2023 R2 Fluent Multi-GPU Solver Update

C. Kurt Svihla
Ansys, Inc.

Ansys

Agenda

- Fluids Capability Chart
- GPU Introduction & Background
- Licensing & Hardware
- 23R2 Features
- Demo
- Summary
- Learn More & Stay Connected

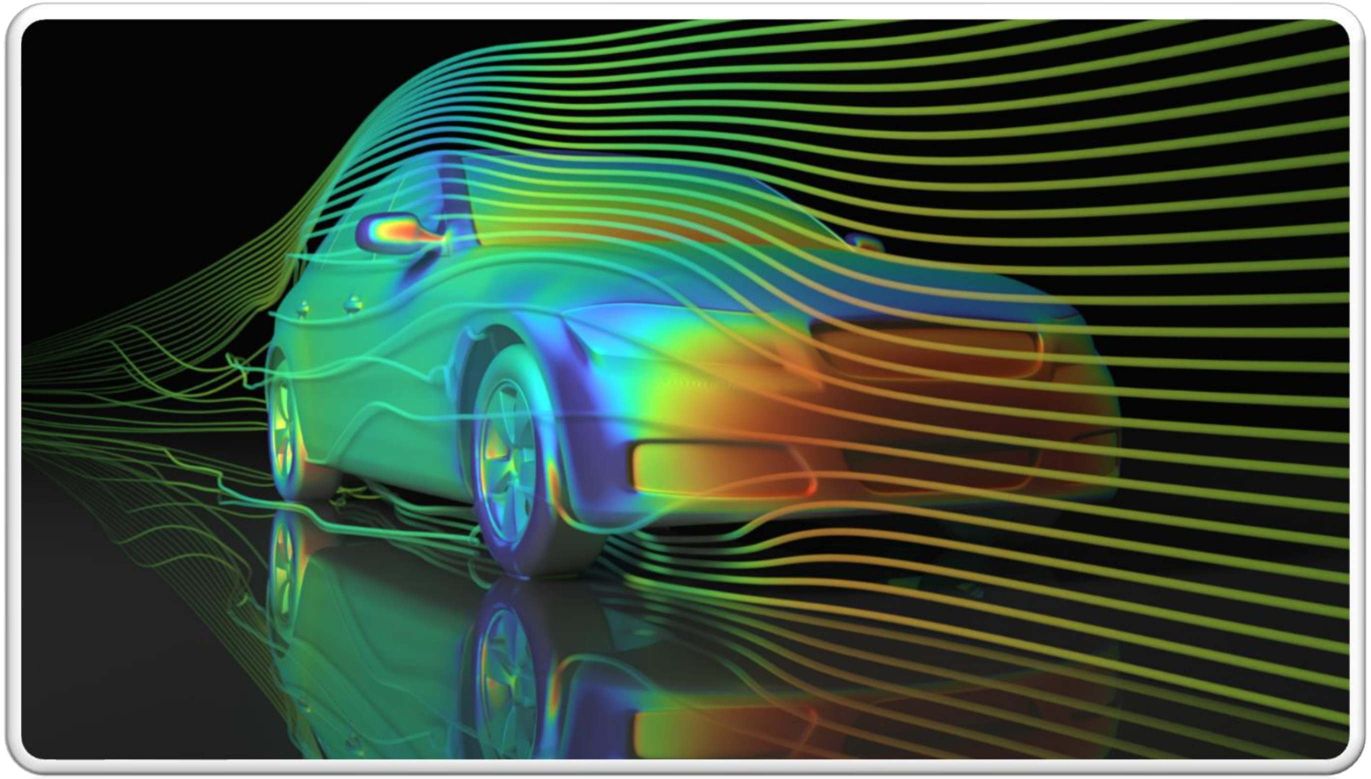

Fluids Capability Chart

	CFD Pro	CFD Premium	CFD Enterprise	CFD PrepPost	CFD Premium Solver	CFD Enterprise Solver
Fluent (CPU)	Included (limited capabilities)	Included	Included	Included**	Included*	Included*
Fluent GPU Solver			Included			Included
CFX		Included	Included	Included**	Included*	Included*
FENSAP-ICE / Fluent Icing			Included	Included**		Included*
Forte		Included	Included	Included**	Included*	Included*
Chemkin & Reaction Workbench			Included			Included
Model Fuel Library - Encrypted		Included	Included	Included	Included	Included
Polyflow - CFD for polymer, glass, metals and cement processing	Included (limited capabilities)	Included (limited capabilities)	Included	Included**	Included* (limited capabilities)	Included*
Geometry Prep - SpaceClaim / Discovery Modeling		Included	Included	Included		
EnSight - Post Processing		EnSight Enterprise	EnSight Enterprise	EnSight Enterprise		
Meshing	Ansys Meshing, Fluent Meshing[1] [1]limited capabilities	Ansys Meshing, Fluent Meshing, ICEM CFD, TurboGrid, OptiGrid	Ansys Meshing, Fluent Meshing, ICEM CFD, TurboGrid, OptiGrid	Ansys Meshing, Fluent Meshing, ICEM CFD, TurboGrid, OptiGrid		
DesignXplorer - Parameter Optimization		Included	Included	Included	Included	Included
Ansys Customization Suite - Create, share and sell simulation apps			Included	Included		
High Performance Computing - Additional cores optional	4 Cores	4 Cores	4 Cores (CPU) 40 SMs (GPU)		4 Cores	4 Cores (CPU) 40 SMs (GPU)

Fluent Capabilities at Different License Levels

Fluent Capabilities Details	Pro	Premium	Enterprise
CAD Import	✓	✓	✓
Steady-state flow (pressure-based coupled solver)	✓	✓	✓
Transient flow (pressure-based segregated solver)	✓	✓	✓
Heat Transfer: fluid energy model, CHT, shell conduction	✓	✓	✓
Basic turbulence models – inviscid, laminar, k-epsilon, k-omega (standard and SST) and Spalart-Allmaras	✓	✓	✓
HPC: 4 HPC cores included, additional core accessed via ANSYS HPC licensing	✓	✓	✓
Porous media (isotropic, orthotropic and conical)	✓	✓	✓
Incompressible, Compressible (ideal gas), Bousinesq approximation, non-Newtonian fluids	✓	✓	✓
Rotating reference frames, multiple reference frames (without periodic interfaces)	✓	✓	✓
Inert and Massless particle tracking	✓	✓	✓
2D fan model	✓	✓	✓
Multi-stream mixing (multiple non-reacting species)	✓	✓	✓
Fluent Meshing: Watertight Meshing workflow including Polyhedral, Poly-Hexcore with Mosaic technology, Tetrahedral and Prism meshing	✓	✓	✓
Fluent Setup and Post Processing, including Reports and Expressions	✓	✓	✓
Parameters, parametric workflow and design points	✓	✓	✓
Workbench integration		✓	✓
Batch solving		✓	✓
TUI (Text User Interface)		✓	✓
Fluent Meshing: Parallel Meshing and Fault Tolerant Meshing		✓	✓
Steady-state flows (all solvers)		✓	✓
Transient flows (all solvers)		✓	✓
Advanced turbulence models		✓	✓
Density-based solver, pressure-based segregated solvers, adjoint solver, 6-DOF solver for flow-driven solid motion		✓	✓
Multiphase including free surface, cavitation, phase change models, thin film model, transition to/from particles, boiling model, surface tension, reactions, granular model and Dense Particulate Coupling (DDPM)		✓	✓
Full particle tracking including break-up, coalescence, erosion, evaporation, thin wall films, reactions and macroscopic model		✓	✓
Chemical Reactions: combustion and reacting flows including pollutants/soot modeling, FGM and Finite Rate combustion models, effusion cooling models, electro-chemistry, corrosion, fuel cells, electrolysis, battery models		✓	✓
Heat Transfer: radiation models, heat exchanger model		✓	✓
Real fluid properties (e.g. steam, cryogenic fluids)		✓	✓
Turbomachinery and Rotating Models: turbomachinery aerodamping and blade flutter analysis, blade film cooling model, periodic interfaces, turbo interfaces including mixing plane and pitch change		✓	✓
Dynamic mesh, overset mesh and mesh adaption (including Polyhedral Unstructured Mesh Adaption)		✓	✓
Acoustics		✓	✓
1-way and 2-way physics couplings, including Fluid-Structure Interaction and thermo elasticity		✓	✓
Multi-GPU solver			✓
HPC: 40 SMs (streaming multiprocessors) for GPU solver, additional SMs accessed via Ansys HPC licensing			✓
Fluent Icing workspace			✓
Fluent Aero workspace			✓
Polyflow Capabilities Details	**Pro**	**Premium**	**Enterprise**
Extrusion, co-extrusion, fiber spinning, blow molding and thermoforming (no viscoelasticity or yield stress models)	✓	✓	✓
Mixers, screw extruders, filling, 2D/3D forming, viscoelasticity & materials with yield stress			✓

Ansys

GPUs Are Critical to the Next Generation of HPC

Engineers want to use simulation to see how their ideas will perform against millions of variables

- Huge numbers of simulations
- Reduce time to market
- Need drastic increase simulation throughput

→ <u>Without compromising accuracy!</u>

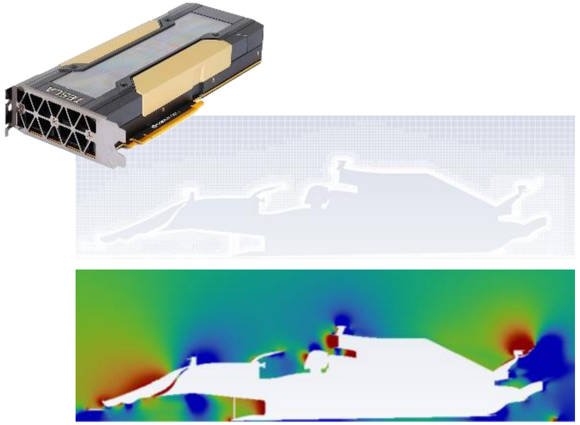

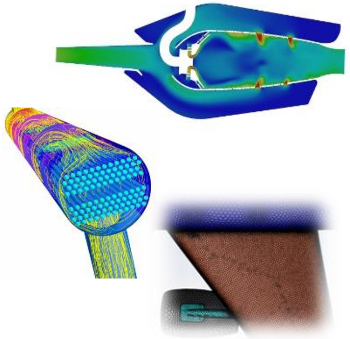

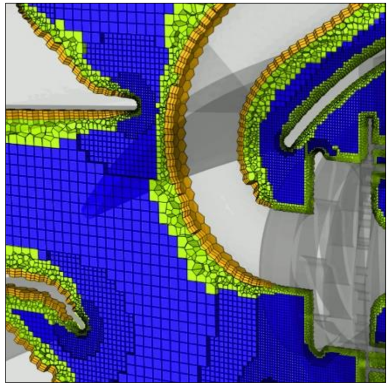

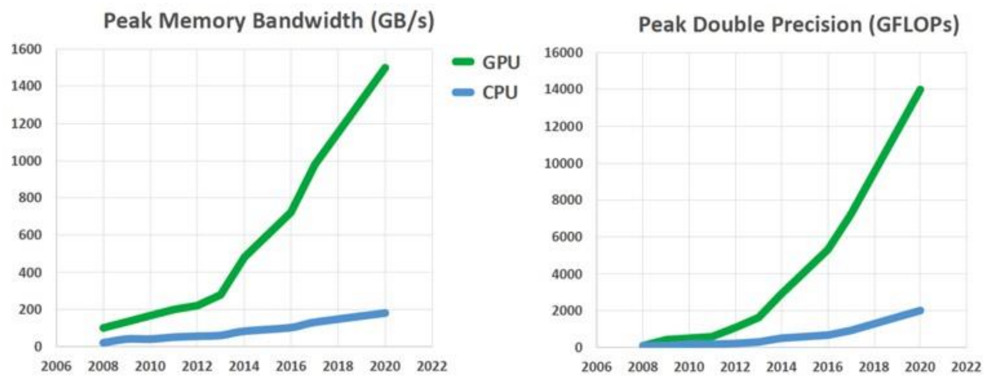

The time is now for general purpose GPUs in scientific computing

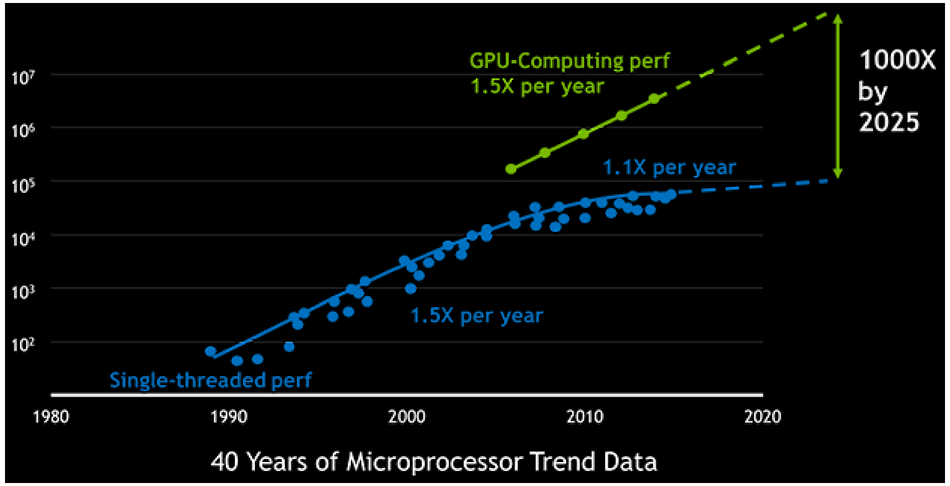

https://www.nextplatform.com/2019/07/10/a-decade-of-accelerated-computing-augurs-well-for-gpus/

Ansys Solutions Have Leveraged GPUs For Years

- Ansys Flagship Products such as Fluent, Mechanical, HFSS and Maxwell use GPUs as accelerators in “offload”
  - GPUs have large bandwidth and computational throughput
  - Good for isolated, expensive computations
  - Solver modules like linear algebra, ray tracing, radiation models can run efficiently on GPUs

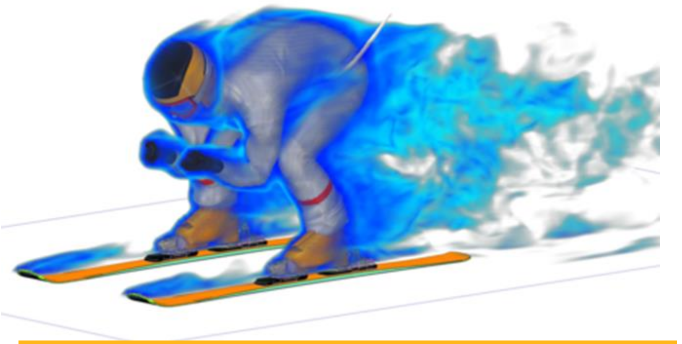

Fluent

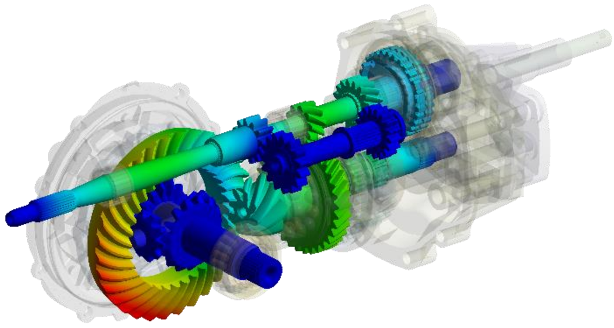

Mechanical

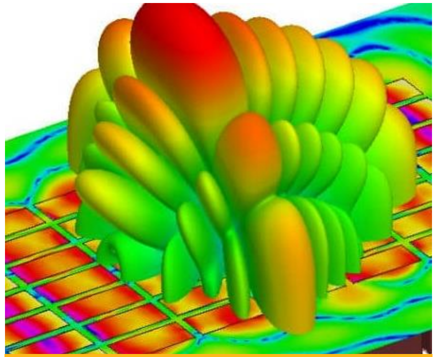

HFSS/SI-Wave

Ansys

Using GPUs to Accelerate CFD Solutions is Not New

AMGX solver

- Linear solving off d to GPU
- Beneficial for ex ro LES with coupled solver
- Enhancement case
- 4X speedup for this case
  - ➢ 28 CPU cores + 4 GPUs per node
  - ➢ 32 CPU cores only per node

