

GPU CFD Acceleration for Turbomachinery Design and Analysis

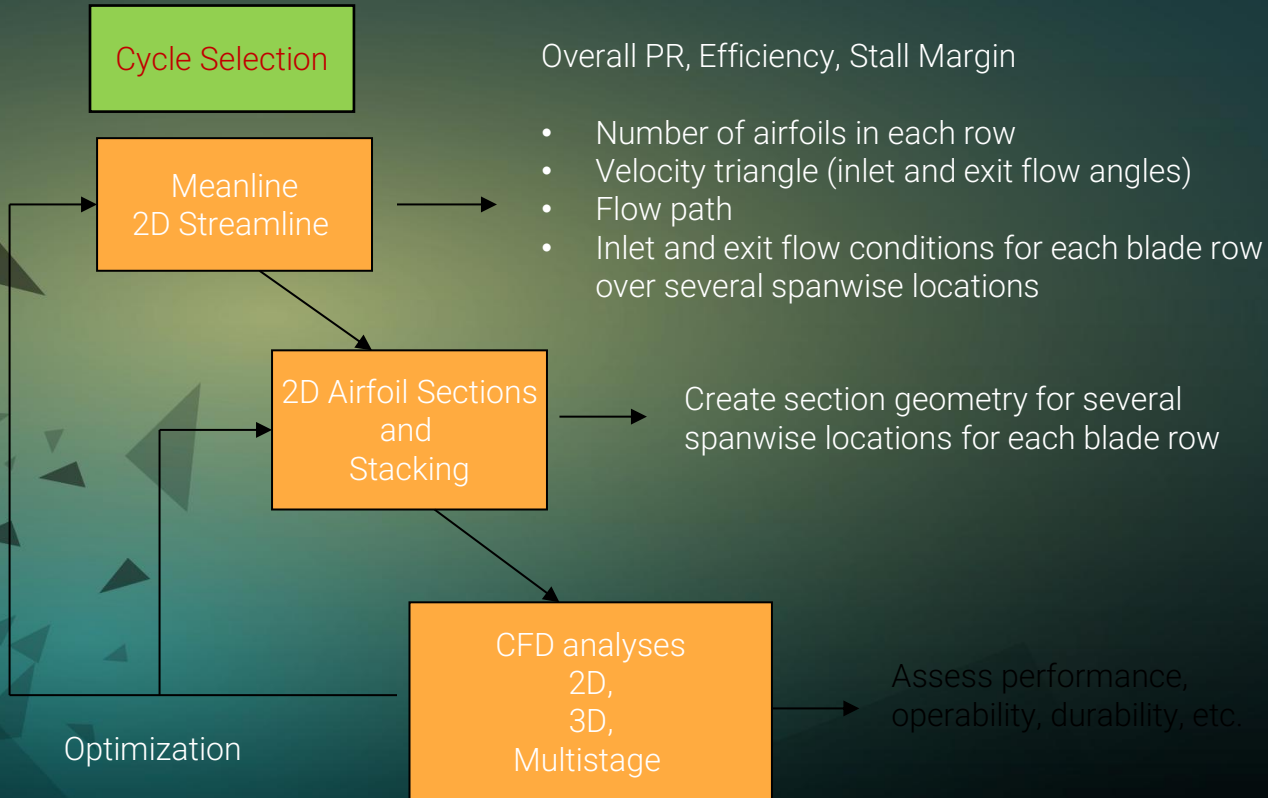
Dr. Hsin-Hua Tsuei
Vice President
ADSCFD

Acknowledgement


Janakiraman Thiyagarajan
for making this presentation possible

All my colleagues at ADSCFD
for contributing to this presentation

Turbomachinery Aerodynamic Design / Analysis



1D Meanline Analysis




- Meanline analysis
 - Provide machine performance estimate and aerodynamic stability of the turbomachinery – blade count, RPM, mass flow rate, efficiency, etc.
 - Preliminary design tool for many design and what-if scenarios
 - Solve 1D mean flow equations with losses based on correlations
 - Hundreds to thousands of runs during preliminary design stage to identify meanline geometry and flow conditions quickly
- Limitations
 - 1D, limited information for 3D geometry
 - Requires massive correlation database
 - Relies on loss models
 - Not reliable for extrapolation

Loss Correlations in Meanline Analyses




- Mach No.
- Reynolds No.
- Airfoil Loading (includes solidity)
- Aspect ratio
- Airfoil max. thickness/chord
- Airfoil leading edge radius
- Airfoil surface roughness
- Shroud roughness
- Rotor OD clearance
- Stator ID clearance
- Axial gap
- Shock wave
- ...



2D streamline and blade to blade analysis

- Multiple streamlines and blade to blade cross sections from hub to tip form a quasi-3D analysis in a design cycle
 - Use two 2D methods to make up a quasi-3D
 - Allows quick inspection of meridional flow path curvature (S1) and blade cross-section (S2) designs
- Streamlines on meridional plane
 - Inviscid analysis with loss models to account for potential loss not modeled
- Blade to blade sections (from hub to tip) on surface of revolution including “true” streamtube thickness for each section
 - Full Navier-Stokes with turbulence model to identify flow separation, pressure gradient, etc.


