NASA/TM–20220001748

Fully Coupled Aeroelastic Stability Analysis of Adaptive Shape Memory Alloy Structural Technologies for Airframe Noise Reduction

Parthiv N. Shah
ATA Engineering, Inc., San Diego, California

Eric L. Blades
ATA Engineering, Inc., Huntsville, Alabama

Michael R. Nucci
ATA Engineering, Inc., San Diego, California

Nicolas D. Reveles
ATA Engineering, Inc., Huntsville, Alabama

Travis L. Turner and David P. Lockard
Langley Research Center, Hampton, Virginia

October 2023

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- Help desk contact information:

https://www.sti.nasa.gov/sti-contact-form/
and select the "General" help request type.

NASA/TM–20220001748

Fully Coupled Aeroelastic Stability Analysis of Adaptive Shape Memory Alloy Structural Technologies for Airframe Noise Reduction

Parthiv N. Shah
ATA Engineering, Inc., San Diego, California

Eric L. Blades
ATA Engineering, Inc., Huntsville, Alabama

Michael R. Nucci
ATA Engineering, Inc., San Diego, California

Nicolas D. Reveles
ATA Engineering, Inc., Huntsville, Alabama

Travis L. Turner and David P. Lockard
Langley Research Center, Hampton, Virginia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

October 2023

Available from:

NASA STI Program / Mail Stop 148
NASA Langley Research Center
Hampton, VA 23681-2199
Fax: 757-864-6500

PREFACE

The research results presented herein constitute a final report prepared by the first author, Parthiv Shah, with contributions from the other authors for collaborative work done under Technical Direction Notice TEAMS3.C2.01.017 of NASA Contract Number 80LARC17C0003.

ABSTRACT

This final report documents work performed by ATA Engineering, Inc., (ATA) to develop computational models and analyze the coupled fluid-structure response of two types of noise treatments applied to the leading-edge-slat component of a high-lift system typical of modern transport aircraft. The first treatment is a slat-gap filler (SGF), which closes the gap between the suction surfaces of a deployed slat and an aircraft main wing, and the second treatment is a slat-cove filler (SCF), which replaces the recirculating flow on the slat cove with a surface that promotes flow attachment.

The representative airframe chosen for this work was NASA's High-Lift Common Research Model (CRM-HL) in a baseline high-lift configuration. Superelastic shape memory alloys (SMAs) have been identified as enabling materials for these structural treatments. Since the technology elements rely upon having a highly reconfigurable structure, designs were assessed for their static aeroelastic deflection as well as their dynamic aeroelastic stability using coupled computational fluid dynamics (CFD) and nonlinear computational structural dynamics (NL-CSD) tools. Specifically, fluid-structure interaction (FSI) problems were solved computationally using the CFD solver Loci/CHEM and the NL-CSD solver Abaqus. As a part of the overall project, a similar capability was implemented using the CFD solver FUN3D coupled to Abaqus, although that work is documented in a separate report and that FSI framework was not used to analyze any of NASA's SGF and SCF configurations.

The technical approach consisted of solving for the flow field around the entire vehicle using a global CFD model, followed by extraction of relevant local subdomain data for CFD and NL-CSD cosimulations. The SGF design was analyzed using both 2D and 3D cosimulations to predict quasistatic aeroelastic deformations and to assess dynamic aeroelastic stability, whereas the SCF was analyzed in 2D only.

SGF static aeroelastic response predictions focused on characterizing the deformed shape, with maximum displacements predicted to be on the order of magnitude of the technology element panel thickness. SGF dynamic aeroelastic response predictions used Partial Floquet analysis of the temporal evolution of selected nodal displacements to quantify the sign and magnitude of aeroelastic damping. Results suggest that the CRM-HL operating conditions would result in a dynamically stable response. The simulated dynamic pressure was also increased up to a factor of about four, and resulting responses suggest that predicted dynamic stability would be achieved with some margin.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 INTRODUCTION AND EXECUTIVE SUMMARY

Aircraft noise reduction is an important goal for the NASA Advanced Air Transport Technology (AATT) Project. During the approach and landing phases of flight, nonpropulsive (or airframe) noise is a priority since its contribution to overall noise is roughly equal to that of the propulsion system. Airframe noise includes aeroacoustic sources produced by the high-lift system, e.g., the leading-edge slats and trailing-edge flaps, and the aircraft undercarriage. The ranking of airframe noise sources is configuration dependent, but model-scale tests [1]–[7] and flyover noise measurements [8] have identified the leading-edge slat as a prominent source of airframe noise during aircraft approach.

The separated, multielement configuration of high-lift systems for typical transport aircraft enables greater lift and lower speeds than single-element systems. However, geometric discontinuities such as gaps, edges, and cavities cause substantial flow unsteadiness and, consequently, aeroacoustic noise. Concepts that have been proposed for leading-edge slat noise reduction include a slat-cove filler (SCF) [9] and a slat-gap filler (SGF) [10], but practical realization of either noise treatment is dependent on having a highly reconfigurable structure. Shape memory alloy (SMA) materials have been identified as enabling for these structural treatment approaches. Previous work demonstrated the feasibility of the SCF [11] and SGF [12] concepts and produced "optimized" designs with respect to sustainment of static aerodynamic loads and minimization of articulation forces for a pseudo-2D model (spanwise-uniform section of very short extent).

1.1. Problem Background

NASA has constructed a 10% scale version of the High-Lift Common Research Model (CRM-HL) developed by Lacy and Sclafani [13] for further development of airframe noise reduction technologies. The original cruise-configuration NASA Common Research Model (CRM) is an open geometry that has been widely used in past AIAA Drag Prediction Workshops [14]. The NASA CRM [15] consists of a contemporary supercritical transonic wing with flow-through nacelles and a fuselage that is representative of a wide-body commercial transport aircraft. The CRM-HL is also an open geometry and has been used in the AIAA Geometry and Mesh Generation Workshop [16] and the third AIAA High-Lift Prediction Workshop [17].

Wind tunnel tests of the CRM-HL were carried out in the NASA Langley Research Center (LaRC) 14×22 tunnel in fiscal year 2021 [18]. The primary goal of the tests was to demonstrate the effectiveness of slat-noise-reduction concepts, such as the SCF and SGF, at a technology readiness level near 5, which is higher than previous tests [9], [19]. SCFs were previously tested on a trapezoidal-wing model [9] and the 26% 777 STAR model [19], but those treatments were solid metallic parts that were incapable of being stowed. Computational fluid dynamics (CFD) and computational aeroacoustics (CAA) calculations using the CRM-HL geometry were also previously used to prescribe suitable SCF and SGF geometries for the airframe. Prior SCF and SGF structural design and analysis work [11], [12] was extended to the CRM-HL in 2D and 3D to develop structural implementations that represent functioning flight hardware. Those

structural implementations were then scaled to the 10% CRM-HL wind-tunnel model and simplified to be nonarticulating but retained boundary conditions, dimensions, and materials to meet the requirements for dynamic scaling. Accordingly, SCF and SGF treatment integration features were designed into the wind-tunnel model, all dimensions were at 10% scale (including thicknesses), and superelastic Nitinol (NiTi) material was used for the treatments.

Dynamic scaling of the SCF and SGF treatments to the CRM-HL for an aerodynamic and aeroacoustic wind tunnel study was also done to assess the impact of treatment dynamics on aeroacoustic performance. Risks associated with testing the dynamically scaled components were mitigated in the wind-tunnel model via the treatment integration details. However, coupling between the flow and elastic treatments was anticipated, with commensurate concerns of static and dynamic stability. It is recognized that incorporation of dynamically scaled components in anything but an aeroelastic test is relatively uncommon. Therefore, it is anticipated that the overarching goal of developing technology for reconfigurable/adaptive/morphing aerospace structures will make studies like the 14×22 tunnel test of the CRM-HL more commonplace in the future. This understanding motivates the need for fully coupled aeroelastic stability analyses of such adaptive SMA structural technologies for airframe noise reduction.

1.2. Project Goal

The goal of the present work performed by ATA Engineering, Inc., (ATA) was to (1) develop high-fidelity fluid and structure models for a representative subdomain of the anticipated baseline test conditions for the CRM-HL in the 14×22 tunnel and (2) to study the coupled fluid-structure response of the dynamically scaled SCF and SGF for substantiation of tunnel entry. Response includes the quasistatic aeroelastic deformation and the dynamic stability. An ancillary goal of the project was to transfer knowledge and models to NASA so that they can perform such coupled analyses in the future when maturing such technology elements.

The technical approach followed these steps: (1) CFD simulation of the global flow field under the assumption of rigid geometry, (2) extraction or generation of a subdomain mesh in some neighborhood of the technology element, (3) quasistatic cosimulation between CFD (Loci/CHEM) and nonlinear computational structural dynamics (NL-CSD) (Abaqus) to obtain the mean deformed structural shape, and (4) dynamic cosimulation between CFD and NL-CSD to assess aeroelastic stability.

Although both the SCF and SGF designs were planned for testing in the 14×22 tunnel, 2D and 3D SGF analysis and results became the focus of this study because considerable computational and experimental experience was gained on a related 2D SCF treatment in synergistic work [20] going on in parallel with this study. A dimensional analysis, similitude, and aeroelastic stability assessment study [21] was performed with information from the present work effort and that represented by [20] to substantiate the safety of the SCF and SGF treatments to the 10% CRM-HL for testing in the 14x22 tunnel. Additionally, the results of the 14x22 tunnel test of the treated 10% CRM-HL showed that the SGF performed considerably better than the SCF. Reasons for the SCF's poorer aeroacoustic performance are still under

investigation, so further post-test analysis of the SCF case in 3D was not taken on in this study and only 2D SCF analysis results are presented in this report.

1.3. Key Results and Findings

The next subsections discuss the key findings that came out of each of the major tasks discussed in this report.

1.3.1. 2D Aeroelastic Analysis

ATA successfully applied fluid-structure interaction (FSI) tools to the problems of the 2D SGF and SCF. The work was completed relatively early in the project and encompassed several exploratory efforts. For the SGF, this included aeroelastic parameter studies to bound estimates of stability, such as varying both aerodynamic and structural parameters, increasing dynamic pressure to up to six times the nominal condition, and varying the SGF stiffness via the thickness and modulus.

SGF quasistatic deflection was found to increase with dynamic pressure, and all dynamic responses were found to be stable. One noteworthy case that used an intentionally softened modulus case exhibited a unique limit-cycle oscillation (LCO) response. Dynamic FSI analyses made use of the Partial Floquet technique for damping estimation.

For the SCF, the dynamic response was evaluated with and without an initial perturbation applied to the model. It was found that while some inherent unsteadiness was present in the baseline FSI simulation run in a time-accurate model, applying an impulsive pressure to the structure allowed the response to ring down in a way where the damping was more numerically evident and quantifiable using Partial Floquet.

1.3.2. Subdomain Modeling

Subdomain modeling of FSI was found to be a viable path forward for static and dynamic stability assessment and was successfully applied to inboard and outboard sections of CRM-HL configuration meshes. It was found that the location of subdomain far-field boundaries (i.e., the size of the subdomain) can affect the solution, and some subdomain results can be sensitive to mesh quality. Meshes on the outboard wing produced excellent agreement between full semispan CFD and subdomain CFD in terms of pressure coefficient.

For inboard subdomains, solutions were obtainable, but the presence of the nacelle and wing–nacelle flow interactions presented greater challenges to numerical stability. Based on a grid-size study that was performed, use of the largest of the three meshes (shown in Figure 1-1) was recommended. The figure demonstrates that solutions could be obtained with generally close agreement between the subdomain and full-mesh solution (Cp and t) over areas of interest.

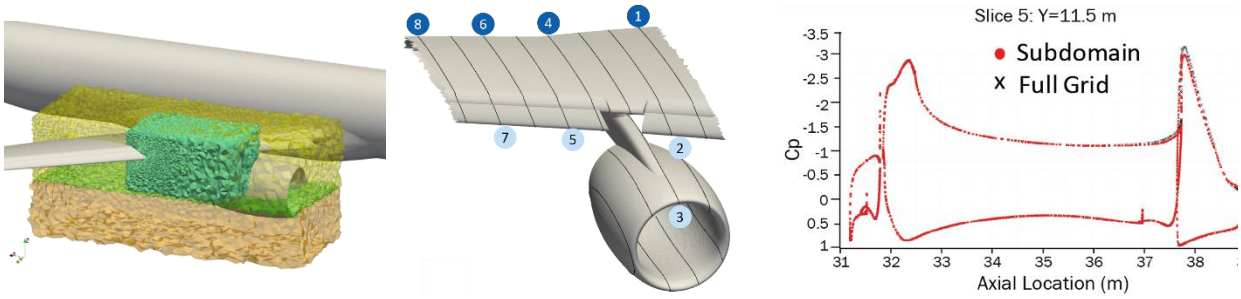

Figure 1-1. Largest subdomain (left) was shown to reproduce section 5 (middle) full semispan mesh pressure coefficient contours (right) over the wing in the region of the SGF technology element.

1.3.3. 3D Aeroelastic Analysis

Quasistatic and dynamic simulations were completed for the CRM-HL at 8° angle of attack, Mach 0.2 flight speed spanning physical and elevated dynamic pressures (q). The baseline configuration corresponded to a conventional landing scenario, and a high-q condition corresponded to a gust scenario. For quasistatic FSI, baseline and high-q loading results in deformations equal to two times and three times panel thicknesses, respectively, for the largest SGF segment analyzed. For dynamic FSI, baseline and high-q loadings excited by time-accurate perturbation loading suggested stable responses indicated by positive damping from Partial Floquet analysis. Figure 1-2 summarizes the modeling approach and typical outputs.

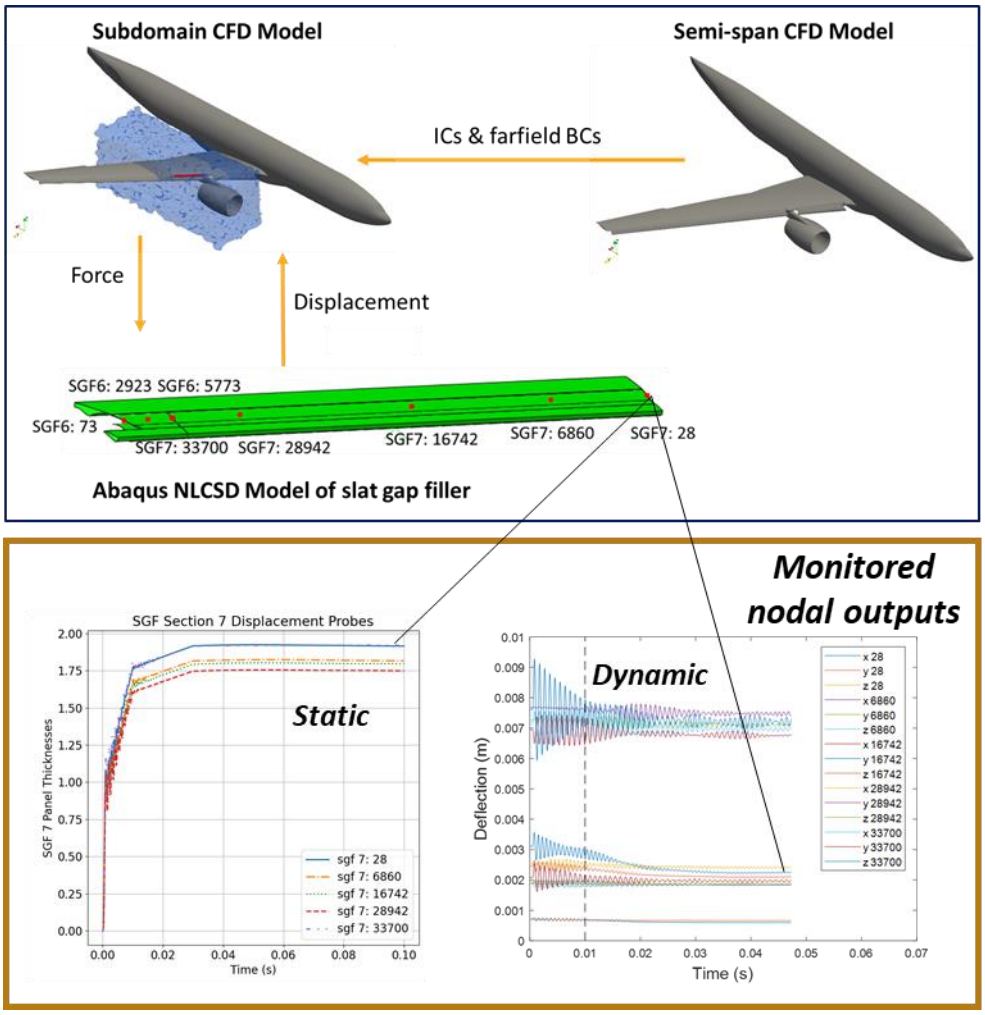

Figure 1-2. The 3D SGF cosimulation between CFD and NL-CSD (top) produced nodal quasistatic (lower left) and dynamic (lower right) panel deformations that were monitored at specific nodes to assess stability.

Successful FSI cosimulation required careful attention to mesh generation in the region of SGF technology element. As part of this work, several capabilities were developed to facilitate the analysis process, including the following: (1) the `scalePut` module for Loci/CHEM to apply boundary and initial conditions with increased dynamic pressure but constant Mach and Reynolds numbers, (2) the `CellqualityRemediation` module for solution stabilization on meshes with locally marginal or poor quality control volumes, and (3) a force-ramping capability to prevent quasistatic deformation overshoots and stably converge to the solution without negative volumes from mesh morphing.

2 TECHNICAL APPROACH

This section provides an overview of the technical approach taken in execution of the project. The content from this section borrows from, and updates, the project proposal documents which evolved over a number of revisions between fiscal years (FYs) 2016 (FY16) and 2021 (FY21). The next three subsections present a statement of technical objectives, a detailed description of the work scope in terms of tasks, and a summary of the main deliverables.

2.1. Technical Objective

The objective of this work was development of computational models and analysis of the coupled fluid-structure response of SCF and SGF noise treatments applied to the leading-edge-slat component of a high-lift system typical of modern transport aircraft. The representative airframe chosen for this work was NASA's CRM-HL in a baseline high-lift configuration. The operational condition corresponds to the approach and landing phases of flight: nominally Mach 0.2 at sea level at 8° angle of attack. The designs needed to be assessed for static aeroelastic deflection as well as dynamic aeroelastic stability. Acceptable response characteristics for both types were necessary if the treatments were to be tested in the 14×22 tunnel and for similar treatments to be considered practical for further study in flight tests or implementation in a service environment.

FSI is a concern for these components, and the analyses in this project had the goal of determining whether the SCF and SGF designs under consideration possess static and dynamic aeroelastic stability. In general, FSI becomes important when a product or structure is operating in a fluid flow regime where the flow affects or excites a structural response or vice versa. In this situation, the flow and the displacement of the structure are not independent problems but are intimately tied together. As a result, a coupled solution is necessary to understand the performance of the system under fluid and mechanical loading. Often, these problems must consider thermal loading as well.

2.2. Work Scope

The overall scope of this effort was conceived in FY16 under the TEAMS2 contract vehicle (FY16–FY17) and continued under the TEAMS3 contract during FY18 through FY21. The tasks were directed toward the development of computational models and analysis of the coupled fluid-structure response of SCF and SGF noise treatments applied to the leading-edge-slat component of a high-lift system typical of modern transport aircraft. Specific run cases that were chosen evolved over the course of the project, both as different technology elements were prioritized or deprioritized and as technical insights were gained about the fundamental FSI response of different elements.

The representative airframe chosen for this work was the CRM-HL in the baseline high-lift configuration (see Figure 2-1), and the conditions of interest corresponded to the approach and landing phases of flight: nominally Mach 0.2 at sea level with an 8° angle of attack. Baseline geometry for the CRM-HL was developed by Boeing under contract with NASA. The baseline CRM-HL is an open geometry that is

freely distributed, and that baseline geometry was modified as necessary by NASA LaRC to include brackets/actuators and kinematics for the high-lift devices. Geometries for the SCF and SGF noise treatments to the slat were developed by NASA LaRC.

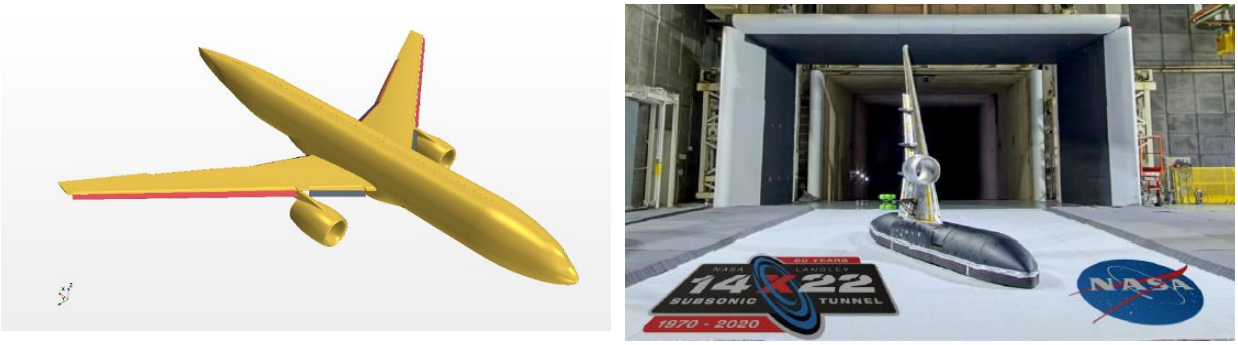

Figure 2-1. CAD representation of baseline CRM-HL geometry (left) and half-symmetry model testing in NASA LaRC 14×22 tunnel (right, from Lockard et al. [22]).

The original (FY16) plan entailed first analyzing specific cases of interest for the wind tunnel test of the CRM-HL via an existing FSI framework consisting of Loci/CHEM and Abaqus coupled through the SIMULIA Co-Simulation Engine (CSE). This was to be followed by development of an FSI framework and capability similar to that for the existing Loci/CHEM–Abaqus framework, via coupling FUN3D to Abaqus also through the use of the SIMULIA CSE. However, geometry and structural design details for the SCF and SGF treatments were not available at the start of the contract (in FY16), so the order of execution of these tasks was reversed.

For this reason, subtask I pertained to the development of the FUN3D-to-Abaqus coupling capability, whereas subtask II pertained to Loci/CHEM–Abaqus engineering simulations in support of the two technology elements. Details of each subtask are described in subsections 2.2.1 and 2.2.2. The outcome of subtask I is documented in a separate report that was originally delivered as a user guide in FY16 and revised in FY21 [23]. The outcomes of subtask II are documented in the current report.

2.2.1. Subtask I

The first phase of work (subtask I) was development of an FSI framework and capability similar to that existing for Loci/CHEM and Abaqus through the use of the SIMULIA CSE with FUN3D as the CFD solver. Abaqus was the target structural solver because of its many advanced simulation capabilities and because of its ubiquitous and widely used user material (UMAT) and user element (UEL) capabilities, which have made it the dominant code worldwide for simulation of smart materials in general and SMAs in particular. The final goal of subtask I was to achieve a strongly coupled FUN3D–Abaqus FSI analysis framework analogous to the Loci/CHEM–Abaqus framework that would be able to achieve acceptable agreement between the FUN3D–Abaqus and the Loci/CHEM–Abaqus results for the set of baseline cases described in subtask II.

The original requirements for the development of the FSI framework (subtask I) included the following list. The consultation and training requirements were generally not completed due to the prioritization of more engineering analysis work in subtask II. Status updates associated with each requirement are given where appropriate.

List and Status of Subtask I Requirements

1. Modification of FUN3D so that it is strongly coupled with Abaqus through the SIMULIA CSE. ATA delivered the modifications of the FUN3D source code on May 31, 2016; therefore, this requirement was achieved in FY16.
2. Modification of the FUN3D–Abaqus framework to be compatible with FUNtoFEM to facilitate maintenance of the FSI framework through updates to FUN3D and/or Abaqus. This requirement was not addressed as part of the current project; however, as part of ATA's support to a NASA University Leadership Initiative (ULI) project, ATA developed a Python-based FUNtoFEM wrapper for the SIMULIA CSE, effectively addressing this requirement. ATA has offered to deliver this code to NASA as a ULI deliverable.
3. Consultation to NASA LaRC on development of FSI models and generation of results via the FUN3D–Abaqus framework for the benchmark cases described for subtask II and on comparison of the FUN3D–Abaqus results to those obtained in subtask II via the Loci/CHEM–Abaqus framework. This consultation was ultimately scaled back during the course of the project as NASA performed FSI using ATA's code modifications for only the Isogai airfoil verification case [24] and not for any SGF or SCF configurations.
4. Documentation of the FUN3D–Abaqus FSI framework development, including notes on maintenance of the capability through FUN3D and Abaqus updates. ATA initially delivered the documentation [23] for the FUN3D–Abaqus FSI framework on May 31, 2016, with updates to this documentation delivered in this project.
5. One day of training on site at LaRC on how to perform FSI analyses with FUN3D–Abaqus. This training was ultimately descoped during the course of this project.
6. One day of web-based training on how to perform FSI analyses with FUN3D–Abaqus. This training was ultimately descoped during the course of this project.

2.2.2. Subtask II

The second phase of work (subtask II) made use of the FSI analysis framework consisting of the Loci/CHEM CFD code coupled to the Abaqus structural code via the SIMULIA CSE. These tools and associated methods are described in section 3 of this report.

ATA set up and executed specific cases that substantiated the suitability and safety of testing flexible SCF and SGF structural noise treatments on a 10% scale, semispan model of the CRM-HL in the 14×22 tunnel. ATA's documentation also instructs NASA LaRC collaborators on the assemblage and execution

of these specific cases for full understanding of the results and to afford the opportunity to duplicate the results using the new FUN3D–Abaqus FSI analysis framework. Specific requirements for subtask II are listed below, with relevant report sections noted.

List and Status of Subtask II Requirements

1. Demonstration of the capability to obtain 2D (or pseudo-2D with a short spanwise section) coupled FSI solutions before attempting 3D calculations. Modeling and analyses began with uncoupled solutions using 3D models (independent calculations with Loci/CHEM and Abaqus that did not employ the SIMULIA CSE) to assess data transfer between the CFD and finite element model (FEM) meshes, robustness of the CFD mesh deformation, etc. This capability is documented in section 4 of this report.
2. A CFD grid study demonstrating that a grid can be defined for the semispan CRM-HL model with an embedded subdomain that encompasses a region of interest in the model for FSI analysis. The ability to impose suitable aerodynamic and structural boundary conditions on the subdomain and the ability to achieve CFD and FSI computations of acceptable fidelity and accuracy were demonstrated. This capability is documented in section 5 of this report.
3. Performance of strongly coupled (e.g., direct information transfer in memory, not through files) 3D FSI simulations of two configurations with the slat fully deployed on the baseline CRM-HL geometry. For these FSI simulations, only the SCF and SGF and small portions of the slat and main wing were treated as elastic; the remainder of the CRM-HL was considered rigid. The slat and main wing portions considered elastic were important in the definition of accurate boundary conditions on the SCF and SGF treatments. The technical approach made use of subdomain models that were extracted from full, semispan CRM-HL meshes (provided by NASA) using, e.g., the `vogcut` subdomain extraction utility in Loci/CHEM. Earlier work, performed substantially in FY18, indicated that there can be grid-related sensitivity to the convergence of subdomain models; however, using the `vogcut` subdomain extraction capability was shown to provide acceptable results on selected cases and was concluded to be the most promising approach. Results for this task are documented in section 6 of this report.

   The configurations analyzed included the following (based on one of the cases having been replaced by a breakout model study performed in FY18):

   - A 2D optimized SCF design consisting of an SMA segment (nominally 0.005″ thick) integrated with the slat via a hinge at the slat cusp and a hard attachment near the slat trailing edge. The design was evaluated with flow representative of the nominal landing condition: Mach 0.2, standard sea level conditions, and an angle of attack of 8°.
   - 2D and 3D representations of SGF design 1, consisting of a single SMA segment (nominally 0.003″ thick) with the inflected, short-chord geometry, attachment constraints to the main wing, and actuator constraints at the end adjacent to the slat trailing edge. The design was stress-free in

the deployed (test) configuration and was evaluated with flow representative of the nominal landing condition: Mach 0.2, standard sea level conditions, and an angle of attack of 8°.

Each analysis consisted of the following:

- Coupled analysis for static deflection under aero load at the nominal landing condition.
- Static stability analysis, i.e., static divergence assessment: ATA performed a series of static aeroelastic simulations at the nominal landing condition while varying flow parameters to increase the aerodynamic loading from the nominal to examine the structural response and determine whether the aerodynamic loading is causing the SCF/SGF to buckle or exceed the specified deformation criteria.
- Dynamic stability analysis, i.e., flutter assessment: ATA performed a series of unsteady coupled simulations at the nominal landing condition while varying flow parameters to increase the aerodynamic loading from the nominal to assess the dynamic response of the SCF/SGF structure. These simulations included a perturbation of the structure at a given flow condition.

4. One day of training on site at LaRC on how to perform FSI analyses with Loci/CHEM–Abaqus. This training was ultimately not required and was deprioritized during the course of this project.
5. One day of web-based training on how to perform FSI analyses with Loci/CHEM–Abaqus. This training was ultimately not required and was deprioritized during the course of this project.
6. CAD models for the SCF and SGF were developed and provided by NASA LaRC, along with FEMs and CFD meshes. ATA provided guidance to NASA LaRC to ensure that computational models were suitable for performing FSI cosimulations. ATA also provided CFD meshing support for the 3D SGF per NASA request, and prioritization of this work led to deprioritization of other proposed tasks. Details related to CFD meshing for the 3D SGF case are documented in section 6 of this report.

   ATA performed the required FSI calculations on ATA computational resources and on NASA LaRC computational resources (NASA LaRC's K-cluster).

2.3. Report Organization

The remainder of this report is organized as follows. Section 3 presents the key tools used in the multiphysics analyses, including the CFD and NL-CSD solvers and the SIMULIA CSE. Additionally, this section discusses the Partial Floquet method for the quantification of dynamic stability.

Section 4 discusses two-dimensional aeroelastic analyses for both the SGF and SCF. This section contains both a description of the models and a summary of the key results.

Section 5 introduces the subdomain modeling approach used for the 3D aeroelastic simulations. This was a verification exercise on rigid geometry, which demonstrated that the subdomain extraction and

boundary condition application process available in Loci/CHEM was able to reproduce the flow field obtained from a global CFD simulation in a substantially smaller local subdomain.

Section 6 presents the 3D aeroelastic analysis of the SGF. This section discusses CFD meshing work ATA performed in close discussion with NASA to represent the areas of the flow near the technology elements. This section also highlights important aspects of the Abaqus model creation for FSI modeling and presents the quasistatic and dynamic aeroelastic results. Section 6 also presents the final status of the 3D SCF case, including a description of the Abaqus model and important aspects for CFD cosimulations, which can be used in future 3D aeroelastic analysis of that configuration.

Finally, section 7 summarizes the overall effort, including important outcomes, lessons learned, and recommendations for future work.

3 TOOLS AND METHODS

This section describes the key tools and analysis methods used in the FSI analyses of the SCF and SGF test cases. These tools include the CFD code (Loci/CHEM), the NL-CSD code (Abaqus), the coupling tools (ATA's multiphysics engine, which uses the SIMULIA CSE), and the Partial Floquet stability analysis technique.

The general stability analysis procedure used here consisted of the following steps:

1. Perform a rigid CFD simulation to establish initial flow conditions.
2. Perform a quasistatic cosimulation to determine the static equilibrium state. The CFD simulation part of the cosimulation was restarted from the converged, rigid simulation.
3. Perform a dynamic or time-accurate cosimulation where the CFD code and the NL-CSD code use the same time step and communicate in lockstep. The dynamic cosimulation was restarted from the quasistatic solution and perturbed using impulsive load applied to the structure.
4. Perform system identification to estimate aeroelastic damping and frequencies. Structural responses—in this case displacements—from key locations are used as input to the Partial Floquet method to assess stability.

3.1. Multiphysics Framework Coupling Overview

The FSI simulations are conducted using a loosely coupled but tightly integrated domain decomposition approach in which the governing equations of motion for each domain (e.g., the fluid domain and the solid or structural domain) are solved using approaches or methods most appropriate for that domain. A computational fluid dynamics (CFD) code is integrated with an NL-CSD code in a fully coupled FSI form[1]. The CFD code is Loci/CHEM, and the NL-CSD code is Abaqus. The coupling of the CFD and NL-CSD codes is done using SIMULIA's CSE, which is a software infrastructure that enables cosimulation and consists of a runtime environment that synchronizes the applications, coordinates communication, and provides spatial and temporal mapping at the domain interface.

3.1.1. Loci/CHEM

The Loci/CHEM (CHEM) code is a modern multiphysics simulation code that is capable of modeling chemically reacting multiphase high- and low-speed flows. CHEM, developed at Mississippi State University, is written in the Loci framework and is a Reynolds-averaged Navier–Stokes (RANS) finite-volume flow solver with a cell-centered control volume definition for generalized or arbitrary polyhedra [25]. CHEM uses density-based algorithms and employs high-resolution approximate Riemann solvers to solve finite-rate chemically reacting viscous turbulent flows [26]. CHEM contains a robust grid

[1] More information about the FSI coupling framework developed by ATA can be obtained by contacting ATA Engineering, Inc. at https://www.ata-e.com/ or https://github.com/ATAEngineering.

deformation capability allowing it to morph the high-aspect-ratio cells common in viscous flow simulations to resolve boundary layers, without sacrificing cell quality.

CHEM CFD solutions were initially verified against similar solutions that NASA had generated using FUN3D.

The next subsections (sections 3.1.1.1 through 3.1.1.6) document several key features of CHEM that were used in the FSI analysis project, as well as efforts undertaken to verify CHEM against comparable tools.

3.1.1.1. Use of `vogcut` for Subdomain Modeling

The majority of the computational expense in the coupled simulations is due to the high cost of CFD simulation. This expense can be lowered by reducing the size of the computational domain using the `vogcut` tool that comes with Loci/CHEM: `vogcut` extracts a subset of cells from a larger volume mesh, enabling the extracted mesh to be run as a breakout model, reducing the computational cost. The subsets to be extracted can be defined with simple geometric shapes (e.g., box, sphere, cone), and all cells with their cell center inside the defined shape will be extracted.

3.1.1.2. `scalablePrescribed` for Boundary Condition Interpolation

Loci/CHEM has a boundary condition interpolation feature known as `scalablePrescribed`. This feature allows the flow solution from a completed simulation using one mesh to be interpolated onto a second mesh for use with boundary condition specification. This feature works by interpolating the solution from one cloud of points representing the first mesh to a second cloud of points representing the second mesh. It may be used for cases where the second mesh is a subset of the first, like when `vogcut` is used to generate the second mesh, or for cases where the two meshes are entirely dissimilar. This capability is applicable to both inflow and outflow boundary conditions, and only the solution data required for the respective boundary condition is used.

3.1.1.3. `cellQualityRemediation`

Poor-quality cells can hinder or inhibit solution convergence. This can be the case even when the regions of poor quality are extremely localized to a small number of cells and do not affect global quantities of interest such as pressure coefficient or skin friction. As part of the flux calculation for each cell face in the mesh, solution values at cell faces must be reconstructed from solution values at cell centers (since Loci/CHEM is a cell-centered finite volume code). To achieve second-order spatial accuracy, a first-order reconstruction is used. However, Loci/CHEM will by default use a more robust, but less accurate, zeroth-order reconstruction in regions of poor cell quality. This zeroth-order reconstruction results in locally first-order spatial accuracy. Loci/CHEM determines that cells with a volume ratio greater than 50, or a maximum included angle of greater than 150 degrees, constitute poor-quality regions. If the region of poor cell quality is small, this local change can have a negligible effect on solution accuracy while greatly helping robustness.

The `cellQualityRemediation` module [27] extends this functionality in Loci/CHEM by allowing the user to decrease the volume ratio and maximum included angle thresholds from their default values, thereby designating more cells as being of poor quality. It also allows the user to include all cells adjacent to those flagged by the volume ratio and maximum angle thresholds in the poor-quality-cell region.

3.1.1.4. Scaling of `put` Files

During simulation, Loci/CHEM can generate `put` files, which are a collection of solution variables (pressure, velocity, temperature, turbulent kinetic energy, turbulent eddy viscosity, and species mass fractions) located at the cell centers of the CFD mesh. These files can be used to initialize a subsequent simulation and/or to provide boundary conditions via the `scalablePrescribed` functionality. The `scalePut` utility [28] allows the user to scale any of the variables in the `put` file by a user-defined amount. This utility can be used to scale the pressure of a flow solution while keeping the Reynolds number constant, which saves CPU time because the `put` file can be instantly scaled to the alternate pressure condition without the need to rerun the analysis. As long as the Reynolds number is held constant, the character of the flow remains the same and the scaling approach is valid. Figure 3-1 shows the pressure coefficient over the CRM-HL wing section for the baseline flow condition as well as various higher-dynamic-pressure cases. The pressure coefficient distribution is roughly the same for all cases because the Reynolds number is held constant.

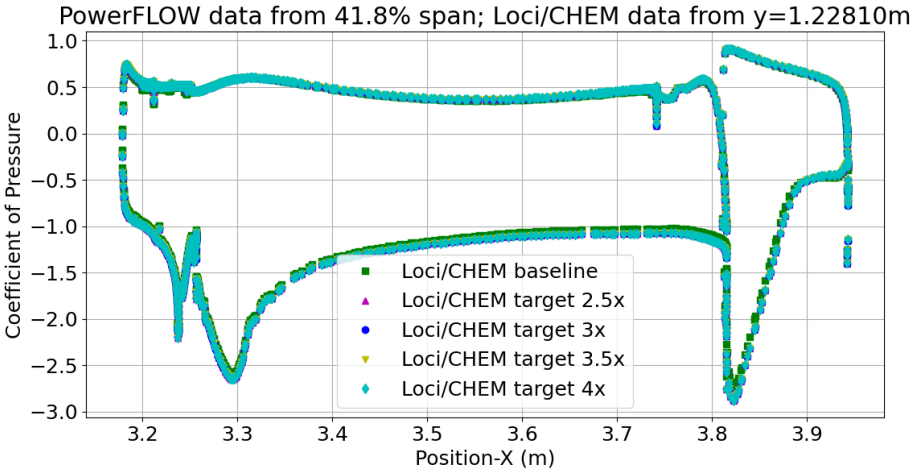

Figure 3-1. Pressure coefficient for a variety of alternate flow conditions.

3.1.1.5. Force Ramping for Quasi-Steady FSI

Cosimulations may involve initial transients as the aerodynamic loading is applied to the FEM. These transients are a result of the loading instantly being “switched on” and can result in a high-amplitude oscillatory displacement response. This transient behavior is neither physically accurate nor important to the analysis when obtaining a quasi-steady result is the goal. Furthermore, it can cause stability problems for the simulation if the transient deformations result in very poor-quality cells or even negative cell volumes. To alleviate these issues, the aerodynamic force coming from Loci/CHEM can be ramped from

0% of the actual value to 100% of the actual value over a user-specified number of iterations. For example, if the user specifies the force to be ramped over the first 100 iterations, the forces are linearly ramped over this period. In this situation, 0% of the force will be applied initially, 10% of the force at iteration 10, and so on until 100% of the force is applied at iteration 100. This has the effect of critically damping or overdamping the transient oscillations, which helps with stability. It has no effect on the end result as long as 100% of the actual force is applied by the end of the simulation.

3.1.1.6. Work Output Monitoring

In a dynamic aeroelastic analysis, a stable oscillation will put work into the flow, whereas an unstable one will extract work from the flow. Therefore, tracking the work done by or to the fluid can be informative for assessing stability. In support of this methodology for assessing stability, a new feature was developed for Loci/CHEM's `abaqus` module (which enables cosimulation) to track the work done by the fluid. This new feature was tested on the canonical Isogai airfoil, which can be made to be stable or unstable by varying the dynamic pressure. Figure 3-2 shows that the work done will have opposite signs for the stable and unstable cases.

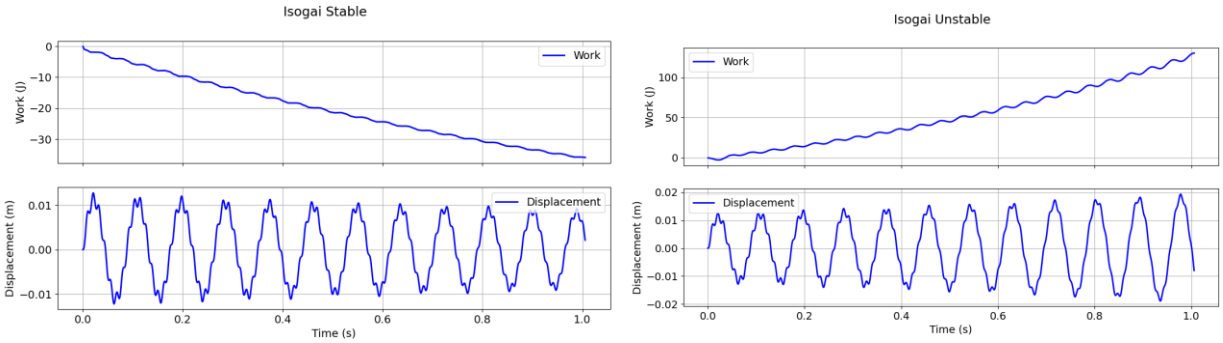

Figure 3-2. Work and displacement time histories for a stable (left) and unstable (right) Isogai airfoil test case.

