AIAA AVIATION Forum
August 2-6, 2021, VIRTUAL EVENT
AIAA AVIATION 2021 FORUM

10.2514/6.2021-2726



# CFD Vision 2030 Roadmap: Progress and Perspectives

Andrew W. Cary*
*Boeing Research and Technology, St. Louis, MO, 63042*

John R. Chawner†
*Pointwise, Inc., Fort Worth, TX, 76104*

Earl P.N. Duque‡
*Intelligent Light, Rutherford, New Jersey, 07070*

William Gropp§
*University of Illinois, Urbana, IL, 61801*

Bil Kleb¶
*NASA Langley Research Center, Hampton, VA, 23681*

Ray Kolonay‖
*Air Force Research Lab, Wright-Patterson Air Force Base, OH, 45433*

Eric Nielsen**
*NASA Langley Research Center, Hampton, VA, 23681*

Brian Smith††
*Lockheed Martin Aeronautics Company, Fort Worth, TX, 76101*



## I. Introduction

The CFD Vision 2030 Study [1], contracted by NASA, was published in 2014 and has become an aspirational objective for focusing CFD development on aerospace applications. The Study identified six main technical domains for aligning research: high performance computing (HPC), physical modeling, algorithms, geometry and grid generation, knowledge extraction, and multidisciplinary analysis and optimization (MDAO). Within each domain, subdisciplines (aka elements) illustrate the maturation of technology during the time period from 2015 to 2030. This maturation is illustrated in Figure 1. Each of the elements have their own timeline to identify key technology milestones and technology demonstrations. This Roadmap provides a compact and informative summary of future objectives and milestones, but some details are necessarily omitted from the graphic due to space constraints. These maturing technologies target the use of CFD for aerospace applications of increasing complexity that are represented in the Study by a series of grand challenge problems. At the 2019 AIAA Aviation forum, five years after the Study was published, a special session was held with invited lectures [2–8] reviewing the progress and future in the six domains.

Recognizing the ongoing utility of the Vision toward improving CFD technology, AIAA formed the CFD 2030 Integration Committee (IC). One of the charters of the IC is to maintain the Roadmap by assessing progress and highlighting the ongoing needs of aerospace CFD. Shortly after the 2019 Aviation Forum special session, the IC formed a Roadmap subcommittee to provide a detailed technical review of the Roadmap and identify progress (and challenges to) following the outlined trajectory. The Roadmap subcommittee released a report in June of 2021 [9] with their findings to provide further guidance to the research community. While this report cannot be exhaustive and there are

*Technical Fellow, Boeing Research & Technology, AIAA Senior Member
†President, AIAA Associate Fellow
‡V.P. of Research and Development, AIAA Associate Fellow
§Professor and Director, NCSA
¶Senior Researcher for Computational Aerothermodynamics, AIAA Associate Fellow
‖Principal Engineer, Aerospace Systems Directorate, AIAA Associate Fellow
**Senior Research Scientist, AIAA Associate Fellow
††Lockheed Martin Senior Fellow, AIAA Fellow

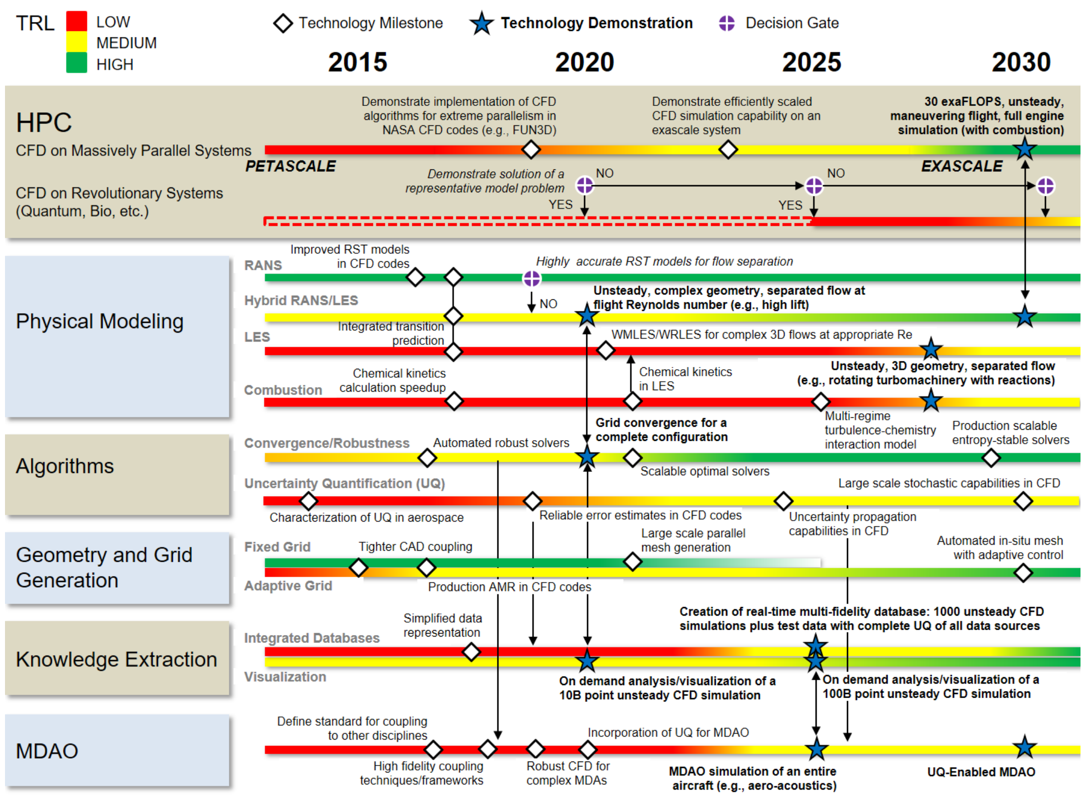


**Fig. 1 Roadmap graphic from original Study [1].**

likely some elements omitted or inadvertently neglected, it represents inputs from a diverse group of experts in the field. This paper describes the approach taken, summarizes the findings, and provides an update to the Vision 2030 Roadmap reflecting the present understanding.