Streamline Analysis



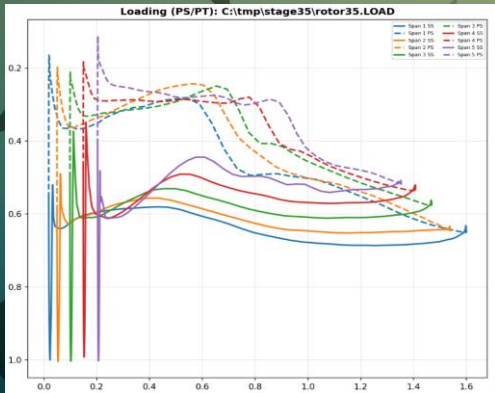

- Streamline analysis
 - Extend 1D meanline to spanwise distributions of Mach, PR, $P_{t_{loss}}$...
 - Solution to 2D Steamline Curvature Equation in meridional plane with loss correlation and flow blockages due to endwalls and wakes
 - *Hundreds of runs per a few candidate meanline outputs*
- Limitations
 - 2D
 - Endwall loss model and clearance effects
 - Blockage, secondary flow

2D blade-to-blade viscous RANS CFD analysis

- Blade to blade RANS CFD analysis
 - 2D cross-section with finite streamtube thickness distribution from LE to TE
 - Blade 2D section meets or exceeds the design goal
 - Incidence effect
 - Camber line distribution effect
 - Thickness distribution effect
 - Streamtube contraction effect
 - Predict losses and deviation
 - Determine if the section will pass the designed flow
 - Fast turn around to optimize design
- Limitations
 - 2D, no spanwise gradient effect





Single Blade Row Full 3D Steady Analysis



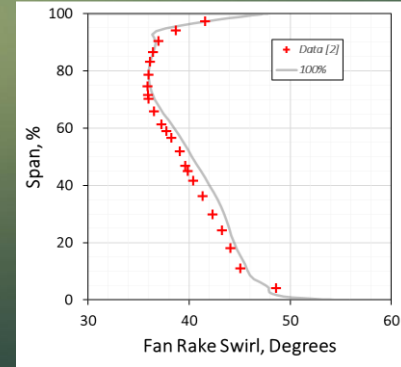
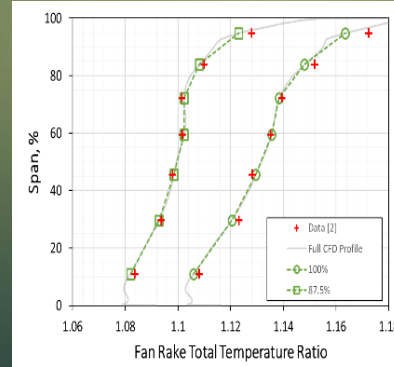
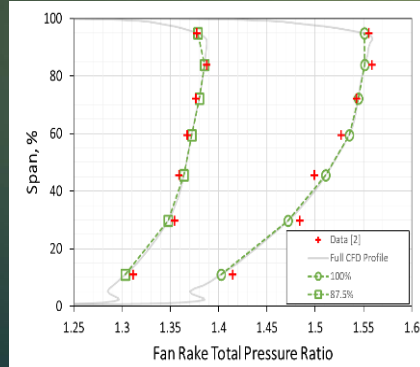
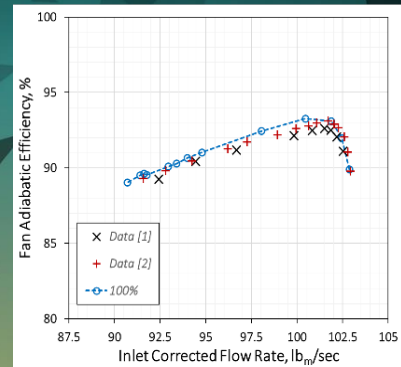

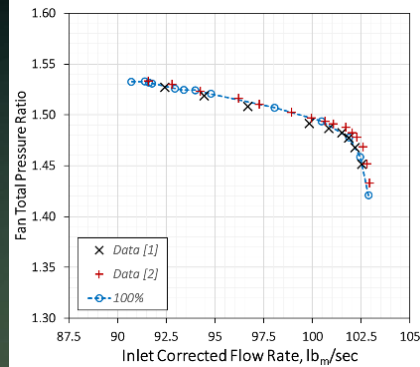

- Full 3D analysis
 - Single blade passage, requires periodic boundary condition
 - Optimize 3D blade geometry and the endwall contours
 - Endwall losses
 - Tip clearance size and numerical resolution
 - Flow range from choke to stall
 - Full span deviation
 - Secondary flow passages can be included in the analysis
 - ~ 1 min on a single Nvidia 4090 GPU laptop
- Limitation
 - Lack of multi blade row with rotor-stator interaction effect

Multiple Blade Row Full 3D Steady Analysis



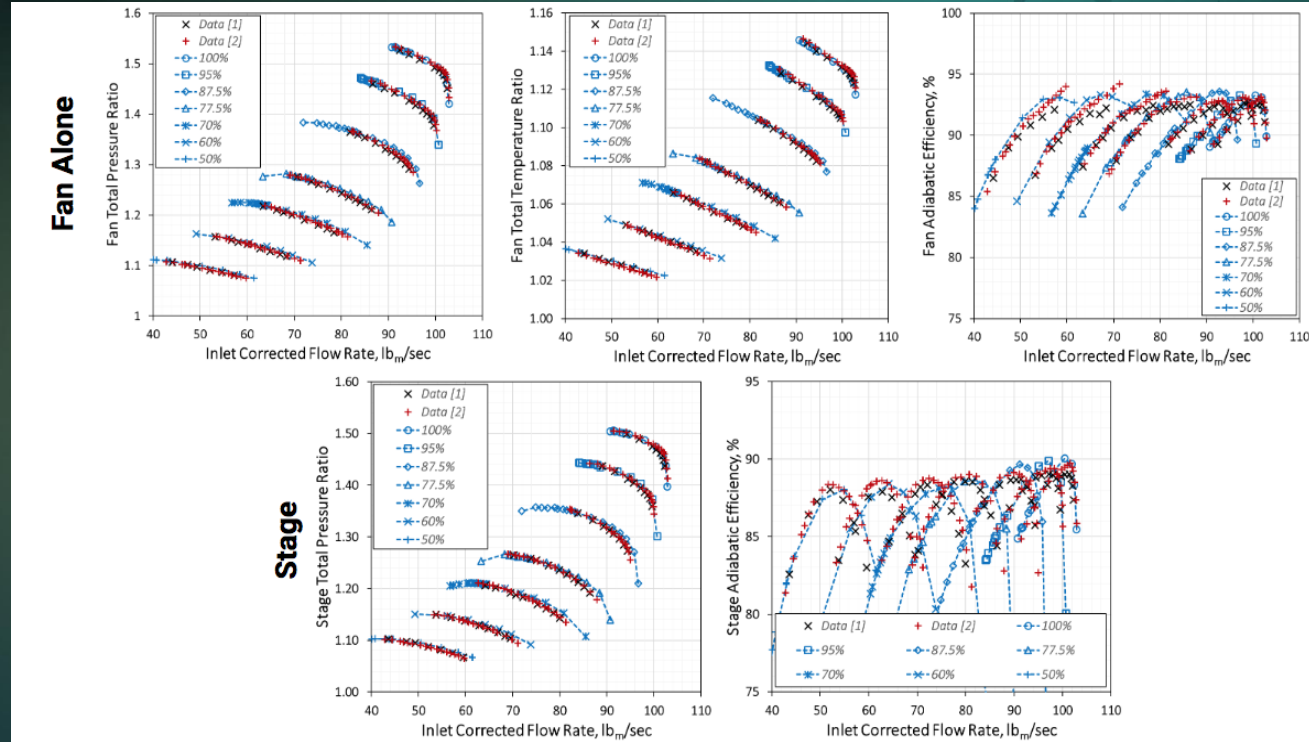
- Steady state multi-blade row analysis
 - Determine overall machine performance
 - Loss distribution per blade row
 - Steady-state rotor-stator interface modeling
 - Mixing plane
 - Frozen rotor
 - Multiple speedlines from choke to stall to create numerical turbomachinery map
 - Each speedline has at least 5 – 7 points, the map can be time-consuming
- Limitation
 - Unsteady 3D flow physics not resolved
 - Wake flow
 - Rotor-stator interaction

