1. Introduction

As described in the main proposal, the code *PlasCom2* solves the compressible Navier-Stokes equations using high-order finite difference discretizations with collocated variables on a series of overlapping, logically structured meshes. MPI is used to pass data in *halo* or *ghost cell* regions to support the discrete derivative operators. OpenMP is used to offload code to attached accelerators. In this document more details on *PlasCom2* are given.

2. High-level design of *PlasCom2*

PlasCom2 is our 2nd generation fluid-thermal-structural solver that built upon lessons learned in our first generation code, *PlasComCM*. The most important lessons learned were:

- Keep separate the domain-specific logic, control flow, and data structures from the computation;
- Centralize and standardize memory management;
- Minimize dependencies and propagation of hidden, global, or implicitly passed data;
- Maintain a thin-as-possible layer between the application and input/output substrate;
- Maintain aggressive testing (unit, system, continuous);
- Build provenance and repeatability into the simulation life cycle;
- Develop an follow software practices that maintain code in a manageable state with many contributors; and
- From the Stampede2 dedicated allocation time (DAT): don't depend critically on external libraries, tools, or packages for input/output (HDF5 and POSIX conflict)

The end result of this high-level design is a predictive science code where the computations and communication are kept separate, and the computations are expressed as hierarchical kernels. Internode communication is handled using MPI-2 routines while intranode communication, thread management, and accelerator offloading uses OpenMP directives available in versions 4.5 and newer. Computation is kept separate from input/output using collective MPI-I/O communication in conjunction with HDF5.

3. Stampede2 performance of *PlasCom2*

During the Stampede2 DAT on April 25, 2018, we had an opportunity to scale *PlasCom2* for a problem size typical of those in the main document using the newly-implemented multispecies WENO method. The result in shown in Fig. 3.1 where it seen that strong scaling drops off when the number of points-per-core drops below 2,000. For larger per-core problem sizes the parallel performance is excellent.

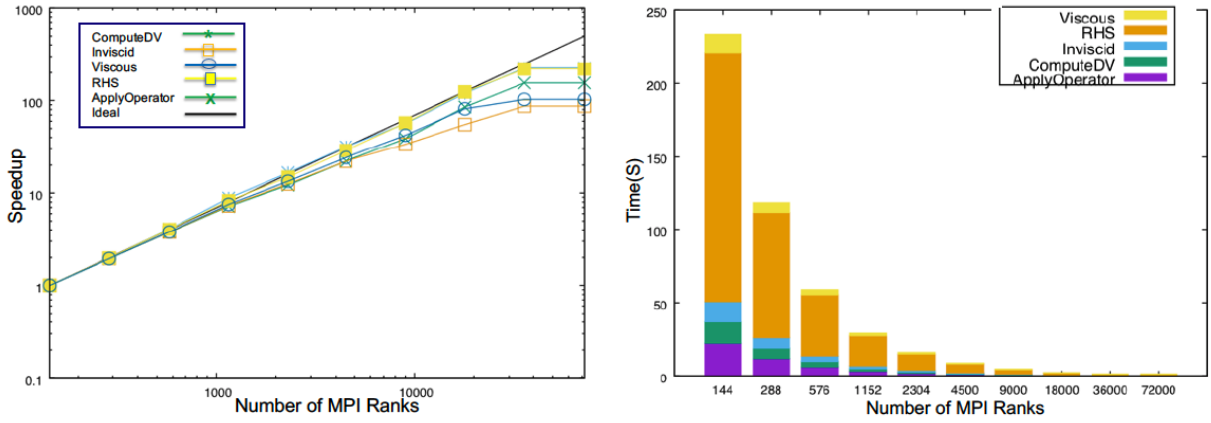

Figure 3.1: Strong scaling of *PlasCom2* on Stampede2 on a problem with 72M mesh points, WENO-based advection, and a 7-species Hydrogen/air gas model.

4. Stampede2 dedicated allocation time (DAT)

On April 25, 2018, our group was offered a 24 hour window of time during which we could have dedicated allocation time (DAT) on Stampede2 with no other users. TACC personnel were instrumental in making this happen and we are grateful to Dr. Dan Stanzione, TACC's director, for encouraging us to pursue it. In this portion we document some of the results obtained during the DAT using *PlasCom2*.

Weak scaling across KNL or SKX nodes Using the same problem described in Fig. 3.1 a weak scaling study was conducted across the KNL nodes and, in a separate set of runs, across the SKX nodes. The smallest node count required to support the 72M point mesh was 8 KNL nodes and 3 SKX nodes, respectively. In keeping the number of points per node constant, the problem was scaled by factors of 8 from the minimum, with data shown in Fig. 4.1. It was observed that the performance on the KNL nodes slightly suffered as the amount of data transferred between nodes increased (see Halo and Viscous items). On the SKX nodes, however, the I/O system had severe problems with the mesh and solution input/output at small scales. This I/O behavior was confirmed in a set of subsequent runs with I/O turned off. It was discovered during this run that POSIX compliance was interfering with the storage targets by requiring global file locking. When a new `$SCRATCH` filesystem mount point was created with local file locking (`local_flocked`) enabled, the I/O performance on the SKX nodes improved considerably.

Running *PlasCom2* across the KNL-SKX nodes A main purpose of the Stampede2 DAT was to test *PlasCom2*'s load balance utility when a single job was run across the KNL and SKX nodes. For this experiment TACC personnel set up a special queue that would launch jobs across both node types, rather than the standard node type-specific queues. The objective of the test was to evaluate whether *PlasCom2*'s load balance utility could identify that the KNL nodes have less performance than the SKX nodes and move work from the KNL nodes to the SKX nodes to improve performance. Figure 4.2 shows that this was indeed the case for the same problem described in Fig. 4.1 using only 512, 1024, or 2048 total nodes. (In each test, 1/2 of the nodes were KNL and 1/2 were SKX.) The slow loss of parallel efficiency as the problem size increases from 512 to 2048 nodes is shown for the case without load balance ('No LB') When load balancing is turned on the time-per-timestep decreases and the performance at larger node counts improves. Two load

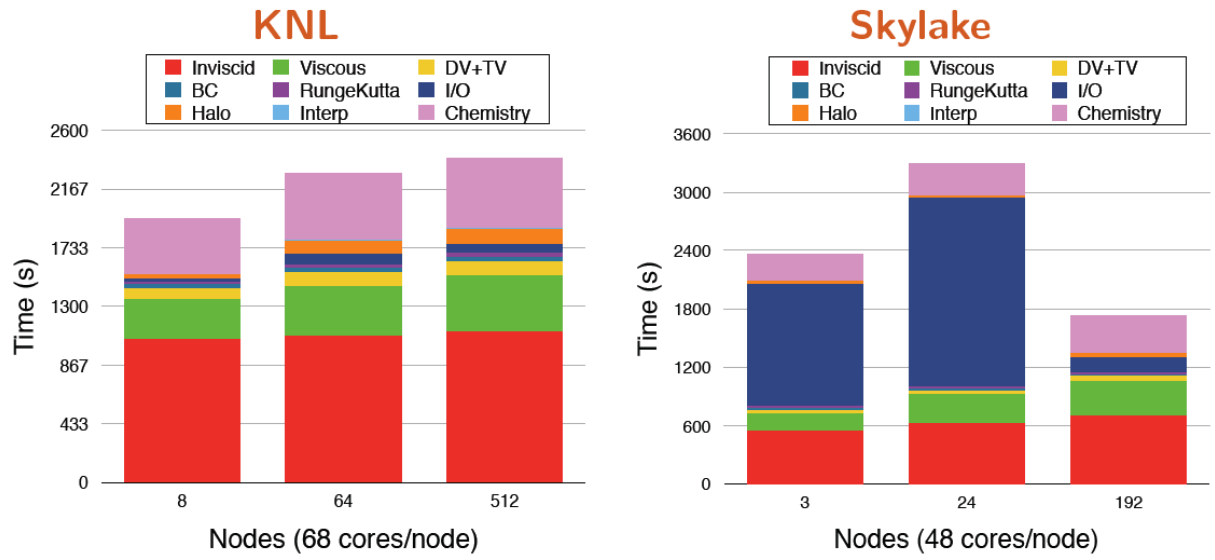

Figure 4.1: Weak scaling for problem described in Fig. 3.1 on KNL and SKX nodes on Stampede2. Observe the I/O challenges on the SKX nodes.