## II. Approach

The Roadmap subcommittee assessed progress by polling experts within each domain and by reviewing conference and journal publications with a focus on the aerospace field. Specific focus was placed on the milestones identified in the original Roadmap. In addition, related technologies that have emerged since the original Study was published were identified. These milestones were assessed for level of completeness and maturity for the industry as a whole. In order to provide a consistent and clear grading scale for each of the identified technologies, a technology readiness level (TRL) scale was utilized with specific interpretations for each of the nine levels. As depicted in Table 1, there is a focus on publicly demonstrated or documented applications of a given technology to provide clear evidence of the maturity for the entire community. This scale is utilized to help assess what technologies are maturing with respect to the Study forecast and which ones are lagging.

The intent of this review is to help focus research and development on the technologies required to achieve the Study's vision of CFD in 2030, while identifying those that are particularly at risk. With this update representing a little over five years since the release of the Roadmap, some milestones are adjusted to reflect the development that has taken place. A TRL score is associated with each of the milestones that will be used to track progress annually.

## III. Roadmap Updates

The Roadmap identified three technology demonstrations to be performed in 2020: 1) unsteady, complex geometry, separated flow predictions at flight Reynolds number (e.g., high lift); 2) grid convergence for a complete configuration; and 3) on-demand analysis/visualization of a 10 billion point mesh, unsteady CFD simulation. These items help to provide a good indication of the advancements made and, while progress has definitely been made on all of these

**Table 1 Technology Readiness Level standard definitions and interpretation for assessing CFD technology development.**

| TRL | NASA Definition [10] | DOD DAG [11] Description | Present Interpretation |
|---|---|---|---|
| 1 | Basic principles observed and reported | Lowest level of technology readiness. Scientific research begins to be translated into applied research and development. | Conference article describing concept/underlying principles |
| 2 | Technology concept and/or application formulated | Invention begins. Once basic principles are observed, practical applications can be invented. Applications are speculative and there may be no proof or detailed analysis to support the assumptions. | Peer-reviewed article describing results from feasibility study. |
| 3 | Analytical and experimental critical function and/or characteristic proof of concept | Active research and development is initiated. This includes analytical studies and laboratory studies to physically validate analytical predictions of separate elements of the technology. | Article or paper demonstrating prototype of capability (with limited scope) |
| 4 | Component and/or breadboard validation in laboratory experiment | Basic technological components are integrated to establish that they will work together. This is relatively “low fidelity” compared to the eventual system. | Capability evaluated/implemented by a CFD team; basic demonstration |
| 5 | Component and/or breadboard validation in relevant environment | Fidelity of breadboard technology increases significantly. The basic technological components are integrated with reasonably realistic supporting elements so it can be tested in a simulated environment. | Successful demonstration of capability on a production-level case |
| 6 | System/subsystem model or prototype demonstration in a relevant environment | Representative model or prototype system, which is well beyond that of TRL 5, is tested in a relevant environment. Represents a major step up in a technology’s demonstrated readiness. | Capability used multiple times by a single CFD team for purposes beyond demonstration (application) |
| 7 | System prototype demonstration in an operational environment | Prototype near, or at, planned operational system. Represents a major step up from TRL 6, requiring demonstration of an actual system prototype in an operational environment such as an aircraft, vehicle, or space. | Use/Evaluation of capability by independent organizations (perhaps in different implementations). This is typically inspired by the successful demonstration of some significant milestone in terms of efficiency, ease of use/robustness, or accuracy. |
| 8 | Actual system completed and qualified through test and demonstration | Technology has been proven to work in its final form and under expected conditions. In almost all cases, this TRL represents the end of true system development. Examples include developmental test and evaluation of the system in its intended weapon system to determine if it meets design specifications. | Application of capability (beyond demonstration) by independent organizations. This implies sufficient robustness for use (value achieved exceeds investment required to obtain it) |
| 9 | Actual system proven through successful mission operations | Actual application of the technology in its final form and under mission conditions, such as those encountered in operational test and evaluation. | Routine/expected use of capability by multiple organizations. OR Acceptance of results by multiple teams |

fronts, there are still enhancements that are required to meet the proposed demonstrations. In the Physical Modeling Domain, there has been steady, incremental progress in the development of unsteady methods for high Reynolds number flight over complex configurations. Test cases from the Third AIAA CFD High Lift Prediction Workshop [12] spurred multiple instances of contributors generating time dependent simulations on the JAXA standard model with promising results using a number of different approaches [13–15]. These results have been positive and a more concentrated study is taking place in preparation for the fourth workshop with significant efforts including time dependent simulations using both wall modeled LES approaches as well as hybrid RANS/LES methods. For the Algorithms Domain technology demonstration, grid convergence has been demonstrated [16–18] for simple configurations (such as the ONERA M6 wing) in different tools with different algorithms. Finite element discretizations have used strong solvers and adaptive meshes to demonstrate progress toward grid convergence. However, these techniques also identify concerns with multiple grid-converged solutions existing for the RANS equations [19]. A large step toward completing a technology demonstration in the Knowledge Extraction Domain was made at SC19. During this event, a large-scale interactive demonstration was presented for a Mars lander configuration [20] simulated using NASA's FUN3D CFD solver on a 6 billion element grid with 200,000 time steps. The interactive visualization of unsteady flow used state-of-the-art hardware (4 DGX-2™ systems each with 16 V100™ GPUs and 16 SSDs for GPUDirect™ storage) that is not typically available to industry, but does represent a significant capability demonstration.

The following subsections review highlights of the achievements in each of the domains through 2020 and provide updates to the Roadmap. In each section, an evaluation of the progress toward milestones is made and milestone additions or adjustments are made based on perceived progress in order to help provide an updated version of the progress toward the Vision outlined in the original Report. A graphical subset of the original and modified Roadmap is presented in each subsection for comparison. This is followed by a table of all of the milestones on the modified Roadmap with an indication of the assessed 2020 TRL (as described above) and an indication of the proposed milestone date. The evolution from the current status to the end state TRL at the deadline is depicted using a color gradient; where the gradient is rapid, more rapid development is suggested. Note that while a milestone should reach TRL 9 to be complete, there are some past milestones that are accepted at a lower TRL level. The symbols in the Roadmap for these are shaded gray to reflect a partial completion.