NASA SDT



NASA SDT – Compressor Map

- 119 Points @ 2 m / Point on a single Nvidia 4090 GPU laptop
- 3 hours 55 minutes to generate the whole map



GPU CFD for Turbomachinery Design and Analysis

- What's required for GPU CFD so far ?
 - Single blade passage analysis, requires periodic boundary condition
 - Steady-state rotor-stator interaction models for selected blade rows or stage analysis
 - Mixing plane or frozen rotor
 - Fast turnaround time to understand stage coupling effects
 - Based on 50 years of turbomachinery design and analysis experience
 - GPU CFD for turbomachinery must include
 - Existing turbomachinery CFD methods and models
 - Provide speedup in all parts of turbomachinery CFD
 - Provide identical results between CPU and GPU codes
 - Retain past validated modeling methods and experience in turbomachinery CFD
 - Simply can not re-validate a GPU code with all the necessary turbomachinery CFD models

How about Turbulence Modeling ?


- Most of turbomachinery CFD integrated in a design cycle use steady RANS turbulence model, mostly engineering level two equation turbulence models
 - $k - \omega$, $k - \omega$ SST, etc.
- Steady RANS provides efficient turnaround time for a product development cycle when hundreds of design variations are required over a short period of time, some are integrated with an optimizer
- Unsteady RANS provides in depth 3D flow physics in an unsteady turbomachinery stage analysis
 - with **sliding mesh** for rotor-stator interface
 - Wake flow, tip clearance, flow separation, vortex rollup physics full captured

Large Eddy Simulation or not ?


- LES requires much larger models (full wheel for every blade row, or large mesh) and much longer turnaround time
 - Often weeks to months to run one LES
 - Not practical in a design cycle, largely for research purposes in turbomachinery
 - With GPU CFD, still requires about a day to run a single stage analysis using 8 or more GPUs, especially when periodic BC is not implemented
 - This is a costly proposition for turbomachinery design
 - Need to acquire sizable GPU workstations
 - How does LES compare with RANS or unsteady RANS ?
 - Long turnaround time prohibits more validated modeling approach, especially for multiple blade rows
 - If LES deviated from data and RANS, should we abolish 50+ years of experience in turbomachinery CFD modeling and adopt a new method that is not well validated for turbomachinery analysis yet ?
 - RANS remains the workhorse in turbomachinery design

Large Eddy Simulation

- GPU acceleration is the great enabler for Large Eddy Simulation Release 10 introduces multiple capabilities for improved LES analysis
- Ex. 250M element LES analysis completed in 28 hours on 8 Nvidia L40 GPUs
- Simulations can be reduced from several months to a little over a day




Skin Friction from a LES simulation of a HPT vane





Skin Friction from an LES simulation of a cooled turbine vane

LES – Cooled Turbine Vane

- LES for cooled turbine vane
conjugate heat transfer analysis
- 250m elements
 - fluid domain ~ 170m elements
 - solid domain ~ 80m elements
- 20,000 time steps
- 28 hrs on 8 Nvidia L40 GPUs
- 14 hrs on 16 Nvidia L40 GPUs