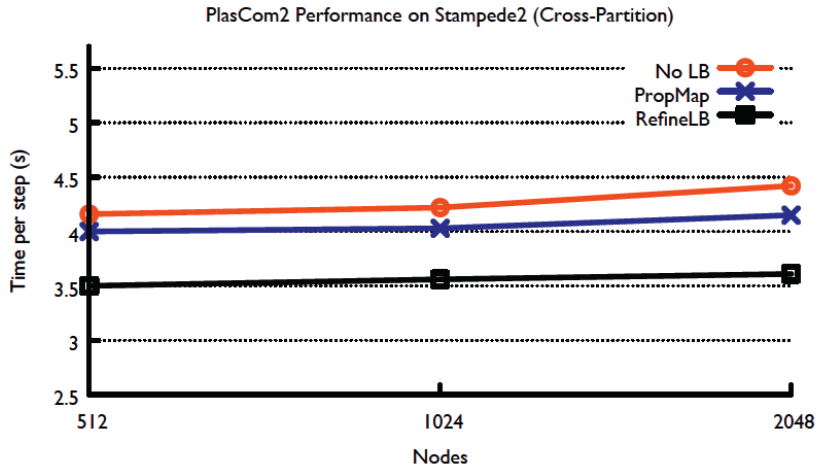

Figure 4.2: Cross KNL-SKX performance using *PlasCom2*'s load balancer.

balancers were tested: 'PropMap' measures the CPU frequency and estimates a node peak FLOP rate, and remaps the problem once. In the 'RefineLB' dynamic measurements are taken every 10 iterations and the work is continually redistributed until a minimum time-per-timestep is achieved, after which the load balancer stops interfering.

Bibliography

[1] Thomas F. Brooks, D. Stuart Pope, and Michael A. Marcolini. Airfoil self-noise and prediction. NASA reference publication 1218 65-1907, NASA, 1989.

[2] E. Envia. Fan noise reduction: an overview. Technical report, Reno, Nevada, 2001.

[3] E. Envia. Fan noise source diagnostics test - vane unsteady pressure results. In *Proceedings of the 8th AIAA/CEAS Aeroacoustics Conferences and Exhibit*, 2002. AIAA Paper 2002-2430.

[4] E. Envia, D. L. Tweedt, R. P. Woodward, D. M. Elliott, E. B. Fite, C. E. Hughes, G. G Podboy, and D. L. Sutliff. An assessment of current fan noise prediction capability. Presented at the 14th AIAA/CEAS Aeroacoustics Conference and Exhibit, AIAA Paper 2008-2991, May 2008.

[5] Edmane Envia, Alexander G. Wilson, and Dennis L. Huff. Fan noise: A challenge to CAA. *International Journal of Computational Fluid Dynamics*, 18(6):471–480, 2004.

[6] Ana Garcia-Sagrado and Tom Hynes. Wall pressure sources near an airfoil trailing edge under turbulent boundary layers. *Journal of Fluids and Structures*, 30:3–34, 2012.

[7] S. P. Gravante, A. M. Naguib, C. E. Wark, and H. M. Nagib. Characterization of the pressure fluctuations under a fully developed turbulent boundary layer. *AIAA Journal*, 36(10):1808–1816, 1998.

[8] William Wolf. *Airfoil aeroacoustics: LES and acoustic analogy predictions.* PhD thesis, Stanford University, 2011.

1. Introduction

Improving our utilization of fluidic-based devices impacts several areas of society, including transportation, energy, human physiology, and the environment. A three-pillared approach using experimentation, simulation, and theoretical modeling is required to understand how fluidic systems work and how design changes impact their performance because of the conceptual difficulties that arise due to the governing dynamics being multi-dimensional and non-linear. When the fluid systems involve subcomponents for which experimental techniques either do not exist or are not suitable then simulation is the only useful tool with which to examine the system.

This proposal considers five such scenarios where the direct numerical simulation of a compressible, viscous fluid is required to understand the system behavior. Each scenario is supported by an external grant and is conducted by a separate graduate student or postdoctoral fellow. Two of the five projects involve generated noise from an aerodynamic or structural origin while two are fundamental studies of boundary layers. The fifth project seeks to develop adjoint-based methods for shock-laden flows. In particular the projects to be supported by this XSEDE request are:

1. Sources of broadband fan noise in the NASA/GE source diagnostic test fan
2. Fluid-structure interaction in hypersonic flows
3. Prediction simulations of transient ignition kernels in supersonic plasma-coupled combustion

In all situations, experiments are limited in the data they can provide and reduced-order models are not available and/or not useful. Large-scale simulations of the underlying physical systems are positioned to play a critical role in developing a foundational understanding of the physical processes involved and in guiding future experiments, for both the proposed applications in particular and for compressible turbulent flows in general.

All of the projects use an advanced complex geometry, high-order, finite-difference code, described in §6, and developed by the PI and his research team. It has been demonstrated to run with excellent scalability on TACC's Stampede-1 using target core-counts on representative problems. Significant performance improvements, both in communication and floating point algorithm speed, were found as part of our previous XSEDE allocations, and are now have the ability to run on the Intel Xeon Phis. In particular, we are making a concerted development effort to evolve the code to run profitably on the KNL-enabled Stampede-2 machine. In the following, we provide background on the basic problems whose study is to be supported by this proposal. Further details on the specific proposed simulation efforts are given in the next sections. Code-specific details are provided in §6 while a summary of the performance is given in the separate *Code Performance and Scalability* document. Preliminary results obtained in our previous XSEDE allocations are detailed in the separate *Progress Summary*, while XSEDE-supported publications are listed in the *Publications* document. All requested resources are given in Stampede-2 node hours.

2. Utilization of Current Allocation

Our current utilization of project TG-CTS090004 is below our historical standard and requires explanation. On July 1, 2017, our 2017-2018 allocation of 1,307,293 Stampede2 node hours began and supported the projects listed in our 2017 allocation request. In March, 2018, an opportunity to utilize the entire Stampede2 machine, including both the KNL and SKX nodes, for 24 hours became available during an extended downtime of the machine. This dedicated allocation time (DAT) happened on April 25, 2018, and was extremely useful to our group and is documented in the *Code Performance and Scaling* document.

To support the DAT, we budgeted our allocation to allow for all 5,932 nodes of Stampede2 to be used for 24 hours, for a total of about 143,000 node hours. In the days leading up to the DAT our remaining allocation was increased by approximately 250,000 node hours by TACC as a kind gesture. For reasons we don't fully understand our allocation was only charged approximately 10,000 node hours for the entire April 25th DAT. So, by the end of the DAT, our remaining allocation was approximately 383,000 node hours *higher* than we anticipated having.

Our current utilization, at the time of this proposal as documented on `portal.xsede.org`, is at 48%, showing us using 630,000 node hours out of a total of 1.3M node hours. In actuality we have burned nearly 80% of our original allocation, but through TACC supplements and less usage during the DAT, our remaining balance shows much less utilization. By our own accounting, we should have roughly 247,000 node hours remaining which would cover our usage for the three months between proposal submission and the beginning of the 2018-2019 allocation on September 1, 2018. We respectfully request the XSEDE Resource Allocation Committee to consider this explanation for our remaining utilization when evaluating the documented resource need.

3. Sources of broadband fan noise in the NASA/GE source diagnostic test fan (funding: AARC)

3.1 Background

Developing quieter, more fuel efficient jet aircraft has been a focused goal since the 1950s. The movement of engine designs from the original turbojet to low-bypass ratio turbofans to today's high-to-ultra high bypass ratio turbofans has sustained noise and fuel use reduction goals. Future growth in aircraft utilization and more stringent requirements on acceptable noise suggest, as concluded by the FAA [39], that "environmental impacts may be the fundamental constraint on air transportation growth in the 21st century."

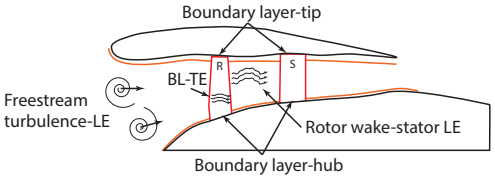

Figure 3.1: Schematic representation of broadband fan noise originated from the interaction of turbulence with different sections of a rotor-stator-nacelle assembly.

The dominant noise source has, however, changed. Large bypass ratio fans produce noise that predominantly originates from the fan itself, rather than from the exhausts streams from the core and fan. While fan noise had been reduced by improved fan design and lower tip speeds, the broadband component of fan noise is still an issue [8, 10, 9]. Sources of the broadband fan noise are mostly due to the interaction of small unsteady flow perturbations with the fan (rotor) and stationary guide vane (stator) blade rows. As shown in Fig. 3.1, the broadband fan noise has the following contributors:

1. *Free-stream turbulence:* Noise generated by the inflow turbulence ingested by the fan
2. *Rotor self-noise:* Noise generated by the passing of the turbulent boundary layer over the trailing edge of the rotating blades.
3. *Rotor-casing boundary layer interaction:* Noise generated by the rotor blades intercepting the turbulent boundary layer on the casing wall.
4. *Rotor-stator interaction:* Noise generated by the turbulent component of the wakes leaving the rotor blades and then impinging the stator blades.

Analysis and modeling of the sound generation process is challenging since the source of the sound, which is fundamentally anisotropic, arises from the interaction of turbulence with loaded and possibly rotating blades. Several approaches have been pursued to analyze the scattered noise.

Early work of Ffowcs Williams and Hawkings [12] leveraged Lighthill's acoustic analogy [26, 27] applied to surfaces in arbitrary motion [11] where the unsteady blade pressures were approximately integrated. More recently, Glegg and Jochault [17] used the relationship of Amiet [1] to reconstruct the surface pressure spectrum of an isolated blade based on the corresponding measured acoustic field. The surface pressure spectrum was then applied to the Lighthill's acoustic analogy to specify the broadband noise from a ducted fan.