This capability was developed and verified for the Isogai airfoil, which was simulated without a turbulence model. It was discovered near the end of the project that the work output module had a conflict when used with turbulence models, and given the time and resources available at that point, ATA was unable to monitor work output for any of the 2D or 3D FSI cases for the SGF or SCF.

3.1.2. Abaqus

The Abaqus Unified finite element analysis (FEA) product suite offers powerful and complete solutions for both routine and sophisticated engineering problems encompassing a vast spectrum of industrial and aerospace applications. Abaqus/Standard employs solution technology ideal for static and low-speed dynamic events where highly accurate stress solutions are critically important. Abaqus/Explicit is an FEA product that is particularly well suited to simulating brief transient dynamic events. Both products provide nonlinear analysis supporting nonlinear materials, large displacements and finite rotations, and general contact. The solvers are parallel and run on shared memory and distributed memory clusters.

3.1.3. SIMULIA CSE

As described previously [29], [30], the coupled multiphysics analysis framework utilizes SIMULIA's CSE, which is a software component for enabling cosimulation between discrete and continuum systems. The CSE consists of several software components. The controller is a customizable state machine in charge of marshaling the client programs (the "clients"), negotiating time, and ensuring that the physics fields are exchanged in a synchronized manner. The mapper provides both spatial and temporal mapping of the physics fields. Spatial mapping includes surface-based coupling in which the domains remain distinct and are coupled through a common surface, and volume-based mapping in which the physical domains overlap. The spatial mapping between the domains is conservative; for example, Loci/CHEM may provide traction vectors on face centers that are then interpolated to nodal forces on a quadrilateral interface mesh on the Abaqus model. Furthermore, the CFD model may use a much finer discretization than does the computational structural dynamics (CSD) model, which is more appropriate to its numerical treatment. The only restriction is that the two models need to be colocated in space. The communication layer is in charge of communicating both control and data messages between the controller and the clients. Finally, the application programming interface (API) layer provides the interface to the clients.

3.1.4. Fluid-Structure Analysis Framework

In a domain-decomposition approach to a coupled aerothermostructural interaction problem, the CFD code will compute the traction (or force) vector and the heat flux at the fluid-structure interface (or wetted surface), and the NL-CSD code will compute the resulting temperatures and structural displacements due to those loads. The overall procedure for a typical problem is illustrated in Figure 3-3.

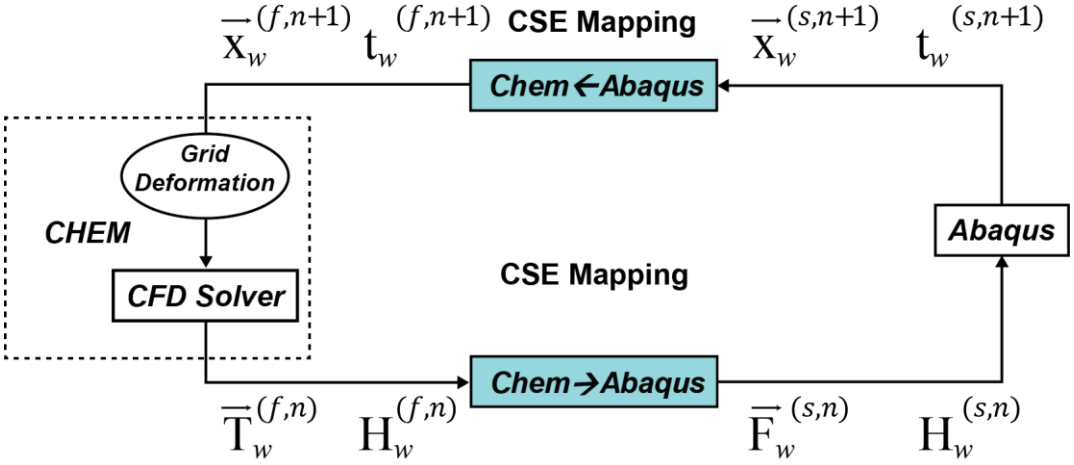

Figure 3-3. Fluid-structure cosimulation data-mapping procedure.

Since a domain-decomposition approach is utilized, each domain will be suitably discretized for the appropriate solver. As such, the interfaces of the two domains will almost certainly be discretized differently and will not match, so the data transfer at the wetted surface interface is done using the mapping tools provided by the CSE, and once the CFD wetted surface is moved to conform to the structural deformations, the flow solver adapts the volume grid to reflect these changes [31]. Note that for

the SCF and SGF cosimulations in the present work, the thermal quantities—heat flux and temperature—were not exchanged as part of the cosimulations.

In previous work, the computational framework discussed above was applied to validation problems of static and dynamic aeroelasticity in high subsonic and mildly supersonic flows [30] and to a hypersonic aerothermoelastic case [32]. Additionally, the techniques have been applied to supersonic nozzle flows to estimate the side loads on the nozzle of the Space Shuttle Main Engine [33] and to assess the stability of landing gear doors for a business jet [34].

3.1.4.1. Temporal Coupling

The code coupling may be exercised with the following temporal coupling options.

- First-order (one data exchange per time step)
- Higher-order (multiple data exchanges per time step)
- Subcycling (multiple time steps per data exchange)

Figure 3-4(a) illustrates the first-order data exchange scheme. It is so named because, with sufficiently small time steps, it will ultimately reduce the temporal order of accuracy to at most first order, regardless of the temporal accuracy applied to the constituent codes. Data are exchanged after the completion of full time steps within the codes. The codes may employ unequal numbers of subiterations to get to the next time step. For this scheme to work, it is essential that each code apply the same time step.

A higher-order data exchange scheme is pictured in Figure 3-4(b). This scheme applies subiteration convergence to the exchanged data to ensure that consistent data are shared between codes. This convergence removes the first-order lag error of the first-order schemes, and ultimately the order of accuracy of the coupled solution will be that of the lowest temporal accuracy of the constituent codes. As CHEM and Abaqus are both at most second-order accurate, second-order temporal accuracy is the most that can be attained in this application.

When considering problems with disparate time scales, it is possible to exercise a subcycling scheme, as portrayed in Figure 3-4(c). Here, the constituent codes must line up at integer values of time steps, although the numbers of time steps the codes take before data exchange need not be the same. For example, in the figure, CHEM advances four time steps and Abaqus advances just two before data are exchanged. This option may be appropriate when the fluid response is much faster than the recession and the computational expense of running Abaqus in lockstep with CHEM needs to be mitigated.

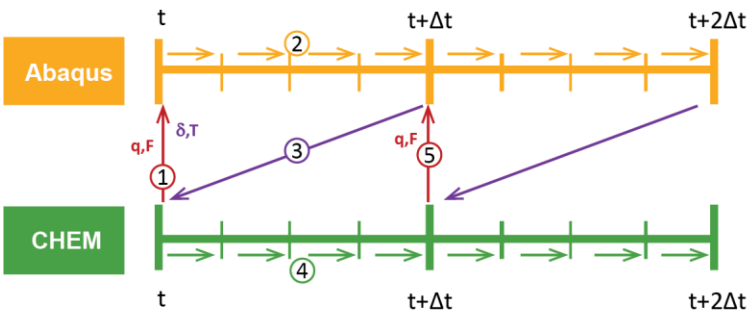

(a)

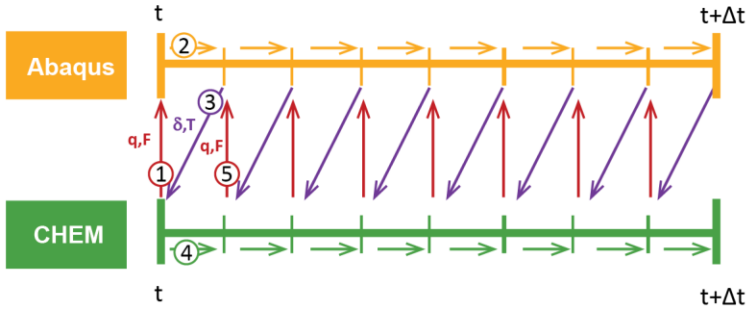

(b)

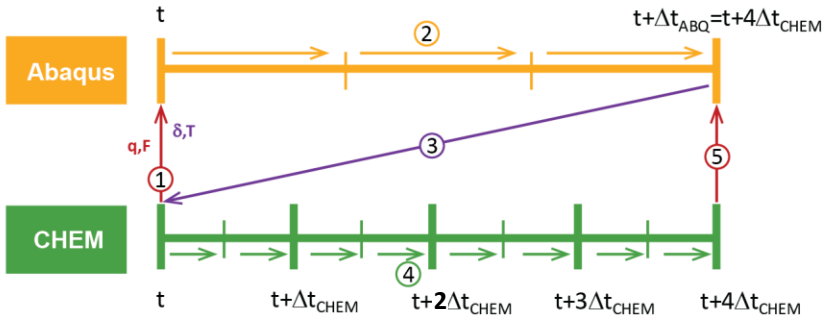

(c)

Figure 3-4. Temporal coupling schemes: (a) first-order data exchange, (b) higher-order data exchange scheme, and (c) subcycling coupling scheme.

3.2. Stability Analysis

Aeroelastic stability is determined via analysis of the time-domain responses after the system has undergone a perturbation designed to excite the modes of interest. A system with positive damping will dissipate energy and return to a near-static-equilibrium configuration, whereas a system with negative effective damping due to aerodynamic forcing will exhibit exponential growth in the response amplitude. Obviously, exponential growth cannot be sustained in nature, and in physical systems this condition ultimately results in the failure of the structure or the presence of an LCO.

In the simplest form, assessing the stability of the system may be done by visual inspection of the response and whether the response envelope grows or decays in time. Growth identified by this simple inspection indicates an instability, and decay represents positive stability. Unfortunately, visual inspection is only able to determine stability as a binary state: either it is stable or it is not; the visual inspection approach is unable to either characterize the degree of an instability (i.e., quantify damping) or ascertain how individual structural modes contribute to stability. Additionally, when the response is complicated with appreciable contribution from many structural modes, an accurate visual inspection may require an undesirable length of simulated time to observe adequate growth or decay in the response, leading to excessive computational expense.

Quantitative assessments of stability that provide estimates of component frequencies and damping values are possible through a variety of computational methods, e.g., Prony's method, the eigenvalue realization algorithm (ERA), fast Fourier transform (FFT)-based methods, and Partial Floquet. Partial Floquet [35] was developed for the more complex case of linear time-periodic systems but works equally well in analysis of linear time-independent systems and was therefore selected to provide the quantitative stability assessments here, due to ATA's extensive familiarity with the methodology.

One challenge with quantitative stability assessment occurs when the structural stiffness is nonconstant, as can occur with nonlinear material properties. This is particularly relevant to SMAs, where modal frequencies are noticeably a function of deformations. To better understand how variable stiffness may affect the stability results, a simple ordinary differential equation (ODE) with a mild nonlinearity is considered:

$$\frac{\partial^2 y}{\partial t^2} + 2\omega(t)\zeta\frac{\partial y}{\partial t} + \omega^2(t)y = 0 \tag{1}$$

$$\omega(t) \equiv 2\pi\left(\Omega_{base} + \Omega_{vary}y\right). \tag{2}$$

Figure 3-5 provides an example solution of the above ODE with a base frequency of 1084 Hz and 100 Hz per unit for the varying nonlinear component, and a damping ratio of 0.1%.

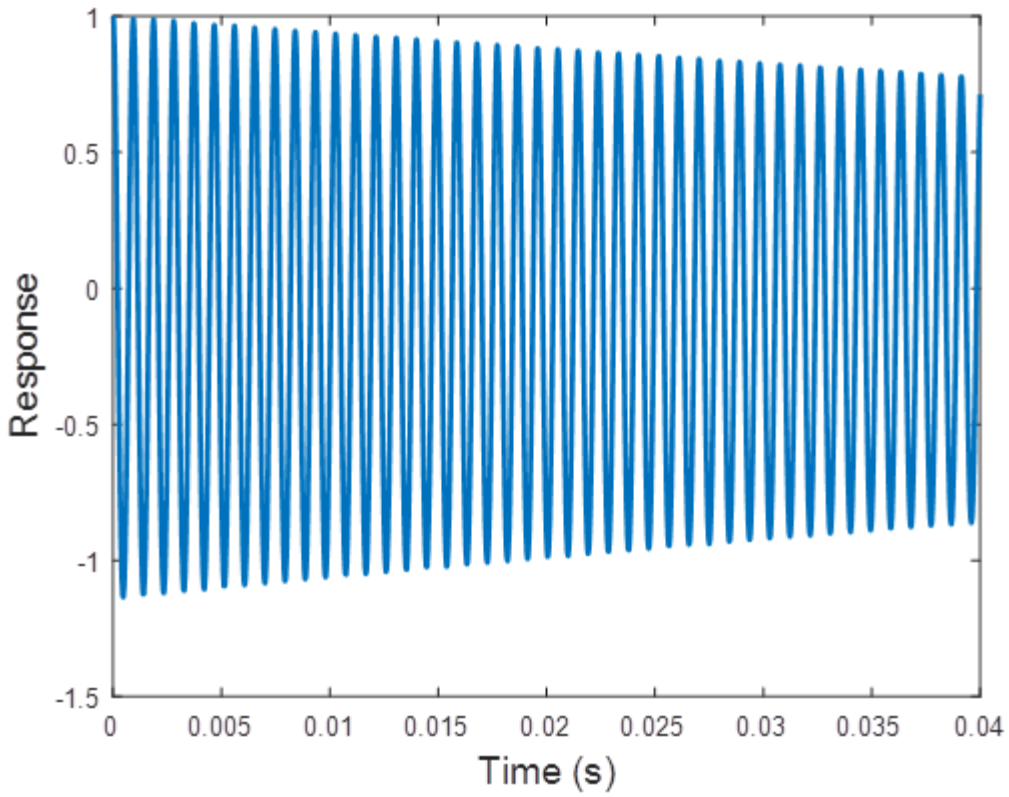

Figure 3-5. Example of response with mild frequency nonlinearity.

Table 3-1 provides the identified results if three modes are extracted. Two frequencies are identified in the region of the base frequency, one with positive stability and the other with negative stability and neither at the prescribed 0.1% damping ratio. It is noted that, despite the detected "instability" with a damping ratio of −0.12%, it is evident that this simple ODE is stable. Additionally, it is noted that the harmonic appears at the correct damping ratio; however, when other modes are present the harmonic is expected to be at a much lower energy level and may not be detected. The takeaway from this simple example is that if a mode with nonconstant frequency is detected as multiple frequencies, the damping ratios may not themselves be accurate.

Table 3-1. Partial Floquet results for three modes.

(Detected Frequency)/(Base Frequency)	**Damping Ratio**
1.9802	0.10%
0.9934	−0.12%
0.9871	0.51%

4 2D AEROELASTIC ANALYSIS

This section discusses the 2D aeroelastic assessments performed for the SGF and SCF.

4.1. Slat-Gap Filler

This section discusses the 2D SGF analyses.

4.1.1. 2D SGF CFD Model

The 2D unstructured, viscous mesh was provided by NASA LaRC and contained 350,955 nodes and 697,151 elements. The CFD domain includes the slat, SGF, and the main wing. Previously, NASA LaRC had performed rigid CFD simulations of the 2D SGF configuration using FUN3D. In the current effort, ATA performed a series of CFD simulations to ensure that Loci/CHEM can adequately replicate the previous results attained using FUN3D. The following parameters were used: Roe's adaptive approximate Riemann solver [36] was used to compute the inviscid fluxes with Venkatakrishnan's limiter [37] applied to these terms, Sutherland's law [38] was used to compute the viscous transport quantities, and the Spalart-Allmaras turbulence model [39] was used to close the fluid dynamic system of equations.

4.1.2. 2D SGF FEM

The 2D SGF FEM was provided by NASA LaRC and is shown in Figure 4-1, and it includes the slat, the main wing, and the SGF technology element. The SGF is modeled using general-purpose shell elements (Abaqus element type `S4`), and the slat and main wing are assumed rigid and represented using rigid body elements (Abaqus element types `R3D3` and `R3D4`). The elastic portion of the FEM, shown in Figure 4-2, originally contained 36 nodes and 17 elements and had a unit width in the Z direction. Two nodes (nodes 1301 and 1302) were added to the FEM so that the FEM geometry would be consistent with the CFD surface boundary definitions. A corresponding pair of nodes were added to the Z = 0 symmetry plane along with the necessary shell elements (elements 1301 and 1302, type = `S4`) to connect them to the upstream and downstream nodes.

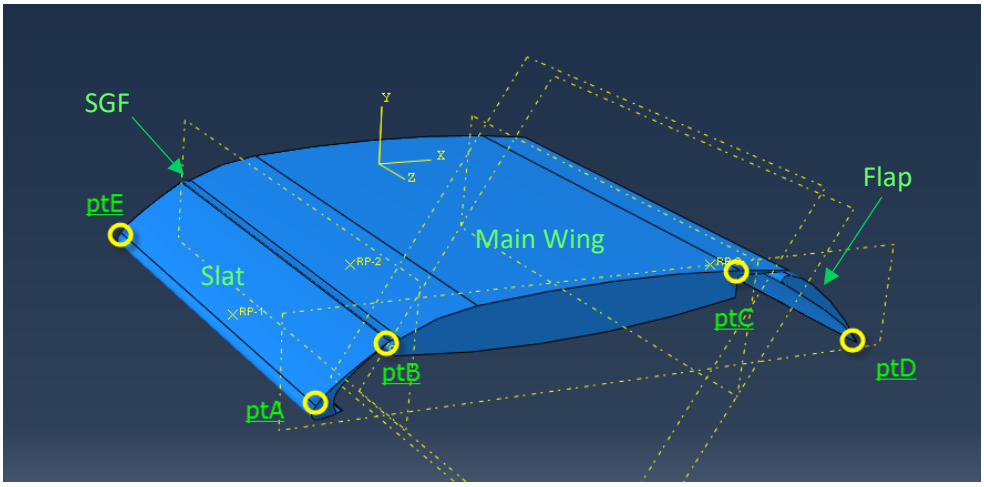

Figure 4-1. Abaqus 2D SGF model.

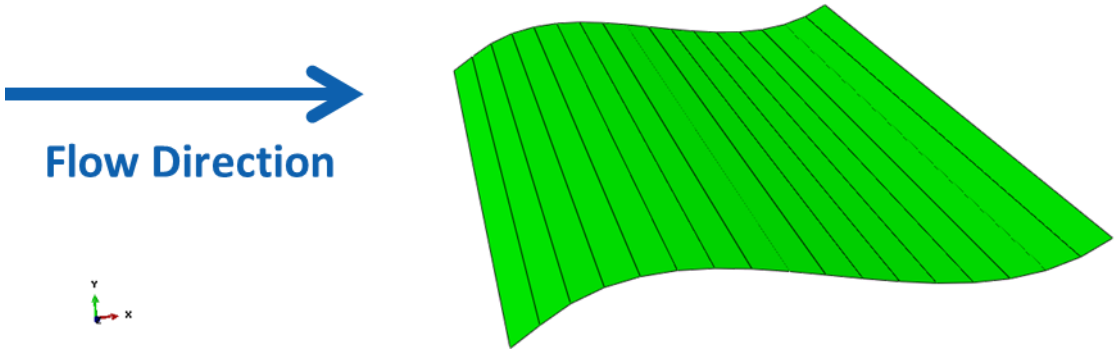

Figure 4-2. SGF FEM using shell elements.

4.1.2.1. FEM Mass Units

Early in the project, there was an apparent confusion over the mass units used in the FEM. The initial natural frequencies ($f_1 \approx 4$ kHz) were deemed to be too high and the structure too stiff. The elastic moduli were well understood, so it was thought that the mass definition was the reason initial natural frequencies were too high. From consultation with NASA LaRC, the mass units were thought to be defined in lbm, so the density units were changed to lbf•s^2/in^4, which lowered the frequencies by a factor of $\sqrt{386}$. Later it was discovered that the initial natural frequencies were correct, so the density was restored to the initial value. A series of quasistatic and dynamic aeroelastic cosimulations were performed using the incorrect density. The density change did not affect the quasistatic results, but it significantly impacted the dynamic cosimulation results since the SGF mass was increased by a factor of 386. For the sake of clarity, the results using the modified, incorrect density are not included except where noted.

4.1.2.2. Wetted Surface Definition

The first step in setting up FSI simulations is to define the wetted surface where the fluid and structure exchange information. Physically, the wetted surface represents the location where the fluid and structure

interact. As mentioned in section 3.1.4, the structural model provides displacements to the CFD solver and receives the aerodynamic force information from the CFD solver. In turn, the CFD solver receives the displacements computed by the NL-CSD code and provides force information to the NL-CSD solver. The CFD wetted surface definition includes surfaces on both the SGF outer mold line (OML) and the inner mold line (IML). As previously mentioned, the SGF FEM is defined using general-purpose shell elements with the nodes located on the OML. However, the shell elements can only be loaded on a single side, i.e., on the face defined using the positive normal vector (`SPOS` in Abaqus terminology) or on the face defined using the negative normal (SNEG). Thus, only the CFD loading computed on either the OML or the IML can be applied. To allow the correct aerodynamic loading to be applied, the FEM was modified to have an additional set of virtual elements on the IML.

A virtual surface was defined on the IML to transfer the loads and displacements on the interior side of the SGF. The geometry for the virtual IML surface was generated by creating a spline using the IML points from the CFD surface mesh, thus ensuring consistent geometry between the CFD model and the FEM. The FEM nodes (on the OML) were projected to the spline to create nodes on the IML, thus creating a spatial one-to-one correspondence between nodes on the OML and the IML that enabled a direct node-to-node connection between the two surfaces. The IML wetted surface definition is shown schematically in Figure 4-3. Note that the node IDs shown are for the nodes on the Z = 1.0 symmetry plane; there is a corresponding set of nodes on the Z = 0.0 symmetry plane. The virtual IML wetted surface elements were defined using `SFM3D4` elements, which have no mass or stiffness associated with them. As such, these elements are nonstructural and are used only to transfer the wetted surface loads and displacements. Therefore, to connect the IML elements to the structural elements on the OML, the IML nodes were rigidly connected to the OML nodes via `*KINEMATIC COUPLING` constraints.

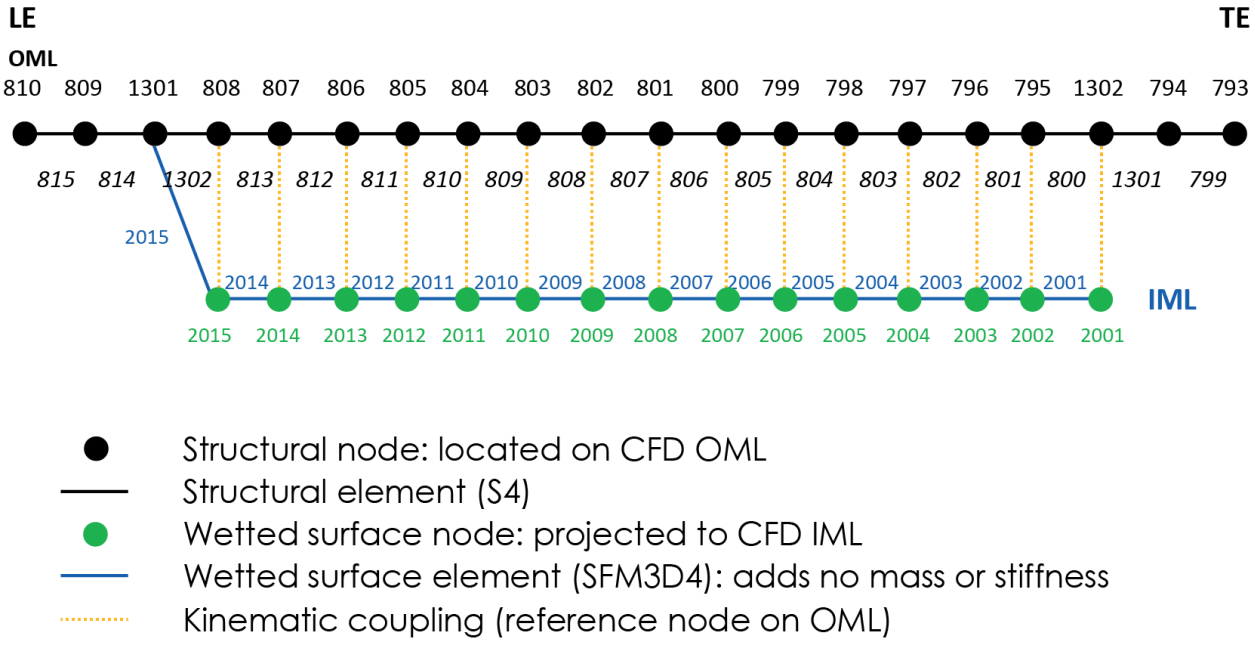

Figure 4-3. Schematic of the 2D SGF wetted surface definition.

The addition of the IML virtual wetted surface was verified to not change the structural response of the SGF FEM via a series of numerical tests. For these analyses, it was assumed that the mechanical restraints for the nodes at the leading and trailing edges were pinned and a Z-symmetry boundary condition was applied to all the nodes. The first test applied a uniform unit pressure load to the original FEM and to the modified FEM with the IML wetted surface. The resulting deformation is shown in Figure 4-4, and a comparison of the horizontal displacement is shown in Figure 4-5(a), with the vertical displacement provided in Figure 4-5(b). The maximum difference in the horizontal component is 1%, with the difference of most nodes being on the order of 0.2%. The maximum difference in the vertical component is 0.8%, with the difference of most nodes being on the order of 0.1%.

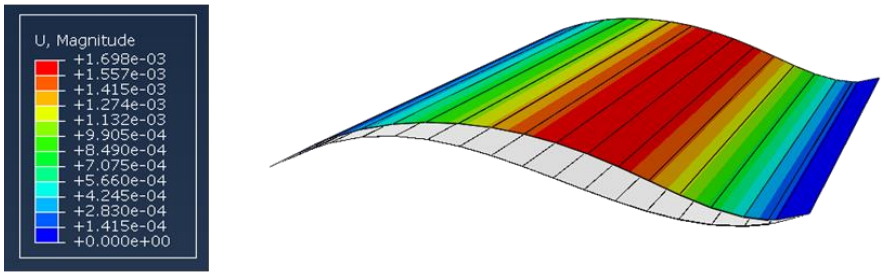

Figure 4-4. SGF displacement contours subject to unit pressure loading.

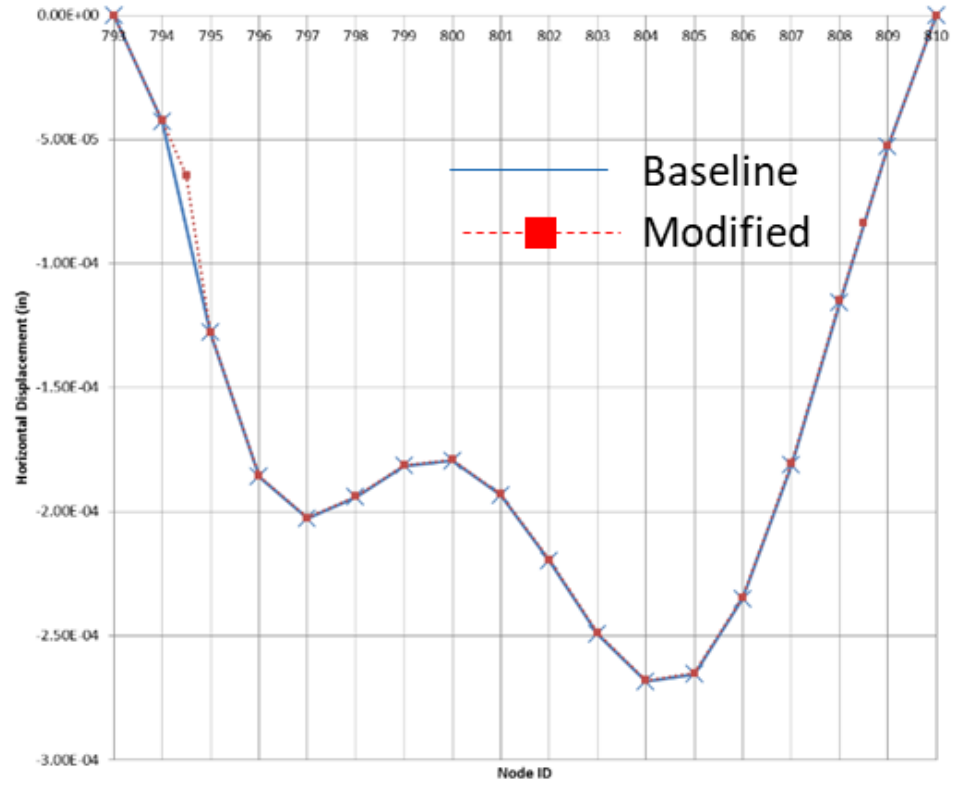

(a)

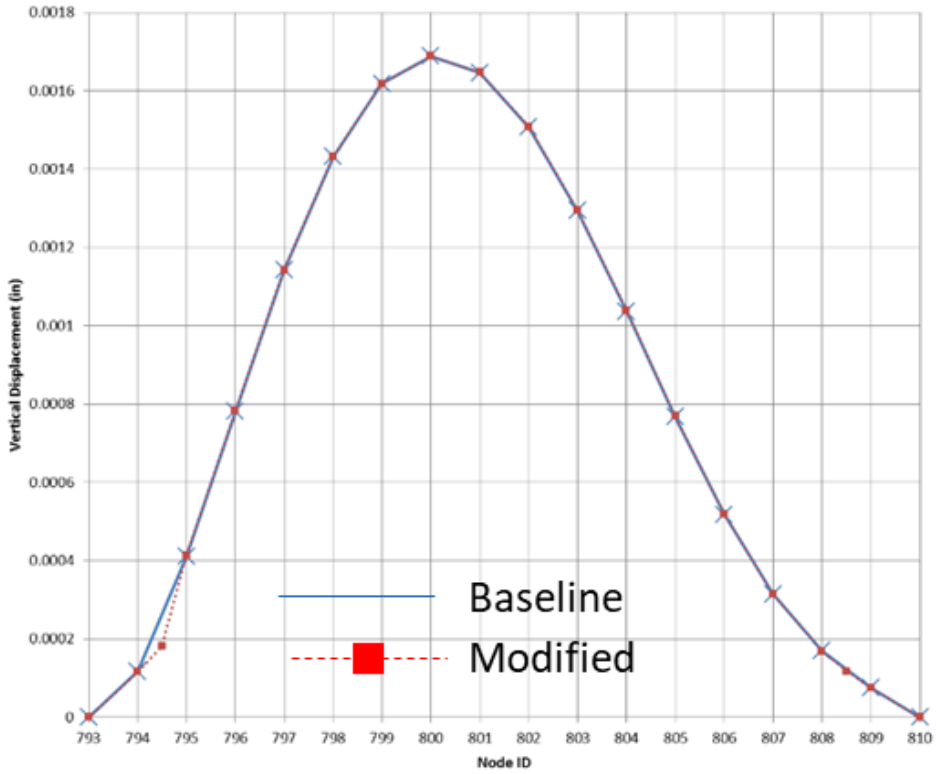

(b)

Figure 4-5. Comparison of the (a) horizontal and (b) vertical displacement components of the baseline and modified SGF FEM subject to a unit pressure loading.

The second test that was performed to ensure that the addition of the virtual IML did not modify the structural dynamics involved comparing the eigensolution for the baseline and modified SGF FEMs. Table 4-1 and Table 4-2 provide comparisons of the natural frequencies and mass, and a comparison of the eigenvectors is shown in Figure 4-6. Due to the small scale in differences between the two FEMs for the static response and eigensolutions cases, it is clear that the virtual IML wetted surface has very little impact on the structural response of the SGF FEM.

Table 4-1. Natural frequency comparison between the original and modified SGF FEMs.

Mode	Baseline (Hz)	Modified (Hz)	% Diff
1	3780.9	3774.7	0.16
2	14832.0	14764.0	0.46
3	26076.0	25842.0	0.90
4	33598.0	33603.0	-0.01

Table 4-2. Mass comparison between the original and modified SGF FEMs.

Baseline ($lb_f s^2/in$)	Modified ($lb_f s^2/in$)	% Diff
3.1665714E-07	3.1669129E-07	0.01

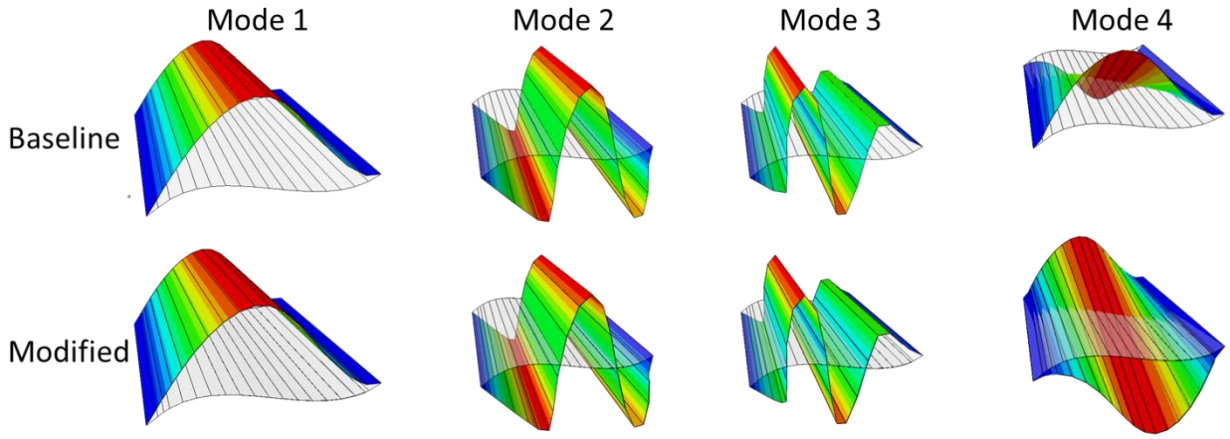

Figure 4-6. Eigenvector comparison between the original and modified SGF FEMs.

4.1.2.3. Wetted Surface Loading

There are two options for exchanging force information between CHEM and Abaqus. The first is to exchange the traction vector, which includes the normal pressure component and has units of pressure, and the second option is to exchange the force vector. As part of the solution of the fluid dynamics equations of motion, CHEM natively computes the surface pressure and shear stress, i.e., the traction vector. The first option essentially maps these pressures from the CFD wetted surface to the FEM wetted surface. Given the pressures mapped to the wetted surface, Abaqus then computes the resulting nodal forces by integrating the pressure over the wetted surface area. However, due to differences in discretization, the area of the CFD wetted surface may not be equal to the area of the FEM wetted surface; therefore, the total force acting on the FEM may not equal the total force acting on the CFD wetted surface. For the force transfer option, CHEM performs the integration to compute the forces acting on the wetted surface, and then these forces are mapped to the FEM wetted surface. This guarantees conservation of force.

Quasistatic cosimulations were performed to examine the difference, if any, between the use of force vectors and traction vectors for cosimulation. The displacement of the center node (node 801) was monitored, and the displacement was 2.6136e-4″ when mapping the nodal force vector and 2.7666e-4″ when mapping the traction vector, resulting in a difference of about 5.8%. Therefore, it was decided that the force transfer option would be used for all the subsequent cosimulations.

A schematic of the FEM wetted surface loading is shown in Figure 4-7. The purple arrows represent the CFD loading on the FEM OML, and the red arrows represent the loading on the FEM IML. The two elements shown with the blue arrows represent part of the FEM that is overlapping the main wing. In the CFD boundary definition, the entire surface of the main wing was grouped as a single CFD surface patch. Thus, these two elements are not cosimulating with the CFD model because there is not a separate CFD boundary surface defined for this small portion of the main wing. To avoid a nonphysical response arising from the situation where part of the FEM wetted surface is loaded and part is not, a pressure loading was manually applied to these two elements in Abaqus as an external load. For each flowfield condition that was simulated, the average pressure computed by the steady, rigid CFD flow field at the location shown in Figure 4-8 was used.

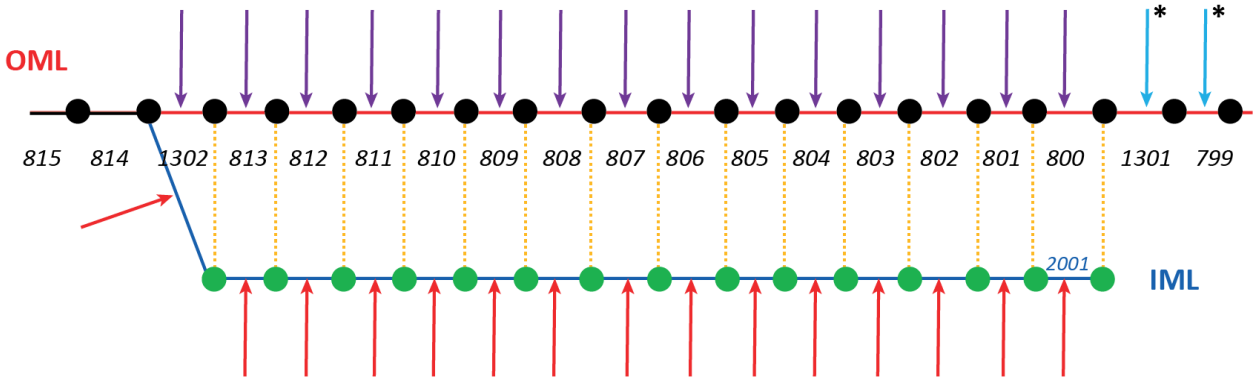

Figure 4-7. Schematic illustrating the SGF FEM wetted surface loading.

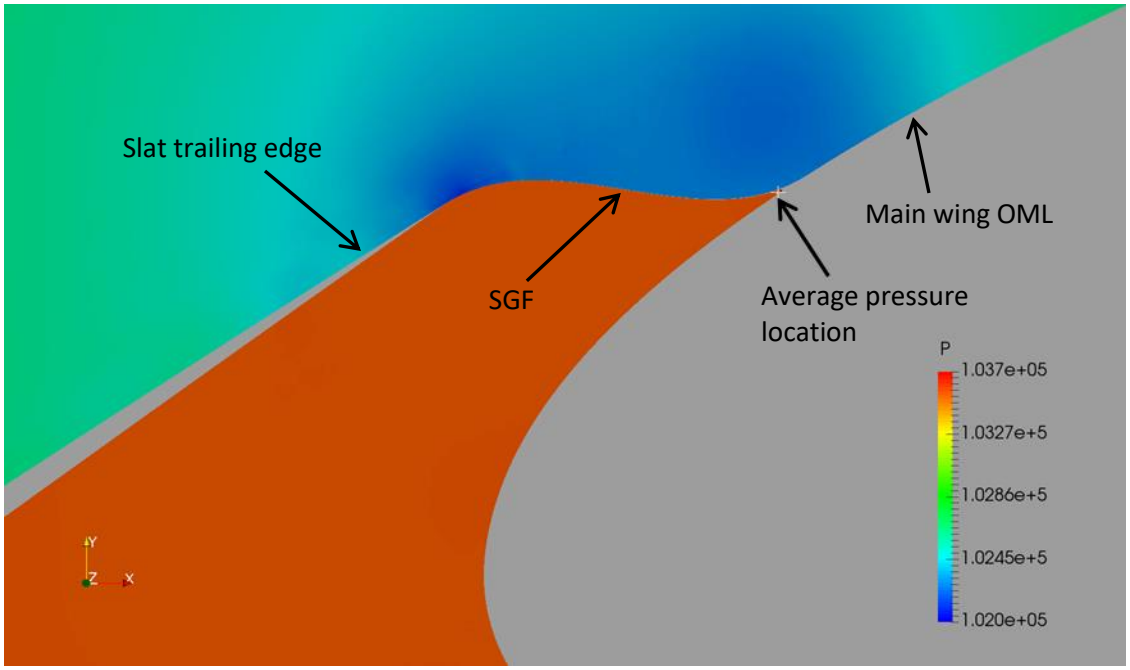

Figure 4-8. Location used to extract the pressure applied to the part of the SGF FEM not cosimulating with the CFD model.

4.1.3. Aeroelastic Stability Analysis Parameters

An aeroelastic parameter study was performed to bound the SGF stability estimates. Five different aeroelastic simulations were performed in which both fluid and structural properties were varied. The dynamic pressure was varied from the nominal value (q = 2884.8 Pa) to six times the nominal value, and Table 4-3 lists the different aerodynamic conditions analyzed. The dynamic pressure was increased by increasing the fluid density and pressure and keeping the velocity and temperature constant. This has the effect of keeping the Mach number fixed and not changing the fluid compressibility effects while increasing the dynamic pressure.

Table 4-3. Conditions analyzed in aeroelastic parameter study.

Mach	0.200	0.200	0.200	0.200	0.200
Q (Pa)	721.20	2884.80	4507.49	11539.19	17308.78
Q_ratio (inc / nom)	0.2500	1.0000	1.5625	4.0000	6.0000
V (m/s)	69.445	69.445	69.445	69.445	69.445
T (K)	300	300	300	300	300
rho (kg/m^3)	0.299093	1.196372	1.869331	4.785488	7.178232
P (Pa)	25757.112	103028.446	160981.947	412113.784	618170.676
a (m/s)	347.224	347.224	347.224	347.224	347.224

In addition to changing the aerodynamic parameters, the stiffness of the SGF was also varied to result in a 50% increase and 50% decrease in the first natural frequency. The decrease in stiffness was achieved by decreasing the modulus, and the increase in stiffness was achieved by increasing the thickness. Figure 4-9 shows a comparison of the first mode shape and frequency for each of the SGF models.