LES – Cooled Turbine Vane




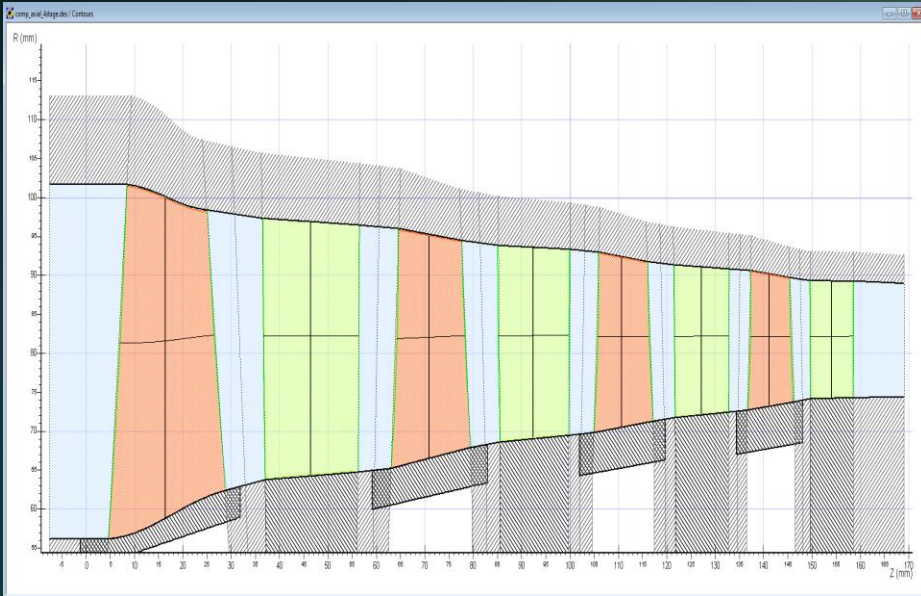
Unleash the Power of GPU CFD

Many time-consuming, large scale, and/or complex 3D flow physics analysis using steady/unsteady RANS to extend current modeling methodologies and enhance understanding of 3D unsteady flow physics

- Multi blade row unsteady analysis often includes many blade passages across several stages
- Asymmetric turbine vane
- Compressor casing treatment
- Cooled turbine conjugate heat transfer analysis
- Inlet distortion
- Airframe – engine coupling
- Rotating stall
- ...