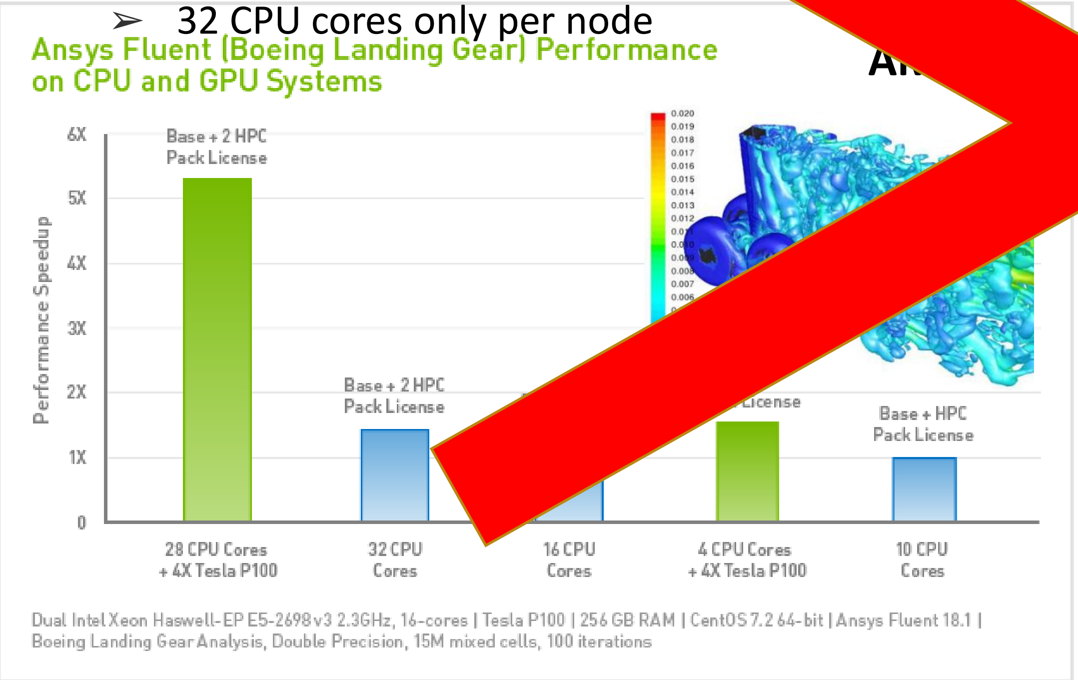

DO radiation solver

- Mixed CPU/GPU co puting
- General benefi rved
- 3X speedu O 14M cell case
  - ➢ 28 GPUs per node vs
  - ➢ es only per node

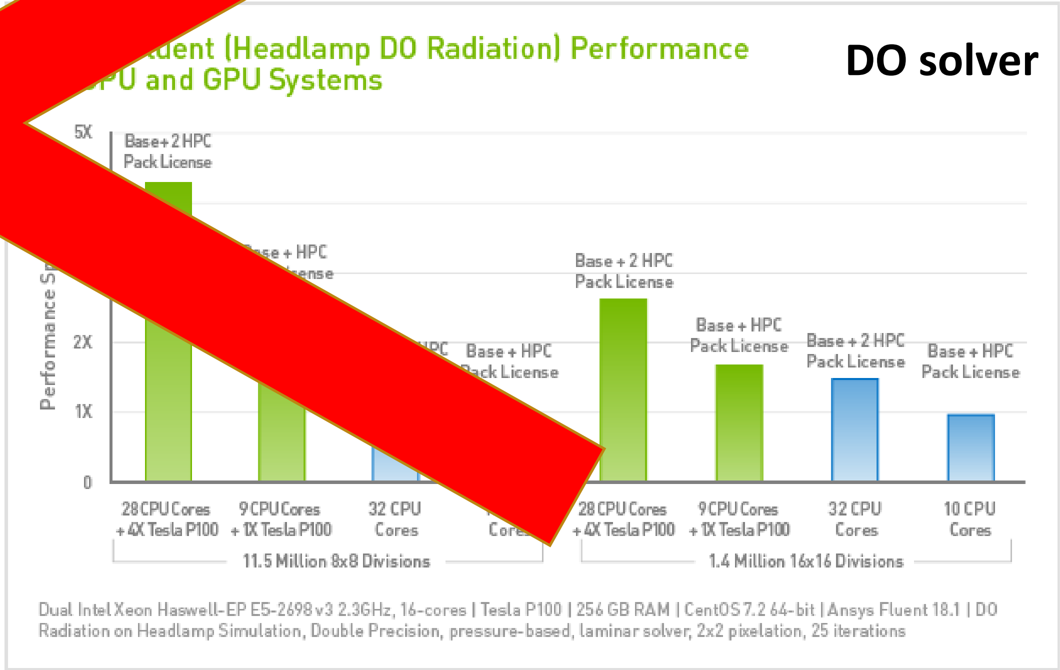

Native GPU Implementation Shows Outstanding Performance Gains

Ansys / FLUENT

DrivAer 250M LES Speedup with 176 CPU cores as base
One A100 40GB replaces ~540 CPU cores, or ~12 nodes
One 8-GPU node replaces ~96 high-end CPU servers

- 2 typical GPUs > 1,000+ CPU cores
- ~7x cheaper hardware purchase cost
- ~4x lower power consumption

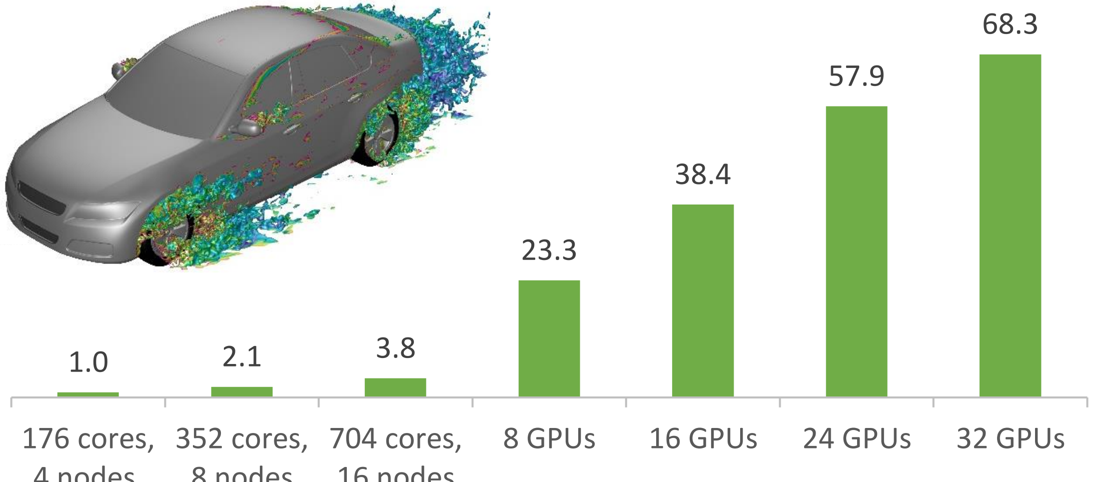

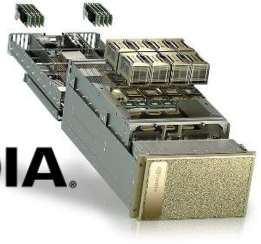

NVIDIA

Native GPU Benefits Go Beyond Fast Turnaround

- Metrics based on Cloud compute power consumption and costs from a mainstream CSP
  - CPU cluster with 1024 Intel® Xeon® Gold 6242 cores
  - GPU server with 6x NVIDIA® V100 cards

- **4X Power Consumption Reduction**
  - CPU consumption = 9,600 W
  - GPU consumption = 2,400 W

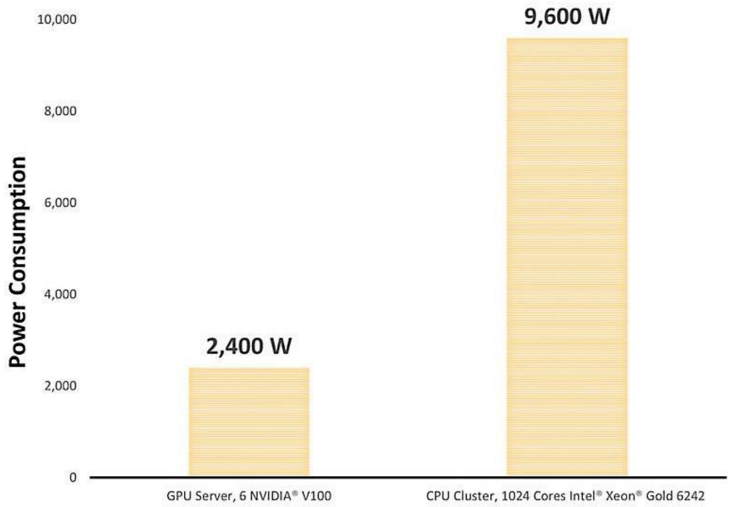

- **7X Hardware Cost Reduction**
  - HW cost for CPU ≈ $500K
  - HW cost for GPU ≈ $71K

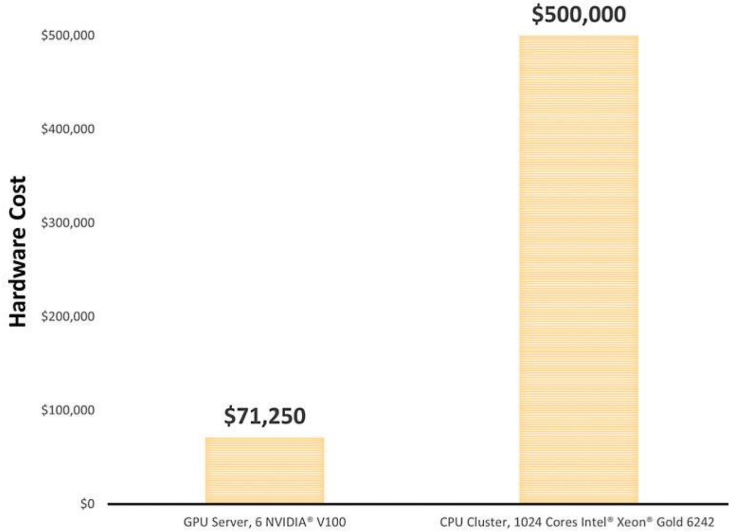

Fast & Memory-Efficient

- *“One DGX replaces 30+ high-end dual CPU servers”*
- 1 – 1.5 GB/M cells Single Precision; 1.5 – 2 GB/ M cells Double Precision

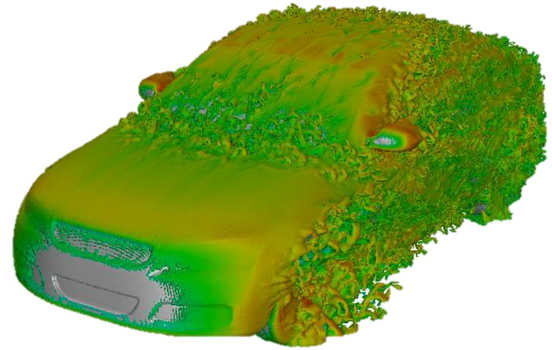

Low Memory footprint
100M cell case fit into 1 single A100 !

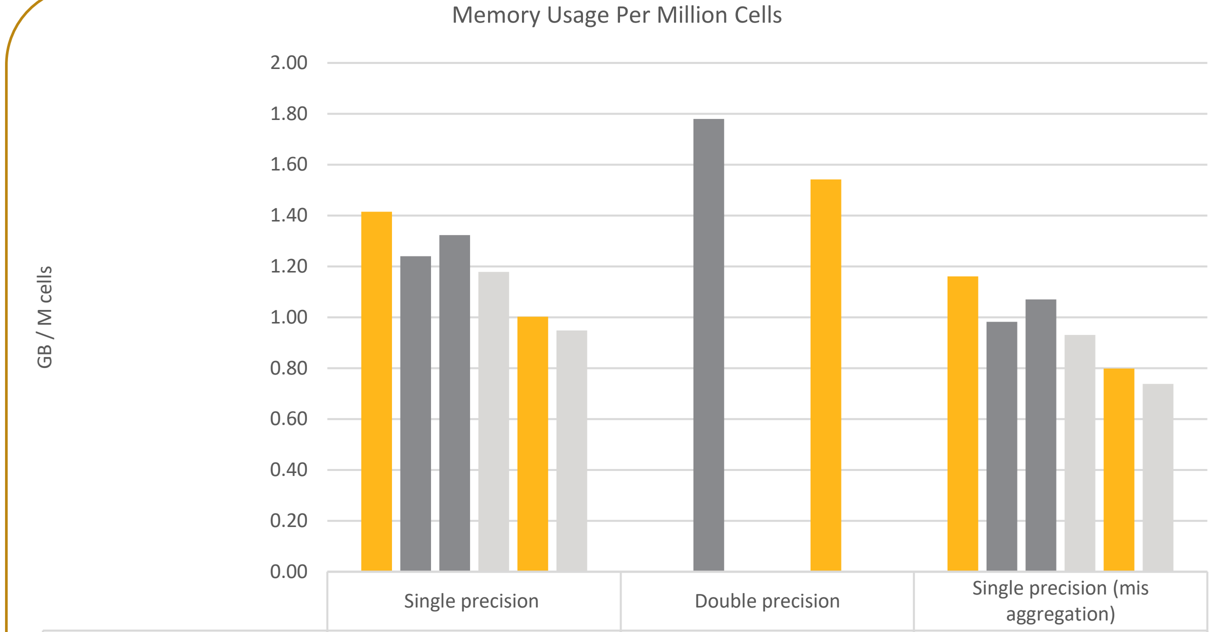

	Single precision	Double precision	Single precision (mis aggregation)
DrivAer 5M hex-poly	1.42		1.16
Airplane 6M polyhedra	1.24	1.78	0.98
DrivAer 10M hex-poly	1.32		1.07
DrivAer 20M hex-poly	1.18		0.93
Winsor Car 41M rapid-octree	1.00	1.54	0.80
F1_Racecar 140M hex-core	0.95		0.74

Ansys

Multi-GPU Solver in Fluent: Release Timeline

UTILIZE THE POWER OF MULTIPLE GPUS TO ACCELERATE YOUR CFD SIMULATIONS

2022 R1

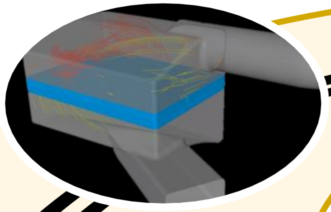

Beta support for :
- Single/multi-GPU (shared / distributed memory)
- Steady & transient simulations
- Incompressible & subsonic compressible flows
- All mesh types
- Ideal Gas and Materials with variable properties
- Turbulence: standard k-e, SST, GEKO, RKE, SBES
- Solid conduction and CHT
- Porous media
- Windows and Linux

2022 R2

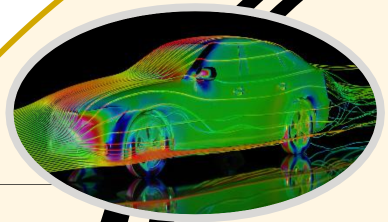

Beta support for :
- Transient scale-resolving simulations
- Non-conformal mesh interfaces
- Moving walls & Moving Reference Frame

2023 R1

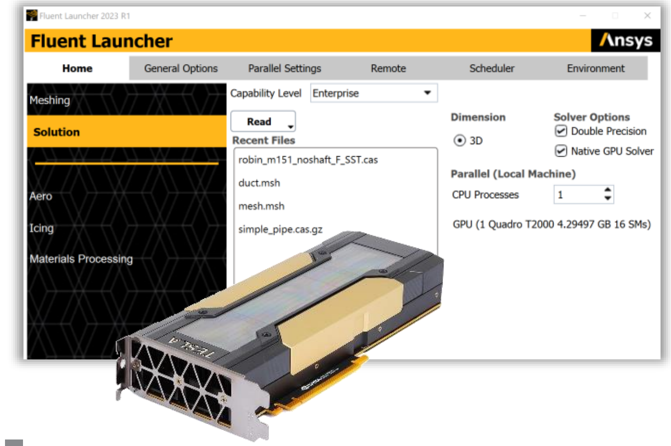

All previous features released, along with :
- Various numerics improvements
- Enhanced robustness and accuracy for RANS
- Species transport
- Enhanced LES numerics
- Improved UX and launcher enhancement
- …

2023 R2

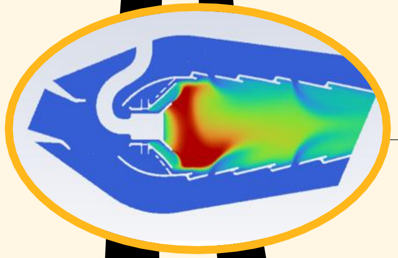

All previous features released, along with :
- Sliding meshes
- Non-stiff reacting flows
- Compressible flows
- …

Ansys

Can I Run the Fluent GPU Solver in Workbench?

The Fluent GPU Solver under Workbench can be turned on with an environment variable:

- Windows
  - `set FLUENT_GPU_DEVICES=all`
- Linux
  - Csh
    - `setenv  FLUENT_GPU_DEVICES all`
  - Bash
    - `export FLUENT_GPU_DEVICES=all`

Licensing & Hardware

What are Streaming Multiprocessors (SMs)?