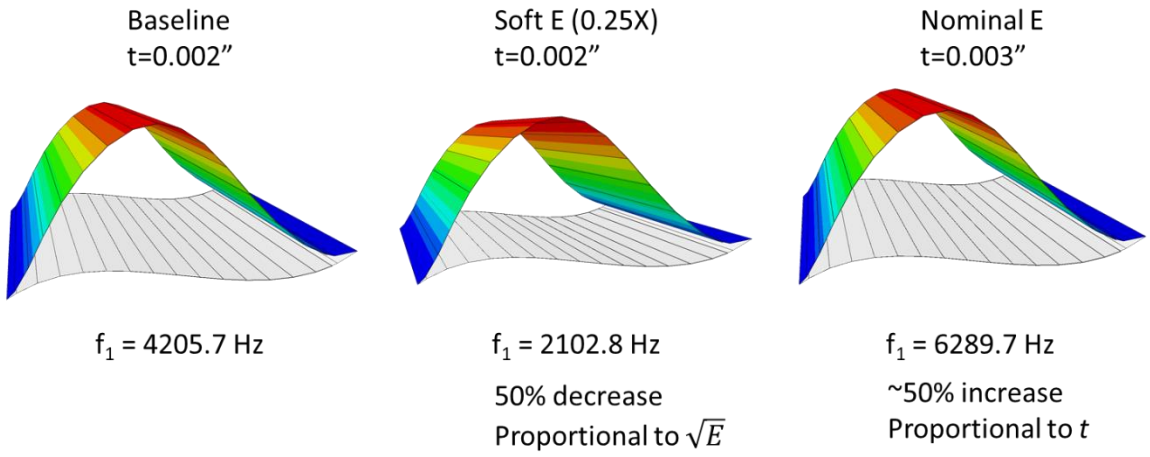

Figure 4-9. Comparison of natural frequencies for varying modulus and thickness cases.

4.1.4. Results

This section documents the key 2D SGF FSI results.

4.1.4.1. Rigid Flow-Field Comparison

The nominal approach condition for the CRM-HL test was Mach 0.2, but a rigid CFD study was performed to examine the effect of increasing Mach number on the local flow around the SGF technology element. Figure 4-10 shows selected velocity vectors and Mach number contours in the region between the slat and the main wing. Recall from the general stability analysis procedure described in section 3 that the rigid CFD simulations serve to establish initial flow conditions for the aeroelastic cosimulations.

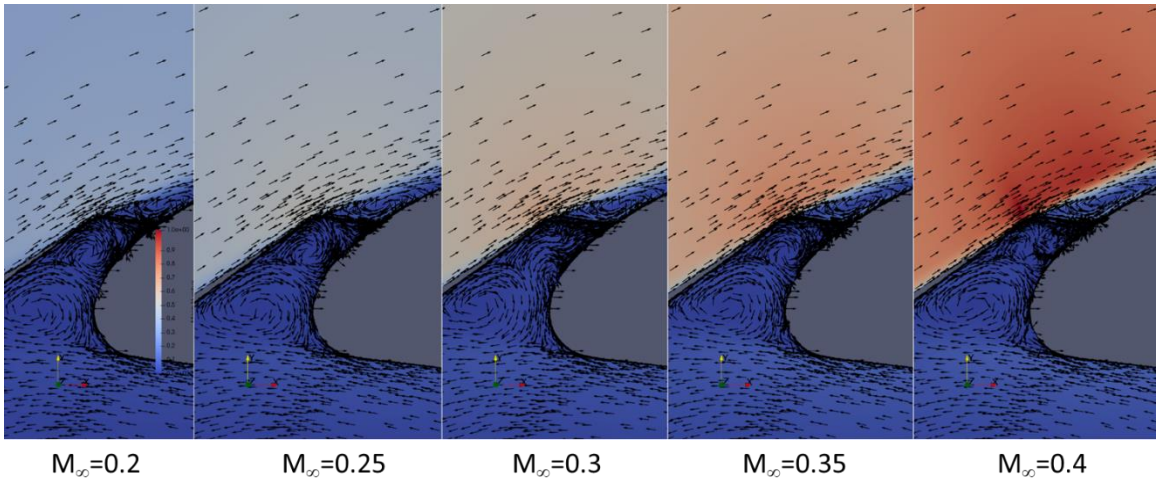

Figure 4-10. Rigid flow solutions for different Mach numbers with velocity vectors superimposed.

4.1.4.2. Quasistatic Cosimulation Results

As indicated in section 4.1.1, the CFD model includes the slat, the main wing, and the SGF technology element. Only the SGF is considered to be elastic; both the slat and main wing are assumed to be rigid. The mechanical restraints used for the aeroelastic analyses included pinning the nodes at the slat–SGF interface (nodes 810, 809, and 1301), clamping the nodes at the SGF–wing interface (node 793), and enforcing a Z-symmetry boundary condition for all the nodes.

The second step in the stability analysis procedure is to perform a quasistatic cosimulation to determine the static equilibrium state. For the quasistatic cosimulations, a dynamic simulation (`*DYNAMIC`) was performed using the Abaqus implicit solver with the quasistatic option (`APPLICATION = QUASI-STATIC`) to add artificial numerical damping, and Loci/CHEM was run using a first-order backward Euler time-integration scheme and only a single Newton subiteration. The converged, rigid CFD simulations served as the initial conditions for the quasistatic cosimulations. Following the procedure outlined in section 3.1.4, fifty cosimulation iterations or data exchanges were performed to ensure that the SGF reached an equilibrium position and the solution converged; both the force and displacement were monitored for convergence. Furthermore, nonlinear geometric effects were included in all Abaqus simulations (`NLGEOM=YES`), and no material damping was included.

Figure 4-11 shows the quasistatic aeroelastic deformation of the SGF with increasing freestream dynamic pressure. As expected, the deformation increases with increasing dynamic pressure, but due to the nonlinear geometric effects, the increase in displacement is not linearly proportional to dynamic pressure. The location of the maximum deflection occurs at approximately the midchord location (node 801) and shifts slightly aft (node 800) as the deflection further increases. Based on these observations, the dynamic response of node 801 was monitored for the dynamic results presented in section 4.1.4.3.

A summary comparison of the quasistatic analyses is provided in Figure 4-12, where the magnitude has been nondimensionalized by the SGF thickness. As can be seen, the displacement includes visible nonlinear behavior due to the large displacement geometric effects that occur at the higher dynamic pressure conditions; at the highest loading condition, the displacement is approximately 4.5 times the thickness. Using the softened modulus, the displacement is 3.3 times larger compared to the nominal modulus. Increasing the thickness to 0.003″ decreases the response by a factor of 4.23 compared to the response using the nominal 0.002″ thickness. Note that the black curve labeled “2 mil” corresponds to the response of the FEM using the incorrect material density and is identical to the red curve using the correct density, thus demonstrating that the material density does not affect the quasistatic response.

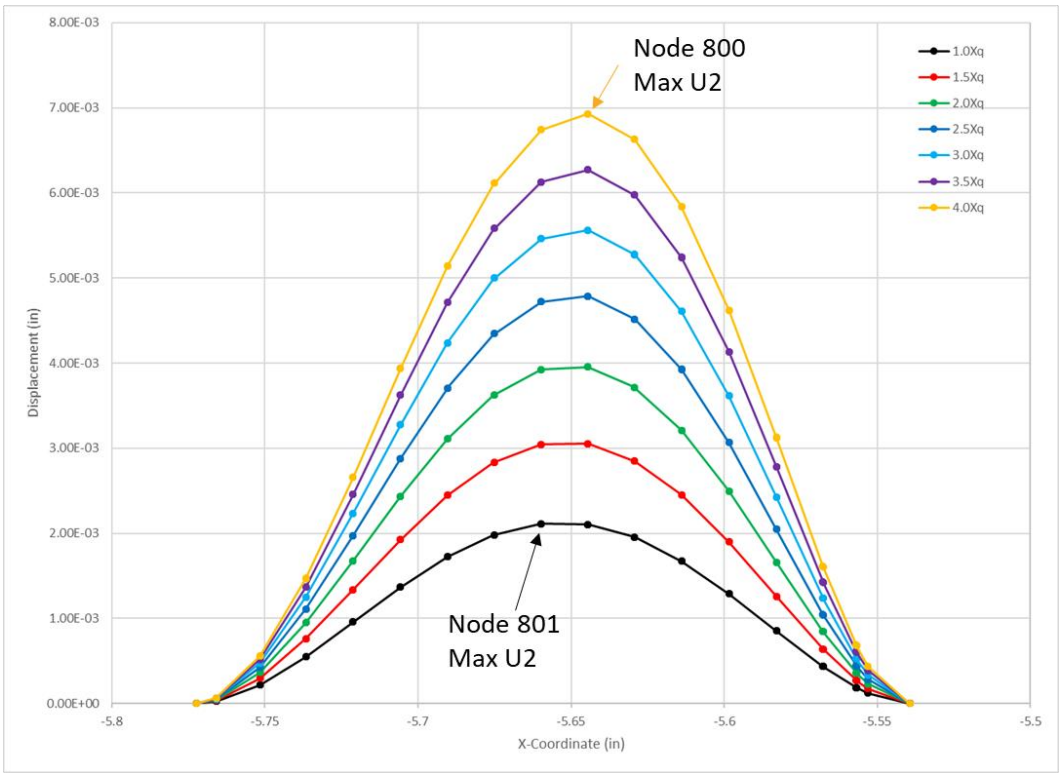

(a)

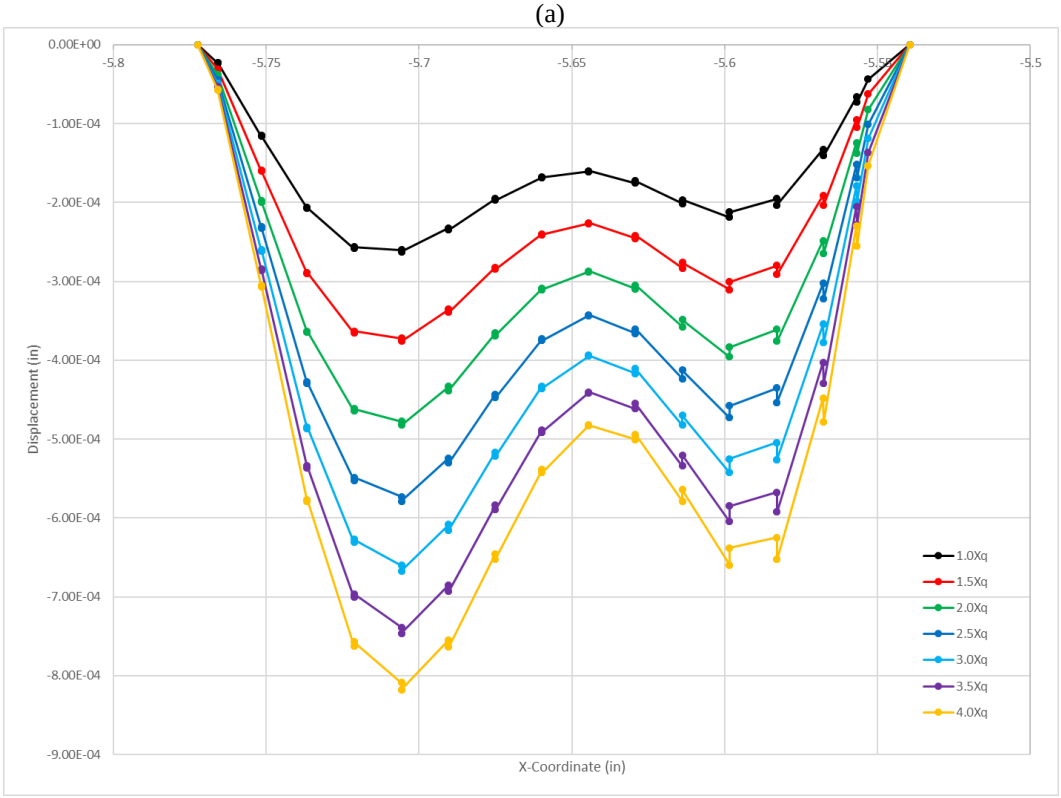

(b)

Figure 4-11. The quasistatic displacement of the 2D SGF increases with dynamic pressure: (a) vertical (U2), (b) horizontal (U1).

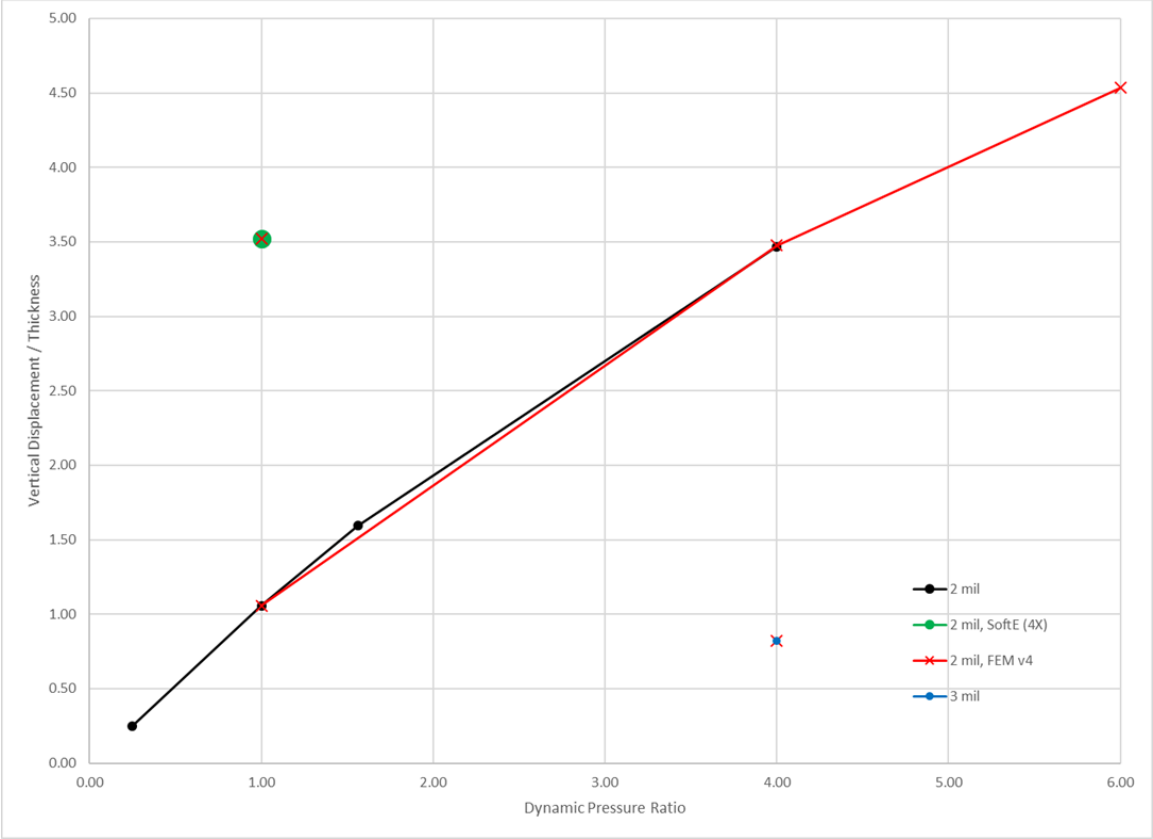

Figure 4-12. Quasistatic displacement summary, node 801.

4.1.4.3. Dynamic Aeroelastic Cosimulation Results

The third step in the stability analysis procedure is to perform a dynamic or time-accurate cosimulation to determine the dynamic response of the SGF. The 2D analyses were a stepping stone toward the 3D analyses described in sections 5 and 6. A dynamic simulation (`*DYNAMIC`) was performed using the Abaqus implicit solver and CHEM running in time-accurate mode. Abaqus and CHEM used the same time step and communicated multiple times per time step: a time step of 1.25×10^{-6} seconds was used, and the two codes exchanged data ten times per time step. The process of selecting the time step and data exchange frequency is discussed in the following paragraph.

The time step was based on resolving the first natural frequency. A time step resolution study was performed using the SGF FEM with the incorrect material density (with $f_1 = 192$ Hz) and examined time steps corresponding to 100, 200, and 400 points per period, and 200 points per period was deemed sufficient. For use with the FEM having the correct density, the converged time step was decreased by $\sqrt{386}$, which results in a $\Delta t = 1.25 \times 10^{-6}$ seconds, which corresponds to approximately 190 points per period of the fundamental mode. In addition to the time step resolution study, subiteration convergence studies were also performed using 6, 10, and 20 subiterations, and 10 subiterations was found to be sufficient.

Using the unsteady parameters described above, the dynamic cosimulation was restarted from the quasistatic solution and perturbed using an impulsive load applied to the structure. The impulse function was of short duration, 20•Δt, such that it had minimal interaction with the mean flow. Given the temporal nature of the impulsive load all the modes will be excited, but the spatial definition is shown in Figure 4-13 and was designed to particularly excite mode 1. Additionally, a lag of 0.00025 s (200•Δt) was incorporated before the impulse was applied.

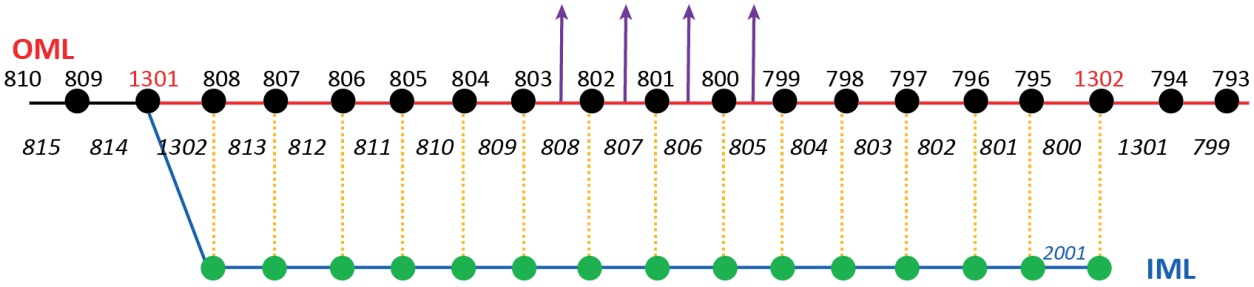

Figure 4-13. Impulse load applied to specific finite elements (805–808) in the middle of the SGF to excite a response similar to the first structural mode.

The transient response of node 801 is plotted in Figure 4-14 for the various cases conducted for the aeroelastic stability parameter study. The response for the sixfold increase in dynamic pressure and nominal SGF thickness was qualitatively similar to the four-times dynamic pressure case and is not shown. The response at the nominal dynamic pressure, Figure 4-14(a), shows that the oscillations initially occur about the quasistatic equilibrium configuration and then damp to small amplitude while the mean value drifts to a slightly lower value. At the fourfold increase in dynamic pressure, Figure 4-14(b), the vibratory response also initially occurs about the quasistatic equilibrium and quickly damps to a smaller amplitude, but the mean value exhibits an asymptotic increase with the smaller amplitude vibration exhibiting a bursting behavior about the mean. Similar trends are also observed when the thickness is increased to 0.003″, as shown in Figure 4-14(c). The softened modulus case, Figure 4-14(d), exhibits a dramatically different response. Initially it follows the same trend as the nominal stiffness case, but after thirty periods of oscillation, it begins to respond in an LCO state.

Although it is not shown in the results above, it is well known that the deflected shape from a coupled fluid-structure solution often differs from an uncoupled solution due to the influence of the structural deflection on the flowfield and vice versa. As a corollary to that observation, the changes in response character demonstrated in Figure 4-14 are attributable to differences between the quasistatic and time accurate FSI solutions. The changes in the mean value result from inertial effects absent from the quasistatic solution, and the unsteadiness in the small amplitude vibrations about the mean are mainly attributable to larger scale unsteady flow phenomena resolved by the URANS CFD solver.

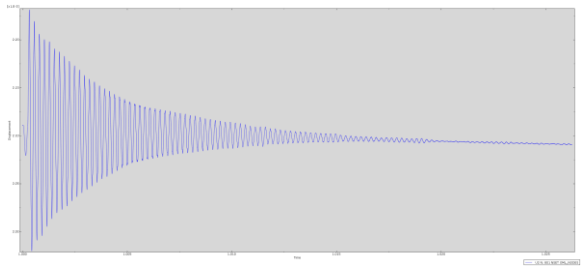

(a) Nominal Q, nominal thickness

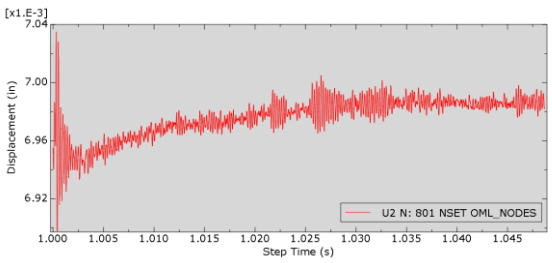

(b) Four times Q, nominal thickness

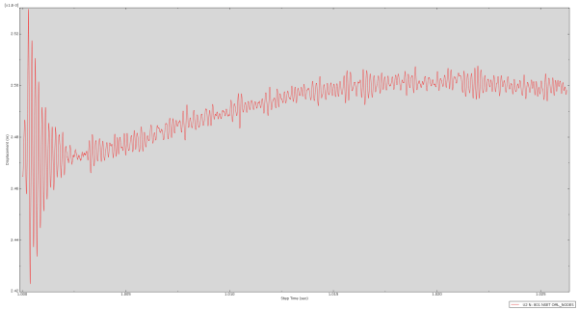

(c) Four times Q, 3 mil thickness

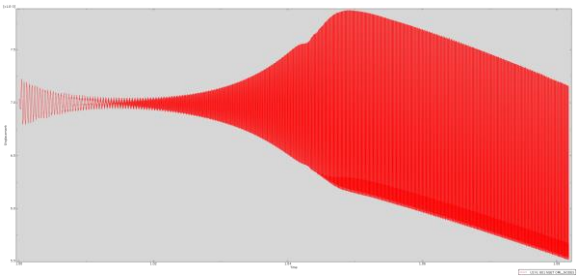

(d) Nominal Q, nominal thickness, soft E

Figure 4-14. The 2D SGF dynamic node 801 responses.

The rapid decay at the nominal dynamic pressure is clearly evident in Figure 4-14(a) and from visual observation can be determined to be stable. The response at the higher dynamic pressure conditions is considerably more complex and makes visual inspection of stability more difficult; stability was ascertained using the Partial Floquet system identification described in section 3.2. Before the system identification technique could be applied, the input signal had to be detrended; i.e., the mean value was removed and an appropriate portion of the signal selected. The nominal dynamic pressure, nominal stiffness, and higher dynamic pressure cases used the entire signal. For the softened modulus case that exhibits an LCO response, the last forty to fifty periods of the response were used, as shown in Figure 4-15.

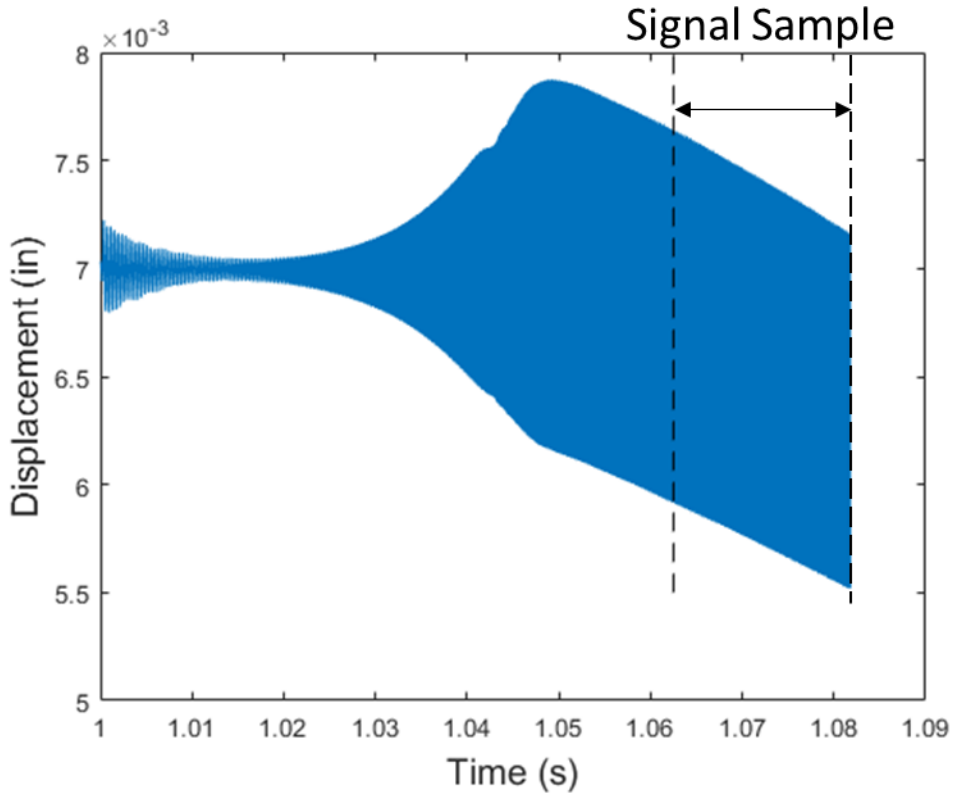

Figure 4-15. LCO signal duration used for the system identification.

The system identification of the first mode for all the cases is summarized in Figure 4-16. As mentioned, the detrended displacement time history was used for the system identification, and the first mode was selected because it had the highest energy content. The system identification shows the frequency of the first mode increasing with dynamic pressure. Damping determines stability, and the identified damping for the first mode is shown in Figure 4-17. All cases were identified as stable (the sign convention used here is that negative damping is stable). Closer inspection of the LCO case shows a subtle decay and confirms the damping prediction and stable response, as shown in Figure 4-18.

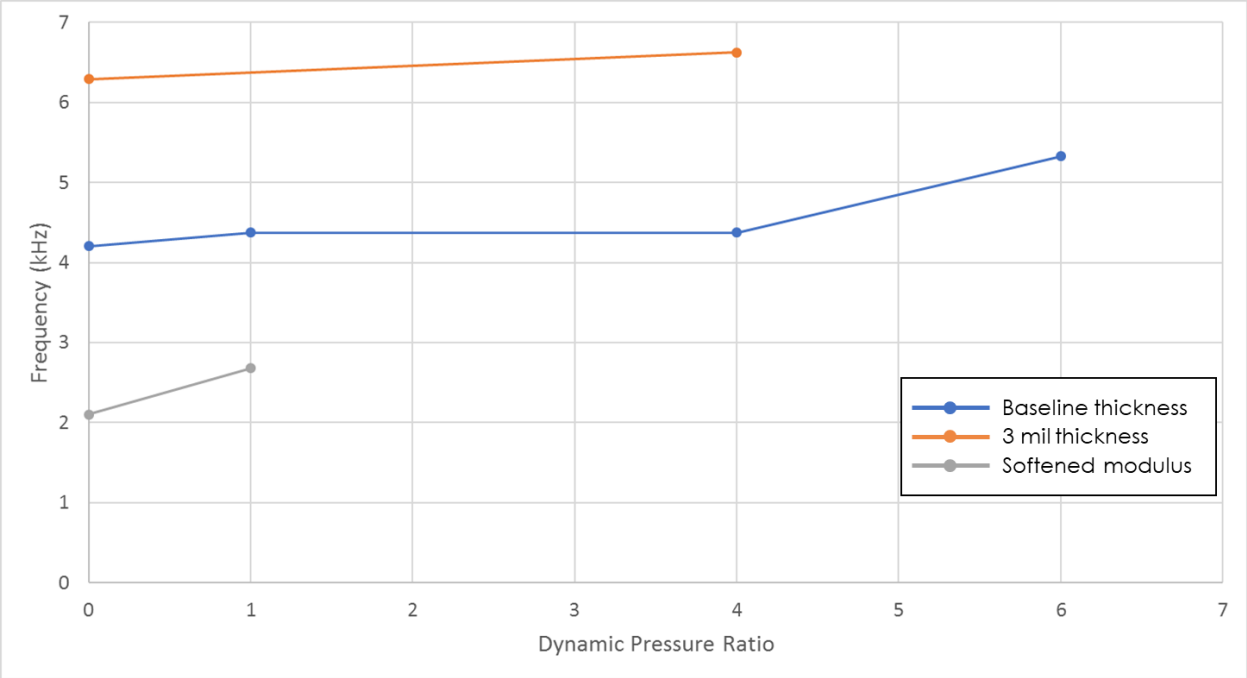

Figure 4-16. System identification and first mode frequency.

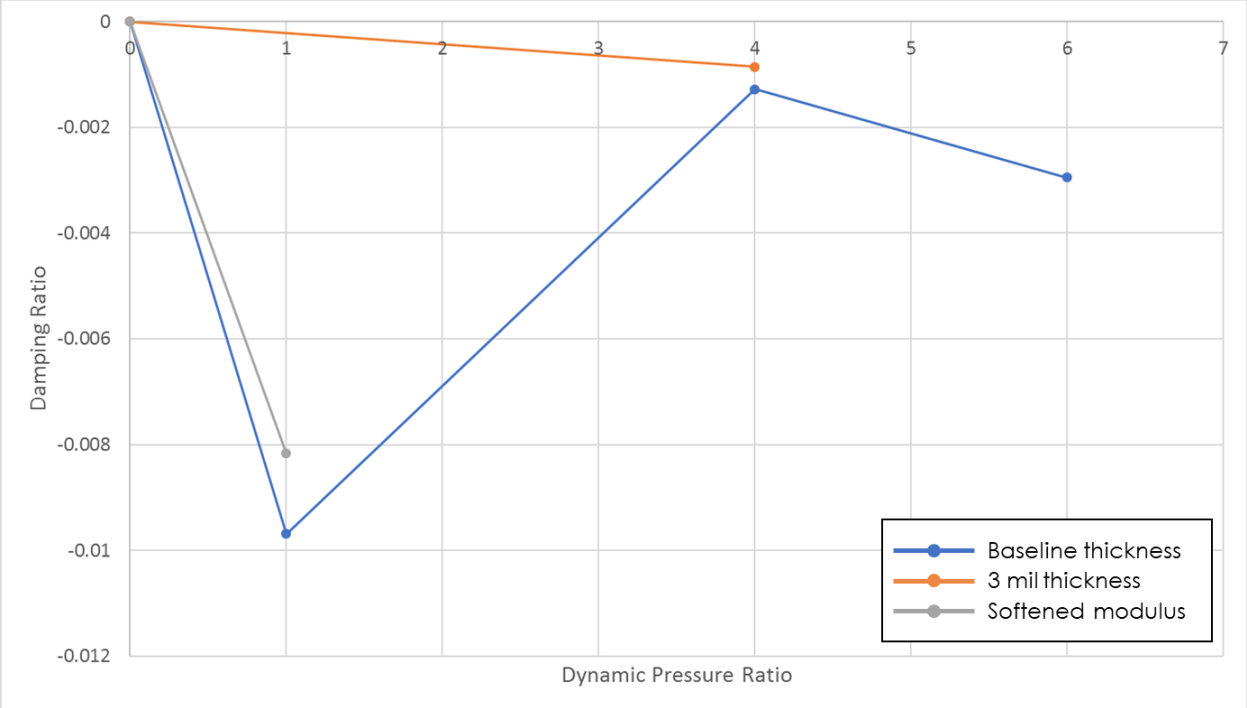

Figure 4-17. System identification and first mode damping ratio.

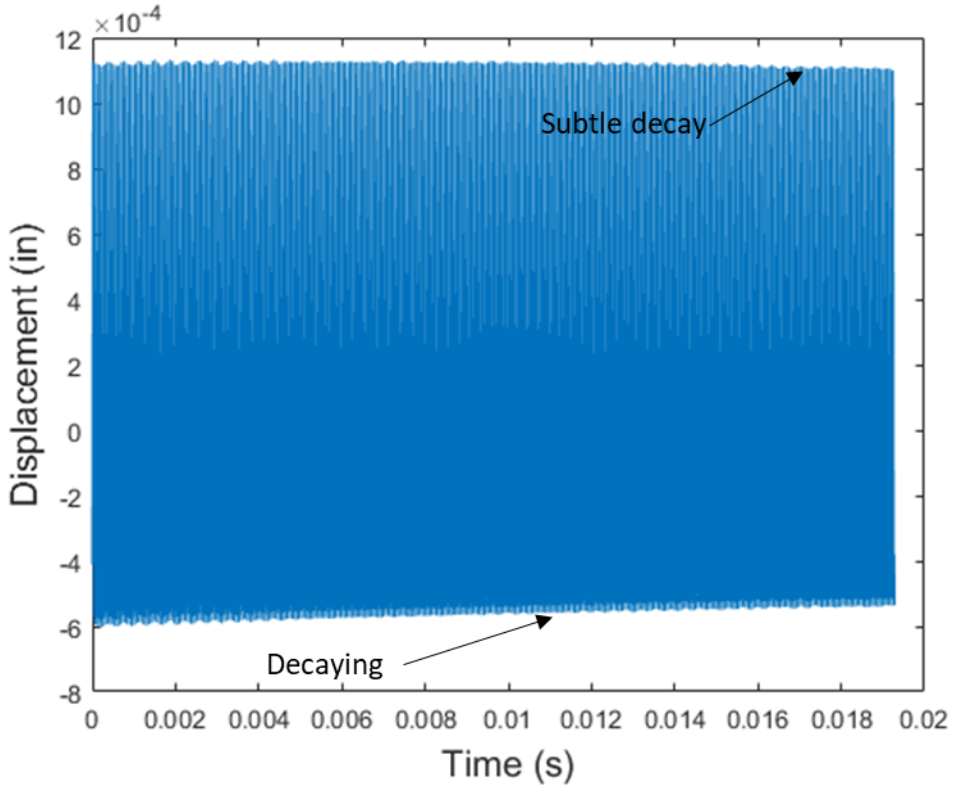

Figure 4-18. Detrended LCO signal exhibits a subtle decay and a lightly damped response.

4.1.4.4. Aerodynamic vs. Numerical Damping Demonstration

The initial damping estimates using the FEM with the incorrect density revealed very low aerodynamic damping, and a question was raised about the relative order of magnitude of the aerodynamic damping to the numerical damping. To demonstrate that the aerodynamic damping arising from the cosimulations is larger than the numerical damping, ATA performed a pair of Abaqus-only, or uncoupled, simulations designed to mimic the stability analysis solution process, using the same aerodynamic loading as the cosimulations. Since the aerodynamic load acts only as an external load and is not cosimulating, and there is no material or structural damping applied, the only damping in the response is due to the numerical damping associated with the numerical algorithms used to solve the structural equations of motion and that associated with the finite discretization of the SGF structural FEM. For these simulations, the SGF used the nominal thickness (t = 0.002″) and the nominal moduli. Despite the FEM using the incorrect density, the comparison of the relative magnitudes of the aerodynamic and numerical damping is still valid.

The aerodynamic load from the converged quasistatic cosimulation was extracted from the cosimulation output files (the `.SRE` file) and applied as a static external load to the uncoupled static simulation. The aerodynamic loading corresponds to the four-times-nominal dynamic pressure condition. Figure 4-19 shows a comparison of the vertical (U2) and horizontal (U1) displacements from the quasistatic

cosimulation and the uncoupled static simulation. As expected, the results match almost identically, and the percent difference between the two solutions is shown in Figure 4-20.

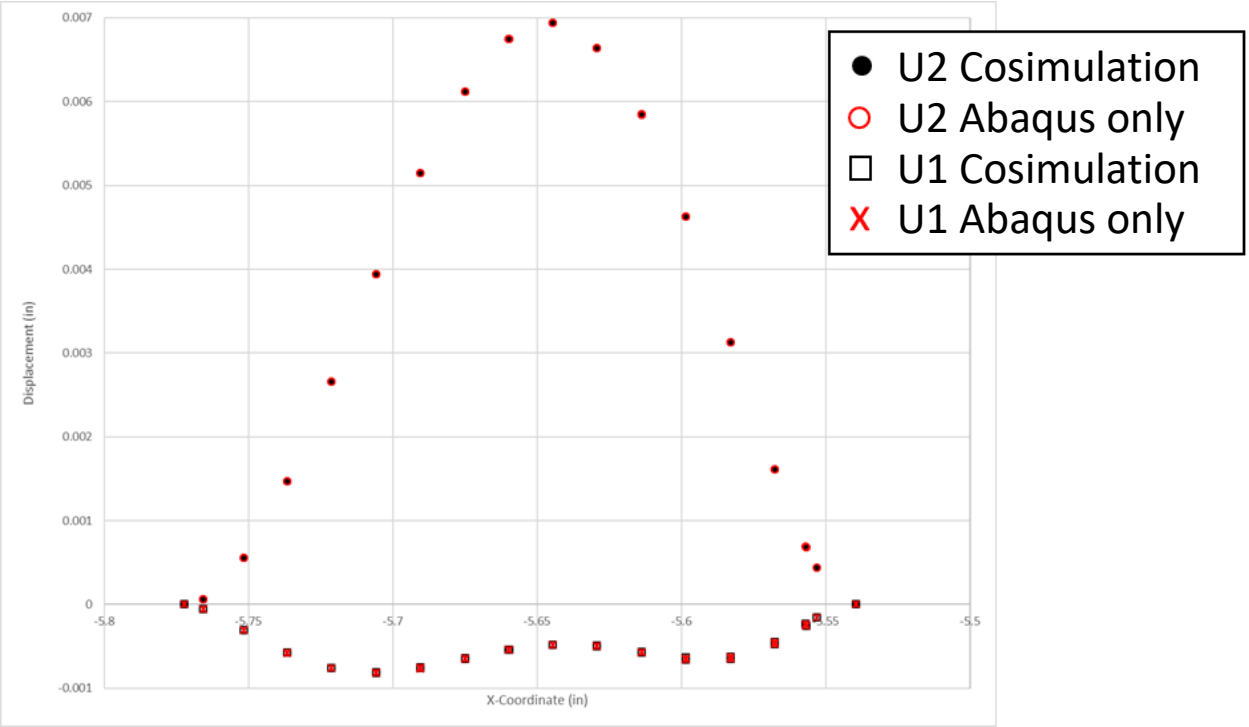

Figure 4-19. Comparison of the coupled and uncoupled static response.

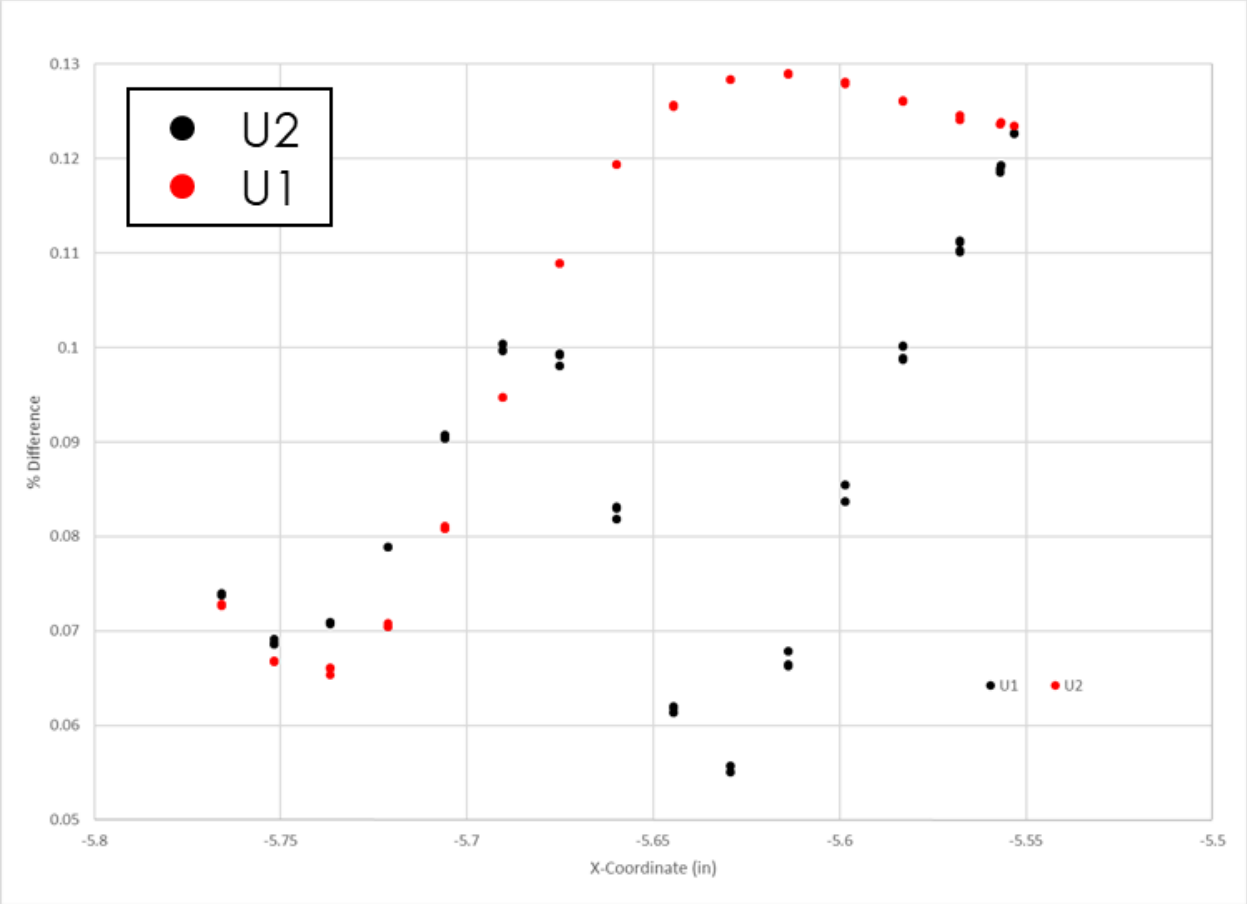

Figure 4-20. Percent difference in the uncoupled versus coupled static response.