4 Stage Axial Compressor


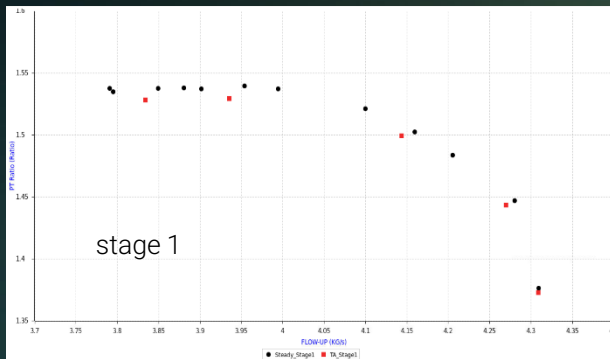

Concepts NREC Axcent example



4 Stage Axial Compressor

Steady state analysis on a single A4090 GPU laptop

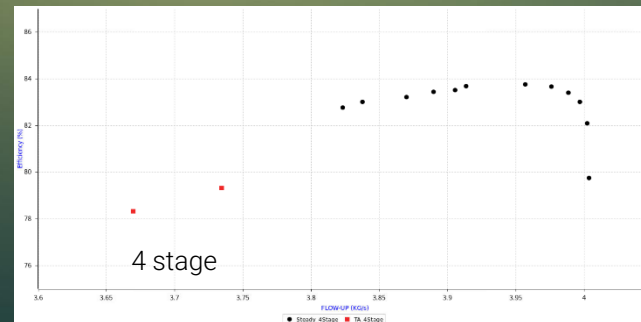

- stage 1 – 1m 2s
- 4 stage – 3m 30s



pressure ratio

- steady
- unsteady, time-averaged

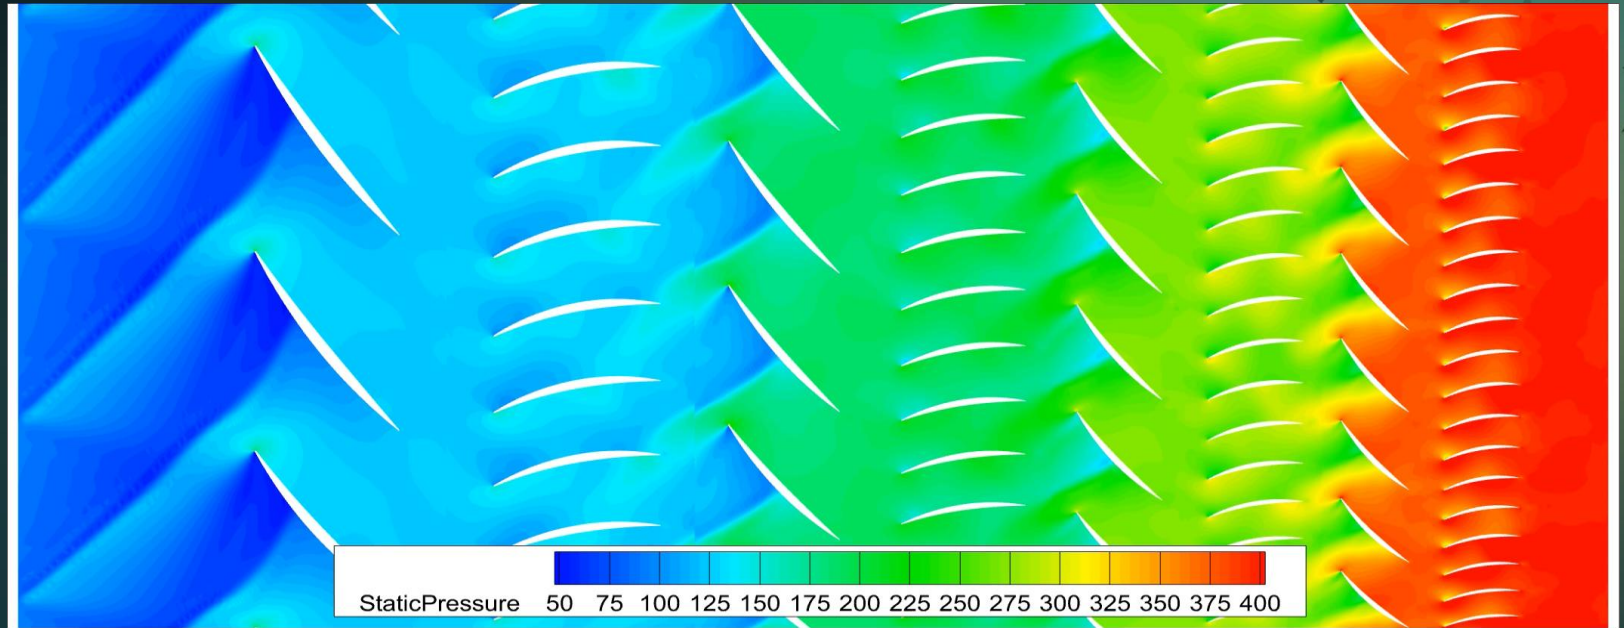
efficiency



Static Pressure @ 90% Span

Unsteady analysis on 8 Nvidia L40 GPUs


- stage 1 – 33m per rev
 - 4 stage – 3hr 23m per rev
- 3,300 time-steps per rev



Mach Number @ 90% Span

Unsteady analysis on 8 Nvidia L40 GPUs

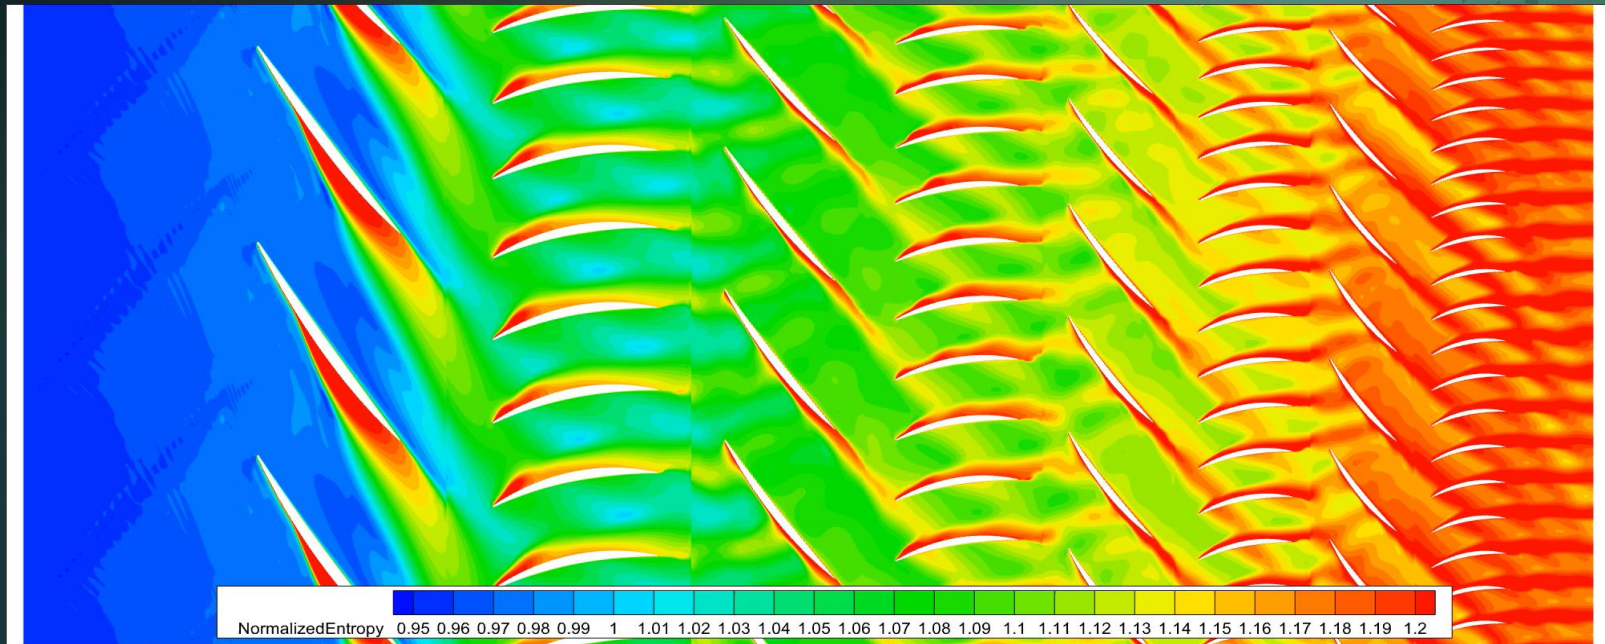
- stage 1 – 33m per rev
 - 4 stage – 3hr 23m per rev
- 3,300 time-steps per rev



Entropy @ 90% Span



Unsteady analysis on 8 Nvidia L40 GPUs

- stage 1 – 33m per rev
 - 4 stage – 3hr 23m per rev
- 3,300 time-steps per rev



Turbocharger

Concepts NREC Axcent example





Turbocharger Unsteady Analysis

Integrated impeller – volute analysis



- ~ 8M elements
- Steady mixing plane
7 min / 6,000 iterations on a single A4090 GPU laptop
- Unsteady sliding mesh
2400 timesteps/rev
20 min per revolution
On 5 Nvidia L40 GPUs

Entropy and Mach Number

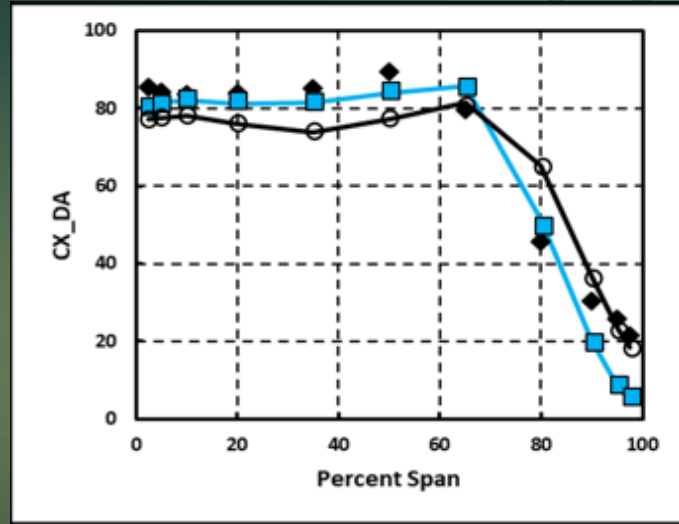



NASA LSR 4.5 Stage Compressor Analysis

- Investigate flow effects from a rotor with a large tip clearance at 4.4% of span (~10X normal clearance size)
- Current steady design methods have difficulty predicting the full performance characteristic
- Full annulus unsteady using dual time stepping
- 408 passages
- 122 million elements
- 2,600 timesteps / revolution