In the modeling attempts, successful broadband noise prediction depends on the turbulent flow statistics that need to be known *a priori*. In addition, these models are generally based on the Lighthill's acoustic analogy applied to a simplified geometry with significantly reduced flow complexity. Therefore, in such models, many of the underlying fluid dynamics contributing to the scattered fan noise may have been excluded or over-approximated.

Figure 3.2: NASA/GE SDT fan.

Because of the flow and geometrical complexity, an alternative approach is to use high-fidelity numerical simulations as an investigative tool, and directly measure the acoustic pressure waves from the computational domain. Recent works of Casalino et al [7] and Li et al [25] have taken initial steps towards direct noise computation (DNC) using well-resolved large eddy simulation (WRLES) and direct numerical simulation (DNS). Resolving all scales of turbulence in a large-scale geometry, such as a rotor-stator-nacelle assembly, is currently unfeasible. However, recent developments in *wall-modeling* approach [2, 40, 20, 21, 35, 15] have offered the promise of DNC for large-scale geometries.

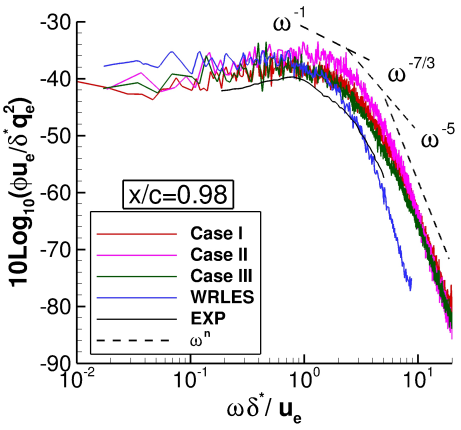

Figure 3.3: Comparison of the pressure power spectral density normalized by outer-layer scales on different NACA0012 airfoil grids at $x/c = 0.98$.

In wall-modeling approach, the near-wall small scales of turbulence are represented by either an algebraic function, or by solving an ordinary/partial differential equation on a much coarser grid near the solid boundary. The wall-model then provides the boundary condition for the bulk of the flow, the scales of which are not affected by the presence of solid boundaries. As a consequence, the demanding spatial and temporal resolutions for the near-wall viscous scales substantially decreases. In the current study, we intent to use the equilibrium wall-model of Kawai and Larsson [20] to perform wall-modeled large eddy simulation (WMLES) of rotor-stator-nacelle assembly for the NASA/GE Source Diagnostic Test (SDT) fan [9] and use the DNC technique to measure the sound scattered from the assembly.

3.2 Objectives

To improve the information available on the two primary sources of modeling uncertainty of broadband fan noise–the relative importance of the individual sources sound and the turbulence characteristics on the blades and in the intrastage wake–simulations motivated by the NASA/GE Source Diagnostic Test (SDT) fan with a 22-blade R4 rotor and 26-vane radial OGV depicted in Fig. 3.2

will be conducted. Several geometry and flow condition changes will be assumed to make the simulations computationally feasible. The simulations will specifically seek answers to these fundamental issues that arise in modeling broadband fan noise:

Q1 What is the relative acoustic importance of blade self-noise, rotor wake-stator interaction noise, and rotor wake 'quadrupole' noise?

Q2 What is the statistical description of the intrastage wake turbulence?

The simulation methodology is WMLES based on the equilibrium wall model implementation of Kawai Larsson [20] which has been used to predict wall flow statistics in canonical geometries.

Prior to the production simulations, the capability of the WMLES approach together with the wall model to predict scattered sound will be tested for flow over an isolated NACA0012 airfoil with zero angle of attack, free stream Reynolds number of 408,000 and Mach number of 0.115. The airfoil configuration is chosen based on the similarity to the NASA/GE SDT fan geometry. We used a portion of the last last year's allocation to perform WMLES of the targeted airfoil flow problem for three different grid resolutions. We have also started a DNS of the same flow configuration to produce a model-free data set for bench-marking (cf. the attached progress report for details). The spanwise-vorticity isosurface of the flow has been shown in Fig. 3.4 showing the development of a turbulent boundary layer over the airfoil surface from the WMLES of three different grids.

Also corresponding pressure density power spectra upstream of the trailing edge of the airfoil in Fig. 3.3. Although consistent with available data sets, our analysis indicated that we need to complete our DNS in order to accurately evaluate the performance of our WMLES implementation.

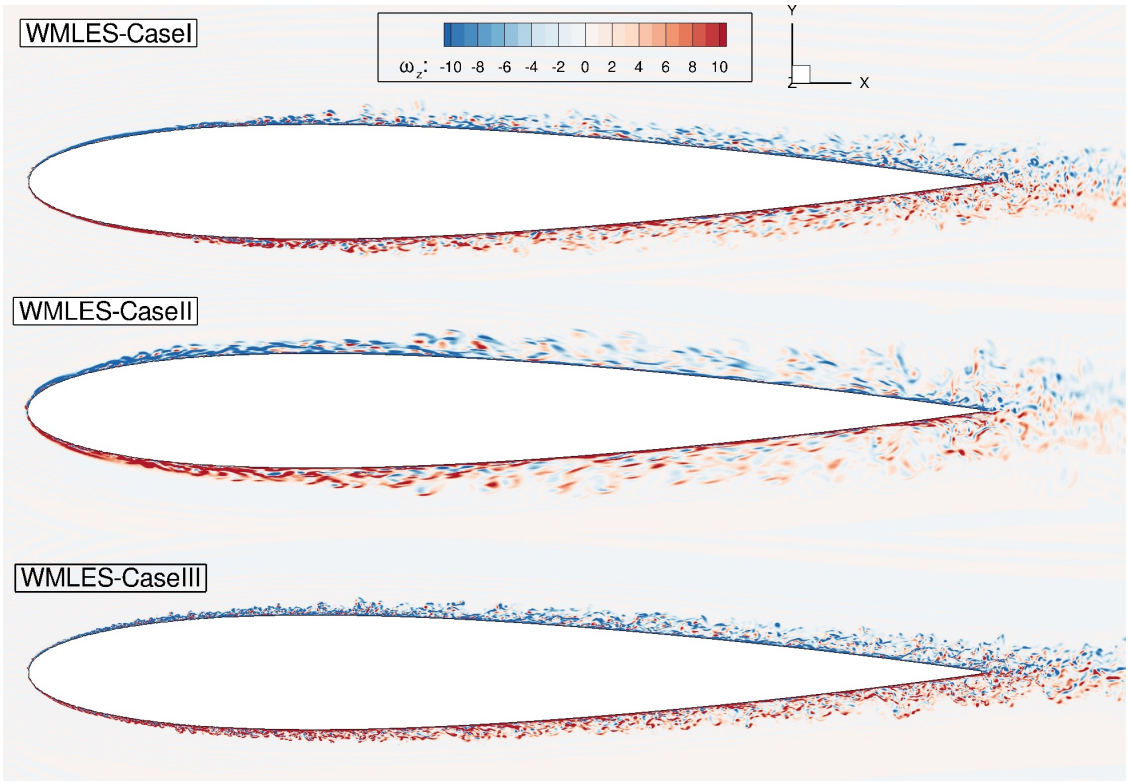

Figure 3.4: Instantaneous streamwise vorticity contour plots. Case I represents the base mesh with 35M grid points, Case II represents the coarse mesh with 16M grid points, and Case III represents the finest mesh with 79M grid points.

Description	Simulation type	Grid points	SUs	Runs	Total
Isolated airfoil	DNS	600M	937K	1	937K
NASA/GE SDT	WMLES, 1/22	100M	21K	1	21K
	WMLES, 1/2	1.5B	313K	1	313K
Total node hours for this project				3	1,271K

Table 3.1: Estimate of the computational resources required for the simulations of WMLES validation and NASA/GE SDT fan.

In addition, the SDT fan geometry grid has been recently generated for one blade of the rotor and one blade of the stator from the engine inlet to exhaust for $1/22^{\text{nd}}$ assuming the azimuthal periodicity. It should be noted that with this selection, the solidity of the OGV blades is modified by reducing the total number of OGV blades from 26 to 22. Nevertheless, in the absence of a detailed quantification of the flow for the SDT fan, WMLES of this reduced-size geometry provides a preliminary basis to design an appropriate grid for the production simulation of SDT fan with 11 rotor and 13 OGV blades (half of the actual geometry).

3.3 Cost estimates

The DNS of isolated airfoil is performed using multi-block grids that are fine very close to solid boundaries and then appropriately coarsen to the far field. The estimated number of grid points for DNS simulations is 600M. The simulation is required to be decomposed on 6,144 cores on 128 Skylake nodes and run for approximately 3 flow through times (FTs). With the computation time for the DNS estimated as 100 micro-sec/point/timestep and the total number of 900k timesteps for 1 FT, the DNS cost is 937K node hours on Stampede2. Similarly, for the WMLES of the SDT fan, with 100M/1.5B grid points each demanding 120 micro-sec/point/timestep and 100k timesteps required for 1 FT, the WMLES of NACA0012 airfoils requires 21K/313K SUs for about 3 FTs on Skylake nodes. **In summary, this project requires approximately 1,300,000 node hours on Stampede2.** Additional details on the DNS and WMLES calculations are given in the accompanying *Progress Report*.