Next, an uncoupled dynamic simulation was performed using the same impulsive load as was applied to the dynamic cosimulations. Similar to the cosimulation process, the transient simulation was restarted from the static analysis. A comparison of the transient response of node 800 is shown in Figure 4-21, in which it is clearly evident that the response from the uncoupled (Abaqus-only) simulation has very little damping while the cosimulation response decays significantly, thus confirming that the aerodynamic damping derived from the coupled cosimulations is larger than the numerical damping.

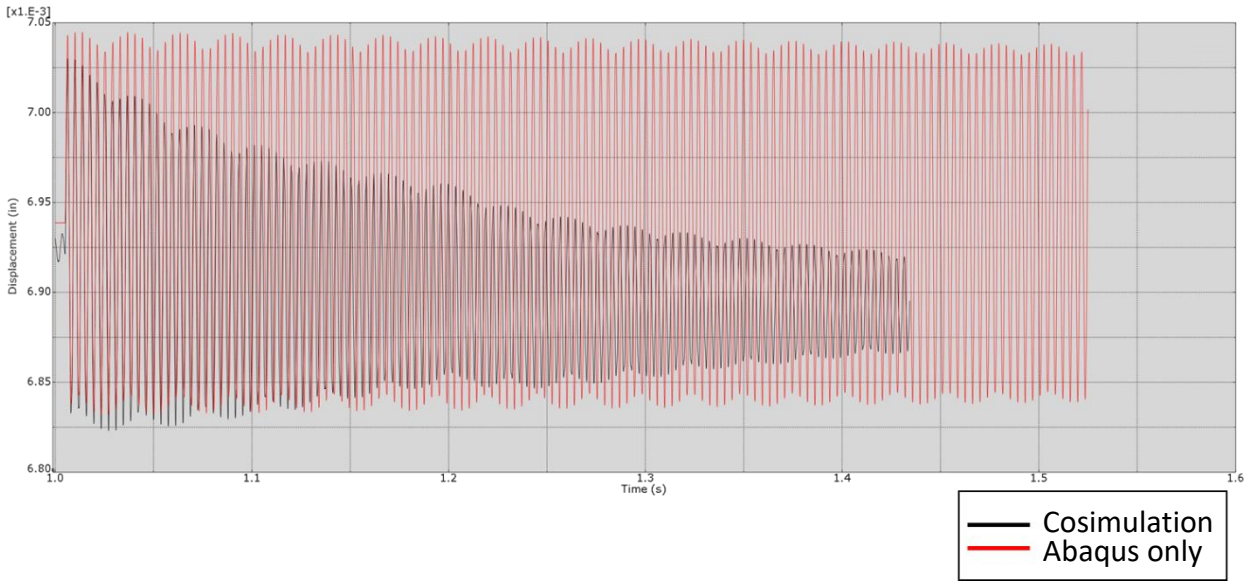

Figure 4-21. Coupled vs. uncoupled transient response comparison for the vertical displacement of node 800, confirming that the aerodynamic damping of the cosimulations is larger than the numerical damping.

4.1.5. Summary and Lessons Learned

The FSI tools were successfully applied to the 2D SGF. An aeroelastic parameter study was performed to bound stability estimates and varied both the aerodynamic and structural parameters. The dynamic pressure was increased to up to six times the nominal condition and SGF stiffness was varied ±50% via changes to the thickness and elastic modulus. The quasistatic deflection increased with dynamic pressure in a nonlinear manner due to the large displacement geometric effects that occur at the higher dynamic pressure conditions; at the highest loading condition, the displacement is approximately 4.5 times the thickness. When the softened modulus is used, the displacement is 3.3 times larger compared to the nominal modulus. Increasing the thickness to 0.003″ decreases the response by a factor of 4.23 compared to the response using the nominal 0.002″ thickness. All dynamic responses were found to be stable using the Partial Floquet system identification technique. The response at the nominal dynamic pressure shows that the oscillations occur about the static equilibrium condition. At the fourfold increase in dynamic pressure, the vibratory response initially occurs about the original equilibrium but then shifts to the response about another equilibrium state. This trend was also observed when the thickness was increased to 0.003″. The softened modulus case, however, exhibits a dramatically different response. Initially it follows the same trend as the nominal stiffness case, but after thirty periods of oscillation it begins to respond in an LCO state with very light damping.

Because the SGF FEM was constructed using shell elements, considerable effort was spent on constructing a wetted surface for the FEM so that the loading and displacement could be transferred to the CFD wetted surface. Future structural modeling efforts for these types of thin structures should consider elements such as continuum shell elements that look like three-dimensional continuum solid elements but have kinematic and constitutive behavior similar to conventional shell elements. The use of continuum

shell elements allows for the specification of the OML and IML surface consistent with the CFD wetted surface definition. When the FEM and CFD meshes are constructed, care should also be taken to ensure consistent boundary or surface patch definitions.

4.2. Slat-Cove Filler

This section discusses the 2D SCF analyses performed between November 2016 and January 2017 under the original TEAMS2 contract.

4.2.1. Description of Models

This section describes the model boundary conditions, mass properties, dry modes, and CFD mesh.

4.2.1.1. Boundary Conditions

ATA began the SCF analysis effort by reviewing the FEM and boundary conditions supplied by NASA LaRC. An overview of the 2D SCF FEM is shown in Figure 4-22. The SCF possessed Z-symmetry, and TIE constraints were placed at the leading and trailing edges. At the trailing edge, the SCF is tied to the slat which involved slat elements 496 and 497 and SCF elements 542 through 545. At the leading edge, the SCF is tied to a rigid hinge that is joined to the slat near the cusp, which involved slat elements 540 and 541 and SCF elements 622 through 624, as depicted in Figure 4-23.

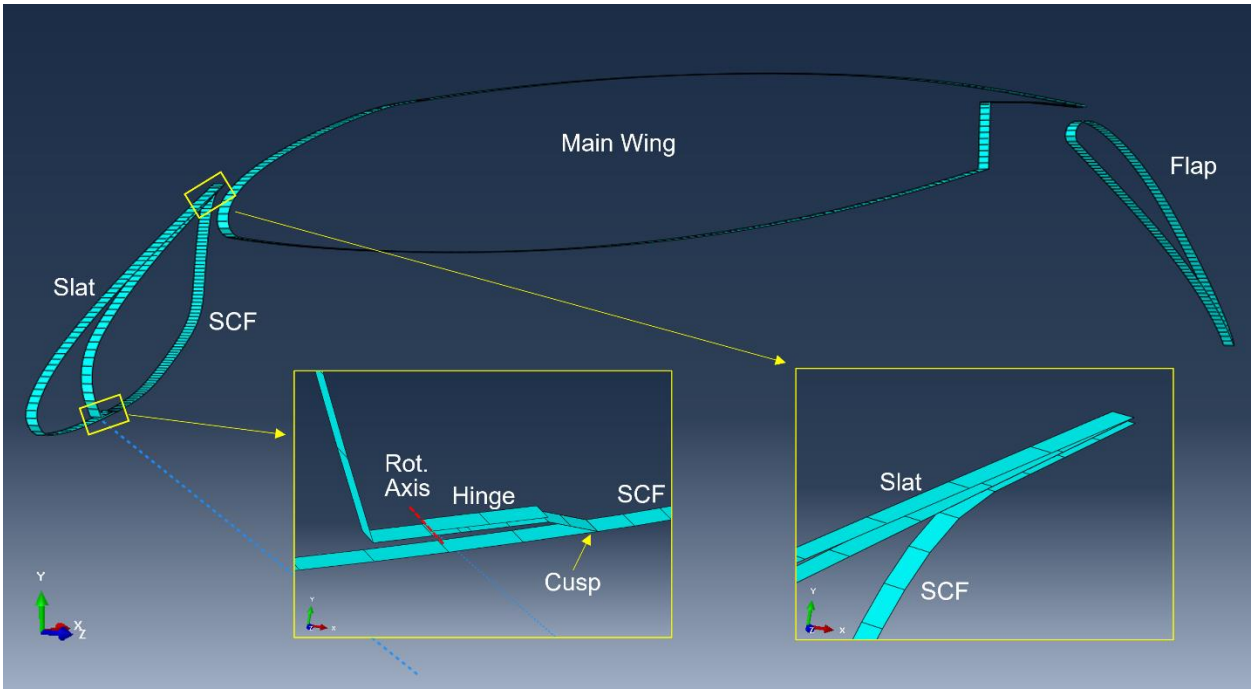

Figure 4-22. SCF FEM showing detail at leading and trailing edges.

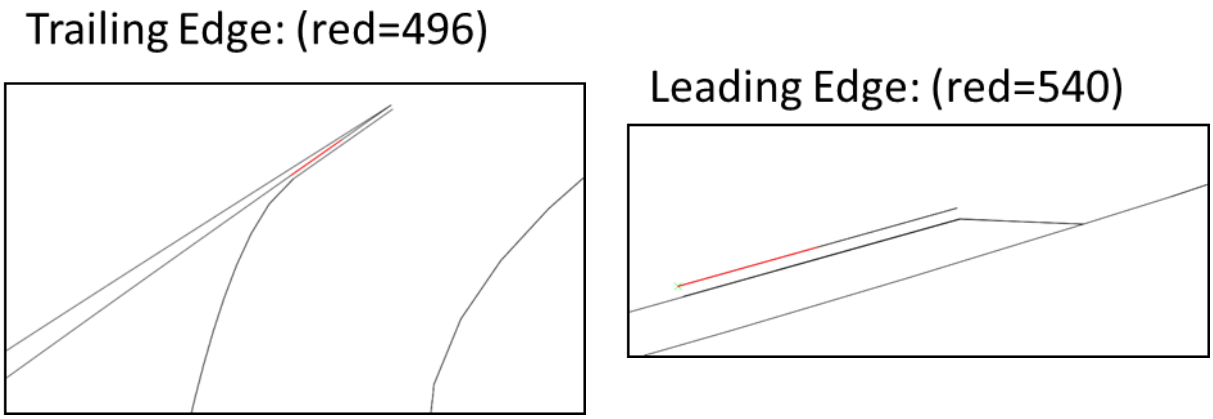

Figure 4-23. SCF TIE boundary conditions at leading and trailing edges.

4.2.1.2. Modal Mass Properties

ATA next reviewed the modal mass properties of the SCF, which are described below

- Mass: 6.46465e-006 lbf•s^2/in^4, equivalent to 0.002496 lbm or 1.13 grams
- Center of mass (x,y,z) = (−0.624078, −0.222308, 0.05) inches
- Moments of inertia (Ixx, Iyy, Izz, Ixy, Iyz, Izx)
    - About the origin: (3.27770e-006, 1.38226e-004, 1.41460e-004, −6.13125e-006, 7.18572e-008, 2.01723e-007) lbf•s^2/in^2
    - About the center of mass: (2.94205e-006, 1.35692e-004, 1.38623e-004, −5.23436e-006, 6.83100e-015, 3.65987e-014) lbf•s^2/in^2

4.2.1.3. SCF Modes

ATA ran a dry-modes analysis in Abaqus with the NASA FEM. The first six modes are shown in Figure 4-24. Of note is that the first SCF mode (lowest frequency) has a frequency of 188.6 Hz and appears to have a shape typical of second bending. Dynamic FSI responses were found to be dominated by a combination of the lowest modes.

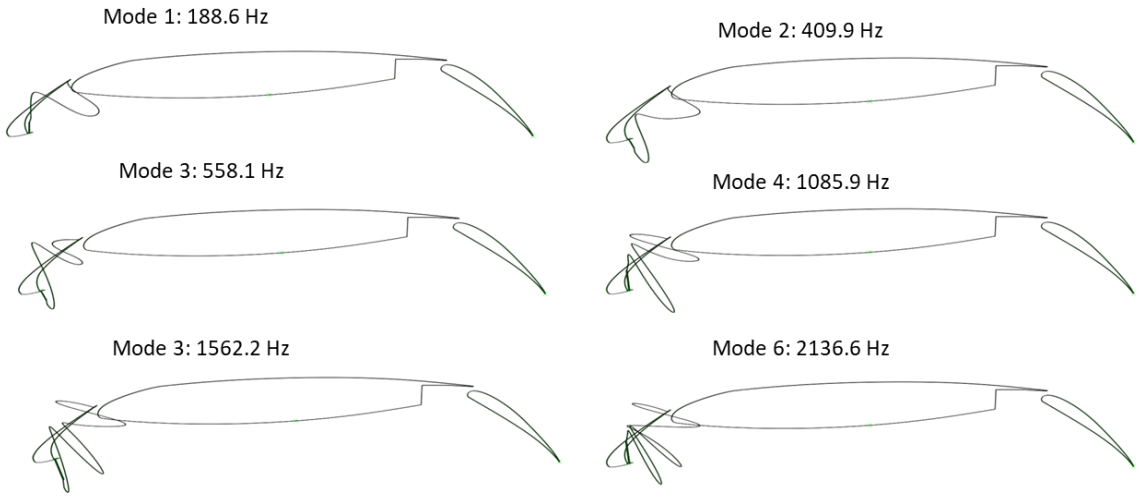

Figure 4-24. First six dry modes of SCF obtained from Abaqus FEM.

4.2.1.4. Mesh

ATA first reviewed the CFD mesh provided by NASA, which is pictured in Figure 4-25. Since the problem is 2D, the XZ plane is shown, and a symmetry boundary is assumed on planes normal to the Y direction. The mesh contained 525,724 nodes and 522,001 elements. All elements were tetrahedral.

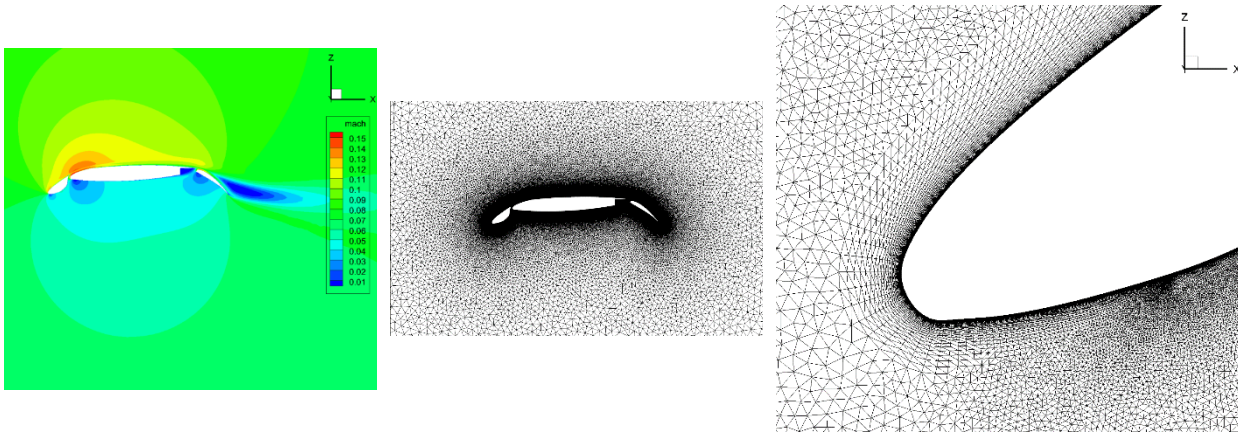

Figure 4-25. CFD mesh was inspected.

A scale factor was applied to the CFD mesh so that the FEM and CFD discretizations would overlay. The received CFD mesh was converted from feet to inches, and the spanwise extent (Z direction) was scaled to 0.1″. ATA then split the symmetry plane into two surfaces, i.e., a left and right plane. The slat and SCF were split into separate surfaces. Figure 4-26 shows an image of the FEM and CFD models together. Scaling was found to help considerably but did not produce an exact match, as shown in Figure 4-27.

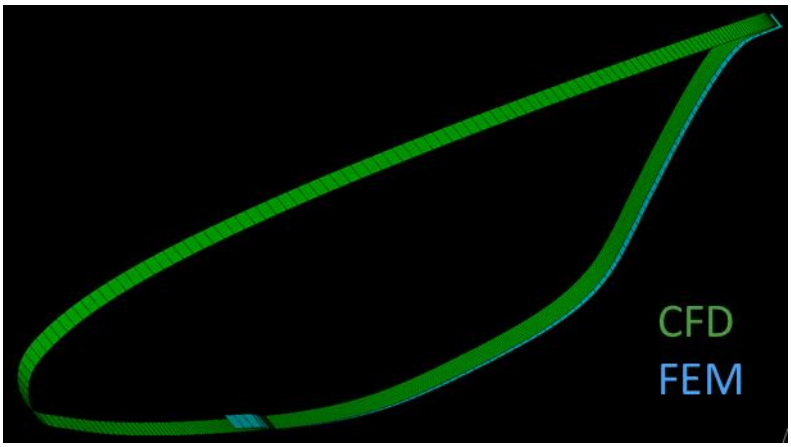

Figure 4-26. CFD mesh conversion results in CFD and FEM overlay.

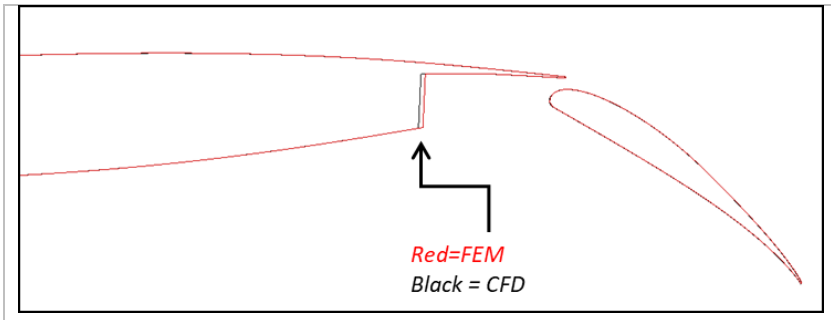

(a) Details near flap cove

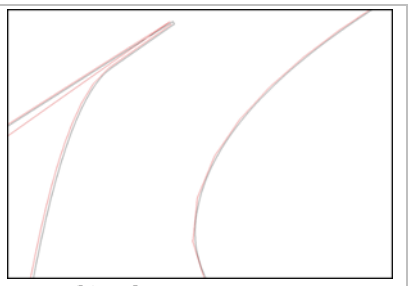

(b) Slat / main wing suction side detail.

Figure 4-27. FEM/CFD mesh overlay does not produce exact overlap between models. FEM shown in red and CFD lines shown in black.

4.2.2. Results

Figure 4-28 presents contours of the Mach number field around the entire wing, as well as pressure coefficient contours near the SCF. After rigid CFD was completed using CHEM and compared to FUN3D (not shown), both quasistatic and dynamic FSI simulations were performed for the SCF. Dynamic FSI simulations were analyzed for stability using the Partial Floquet method.

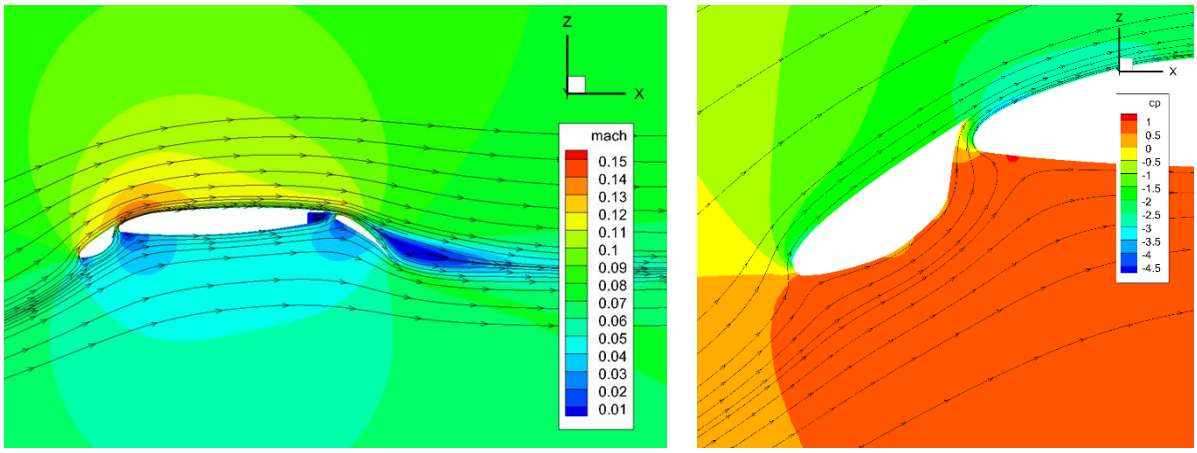

Figure 4-28. FUN3D rigid CFD solution Mach number contours (left) and pressure coefficient contours near SCF (right).

4.2.2.1. Quasistatic Aeroelastic Cosimulation

A quasistatic FSI simulation was initially performed. Abaqus was run using the implicit solver under the Application = QUASI-STATIC setting with a variable time step that was bounded between 1 μs and 0.25 ms. CHEM was run using a first-order time-integration scheme and no additional subiterations. The quasistatic deformation was found to exhibit oscillatory response about a mean value, as shown in Figure 4-29. The maximum deformation was found to be about 0.002″. Given the inherent oscillatory response, the first dynamic FSI simulation was run without a perturbation to determine whether the structure would be able to self-excite from numerical disturbances in the coupled solution.

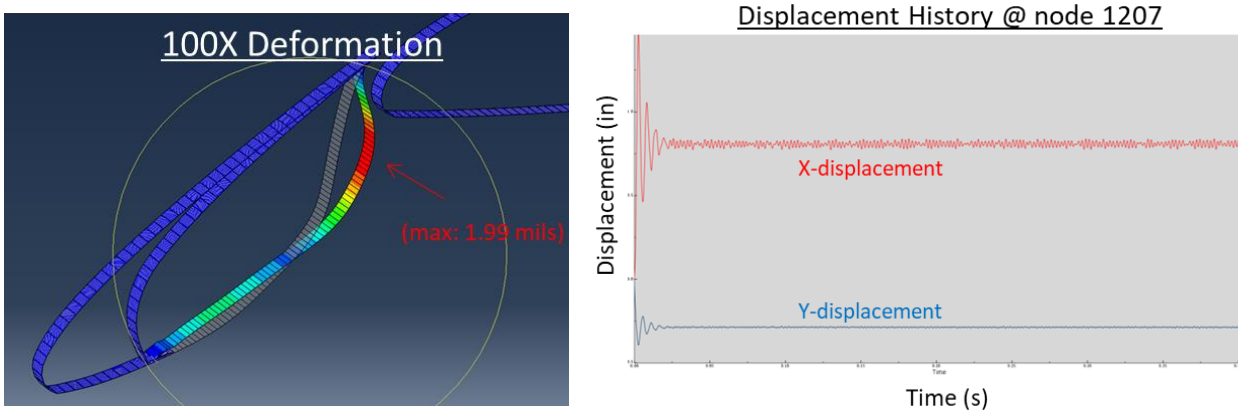

Figure 4-29. Quasistatic response of SCF showed oscillatory behavior.

4.2.2.2. Dynamic Aeroelastic Cosimulation

Two time-accurate FSI simulations of the SCF were performed: an initial study with no excitation (perturbation) of the structure, followed by a second study that used an impulsive perturbation.

4.2.2.2.1 Time-Accurate Investigations—No Perturbation

Based on the oscillatory response of the quasistatic simulation, a dynamic FSI simulation of the SCF was first performed by restarting the quasistatic solution in a fully time-accurate manner at a time step of 100 μs. A snapshot of the instantaneous pressure comparing rigid steady and aeroelastic unsteady cases in Figure 4-30 shows few easily discernable differences. However, an animation was able to illustrate that the unsteady case possessed a fluctuating flow field.

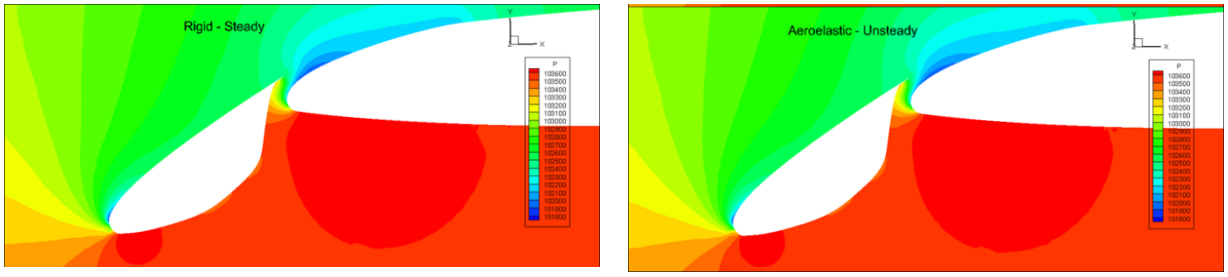

Figure 4-30. Few differences are easily discernable between the rigid steady (left) and aeroelastic unsteady (right) pressure fields.

An oscillatory response of the structure was observed as shown in Figure 4-31 and was investigated using the frequency-domain response spectrum, and with Partial Floquet analysis. The primary response was near the frequency of the first mode, and damping values were somewhat inconclusively found to be both positive and negative. A statistical assessment of the oscillations suggested an X-deformation standard deviation of 0.8% and a Y-deformation standard deviation of 0.4%. It was suspected that this time-accurate simulation was subject to numerical background noise, so the simulation was repeated with a time step one order of magnitude smaller (10 μs). As shown in the same figure, the refined time step removed the background noise that was attributed to the underlying (rigid) CFD solution. It was thought that a perturbation that exceeded the background noise would provide a more robust determination of system stability, which led to the next time-accurate investigation using an impulsive excitation.

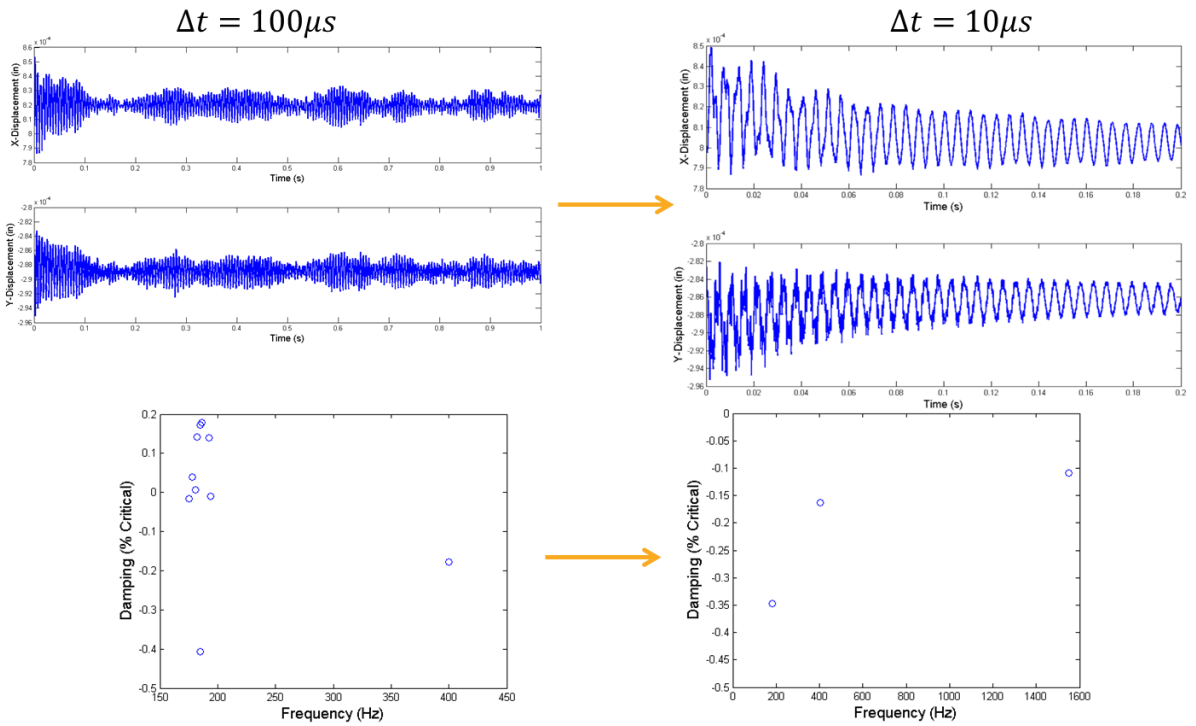

Figure 4-31. Time-accurate investigation with no perturbation resulted in LCO response.

4.2.2.2.2 Time-Accurate Investigations—With Perturbation

Time accurate simulations were repeated using an impulsive pressure applied to element 592 of the FEM. A "coarse" time step of 100 μs was used. This impulse produces a positive deflection for most modes. Figure 4-32 shows the first three modes and the location of the element that receives the impulse; these modes correspond to the ones shown in Figure 4-24. The impulse magnitude was set to 0.25 Pascals and its duration was 0.05 s, approximately equal to the period of the first mode.

Figure 4-33 indicates that the dynamic perturbation results in a heavily damped response. Additionally, the initial response to the perturbation excitation is seen to overwhelm the background oscillations that were seen in the case without a perturbation. A Partial Floquet reconstruction of the response shown in Figure 4-34 suggests that the response to the excitation is primarily composed of the first mode.

Table 4-4 indicates that Partial Floquet reconstruction produces accurate modal parameters. The first two modes compare well between unperturbed and perturbed responses. It is noted that the Partial Floquet method cannot pick out all modes. Unresolved modes (and dry mode damping) are marked as N/A. Visually, all modes are found to be stable.

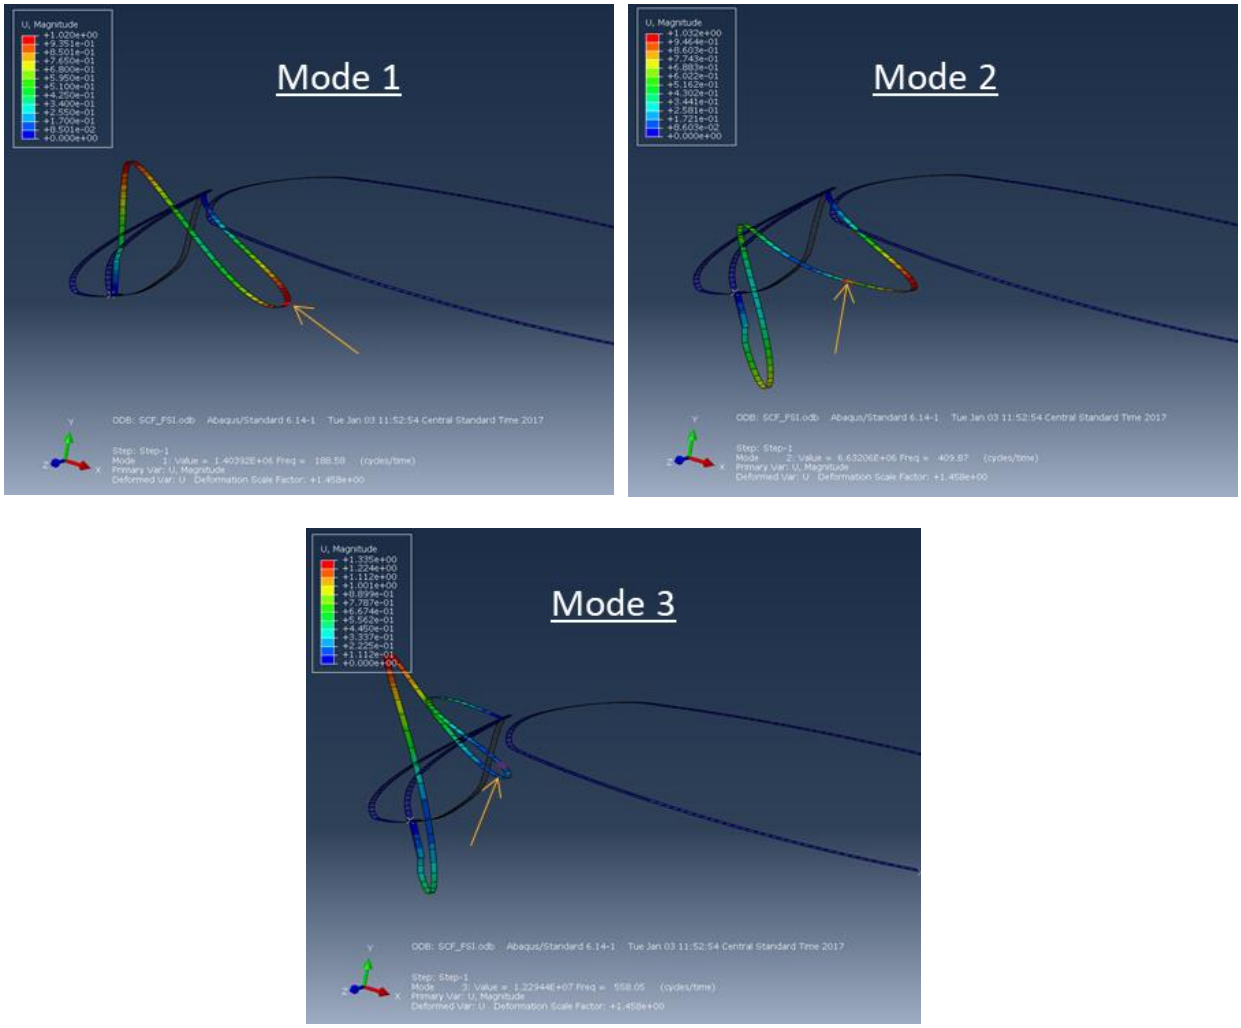

Figure 4-32. Perturbation applied at element 592.

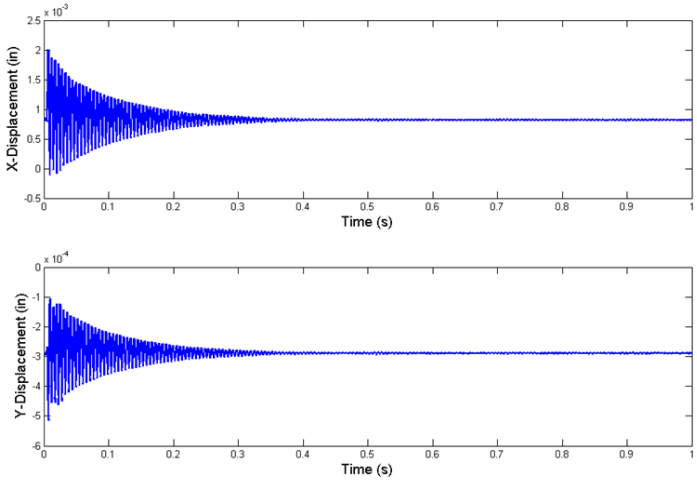

Partial Floquet

- Freq: 183.26 Hz
- Damping: -0.84%

- Freq: 404.6 Hz
- Damping: -0.76%

- Freq: 529.9 Hz
- Damping: -0.77%

Figure 4-33. Damped response in X (top) and Y (bottom) shows heavy damping.

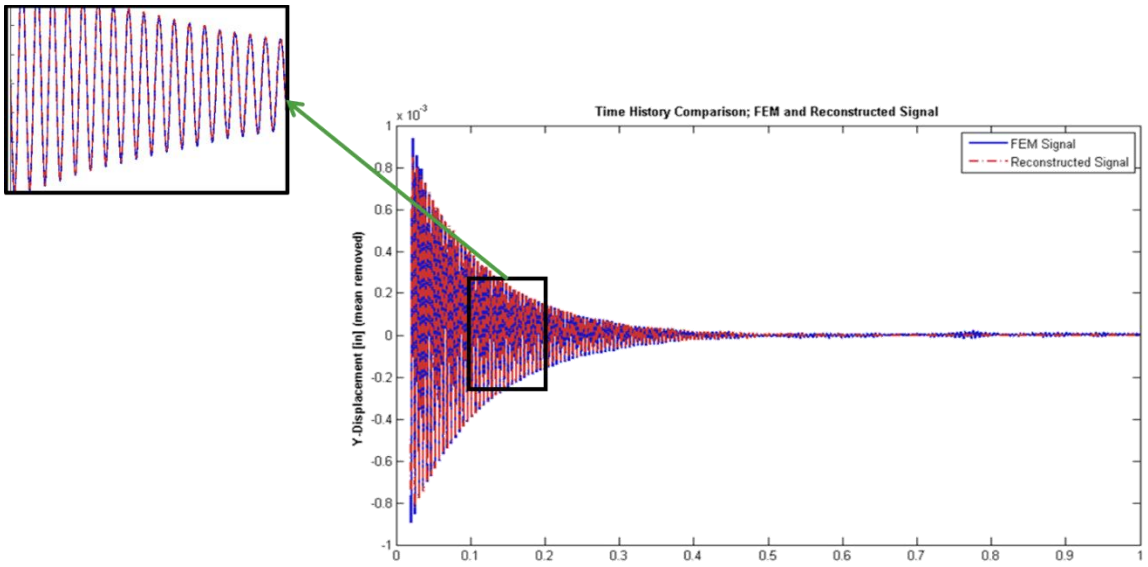

Figure 4-34. Partial Floquet reconstruction using one mode suggests that the response is dominated by the first mode.

Table 4-4. Agreement of modal parameters seen with or without perturbation for leading two modes.

		Perturbation	Steady (Time-Accurate)	Dry Modes
Mode 1	Frequency (Hz)	183.3	183.7	188.6
	Percent Damping	-0.84	-0.35	*N/A*
Mode 2	Frequency (Hz)	404.6	403.6	409.9
	Percent Damping	-0.76	-0.16	*N/A*
Mode 3	Frequency (Hz)	529.9	*N/A*	558.1
	Percent Damping	-0.77	*N/A*	*N/A*
Mode 5	Frequency (Hz)	*N/A*	1551	1562
	Percent Damping	*N/A*	-0.11	*N/A*

4.2.3. Summary and Lessons Learned

ATA successfully applied FSI tools to the problem of the 2D SCF. The quasistatic deformation was found to exhibit oscillatory response about a mean value with a maximum deformation of about 0.002″. Given the inherent oscillatory response, the first dynamic FSI simulation was run without a perturbation to determine whether the structure would be able to self-excite from numerical disturbances in the coupled solution, and a second dynamic FSI simulation was run using a prescribed impulsive perturbation. It was

found that applying an impulsive pressure to the structure allowed the response to ring down in a way such that the damping was more numerically evident and quantifiable with Partial Floquet analysis.

5 3D SUBDOMAIN MODELING STUDIES ON CRM-HL BASELINE GEOMETRY

For the 3D aeroelastic simulation results presented in section 6, a subdomain modeling approach was used to minimize the computational expense. In the subdomain modeling approach, the flow solution is first obtained for the entire domain, which in this case is a semispan model of the CRM-HL configuration, and the entire vehicle is assumed to be rigid. Next, a smaller subdomain mesh is either extracted from the full domain mesh or constructed afresh.

Subdomain modeling for FSI offers substantial computational savings, as it offers the potential of a significantly reduced computational domain size. Figure 5-1 presents an example of a bounding box which can be used for a subdomain mesh for FSI of a technology element on the slat of a wing in a high-lift configuration. Keys aspects of effective subdomain modeling include proper selection of subdomain size, i.e., subdomain boundaries sufficiently far from location of interest and yet small enough to realize computations savings, and appropriate definition of boundary conditions at the subdomain farfield boundaries.

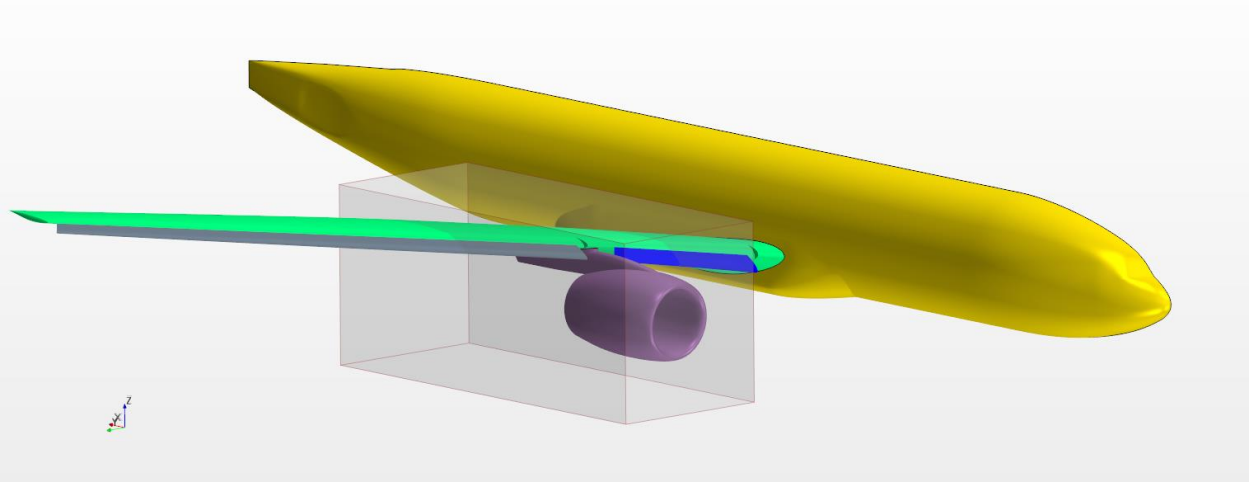

Figure 5-1. Bounding box illustrating the extents of a local subdomain used in the subdomain modeling approach.