## A. High Performance Computing

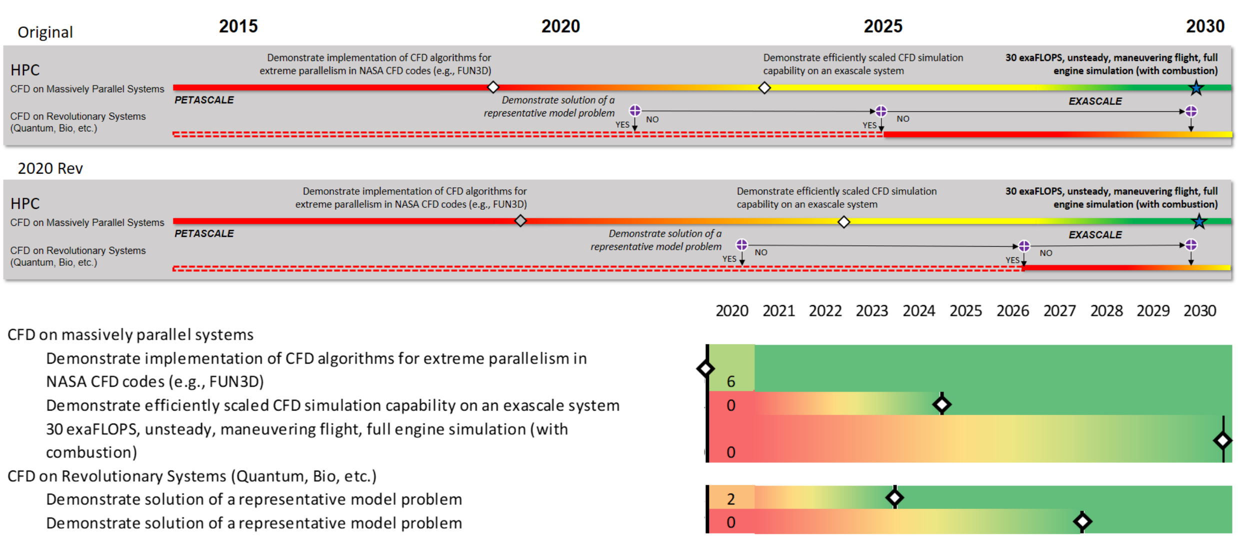


The High Performance Computing (HPC) component of the Study is considered an enabling technology across all domains of the Roadmap. The HPC Domain is further organized into a primary element aimed at an evolutionary progression of more conventional hardware technologies, as well as a secondary element intended to monitor community progress toward potential game-changing use of revolutionary hardware technologies such as quantum and neuromorphic computing for CFD applications. While progress has been made in the field of quantum computing, this technology has yet to penetrate into aerospace CFD. Accelerators based on Graphics Processing Units (GPUs) are an enabling near-term technology to realize exascale computing for practical applications, but generally require a substantial refactoring effort.

The first of the HPC Domain milestones appeared in 2019 and called for a demonstration of the implementation of algorithms for extreme parallelism in a NASA CFD application. At that time, the Department of Energy (DOE) Summit system located at Oak Ridge National Laboratory (ORNL) held the top ranking as the world's most powerful HPC system, enabled predominantly by its six NVIDIA® GPUs available on each compute node [21]. In Refs. [22] and [23], computational campaigns on Summit using the NASA FUN3D unstructured-grid solver are described. In these efforts, a comprehensive port to NVIDIA GPUs based on the CUDA™ programming model was used to demonstrate a performance advantage of 4.5x and 6.5x for the NVIDIA Tesla™ V100 GPU over dual-socket Intel® Xeon™ Gold 6148 (40 cores total) and IBM® POWER9® (44 cores total) CPUs, respectively. Excellent scaling to 1,024 nodes of Summit (6,144 GPUs) was observed, with absolute performance equivalent to approximately 1.2 million Intel Xeon Gold 6148 cores and a nominal per-node performance advantage of 36x for GPU- versus CPU-based simulations on the Summit system. This capability was used throughout 2019 to perform parametric studies of long duration, high-resolution simulations of a supersonic retropropulsion concept for entry, descent, and landing operations of a human-scale Mars lander, using spatial meshes of 6 billion elements and with each simulation producing several hundred terabytes of output data. The port of FUN3D to NVIDIA GPU architectures was the result of several years of workforce development through strategic partnering, extensive software modifications, and the adoption of a steady progression of available hardware features. Minimizing data motion across complex memory hierarchies and identifying approaches to substantially increase node-level parallelism were critical.

While the fundamental hardware technologies for HPC have been relatively stable for the past two decades with periodic hardware refreshes bringing additional processing cores, vectorization support, and improved memory performance, there has not been a pressing need for large-scale application updates. Compilers for common high-level languages such as Fortran, C, and C++ have delivered reasonable performance with minimal developer effort. Node-level parallelism has generally called for O(100) degrees of concurrency, which could be readily achieved through popular shared-memory programming models such as OpenMP [24, 25] or POSIX™ Threads, or the ubiquitous message-passing model of MPI [25].

As the HPC community prepares for the imminent generation of exascale systems, the landscape is undergoing a fundamentally disruptive paradigm shift in the technologies driving today's most powerful computing architectures. Looming manufacturing constraints and power requirements that grow as a strong function of clock speed have forced vendors to seek improved performance through vastly higher levels of concurrency, or parallelism, using processing elements that often operate at reduced clock speeds compared to prior architectures. Increasingly elaborate memory hierarchies abound and often include High Bandwidth Memory (HBM) ideal for the proliferation of memory-bound applications characterized by motifs with low arithmetic intensity. Trends suggest a steady increase in heterogeneity; that is, architectures over the next decade are likely to leverage highly diverse arrays of processing elements most amenable to very specialized tasks.

The RIKEN Center for Computational Science in Japan recently debuted Fugaku, the world's new top system, which delivers a 442-petaflop LINPACK rating [21] and is based on a new ARM processor from Fujitsu. Within the United States, the DOE will accept delivery of three state-of-the-art HPC systems over the next few years starting in late 2021 with the delivery of the Frontier system at Oak Ridge National Laboratory. Frontier will use AMD CPUs and GPUs. Argonne National Laboratory will follow with the Aurora system, which will use Intel CPUs and new Intel GPUs. In 2023, the El Capitan system will arrive at Lawrence Livermore National Laboratory. Like Frontier, this system will also be based on AMD CPUs and GPUs and is expected to realize two exaflops in LINPACK performance. The European community hosts a number of systems incorporating similar CPU and GPU technologies, but also aims to field a new architecture of their own in the next few years through the European Processor Initiative [26].