4. Fluid-Structure Interaction in Hypersonic Flows (funding: Air Force Office of Scientific Research)

4.1 Background

Sustained hypersonic flight imposes structural and thermal loads on the aircraft exterior body panels and control surfaces that are difficult to predict using design-appropriate computational tools. Current design methods rely on conservative estimates for the fluid-generated loads, leading to overbuilt and heavy vehicles or, worse, to unanticipated structural failures. One such example is shown in Fig. 4.1 where, in 1968, a modified X-15 hypersonic research vehicle was flown with a dummy scramjet engine mounted on an underbody pylon. The engine-generated conical shock interacted with the pylon-generated normal shock to create an intense, localized heated supersonic jet formed between two shock-shock interaction slip lines, causing the pylon to thermally weaken and fail after approximately 30 seconds of flight at Mach 6.7 [41]. The vehicle safely landed but the dummy engine separated from the vehicle and was lost while the pylon contained severe damage.

4.2 Objectives

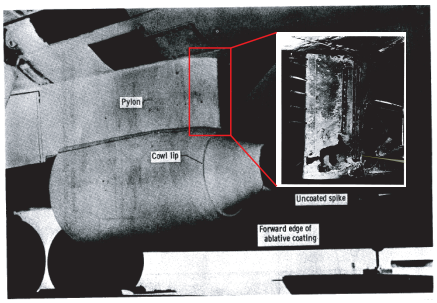

Figure 4.1: Main image: bottom of X-15-2 research aircraft showing pylon-mounted dummy scramjet engine. Inset: post-flight damage found on pylon caused by shock-shock interaction. Images from [41].

Preventing unanticipated failures while reducing design conservatism requires the development of new, robust reduced order models for several hypersonic flow scenarios similar to that encountered in Fig. 4.1. In this project we seek to use direct simulation techniques in the fluid and solid domains to understand when and where aerodynamic thermal and structural loads create a coupled response that are more severe than would be estimated by single discipline analysis. Our first steps towards this goal concern the validation of our flow solver at hypersonic relevant conditions with shock-laden flows and the fluid-structure interaction that occurs in a compression ramp geometry.

The fluid validation we consider is shown in Fig. 4.2 where a tunnel-mounted flat plate with an unknown but nearly sharp leading edge precedes a flat ramp. The tests were conducted in the RWTH Aachen TH-2 shock tunnel where compressed air is generated and pushed through a Mach 7.7 nozzle using an upstream shock tube with Helium and Air as working driver and driven gases, respectively [33]. Using the code described in §6 we simulate the impulsively started gas over the nozzle until it reaches steady state and subsequently determine surface pressures for comparison against available data. Figure 4.2 (bottom) indicates that while we have correctly predicted the upstream extent of the boundary layer separation our on-ramp pressure is not consistent with experimental measurement because of an incorrect leading edge model radius.

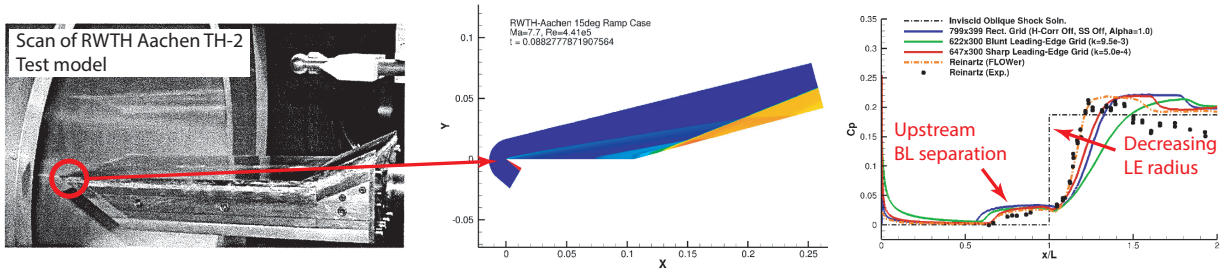

Figure 4.2: RWTH Aachen TH-2 validation case showing experimental model with unknown leading edge radius (top), computational model with estimated leading edge radius and predicted pressure field (middle), and surface pressure coefficient $C_p = 2(p - p_\infty)/(\rho_\infty U_\infty^2)$ comparisons. Observe the strong dependence on the leading edge radius. Data and scan figure from [33].

Simultaneous with our fluid validation study we are providing computational support in an AFOSR-UMD-UIUC-NASA collaboration using the Mach 6, 20 inch tunnel at NASA Langley Research Center. In this experiment, a flat plate model holds a 35 degree compression ramp fitted with a compliant panel of 4130 stainless steel. The Mach 6 freestream flow generates an oblique shock that sits on top of the flexible panel and generates a strong fluid-structure dynamical response that causes the panel to deform by an amount approximately twice the thickness (see Fig. 4.3) for the given conditions. Our study will include scenarios similar to Fig. 4.3 over a parameter range given below.

4.3 Cost estimates

This project involves two parts: the validation calculations of the RWTH Aachen TH-2 facility and the fluid-structure interaction simulations from the AFOSR-UMD-UIUC-NASA collaboration. Cost estimates are given separately below.

For the RWTH Aachen calculations, the freestream Mach number is 7.7 and the unit Reynolds number is $4.2 \times 10^6\,\mathrm{m}^{-1}$. Based on a target y^+ value on the flat plate of $y^+ = 0.5$, the first layer cell height is approximately $\Delta y = y^+ * (0.01179)^{-1/2} Re_L^{-13/14} = (0.5)(0.01179)^{-1/2}(10^7)^{-13/14} = 1.5 \times 10^{-6}$ m. Based on the expected size of the domain and the limitation on maximum cell aspect ratio, this would put the 2D grid at $1500 \times 400 = 600{,}000$ grid points. Considering the unsteady nature of the problem, a solution time of at least 3 through-flow times is expected to be needed to achieve statistical convergence. With a reference length of $L_{ref} = 0.30m$, a freestream Mach number of $M_\infty = 7.7$ and a stagnation temperature of $T_0 = 520$ K, the through-flow time is estimated at $\Delta t_{flow} = 3.1 \times 10^{-4}$ s. Taking CFL = 0.5, a single time step is estimated at $\Delta t = 7.5 \times 10^{-10}$ s, and 5×10^6 timesteps are needed for several flow through times worth of data. Using the measured computational time of 400μs per grid point per time step for the WENO-enabled fluid-only validation solver on the KNL nodes, we require 7,000 node hours for each leading radius. We anticipate needing two more simulations thus 14,000 node hours are needed for the validation.

For the AFOSR-UMD-UIUC-NASA collaboration simulation, a similar analysis to that in the preceding paragraph with the WENO-enabled fluid solver coupled to the thermo-mechanical solver shows that approximately 2,000 KNL node-hours are needed for each unit flow through time of approximately 8×10^{-4} s; however, the panel response times are much longer than a flow through time as Fig. 4.3 indicates. We estimate approximately 15 flow through times are needed to acquire meaningful structural response statistics leading to each simulation requiring 30,000 KNL node hours. Our parameter study is motivated by the experimental measurement campaign where ramp angles of 10°, 20°, 30°, and 35°were considered for a single panel thickness of 0.032 inches, the minimum considered failure-safe by NASA LaRC personnel. Our simulations will expand on the experimental panel thickness to include values of 0.008 inches, 0.016 inches, and 0.064 inches to explore panel responses on either side of the measured values. Thus we will need to simulate 16 separate configurations (4 ramp angles, each at 4 thicknesses) at a cost of 30,000 node hours each. The total requirement for the AFOSR-UMD-UIUC-NASA collaboration is thus 480,000 node hours. The total estimated resource for this project is **494,000 node hours.**

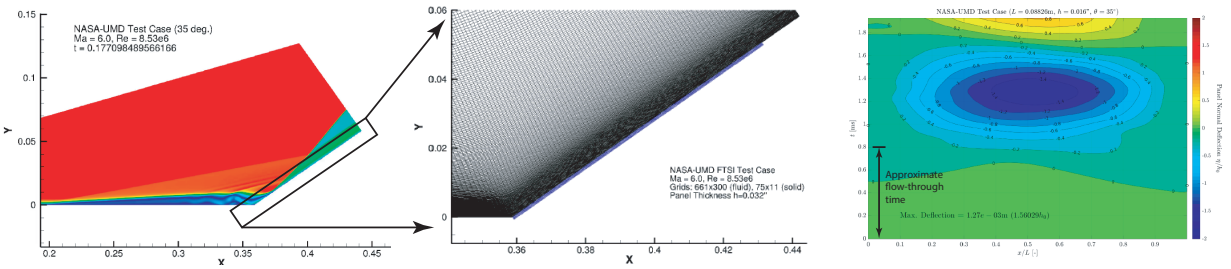

Figure 4.3: Computational prediction of the Mach 6 flow interacting with a 35 degree ramp fitted with a 4130 stainless steel panel. Left panel shows the flow, including the unsteady upstream separated boundary layer that impacts on the panel (middle image) and generates the space-time response shown in the right image. For a panel thickness of 0.016 inches the panel deflects approximately twice the thickness.