ATA has successfully used this subdomain modeling approach for other FSI applications [40]. In those cases, the flow field was supersonic, and any pressure perturbations introduced by the elastically deforming structure could not propagate upstream but could only convect downstream, thus making the selection of the bounding box and definition of farfield boundary conditions straightforward. For the CRM-HL configuration, the flow is subsonic, allowing disturbances to propagate in all directions, which makes balancing the subdomain size requirements more challenging. Therefore, ATA decided to investigate the suitability of using the subdomain approach on the CRM-HL configuration by conducting a parameter study to determine the necessary size of the subdomain. Two studies were performed: the first on an outboard section of the wing, and the second on an inboard section of the wing that includes the nacelle and the region where the technology elements are most highly loaded. This subdomain modeling study and the CFD results used in the study were all conducted at full scale.

A subdomain mesh may be extracted from a larger mesh using the `vogcut` utility (as described in section 3.1.1.1), and boundary conditions may be extracted from a solution on the larger mesh using the `scalablePrescribed` utility (as described in section 3.1.1.2). In this project, the full-domain meshes pertain to semispan domains of the CRM-HL. NASA provided ATA with a number of semispan meshes that were used for subdomain studies.

- **l10wu5** – Initial mesh provided by NASA to ATA, generated using the VGRID meshing software.
- **crm-hl-2block-2018_04_13** – This mesh had finer resolution than l10wu5. It was generated with the Pointwise meshing software and contained embedded boundaries used to define and extract subdomain models.
- **crm-hl-2018_06_26** – The surface mesh was nearly identical to crm-hl-2block-2018_04_13; however, the volume mesh differed because it did not contain embedded boundary surfaces.

The boundary conditions on the subdomain outer boundaries are prescribed by using the flowfield solution from the full-domain model. In Loci/CHEM, the `inflow`, `supersonicInflow`, and `farfield` boundary conditions allow nonuniform flow conditions to be specified from a user-defined file. For all the subdomain model simulations presented here, the boundary condition for the subdomain outer or "farfield" boundary was defined using Loci/CHEM `farfield` boundary condition. The boundary condition values were obtained from the flowfield solution from the full-domain model. In essence, the `scalablePrescribed` utility allows the subdomain model to interpolate the flow solution from the full-domain model to enforce the appropriate boundary condition on the "farfield" boundaries.

Before proceeding with the subdomain studies, ATA verified that the Loci/CHEM solution was substantially in agreement with the NASA-provided FUN3D solution provided for the *l10wu5* grid. ATA then performed subdomain modeling studies for both inboard and outboard subdomains, which are described in the next sections. The key metrics for comparison between the subdomain and full model are pressure coefficient (C_p) and wall shear stress profiles at various spanwise locations.

5.1. Outboard Subdomain

Subdomain models were first compared to a semispan solution over the outboard section of the CRM-HL. The outboard case is simpler than the inboard one as it contains no nacelle geometry embedded in the domain. For all the subdomains that were either extracted with `vogcut` or created in a meshing tool, C_p and wall shear stress comparisons were made between the subdomain solution and the semispan solution at eleven spanwise locations, as pictured in Figure 5-2.

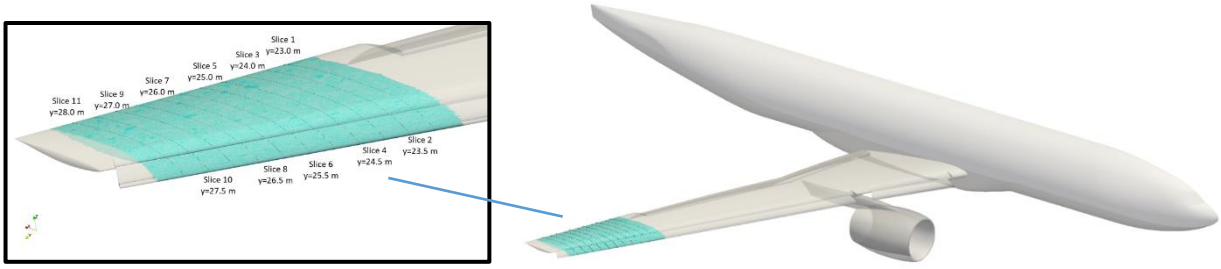

Figure 5-2. Region on semispan domain (shown shaded) where outboard subdomain results were compared to semispan CFD results.

5.1.1. Initial Outboard Subdomain Analysis: *l10wu5* mesh

Initial comparisons were made between the full (semispan) *l10wu5* grid and breakout models extracted using `vogcut` and for which `scalablePrescribed` boundary conditions were used to enforce boundary data at the farfield extents of each subdomain model. Three breakout subdomains were extracted from the merged grid *l10wu5* using the `vogcut` utility. As shown in Figure 5-3, `vogcut`-extracted subdomain grids for the outboard span section reduced problem size by a factor of ten relative to the full semispan model. Since `vogcut` is a precise removal of a subset of cells from the semispan domain, each subdomain mesh still preserves the boundary layer mesh properties over the multielement wing.

Grid	Domain Size (m)			No. of Nodes	No. of Cells
	X	Y	Z		
Grid1	12	6	8	1,329,545	4,855,817
Grid2	18	6	16	1,358,739	5,021,641
Grid3	24	6	24	1,369,465	5,080,469
Full	-	-	-	16,182,990	53,038,427

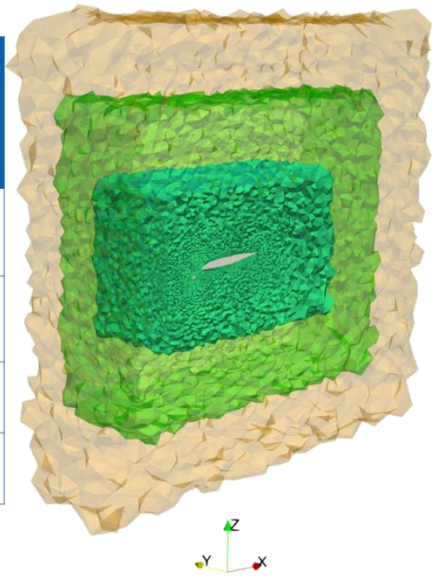

Figure 5-3. Summary of initial outboard subdomains extracted using `vogcut` utility.

Pressure coefficient line plots shown in Figure 5-4 demonstrate that excellent agreement was obtained between `vogcut` subdomains and full semispan solutions in the middle and the two spanwise extremes of the subdomain (see typical spanwise cut planes in Figure 5-2).

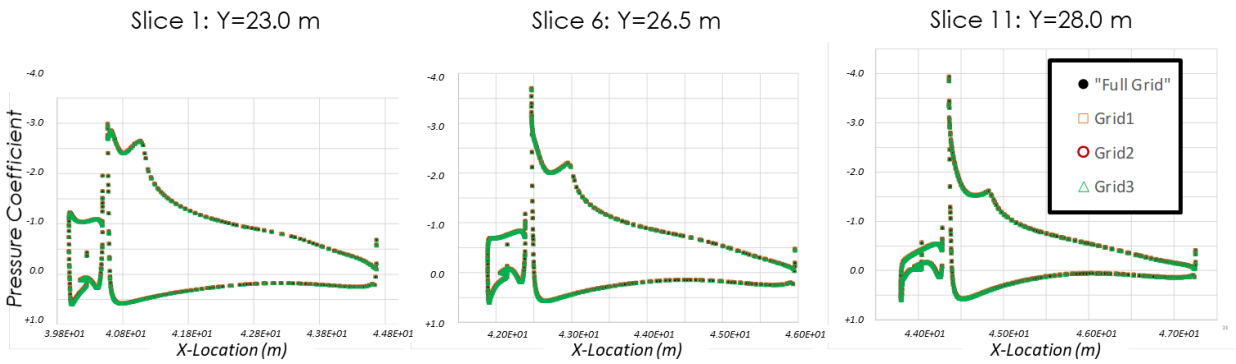

Figure 5-4. Subdomain model pressure coefficient plots reproduce semispan CFD solution ("Full Grid") nearly identically.

5.1.2. Non-Point-Matched Subdomain Meshes

An alternative approach for creating the subdomain models that was examined was to explicitly generate, rather than extract, subdomain meshes. In this study, ATA created the subdomain meshes using Pointwise with a structured mesh topology. The subdomains extracted using `vogcut` provide a precise point-matched subdomain boundary with the full domain, making the interpolation of the flow solution from full domain to the subdomain boundary trivial. The objective of this effort was to determine whether a non-point-matched subdomain mesh affects the interpolation and thus the solution or if a precise match of subdomain boundaries is needed to obtain good agreement between the subdomain CFD solution and the full semispan CFD solution. As shown in the upper-right corner of Figure 5-5, structured meshes that were generated used the same dimensions as the `vogcut` subdomain models and were of a size similar to their `vogcut` counterparts.

Figure 5-6 suggests there was good agreement between non-point-matched subdomain grids and full-domain solutions generated using the unstructured *l10wu5* grid. There is a subtle discrepancy for the non-point-matched breakout model geometry observed at the leading edge, as shown in Figure 5-7 at spanwise location Y = 23 m. Upon further inspection, a discrepancy in the model geometry was found and was thought to be the reason for differences in the pressure coefficient and wall shear stress, i.e., a more consistent geometry would have yielded excellent agreement in a fashion similar to the `vogcut` subdomains.

Grid	Domain Size (m)			No. of Cells
	X	Y	Z	
Breakout1	12	6	8	4.5M
Breakout2	18	6	16	5M
Breakout3	24	6	24	5.5M

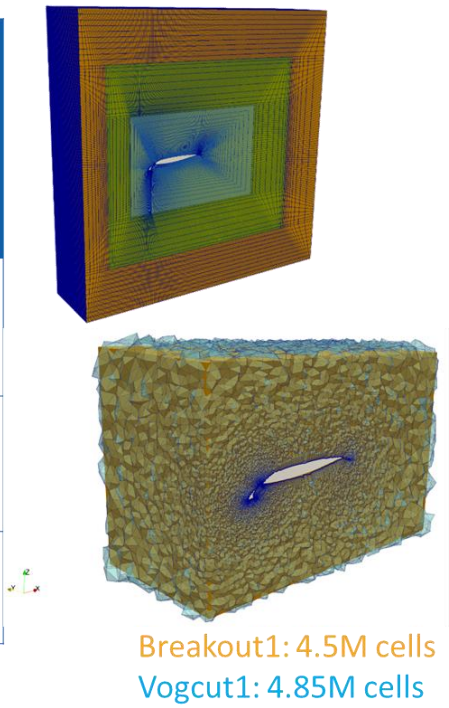

Figure 5-5. Summary of structured outboard non-point-matched subdomains.

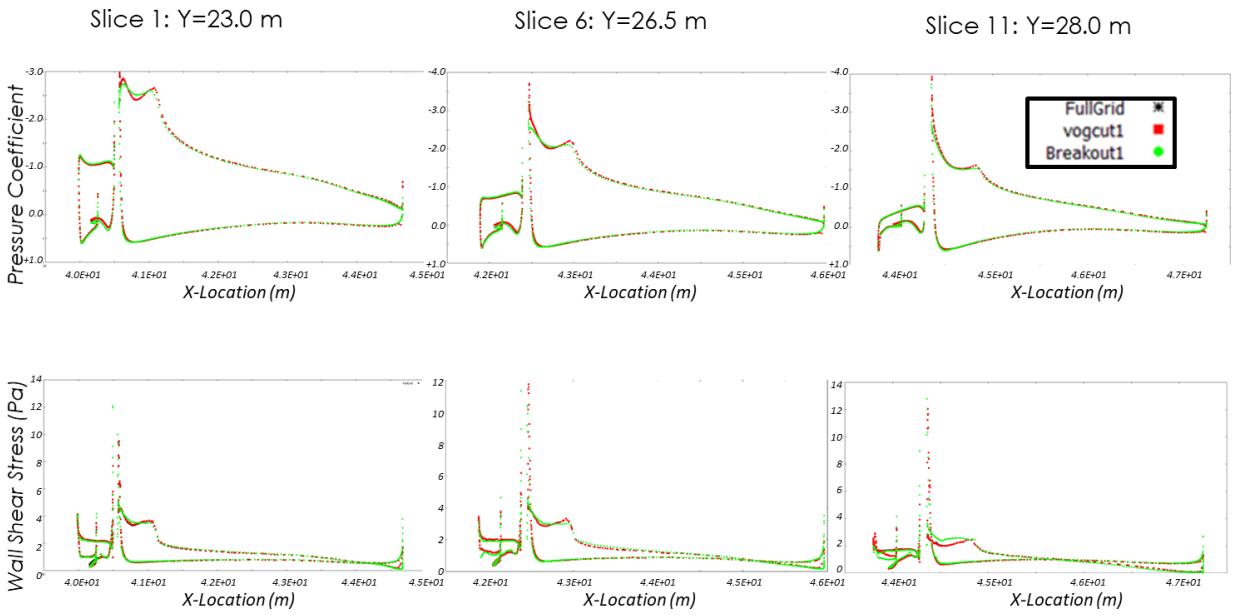

Figure 5-6. Good agreement in pressure coefficient and wall shear stress distribution at three spanwise locations. Comparison is between full semispan model ("FullGrid"), `vogcut` mesh ("vogcut1") and non-point-matched breakout model ("Breakout1").

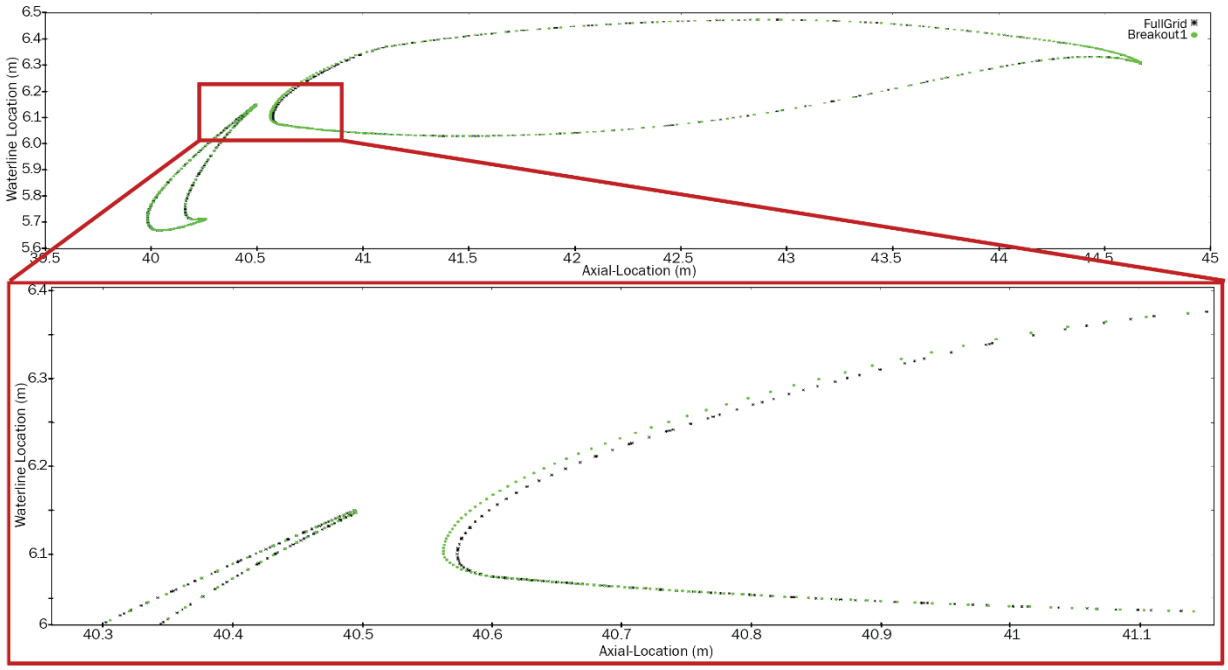

Figure 5-7. Model geometry discrepancy at Y = 23.0 m spanwise location. Black symbols pertain to the full grid; green symbols are from the Pointwise breakout model.

5.1.3. Refined Semispan Mesh Outboard Subdomain Analysis: *crm-hl-2018_04_13*

The outboard subdomain study was repeated using the *crm-hl-2018_04_13* mesh that contained embedded subdomains, which were rectangular with planar boundaries. However, the simulations using the embedded subdomains were unsuccessful, as the subdomain solutions developed local hotspots near the planar subdomain boundaries, and the solution diverged. A variety of mitigation strategies were attempted, including changing the size of the interpolation stencil used at the subdomain boundaries, switching to the shear stress transport (SST) turbulence model, decreasing the CFL (Courant number), and performing additional iterations on the full-domain solution before attempting the subdomain simulations, but none were successful.

To help ascertain whether the problem was due to the subdomain grid or to the full-domain solution, various combinations of subdomains with full-domain solutions were attempted. Subdomain simulations using the *10wu5* subdomains (generated using `vogcut`) and the *crm-hl-2018_04_13* full-domain solution were successful. However, the simulations using the *crm-hl-2018_04_13* subdomains with the *10wu5* full-domain solution were unsuccessful. This suggests that the problem is with the embedded *crm-hl-2018_04_13* subdomain meshes and not the *crm-hl-2018_04_13* full-domain solution.

To further test this hypothesis, `vogcut` was used to generate a subdomain mesh from the *crm-hl-2018_04_13* mesh. A spherical domain was defined, and the diameter was selected so that the farfield subdomain boundaries would not be near the embedded planar boundaries. A comparison of the spherical subdomain and the planar embedded boundaries is shown in Figure 5-8. The subdomain simulation using

the `vogcut`-generated spherical subdomain and the *crm-hl-2018_04_13* full-domain solution was successful, providing further evidence that the problem was with the embedded subdomain meshes. Figure 5-9 shows a comparison of the various subdomain simulations using different full-domain solutions. Some minor discrepancies (indicated by arrows) were observed and are thought to be due to mesh resolution differences.

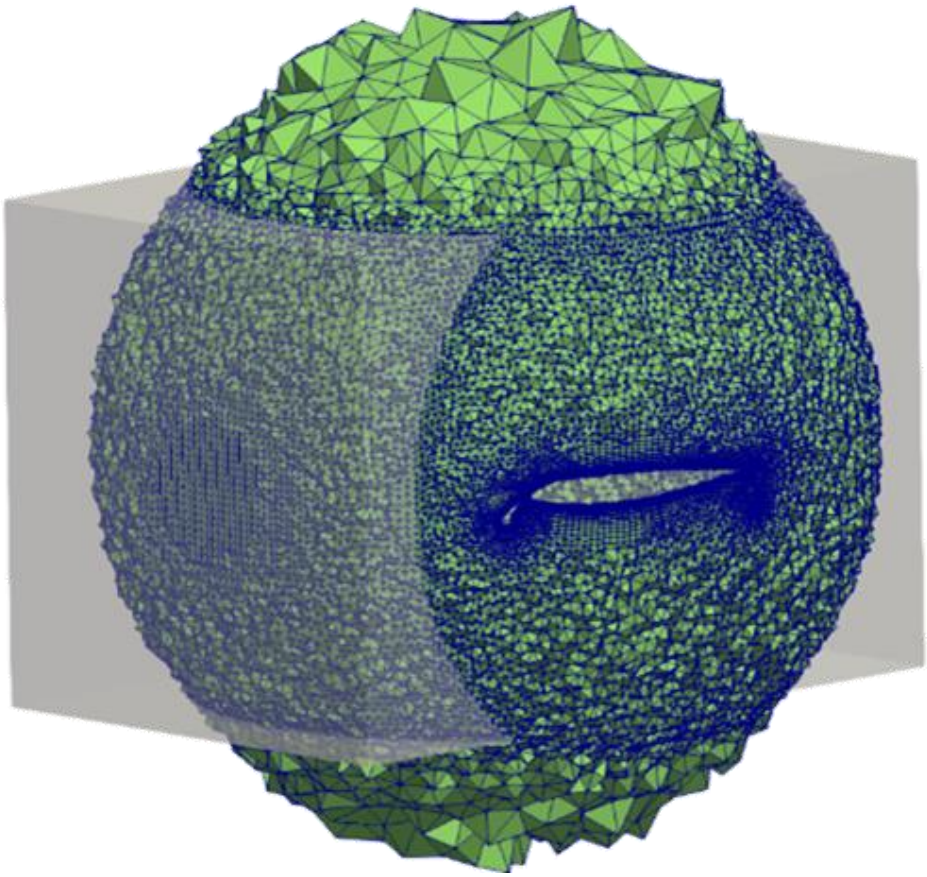

Figure 5-8. Comparison of the outboard subdomain generated using `vogcut` with the *crm-hl-2018_04_13* embedded subdomain mesh.

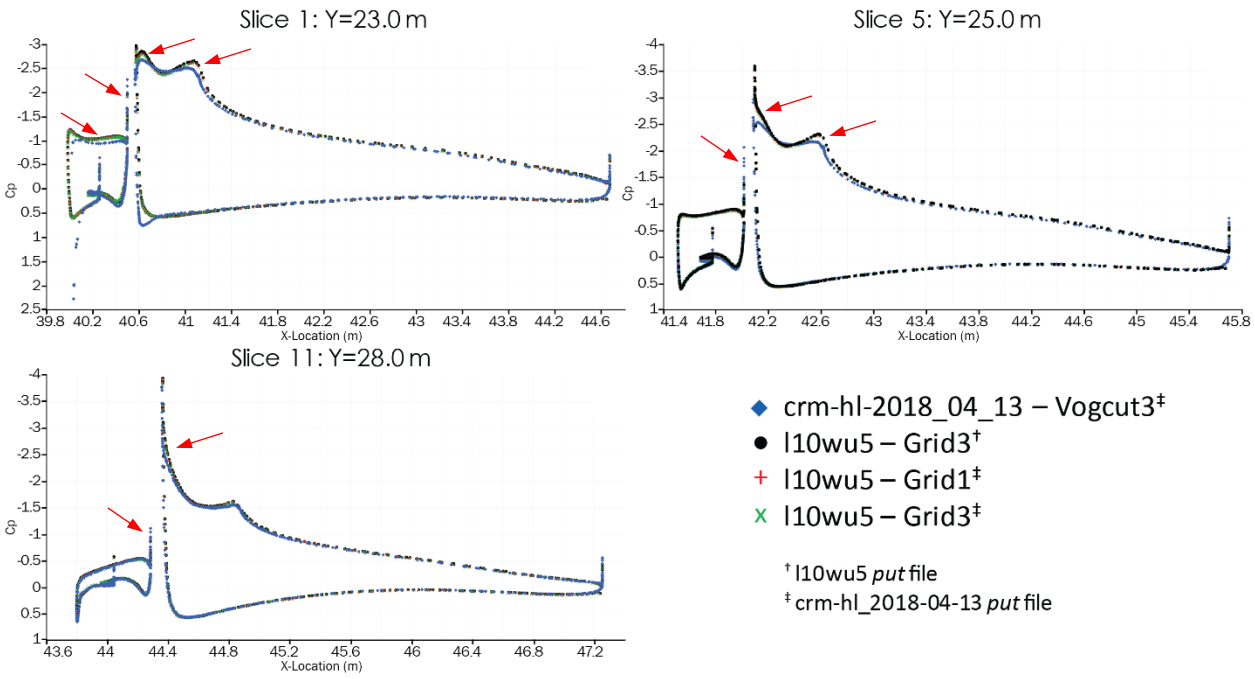

Figure 5-9. Similar good agreement seen using spherical `vogcut` subdomain on refined semispan *crm-hl-2018_04_13* mesh. Discrepancies indicated by arrows.

5.2. Inboard Subdomain

The inboard subdomain was both a more relevant and a more challenging subdomain example. It is more relevant because the 3D SGF technology element is most highly loaded in this region between sections 6 and 7 of the slat (see section 6.1.1.3). The challenges stem from the fact that this section is located in proximity to the engine nacelle, leading to a more locally three-dimensional flow field as well as increased areas of transition between prismatic boundary layer cells and isentropic core mesh cells.

5.2.1. Initial Inboard Subdomain Analysis: *l10wu5* mesh

A series of studies were conducted to establish the appropriate limits of the bounding box so that they are far enough away from the flexible SGF to capture important flow phenomena such as separation off the nacelle and flaps, while being close enough to result in a significant reduction in mesh size. The outboard subdomain study described in section 5.1.1 was repeated for the inboard region of the model where the most important (heavily loaded) technology elements are located. As part of the inboard subdomain study, three subdomain sizes were examined. Similar to the outboard subdomains, the inboard subdomains were generated from *l10wu5* grid using `vogcut` and are shown in Figure 5-10.

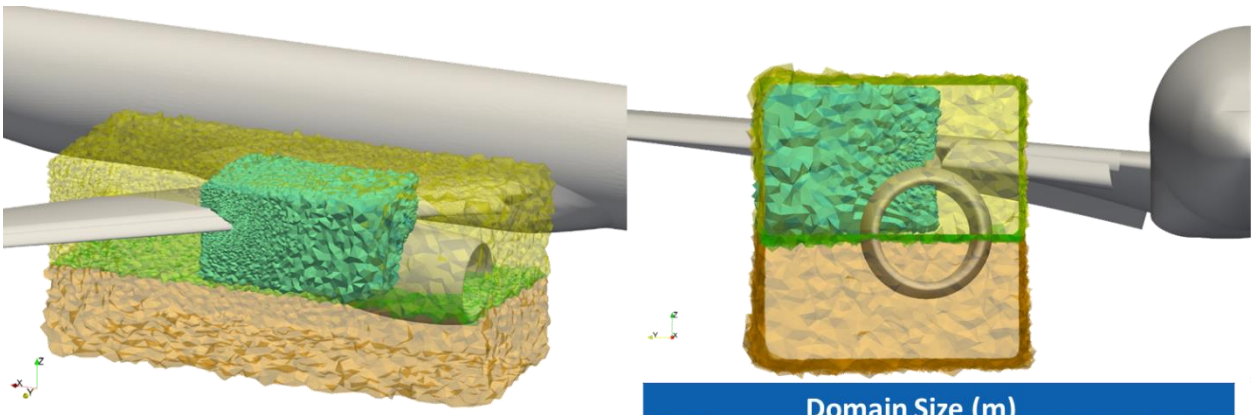

Grid 1: Cuts through nacelle and wing
Grid 2: Cuts through nacelle
Grid3: Contains entire nacelle

Grid	Domain Size (m)			No. of Cells
	X	Y	Z	
Grid1	7.6	5.1	4.0	4.5 M
Grid2	19.5	7.6	4.5	5.0 M
Grid3	19.5	7.6	8.2	5.5 M
Full	-	-	-	53 M

Figure 5-10. Three inboard subdomains obtained using `vogcut` from the l10wu5 mesh.

To guide the assessment, the pressure coefficient and shear stresses were computed at multiple spanwise locations. Locations are given in Table 5-1 and illustrated in Figure 5-11. Each subdomain has different intersection boundaries with the wing and nacelle surfaces. It is also noted that the chordwise extent of *Grid1* does not extend to the trailing edge. The spanwise extent of the SGF technology element is approximately between locations 4 and 6 shown in Figure 5-11.

Comparisons between the pressure coefficients obtained using the full domain and the subdomain solutions are shown in Figure 5-12. In general, very good agreement is observed at the leading edges and over the main portion of the wing at all three locations. In the most outboard section, spanwise location 8 in Figure 5-12, there are some discrepancies in pressure coefficient near the flap trailing edge. This pressure difference was not of concern since it is on the trailing edge of the wing and sufficiently far downstream of the technology elements. Additionally, there are discrepancies associated with the upstream and downstream boundaries of *Grid1*, suggesting it is favorable to have a mesh whose upstream and downstream boundaries do not intersect the wing or nacelle.

Wall shear stresses at the same spanwise stations are shown in Figure 5-13. In general, the inboard shear stress comparison follows same trends as the C_p comparison, in that they agree well. As shear stress relies on the velocity gradient in the boundary layer, it is to be expected that some additional discrepancies may be seen over areas of high curvature. Again, these discrepancies were not of concern as they were deemed unlikely to impact the aeroelastic response of a technology element.

Table 5-1. Spanwise locations for pressure coefficient comparison.

Slice	Y-Location (m)	Slice	Y-Location (m)
1	7.5	5	11.5
2	8.5	6	12.5
3	9.5	7	13.5
4	10.5	8	14.5

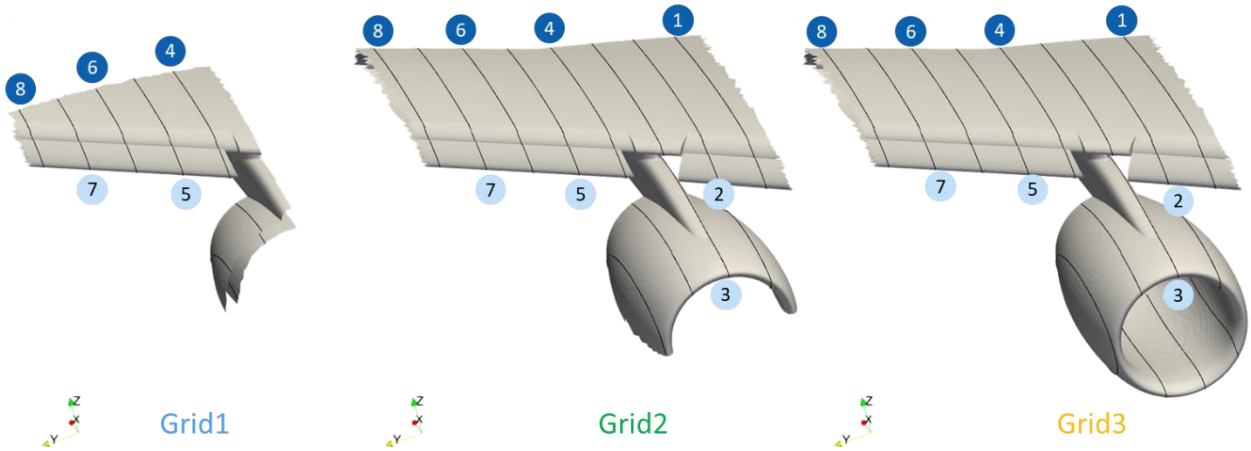

Figure 5-11. Spanwise locations used for the pressure coefficient comparison between the full domain and the subdomain solution.

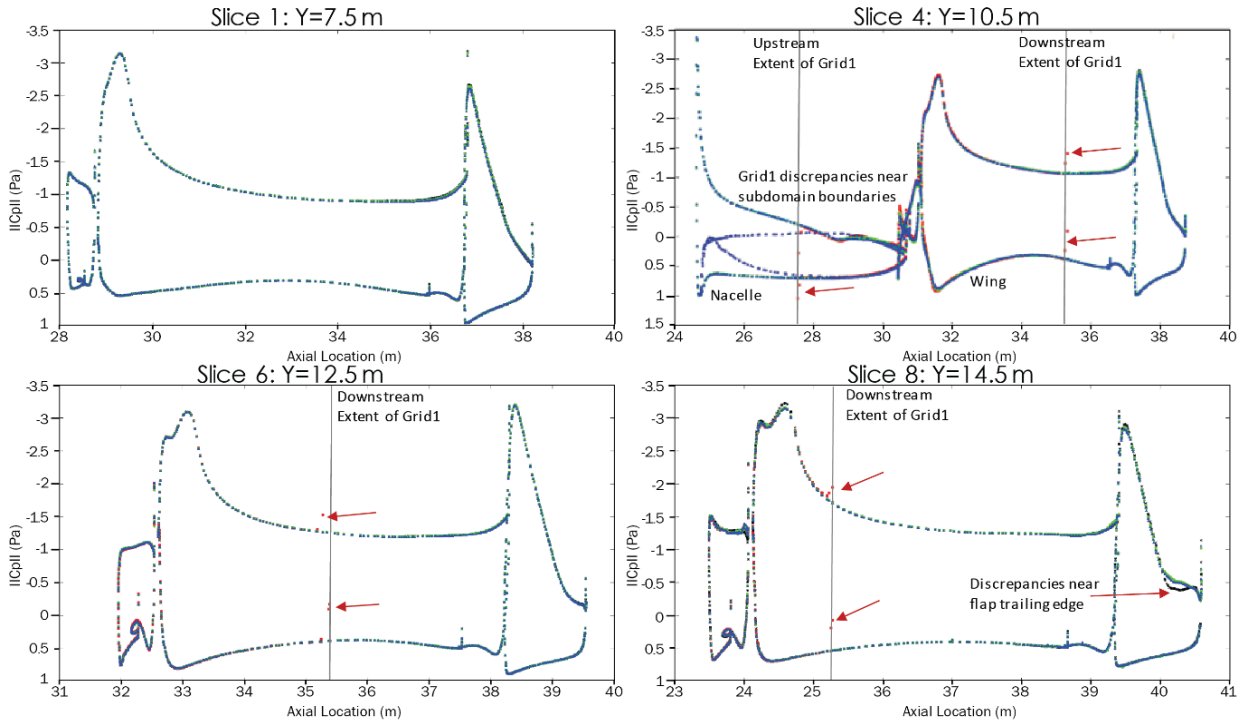

Figure 5-12. Pressure coefficient comparisons between the full grid and subdomain models.

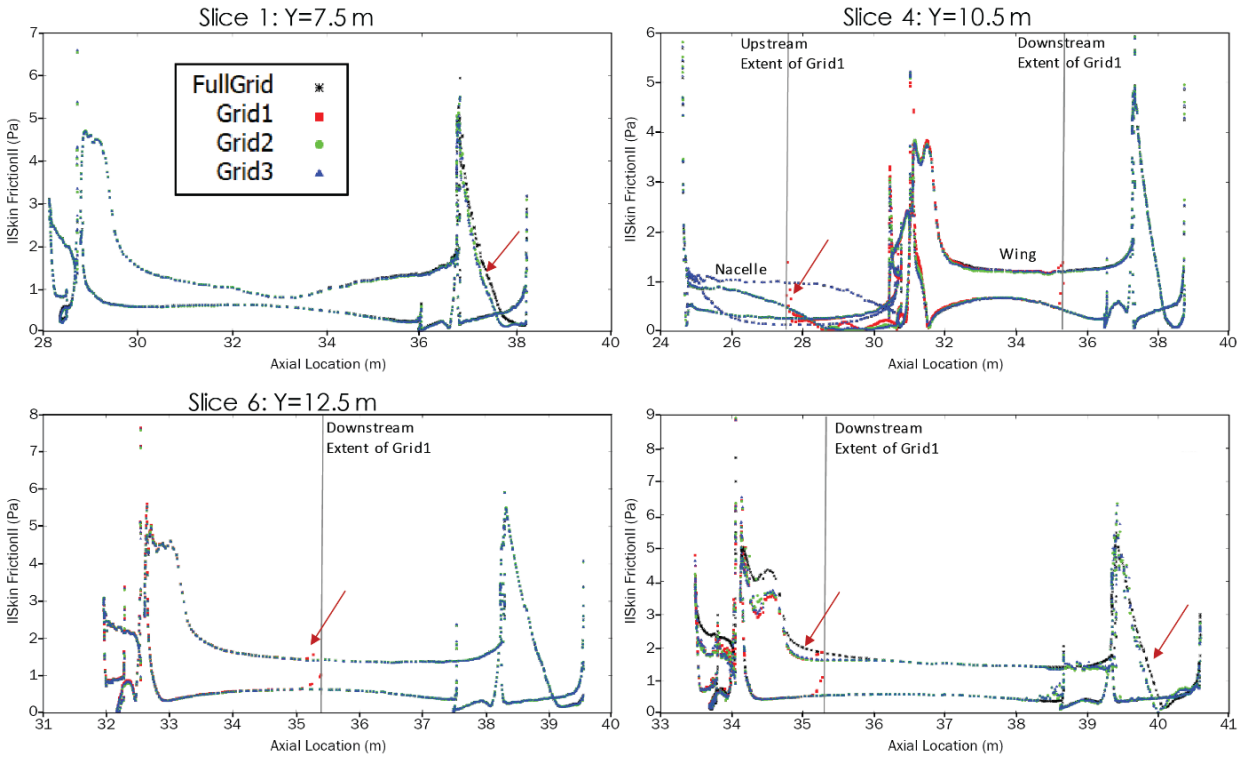

Figure 5-13. Wall shear stress comparisons between the full grid and subdomain models.

5.2.2. Refined Semispan Mesh Inboard Subdomain Analysis: *crm-hl-2018_04_13*

The inboard subdomain analyses were attempted using the refined *crm-hl-2018_04_13* mesh. Similar to the outboard subdomain simulations using the embedded subdomains, the simulations using the embedded inboard subdomains were unsuccessful. Due to the difficulty in obtaining subdomain solutions using this mesh, a new semispan mesh, referred to as *crm-hl-2018_06_26*, was provided by NASA. The *crm-hl-2018_06_26* mesh has a surface mesh resolution similar to the *crm-hl-2018_04_13* mesh but contains fewer nodes (31,950,320 vs. 33,553,419) and fewer elements (70,263,900 vs. 80,583,826). The most notable difference between the two meshes is that the *crm-hl-2018_06_26* mesh does not contain embedded boundaries. A comparison of the field resolution is provided in Figure 5-14. The subdomain analysis using the *crm-hl-2018_06_26* mesh is described in the following section.

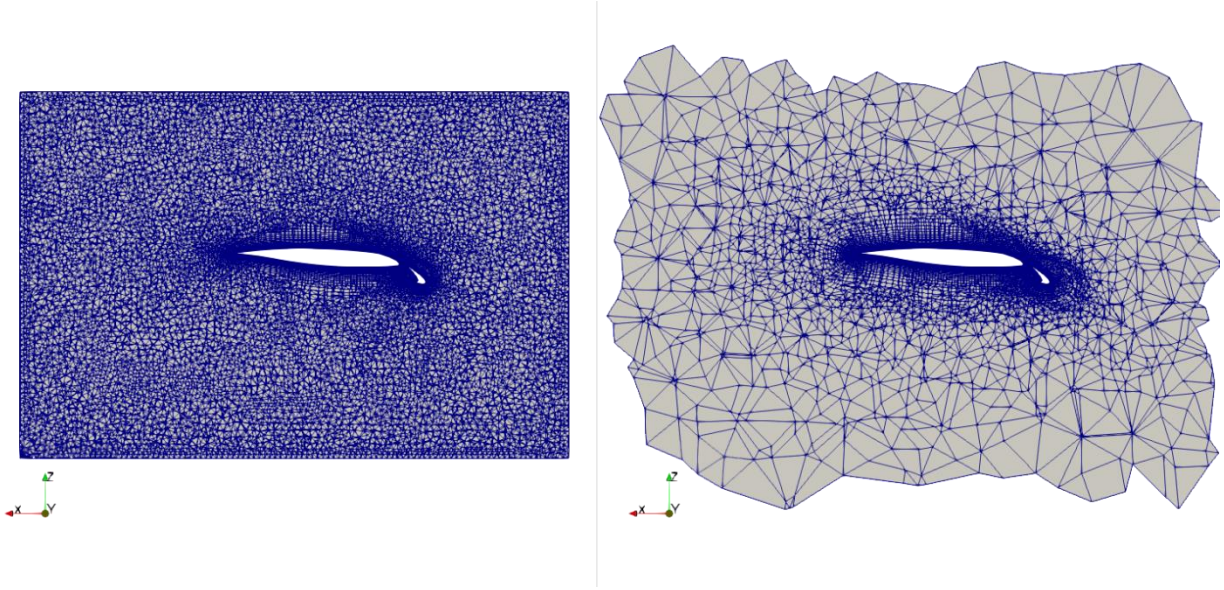

(a) (b)

Figure 5-14. Comparison of the field resolution between the (a) *crm-hl-2018_04_13* mesh and the (b) *crm-hl-2018_06_26* mesh.

5.2.3. Refined Semispan Mesh Inboard Subdomain Analysis: *crm-hl-2018_04_26*

The inboard subdomain analyses were repeated using refined mesh *crm-hl-2018_06_26*. A subdomain was generated using `vogcut` with dimensions of the original inboard *Grid3* subdomain (19.5 m × 7.6 m × 8.2 m). As shown in Figure 5-15, pressure coefficient profiles agree well between subdomain and full models at slices 1, 5, and 6. Figure 5-16 shows the same profiles for slices 7 and 8. Here, minor differences are seen on the pressure side of inboard flap near the subdomain boundaries; otherwise, the agreement is very good and suitable for FSI simulations of the SGF technology element. It should be noted that simulations for the outboard subdomain were also successful but, for brevity, the results are not included in this report.