- HPC licensing for the native GPU solver is based on the total number of streaming multiprocessors (SMs) used across all GPU cards
- GPU cards are made up of many SMs, and each SM contains many CUDA cores
- More powerful GPU cards typically contain more SMs

Tesla P100 Layout: 56 SMs, 3584 CUDA cores

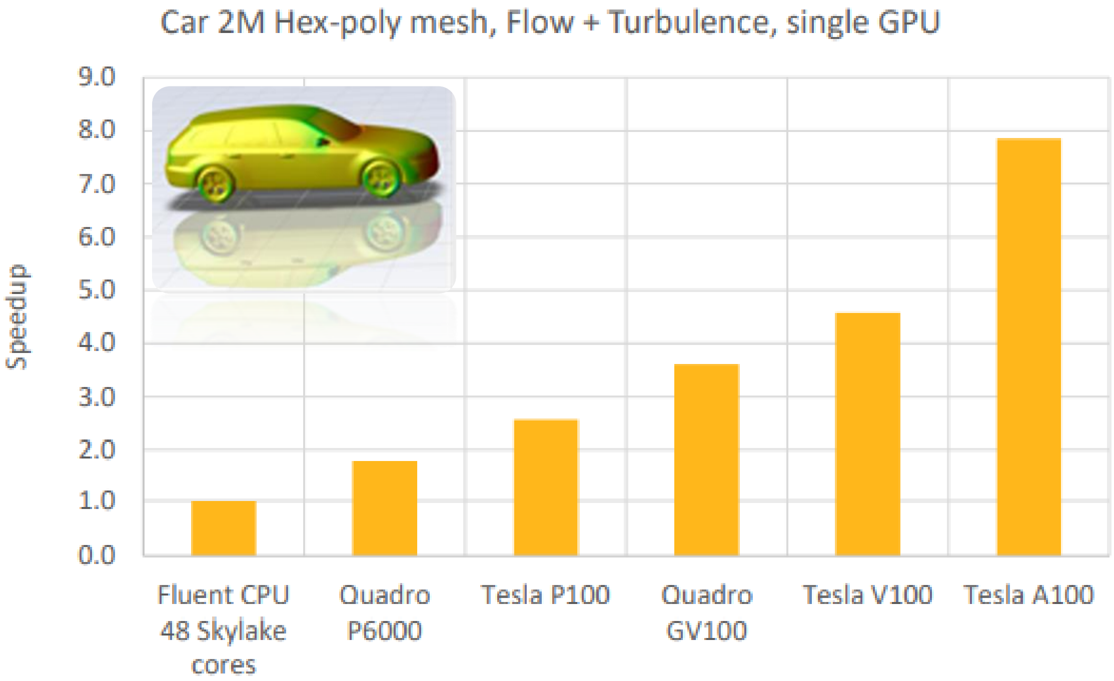

	Fluent CPU 48 Skylake cores	Quadro P6000	Tesla P100	Quadro GV100	Tesla V100	Tesla A100
# SMs		28	56	80	80	108
# CUDA		3840	3584	5120	5120	6912

Licensing

- 40 SMs are included with CFD Enterprise
- Additional SMs enabled with Ansys HPC/Packs/Workgroups
  - 1 HPC task = 1 additional SM
  - HPC Packs scale as usual (e.g. 3 Packs enables 128 additional SMs)

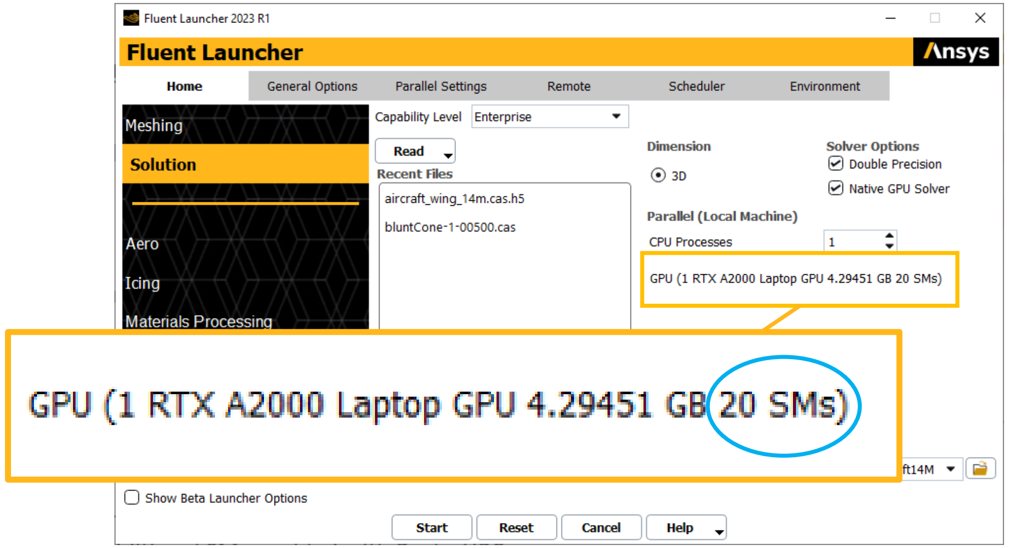

Note:

- Fluent pre/post-processing for the GPU solver can use multiple CPUs without consuming additional HPC tasks
- Licensing is based on the total number of SMs across all GPUs, irrespective of the number of GPUs
- The Fluent Launcher and the Fluent TUI display the number of SMs in use
- All available SMs are used on a GPU card. It is not possible to restrict usage to a subset of SMs

EXAMPLE

- **A100 GPU card contains 108 SMs**
- **4 x A100 GPU cards requires: (4 * 108) - 40 = 392 HPC tasks**
- **4 HPC Packs or 392 Workgroup tasks**

# SMs	HPC Workgroup	HPC Packs
1 – 40	0	0
41 – 48	1 – 8	1
49 – 72	9 – 32	2
73 – 168	33 – 128	3
169 – 552	129 – 512	4
553 – 2088	513 – 2048	5

HPC Requirements for Common GPUs

GPU Card	Description	#SMs	RAM (GB)	~CPU Core Equivalent	Max Problem Size* (1 GPU)	#HPC Packs
RTX A2000 Mobile	Released 2021. Standard laptop GPU.	20	4	6 – 12	3M cells	0
Quadro P2000 Mobile	Released 2017. Standard laptop/desktop GPU	6	4	10 – 15	3M cells	0
Quadro T2000	Released 2019. Standard laptop GPU	16	4	10 – 15	3M cells	0
Quadro P6000	Released 2016. Older high-end workstation GPU	30	24	50 – 90	20M cells	0
Quadro RTX 4000	Released 2018. Typical workstation GPU	36	8	60 – 100	7M cells	0
Tesla P100	Released 2016. Older server GPU	56	16	100 – 130	13M cells	2
Quadro RTX 6000	Released 2018. High-end workstation GPU	72	24	120 – 200	20M cells	2
Tesla V100	Released 2017. Previous gen server GPU	80	32	150 – 220	27M cells	3
Quadro RTX A5000	Released 2021. Top-end workstation GPU	64	24	130 – 220	20M cells	2
Tesla A100	Released 2020. Top-end server GPU	108	40, 80	200 – 400	33M, 67M cells	3
Tesla H100	Expected Q3 2022. Next-gen server GPU	132	80	Not yet tested	67M cells	3

** Assuming 1.2GB RAM per millions cells. Actual RAM requirements will be case specific and will depend on the mesh type, physics solved, single vs double precision and other factors. Larger mesh sizes can be solved using multiple GPUs.*

Ansys

23R2 Features

23R2 Feature Status

Release

Low speed compressible solver
Sliding mesh
EDM combustion model
Extended monitors - Async monitors
Extended monitors - Point/cut plane monitors
Extended monitors - Mass averaged and Sum
Expanded profiles - species + cell zones
GPU/CPU remapping - Invoke with -gpu_remap

Beta

Coupled solver
Rotational periodic
Anisotropic heat conductivity
UDS - Support constant or profile conditions
Morphing and parametric
Poor mesh numerics
DVS export directly
Fast restart with lightweight Fluent

Ansys

GPU Solver: Sliding interfaces

Model rotating and stationary components

- An efficient and robust sliding mesh algorithm
- Enhanced robustness of dramatically changing intersections
- Test case: F1 140M rotating wheels, URANS SST

1 GPU ~ 16 CPU nodes (704 cores)
With 8 GPUs, solve 1000 steps in 40 minutes!

Intel Platinum 8168 dual-socket cluster
Nvidia A100 40GB 8-GPU per node

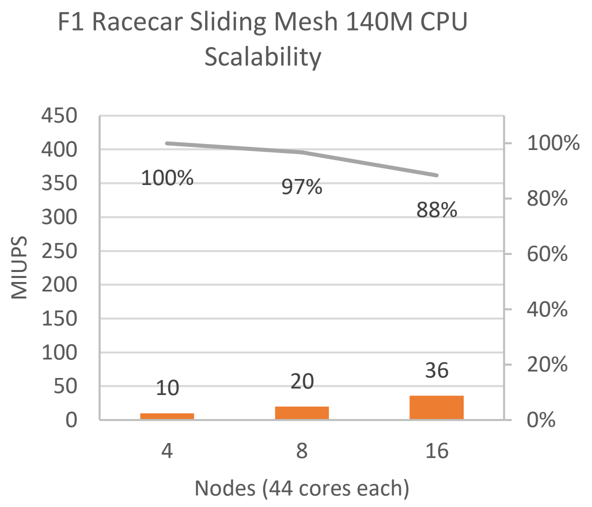

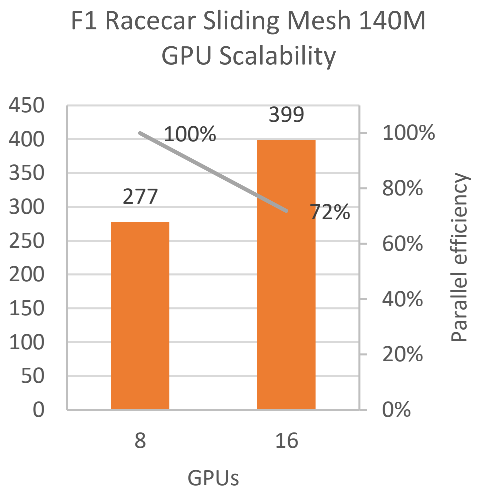

GPU Solver: Compressibility

The Fluent GPU solver now supports low-speed compressibility in addition to incompressible flow

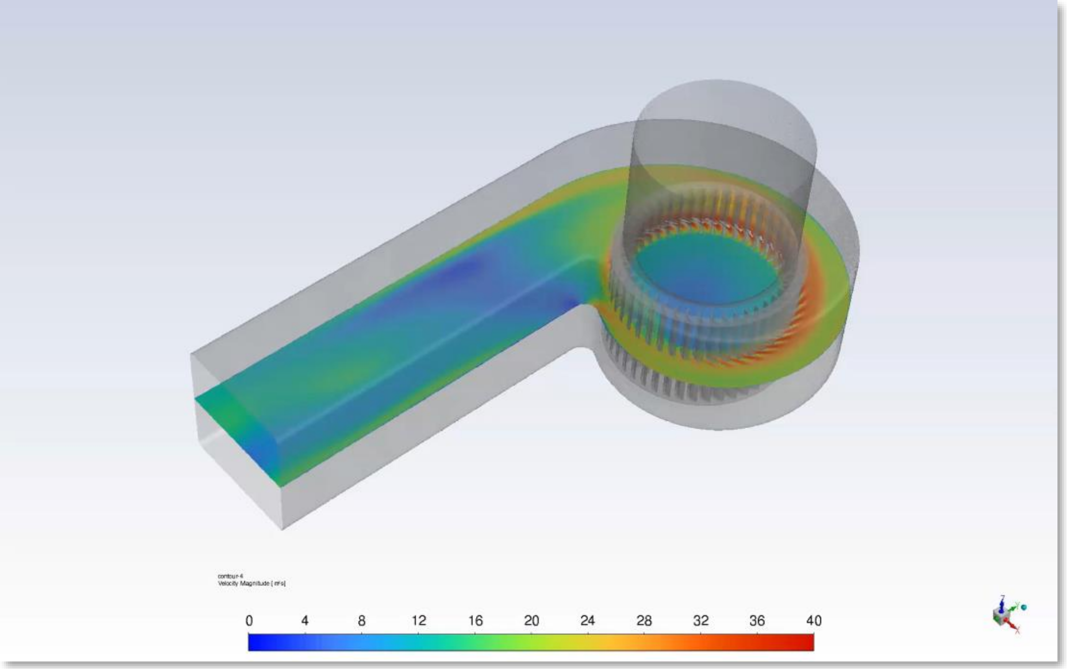

Centrifugal blower with compressible flow and sliding mesh

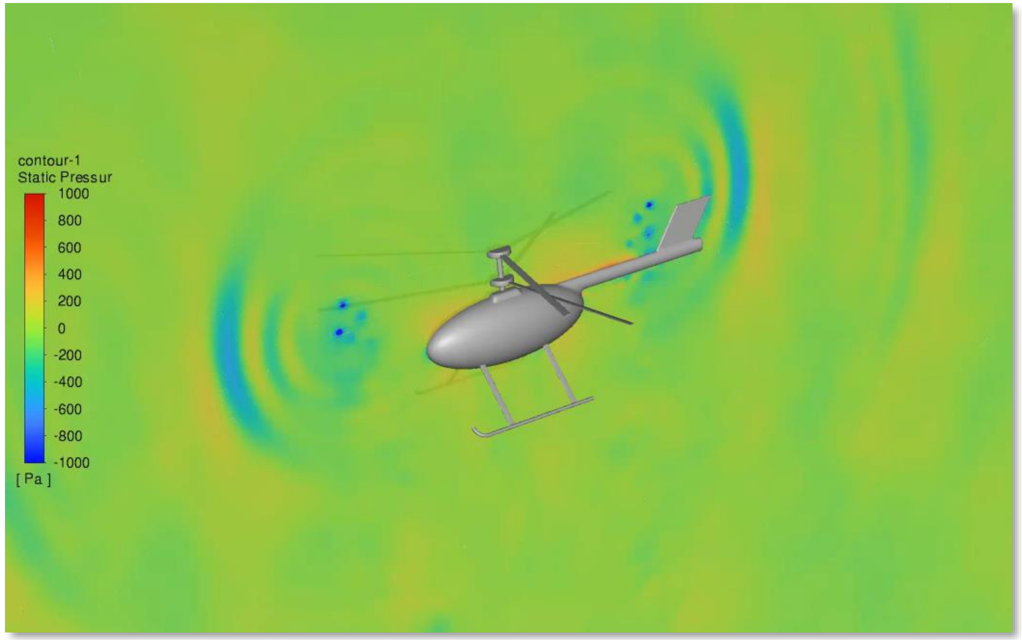

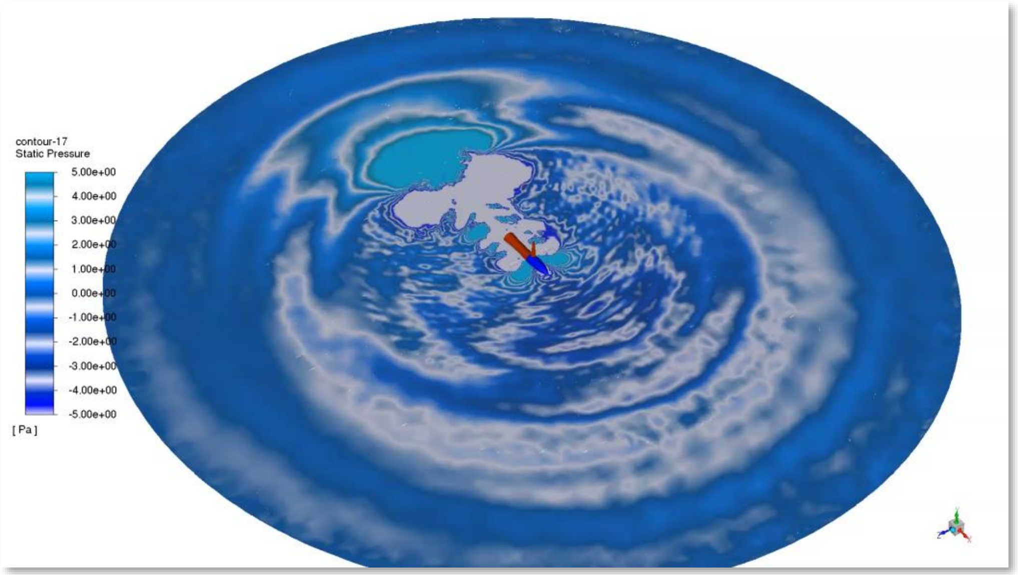

GPU Solver: Combustion

The Eddy Dissipation Model is now available with the GPU Solver

- Reaction rates obtained from turbulence parameters
- Improved robustness for variable density flows
- Test case
  - Generic Combustor, 12M polyhedral cells
  - 5 species, 1 reaction
  - EDM combustion model
  - Realizable K-epsilon turbulence
  - Intel(R) Xeon(R) CPU E5-2690 v3 @ 2.60GHz vs NVIDIA A100 80 GB

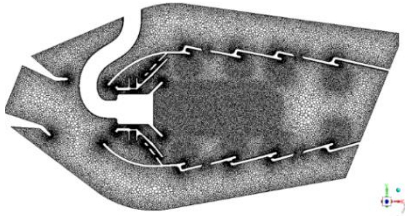

Intel(R) Xeon(R) CPU E5-2690 v3 @ 2.60GHz vs NVIDIA A100 80 GB

Single GPU ~=576 CPU cores

Two GPU ~=860 CPU cores

Matching static temperature profiles

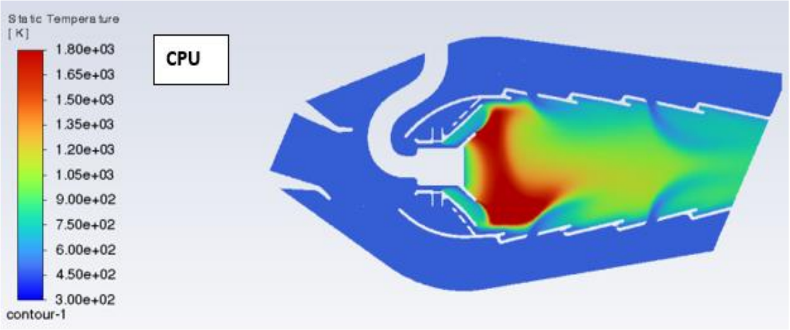

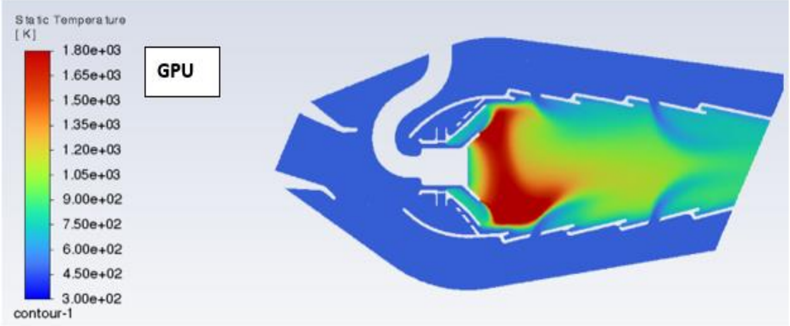

Matching maximum and average temperature and mass fractions

	CPU	GPU
Area-Weighted Average Static Temperature [K] — outlet_hot	814.90027	816.33144
Maximum of Facet Values Static Temperature [K] — outlet_hot	1120.1417	1109.9822
Maximum of Facet Values Mass fraction of co2 — outlet_hot	0.046799511	0.045978013
Area-Weighted Average Mass fraction of co2 — outlet_hot	0.023473894	0.023566746