Migrating or developing large-scale applications to deliver portable performance across a broad swath of diverse, complex architectures generally calls for more expressive programming models beyond the conventional approaches of Fortran, C, and C++. Today's developer has a plethora of options available offering a variety of advantages and disadvantages, including high-level abstractions such as directive-based approaches and abstraction models, entirely new languages specific to a particular hardware architecture, and machine-specific intrinsics sometimes necessary to achieve peak performance. The mapping of application data and algorithms to hardware, as well as latency hiding associated with memory accesses, network traffic, and I/O subsystems is of paramount importance. Mixed-precision algorithms performed on specialized hardware driven by the rapid growth of the machine learning community can yield substantial performance benefits. The use of asynchronous task-based models leveraging knowledge of algorithmic graph dependencies and a run-time scheduler to dynamically allocate work to idle processing elements offers great promise but may also require a complete overhaul of the target application.

The next Roadmap milestone, originally scheduled for 2023, calls for the demonstration of an efficiently-scaled CFD simulation on an exascale system. The third and final milestone in the Roadmap calls for a 30 exaflop simulation of, for example, maneuvering flight including a comprehensive analysis of the propulsion system accounting for full combustion physics. Although Japan's Fugaku boasts an impressive 442-petaflop LINPACK rating, the more appropriate metric may be its HPCG (High Performance Conjugate Gradients) rating [21], widely considered to be more relevant in the context of today's large-scale science and engineering applications. In this measure, Fugaku achieves 16 petaflops. Viewed in this light, the 2030 milestone calling for a 30-exaflop simulation lies some three orders of magnitude beyond the current state of the art. While this considerable gap may appear daunting at the present time, such a rate of growth is remarkably consistent with the historical trend of the past 25 years [21] and speaks to the meticulous deliberations and foresight of the authors of the Study.

Finally, as noted earlier, the revolutionary systems timeline is intended to monitor developments in game-changing hardware technologies, which may hold the potential to bring disruptive change across the fields of science and engineering. The field of quantum computing holds the promise of radically more efficient computation enabled by basic physical properties of quantum physics. Basic research has led to tremendous strides in this field in recent years and large efforts are actively funded by the Department of Energy, the Department of Defense, and the National Science Foundation, among others in the United States, as well as private industry and other nations around the world. NASA has long sponsored the Quantum Artificial Intelligence Laboratory, or QuAIL, project at the NASA Ames Research Center. Recent examples of quantum computing applications include linear algebra, machine learning, differential equations, hybrid approaches combining quantum computing with classical HPC, and recent research simulating quantum computation of a Poisson solver as applied to CFD [27]. Maturing the use of this technology for CFD will hinge on increased interactions between aerospace scientists and engineers and quantum computing researchers.

## B. Physical Modeling

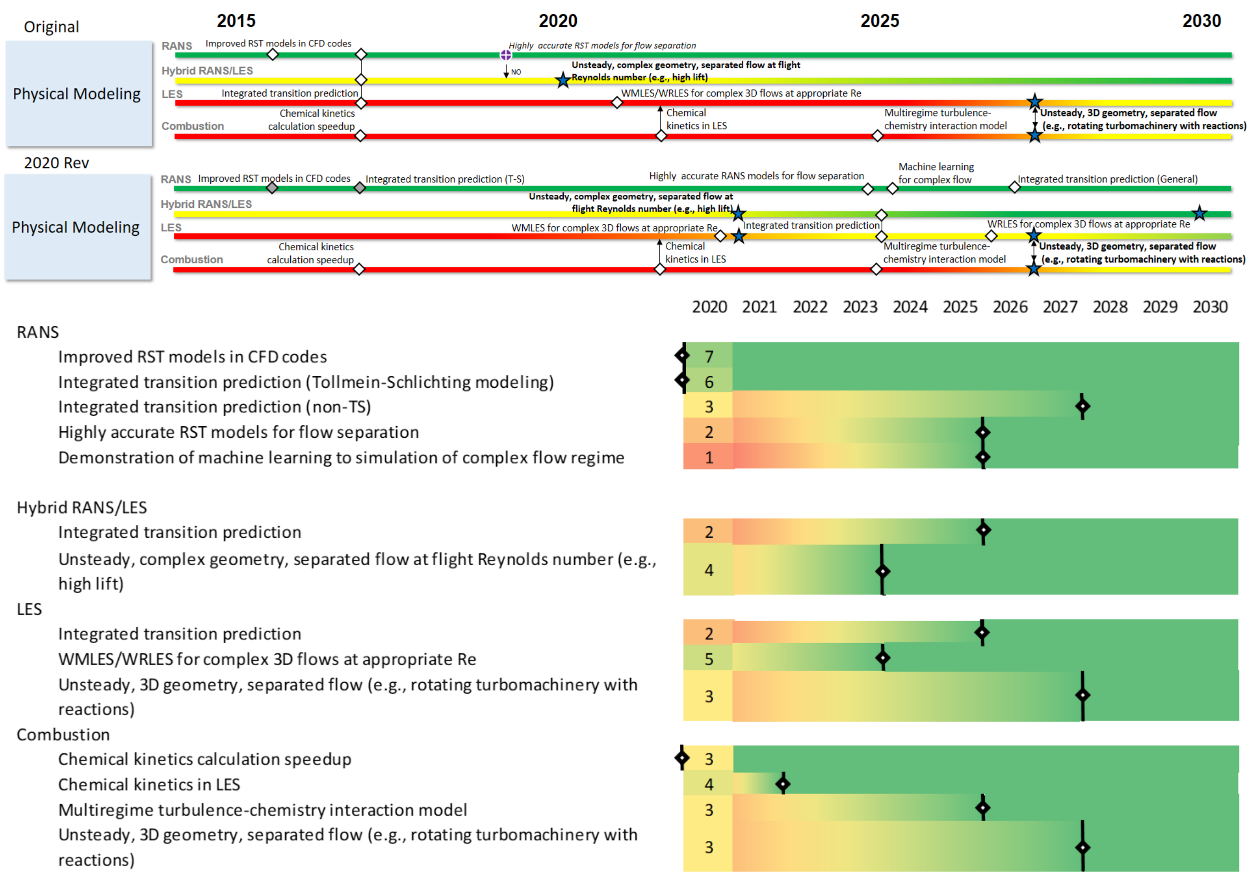

The Physical Modeling Domain encompasses the key modeling technologies required to represent complex physical phenomena for air vehicles including turbulence, transition to turbulence, and complex chemical reaction phenomena related to combustion. The domain is subdivided into four elements: RANS, Hybrid RANS/LES, LES and Combustion. The first three elements related to computational simulation of turbulence are not completely distinct, as aspects of RANS modeling and LES modeling are contained in hybrid models. It is acknowledged that there is a long list of additional physical phenomena that are not included in the Roadmap but that are important for many applications. This list includes icing phenomena, two phase flows in turbulence, real gas effects and plasma phenomena, and modeling for high altitude rarefied gas applications. Key milestones included in the original Roadmap are Improved RST Models in CFD Codes (2016), Highly accurate RST models for flow separation (2019), Integrated transition prediction (2017), Unsteady complex geometry separated flow at flight Reynolds Number (e.g., high lift) (2020) and Chemical kinetics calculation speedup (2017). Progress has been made in many of these areas but the full potential expected from these technologies has not yet arrived.