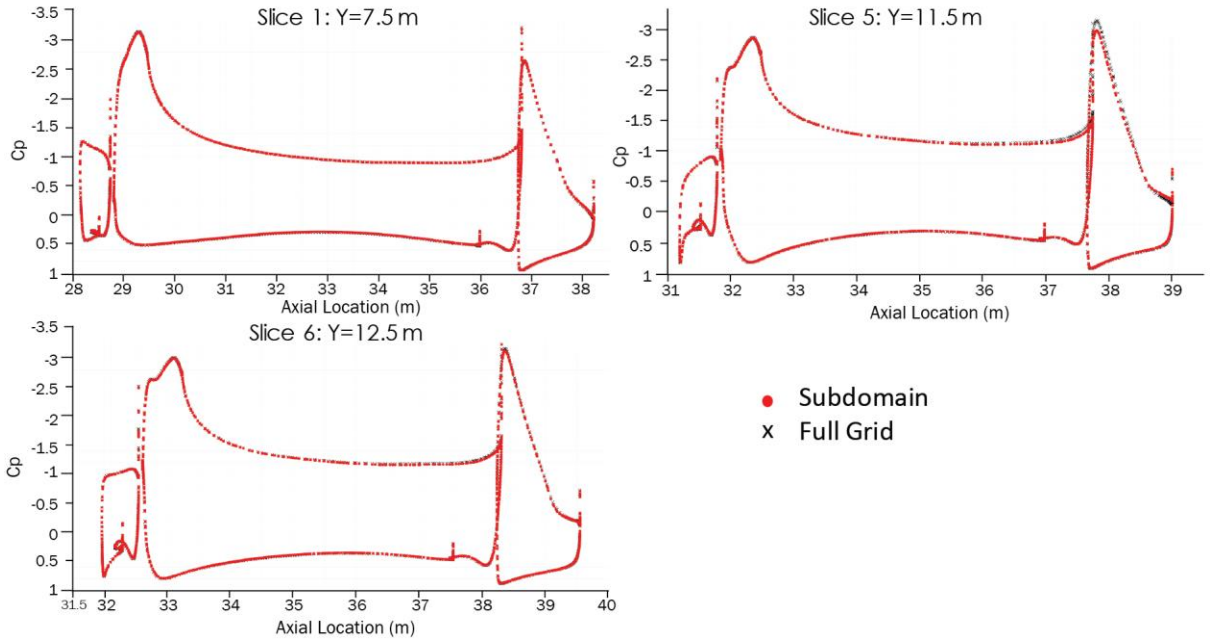

Figure 5-15. Pressure contours on slices 1, 5, and 6 of *Grid3* subdomain model using crm-hl-2018_06_26 mesh.

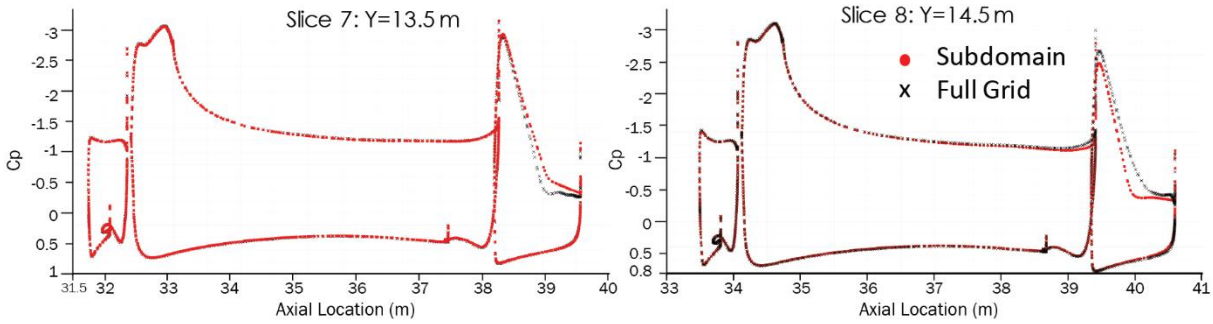

Figure 5-16. Pressure contours on slices 7 and 8 show minor differences on the pressure side of inboard flap near the subdomain boundaries, but otherwise good agreement in the slat region.

5.3. Unsteady Verification

In addition to the steady-state subdomain simulations, a series of unsteady simulations were performed using both the *l10wu5* inboard and outboard subdomains. The goal of the unsteady simulations was to verify that the subdomain solution would exhibit no hysteresis and would return to the same nominal solution as a full semispan model. The subdomain models were tested with an unsteady simulation with a localized pressure release as a time-dependent source term. As a further stress test, the smallest subdomain models (*Grid1)* for both the inboard and outboard subdomains) were selected, as the subdomain farfield boundaries are closer to the wing and source of the pressure disturbance, so any potential reflections or interference with the subdomain boundaries would be most apparent. The location of the pressure disturbance in each subdomain is shown in Figure 5-17.

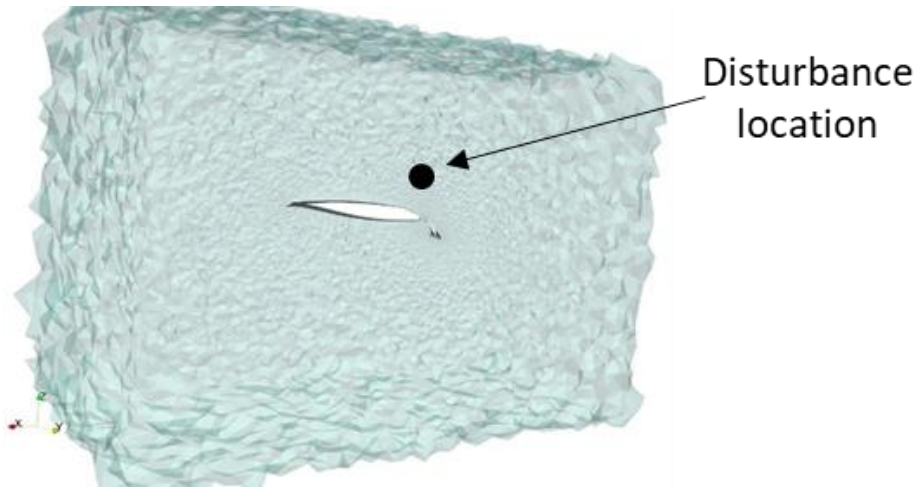

(a)

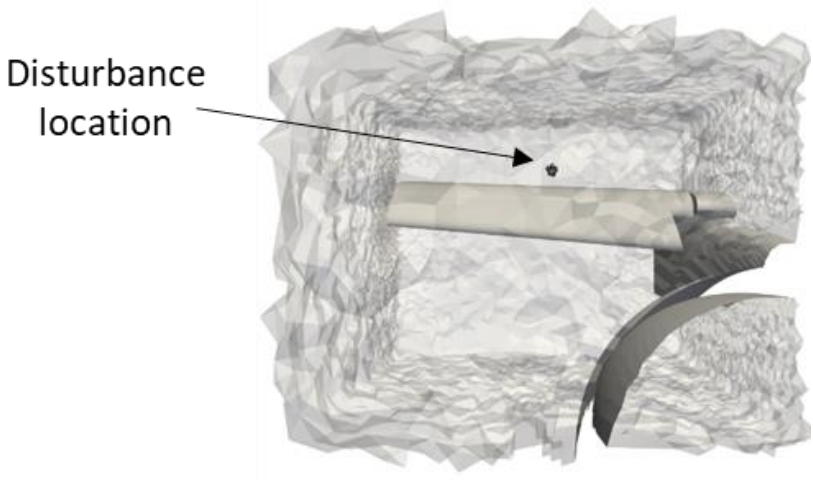

(b)

Figure 5-17. Location of the localized pressure release disturbance in the (a) outboard and (b) inboard subdomains.

To establish a baseline, unsteady simulations were performed for each subdomain without the pressure disturbance to measure the inherent unsteadiness in flow. It was observed that there was no substantial unsteadiness at the outboard location and only a small amount of unsteadiness at the inboard location. Based on this observation, it was concluded that the inherent unsteadiness would not affect the quasistatic or dynamic aeroelastic responses.

To verify that the subdomain solution would exhibit no hysteresis and would return to the same nominal solution, the unsteady simulations were repeated with the small pressure disturbance that was generated above the wing suction surface. The disturbance provided about 40 Pa of over- and under-pressure. The disturbance was initiated at 0.1 s and had a duration of 0.076 s. The flow fields were monitored after the disturbance passed as an acoustic disturbance; specifically, points directly below the pressure disturbance

on the suction side of the wing near the wing leading edge in both the inboard and outboard subdomains were monitored. The pressure time history for the point in the outboard subdomain is shown in Figure 5-18. The pressure fluctuation due to the disturbance is evident, but the pressure returns to the initial steady-state value. The pressure time history for the point in the inboard subdomain is shown in Figure 5-19. As can be seen, there is an initial unsteadiness in the flow due to the monitored point being immediately downstream of the nacelle. Similar to the point monitored in the outboard subdomain, the pressure disturbance is clearly evident, and the flowfield solution eventually returned to its pre-disturbance response. The unsteady pressure disturbances were shown not to interact with subdomain boundaries, so there is no concern of the pressure fluctuations caused by the deforming technology elements having any sort of interference with the subdomain boundaries.

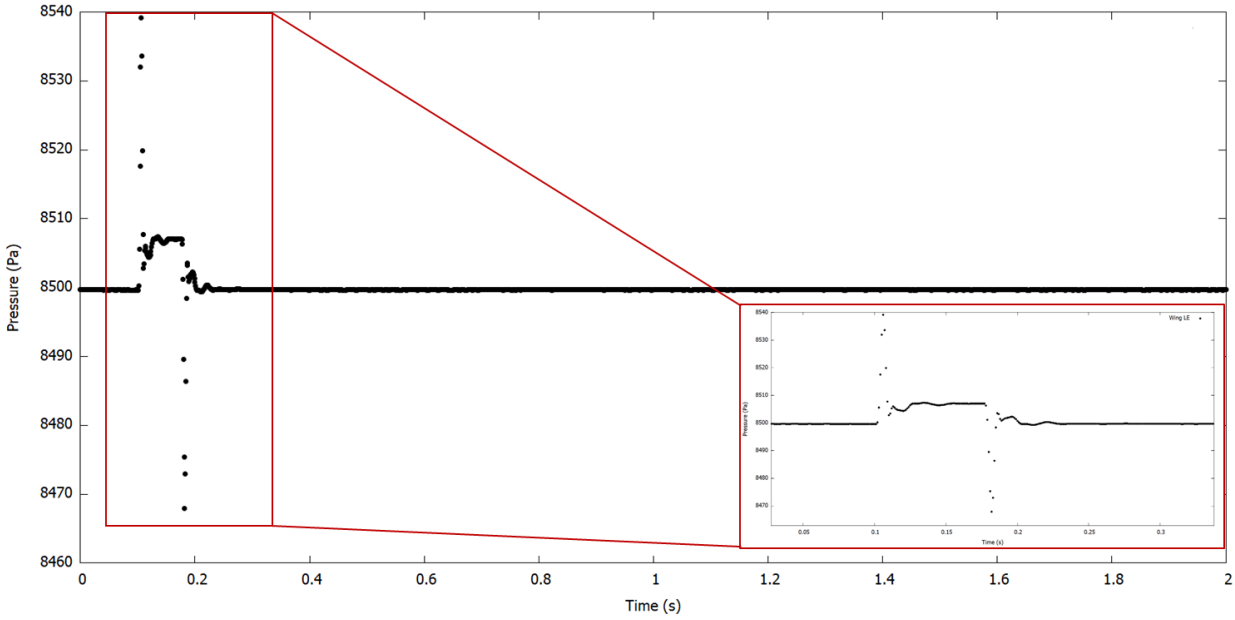

Figure 5-18. Outboard subdomain pressure time history subject to a localized pressure disturbance.

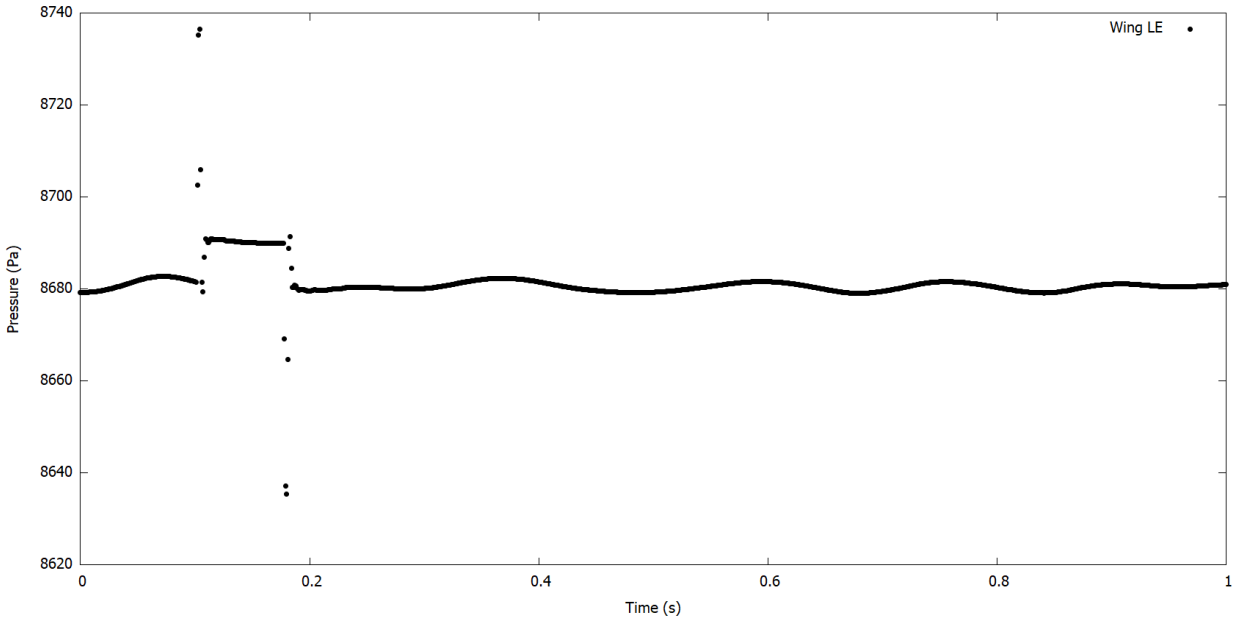

Figure 5-19. Inboard subdomain pressure time history subject to a localized pressure disturbance.

5.4. Summary and Lessons Learned

Modeling studies using rigid geometry suggest that the subdomain approach is a viable path forward for FSI simulations. ATA was able to successfully apply the methodology to inboard and outboard sections of CRM-HL configuration meshes. Both steady and unsteady behavior were examined using subdomain models.

In general, subdomain results (Cp and τ) using `vogcut` compare well to full-grid solutions. Notable discrepancies were found near subdomain boundaries, suggesting that the location of the subdomain farfield boundaries (i.e., the size of the subdomain) could affect the solution. Based on the grids studied, the largest grid (*Grid3*) was selected as the inboard subdomain for FSI simulations.

The subdomain modeling strategy was also found to be unaffected by unsteady disturbances from a local pressure release. Flow fields were found to return to their pre-disturbance solution without any evidence of hysteresis.

Several conclusions and recommendations are proposed for subdomain modeling in general:

- It is recommended to use a subdomain similar to the largest (*Grid3*) used in these studies, which was well upstream of the nacelle and downstream of the wing trailing edge, and whose only intersection with surface geometry was at spanwise stations along the wing far from the technology element.
- It is preferable to have good mesh quality in regions near subdomain boundaries.

- There is no specific advantage to using embedded surfaces within the full domain volume mesh to define subdomains; if embedded surfaces are to be used, it may be advantageous to define the `vogcut` bounding box inside the prescribed internal boundary surfaces.

6 3D AEROELASTIC ANALYSIS ON CRM-HL WITH TECHNOLOGY ELEMENTS

This section documents the final model and results for the SGF FSI analyses and also documents the final status for the continuation of work on the 3D SCF FSI assessment.

6.1. Slat-Gap Filler

This section summarizes the full 3D aeroelastic analysis that was performed for the SGF technology element. This effort began with receipt of NASA-generated models of the structure and semispan CFD domain. ATA first reviewed the models to ensure that they were collocated in space and geometrically similar enough to allow data mapping in cosimulation. Through this process, ATA determined that the CFD mesh required modifications for robust mesh deformation. ATA performed these modifications in the semispan domain. ATA also identified modifications of the NL-CSD model necessary to ensure robust FSI simulation. These modifications were made by Dr. Turner of NASA.

6.1.1. Description of Models

CFD and NL-CSD models for the CRM-HL geometry were created in Loci/CHEM and Abaqus. The models were used to evaluate the dynamic stability of the SGF under two flow conditions as shown in Table 6-1. The baseline condition corresponds to a nominal landing condition, whereas the high-q case corresponds to a landing condition in which there was a wind gust increasing the dynamic pressure. The dynamic pressure of the high-q case was chosen to nominally result in a deformation of four panel thicknesses using an estimated pressure differential between the pressure and suction sides of the SGF and assuming a uniform pressure distribution.

Table 6-1. Flow conditions for SGF aeroelastic assessment.

Case Name	Dynamic Pressure (Pa)	Mach	Reynolds Number (m^{-1})	Angle of Attack (degrees)
Baseline	2873.10	0.2	4.6e6	8
High-q	5390.49	0.2	4.6e6	8

Figure 6-1 presents an overview of the cosimulation setup. The NL-CSD model cosimulates with a subdomain CFD model that is informed by a vehicle-level semispan rigid CFD solution. The semispan model is run a priori and is used to supply the subdomain model with initial conditions and boundary conditions. The subdomain CFD and NL-CSD models are run concurrently during the cosimulation.

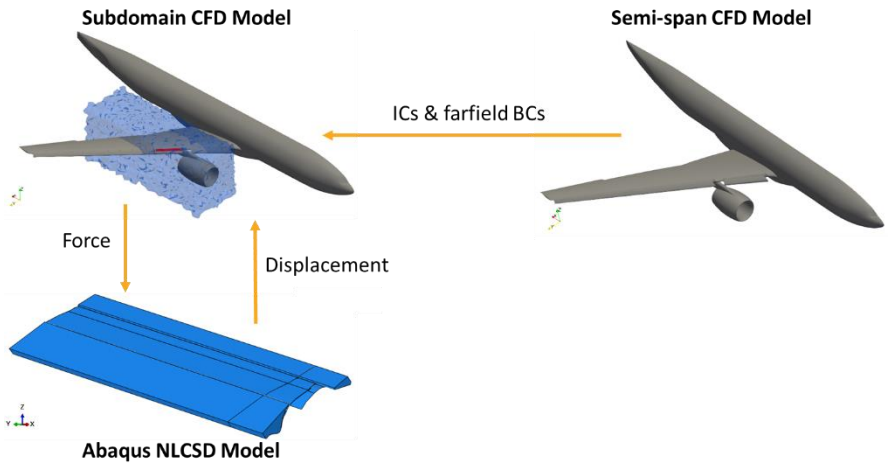

Figure 6-1. Cosimulation setup with semispan/subdomain CFD models and NL-CSD FEM.

6.1.1.1. Vehicle Semispan Flowfield CFD Computation

The CRM-HL semispan model treated with the SGF technology elements was used for a RANS CFD simulation at the baseline flow conditions using the Spalart-Allmaras turbulence model. The detailed geometry of this model with its myriad small edges and features made generating a high-quality unstructured CFD mesh challenging. Several problem areas were identified in the initial mesh developed by NASA. The small edges and features led to a collapse of the prismatic layers generated off the vehicle body during the meshing process. This situation in turn resulted in large volume ratios at the interface between the anisotropic prismatic layers and the isotropic core mesh, which resulted in stability problems for Loci/CHEM. Loci/CHEM's `vogcheck` utility graded the mesh as unusable due to high volume ratios between cells and nonplanar (twisted) cell faces as shown in Table 6-2.

Table 6-2. Assessment of initial CRM-HL semispan mesh from `vogcheck`.

	Volume Ratio	Face Angle	Twist	Shear Twist
Excellent	$x < 10$	$x < 90$	$x < 0.1$	$x < 0.1$
Good	$10 \le x < 50$	$90 \le x < 150$	$0.1 \le x < 0.2$	$0.1 \le x < 0.2$
Poor	$50 \le x < 100$	$150 \le x < 170$	$0.2 \le x < 0.45$	$0.2 \le x < 0.45$
Marginal Utility	$100 \le x < 1000$	$170 \le x < 179$	$0.45 \le x < 0.8$	$0.45 \le x < 0.8$
Unusable	$x \ge 1000$	$x \ge 179$	$x \ge 0.8$	$x \ge 0.8$

To alleviate this problem, a cell-quality remediation [27] module was developed for Loci/CHEM. The module will locally switch the solution to first-order spatial accuracy (instead of the default second order) in regions of poor cell quality. The local cell quality is judged by comparing the cell volume ratio and face maximum included angle to user-defined threshold values. In general, the regions of poor cell quality contain a very small number (<< 1%) of cells compared to the total number of cells in the domain. Furthermore, it was demonstrated that the local reduction in spatial order of accuracy has a negligible

effect on the key outputs of the simulation (e.g., pressure coefficient and skin friction coefficient on the vehicle body).

The pressure and skin friction coefficients on the wing surface were compared to a CFD solution previously generated in PowerFLOW, as shown in Figure 6-2. Qualitatively, the results showed good agreement and demonstrated Loci/CHEM's ability to accurately model the highly separated flow near the wing-nacelle junction.

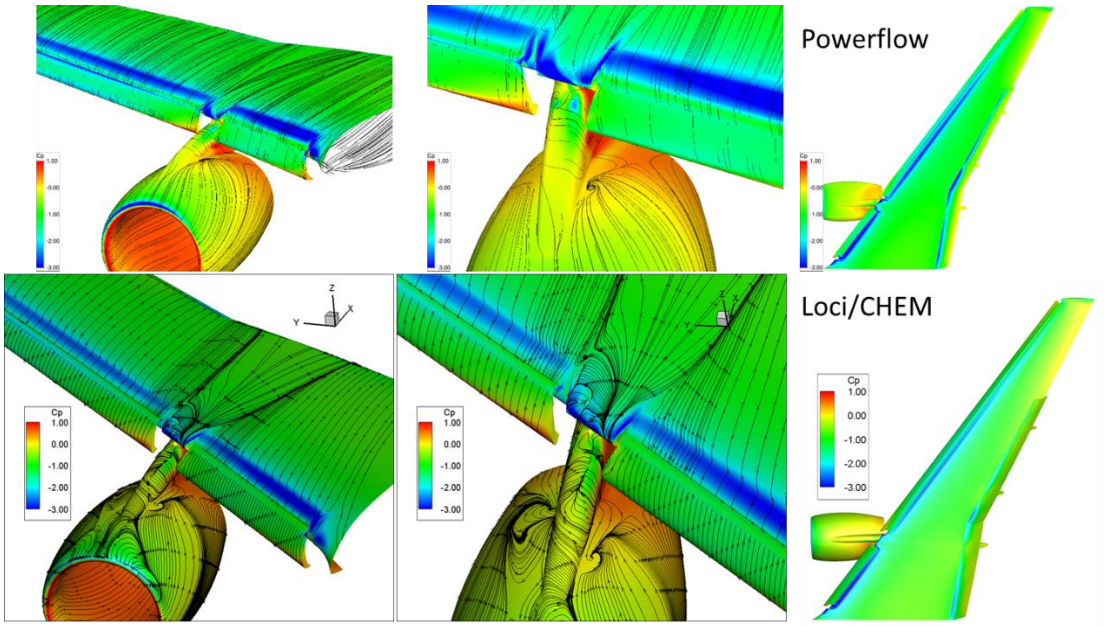

Figure 6-2. Pressure and skin friction coefficient comparison between Loci/CHEM and PowerFLOW.

6.1.1.2. Subdomain Breakout CFD Model Creation

The initial CRM-HL semispan mesh generated by NASA was unsuitable for FSI simulation because (1) it was a very large and therefore computationally expensive mesh and (2) the mesh quality near the SGF was not sufficient to support large deformations on the order of four panel thicknesses. In an effort to improve the local cell quality to support large mesh deformations, ATA created a structured mesh around the flexible portion of the SGF. The structured mesh interfaced to an unstructured mesh that covered the rest of the domain. To allow each mesh region to have the highest-quality cells possible, the interface at the structured/unstructured boundary was not point-matched; during simulation, Loci/CHEM interpolates the solution across the interface.

To reduce the cell count, the subdomain modeling approach was used to minimize the computational expense. For the SGF application, Loci/CHEM's `vogcut` utility was used to extract all cells that have nodes within a box defined around the flexible portion of the SGF. The subdomain, shown in Figure 6-3, results in an approximately threefold reduction in the total number of cells.

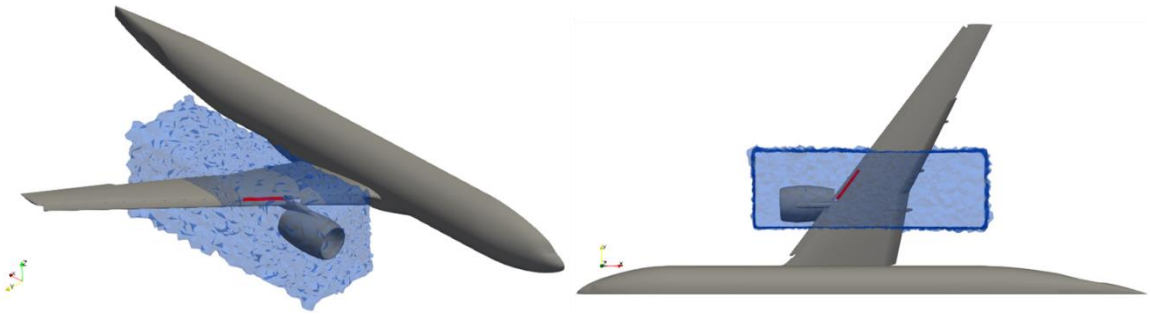

Figure 6-3. Flexible SGF region (red) and the CFD subdomain around it (blue).

To ensure that the subdomain model was accurately modeling the flow, the pressure coefficient along a spanwise section was compared between the semispan PowerFLOW simulation and the Loci/CHEM subdomain models. The subdomain models were run with the Spalart-Allmaras turbulence model as before, and with the k-ω SST model,[2] which was observed to have better stability for this simulation. In Figure 6-4, both the Spalart-Allmaras and k-ω SST subdomain simulations show good agreement with PowerFLOW. For numerical stability reasons, ATA decided to use the k-ω SST model for the FSI simulations.

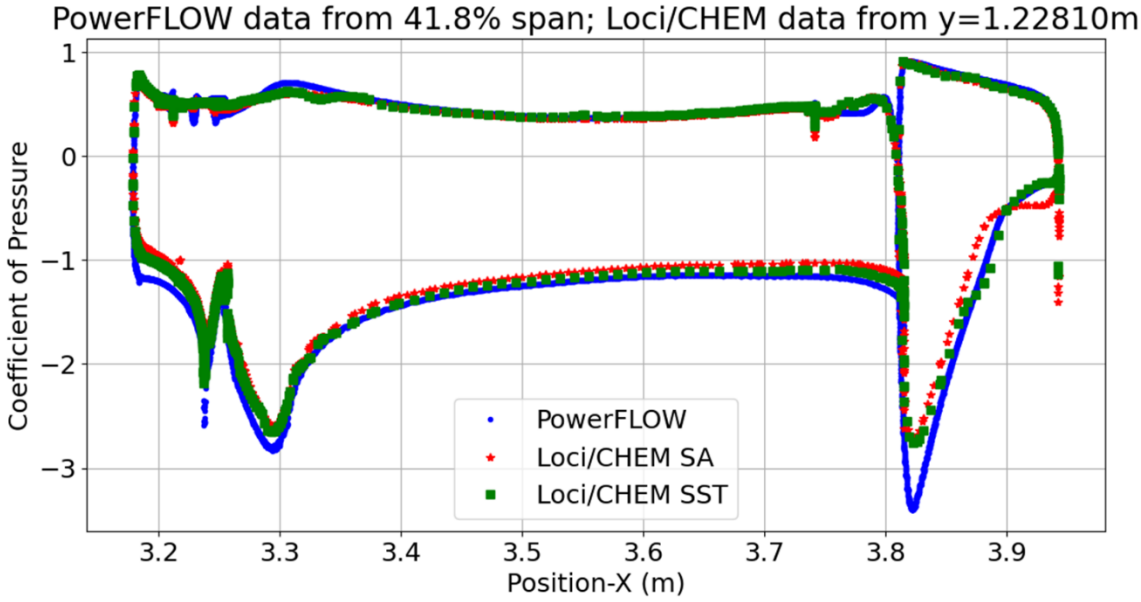

Figure 6-4. Loci/CHEM subdomain simulation with SST model suggests reasonable agreement with PowerFLOW.

[2] The additional transport equations in the model track the turbulent kinetic energy (k) and the specific dissipation rate (ω).

6.1.1.3. Description of the 3D FEM for the CRM-HL SGF Treatment in the Subdomain for FSI Assessment

The purpose of the FEM described in this section was to capture the essential behaviors of the flexible SGF treatment to acceptable fidelity without overburdening the FSI calculation. The main driving factor for simplification of the FEM stemmed from the requirement to limit the FSI computations, and thus the deformable structure, to within the boundaries of the subdomain. The computational subdomain encompasses the inboard end of the outboard slat including the first two sections (identified as 6 and 7 for the entire model) of SGF treatment on the outboard slat component. This region was selected because of the complex flowfield in this region, largely stemming from the presence of the nacelle and pylon, and the resulting increase in loading and unsteady phenomena around the slat and treatment components there. A CAD image showing the relevant section of the CRM-HL wing in the subdomain is shown in Figure 6-5.

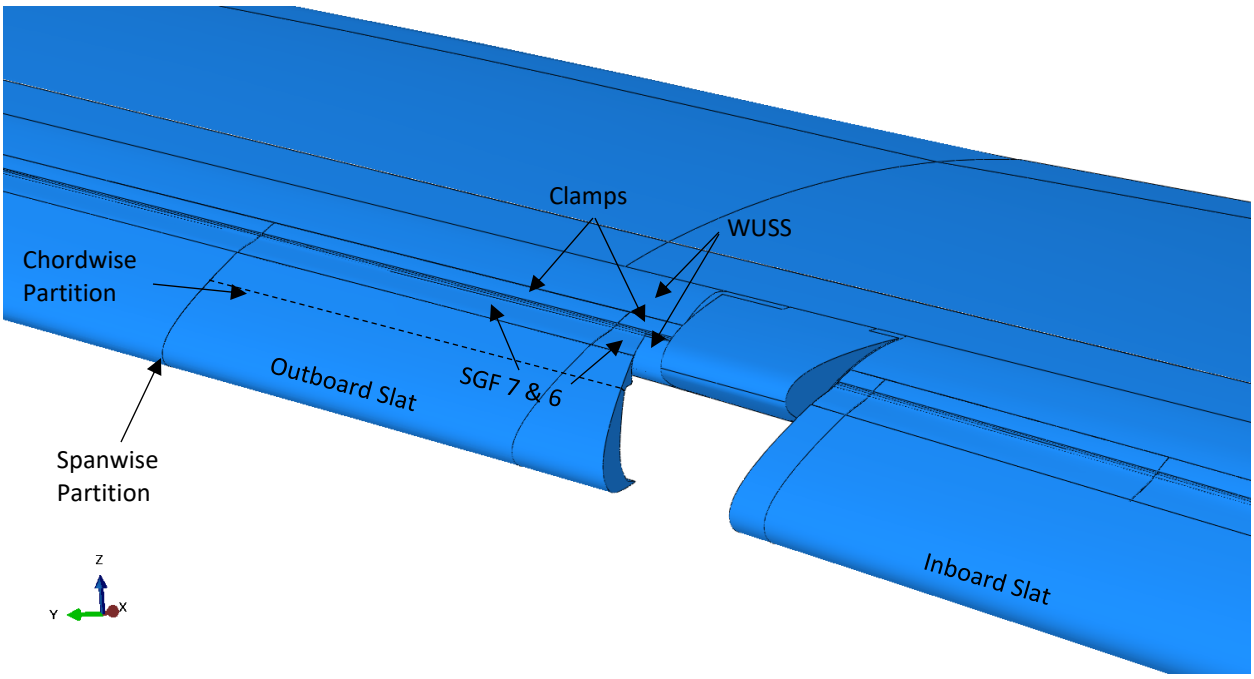

Figure 6-5. CAD image of CRM-HL wing in subdomain region.

The physical model was composed of a single contiguous slat component inboard of the nacelle and pylon and another outboard. The subdomain approach required grid deformation to be zero at its boundaries. Although several simplifications to the slat and wing under slat surface (WUSS) behavior were considered, the approach selected entailed partitioning the outboard slat, as shown in Figure 6-5, in the spanwise direction within the outboard limit of the subdomain. The partition was taken at the outboard end of SGF section 7: the section of outboard slat within the subdomain was considered for development of the FEM, and everything outboard was neglected. The loss in mass and stiffness influence of the slat section outboard of the subdomain was recognized, and the resulting model was considered conservative from a static deflection and stability standpoint. The outboard slat section within the subdomain was further simplified by neglecting all but the trailing section, as defined by the chordwise partition near and adjoining the joint with the SGFs in Figure 6-5, because it was believed that the elasticity of that section

of the slat had the greatest influence on the deflection and stability of the SGF treatments. All degrees of freedom were constrained at the root end of that elastic aft section of the slat to represent its connection to the rest of the slat. The WUSS was modeled as rigid, and only the outer surface in the region of possible interaction with the SGFs was retained in the FEM. Figure 6-6 depicts a cross-section of the model through SGF 7, showing the essential features captured in the FEM.

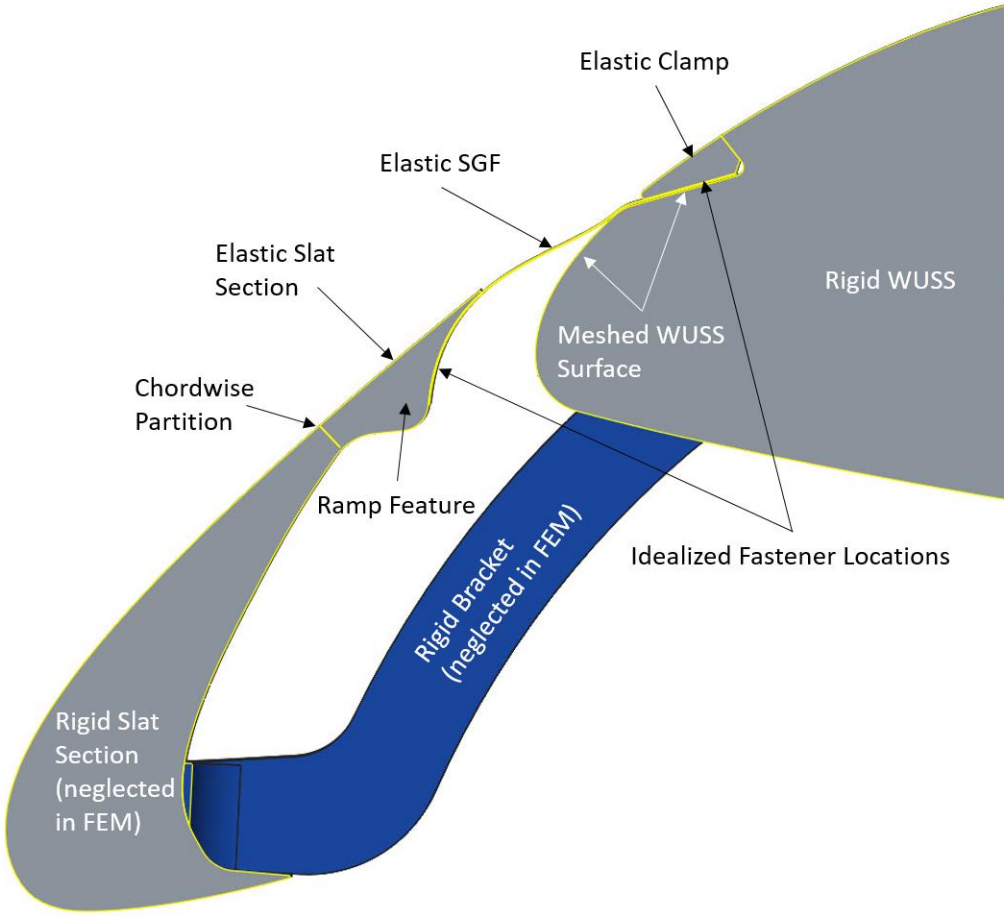

Figure 6-6. Cross-section through SGF 7 showing essential features included in FEM.

The physical model required fastened joints between the SGF treatment and the rest of the model for fail-safety. Limitations imposed by the small size at 10% scale forced inclusion of some features not characteristic of flight hardware, such as the ramp in the cove of the slat and the clamps on the WUSS for securement of the SGF sections, shown in Figure 6-6; however, these features were deemed unimportant because they do not significantly affect the flow. It was believed that adequate fidelity could be achieved by simplifying the computational model by deleting fastener holes and introducing point-based fasteners at the relevant locations.

Table 6-3. Properties of materials used in CRM-HL and FEM thereof. See Auricchio and Sacco [41] for definition of parameters.

Stainless Steel	
Property	Value
E	206.9 GPa
ν	0.26
ρ	7915 kg/m^3
Nitinol	
E_A, E_M	48.9 GPa, 40 GPa
$\nu_A=\nu_M$	0.33
H, Transf. ε	4.40%
σ^{Ms}, σ^{Mf}	471 MPa, 514 MPa
σ^{As}, σ^{Af}	196 MPa, 166 MPa
Tens./Comp.	Symmetric
Ref. Temp.	24°C
$C^A=C^M$	10 MPa/K
ρ	6480 kg/m^3

The FEM resulting from the simplifications described above is shown in Figure 6-7. The slat trailing section had a chord of approximately 24.94 mm (0.982″) along the outer surface at the juncture between SGF sections 6 and 7, it was meshed with linear hexahedral elements (Abaqus C3D8R) having an approximate global size of 1.143 mm (0.045″), and it was assigned the material properties of stainless steel listed in Table 6-3. Similarly, the clamps securing the aft section of the SGFs to the WUSS had a top-edge length of approximately 8.79 mm (0.346″, between SGF sections 6 and 7), it was also meshed with C3D8R elements having an approximate global size of 1.143 mm (0.045″), and it was assigned properties of stainless steel. The outboard edge of SGF section 6 and, similarly, the inboard edge of SGF section 7 had an unsupported length of approximately 18.26 mm (0.719″). The SGFs were meshed with continuum shell elements (Abaqus SC8R) having chordwise/spanwise dimensions of approximately 0.635/1.27 mm (0.025″/0.05″), and they were assigned the material properties of an SMA (NiTi) listed in Table 6-3. The built-in model in Abaqus for superelastic materials, which is based on the model developed by Auricchio and Sacco [41], was used to capture the large deformation behavior of the SMA.

Point-based fasteners were defined at the locations indicated in Figure 6-7 to join the SGFs with the slat trailing section. Similar fasteners were defined at the relevant locations between the clamps, SGFs, and WUSS surface. A contact-tie constraint was defined between the bottom of the clamps and the top of the SGFs. Surface-to-surface contact was defined between the SGFs and the rounded nose of the clamps on top and between the SGFs and the WUSS surface below. Finally, surface-to-surface contact was also defined between the tops of the SGFs and the underside of the slat trailing section.

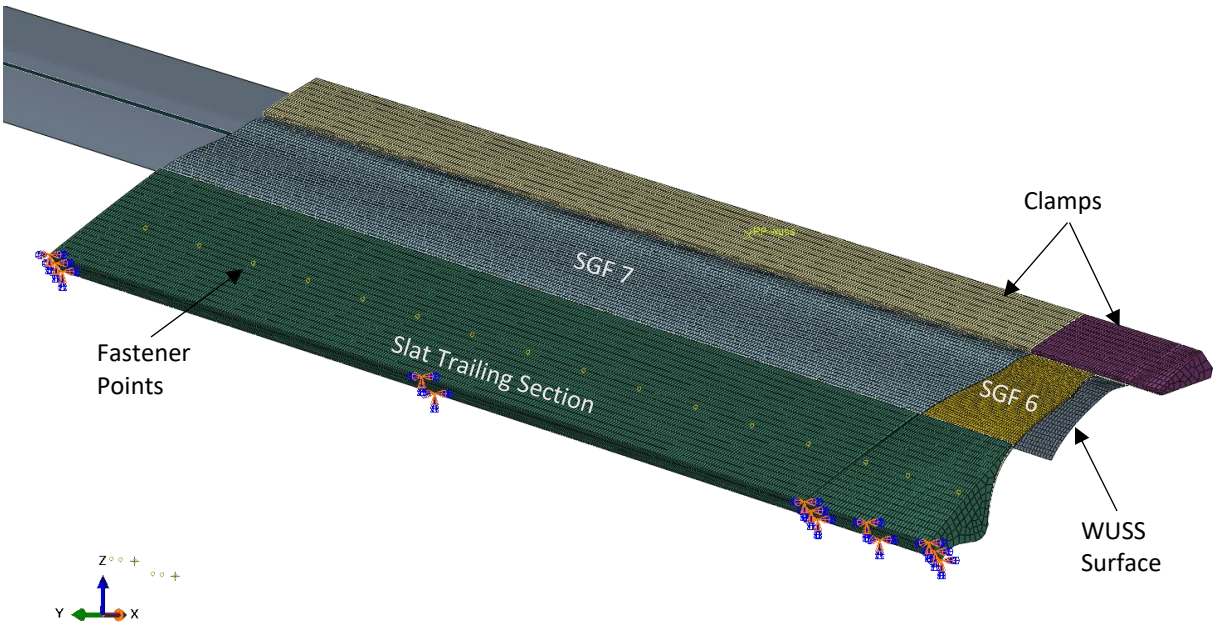

Figure 6-7. FEM of simplified SGF assembly within the computational subdomain.