GPU Solver: Turbulence

- Enhanced robustness
- RANS equation solution frequency controls for SBES

400 steps	Timing standard	Timing when updating k+ω every 5th step only
CPU (Milan 128 cores)	11792.5 [29.5 s/dt]	6533.2 [16.3 s/dt]
GPU (DGX 8 x A100)	357.6 [0.9 s/dt]	**255.7 [0.6 s/dt]**
Speedup	32.8 (1 GPU ~ 512 CPU cores)	27.2 (1 GPU ~ 435 CPU cores)

Auto external aero case, 156M with SBES, DGX: 8 x A100 80GB, enables less than 1 day turn-around

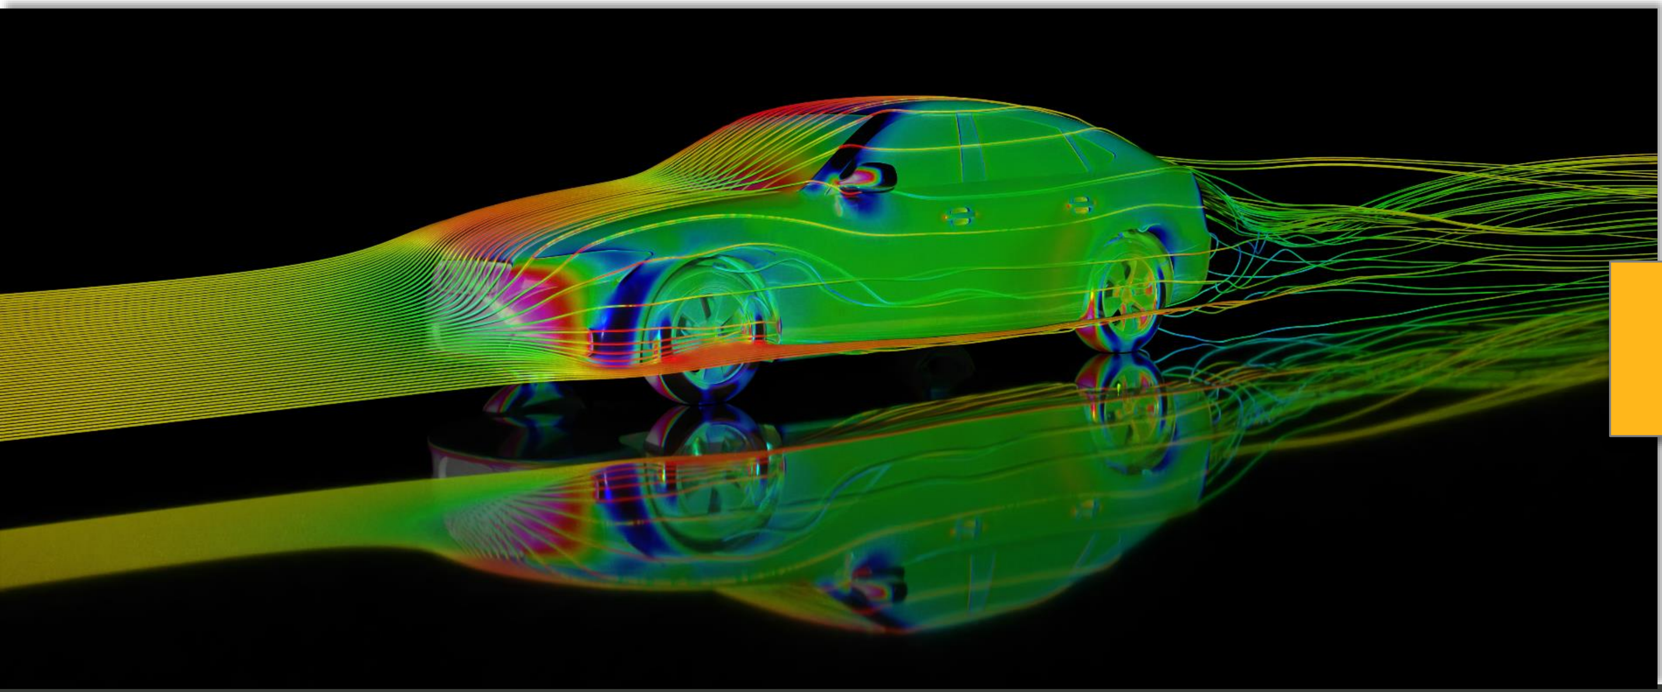

0.6 s per time step for 156M cell external aero calculation

Ansys

Monitors and residuals

Asynchronous monitors and residuals

- GPU Solver running asynchronously from UI
- Significantly enhanced throughput
- Invoke with -gpu_async
- Test case: Proprietary auto external case, 160M LES
  - Batch mode, 20~40% gains, eliminated overhead of monitors
  - GUI mode, 100+% gains

Support for point, plane and line monitors

Auto external aero case, 160M LES, **time per step**, and performance gains with async monitors and residuals

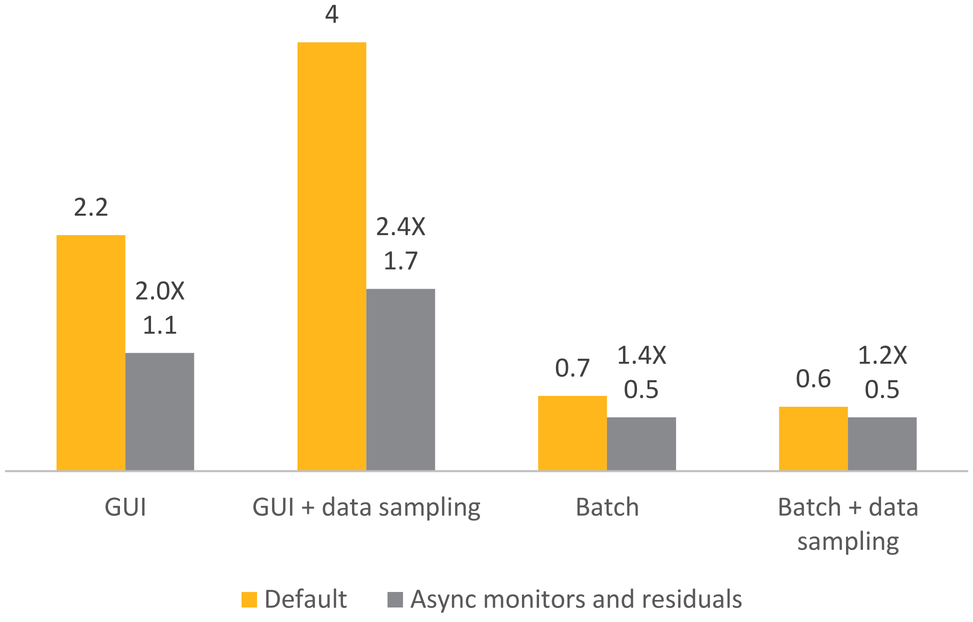

Extended monitors

- Added mesh object support
  - Point monitors
  - Cut plane monitors
    - All three plane types – coordinate plane, point and normal, three point
    - Both per-surface? selection and cumulative selection (multiple surfaces)
- Added support for new field variables:
  - Dynamic Pressure
  - Total Pressure
  - Pressure Coefficient
  - Velocity Magnitude
  - Mach Number
  - Q Criterion Normalized
  - Lambda 2 Criterion
  - Wall Shear Stress
- Added averaging support
  - Mass averaged
  - Sum
- Optimized Fluent coupling with GPU monitors
  - Removed CPU side overhead, the default sync mode is improved significantly from before
    - E.g. 52” → 34” for one case

CPU/GPU remapping

- Enable configurable CPU to GPU count mapping, instead of fixed 1:1
- Help with initialization, post processing, I/O and other CPU intensive operations
- Invoke with command-line -gpu_remap
- Test case: Generic_F1_310M
  - t8 original – t8 with 8 GPUs
  - t32 remap – t32 with 8 GPUs remapped
  - I/O time significantly reduced
  - Setup time moderately increased
  - Iteration time lightly increased
  - Overall, 25% performance gain

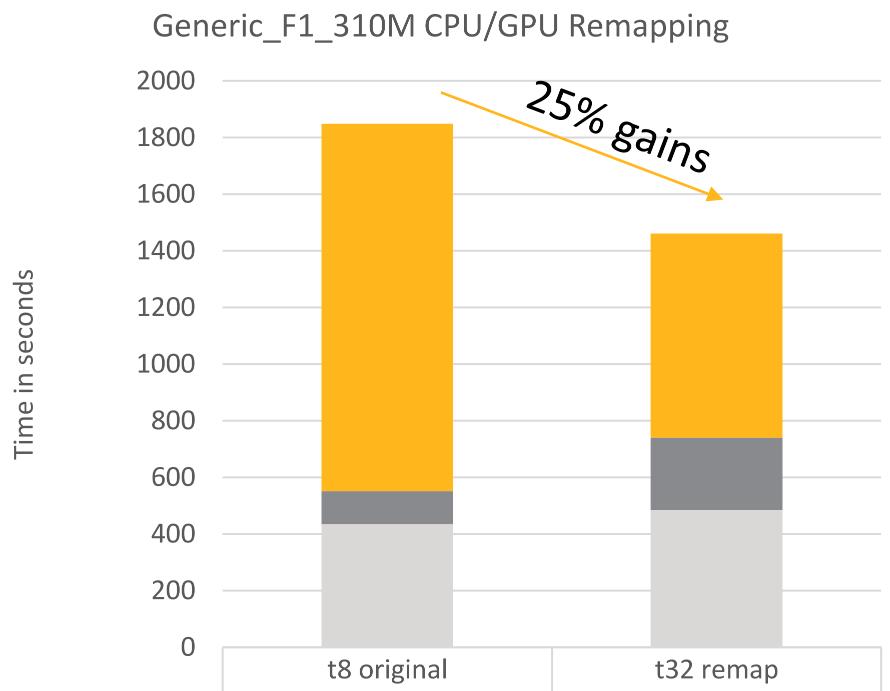

	t8 original	t32 remap
Case read	1297	721
Setup	117	255
500 iterations	435	485

Extension to the latest NVIDIA architectures

- Extended GPU Solver build to latest architectures, e.g., Hopper and future architectures
- The 2023 R2 release runs on NVIDIA H100 and RTX 6000 Ada
- H100 has about 18~25% performance gains over A100
  - H100 1944 GB/s, 114 SMs (PCIe version)
  - A100 1944 GB/s, 108 SMs

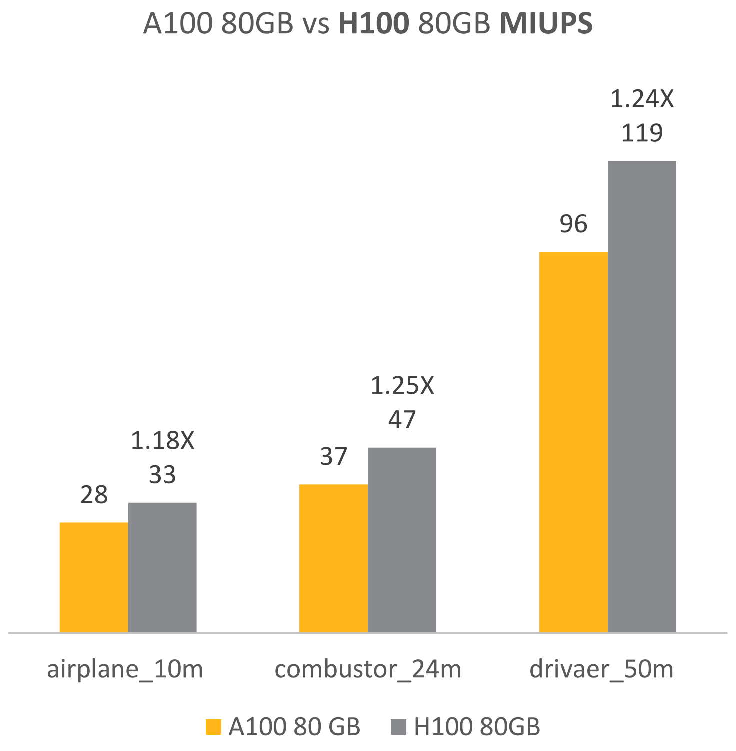

MIUPS = Million cell Iteration Updates Per Second (i.e. Iteration Updates Per Second Per Million Cells)

Benchmark Suite

10 GPU benchmark cases are now available

- car_2m, sedan_4m, airplane_10m, combustor_24m, exhaust_system_33m, winsor_41m, drivaer_50m, f1_racecar_140m, drivaer_250m, open_racecar_280m
- Different mesh sizes, types
- Segregated, coupled
- Incompressible, compressible
- RANS, LES
- Steady, transient

• Simple python scripts to conveniently run the benchmarks

Fluent GPU Solver Benchmarks

For Fluent 2023 R2 and newer

1 Where to download files

Get latest Fluent 2023 R2 or newer versions and install it

1.1 Get the new case files

OneDrive for these cases: GPU Solver Additional Benchmarks

```
car_2m, airplane_10m, combustor_24m, winsor_41m, drivaer_50m, drivaer_250m
```

1.2 Get the original standard benchmark files

Ansys customer portal for these cases:
https://support.ansys.com/TrainingAndSupport/ANSYSFluentBenchmarks

```
sedan_4m, exhaust_system_33m, f1_racecar_140m, open_racecar_280m
```

1.3 Placement

copy the .tar files to Ansys installation under [path]/ansys_inc/v<version>/fluent

```
tar -xvf <case>.tar
```

Repeat above steps for each individual case packages.

2 Commands to run

Please first make sure you have python command available from Python3 package

2.1 Example to run 8, 4, and 2 GPUs for cases airplane_10m and combustor_24m

```
python ansys_inc/v232/fluent/bench/bin/fluent_benchmark_gpu.py -gpu -cores
8,4,2 -cases airplane_10m,combustor_24m
```

2.2 Example to run 4 GPUs for case airplane_10m on AMD GPUs

```
python ansys_inc/v232/fluent/bench/bin/fluent_benchmark_gpu.py -gpu.amd -
cores 4 -cases airplane_10m
```

3 Case description

3.1 car_2m: External Flow over DrivAer car model

Size: 2M cells

Cell Type: Projected Octree

Solver: Pressure based segregated solver, Least Square cell based, steady

Models: GEKO k-omega Turbulence

GPU Solver: Expanded profiles and added UDS

- Profiles
  - Cell zone source as profiles
  - Boundary condition profiles for species and UDS
- UDS (**beta**)
  - User defined scalar transport equations with associated boundary conditions
  - Steady and transient mode
  - Support for profiles (BC and cell zone source)
  - Up to 5 UDSs
- Test case for profiles
  - Parabolic species profile for o2 and h20 at the right-end (velocity-inlet) BC