Reynolds stress transport (RST) models have seen continued development since the publication of the Roadmap. Significant efforts have been made at DLR [28], NASA [29], China [30, 31], and Superior Tech in Lisbon and Maritime Research Institute in the Netherlands [32]. While these models are available for use in several codes including popular commercial flow solvers, outside of DLR, they have not penetrated extensively into aerospace industry applications. These models have shown significant benefit in flows with curvature and swirl as well as corner flows, but have not demonstrated significantly improved predictions of separated flows, a significant focus of the Roadmap. The Roadmap included a decision point in 2019 for RANS turbulence models, particularly with respect to flow separation predictions. While there have only been incremental enhancements in RANS models for this purpose, it is recognized that RANS methods will continue to play an important role in many aircraft industry applications including conceptual design, optimization, and loads prediction. As an example of ongoing development, machine learning (ML) methods have seen a steady increase in research over the last five years. This development, in its infancy when the Study was released and thus not mentioned, has multiple potential implications and has led to the addition of a new milestone.

Spanning all of the turbulence modeling elements was the 2017 milestone for integrated transition prediction, which is recognized as an area of major importance to expanding the fidelity of CFD for many applications. Prior to 2014, a transport model-based method for prediction of Tollmien-Schlichting 2D transition was available through the two-transport equation Langtry-Menter model. Over the past six years there has been active research resulting in the development of additional transport-based transition prediction models, but there has been little convergence among groups or CFD codes around a single model or approach with these new models. For many classes of application, including swept transonic wings, laminar flow designs, low pressure turbines, hypersonic flows, and lower Reynolds number flows representative for small and moderate scale UAVs, these methods are insufficient and accurate and automated predictions of transition are not yet available. These milestones are of key importance but sufficiently far from ready that they must be delayed.

Scale-resolving methods have emerged with uneven success for reaching the technology demonstration for 2020 of modeling unsteady complex geometry separated flow at flight Reynolds number. Some progress in this front has already been alluded to in the introduction. While there has been success in applications of hybrid RANS-LES methods to high angle of attack tactical fighter applications, these are simpler than others because locations of separation onset are largely dictated by sharp leading edges in the thin, swept wings. For other applications, particularly off-design flight regimes such as high lift and lift break conditions, highly accurate simulations are elusive. Hybrid methods are currently most effective where the point of flow separation is fixed by a sharp edge or a shock. Unsteady methods such as Lattice-Boltzmann and wall modeled LES (WMLES) show promise for these flows, but additional computational cost is a barrier, and their accuracy in predicting smooth-body separation with current grid counts is even more debatable than that of hybrid methods. NASA has committed to "develop and demonstrate computationally efficient, eddy-resolving modeling tools that predict maximum lift coefficient ($C_{L,max}$) for transport aircraft with the same accuracy as certification flight tests" (by 2025). These could be either hybrid RANS/LES or WMLES. Given the current state of the art, this is considered an ambitious goal. While the 2020 demonstration was not met, the emphasis in this area by NASA suggests that it may be realized with only a moderate delay.

Carefully controlled validation experiments are also important for the advancement of physical modeling, in particular for turbulent separated flows. The NASA Juncture flow test case evaluated 3D separation in the adverse pressure gradient region of a wing-fuselage juncture [33]. While the initial test campaign was completed in the fall of 2019, post processing and simulation of data continue into 2020 as well as a planned second test entry. Experimental programs in progress to investigate smooth body separation, reattachment and recovery include the VT bump experiment [34] and the Boeing speed bump [35]. Transonic shock induced separation evaluation of a bump on an axisymmetric cylinder, first tested in the 1970s, is being reevaluated with carefully defined tunnel wall boundaries at a range of Reynolds numbers by Lynch et al. [36]. These experimental programs are a critical part of the efforts to improve the fidelity of physical models for CFD.

The final element in the Physical Modeling Domain examines combustion and includes an early milestone for chemical kinetics calculations speedup. For several time-averaged combustor performance parameters, such as overall heat release and spatial temperature distribution, it is sufficient to include very simple (1- to 4-step) chemical kinetics models, or pretabulated chemistry models such as used in flamelet combustion models where the flamelets themselves are computed a priori with detailed chemical kinetic models. These can be done in the RANS or LES frameworks, with LES simulations being more typical for combustor designs that include large recirculation zones and/or strong jet-in-cross flow phenomena. These computations are not very compute-intensive and are performed routinely and chemical kinetics calculation speedup is not a roadblock to performing these simulations. However, other combustor phenomena, including predictions of emissions (oxides of Nitrogen [NOx], carbon monoxide [CO], soot), ignition and extinction events, and assessing impacts of unconventional fuels with significant difference in composition relative to conventional fuels, do require more detailed chemistry approaches and chemical kinetics calculation speedup is required. Prediction of soot entails more complexities, not only in chemical kinetic models, but also in the models that describe the physics of soot production and oxidation. There are multiple approaches being leveraged to provide better inclusion of chemical kinetics in CFD simulations. These include "laminar chemistry", which neglects the effects of turbulence, stochastic methods to include the turbulence-chemistry interactions, and adaptive methods that employ different models within different regions of the simulation. In addition to improving the integration of chemical kinetics into CFD, two additional outstanding issues need to be added to the Roadmap. First, modeling of soot and other nonvolatile particulate matters is becoming more important due to increased environmental concerns. Second, modeling of atomization at relevant operating conditions is key to accurate descriptions of the spray droplet distributions and other properties are key to obtaining accurate predictive simulations of emissions and unsteady phenomenon such as ignition and blowout.