There were regions of the fluid domain in the vicinity of the model where it was either difficult to grid or a gap between adjacent structural components could close or open due to structural deformation during the solution. These regions correspond mainly to the interactions between the SGFs and the slat trailing section, the SGFs and the WUSS surface below, and at the inboard ends of the aforementioned regions. These regions are highlighted in Figure 6-8. Fictitious surfaces were created in these areas, they were meshed with surface elements (Abaqus SFM3D4), and the meshed surfaces were tied along their edges to the adjacent structures to provide a deforming but ever-present boundary for the fluid domain. Steady pressure loads were defined between the SGF sections and the adjacent slat trailing section with the magnitude and spanwise distribution dictated by the steady pressure distribution at the trailing edge of the slat from the steady, semispan CFD solution. Steady pressure distributions were similarly defined between the SGF sections and the WUSS with the magnitude and distribution again coming from the steady, semispan CFD solution. Finally, pressure loads were defined for the top and bottom surfaces of the SGF components for transfer of transient pressure loads from the CFD solver.

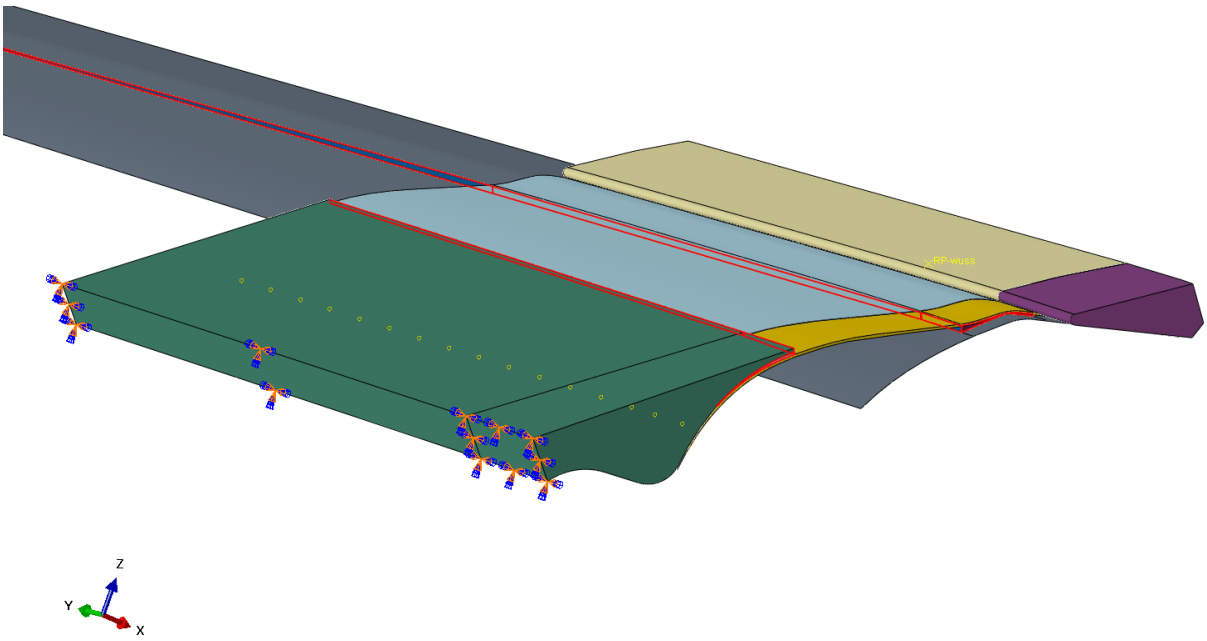

Figure 6-8. FEM with fictitious flow-impermeable surfaces highlighted in red.

The model used in the study is depicted in Figure 6-9, with monitored nodes on the two SGF sections marked with red dots.

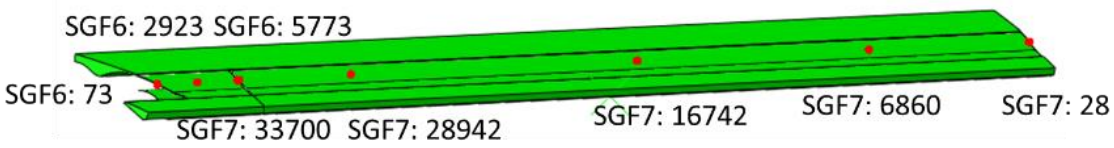

Figure 6-9. Nodes monitored in Abaqus model.

The first mode of the structure is illustrated in Figure 6-10. Note that the mode shape is essentially consistent as a function of static deformation aside from the arbitrary sign convention between mode shape comparisons. The modal frequency is a strong function of deformation quantity.

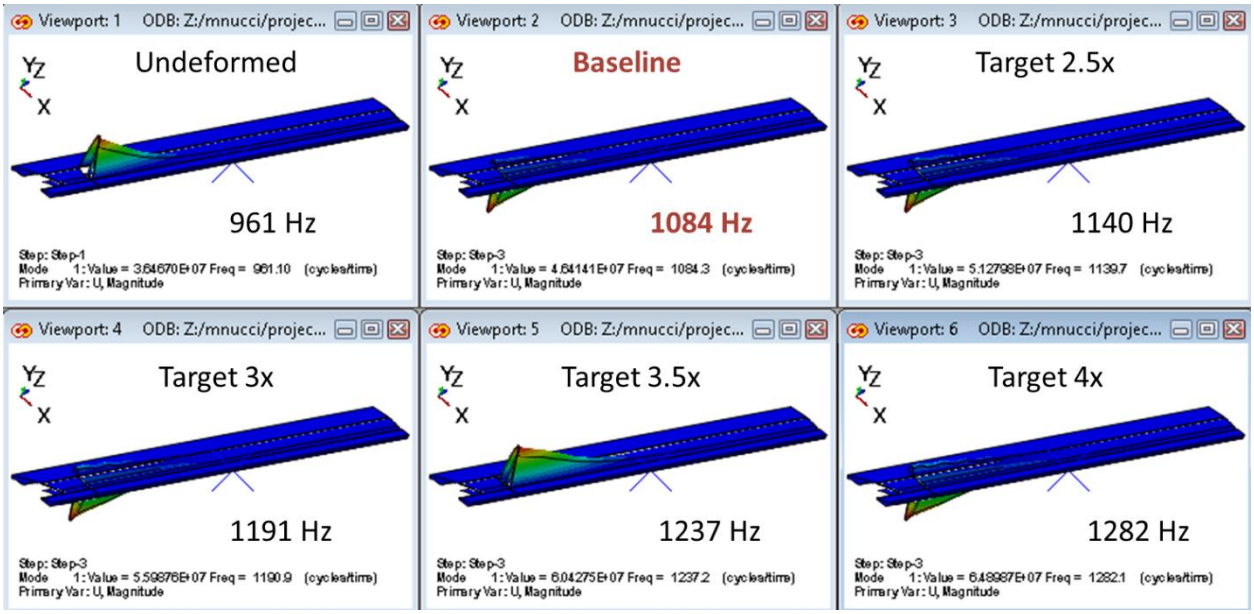

Figure 6-10. First mode of SGF.

6.1.1.4. Model Creation Lessons Learned

CFD mesh creation was the most time-consuming and critical part of the analysis process. For future analyses, vast amounts of time may be saved by considering the needs and requirements of cosimulation at the start of the model creation process. The biggest lessons learned during this process were that the CFD mesh should be designed to support large deformations and that it is critical for the CFD and FEM models to geometrically match.

6.1.1.4.1 Mesh Must Support Large Deformations

To simulate large deformations, the grid deformation algorithm must be capable of deforming the mesh without encountering negative volumes. For meshes that in the undeformed state are graded as unusable by `vogcheck`, this can be a challenge because poor-quality cells are more likely to encounter a negative volume during deformation. If complex geometry does not allow normal unstructured gridding methods to generate high-quality cells near the regions of large deformation, as with the CRM-HL semispan mesh, a hybrid meshing strategy may be used. The hybrid strategy involves creating a locally structured mesh near the regions of large deformation. This approach can be beneficial if it allows for higher-quality cells in the region. The structured block can interface with the surrounding unstructured mesh through either point-matched (preferred) or non-point-matched boundaries. If the boundaries are non-point-matched, the solution will have to be interpolated across them. However, point-matching the boundaries can cause poor-quality cells if the near-wall spacings between the structured and unstructured regions are dissimilar. With the CRM-HL model, a locally structured mesh was created around the gap fillers, as this was where the largest deformations would occur. As shown in Figure 6-11, the structured mesh was propagated away from deforming regions and interfaced with the unstructured mesh via non-point-matched boundaries.

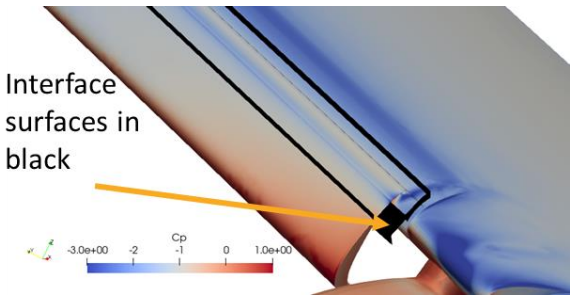

Figure 6-11. Structured-to-unstructured grid interface for CRM-HL mesh.

6.1.1.4.2 CFD and FEM Geometry Must Match

At the start of the cosimulation, a mapping is calculated between the cosimulating surfaces in the fluid and structural models. For the mapping to succeed, the fluid and structural surfaces need to be colocated in space. Furthermore, they should contain *identical* geometry (to the extent possible). If this is not the case (as it was not with the CRM-HL cosimulations), the mismatch can cause mapping failures and/or difficulties for grid deformation.

Figure 6-12 shows two examples of mismatched geometry between the CFD and FEM models which caused problems for grid deformation. On the left, the CFD geometry is such that there can never be a gap between the slat trailing edge and the slat-gap filler. In the FEM, a gap is possible because the connection is modeled with contact. The slat trailing edge is not a cosimulating surface in either the CFD or FEM models, but it is a deformable surface in the CFD model to allow it to grow and shrink with the gap opening/closing in the FEM. Depending on the level of deformation, the grid deformation algorithm may experience problems near the slat trailing edge. A suggested workaround is to put a dummy surface (shown in orange) in the FEM to explicitly model the slat trailing edge, as was done in the FEM for this cosimulation case and discussed in section 6.1.1.3. This way, the slat trailing edge can be included in the cosimulation, which will help the grid deformation algorithm but avoids causing any change to the structural response since the added elements lack mass and stiffness.

On the right in Figure 6-12 is an example of another geometrical mismatch between the models. There is a WUSS underpart in the CFD mesh, but it does not exist in the FEM. This situation can be problematic because the FEM may be allowed to deform into the space occupied by the WUSS underpart in the CFD mesh, which presents a difficult challenge for the grid deformation algorithm and usually leads to a negative volume. A suggested workaround would be to include the WUSS underpart in the FEM so that the geometries of the two models are identical. This was accomplished in the FEM of this cosimulation model by defining fictitious surfaces corresponding to the boundaries of the “WUSS under” region, as discussed in section 6.1.1.3.

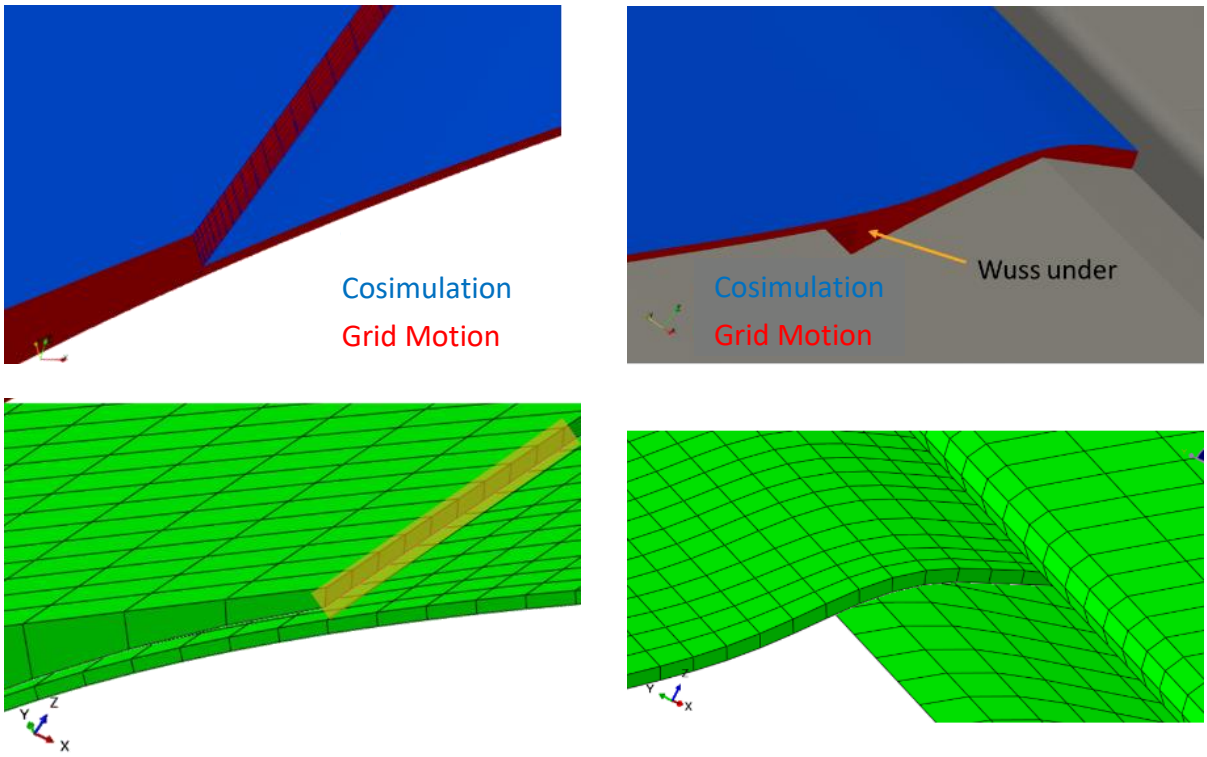

Figure 6-12. FEM (lower left) can have gaps between slat and SGF, but CFD cannot (upper left); CFD contains WUSS under geometry (upper right), but FEM does not (lower right).

6.1.2. Results

Results of quasistatic and dynamic aeroelastic cosimulations are described in these two sections.

6.1.2.1. Quasistatic Aeroelastic Cosimulation

An aeroelastic cosimulation was conducted for the baseline and high-q cases, with the goal of determining the quasistatic deformed shape of the SGF. The Loci/CHEM models were restarted from the previously generated subdomain results. As previously mentioned, the semispan solution file was again used to provide boundary conditions on the surfaces of the subdomain box. Since the generation of the semispan solution was a computationally intensive process, the baseline solution file was scaled to generate the solution file for the high-q case, using the newly developed scalePut utility [28]. This utility was used to scale the pressure and density by a factor of 1.9 (the ratio of the dynamic pressures of the two cases). Sutherland's law was used to calculate viscosity, and the leading coefficient in Sutherland's law (when in the form $\mu = \frac{a_1 T^{a_2}}{T+a_3}$) was scaled by the same factor. This scaling had the effect of increasing the dynamic pressure without altering the key nondimensional parameters of the flow (Mach and Reynolds numbers) so that the compressibility characteristics of the flow remained unchanged.

A quasistatic cosimulation was defined using the Loci/CHEM and Abaqus models. On the flexible SGF surfaces, Loci/CHEM imports displacement from Abaqus and exports nodal force to Abaqus. Since the

simulation is quasistatic, the time integration mode in Loci/CHEM was set to steady state and the Abaqus `*DYNAMIC, APPLICATION=QUASI-STATIC` solution procedure was used. The duration of the quasistatic cosimulation was defined to be 0.1 second. The cosimulation time step, which determines how often Loci/CHEM and Abaqus exchange data, was allowed to vary over the course of the simulation, as shown in Table 6-4. At the beginning of the simulation, the cosimulation time step was a relatively small 1e-4 seconds. As the simulation converges toward a quasistatic state, the solutions changed less, and a larger cosimulation time step was used.

Table 6-4. Quasistatic cosimulation time stepping.

Simulation Time, t (s)	$0 \leq t < 0.001$	$0.001 \leq t < 0.01$	$0.01 \leq t < 0.1$
Cosimulation Time Step (s)	0.0001	0.001	0.01

To avoid large impulse forces at startup, the force passed from Loci/CHEM to Abaqus was linearly ramped from passing 0% at t = 0, to passing 100% at t = 0.01. Figure 6-13 and Figure 6-14 show the displacements of the tracked nodes (see Figure 6-9) on SGF sections 6 and 7. The nodal displacements are nondimensionalized by their respective panel thickness. The baseline case results in SGF section 7 displacements of slightly under two panel thicknesses, whereas the displacements for the high-q case reach three panel thicknesses.

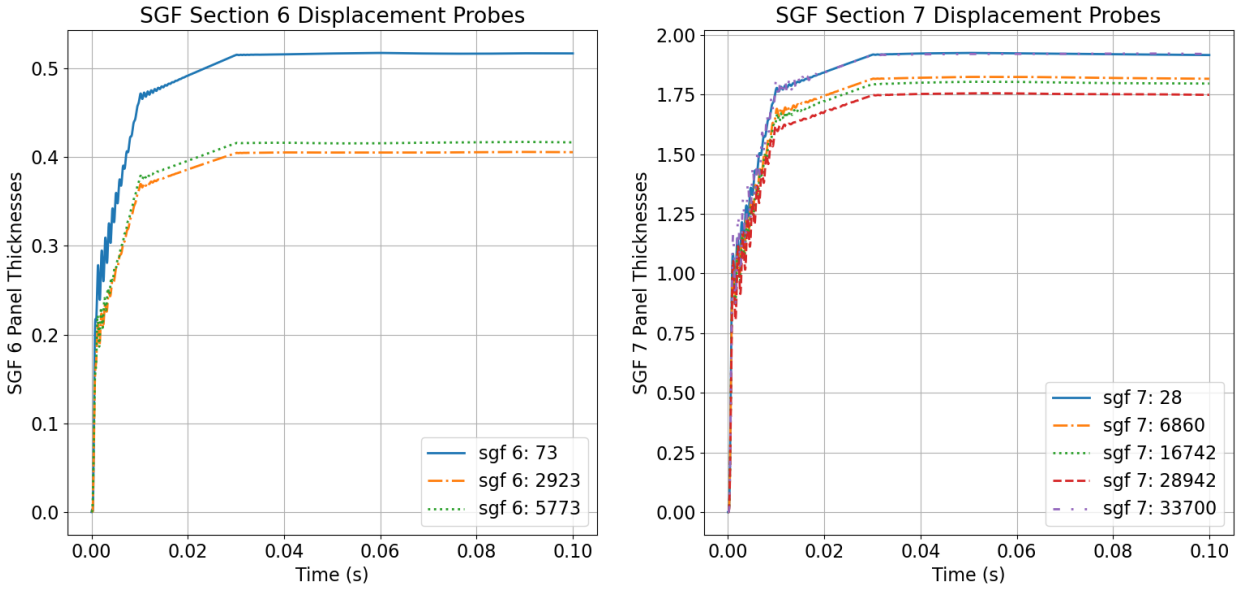

Figure 6-13. Nondimensional displacement magnitude for baseline quasistatic cosimulation.

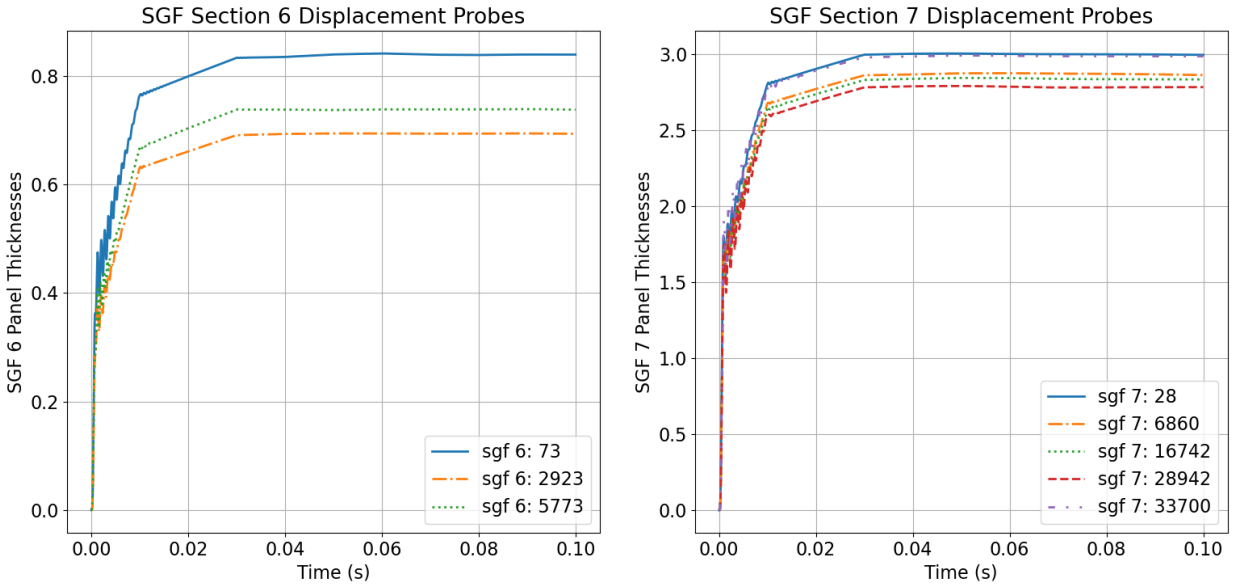

Figure 6-14. Nondimensional displacement magnitude for high-q quasistatic cosimulation.

The quasistatic deformation stiffens the SGF and results in a shift of the modes to higher frequencies. Figure 6-15 shows the first mode for the unloaded, baseline, and high-q cases. Figure 6-16 shows the frequency of the first ten modes versus dynamic pressure for several quasistatic loading cases. These results confirm that the SGF modes shift to higher frequencies with higher loading/deformation and demonstrate that the frequency spacing between sequential modes may change, which indicates that the dynamic response of the SGF will be sensitive to the deformation level.

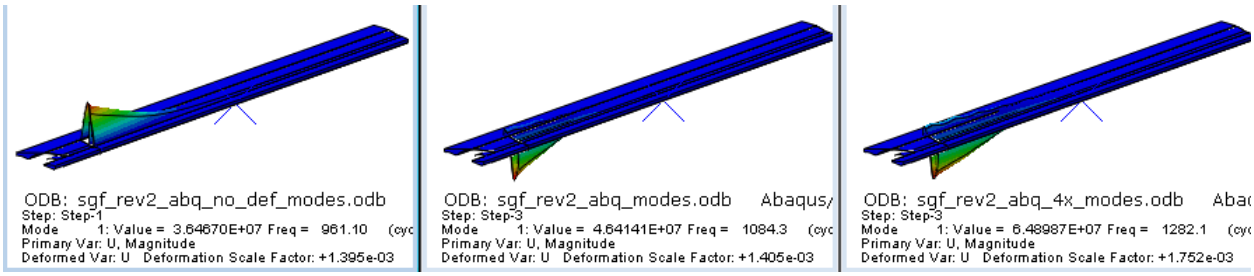

Figure 6-15. Mode 1 across unloaded (left), baseline (center), and high-q (right) cases.

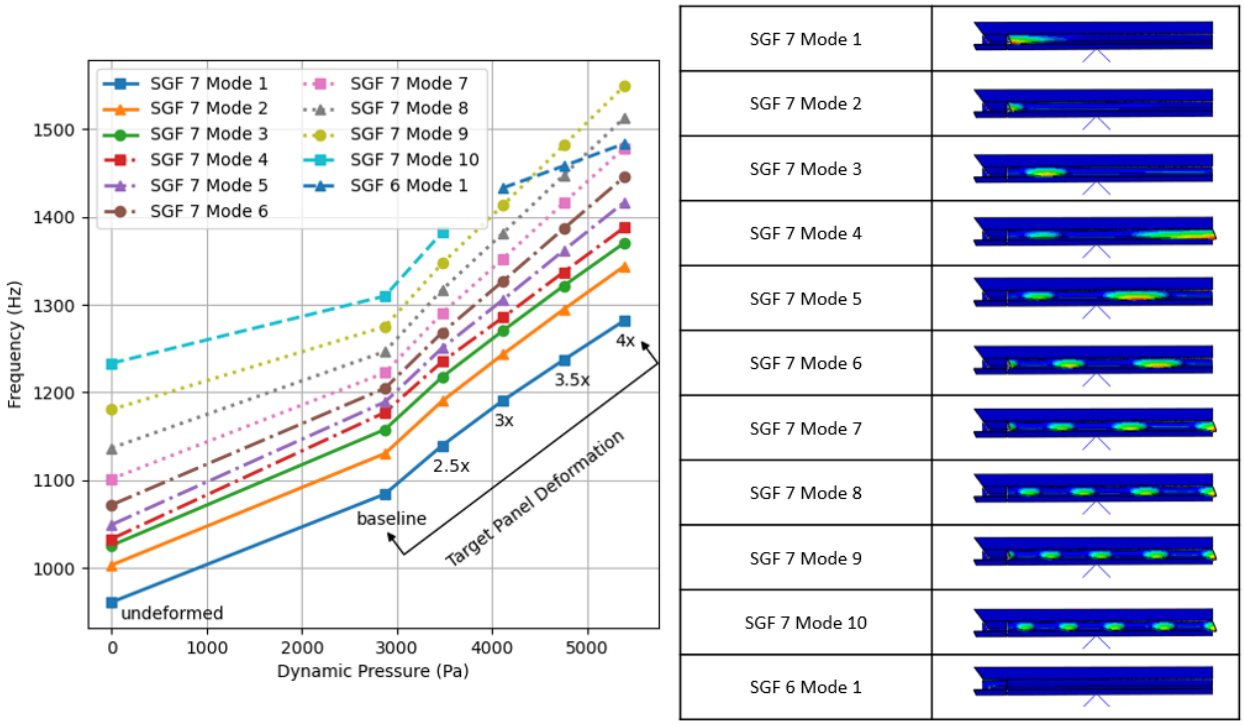

Figure 6-16. Mode numbers grouped by shape (selected modes ranked by frequency on previous figures).

6.1.2.2. Dynamic Aeroelastic Cosimulation

Dynamic aeroelastic cosimulations were restarted from the quasistatic results with a purpose of assessing the dynamic stability of the SGF. Both Loci/CHEM and Abaqus used the same 1e-5 second time step, and the cosimulation was set up to have second-order temporal accuracy with the codes communicating multiple times per time step, i.e., using the higher-order data exchange scheme illustrated in Figure 3-4(b). The Abaqus model switched to the `*DYNAMIC, APPLICATION=MODERATE DISSIPATION` solution procedure to reduce numerical dissipation, and the SGF FEM was perturbed with a spatially distributed force that was defined as a superposition of its first four mode shapes. The force was applied as a step function at the beginning of the cosimulation, active over the initial 2.5e-4 seconds. The magnitude of the force was small, 2.5e-3 psi for the baseline case and 2.5e-4 psi for the high-q case, and the purpose of the force was to excite the structure and cause it to oscillate. Stability was assessed by examining the growth or decay of the ensuing structural response. If the oscillations damped out over the course of the simulation, the SGF could be deemed stable; if they grew, it was unstable. The displacements were monitored at the same eight nodes as with the quasistatic cases (see Figure 6-9). Figure 6-17 and Figure 6-18 show the normalized displacements on SGF sections 6 and 7 for the baseline case. Visual inspection indicates that the oscillations are growing smaller, but there is some evidence of a beating phenomenon, especially in the SGF section 7 results. Given the restricted length of the time history available due to computational resources, it is difficult to reach a conclusion on stability via visual inspection alone, so Partial Floquet analysis was used to further assess stability. The baseline simulation obtained 0.0472 seconds of data.

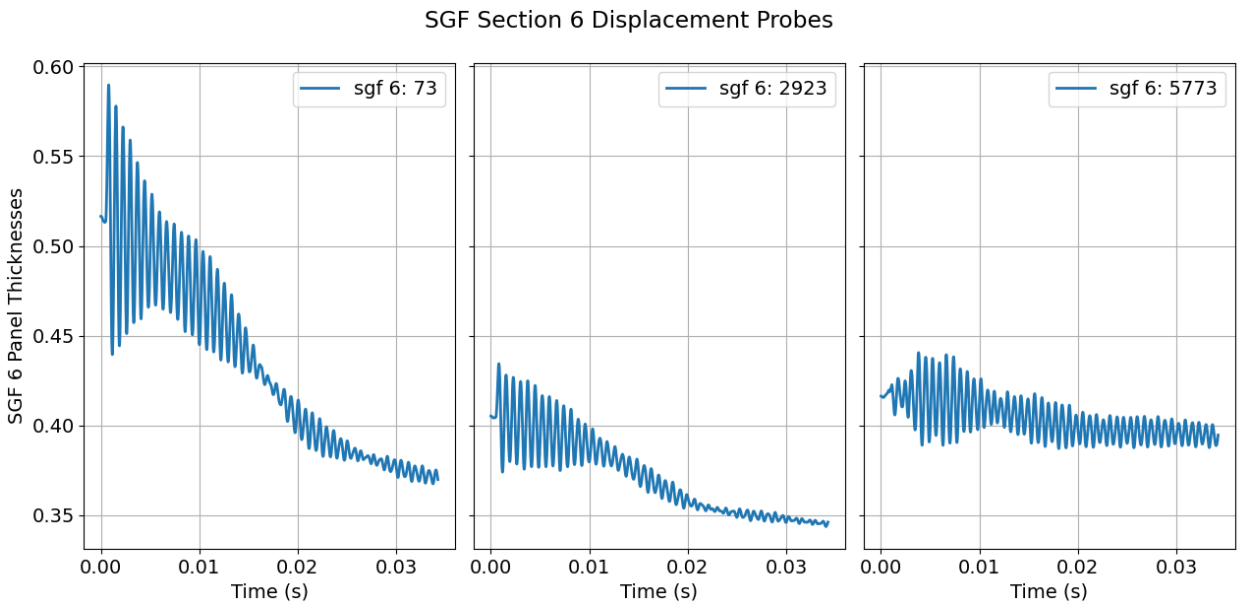

Figure 6-17. Baseline dynamic pressure displacement-time histories on SGF section 6.

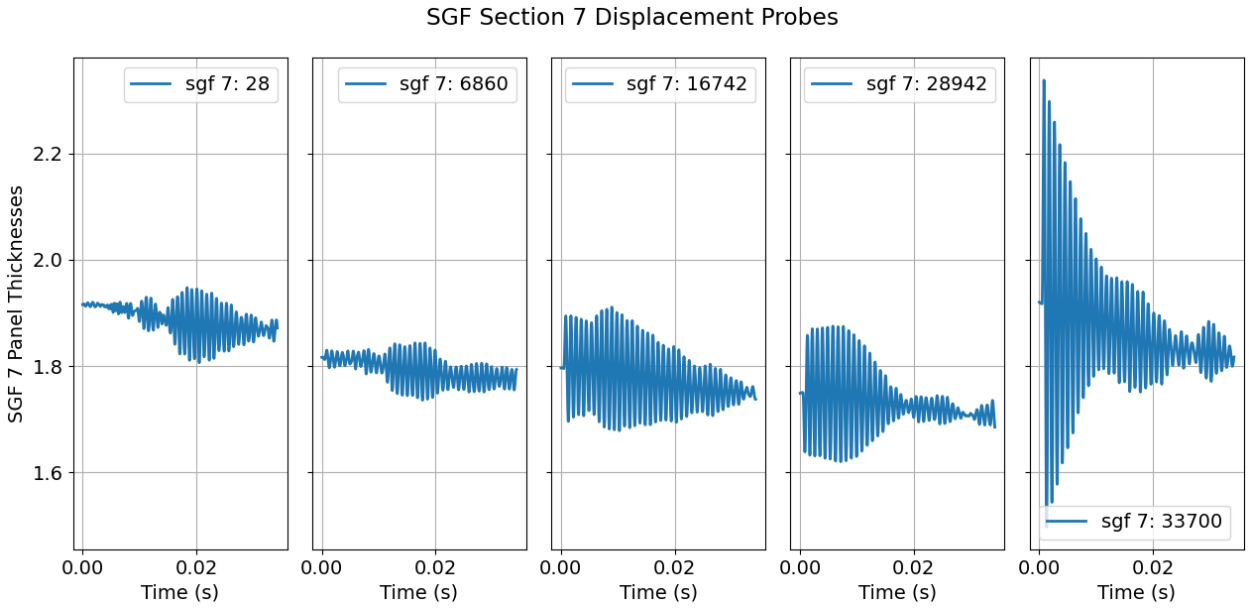

Figure 6-18. Baseline dynamic pressure displacement time histories on SGF section 7.

Table 6-5 provides the damping ratios for the detected frequencies obtained from a Partial Floquet analysis from section 7 nodes (28, 6860, 16,742, 28,942, 33,700) after the first 1000 time steps were discarded to avoid the initial transient around the perturbation. All frequencies detected were observed to be stable, consistent with visual inspection.

Table 6-5. SGF section 7 baseline dynamic pressure case Partial Floquet analysis suggests stable response.

Frequency (Hz)	Damping Ratio
1070.6	0.07%
1,101.4	1.00%
1,165.5	0.007%

Note that the first two frequencies appear to correspond to the first mode (approximately 1,084 Hz at baseline deformation). This is attributed to the dependence of modal frequency on deformation (see Figure 6-10). Consistent with the previous contrived example discussed in section 3.2, this results in duplicate modal detections. Also similar to the example, if the length of the simulation is restricted, the duplicate mode is observed to be negative. For example, if only the first 0.0422 seconds of the simulation is considered, the first two modes are observed as occurring at 1069.5 Hz and 1101.6 Hz with critical damping ratios of −0.26% and 0.98%, respectively. As discussed in the contrived numerical example, an identified instability can be detected even when no instability is possible. In the baseline dynamic pressure results, both damping ratios from the duplicate first mode are detected as stable after approximately 0.045 seconds.

The high-q case was run in the same manner as the baseline case. Figure 6-19 and Figure 6-20 show the displacement time histories for the dynamic cosimulation. Visually, the oscillations appear to be damping out, and there is less evidence of the beating phenomenon than in the baseline case. A Partial Floquet analysis was used to confirm the assessment that the high-q case is stable.

The increased dynamic pressure simulation obtained 0.0332 seconds of data. As expected, the steady component of the deformation is higher due to the increased dynamic pressure. However, the amplitude of dynamic deformation is noticeably reduced as compared to the baseline dynamic pressure.

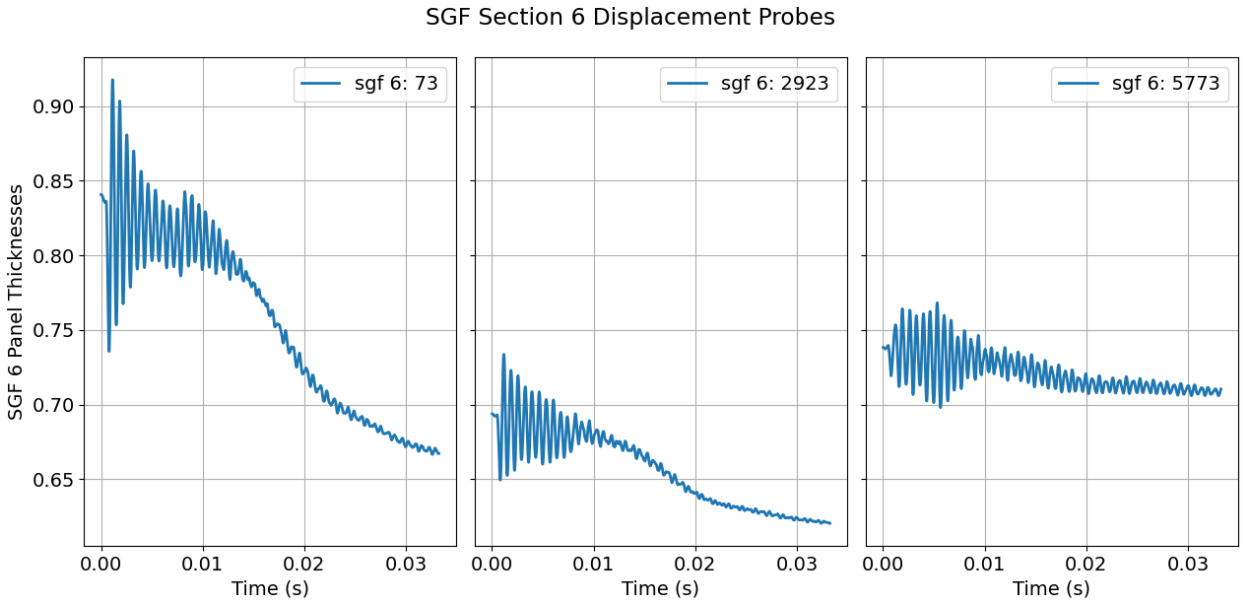

Figure 6-19. Elevated (high-q) dynamic pressure displacement time histories on SGF section 6.

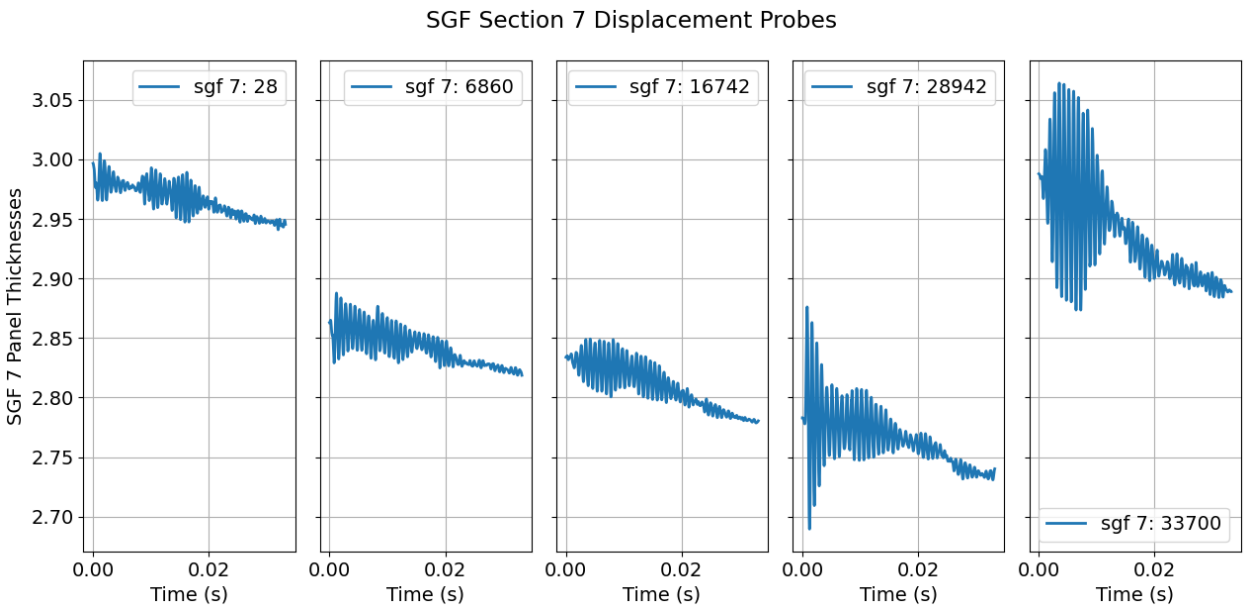

Figure 6-20. Elevated (high-q) dynamic pressure displacement time histories on SGF section 7.

Table 6-6 provides the damping ratios for the detected frequencies obtained from a Partial Floquet analysis from node 6860 after the first 1000 time steps were discarded to avoid the initial transient around the perturbation. All frequencies detected were observed to be stable, consistent with visual inspection. The system identification technique appears to find three distinct modes in this case, and duplicate modes are not observed at the higher dynamic pressure. The distinct modal content could have a variety of causes, but it is worth noting that the system identification technique preferentially finds modes with the highest energy content first. Thus, it is possible that this case had a more even distribution of energy across the structure's modes. Another possible explanation is that the reduction in dynamic deformation as compared to the baseline case results in the nonlinear structure behaving more linearly due to increased stiffness. As discussed in section 3.2, a nonlinear response in the frequency was observed to result in duplicate modes.

Table 6-6. SGF section 7 four times dynamic pressure case Partial Floquet analysis suggests stable response.

Frequency (Hz)	Damping Ratio
1171.8	1.75%
1268.8	1.25%
1505.0	0.85%

6.1.3. Summary and Lessons Learned

There were two main lessons learned during the cosimulations. The first deals with running on the NASA LaRC K-cluster and is detailed in the Appendix, and the second deals with the choice of turbulence model and is discussed here. The cosimulations were initially run with the Spalart-Allmaras turbulence model due to some historical comparisons with FUN3D simulations using the same model. This model worked

fine for the rigid and quasistatic simulations but would consistently crash during the dynamic cosimulations. When the turbulence model was switched to SST, the dynamic cosimulations ran without issue. For completeness, the rigid and quasistatic simulations were rerun with the SST model, and the results did not significantly change, as shown in Figure 6-4. For this combination of grid and flow conditions, the SST model proved more robust than Spalart-Allmaras while providing equal accuracy.

6.2. Final Status for Continuation of SCF to 3D FSI Modeling

Modeling the 3D SCF was deprioritized during the project based on unexpected results after wind tunnel testing which implied that more work may be needed to redesign the deployed profile. NASA has expressed a need to continue SCF FSI analysis once the redesign is complete. That leaves this as a topic of future work utilizing the best practices developed in this project. This section discusses the status of the SCF FEM for continuing aeroelastic assessment.

6.2.1. Description of the 3D FEM for the CRM-HL SCF Treatment in the Subdomain for FSI Assessment

Similar to the SGF FEM, the SCF structural FEM sought to exhibit essential behavior for the FSI calculation while minimizing computational burden. Therefore, all structure outside the subdomain was neglected, and within the subdomain only features directly affecting the behavior of the flexible SCF treatment were included in the FEM. The SCF treatment was designed and applied to the outboard slat of the CRM-HL only, with the inboard slat being treated with the SGF. Additionally, the SCF treatment was necessarily broken into sections between the slat brackets, as shown in Figure 6-21, which depicts the model as viewed from outboard and below the wing.