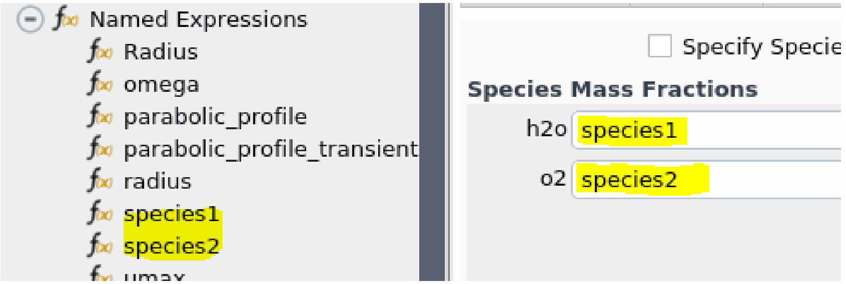

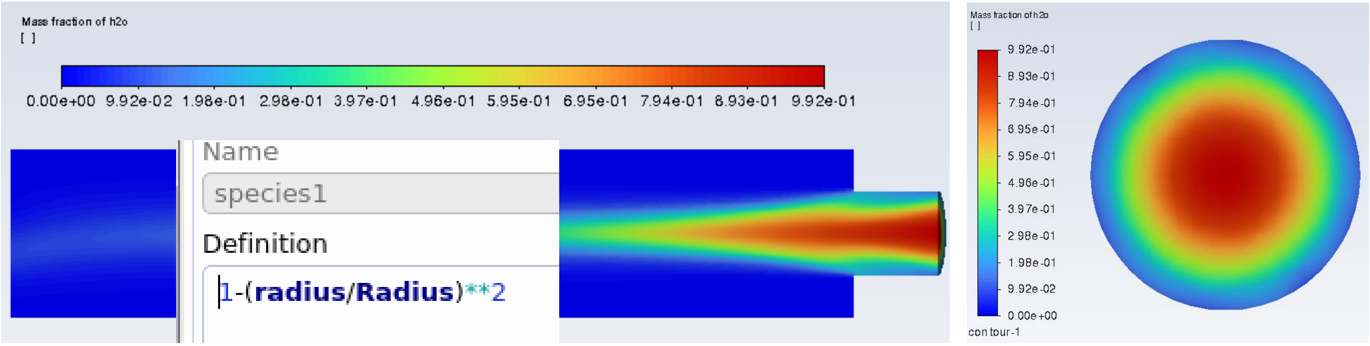

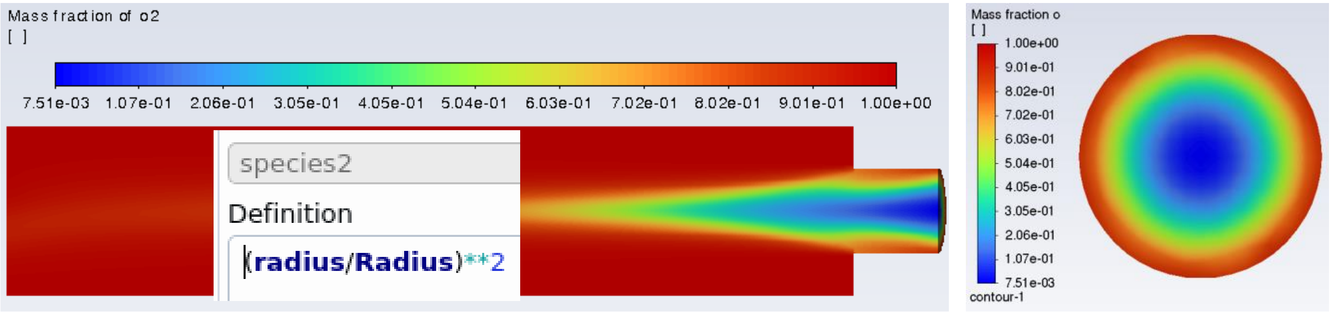

Coupled solver (beta)

- Steady-state (beta) available for testing
- Pseudo-transient (alpha) not yet ready for testing

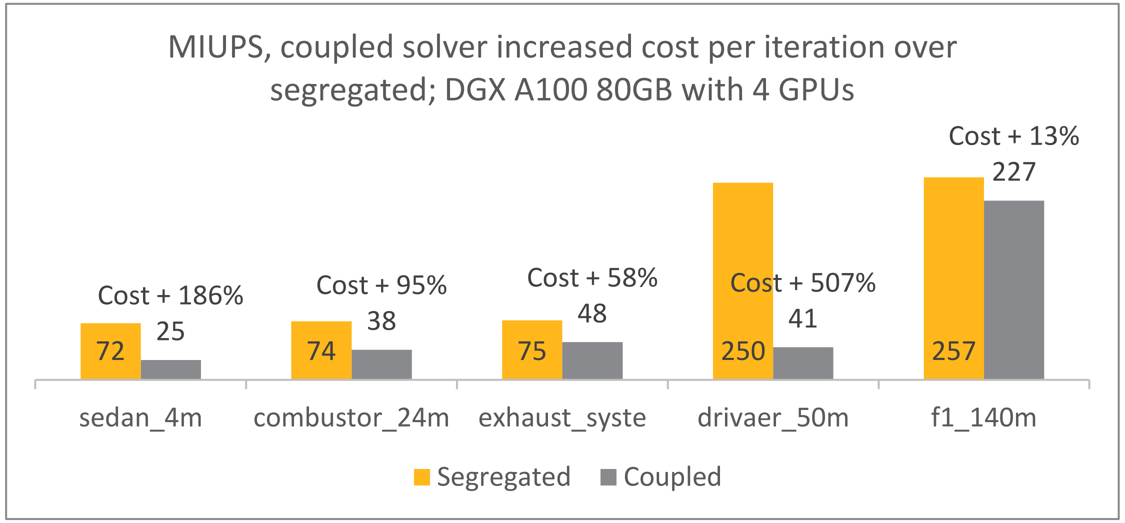

49m DrivAer Performance		
	Time/it	Time/it/mi cells
CPU (-t40)	19.6 s	0.40 s
GPU (4 cards)	1.2 s	0.024 s

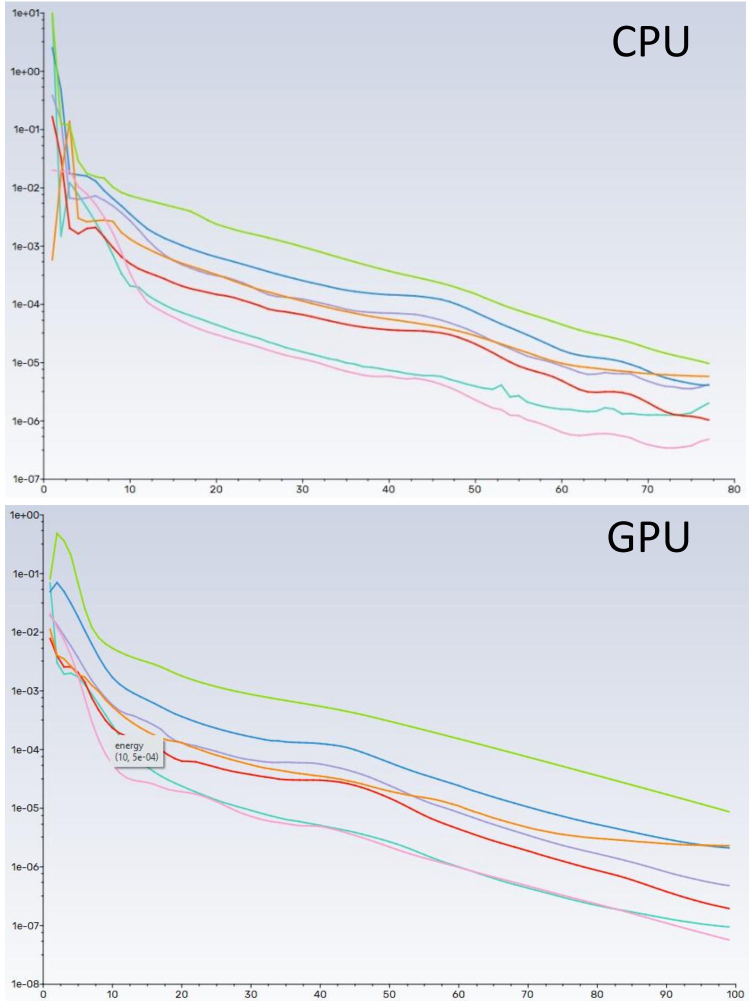

Hannover Inlet Guide Vane

Rotational periodicity (beta)

- **Rotational periodicity added**
  - Segregated, coupled, NCI, …

Test case

- Hanover IGV
- 310K cells
- Compressible
- Pressure inlet, pressure outlet
- SST k-omega

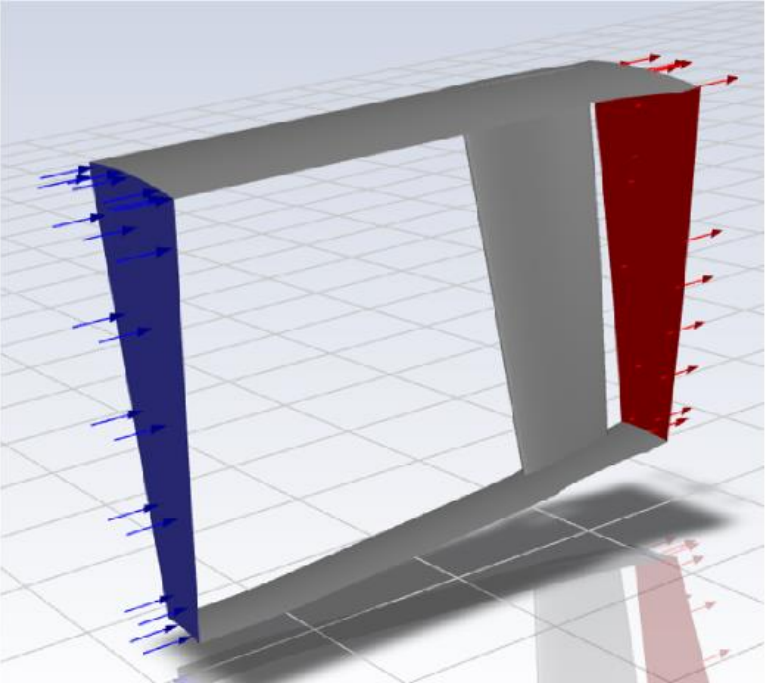

CPU – Static pressure

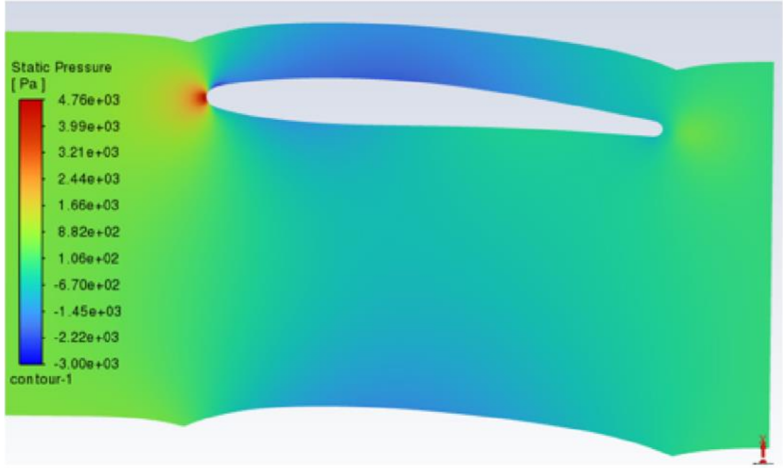

GPU – Static pressure

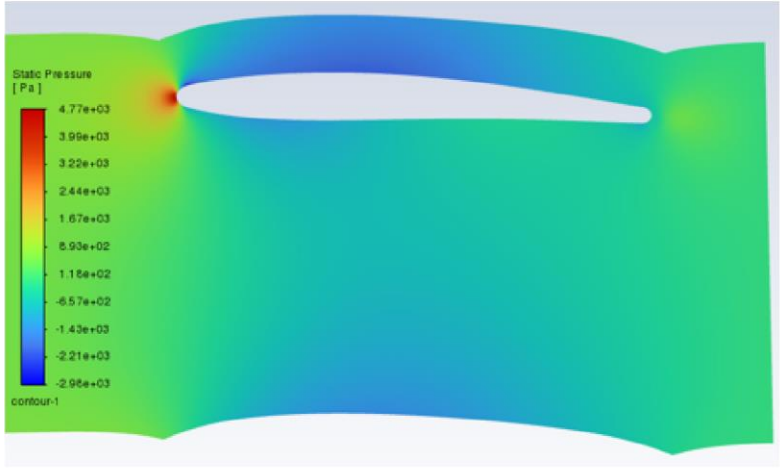

Poor mesh handling (beta)

Example: single precision run on mesh with poor elements and extreme stretching ratios

- **Poor cells identified using four criteria**
  - Minimal orthogonal skewness (<1e-04)
  - Maximum face warpage (>0.5)
  - Number of neighboring elements (>2)
  - Left handedness (<-1e-16)
- **Delete all marked cells**
  - Newly created boundary faces are set as wall boundaries
  - Marked cells visualized in poor cell marker in Fluent

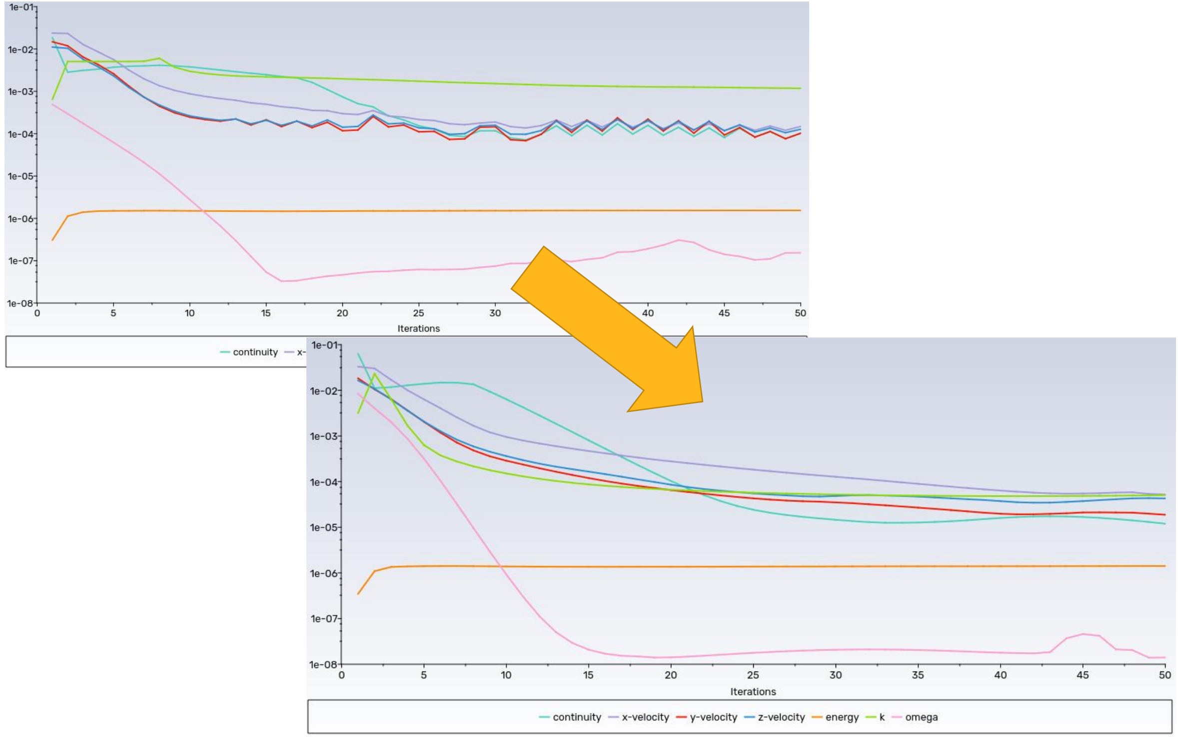

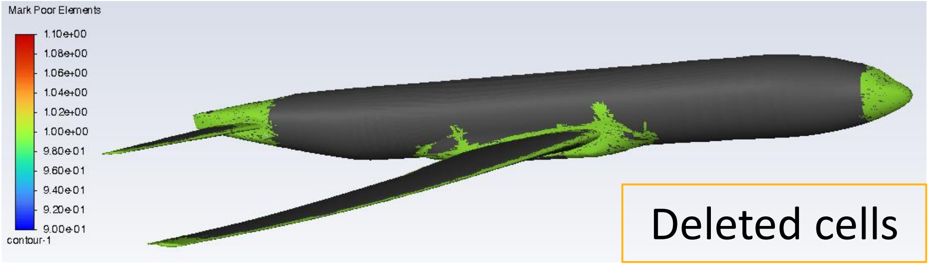

Deleted cells

Anisotropic thermal conductivity (beta)

- Solid thermal conductivity specified as a matrix:
  - Anisotropic
  - Orthotropic
  - Cylindrical orthotropic
- Matching GPU/CPU results