5. Prediction simulations of transient ignition kernels in supersonic plasma-coupled combustion (funding: DOE)

5.1 Background

True simulation predictions with quantified uncertainty are made for a plasma-coupled combustion flow (shown in figure 5.1). Ignition of a novel slotted-hydrogen jet configuration in $M = 2.7$ crossflow is seeded using optical laser induced breakdown (LIB) for 300 milli-Joule per laser pulse. The details of the experimental configuration which the computational model is based is described in full elsewhere [23]. The LIB hotspot dissociates the local H_2-O_2 mixture into elemental species, generates vorticity, and heats the gas to over 10^4 Kelvin [30], activating the chemical mechanisms within $\approx 10^{-9}$ seconds. Local velocities due to the LIB expansion exceeds 10^3 m/s [30]. Inviscid fluxes of the compressible, reacting flow equations were discretized with WENO, and viscous fluxes used sixth-order repeated first-order finite differences with summation-by-parts properties of their discrete operators. A detailed hydrogen chemical mechanism is used to model the radicals. Comparison is made with corresponding experiments. We focus particularly on the kernel's position, size and orientation based on its numerical chemiluminescence distribution (in part from the generation of hydroxyl radical OH). An example of the prediction is shown in figure 5.2.

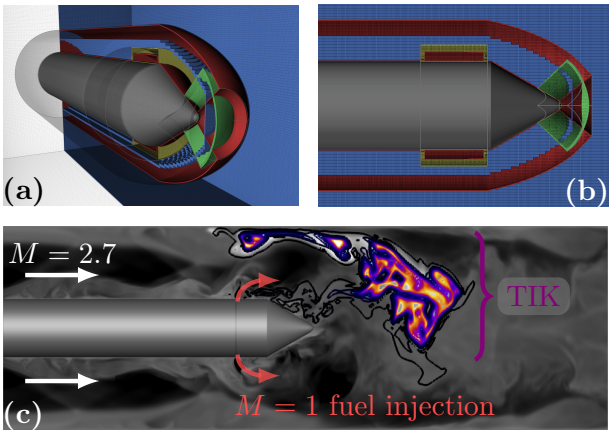

Figure 5.1: (a) Overset meshes ACT-II test section [23] and side view in (b). Streamwise plane of a transient ignition kernel (TIK) colored by $\log(Y_{OH})$ with temperature in grays.

5.2 Objective

For this task, the primary use of Stampede2 will be to conduct the following simulation tasks in the newly developed WENO solver in *PlasCom2* on the 26 million grid point overset mesh in figure 5.1 for the supersonic flow as discussed. Specific tasks involve:

- Preliminary baseline simulations establishing the turbulence simulating 300μs (10^5 time steps). This entails four separate simulations with different inflow profiles due to uncertain inflow conditions from the experimental setup.
- The primary prediction (shown in figure 5.2) for the reacting flow is the development of the initial flame kernel at discrete times up 6μs (18×10^3 time steps). The prediction involves two separate breakdown positions: one located upstream of the jet exit and the other downstream. These locations correspond to experimentally-measured positions. Four simulations will be used at each location to assess the aleatoric turbulence effect on the kernel characteristics (as shown by error bars in figure 5.2). Uncertainty propagation requires approximately 40 additional three-dimensional simulations.
- Simulations to 95μs (10^5 timesteps) are then used to examine longer-time kernel dynamics leading to a sustained flame. Comparisons with experiments will follow.

5.3 Cost estimates

For the 2018-19 Stampede2 allocation, the prediction simulations will use the methods and codebase implemented in *PlasCom2*, on a geometrically different configuration – different from last year – which includes a cavity to help stabilize the flame. The mesh is larger than last year, currently by a factor of 2.5 (totaling 68 million grid points before mesh independence is confirmed). Ethylene will be simulated as a prospective scramjet fuel due to its large energy per unit volume and advantages of storage[4]. This change in fuel will entail an additional cost of up to 24× the previous simulations due to a more restrictive time scales (reducing it by $\frac{1}{6}\Delta t$) and transporting 4× the number of reacting species.

The *PlasCom2* 120 μs time per grid point per timestep estimate is valid for WENO methods with single species gases without combustion. For last year's calculation with a 7-species hydrogen / air gas model and reactions, the corresponding time per timestep increased to 250 μs. Using Ethylene increases that time again to apporixmately 400 μs. We thus estimate that the resources required for this project are $(10^5 + 40 \times 18 \times 10^3 + 10^5) \times 68 \times 10^6 \times 400 \times 10^{-6} \implies$ 145,000 node hours. Thus, under an equivalent simulation campaign as the previous year, the number of anticipated node hours would be **145,000 node hours**.

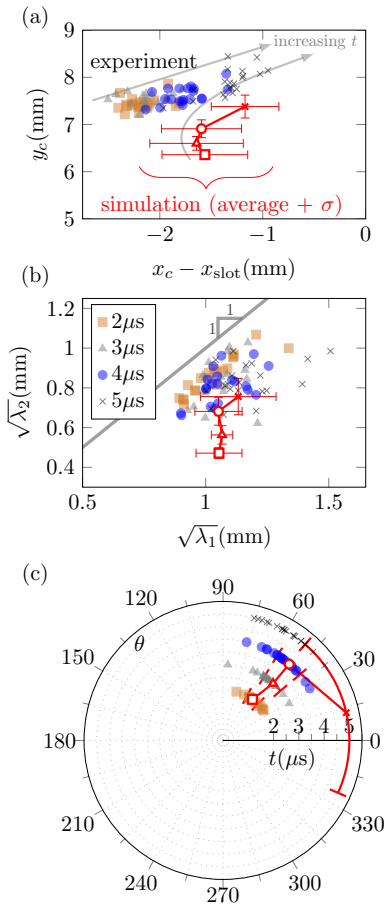

Figure 5.2: Transient ignition kernel prediction: (a) centroid, (b) size, and (c) orientation. The bars on the simulation data represent the aleatoric uncertainty from instantaneous turbulence realizations.

6. Solver

Please see the separate "Code Performance and Scalability" document for performance-related information.

6.0.1 Fluid solver The fluid solver to be used for the proposed studies solves the compressible Navier-Stokes equation in non-dimensional form for calorically and thermally perfect mixtures of gases. The equations are written in conservative form for the uniform computational coordinates $\boldsymbol{\xi} = (\xi, \eta, \zeta)$ which are related to the Cartesian coordinates $\boldsymbol{x} = (x, y, z)$ by a time-dependent mapping $\boldsymbol{\xi} = \boldsymbol{\Xi}(\boldsymbol{x}, t)$. In the $\boldsymbol{\xi}$-coordinates the equations of motion are written as

$$\frac{\partial}{\partial t}\left(\frac{Q}{J}\right) + \frac{\partial(\hat{F}_i - \hat{F}_v)}{\partial \xi} + \frac{\partial(\hat{G}_i - \hat{G}_v)}{\partial \eta} + \frac{\partial(\hat{H}_i - \hat{H}_v)}{\partial \zeta} = \frac{Z}{J} \tag{6.1}$$

where the $\{\hat{F}_i, \hat{G}_i, \hat{H}_i\}$ are the inviscid fluxes in the ξ, η, and ζ directions, respectively. The $\{\hat{F}_v, \hat{G}_v, \hat{H}_v\}$ are the corresponding viscous fluxes. The fluid is Newtonian with a Fourier law of heat conduction. The spatial derivative operators use either explicit diagonal summation-by-parts finite (SBP) difference operators[36] or energy stable WENO (ESWENO) operators modified for curvilinear meshes and improved dispersion properties [14, 31]. Time advancement is achieved by

a suite of s-stage additive Runge-Kutta methods [22]. Currently implemented are first order fully implicit, fourth order explicit-implicit, and fourth-order explicit.

The boundary conditions supplied with the code primarily use the simultaneous-approximation-term method (SAT) [6, 37, 38, 32] that combine the SBP or ESWENO operators mentioned above with a specific enforcement of the boundary conditions through a penalty term which is added to the right-hand-side of Eq. (6.1), namely $\sigma^{I1}P^{-1}E_1A^+(\boldsymbol{q}-\boldsymbol{g}^{I1})+\sigma^{I2}\mathrm{Re}^{-1}P^{-1}E_1I(\boldsymbol{q}-\boldsymbol{g}^{I2})$. This term drives the boundary solution Q to the boundary data $(\boldsymbol{g}^{I1},\boldsymbol{g}^{I2})$ in a provably stable manner provided $\sigma^{I1}\leq -2$ and $\sigma^{I2}\leq -\frac{1}{4p_0}\max\left(\frac{\gamma\mu}{\mathrm{Pr}\rho},\frac{5\mu}{3\rho}\right)$ where P and p_0 are constants dependent on the specific ESWENO/SBP operators. The boundary data are given by $\boldsymbol{g}^{I1}=[\rho,\ \rho(\boldsymbol{u}-\dot{\boldsymbol{x}}_w),\ \rho e(T_w)+\frac{1}{2}\rho|\boldsymbol{u}-\dot{\boldsymbol{x}}_w|^2]^T$ and $\boldsymbol{g}^{I2}=[\rho,\ \rho\dot{\boldsymbol{x}}_w,\ \rho e(T_w)+(1/2)\rho|\dot{\boldsymbol{x}}_w|^2]^T$ for no-slip walls with temperature T_w and velocity $\dot{\boldsymbol{x}}_w$, both of which are computed by the nonlinear thermo-mechanical solver of §6.0.2 for interacting surfaces. The ESWENO/SBP-SAT implementation yields accurate and robust discretizations for complex-geometry problems and is our preferred method [5]. These equations are supplemented with a compressible dynamic Smagorinsky LES model, [16, 28, 29] shock capturing models [13, 19] and compact filters [24].