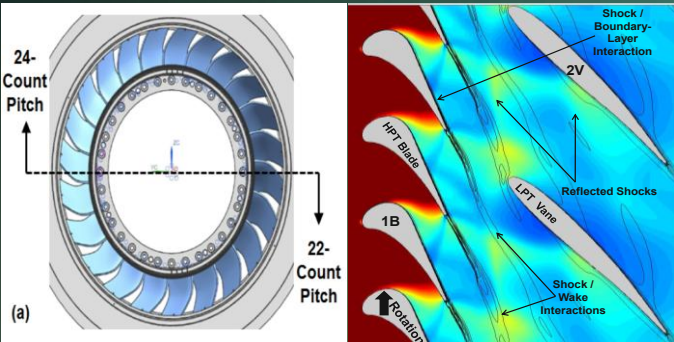
NASA LSR 4.5 Stage Compressor Analysis



	80 2.5 GHz Intel Xeon Cores	8 Nvidia L40 GPU	GPU to CPU performance
Time for 3 Revolutions	12 days	14 hrs 10 min	41X Speedup

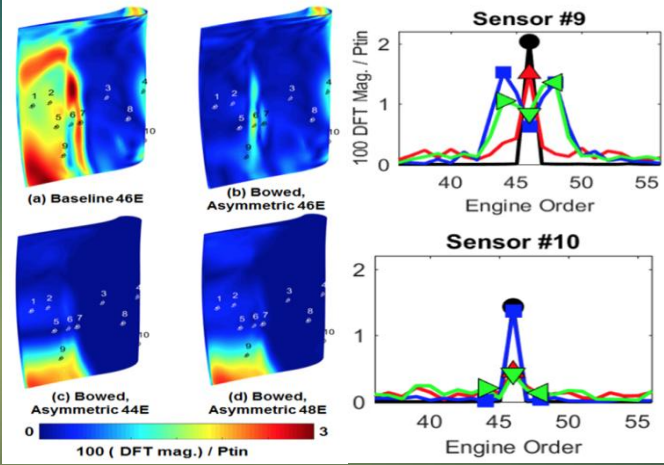
AFRL Asymmetric Turbine Vane

- Analysis can predict qualitatively and quantitatively the redistribution of unsteady signal power from the baseline vane to neighboring frequencies.
- With ADS GPU acceleration, a month long analysis on CPU clusters can be reduced to less than a day



Baseline Vane
 — Simulation
 — Experiment

Bowed, Asymmetric Vane
 — Simulation
 — Experiment

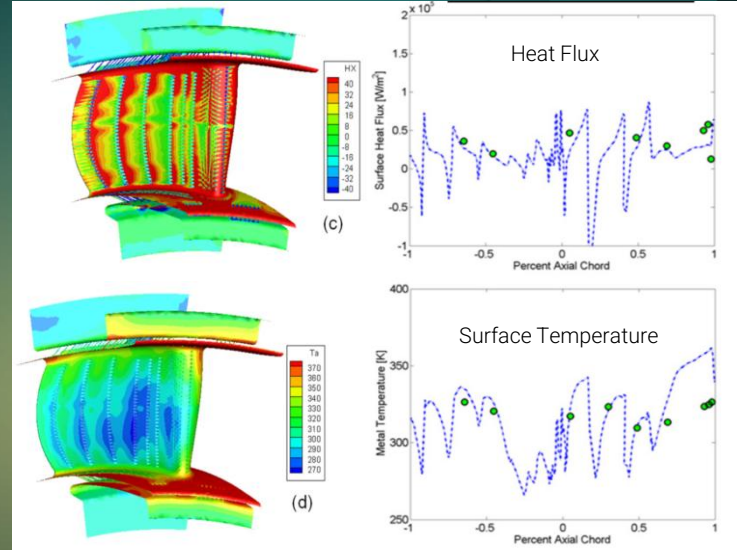



	92 2.2 GHz Intel Xeon Cores	8 Nvidia L40 GPUs	GPU to CPU comparison
Time for 3 Revolutions	~30 days	15 hrs 52 min	45.5X Speedup

Clark, J P, Anthony, R J, Ooten, M K, Finnegan, J M, Johnson, P D, Ni, R H, "Effects of Downstream Vane Bowing and Asymmetry on Unsteadiness in a Transonic Turbine" GT2018-76735

Cooled Turbine Vane (HIT 1V)

- Film cooled airfoils are incredibly complex; both to simulate and to mesh. The cooled 1V of the HITRT has over 650 cooling holes and 2 plenums.
- Unstructured meshes tend to be more memory intensive and less computationally efficient
- Steady analysis with conjugate heat transfer
- Includes a kapton layer
- 12M elements fluid domain
- 11M elements in the solid domain