Test case

- E-motor, 8M cells
- 2 fluids (air and ethylene)
- MRF
- heat sources

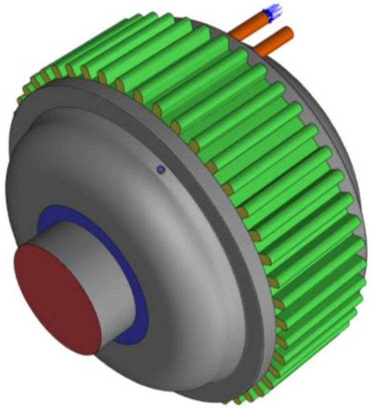

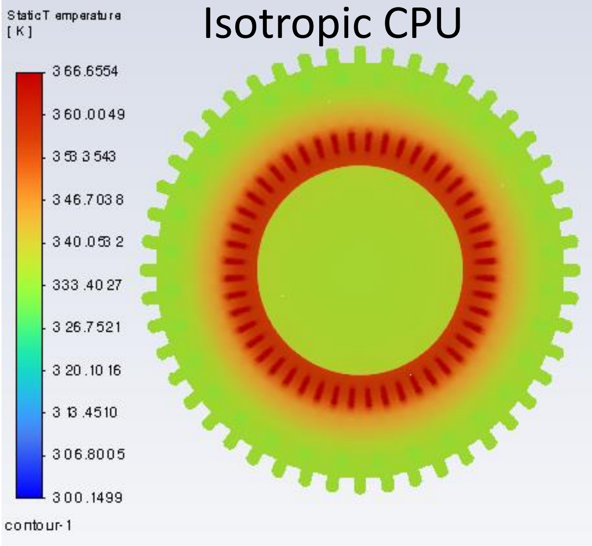

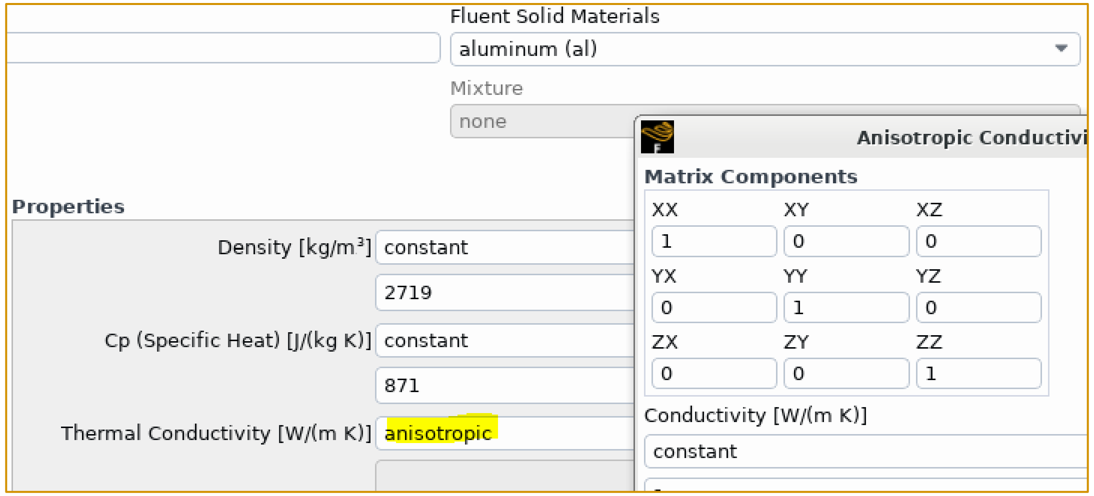

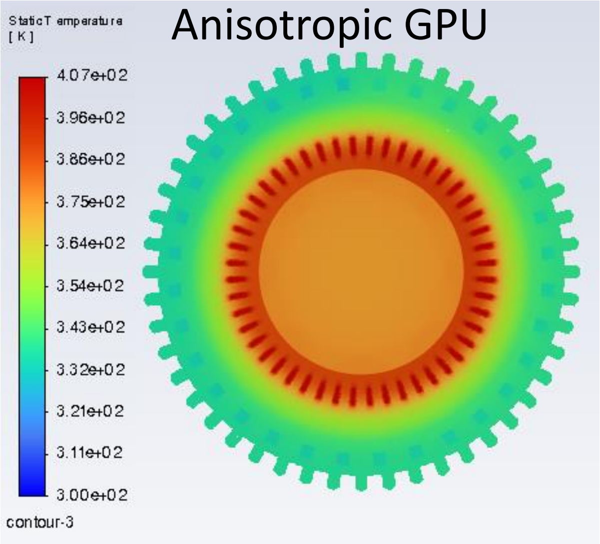

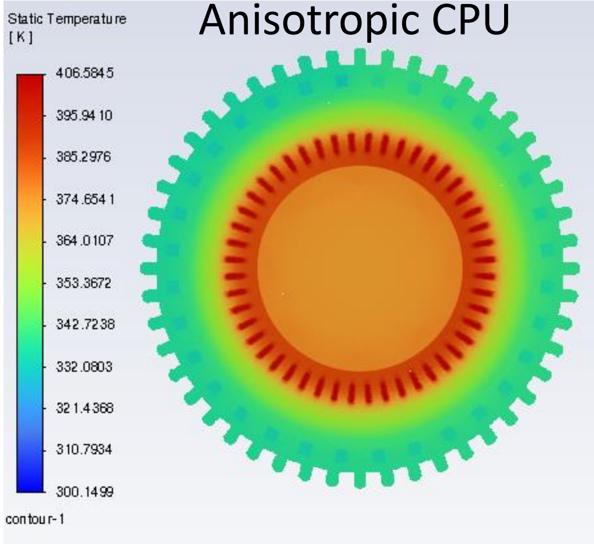

DVS export for post processing (beta)

- Export DVS files directly from GPU Solver
  - Transient and steady, GUI and TUI
  - Combination of cell and face zones and variables
  - Limitations: User-created surfaces not supported

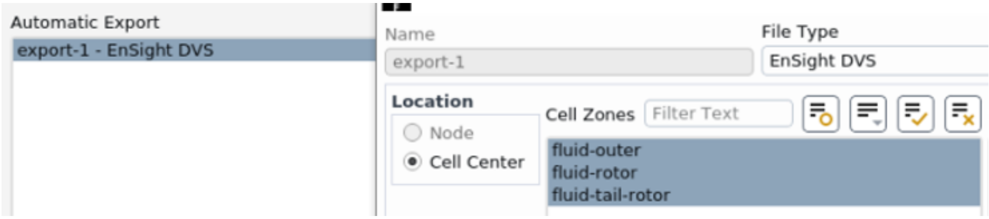

- **Test case**
  - Helicopter, 4.43 million, two sliding zones, with 2 GPUs, 200 steps, export every 10 steps
  - DVS file size on disk: 15 GB for 20 exports
  - 5 minutes w/o DVS export
  - 6 minutes with DVS export, ~20% more time than without

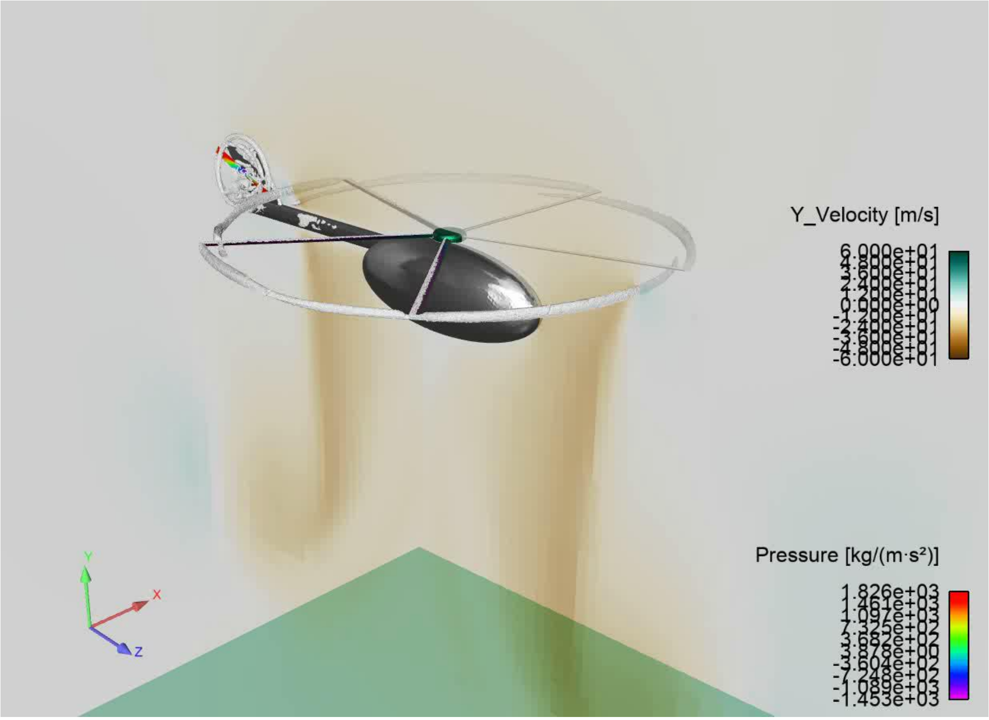

Demo: Full Wheel Axial Fan

Kurt Svihla

Axial Fan Geometry and Mesh

- Based on Axial Fan Workshop in Fluent Rotating Machinery Course
- Geometry: 30 rotor blades, 20 stator blades (full 360° modeled)
- Mesh: hex with inflation layers, non-conformal interfaces

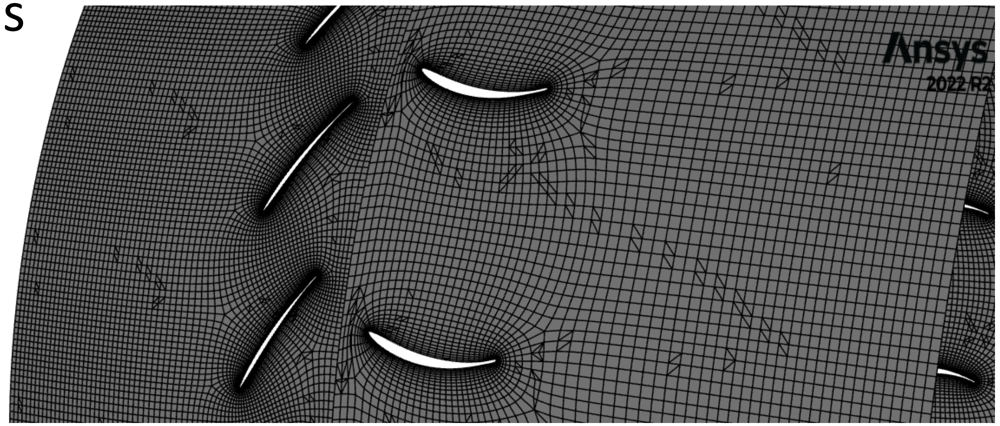

Mesh type	Mesh Count	Min Orth. Qual.	Max Aspect
Hex	2,529,960	0.064	553

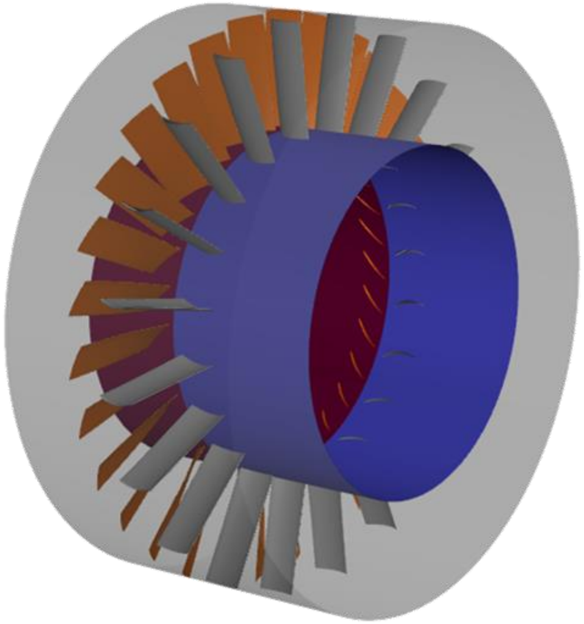

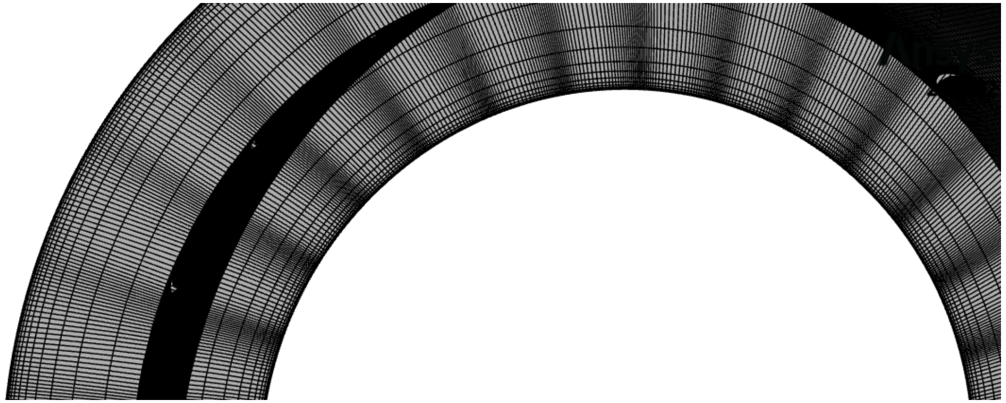

Hardware

- Dell Laptop CPU: 11th Gen Intel(R) Core(TM) i7-11850H @ 2.50GHz, 64 GB RAM, 8 cores
- Dell Laptop GPU: RTX A5000 Laptop GPU, V8.6, 16 GB, 366.272 GB/s, 48x128 cores   48 SM's

Starting the Launcher

- Enterprise license level should be set in order to select the native GPU solver
- Problems are read in and set up on CPU's, when the solver is launched it switches to the GPU