6.0.2 Structural solver In the solid domain, the structural state is found using a nonlinear, finite-strain finite element solver. In the finite-strain formulation, a body in the initial configuration B_0 is defined by the vector $\boldsymbol{X}$. An applied load will cause a deformation $\boldsymbol{u}$ after which the body will move to the current configuration B defined by $\boldsymbol{x}=\boldsymbol{X}+\boldsymbol{u}$. Conservation of linear momentum gives the strong form of the governing equation in the current configuration B as,

$$\rho\frac{\mathrm{d}^2\boldsymbol{u}}{\mathrm{d}t^2}=\nabla\cdot\boldsymbol{\sigma}+\rho\boldsymbol{b} \tag{6.2}$$

where $\boldsymbol{b}$ is a field of body forces, $\boldsymbol{\sigma}$ is the Cauchy stress tensor and ρ is the density. Using the principle of virtual work and pulling all quantities back to the reference configuration B_0 yields the weak form as

$$\delta W=\int_{B_0}\beta^2(\Theta)\hat{\boldsymbol{P}}:\delta\boldsymbol{F}\,\mathrm{d}V+\int_{B_0}\ddot{\boldsymbol{u}}\cdot\delta\boldsymbol{u}\,\mathrm{d}V-\int_{\partial B_0}\boldsymbol{t}_0\cdot\delta\boldsymbol{u}\,\mathrm{d}A-\int_{B_0}\boldsymbol{b}_0\cdot\delta\boldsymbol{u}\,\mathrm{d}V=0 \tag{6.3}$$

where the subscript $(\cdot)_0$ denotes a quantity in the reference configuration and $\hat{\boldsymbol{P}}$ is the first Piola-Kirchhoff stress tensor. The $\beta^2(\Theta)$ term is a thermal stretch ratio used for thermal gradients. In a nonlinear problem, if a body in the current configuration $\boldsymbol{x}$ does not satisfy $\delta W(\boldsymbol{u})=0$, then linear correction is made, *i.e.*,

$$\delta W(\boldsymbol{u}+\Delta\boldsymbol{u})=0\implies\delta W(\boldsymbol{u})+D\delta W(\boldsymbol{u})\Delta\boldsymbol{u}\approx 0 \tag{6.4}$$

where $D\delta W(\boldsymbol{u})\Delta\boldsymbol{u}$ is the directional derivative of $W(\boldsymbol{u})$ along the direction $\Delta\boldsymbol{u}$ [34]. Discretization of Eq. (6.4) uses 20 node quadratic elements for the displacement, velocity, and acceleration terms. After algebra the discretizated, linearized equation version on Eq. (6.4) can be written as

$$M\ddot{\boldsymbol{u}}+K\Delta\boldsymbol{u}=\boldsymbol{R}^{\mathrm{ext}}-\boldsymbol{R}^{\mathrm{int}} \tag{6.5}$$

where $\boldsymbol{R}^{\mathrm{ext}}$ and $\boldsymbol{R}^{\mathrm{int}}$ the external and internal load vectors, respectively, and K and M are the tangent stiffness and consistent mass matrices, respectively.

Equation (6.5) is integrated in time from time t_n to t_n+h_s using the second order, A-stable Newmark-β method [18]. The linear system that must be solved for $\Delta\boldsymbol{u}$ using a Newton-Krylov solver from PETSc [3]. The solver is fully parallel and uses MPI for the message passing. Note that the finite element solver is second order in space and in time and solves the fully nonlinear equations for the solid.

6.0.3 Fluid-solid interface The fluid and solid solvers are weakly coupled at the fluid-solid interface using a multirate time integration technique. The spatial coupling assumes that the fluid and solid nodes coincide. The information shared across the interface is as follows: the fluid solution passes the fluid stress tensor and heat flux vector, τ_{ij} and $\dot{q}_i$, to the solid while the solid passes updated surface coordinates $\boldsymbol{x}_w$, velocity $\dot{\boldsymbol{x}}_w$, and temperature T_w. The full fluid stress tensor and heat flux vectors are passed rather than the traction vector $\tau_{ij}n_j$ and normal heat flux component $\dot{q}_i n_i$ so that the updated fluid load on the structure during the Newton-Rhapson iterations can be calculated without further communication between the fluid-solid solvers across the interface.

The temporal integration of the FTSI system uses a multirate method that assumes the thermal, structural, and fluid timesteps are ordered according to $h_t = h_s > h_f$. For simplicity the thermal and structural timesteps are taken to be an integer multiple of h_f. The algorithm for marching from time t_n to $t_n + h_t$ is: (1) use a Runge-Kutta 2 method to estimate the fluid solution at times $t_n + h_t$, (2) determine the thermo-structural solution at time $t_n + h_t$ using the thermo-mechanical solver and the estimated fluid data, (3) integrate the fluid solution from time t_n to $t_n + h_t$ using interpolated thermo-mechanical data. The method has been shown to be stable and 2nd order accurate in time.

7. Local and other external computing resources resources

The work proposed involves several separate projects. Each project is (a) externally funded and (b) has a dedicated graduate student or postdoctoral fellow who has experience running on XSEDE systems. Each student or postdoc has an up-to-date desktop computer with several terabytes of local storage, at least six CPU-based cores, and several gigabytes of memory on which to perform local data processing. In addition, the students have shared access to two local clusters—one dedicated to the PI and fitted with 208 cores with another shared between four other PIs and fitted with 1,056 Intel Haswell cores—each with infiniband network, and nearly 100+ terabytes of RAID storage. In short, our local resources are sufficient to support the work proposed.

The PI has access to several non-XSEDE resources, though none of these are available to the proposed projects. For full disclosure, the PI has several million CPU hours on Department of Energy machines for DOE-funded projects, and 400,000 node hours on Blue Waters as part of the PI's 'Blue Waters' professorship. For the DOE-funded project in §5 we request NSF resources because the DOE does NOT have a production scale system with Intel Xeon Phis installed and the student is not eligible to access the available DOE hardware because of his citizenship. **None of these external allocations are to be used for the projects discussed in this proposal.**

8 References

[1] R. K. Amiet. Noise due to turbulent flow past a trailing edge. *J. Sound and Vib.*, 47(3):387–393, 1976.

[2] Elias Balaras, Carlo Benocci, and Ugo Piomelli. Two-layer approximate boundary conditions for large-eddy simulations. *AIAA Journal*, 34(6):1111–1119, 1996.

[3] Satish Balay, Jed Brown, Kris Buschelman, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc Web page, 2012. http://www.mcs.anl.gov/petsc.

[4] F. S. Billig. Research on supersonic combustion. *Journal of Propulsion and Power*, 9(4):499–514, 1993.

[5] D. J. Bodony. Accuracy of the simultaneous-approximation-term boundary condition for time-dependent problems. *Journal of Scientific Computing*, 43(1):118–133, 2010.

[6] M. H. Carpenter, D. Gottlieb, and S. Abarbenel. Time-stable boundary conditions for finite difference schemes involving hyperbolic systems: Methodology and application for high-order compact schemes. *Journal of Computational Physics*, 111:220–236, 1994.

[7] D. Casalinio, A.F.P. Riberio, E. Fares, S. Nolting, A. Mann, F. Perot, Y. Li, Y. Lew, P.T. Lew, C. Sun, P. Gopalakrishnan, R. Zhang, H. Chen, and K. Habibi. Towards lattice-Boltzmann prediction of turbofan engine noise. 2014.

[8] E. Envia. Fan noise source diagnostics test - vane unsteady pressure results. In *Proceedings of the 8th AIAA/CEAS Aeroacoustics Conferences and Exhibit*, 2002. AIAA Paper 2002-2430.

[9] E. Envia, D. L. Tweedt, R. P. Woodward, D. M. Elliott, E. B. Fite, C. E. Hughes, G. G Podboy, and D. L. Sutliff. An assessment of current fan noise prediction capability. Presented at the 14th AIAA/CEAS Aeroacoustics Conference and Exhibit, AIAA Paper 2008-2991, May 2008.

[10] Edmane Envia, Alexander G. Wilson, and Dennis L. Huff. Fan noise: A challenge to CAA. *International Journal of Computational Fluid Dynamics*, 18(6):471–480, 2004.

[11] J.E. Ffowcs Williams and Hawkings D.L. Sound generation by turbulence and surfaces in arbitrary motion. *Proc. R. Soc. London A*, 264(151):321 – 342, 1969.