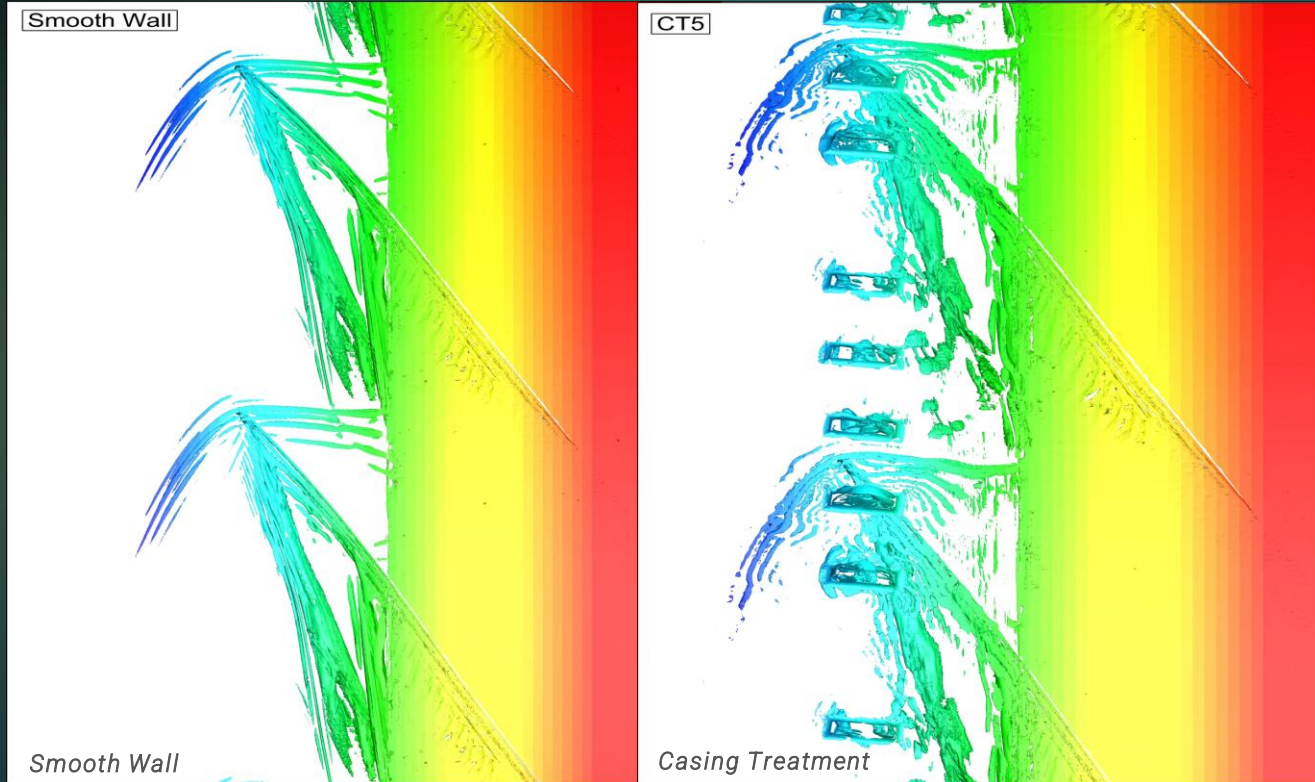


Ni, R N, Humber, W H, Fan, G, Clark, J P, Anthony, R J, Johnson, J J,, "Comparisons of Predictions from Conjugate Heat Transfer Analysis of a Film-cooled Turbine Vane to Experimental Data" GT2013-94716

	16 2.5 GHz Intel Xeon Cores	8 Nvidia L40 GPUs	GPU to CPU comparison
Time to Convergence	6 hrs 40 min	4 min 42 sec	84.5X Speedup

Novel Casing Treatments

- Mixed structured and unstructured mesh
- ~ 8m elements
- 1 hr. 52 min. per rev on 8 Nvidia L40 GPUs



Novel Casing Treatments

- Mixed structured and unstructured mesh
- ~ 8m elements
- 1 hr. 52 min. per rev on 8 Nvidia L40 GPUs



Rotating Stall

ECL5/Catana Geometry



Number of blades	16
Diameter	508 mm
Design speed	11000 rpm
Design massflow	36.0 kg/s
Design pressure ratio	1.35
Design M_{tip}^r	1.0