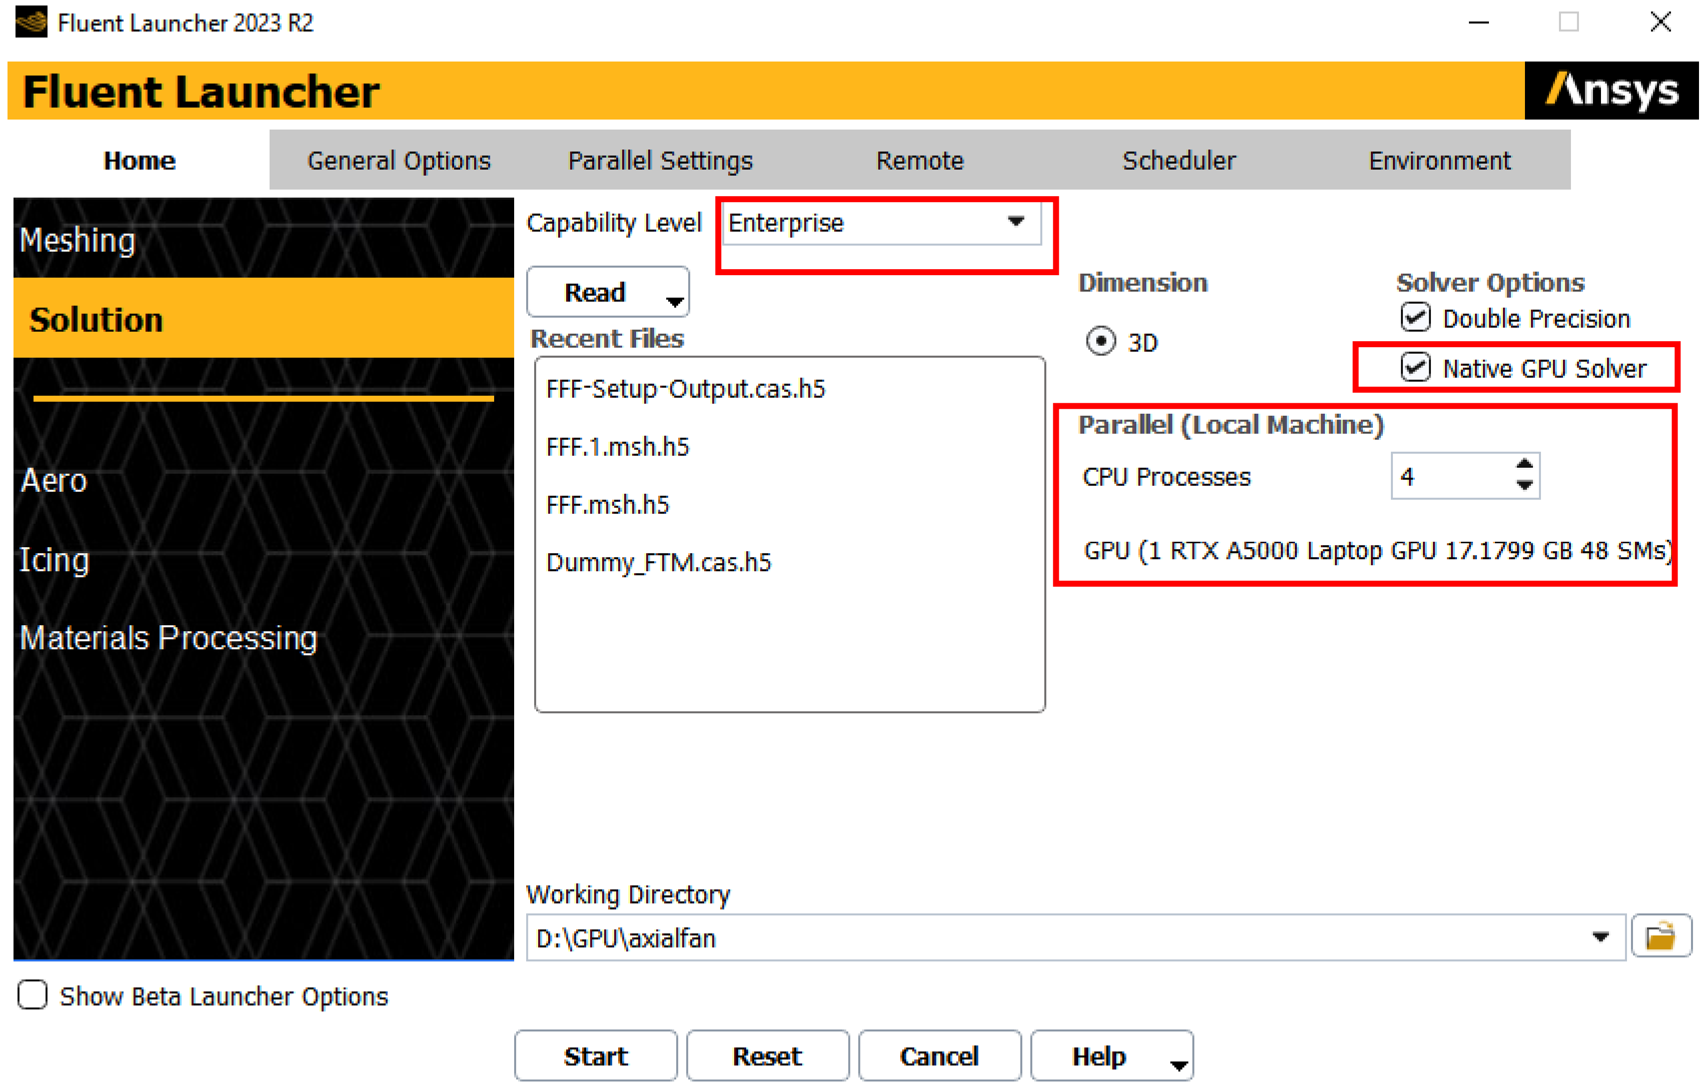

Switching to the GPU Solver

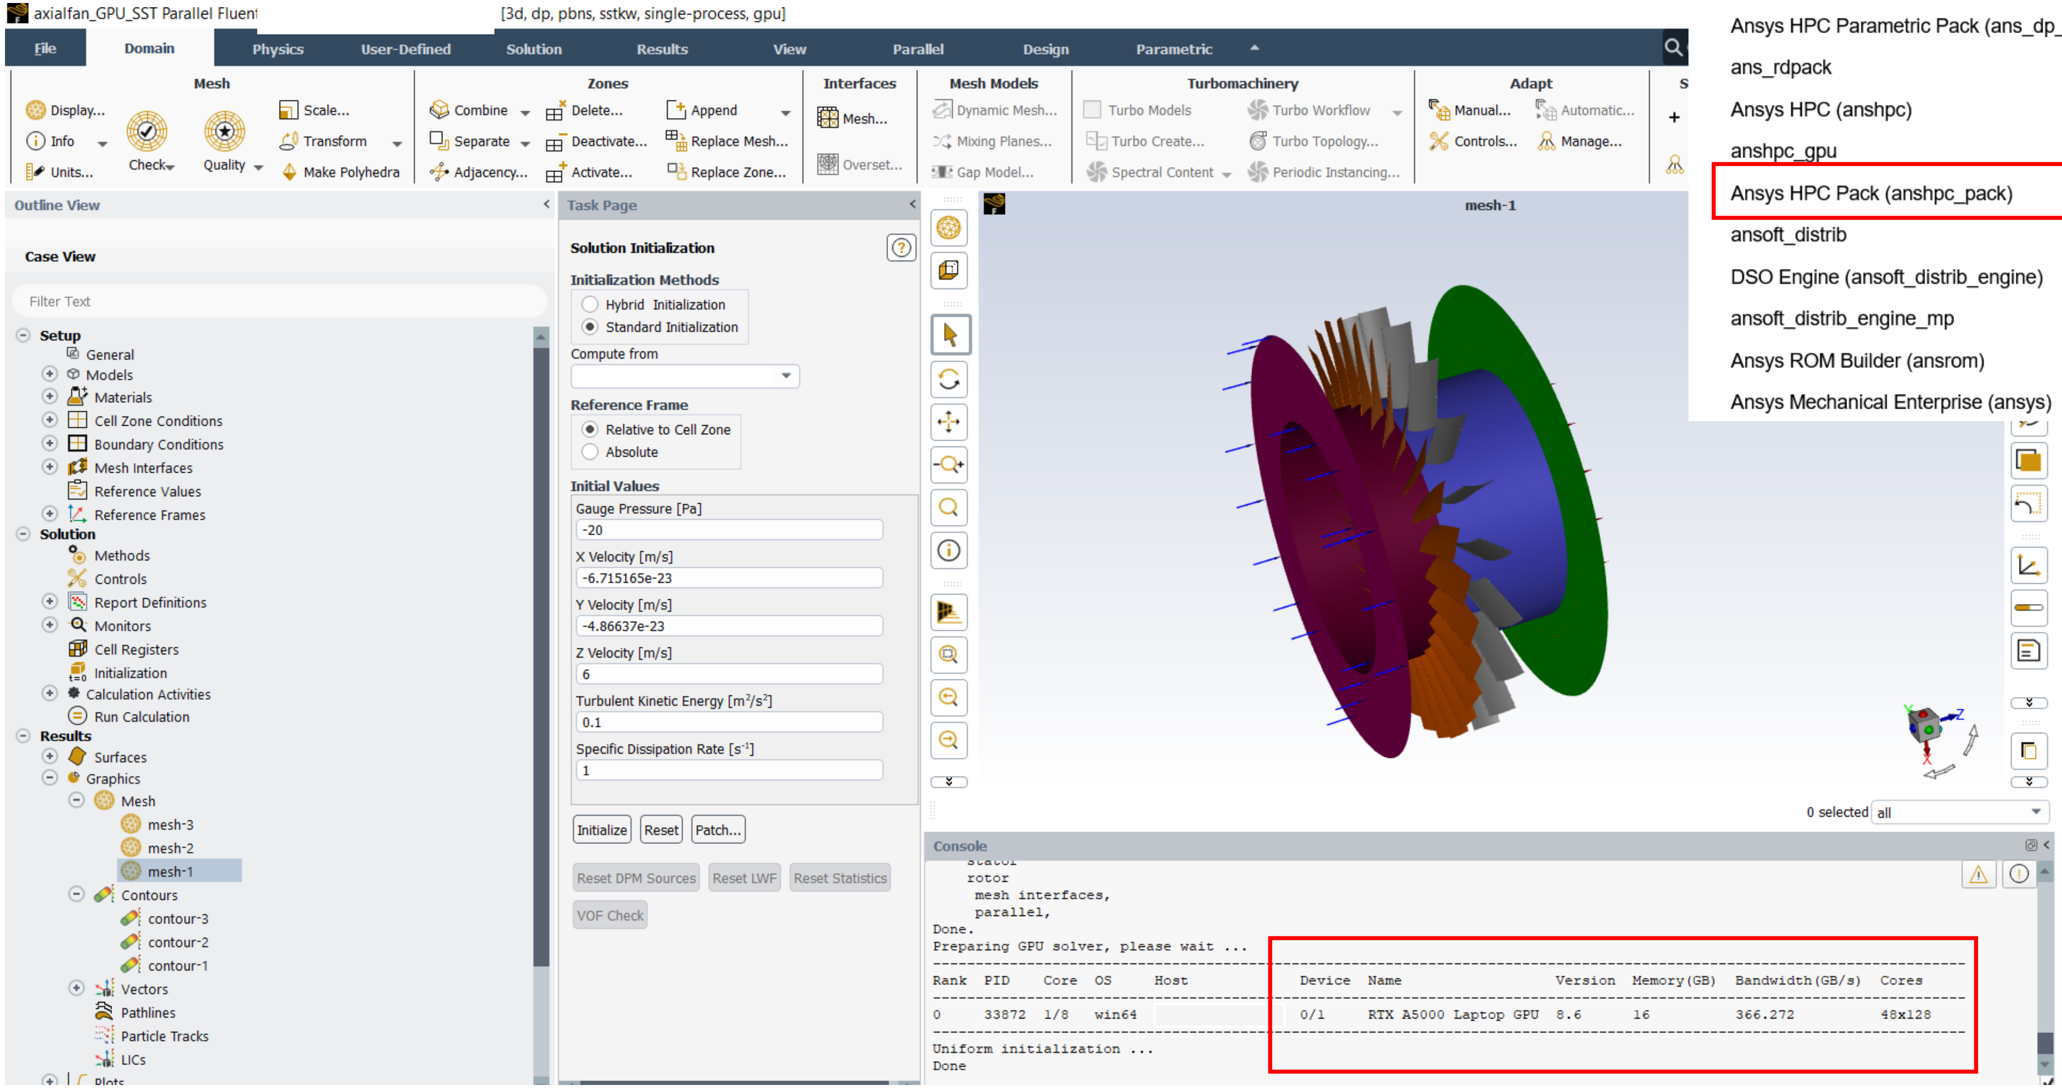

Current License Usage

Ansys HPC Parametric Pack (ans_dp_pack)	0/10
ans_rdpack	0/10
Ansys HPC (anshpc)	0/80
anshpc_gpu	0/10
Ansys HPC Pack (anshpc_pack)	1/80
ansoft_distrib	0/10
DSO Engine (ansoft_distrib_engine)	0/10
ansoft_distrib_engine_mp	0/10
Ansys ROM Builder (ansrom)	0/10
Ansys Mechanical Enterprise (ansys)	0/10

48 SM GPU Card = 1 HPC Pack

Laptop Resource Usage

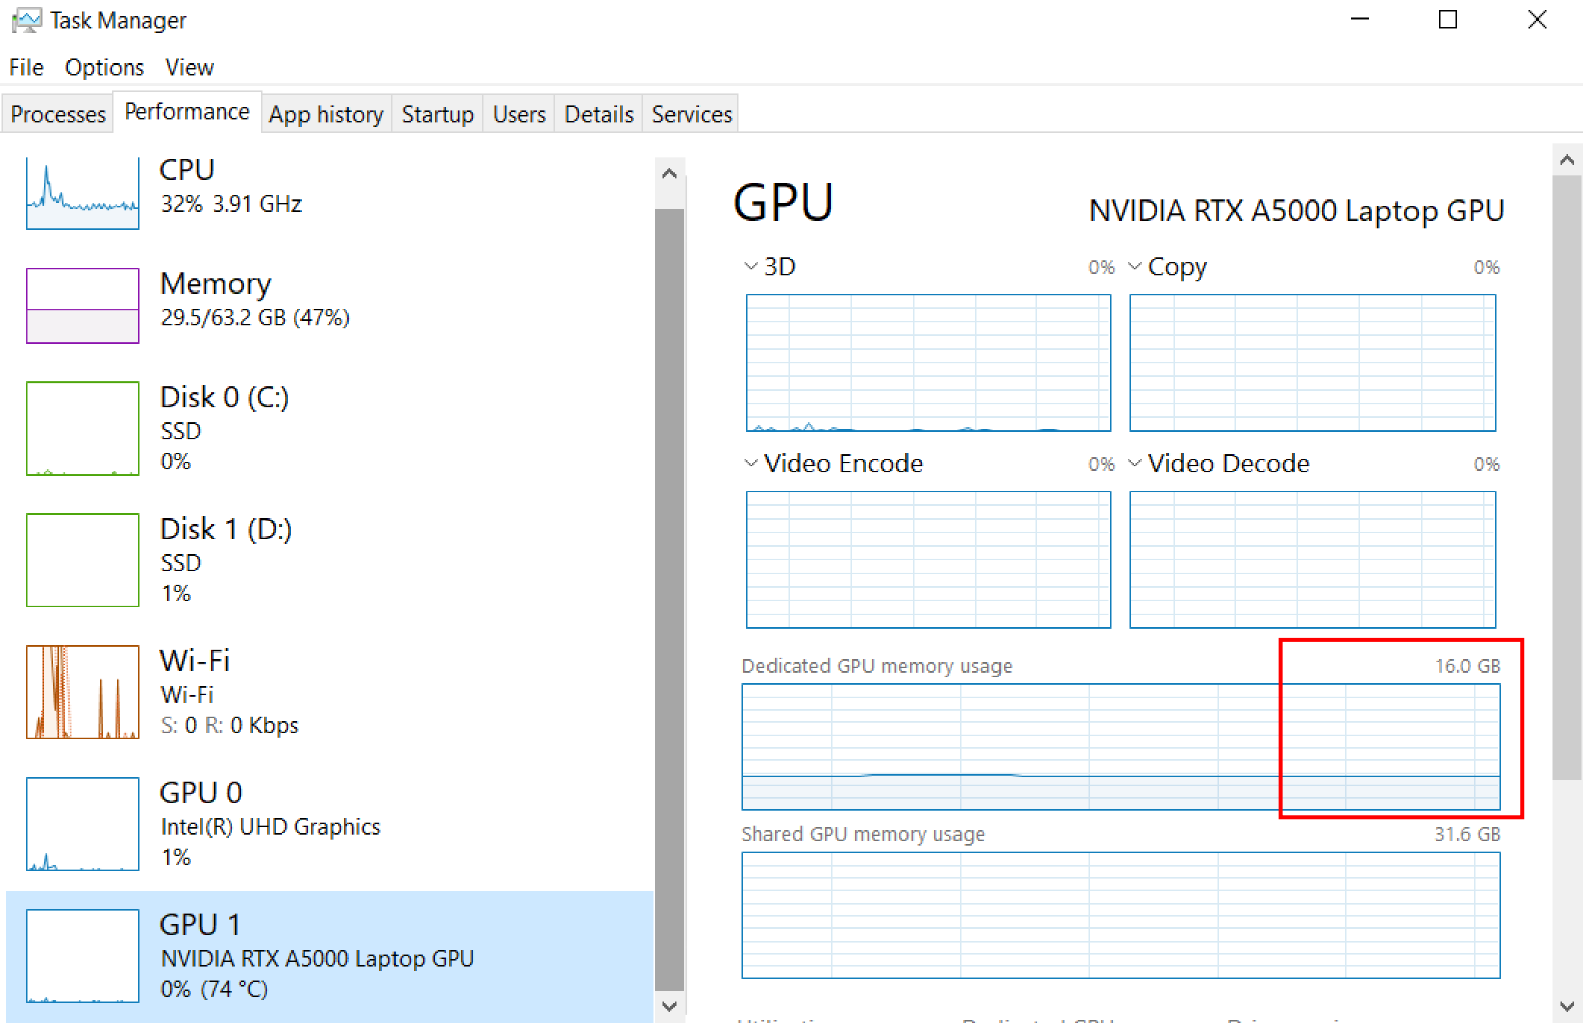

2.5 M Cell Hex Mesh Case Requires 5 GB GPU Memory in Double Precision

Axial Fan (Mass Flow Inlet, Pressure Outlet)

ACE: Kurt Svihla

- **Geometry:**
  - Full 360° with rotor/stator cell zones
- **Mesh:**
  - Hex mesh (Turbogrid)
  - Prism layers, Max Aspect Ratio = 553
  - Cells: 2,529,560
- **Modeling:**
  - Standard and realizable k-ε, k-ω SST, GEKO
  - Constant density air
  - Nonconformal MRF for Frame Change
  - 2880 rpm
- **Numerics:**
  - Steady
  - Simple (P-V coupling)
  - Flow 2nd order upwind, turbulent spatial discretization: 1^{s} and 2nd order upwind
  - Standard Initialization
- **Tested features:**
  - CPU vs GPU predictions
  - Solver robustness
  - Speed

SST, CPU

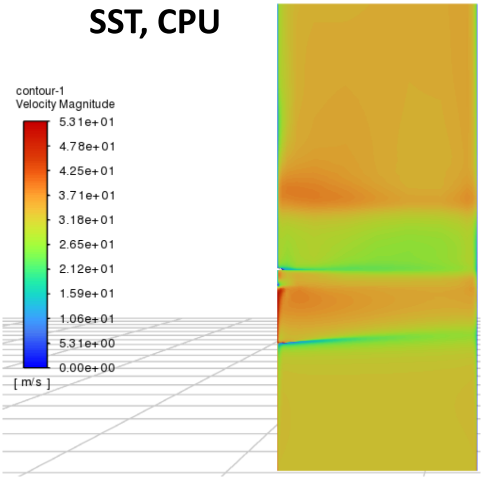

SST, GPU

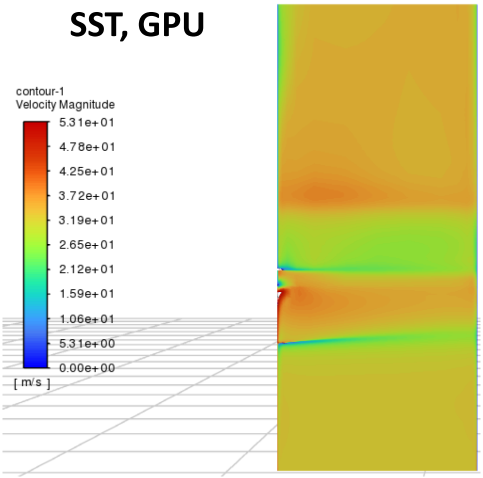

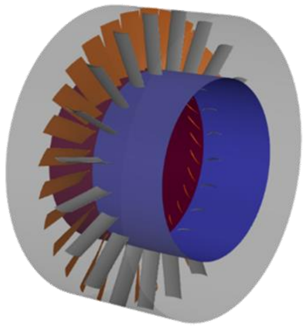

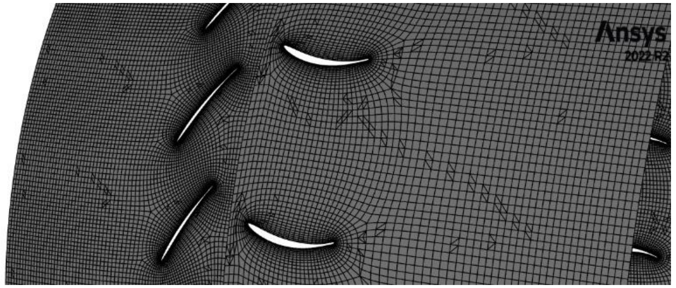

SST, GPU

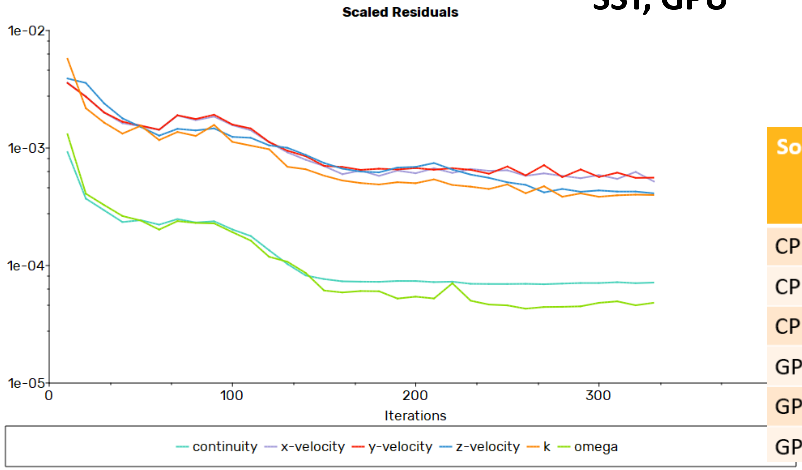

Conclusions:

- Overall results and solver performance satisfactory
- Significant speedup with GPUs

Solver	Cores	Turbulence Model	P-V Coupling	Fan Torque (Nm)	ΔP (Pa)	Iterations, Time to Converge
CPU	4	k-ω SST	SIMPLE	-20.01	1042.9	332, 975 s
CPU	4	k-ε Realizable	SIMPLE	-20.49	1075.8	332, 764 s
CPU	8	k-ε Standard	SIMPLE	-20.77	1112.8	332, 641 s
GPU	1	k-ω SST	SIMPLE	-19.77	1035.2	332, 94.8 s
GPU	1	k-ε Realizable	SIMPLE	-21.09	1131.0	332, 79.8 s
GPU	1	k-ε Standard	SIMPLE	-19.73	1052.0	332, 79.0s

Summary

Unique tool for Ansys CFD

- The first of its kind in the marketplace
  - Most other GPU CFD codes are less generalizable (e.g., LBM and SPH codes)
  - Other Navier-Stokes GPU codes do offloading/acceleration

Exposed in Fluent

- Available with CFD Enterprise license
- Runs out-the-box without HPC licenses*

HPC compatible

- Compatible with Workgroups and Packs
- Multi-GPU Cloud configurations available

Ansys **CFD Enterprise**

No compromise on accuracy
End-to-end workflow
Advanced physics
GPU speed & cost

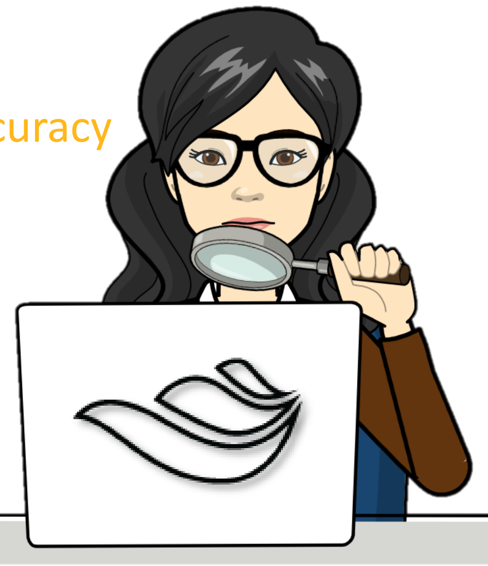

**Depending on the type of card – see Solver Licensing section*

Stay Connected

Educational Resources

- Blog series
  - **Blog 1**: Introducing GPUs for CFD
  - **Blog 2**: Laminar and Turbulent flows
  - **Blog 3**: Small models with CHT and porous media

- White paper focused on speed and accuracy

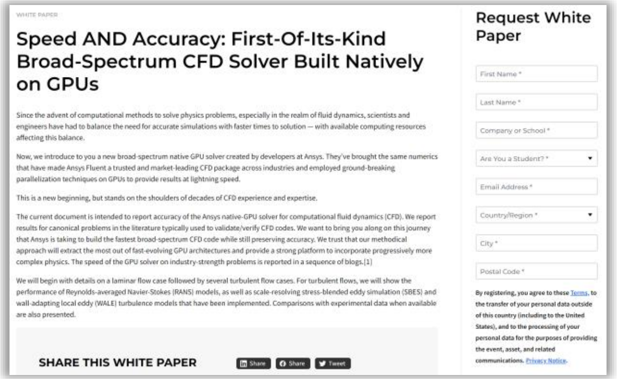

WHITE PAPER

Speed AND Accuracy: First-Of-Its-Kind Broad-Spectrum CFD Solver Built Natively on GPUs

Since the advent of computational methods to solve physics problems, especially in the realm of fluid dynamics, scientists and engineers have had to balance the need for accurate simulations with faster times to solution — with available computing resources affecting this balance.

Now, we introduce to you a new broad-spectrum native GPU solver created by developers at Ansys. They've brought the same numerics that have made Ansys Fluent a trusted and market-leading CFD package across industries and employed ground-breaking parallelization techniques on GPUs to provide results at lightning speed.

This is a new beginning, but stands on the shoulders of decades of CFD experience and expertise.

The current document is intended to report accuracy of the Ansys native-GPU solver for computational fluid dynamics (CFD). We report results for canonical problems in the literature typically used to validate/verify CFD codes. We want to bring you along on this journey that Ansys is taking to build the fastest broad-spectrum CFD code while still preserving accuracy. We trust that our methodical approach will extract the most out of fast-evolving GPU architectures and provide a strong platform to incorporate progressively more complex physics. The speed of the GPU solver on industry-strength problems is reported in a sequence of blogs.[1]

We will begin with details on a laminar flow case followed by several turbulent flow cases. For turbulent flows, we will show the performance of Reynolds-averaged Navier-Stokes (RANS) models, as well as scale-resolving stress-blended eddy simulation (SBES) and wall-adapting local eddy (WALE) turbulence models that have been implemented. Comparisons with experimental data when available are also presented.

SHARE THIS WHITE PAPER Share Share Tweet

Request White Paper

First Name *

Last Name *

Company or School *

Are You a Student? *

Email Address *

Country/Region *

City *

Postal Code *

By registering, you agree to these Terms, to the transfer of your personal data outside of this country (including to the United States), and to the processing of your personal data for the purposes of providing the event, asset, and related communications. Privacy Notice.

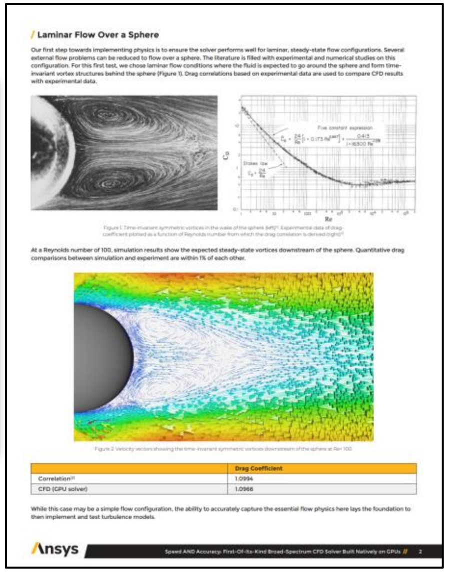

/ Laminar Flow Over a Sphere

Our first step towards implementing physics is to ensure the solver performs well for laminar, steady-state flow configurations. Several external flow problems can be reduced to flow over a sphere. The literature is filled with experimental and numerical studies on this configuration. For this first test, we chose laminar flow conditions where the fluid is expected to go around the sphere and form time-invariant vortex structures behind the sphere (Figure 1). Drag correlations based on experimental data are used to compare CFD results with experimental data.

At a Reynolds number of 100, simulation results show the expected steady-state vortices downstream of the sphere. Quantitative drag comparisons between simulation and experiment are within 1% of each other.

	Drag Coefficient
Correlation	1.0994
CFD (GPU solver)	1.0966

While this case may be a simple flow configuration, the ability to accurately capture the essential flow physics here lays the foundation to then implement and test turbulence models.

Ansys Speed AND Accuracy: First-Of-Its-Kind Broad-Spectrum CFD Solver Built Natively on GPUs // 2

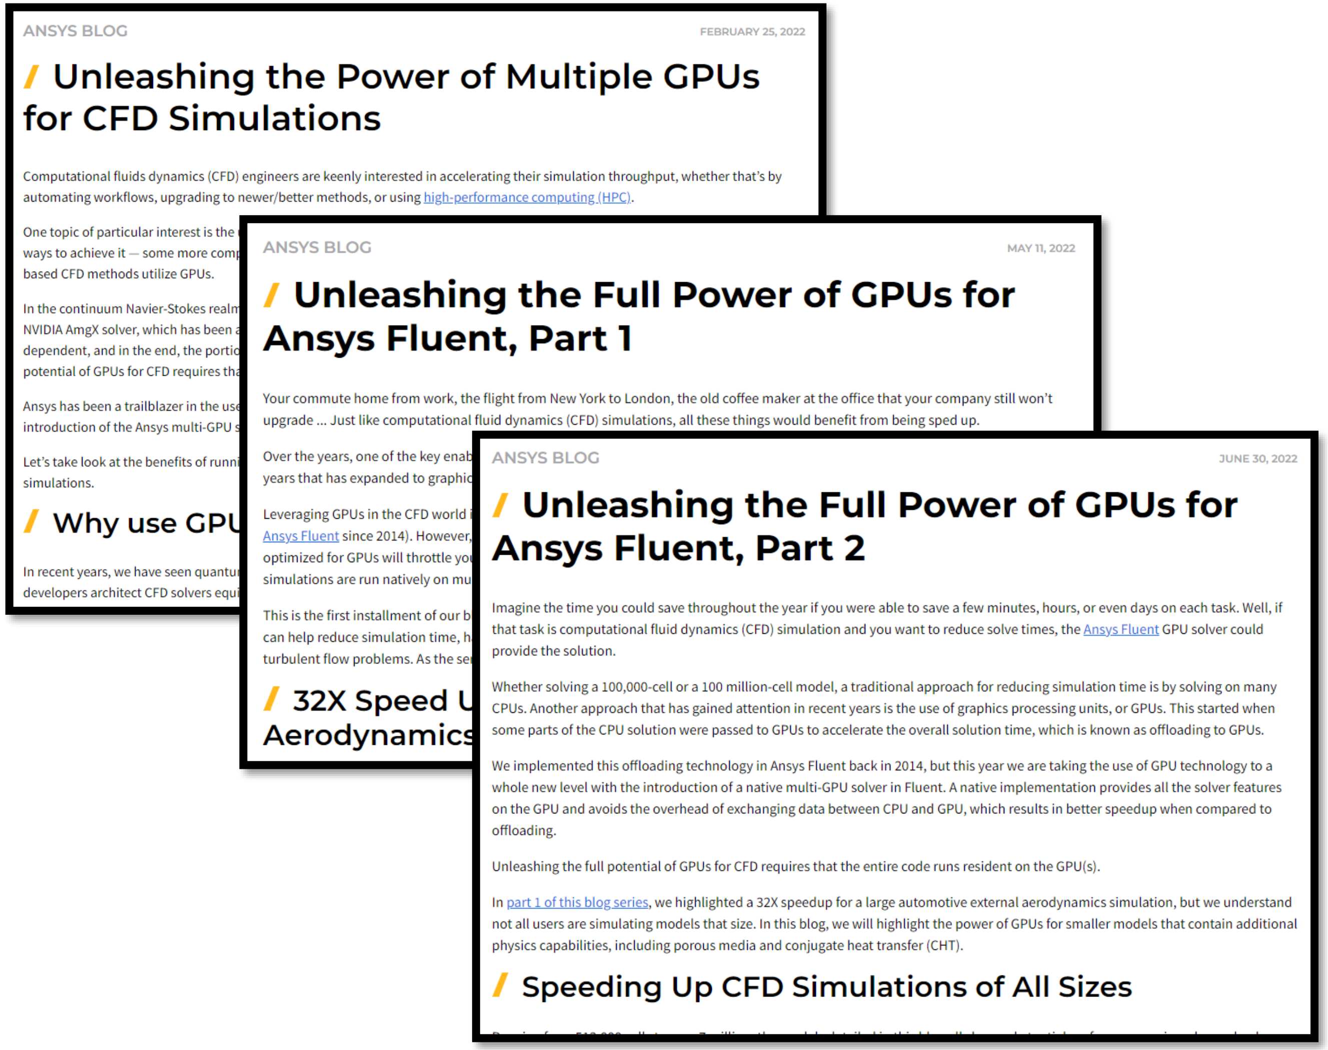

ANSYS BLOG FEBRUARY 25, 2022

/ Unleashing the Power of Multiple GPUs for CFD Simulations

Computational fluids dynamics (CFD) engineers are keenly interested in accelerating their simulation throughput, whether that's by automating workflows, upgrading to newer/better methods, or using high-performance computing (HPC).

One topic of particular interest is the ways to achieve it — some more comp based CFD methods utilize GPUs.

In the continuum Navier-Stokes realm NVIDIA AmgX solver, which has been a dependent, and in the end, the portio potential of GPUs for CFD requires tha

Ansys has been a trailblazer in the use introduction of the Ansys multi-GPU s

Let's take look at the benefits of runni simulations.

/ Why use GPU

In recent years, we have seen quantu developers architect CFD solvers equi

ANSYS BLOG MAY 11, 2022

/ Unleashing the Full Power of GPUs for Ansys Fluent, Part 1

Your commute home from work, the flight from New York to London, the old coffee maker at the office that your company still won't upgrade … Just like computational fluid dynamics (CFD) simulations, all these things would benefit from being sped up.

Over the years, one of the key enab years that has expanded to graphic

Leveraging GPUs in the CFD world i Ansys Fluent since 2014). However, optimized for GPUs will throttle yo simulations are run natively on mu

This is the first installment of our b can help reduce simulation time, h turbulent flow problems. As the se

/ 32X Speed U Aerodynamics

ANSYS BLOG JUNE 30, 2022

/ Unleashing the Full Power of GPUs for Ansys Fluent, Part 2

Imagine the time you could save throughout the year if you were able to save a few minutes, hours, or even days on each task. Well, if that task is computational fluid dynamics (CFD) simulation and you want to reduce solve times, the Ansys Fluent GPU solver could provide the solution.

Whether solving a 100,000-cell or a 100 million-cell model, a traditional approach for reducing simulation time is by solving on many CPUs. Another approach that has gained attention in recent years is the use of graphics processing units, or GPUs. This started when some parts of the CPU solution were passed to GPUs to accelerate the overall solution time, which is known as offloading to GPUs.

We implemented this offloading technology in Ansys Fluent back in 2014, but this year we are taking the use of GPU technology to a whole new level with the introduction of a native multi-GPU solver in Fluent. A native implementation provides all the solver features on the GPU and avoids the overhead of exchanging data between CPU and GPU, which results in better speedup when compared to offloading.

Unleashing the full potential of GPUs for CFD requires that the entire code runs resident on the GPU(s).

In part 1 of this blog series, we highlighted a 32X speedup for a large automotive external aerodynamics simulation, but we understand not all users are simulating models that size. In this blog, we will highlight the power of GPUs for smaller models that contain additional physics capabilities, including porous media and conjugate heat transfer (CHT).

/ Speeding Up CFD Simulations of All Sizes

Stay Connected

Ansys Fluids
Ansys Fluids products are for engineers who need to make better, faster decisions to get the accurate results they need.
Software Development · Canonsburg, PA · 44,240 followers

Rapidly spreading the news via LinkedIn

Ansys Fluids
44,241 followers
7mo • Edited •

The Ansys Fluent **Multi-GPU Solver** is seriously FAST ⚡ ! Here's a sneak peek of a ~6 million-cell **hybrid RANS-LES** case with the upcoming **Sliding Mesh** support 👀. With >10**X** speed up over CPUs, the GPU solver delivers results in hours instead of days. Imagine the engineering insights from an entire Design of Experiments (DOE) with the same lead time of just a single design point! Learn more: https://ansys.me/3OwOWIE

NVIDIA Design and Visualization #nvidiadesign
#simulation #GPU #GPUs #Fan #innovation #NVIDIA #ansys #sustainability

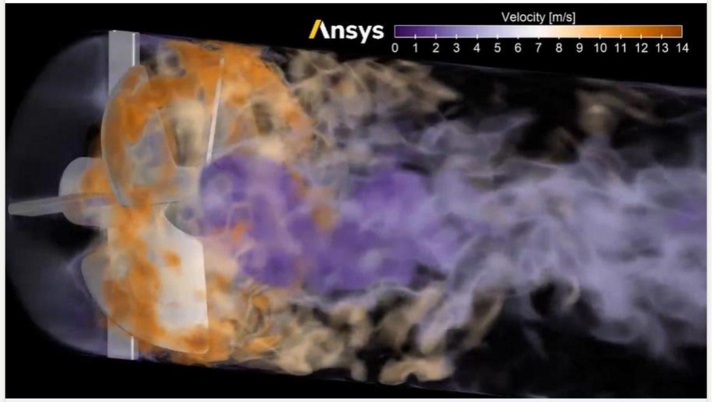

Sliding Mesh in Ansys Fluent Multi-GPU Solver

1,044 · 18 comments · 78 reposts

Ansys Fluids
44,240 followers
7mo • Edited •

Simulation truly affects every facet of our lives these days. Here's another sneak peek 👀 of the ground-breaking Ansys Fluent **Multi-GPU Solver** technology in the context of **urban-air mobility**. Scale-resolving simulations of rotating components on GPUs are order-of-magnitude faster and yet as accurate (same expertise building out the numerics) as when carried out on CPUs! Learn more : https://ansys.me/3irhQOb

NVIDIA Design and Visualization
#simulation #GPU #GPUs #drone #innovation #ansys #sustainability #nvidiadesign #NVIDIA #aerospace #UrbanMobility

Drone Simulation with Ansys Fluent Multi-GPU Solver

977 · 7 comments · 72 reposts

Ansys Fluids
44,240 followers
6mo • Edited •

Here are 6 reasons to learn more about the broad-spectrum native multi-GPU solver, created by developers at Ansys. They've brought the same numerics that have made Ansys Fluent a trusted and market leading CFD software across industries and employed ground-breaking parallelization techniques on GPUs to provide results at lightning speed. This is a new beginning but stands on the shoulders of decades of CFD experience and expertise. Learn more:
https://lnkd.in/e6597ugu

#GPU #NVIDIAdesign #innovation #CFD #technology #fluids #gpus #nvidia

135 · 1 comment · 17 reposts

Ansys