## C. Algorithms

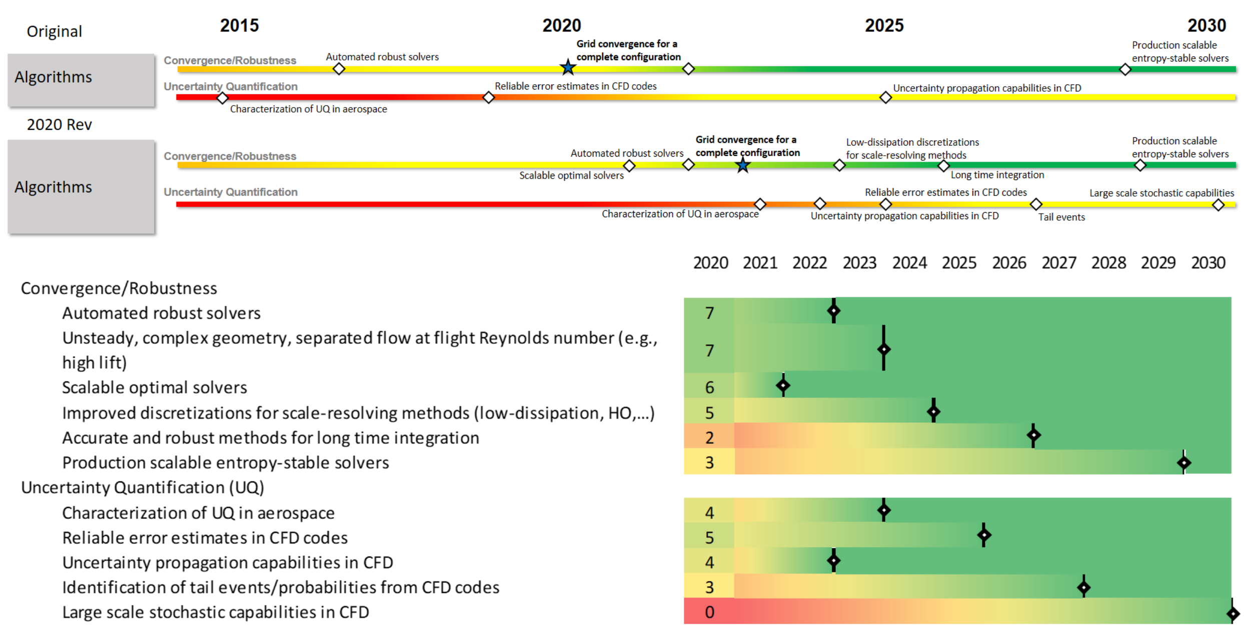


The Algorithm Domain of the Roadmap includes elements for both the convergence and robustness of numerical algorithms and uncertainty quantification (UQ). The overall objectives identified in this domain are associated with estimating the uncertainty in CFD simulations, both through assessment of the sensitivities in the results as well as reducing numerical errors associated with grid resolution and iterative convergence. Technology milestones targeted to have been reached by 2020 include development of technology for automated robust solvers (2016), characterization of uncertainties in aerospace (2015), and reliable error estimates in CFD codes (2019). These were to lead to a 2020 technology demonstration of grid convergence for a complete configuration. As outlined in the introduction to this section, this demonstration was not met for complex configurations, but demonstrated for simpler configurations such as a transonic wing. This case study identified additional considerations associated with achieving grid convergence for RANS simulations as nonunique solutions appear to exist and robust methods for controlling the particular solution have not been developed.

The improvements identified in the Algorithms Domain require significant advancement in several areas of CFD, beginning with enhancing solver robustness leading to improved convergence levels in terms of both iterative convergence and grid convergence. Improvements in grid convergence have been further aided by advances in solution-based mesh adaptation. However, the typically highly-skewed meshes resulting from mesh adaptation increase the challenges for robust solver convergence, leading to further demands for robust and efficient algorithms. The increasing demand for scale-resolving simulations also creates demand for robust, low-dissipation numerical schemes.

A key example of algorithms that improve the solver robustness is the Hierarchical Adaptive Nonlinear Iteration Method (HANIM) developed at NASA Langley in 2016 [37, 38]. HANIM is a highly scalable iterative solver that uses an adaptive pseudotime approach to accelerate iterative convergence. HANIM extends the simple preconditioner of USM3D by providing two additional hierarchies around the preconditioner. The hierarchies are a matrix-free Newton-Krylov linear solver for the exact linearization of discrete RANS equations and nonlinear control of the solution update. This algorithm has also been implemented in FUN3D [39] and has provided robust and dramatically accelerated solutions for NASA configurations, including being able to achieve machine-zero residuals on multiple grid families, while the baseline solver could not attain target low levels of residuals on many of the same grids.

In addition to enhancements for traditional finite volume CFD discretizations, enhancements have been made with other schemes. There are multiple codes successfully using p=1 Streamline Upwind Petrov/Galerkin (SUPG) discretizations for RANS applications [19, 40, 41]. These schemes demonstrate higher accuracy for the same degrees of freedom as finite volume schemes and use robust Newton-Krylov solvers to achieve reliable convergence. Other methods such as Galerkin/Least squares and discontinuous Galerkin formulations are also showing progress.

The results of 2D multielement high-lift computations demonstrate the computational savings potential from using higher-order stabilized continuous Galerkin discretizations in combination with the MOESS output-based mesh adaptation method. High-order methods, in particular p=2 and p=3 continuous Galerkin variational multiscale with discontinuous subscale discretizations, provide accurate outputs with an order of magnitude less computational time than p=1 methods [42]. These approaches lead to high accuracy on coarse grids, but require curved-elements – a new grid generation challenge. Additionally, the solution cost remains an issue for these approaches and improvements in solution techniques will be necessary to make them competitive for steady RANS problems. Furthermore, Lattice Boltzmann codes have been demonstrated to provide industry-level solutions on multiple aerospace applications including aeroacoustics and general unsteady/separated flow applications [14].

UQ has continued to see slow penetration into CFD problems over the last five years. The JANNAF guide [43, 44] from 2015 remains a key reference. In 2016, Barth [45] received the Fluid Dynamic Best Paper award for his overview of application of uncertainty estimates to CFD computations. Although there has been an increase in the number of publications each year, the community is still limited. With the cost of each CFD simulation and the number of uncertain quantities, the use of Monte Carlo techniques has been regarded as infeasible for most applications. This has led to applications of UQ in CFD relying on various types of surrogate modeling including radial basis functions, Gaussian Processes, and nonintrusive polynomial chaos (NIPC) [46]. The AIAA Community of Standards released a preview [47] of their updated standard that provides a general framework for describing CFD uncertainty. This framework provides indications of identifying, characterizing, propagating, and analyzing uncertainties. Important aspects of this characterization are separating random (aleatory) uncertainty from lack of knowledge (epistemic) uncertainty. Some of the key epistemic uncertainties in CFD include the discretization error estimation and model form uncertainty.