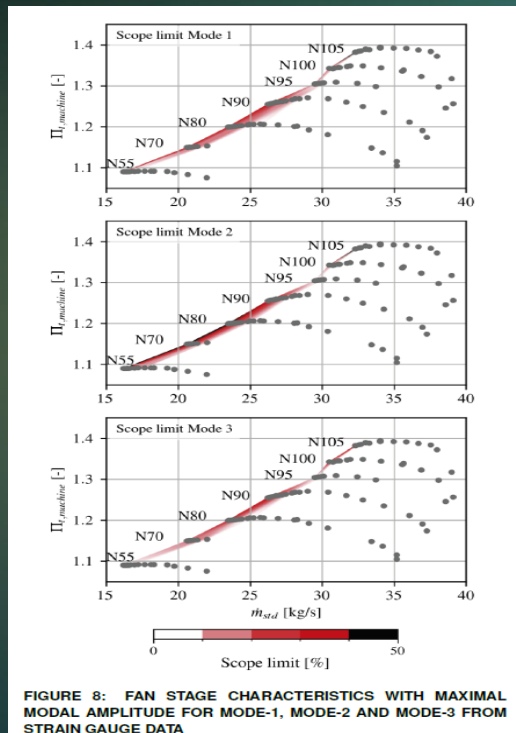


FIGURE 8: FAN STAGE CHARACTERISTICS WITH MAXIMAL MODAL AMPLITUDE FOR MODE-1, MODE-2 AND MODE-3 FROM STRAIN GAUGE DATA

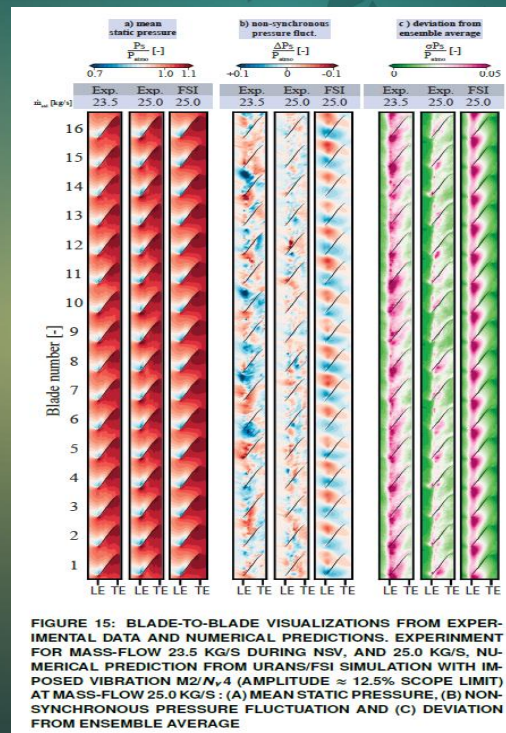






FIGURE 15: BLADE-TO-BLADE VISUALIZATIONS FROM EXPERIMENTAL DATA AND NUMERICAL PREDICTIONS. EXPERIMENT FOR MASS-FLOW 23.5 KG/S DURING NSV, AND 25.0 KG/S, NUMERICAL PREDICTION FROM URANS/FSI SIMULATION WITH IMPOSED VIBRATION $M2/N,4$ (AMPLITUDE $\approx 12.5\%$ SCOPE LIMIT) AT MASS-FLOW 25.0 KG/S: (A) MEAN STATIC PRESSURE, (B) NON-SYNCHRONOUS PRESSURE FLUCTUATION AND (C) DEVIATION FROM ENSEMBLE AVERAGE

Stall Cell moving at $\sim 58\%$ wheel speed




Rotating Stall - ECL5/Catana

- 23 M elements, full wheel
- 30 min / rev on 8 Nvidia L40 GPUs
- Rotating stall cell stabilized after 20 rev
- Unsteady full wheel analysis carried out to over 80 rev
- Rotating stall analysis is realistic with GPU CFD



Inlet Distortion


- Understanding the impacts of distorted inflow is critical to closing the engine design
- Must run full annulus as the distortion pattern may not be axi-symmetric
- Apply a prescribed pattern upstream of the fan face or apply your own pattern from an external airframe analysis
- Run steady state or time accurate and use ADS post processing to quantify engine sensitivity to distortion
- *With GPU acceleration, distortion simulations can be completed in 1 day rather than 1 month*



*Radial Axisymmetric, Circumferentially varying, Arbitrary
Total Pressure, Total Temperature, Swirl, Radial flow
angles, Turbulence Intensity and Length Scale*

Airframe + Engine coupling





- Structured or unstructured meshes can be used for the entire aircraft, including inlet, nacelle, and engine
- Unstructured mesh can be coupled to structured multiblock meshes
- Steady state or time resolved analysis
- *Co-optimize inlet and engine in the design cycle*



*Nacelle with Fan simulated at an
Angle of Attack of 45 degrees*






Total Pressure Distortion Transfer @ AoA 45°

- 67 M elements, full wheel
- 1.5 hrs / rev on 8 Nvidia L40 GPUs



4th HIGH LIFT COMMON RESEARCH MODEL


- Lift curve generated with a custom mesh from Ennova CFD
 - 50M Elements with higher grid density on the high lift surfaces
 - All simulations were executed on a GPU laptop with a single A4090 GPU with 16GB of GPU memory
 - 1 point executes in 5h 30m
 - Full lift curve with cold start completes ~44 hours (8 points)
 - Full lift curve with warm start completes ~22 hours (8 points)
 - An 8 Nvidia A100 system can execute 1 point in under 45 minutes - **Lift Curve in 6 hours**



Multi-Propulsor


- Driven by interest from the VTOL and Electrified Aircraft industries, simulating multiple propulsors is a necessity as the engine become more tightly integrated with the airframe
- Multiple non-traditional operating conditions must be analyzed, take off, hover, transition, and cruise flight
- ADS enabled multiple independent rotors to be simulated at any arbitrary axis of rotation
- Simulation times in under 20 minutes per point

** Completed on an 8GPU Compute Server




Accelerate Everything


- With ADS, all CFD methods run efficiently on the GPU
- Add a new level of fidelity to your design process with accelerated time accurate analysis
- Reducing simulation turnaround time from weeks to hours, new insights can be gained




Large Eddy Simulation




Multi-stage Turbomachinery




Coupled Airframe And Engine



Aircraft



VTOL



Radial Turbomachinery

Summary

- Turbomachinery GPU CFD needs to include methods and models for design level analysis
 - Periodic BC -- 1 min per blade row
 - Mixing plane -- 2 min per stage
 - Frozen rotor -- 2-3 min per stage
 - Sliding mesh ~ 30 min to 1 hr./ rev
 - 30X – 100X speedup over CPU based CFD
- GPU and CPU codes need to provide **identical** solutions
 - Retain 50+ years of turbomachinery CFD experience and validated models
 - Extensive validation required if GPU CFD has a different answer
- Steady and unsteady RANS stay the top turbulence model of choice
 - LES needs to improve turnaround time to become a useful design level tool

Topics of interest

- Turbulence modeling and issues in turbomachinery
- Rotor – stator interaction modeling in turbomachinery
- Steady vs unsteady comparison
- Stage matchup, performance and loss

GPU ACCELERATED CODE LEO

Full solver has been GPU Accelerated for maximum performance

All simulation technologies have been made efficient for the GPU

High memory efficiency – 20M nodes per 16GB GPU memory

Handles all the latest NVIDIA GPUs – L40, A100, A40, A10, RTX A6000, RTX A5000, etc.

GPU ACCELERATED CODE LEO

DELIVERS

50% REDUCTION IN COST PER SIMULATION

50X IMPROVEMENT IN SIMULATION TIME

IDENTICAL RESULTS BETWEEN CPU AND GPU CODE

1000 CPU CORES ON YOUR DESKTOP

GPU ACCELERATED CODE LEO

WHAT THAT MEANS

MONTH LONG SIMULATIONS COMPLETE OVERNIGHT

DELIVER ADVANCED DESIGNS ON TIME AND ON BUDGET

ENGINEERS HAVE FUN AND CAN AFFORD TO BE WRONG

HIGH FIDELITY SIMULATION IN THE DESIGN CYCLE