[12] J.E. Ffowcs Williams and Hawkings D.L. Theory relating to the noise of rotating machinery. *Journal of Sound and Vibration*, 10(1):10 – 21, 1969.

[13] B. Fiorina and S. K. Lele. An artificial nonlinear diffusivity method for supersonic reacting flows with shocks. *J. Comput. Phys.*, 222(1):246–264, 2007.

[14] T. C. Fisher, M. H. Carpenter, N. K. Yamaleev, and S. H. Frankel. Boundary closures for fourth-order energy stable weighted essentially non-oscillatory finite-difference schemes. *Journal of Computational Physics*, 230:3727–3752, 2011.

[15] Kocheemoolayil Joseph George and Sanjiva K. Lele. Large eddy simulation of airfoil self-noise at high reynolds number. 22nd AIAA/CEAS Aeroacoustics Conference, 2016.

[16] M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. A dynamic subgrid-scale eddy viscosity model. *Phys. Fluids A*, 3(7):1760–1765, 1991.

[17] S. A. Glegg and C. Jochault. Broadband self-noise from a ducted fan. *AIAA Journal*, 36(8):1387–1395, 1998.

[18] J. L. Humar. *Dynamics of Structures*. Englewood Cliffs, 1990.

[19] S. Kawai and S. K. Lele. Localized artificial diffusivity scheme for discontinuity capturing on curvilinear meshes. *J. Comput. Phys.*, 227:9498—9526, 2008.

[20] Soshi Kawai and Johan Larsson. Wall-modeling in large eddy simulation: Length scales, grid resolution, and accuracy. *Physics of Fluids*, 24(1), 2012.

[21] Soshi Kawai and Johan Larsson. Dynamic non-equilibrium wall-modeling for large eddy simulation at high reynolds numbers. *Physics of Fluids*, 25(1), 2013.

[22] C. A. Kennedy and M. H. Carpenter. Additive Runge-Kutta schemes for convection-diffusion-reaction equations. *Applied Numerical Mathematics*, 44:139–181, 2003.

[23] Gyu Sub Lee, Qili Liu, Damiano Baccarella, Gregory S. Elliott, and Tonghun Lee. A novel supersonic injection scheme for laser induced breakdown ignition. In *2018 AIAA Aerodynamic Measurement Technology and Ground Testing Conference*, 2018/07/11 2018.

[24] S. K. Lele. Compact Finite Difference Schemes with Spectral-like Resolution. *J. Comp. Phys.*, 103:16–42, 1992.

[25] Q. Li, N. Peake, and M. Savill. Large eddy simulations for fan-ogv braodband noise prediction. 2008.

[26] M. J. Lighthill. On sound generated aerodynamically I. General theory. *Proc. R. Soc. London A*, 211:564–587, 1952.

[27] M. J. Lighthill. On sound generated aerodynamically II. Turbulence as a source of sound. *Proc. R. Soc. London A*, 222:1–32, 1954.

[28] D. K. Lilly. A Proposed Modification of the Germano Subgrid-scale Closure Method. *Phys. Fluids A*, 5:1080–1082, March 1992.

[29] P. Moin, K. Squires, W. Cabot, and S. Lee. A dynamic subgrid-scale model for compressible turbulence and scalar transport. *Phys. Fluids A*, 3:2746–2757, 1991.

[30] Alessandro Munafò, Andrea Alberti, Carlos Pantano, Jonathan Freund, and Marco Panesi. Modeling of laser-induced breakdown phenomena in non-equilibrium plasmas. In *2018 AIAA Aerospace Sciences Meeting*, 2018/07/11 2018.

[31] T. Nonomura, D. Terakado, Y. Abe, and K. Fujii. A new technique for freestream preservation of finite-difference WENO on curvilinear grid. *Computers and Fluids*, 107:242–255, 2015.

[32] J. Nordström, J. Gong, E. Van der Weide, and M. Svärd. A stable and conservative high order multi-block method for the compressible navier-stokes equations. *J. Comput. Phys.*, 228:9020–9035, 2009.

[33] H. Olivier. Hot experimental testing. Final Report WP 5510, Shock Wave Laboratory, RWTH Aachen, 1999.

[34] C. Ostoich, D. J. Bodony, and P. H. Geubelle. Aerothermoelastic response of a panel under a high speed turbulent boundary layer using direct numerical simulation. AIAA Paper 2013-1662, Presented at the 54nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, April 2013.

[35] George Ilhwan Park and Parviz Moin. An improved dynamic non-equilibrium wall-model for large eddy simulation. *Physics of Fluids*, 26(1), 2014.

[36] B. Strand. Summation by parts for finite difference approximations for d/dx. *Journal of Computational Physics*, 110:47–67, 1994.

[37] M. Svärd, M. H. Carpenter, and J. Nordström. A stable high-order finite difference scheme for the compressible navier-stokes equations, far-field boundary conditions. *Journal of Computational Physics*, 225:1020–1038, 2007.

[38] M. Svärd and J. Nordström. A stable high-order finite difference scheme for the compresible Navier-Stokes equations: No-slip wall boundary conditions. *Journal of Computational Physics*, 227:4805–4824, 2008.

[39] I. Waitz, J. Townsend, J. Cutcher-Gershenfeld, E. Greitzer, and J. Kerrebrock. Aviation and the environment: A national vision statement, framework for goals and recommended actions. Report to the united states congress, December 2004.

[40] Meng Wang and Parviz Moin. Dynamic wall modeling for large-eddy simulation of complex turbulent flows. *Physics of Fluids*, 14(7):2043–2051, 2002.

[41] J. D. Watts. Flight experience with shock impingement and interference heating on the X-15-2 research airplane. Technical Memorandum TM X-1669, National Aeronautics and Space Administration, 1968.

1. Sources of Broadband Fan Noise in the NASA/GE Source Diagnostic Test Fan (funding: AARC)

In modern, high-bypass ratio turbofan engines the reduction of jet exhaust noise through engine design has increased the acoustic importance of the main fan to the point where it can be the primary source of noise[2] in the flight direction. While fan noise has been reduced by improved fan designs and lower tip speeds it, and especially its broadband component, still remains an issue.[3–5]

The goal of this project is to apply the high-fidelity wall-modeled large-eddy simulation (WMLES) approach to the NASA/GE Source Diagnostic Test (SDT) fan geometry to resolve the turbulent flow structures within the rotor-stator stage that give rise to the sources of the aeroacoustic broadband fan noise. We used a part of the last year's XSEDE allocation to verify the capability of the WMLES in predicting aerodynamic and aeroacoustic quantities. The chosen test case for the validation purpose was the turbulent boundary layer developed on a NACA0012 airfoil with zero angle of attack in a uniform flow with chord-length based Reynolds number $Re_c = 408,000$ and uniform flow Mach number $M = 0.115$. We performed WMLES of this flow configuration for three different grid resolutions to assess the WMLES capabilities in aerodynamic and aeroscoustic flow predictions. We also started DNS of the same flow problem to obtain a model-free first-principles dataset for the validation purpose.

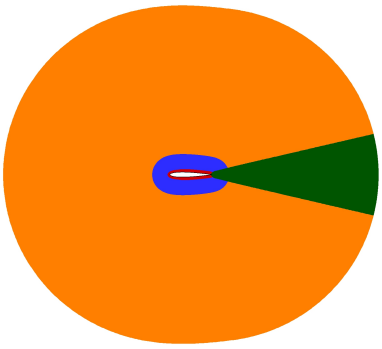

Figure 1.1: Representation of overset meshes for the NACA0012 airfoil in a uniform flow. The wall-model grid is not shown in this figure.

The generated grids exploit the overset mesh layouts to reduce the grid resolution away from the airfoil surface where the flow streamlines are not affected by the airfoil geometry. Figure 1.1 shows a representative of the computational grid with the overset meshes used at different regions of the domain in order to reduce the computational cost outside the boundary layer. Different meshes are distinctly colored the details of which are provided in Table 1.1. It should be noted in Fig. 1.1 that the WM grid is not shown since it covers a very thin region near the airfoil surface. However, the details are given in Table 1.1 in black color.

Table 1.1: Details of the grid resolution for fine and coarse grids. The colors are associated with those meshed in Fig. 1.1. The numbers in black are associated with the WM grid.

	DNS			WMLES-CaseI			WMLES-CaseII			WMLES-CaseIII		
Grid	N_x	N_y	N_z	N_x	N_y	N_z	N_x	N_y	N_z	N_x	N_y	N_z
I	504	1492	205	201	832	85	201	829	43	237	873	85
II	4345	450	205	1381	130	85	681	135	43	613	155	85
III	2173	186	103	691	106	43	681	104	43	1221	155	169
IV	725	121	52	231	91	15	341	135	22	613	155	85
V	–	–	–	1009	31	85	627	31	43	1199	110	43
VI	–	–	–	–	–	–	–	–	–	600	138	22
VII	–	–	–	–	–	–	–	–	–	228	31	169
IIX	–	–	–	–	–	–	–	–	–	1413	31	85
IX	–	–	–	–	–	–	–	–	–	228	31	169
Total	601,171,724			35,598,378			16,012,245			79,328,339		

In order to trip the boundary layer, surface roughness is added to the airfoil surface from the leading edge to the 20% of the chord for all of the grids, as shown in Fig. 1.2, consistent with the experiments of Brooks et al.[1] Figure 1.3 showing the streamwise compoent of the vorticity on the airfoil surface (top view) confirms the success of the this approach since the turbulence boundary layer forms relatively close to the leading edge of the airfoil, although the flow structures appear to be scaled wit the grid spacing. The skin friction coefficient is calculated from our WMLES cases and compared with experimental and WRLES datasets in Fig. 1.4. It is observed that near the trailing edge, the three profiles from our simulations are close to each other and predict similar values as opposed to those of[8] and.[6] In the first half of the leading edge, the differences between our cases are more pronounced which is due to the quantitative and qualitative variations in the transition.