An area that has seen particular focus for uncertainty quantification is assessing the uncertainty of turbulence model closure coefficients ranging from simple configurations [48–51] through full configurations such as the Common Research Model [52, 53]. An alternate approach for assessing the impact of turbulence model uncertainty is based on the Lumley triangle and has demonstrated appropriate trends for capturing differences between experiment and simulation [54]. General awareness of UQ techniques beyond Monte Carlo is increasing across the community and multiple organizations are collaborating to provide a consistent terminology and recognition of different viable approaches [55].

Multifidelity techniques offer another key approach to leveraging CFD results for engineering analysis. Examples of this type of approach are given by West [56] and Wendorff [57]. The DARPA-sponsored Scalable Environment for

Quantification of Uncertainty and Optimization in Industrial Applications (SEQUOIA) [58–60] considered several representative applications using multifidelity techniques for both UQ and robust design considerations.

While there have been improvements in the robustness of algorithms for both finite volume and finite element solvers, these have yet to be widely adopted, but are expected to have gained enough general adoption that the milestone will be completed in the near term. These will help to enable both error estimation for mesh adaptation techniques and the ability to reach grid convergence because of degradation of present solvers with increasing grid size. However, the challenges of interpreting/achieving grid convergence if there are multiple possible solutions introduces additional issues. HANIM appears to meet the expectations for a scalable optimal solver but needs increased adoption to achieve the metrics outlined for achieving the milestone. With the increased emphasis on scale-resolving simulations to address accuracy issues in separated flows, there are new milestones added to tailor discretizations for these situations and improve the robustness and accuracy for long time integration. Increases in penetration of UQ have been made since the publication of the Study but it is not yet standard in the aerospace CFD industry. Efforts are underway to increase fluid dynamic community awareness of techniques that have been successfully applied and to develop techniques that are appropriate for CFD to develop reliable error estimates. Although delayed compared to the prediction of the Study, uncertainty quantification methods are expected to play a significant role in the next five years as increased CFD capability and HPC capacity enable both wider and more efficient CFD usage. This change will be enabled by meeting the milestones for having common methods to characterize and propagate uncertainty in CFD analysis, resulting in methods to provide reliable error estimates associated with CFD results.



## D. Geometry Modeling and Mesh Generation

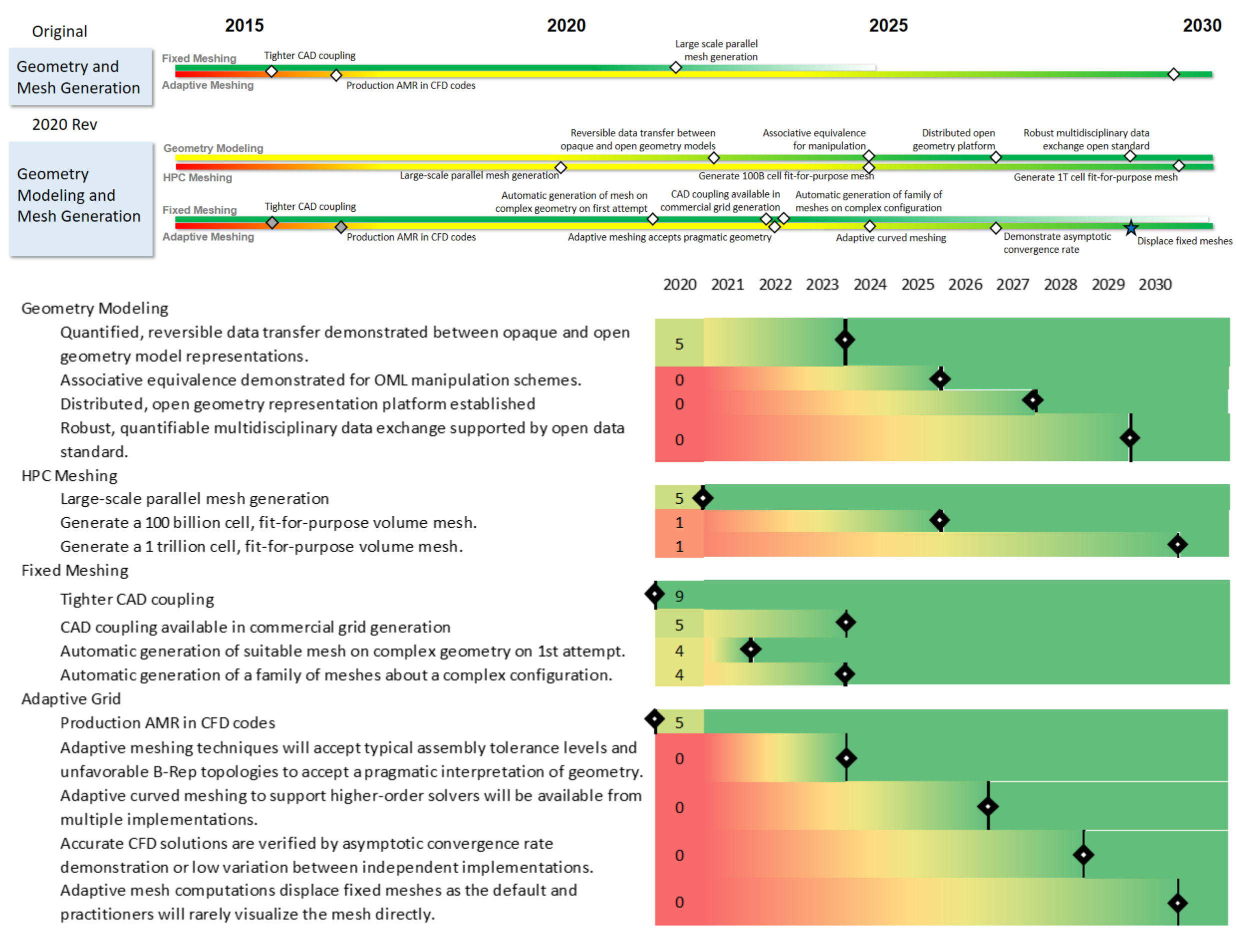


The Geometry Modeling and Mesh Generation Domain of the Roadmap (formerly Geometry and Grid Generation) contains two elements for Fixed and Adaptive Meshing with milestones for achieving tighter coupling with CAD software and the addition to CFD code of production-level adaptive mesh refinement from the period before 2020. Future milestones include large-scale parallel meshing in the 2021 timeframe and automated, adaptive meshing in 2030. Important new elements are being captured by including Geometry Modeling and HPC Meshing as additional elements to provide more refinement of the key tasks necessary to develop robust and automated mesh generation at scales required for the Vision.

