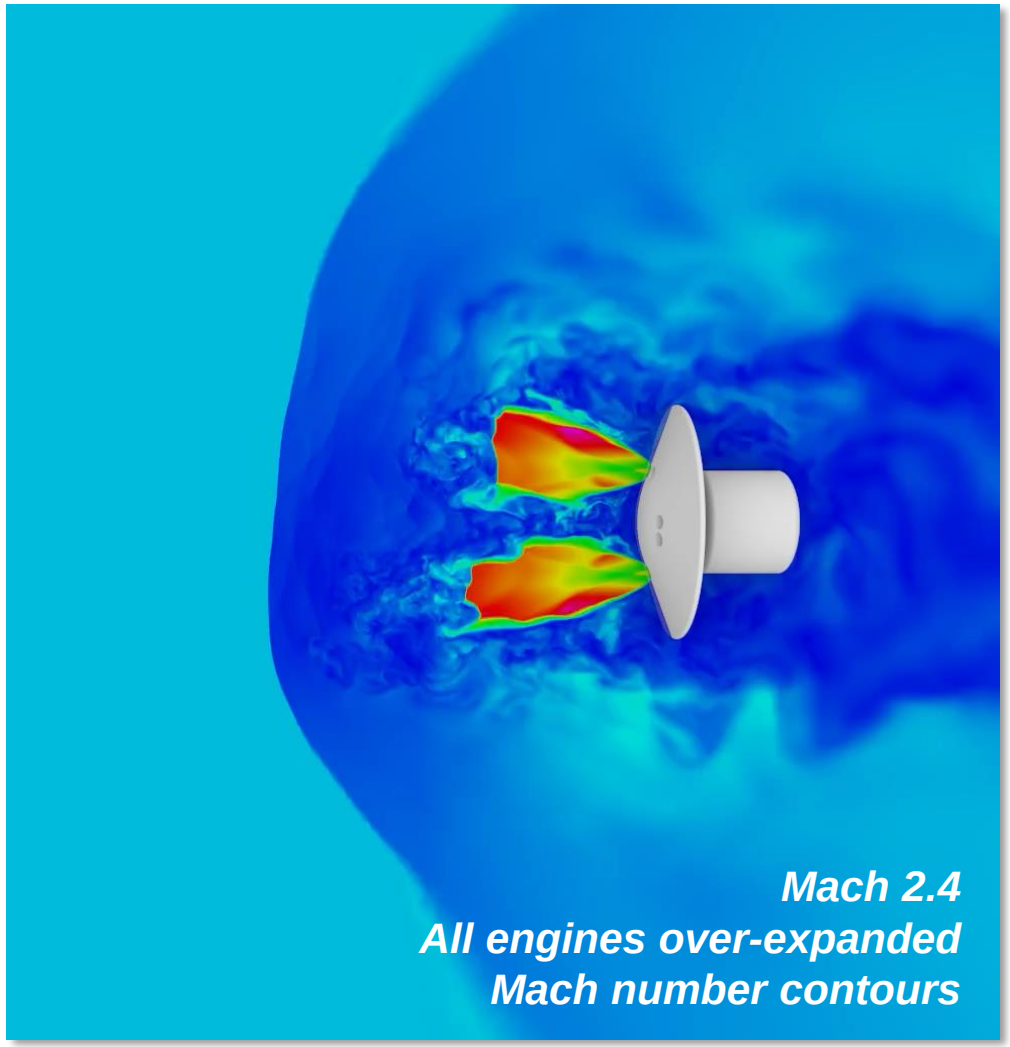
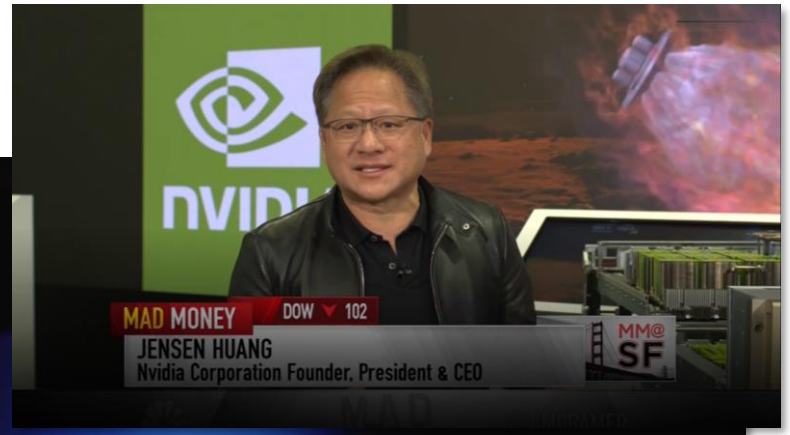
Thank you for having us!