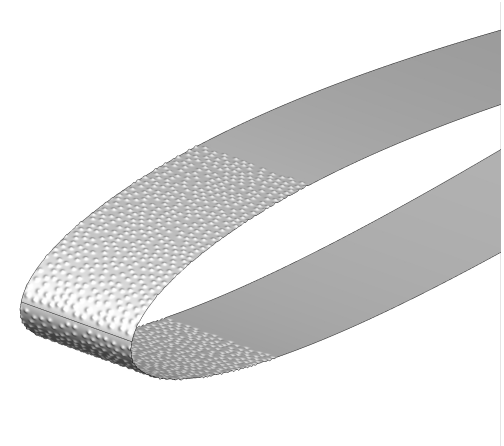

Figure 1.2: Distribution of surface roughness near the leading edge.

The mean streamwise velocity profiles normalized by the viscous scales are plotted in Fig. 1.5 at two different locations along the chord. The profiles from our WMLES are presented for both the WM region as well as the LES region with, respectively, hollow and filled symbols. The plots indicate a smooth transition of the velocity profiles between the two regions. It should be noted that although the LES region is extended to the wall, the corresponding profile is truncated at the intersection with the WM-region profile. The profiles also reveal that the height of the WM region is about 50 viscous length units. The comparisons indicate that in the log-layer, all our WMLES cases under-predict the velocity profile when compared with the experimental data, whereas WRLES of[8] over-predicts the velocity profile in the same region. In the wake region, our cases start to deviate from from each other, with Case III and II to exhibit, respectively, the highest and lowest u^+ values at the same y^+. This disparity in WMLES velocity profiles in the wake region is associated with the tripping and transition dynamics in the upstream. In addition, variations in C_f of WMLES results for a given chord location contribute to the observed variation since $C_f \propto \tau_w$ is the basis of the viscous scaling.

The broadband component of the trailing edge noise is associated with pressure fluctuations on the

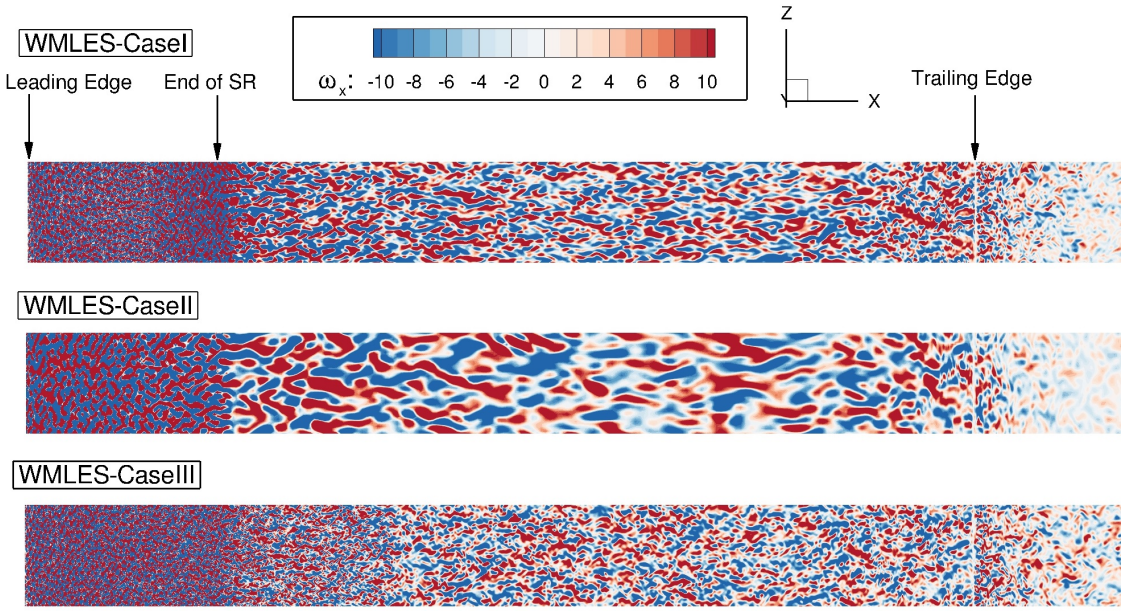

Figure 1.3: Instantaneous streamwise vorticity contour plots.

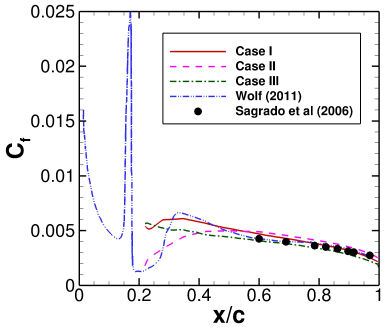

Figure 1.4: Comparison of the skin friction coefficient distribution.

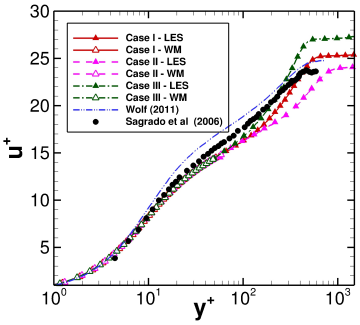

Figure 1.5: Velocity profile as a function of wall-normal distance normalized by viscous scales at $x/c = 0.92$.

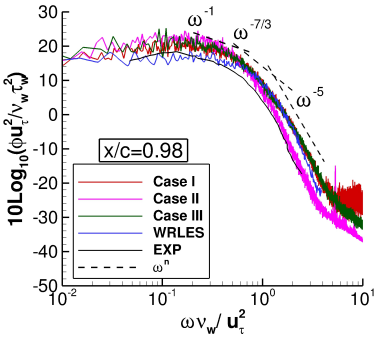

Figure 1.6: Comparison of the pressure power spectral density normalized by viscous scales on the airfoil surface at different locations.

airfoil surface that originate from large-scale to small-scale flow structures in the turbulent boundary layer. Therefore, the success of WMLES approach in predicting the far-field noise depend on the accuracy of the pressure fluctuations prediction in the boundary layer. Accordingly, pressure power spectral density (PSD) is a well-established statistical quantity that characterizes the pressure fluctuations.[7] argued that the PSD exhibits a power-law behavior, i.e. $\phi(\omega) \sim \omega^n$, with power n depending on the range of the frequency spectrum. They suggested three distinct frequency range associated with power $n = -1$, $-7/3$, and -5 that correspond to the pressure sources in, respectively, logarithmic portion of the boundary layer, upper portion of the buffer layer, and bellow the lower portion of the buffer layer. These power-law profiles are depicted with black dashed lines in Fig. 1.6 along with the PSD profiles near the trailing edge of the NACA0012 airfoil ($x/c = 0.98$).

In a similar fashion, the PSD profiles can be evaluated based on the frequency range. First we consider the low-frequency range associated with large-scale pressure containing structures. Figure 1.6 shows that all of our three cases predict higher values in the this frequency range compared with experimental and WRLES datasets. Near the trailing edge at $x/c = 0.98$ as shown in Fig. 1.6, all three cases share the same level in the low-frequency range. This behavior is because of the fact that all cases successfully capture the large-scale structures upstream of the trailing edge.

The behavior of the profiles in the high-frequency range associated with small-scale structures can also be argued from a grid-scaling viewpoint. We use a fine grid spacing near the trailing edge of all WMLES cases to resolve the bluntness of the trailing edge. Therefore, near the trailing edge, all three grids have similar grid spacings, except Case II in the spanwise direction which is coarser by a factor of two. As a result, the comparison of the PSD profiles in the high-frequency region indicates that at the trailing edge, our WMLES profiles exhibit a convergence behavior; Case I and III profiles lie on top of each other, while Case II profile approaches the other two. Also close to the trailing edge, the PSD profiles of Case I and III are in good agreement with that of[8] in the high-frequency range

These quantifications show that the WMLES is able to predict aerodynamic quantities of the flow over the NACA0012 airfoil, provided that the adequate gird spacing is used in the streamwise and spanwise directions. We continue the DNS of the NACA0012 airfoil problem in order to have a model-free dataset for validation of the WMLES predictions. We also plan to set up the SDT fan problem at approach condition during the next allocation cycle, in order to identify the sources of the broadband fan noise arising from the turbulent flow interacting with solid boundaries of the rotor and stator blades as well as the hub and nacelle walls.