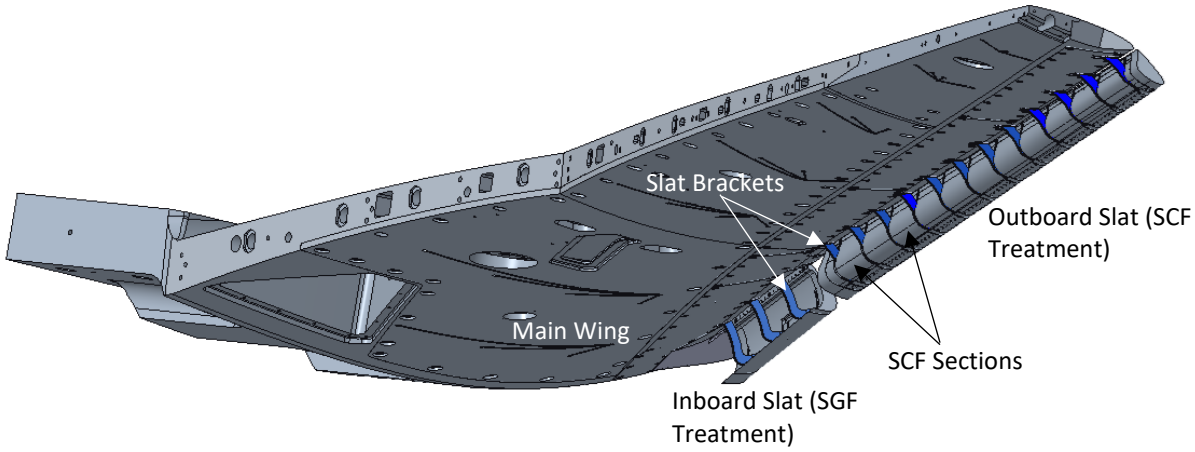

Figure 6-21. CAD image of CRM-HL wing (without trailing-edge section) showing detail of inboard and outboard slats.

The SCF sections were attached to the slat only, not the main wing, with one attachment near the trailing edge of the slat and one near the lower trailing edge (or cusp). The inboard end was treated with a solid structure that represented the most-inboard SCF section and end closure. The extent of the inboard end

treatment and its construction allowed it to be interchangeable with other end treatments. The dynamically scaled SCF treatments started with section 1 just inboard of the first bracket on the outboard slat, as shown in Figure 6-22.

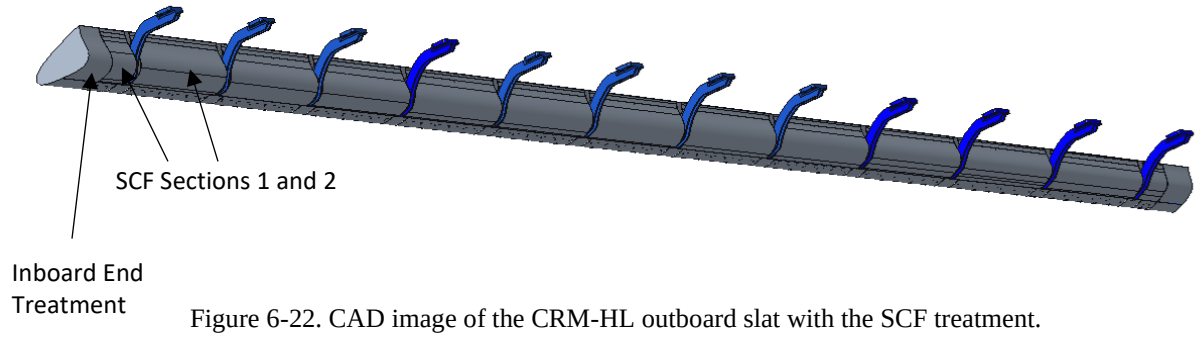

Figure 6-22. CAD image of the CRM-HL outboard slat with the SCF treatment.

The subdomain used for the SGF model was also used for the SCF because of relative flow complexity and unsteadiness in that region and because the SCF is most heavily loaded at the inboard end of the outboard slat. It is again noted that the grid deformation must be zero at the subdomain boundaries. In the case of the SCF treatment, that requirement necessitated limiting the deforming structure to the first two SCF sections at the inboard end of the outboard slat. The outboard slat was partitioned at the outboard edge of a slat feature just outboard of the second bracket, as shown in Figure 6-23. Everything outboard of the indicated partition was neglected. The loss in mass and stiffness influence of the slat section outboard of the subdomain was recognized, and the resulting model was considered conservative from a static deflection and stability standpoint.

Figure 6-23. Detailed view at inboard end of outboard slat, showing partitioning limit for FEM development.

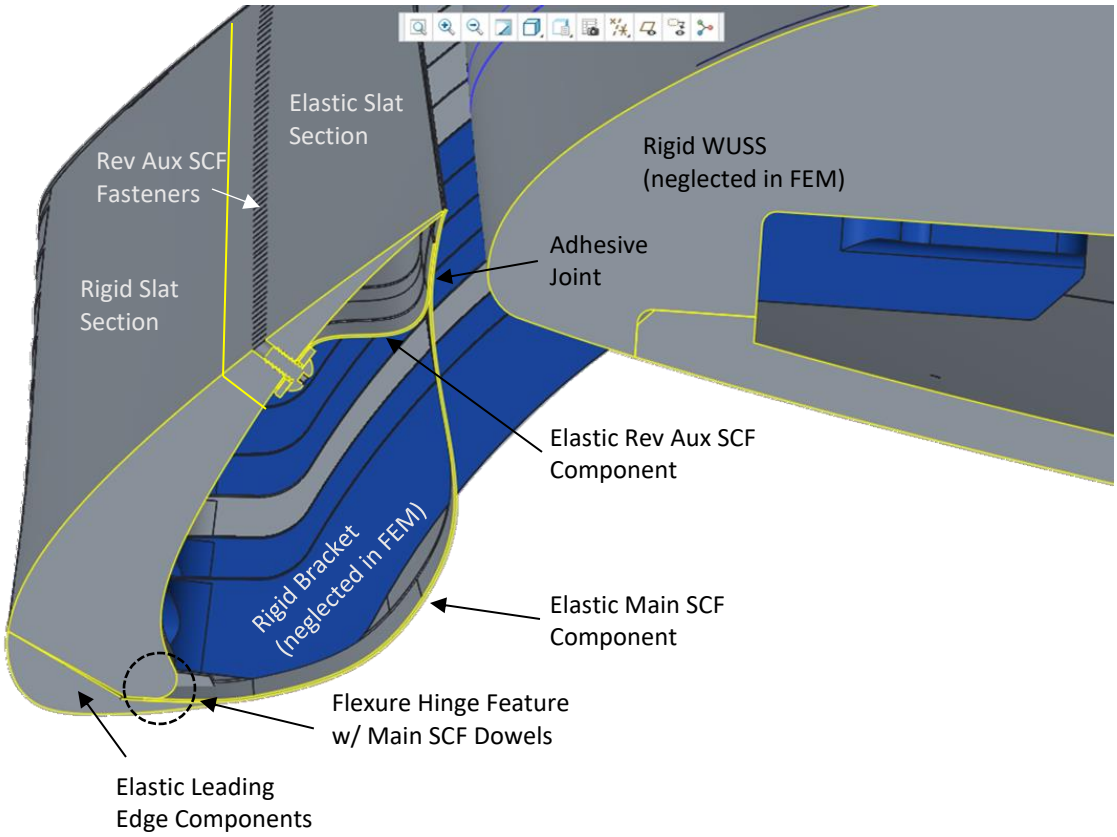

Figure 6-24. Cross-section of slat showing SCF components and integration with model.

Figure 6-24 shows a cross-section of the outboard slat showing the SCF components, their integration into the model, and other simplifying aspects of the FEM. The physical model required fastened joints between the SCF components and the slat for fail-safety. The SCF treatment consists of component pairs: the main and the reverse auxiliary (rev aux) SCF components. The rev aux SCF components were fastened to the trailing section of the slat well forward of the trailing edge. The main SCF components were secured to the slat by leading-edge components with dowel pins and fasteners. The pairs of SCF components were joined in their overlapping region with adhesive. The outboard slat section within the subdomain was further simplified by limiting the deforming model to the trailing section, ending just forward of the rev aux SCF joint, as indicated by the chordwise partition shown in Figure 6-24. The slat becomes very stiff forward of that location, and it was believed that the elasticity of that section of the slat had the greatest influence on the deflection and stability of the SCF treatments. All degrees of freedom were constrained at the root end of that elastic aft section of the slat to represent its connection to the rest of the slat. The leading-edge components, securing the forward end of the main SCFs to the slat, were also treated as elastic in the FEM because they have an influence similar to that of the slat trailing section on the deflection and stability of the SCF treatments. The remainder of the slat and inboard end treatment was treated as a rigid body in the FEM.

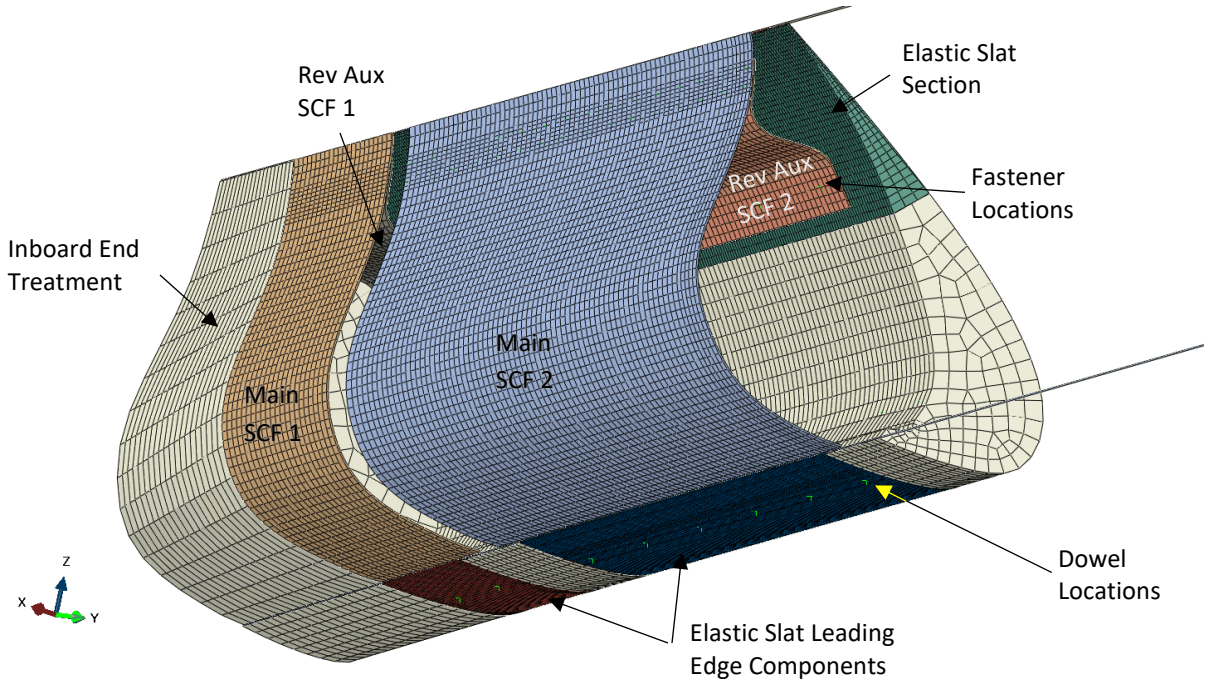

Figure 6-25. FEM of simplified SCF assembly within the computational subdomain.

Figure 6-25 shows the FEM resulting from the simplifications described above. The CAD was defeatured of fastener holes to simplify the model and make the mesh more efficient and regular. The slat trailing section had a chord of approximately 31.47 mm (1.24″) along the outer surface at the inboard edge and juncture with the inboard end treatment, and it was meshed with linear hexahedral elements (Abaqus C3D8R) having a spanwise resolution of 1.27 mm (0.05″), a one-way bias of 1.27 mm (0.05″) to 0.51 mm (0.02″) from root to tip in the chordwise direction, and three elements through the thickness. This meshing strategy was used to provide good resolution and well-behaved contact behavior between the main SCF and the slat trailing edge. The slat trailing section was assigned the material properties of stainless steel listed in Table 6-7. The remainder of the slat was modeled as a rigid body and meshed with surface elements (Abaqus R3D4) having a spatial resolution compatible with the elastic elements in the relevant adjoining regions and coarser in regions away from elastic bodies so as to reduce computational expense. The leading-edge components were assigned the material properties of stainless steel and meshed with linear hexahedral elements (Abaqus C3D8R). The spanwise resolution was also 1.27 mm (0.05″), with the same chordwise resolution in the region where the component adjoins the slat and a 1.27 mm (0.51″) resolution where the component adjoins the main SCF.

The main and rev aux SCF components were meshed with continuum shell elements (Abaqus SC8R) having chordwise/spanwise dimensions of approximately 1.27/2.54 mm (0.05″/0.1″), and they were assigned the material properties of the SMA (NiTi) listed in Table 6-7. The built-in model in Abaqus for superelastic materials, which is based on the model developed by Auricchio and Sacco [41], was used to capture the large-deformation behavior of the SMA.

Table 6-7. Properties of materials used in CRM-HL and FEM thereof. See Auricchio and Sacco [41] for definition of parameters.

Stainless Steel	
Property	Value
E	206.9 GPa
ν	0.26
ρ	7915 kg/m^3
Nitinol	
E_A, E_M	48.9 GPa, 40 GPa
$\nu_A=\nu_M$	0.33
H, Transf. ε	4.40%
σ^{Ms}, σ^{Mf}	471 MPa, 514 MPa
σ^{As}, σ^{Af}	196 MPa, 166 MPa
Tens./Comp.	Symmetric
Ref. Temp.	24°C
$C^A=C^M$	10 MPa/K
ρ	6480 kg/m^3

Point-based fasteners were introduced at the relevant locations for attaching the rev aux SCF components to the elastic trailing section of the slat, as shown in Figure 6-25, and to attach the elastic leading-edge components to the rigid part of the slat (not shown). A contact-tie constraint was defined between the bottom of the main SCFs and the adjoining surface regions of the leading-edge components to model the clamping pressure achieved by the fasteners and constraint provided by the dowels in the physical model. Similarly, contact tie constraints were defined to join each pair of main and rev aux SCF components over the adjoining (adhesive) area. Surface-to-surface contact interactions were defined between each main SCF component and the respective leading-edge component at the lower trailing edge (cusp) and between each main SCF and the slat trailing section at the upper trailing edge. Surface-to-surface contact was also defined in the adjoining regions with fasteners between the rev aux SCF components and the slat trailing section and between the leading edge components and the rigid part of the slat.

6.2.2. Key Features for CFD Cosimulations

The fluid region within the slat cove, i.e., between the main SCF components and the slat, was excluded from the CFD grid used in rigid studies and will be excluded from aeroelastic cosimulations. The CFD gridding approach will accommodate the gaps between main SCF sections by bridging the gaps with deforming grids.

Features had to be introduced into the FEM to handle regions of the fluid domain where gaps could open or close due to structural deformation during the solution. Regions of transient gaps in the SCF model existed at the trailing edges of the slat where the main SCF components interact with the elastic slat section and leading-edge components. These regions, highlighted in Figure 6-26, were closed with fictitious surfaces that were meshed with surface elements (Abaqus SFM3D4), and the meshed surfaces

were tied along their edges to the adjacent structures to provide a deforming but ever-present boundary for the fluid domain. Steady-pressure loads were defined on the inside of the main SCF components and distributed according to the steady pressures at the gaps between main SCF components from the steady semispan CFD solution. Steady pressure distributions were similarly defined between the main SCF and the leading-edge components in the overlap region aft of the clamped joint. Finally, pressure loads were defined for the outer surface of the main SCF components for transfer of transient pressure loads from the CFD solver.

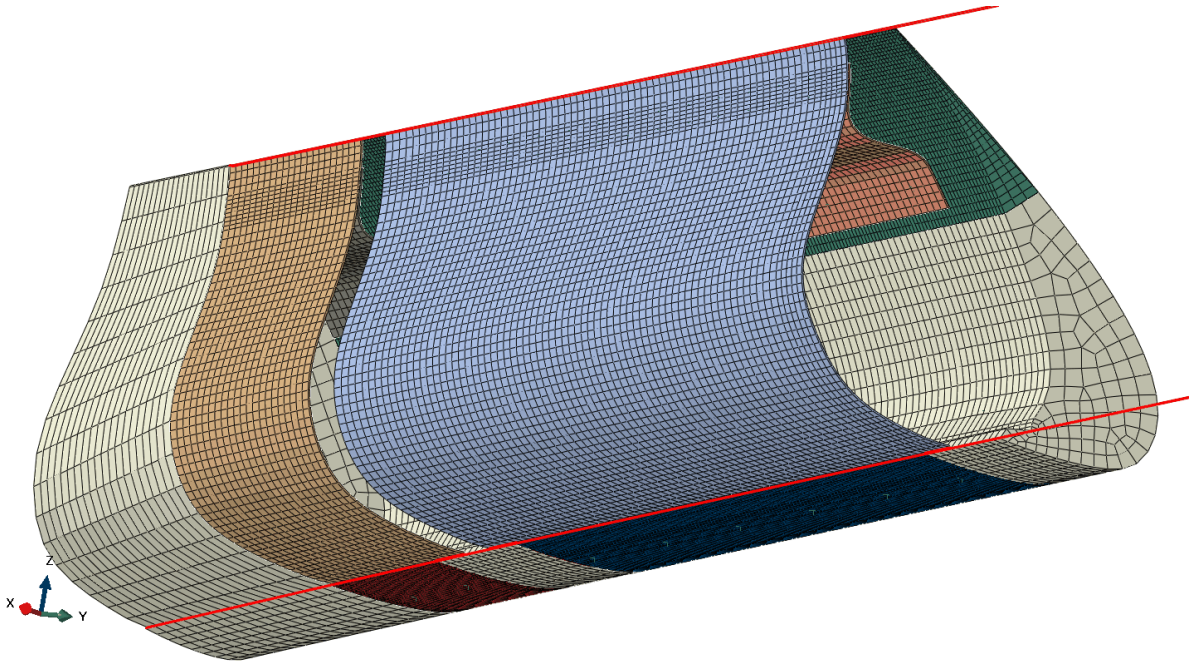

Figure 6-26. FEM with fictitious flow-impermeable surfaces highlighted.

7 SUMMARY AND CONCLUSIONS

Prior work on adaptive SMA SGF and SCF structures demonstrated promising slat noise reduction for aircraft in high-lift configurations. However, successful TRL maturation requires not only acoustic and quasistatic structural assessment, but also aeroelastic stability analysis. This report has described how ATA Engineering and NASA collaborated to apply Loci/CHEM and Abaqus to solve this problem in both a quasistatic and dynamic manner.

FSI analyses included 2D SGF and SCF configurations, along with a 3D SGF geometry. The technical approach followed these steps:

- CFD simulation of the global flow field under the assumption of rigid geometry.
    - For the 2D simulations, this consisted of the flow field around a multielement wing section.
    - For the 3D simulations, this consisted of a semispan model of the CRM-HL.
- Extraction or generation of a subdomain mesh in some neighborhood of the technology element. In general, the neighborhood bounds were defined to achieve a balance between accuracy and computational efficiency.
- Quasistatic cosimulation between CFD (Loci/CHEM) and NL-CSD (Abaqus) to obtain the mean deformed structural shape. A key output was the deflection of the structure, typically measured against panel thickness.
- Dynamic cosimulation between CFD and NL-CSD to assess aeroelastic stability. A key output of this analysis was the quantified damping from Partial Floquet analysis.

Generation of suitable subdomain meshes that could also successfully be morphed on the compliant technology element boundaries required care, and ATA made a specific effort to develop best meshing practices for the 3D SGF case due to the compliant technology element's thin structure and proximity to other rigid surfaces in the simulation domain.

Completing the technical approach also required the development of a number of tools and methods that were not available at the start of the project:

- A cell-quality remediation module in Loci/CHEM to drop the solver to first-order spatial accuracy in regions of poor mesh quality.
- A scaling utility to simulate elevated dynamic pressures while holding Mach and Reynolds numbers fixed.
- A force ramping capability for quasi-steady FSI to avoid large deformation overshoots that resulted in negative-volume mesh cells upon morphing.

The long-term value of this capability is a means to substantiate margins of safety for future lightweight structures used for airframe noise reduction. The methods described herein using Loci/CHEM and

Abaqus have also been implemented for FUN3D to Abaqus. The next three sections describe the key findings, lessons learned, and recommendations for future work.

7.1. Key Outcomes

The following key outcomes were achieved through this effort:

- ATA demonstrated that the Loci/CHEM `vogcut` utility along with the `scalablePrescribed` boundary and initial condition mapping utility could be used on 3D rigid subdomains of the CRM-HL model to produce flow fields substantially equivalent to those obtained with a full semispan solution. This result was demonstrated using comparisons of pressure coefficient and wall shear stress.
- The simulated 3D SGF with two sections (SGF6 and SGF7) had quasistatic aeroelastic deformation measured in terms of the sections' respective thicknesses of about 0.5 and 2.0 under the baseline loading condition.
- The same 3D SGF sections had quasistatic deformations of about 0.8 and 3.0 panel thicknesses under the elevated (high-q) dynamic pressure condition.
- Dynamically, the 3D SGF was found to possess stabilizing aerodynamic damping at baseline dynamic pressures as well as in flight conditions corresponding to the elevated (high-q) dynamic pressure case.
- New methods to improve stability (`cellQualityRemediation`, force ramping) and efficiency (`scalePut`) were developed as an outcome of this project.

7.2. Lessons Learned

Lessons learned are summarized in this section for several major parts of the project.

7.2.1. 2D Aeroelastic Analysis

- When using kinematic coupling to tie structural FEM nodes to wetted surface nodes for CFD, checks should be performed to ensure that the wetted surface does not affect FEM response.
- Exchange of nodal force guarantees force conservation between fluid and structural models.
- Dynamic FSI time-step selection should be based on the resolution of the structural modes of interest.
- For relatively inexpensive 2D models, parameter studies are useful as a means to bound stability estimates.
- Application of impulsive loading in dynamic FSI may improve the signal-to-noise ratio for damping estimation.

7.2.2. 3D Subdomain Modeling

- While good mesh quality is desirable everywhere in the subdomain, it is particularly important in regions near subdomain boundaries and near the compliant technology element.

- It is recommended that this modeling use a subdomain similar to the largest (Grid3) from the studies presented herein, which was well upstream of the nacelle and downstream of the wing trailing edge, and whose only intersection with surface geometry was at spanwise stations along the wing far from the technology element.
- There is no specific advantage to using embedded surfaces within a full model volume mesh to define subdomains; if embedded surfaces are to be used, it may be advantageous to define the `vogcut` bounding box inside the prescribed internal boundary surfaces.

7.2.3. 3D Aeroelastic Analysis

- Meshing with FSI in mind is critical—high-quality cells should be near areas of mesh motion, and boundaries should be split to align with mesh motion boundary conditions.
- The SGF nonlinear structural response is known through Abaqus modal analysis to vary the modal frequencies as a function of deformation. The finite perturbation aeroelastic stability analysis at baseline dynamic pressure appears to have resulted in sufficient frequency variation to result in duplicated first mode detections. These duplicate modes were both detected as stable after approximately 0.045 second or 49 cycles of simulated data.
- The system identification technique did not produce duplicate modes when applied to the higher dynamic pressure simulation. This may be due to the reduced amplitude of dynamic response.
- There is a relevant bug in the Abaqus 2019 CSE, so that version should be avoided.

7.3. Recommendations for Future Work

The following future work recommendations are put forth based on the current effort:

- The baseline dynamic pressure simulation should be repeated with a structural model obtained from linearizing about the baseline deformation. This simulation will enable confirmation or rejection of the hypothesis that duplicate modes obtained with the system identification tool are caused by the nonlinearity in the structural model.
- A 3D static and dynamic FSI analysis should be carried out on the SCF geometry once its performance is better understood.
- NASA LaRC should utilize the LaRC CFD solver FUN3D to perform aeroelastic simulations with Abaqus for nonlinear structures.

REFERENCES

[1] Hayes, J.A., Horne, W.C., Soderman, P.T., and Bent, P.H. "Airframe Noise Characteristics of a 4.7% Scale DC-10 Model." AIAA Paper 1997-1594, 1997.

[2] Storms, B.L., Ross, J.C., Horne, W.C., Hayes, J.A., Dougherty, R.P., Underbrink, J.R., Scharpf, D.F., and Moriarty, P.J. "An Aeroacoustic Study of an Unswept Wing with a Three-Dimensional High Lift System." NASA TM 112222, February 1998.

[3] Dobrzynski, W., Nagakura, K., Gehlhar, B., and Buschbaum, A. "Airframe Noise Studies on Wings with Deployed High-Lift Devices." AIAA Paper 1998-2337, 1998.

[4] Storms, B.L., Hayes, J.A., Moriarty, P.J., and Ross, J.C. "Aeroacoustic Measurements of Slat Noise on a Three-Dimensional High-Lift System." AIAA Paper 1999-1957, 1999.

[5] Choudhari, M.M., Lockard, D.P., Macaraeg, M.G., Singer, B.A., Streett, C.L., Neubert, G.R., Stoker, R.W., Underbrink, J.R., Berkman, M.E., Khorrami, M.R., and Sadowski, S.S. "Aeroacoustic Experiments in the Langley Low-Turbulence Pressure Tunnel," NASA TM 211432, February 2002.

[6] Mendoza, F., and Brooks, T. "Aeroacoustic Measurements of a Wing/Slat Model." AIAA Paper 2002-2604, 2002.

[7] Terracol, M., Manoha, E., and Lemoin, B. "Noise Sources Generation in a Slat Cove: Hybrid Zonal RANS/LES Simulation and Dedicated Experiment." AIAA Paper 2011-3203, 2011.

[8] Chow, L.C., Mau, K., and Remy, H. "Landing Gears and High Lift Devices Airframe Noise Research." AIAA Paper 2002-2408, 2002.

[9] Streett, C.L., Casper, J., Lockard, D.P., Khorrami, M.R., Stoker, R., Elkoby, R., Wenneman, W., and Underbrink, J. "Aerodynamic Noise Reduction for High-Lift Devices on a Swept Wing Model." AIAA Paper 2006-0212, 2006.

[10] Zhang, Y., Cattafesta, L., Pascioni, K., Choudhari, M., Lockard, D., Khorrami, M., and Turner, T. "Slat Noise Control Using a Slat Gap Filler." AIAA-2020-2553, *AIAA Aviation 2020 Forum*, Virtual Event, June 15–19, 2020.

[11] Scholten, W.D., Hartl, D.J., Turner, T.L., and Kidd, R.T. "Development and Analysis-Driven Optimization of Superelastic Slat-Cove Fillers for Airframe Noise Reduction." *AIAA Journal* 54, No. 3, 2016: pp. 1078–94.

[12] Turner, T.L., and Long, D.L. "Development of a SMA-Based, Slat-Gap Filler for Airframe Noise Reduction." AIAA Paper 2015-0730, 2015.

[13] Lacy, D.S., and Sclafani, A.J. "Development of the High Lift Common Research Model (HL-CRM): A Representative High Lift Configuration for Transonic Transports." AIAA Paper 2016-0308, 2016.

[14] AIAA Drag Prediction Workshop website, https://aiaa-dpw.larc.nasa.gov/, retrieved 20 May 2021.

[15] NASA Common Research Model website, https://commonresearchmodel.larc.nasa.gov/, retrieved 17 May 2021.

[16] AIAA Geometry and Mesh Generation website, http://www.gmgworkshop.com/, retrieved 17 May 2021.

[17] AIAA High Lift Prediction Workshop, https://hiliftpw.larc.nasa.gov/, retrieved 20 May 2021.

[18] Lockard, D., Turner, T., Bahr, C., and Hutcheson, F, "Overview of Aeroacoustic Testing of the High-Lift Common Research Model," AIAA-2021-2113, *2021 AIAA Aviation Forum*, Virtual Event, 2-6 August 2021, DOI: 10.2514/6.2021-2113.

[19] Horne, W.C., Burnside, N.J., Soderman, P.T., Jaeger, S.M., Reinero, B.R., James, K.D., and Arledge, T.K. "Aeroacoustic Study of a 26%-Scale Semispan Model of a Boeing 777 Wing in the NASA Ames 40- by 80-Foot Wind Tunnel." NASA TP 2004-212802, October 2004.

[20] Scholten, W., Patterson, R., Eustice, M., Cook, S., Hartl, D., Strganac, T., Turner, T., "Aerodynamic and Structural Evaluation of an SMA Slat-Cove Filler Using Computational and Experimental Tools at Model Scale," *Proc. ASME 2018 Conf. on Smart Matl., Adaptive Struct. and Intell. Sys.*, SMASIS2018-8129, San Antonio, TX, 10-12 September 2018.

[21] Turner, T. L., "Aeroelastic Stability Assessment Methodology and Application to Slat Noise Treatments on the High-Lift Common Research Model" NASA TM, to be published.

[22] Lockard, D.P., et al. "High-Lift Common Research Model Aeroacoustics Test." NASA Acoustics Technical Working Group Meeting Presentation, April 2021.

[23] Reveles, N. "User's Guide for Development and Implementation of the FUN3D / Abaqus Coupling." Revision 1, 2021.

[24] Isogai, K. "On the Transonic-Dip Mechanism of Flutter of a Sweptback Wing." *AIAA Journal* 17, No. 7, 1979: pp. 793–5.

[25] Luke, E., and George, T. "Loci: A Rule-Based Framework for Parallel Multidisciplinary Simulation Synthesis." *Journal of Functional Programming*, Special Issue on Functional Approaches to High-Performance Parallel Programming, Vol. 15, No. 3, 2005: pp. 477–502.

[26] Luke, E. "On Robust and Accurate Arbitrary Polytope CFD Solvers (Invited)." AIAA 2007-3956, AIAA Computational Fluid Dynamics Conference, Miami, 2007. doi: 10.2514/6.2007-3956.

[27] ATA Engineering, Cell Quality Remediation Module, URL: https://github.com/ATAEngineering/cellQualityRemediation

[28] ATA Engineering, scalePut utility, URL: https://github.com/ATAEngineering/scalePut

[29] Blades, E.L., Miskovish, R.S., Nucci, M., Shah, P., Bremner, P.G., and Luke, E.A. "Towards a Coupled Multiphysics Analysis Capability for Hypersonic Vehicle Structures." AIAA-2011-1962, *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2011. doi: 10.2514/6.2011-1962.

[30] Blades, E., Miskovish, S., Luke, E., Collins, E., and Kurkchubashe, A. “Multiphysics Simulation Capability Using the SIMULIA Co-Simulation Engine.” AIAA-2011-3397, *20th AIAA Computational Fluid Dynamics Conference*, 2011. doi: 10.2514/6.2011-3397.

[31] Luke, E., Collins, E., and Blades, E. “A Fast Mesh Deformation Method Using Explicit Interpolation.” *Journal of Computational Physics* 231, No. 2, January 2012: pp. 586–601. doi: 10.1016/j.jcp.2011.09.021.

[32] Blades, E., Shah, P., Nucci, M., and Miskovish, S. “Demonstration of Multiphysics Analysis Tools on Representative Hypersonic Vehicle Structures.” AIAA 2013-1746, *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2013. doi: 10.2514/6.2013-1746.

[33] Blades, E., Luke, E., and Ruf, J. “Fully Coupled Fluid-Structure Interaction Simulations of Rocket Engine Side Loads.” AIAA-2012-3969, *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2012. doi: 10.2514/6.2012-3969.

[34] Blades, E.L., and Cornish, A. “Aeroelastic Stability Predictions of a Business Jet Landing Gear Door using High Fidelity Fluid-Structure Interaction Tools.” AIAA-2015-0173, *56th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015. doi: 10.2514/6.2015-0173.

[35] Bauchau, O.A., and Wang, J. “Efficient and Robust Approaches to the Stability Analysis of Large Multibody Systems,” *Journal of Computational and Nonlinear Dynamics* 3, No. 1, January 2008.

[36] Roe, P.L. “Approximate Riemann Solvers, Parameter Vectors and Difference Schemes.” *Journal of Computational Physics* 43, 1981: pp. 357–72.

[37] Venkatakrishnan, V. “On the Accuracy of Limiters and Convergence to Steady State Solutions.” AIAA Technical Report 93-0880, 1993.

[38] Sutherland, W. “The Viscosity of Gases and Molecular Force.” *Philosophical Magazine*, S. 5, 36, 1893: pp. 507–31.

[39] Spalart, P.R., and Allmaras, S.R. “A One-Equation Turbulence Model for Aerodynamic Flows.” AIAA-92-0439, January 1992.

[40] Blades, E.L., Shah, P.N., Nucci, M., and Miskovish, R.S. “Demonstration of Multiphysics Analysis Tools on Representative Hypersonic Vehicle Structures.” AIAA-2013-1746, 2013.

[41] Auricchio, F., and Sacco, E. “A One-Dimensional Model for Superelastic Shape-Memory Alloys with Different Elastic Properties Between Austenite and Martensite.” *International Journal of Non-Linear Mechanics* 32, No. 6, 1997: pp. 1101–14.

APPENDIX

The K-cluster at NASA LaRC was initially used to run the rigid CFD simulations (semispan and subdomain) as well as the quasistatic and dynamic cosimulations. The rigid CFD simulations involved Loci/CHEM only (no Abaqus or CSE) and all ran smoothly without issue. As the quasistatic cosimulations began, three main issues were encountered.

1. The Abaqus module for Loci/CHEM would either not compile or cause an immediate crash during simulation execution.
2. Cosimulations would crash during the startup portion of the solve, when the SMACseConnect CSE API call was made.
3. Cosimulations would unpredictably crash in a nonrepeatable manner.

The first two issues were found to be repeatable, but the last appeared to be random.

A-1 Repeatable Issues

The first issue was eventually solved through a process of attempting to compile/run simulations using various compilers and environment modules. The first attempted Loci/CHEM installation on K-cluster was compiled with the Intel compiler suite and with the environment modules shown in Table A-1. The modules in Table A-1 are listed in the order that they were loaded. This is important because modules loaded later can overwrite environment variables set by earlier modules.

Table A-1. Modules loaded for compilation of Loci/CHEM.

Module Name	Purpose
intel_2015.0.090	Intel compiler suite, loaded by default
mpt-2.16	MPI installation, default version
intel_2018.0.033	Intel compiler suite, dependency for petsc_3.11.3_mpt
petsc_3.11.3_mpt	PETSc module available on K, needed for Abaqus module

To run cosimulations with Loci/CHEM and Abaqus, the Abaqus module for Loci/CHEM must be compiled. This requires the loading of one of the Abaqus environment modules shown in Table A-2, extracting the CSE API from Abaqus, and linking the Loci/CHEM module to the CSE API. It is important to note that each Abaqus environment module on the K-cluster has a dependency on an Intel compiler suite module, and this module may or may not be the same as the default Intel compiler suite module (intel_2015.0.090) which is automatically loaded for the user on login.

Table A-2. Abaqus environment modules available on K-cluster r, and their dependencies.

Abaqus Module	Dependency
ABAQUS_2017	intel_2015.0.090
ABAQUS_2018	intel_2015.0.090
ABAQUS_2019	intel_2016.3.210
ABAQUS_2020	intel_2019.1.144

The SGF FEM was originally provided in Abaqus/2018, so that version was initially used in an attempt to compile the Abaqus module for Loci/CHEM. As shown in Table A-3, the compilation was unsuccessful due to a missing Intel symbol. It was hypothesized that the source of the problem could be the Intel compiler since the Intel symbol should be provided by the compiler). Therefore, ATA next recompiled Loci/CHEM with the GNU compiler suite and attempted to compile the Abaqus module with GNU. Since PETSc is a dependency of the Abaqus module and there was no available installation of PETSc using the GNU compilers on the K-cluster, PETSc was installed from scratch and a module (petsc/3.13.2) was created for it. As shown in Table A-3, compilation of the Abaqus module for Loci/CHEM was successful, but cosimulations immediately failed with missing library errors. It was then hypothesized that MPT (or another dependency) was compiled with Intel and was causing a conflict with GNU. Another attempt was made, switching the version of the Intel compiler suite back to the default 2015 because the original missing symbol error was for version 2018. As shown in Table A-3, compilation of the Abaqus module for Loci/CHEM was successful, but cosimulation crashed with an error indicating that a pure virtual function was called. In object-oriented programming, a pure virtual function is one that must be overridden by an inheriting class. It should not be possible to call a pure virtual function, and the compiler should flag this error. Another attempt was made, switching to Abaqus/2019. As shown in Table A-2, Abaqus/2019 has a different Intel dependency than Abaqus/2018. The Abaqus module for Loci/CHEM compiled successfully, but cosimulations crashed during the startup phase during an API call to SMACseConnect. This behavior was confirmed as a bug in the 2019 CSE by SIMULIA, who suggested an upgrade to Abaqus/2020, in which the bug is fixed.

The attempt with Abaqus/2019 was very informative because it indicated that the undefined symbol error would go away if the ABAQUS_2019 module was loaded. The reason for this behavior is that ABAQUS_2019 brings in intel_2016.3.210, making it the last Intel module loaded, and therefore taking precedence over the previously loaded intel_2018.0.33. The undefined symbol error was for Intel 2018, but this attempt showed that Intel 2016 did not have the same problem. A final attempt was made using Abaqus/2020 per SIMULIA's suggestion. Table A-2 shows that Abaqus/2020 has dependence on yet another Intel version (2019); however, Intel 2016 can be loaded last to avoid any undefined symbol errors. Table A-3 shows that this attempt was successful in that the Abaqus module compiled and cosimulations were able to run.

Table A-3. Summary of compilation/simulation attempts with different compilers and versions of Abaqus.

Attempt	Modules Loaded	Result
Abaqus/2018 with Intel compilers	intel_2015.0.090 mpt-2.16 intel_2018.0.33 petsc_3.11.3_mpt ABAQUS_2018	Loci/CHEM Abaqus module does not compile with undefined symbol error for __intel_skx_avx512_memcpy in Intel 2018 library
Abaqus/2018 with GNU compilers	intel_2015.0.090 mpt-2.16 petsc/3.13.2 (created by ATA) gcc_9.2.0 ABAQUS_2018	Loci/CHEM Abaqus module compiles without issue, but co-simulation fails with library problems: "GLIBCXX_3.4.26", "CXXABI_1.3.8", etc not found
Abaqus/2018 with only Intel 2015	intel_2015.0.090 mpt-2.16 petsc/3.13.2_intel (created by ATA) ABAQUS_2018	Loci/CHEM Abaqus module compiles without issue, bu co-simulation fails with "pure virtual function called" error
Abaqus/2019 with Intel compilers	intel_2015.0.090 mpt-2.16 intel_2018.0.33 petsc_3.11.3_mpt intel_2016.3.210 ABAQUS_2019	Loci/CHEM Abaqus module compiles, but co-simulation crashes during SMACseConnect API call
Abaqus/2020 with Intel compilers, load Intel 2016 last	intel_2015.0.090 mpt-2.16 intel_2018.0.33 petsc_3.11.3_mpt intel_2019.1.144 ABAQUS_2020 intel_2016.3.210	Loci/CHEM Abaqus module compiles, co-simulations run sucessfully

In summary, the first two issues encountered on the K-cluster were overcome by upgrading to Abaqus/2020 and loading the Intel 2016 module last. Aside from the bug with Abaqus/2019, which was confirmed as a SIMULIA problem, the other compilation/execution problems experienced may be indicative of an abnormal and/or incorrect installation of Abaqus on the K-cluster. The attempts outlined in Table A-3 indicate that successful linking to the CSE involves some sort of dependence to Intel 2016. This seems to be the case regardless of which Intel version the Abaqus environment module is dependent on. These types of issues have not been experienced on other high-performance computing (HPC) systems, further raising a red flag about the configuration on K-cluster.

A-2 Nonrepeatable K-Cluster Issues

The final issue encountered on the K-cluster occurred during cosimulation, in which the simulation would occasionally experience a segmentation fault with signal 11. Furthermore, the problem was so severe that it would crash K-cluster nodes along with the simulation. This issue was not repeatable, meaning that the segmentation fault would not occur at the same point in the simulation and in fact sometimes would not occur at all. This issue was also unique to the K-cluster in that ATA had never experienced anything similar on any of the myriad other HPC hardware used to run these types of FSI problems. ATA engaged

SIMULIA to help debug the problem. Per SIMULIA's suggestion, scripts were set up to monitor ports and memory usage, and CSE debugging environment variables were set to enable diagnostic output. However, the problem could not be resolved quickly, and further investigation was out of scope for this effort.

For the quasistatic cosimulations, the issue was not all that problematic because the run times were short (less than a day). Oftentimes the simulations never encountered the segmentation fault, and even when they did, they could be relaunched without much issue. The dynamic cosimulations, on the other hand, required run times of many days, which increased the odds of experiencing a segmentation fault, and after each segmentation fault the simulation would have to be restarted and the relevant output glued together to maintain a coherent time history. The frequency of the segmentation faults made running the dynamic cosimulations on the K-cluster intractable. The simulations were moved to ATA's HPC cluster, where the segmentation faults did not occur.