The need for simultaneous access to multiple forms of geometric representation are becoming better appreciated throughout the aerospace community [7]. This capability allows various instances of the model to be generated from the same feature-based parameterization. In this way, there is not a single resultant model, but several models may be generated, one for each type of simulation, at the appropriate level of fidelity. Matching and equivalence between models generated from the same Master Geometry is accomplished by attribution on the topological entities. Mesh adaptation fidelity requirements and computationally distributed access to geometry are some of the reasons that traditional geometry kernels need to be redeveloped to meet the requirements for use as part of an integrated system on HPC [61, 62]. Furthermore, a growing awareness of the potential geometric ambiguities inherent in all B-Rep models has led many researchers who are pursuing the Study's goals to either build their own B-Rep geometry modelling kernels (and provide bespoke control of the attendant consequences) [63, 64] or to seek alternative geometry modelling techniques that yield geometrically watertight models [62, 65]. More recently, the focus of research interest has moved away from modelling techniques in which OMLs are parameterized a priori to those where the desired geometry emerges as a result of other, physics-based drivers. The published examples all rely on spatial-occupancy-based geometry modelling techniques – some explicit [66], others implicit [62]. Here, there has been some transference of the lessons originally learned in structural modelling (via topology optimization). Compelling demonstrations are beginning to appear in a wide range of practical applications, not confined to the aerospace sector (or, for that matter, engineering).

The AIAA Geometry and Mesh Generation Workshop series (GMGW) [67, 68] has provided insight into geometry model interoperability issues. In the first two workshops, a model of the NASA High Lift Common Research Model (HL-CRM) was developed and disseminated to participants in multiple formats. all participants for the first workshop identified that modest levels of repair were required but the specific problematic areas varied among mesh generators. This workshop also identified that interoperability issues can be exacerbated when mesh adaptation is used since the maximum permissible tolerances needed to trim the B-Rep surfaces is not known a priori. When a second model used to support manufacture of the wind tunnel model was provided eighteen months later for GMGW-2, there was general concern that the workload involved in preparing it for meshing would likely make the model unsuitable for use in a workshop setting.

Progress in mesh generation has largely addressed the suitability of the mesh for a given analysis that can span from creating a valid mesh, improving cell quality or resolution, and adapting the mesh to the solution and geometry. A large range of techniques are used to support this suitability. In addition to improving suitability, the need for increased sizes of meshes and more integrated analysis processes are driving increased use of HPC for mesh generation.

To facilitate development and assessment of the methods and tools necessary to support unstructured mesh adaptation for the Vision, an open group of researchers formed the Unstructured Grid Adaptation Working Group (UGAWG), published the results of their first benchmarks for adaptivity [69], established an online presence [70], and documented the use of geometry models for mesh adaptation [71]. The working group is notable because its goals are derived directly from the Study and because its membership is cross-organizational, international, and includes representatives from industry, academia, and government.

Another emerging area of mesh generation is the creation of high-order (H-O) meshes for use in H-O flow solvers. To take advantage of the methods, curved meshes are required. While there are a number of techniques that have been applied to deform linear meshes into curved elements, there are particular challenges for high aspect cells consistent with resolving boundary layers in high Reynolds number flows.

Milestones associated with Geometry Modeling and Mesh Generation include continuing the linkage between geometry and mesh, both for automated processes of parametric geometry as well as supporting mesh adaptation. The need for associative equivalence will become more important as multiple representations of models are required for different types of analysis. There are strong desires to increase the level of automation in fixed mesh generation. As adaptive mesh generation becomes more of the default, there will be needs to ensure that it will accept typical assembly tolerances and curved elements.

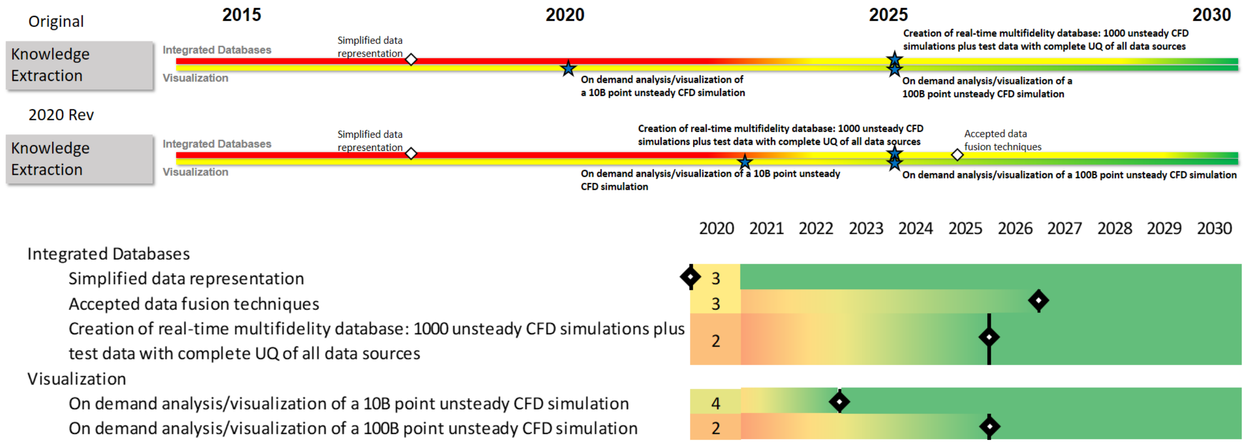


## E. Knowledge Extraction

The Knowledge Extraction Domain focuses on extracting data from simulations to attain engineering results. In this regard, there are several key areas that are included in the Study: database generation and large-scale visualization. Databases of aerodynamic or propulsion results are heavily used to perform more detailed analysis, ranging from estimating trim drag through complete maneuver simulation. The first set of milestones in this domain focuses on being able to generate this type of database, including the development of multifidelity capabilities that incorporate different levels of CFD analysis with time-dependent scale-resolving simulations and experimental data. This timeline has requirements similar to that of a digital thread in ensuring that consistent geometry is maintained and results are able to be assimilated into a database and appropriately combined including estimates of uncertainty. While this is an area of active research, there is no consensus format for representing this rich and varied data. Being able to capture this type of information for aerospace simulations is an emphasis of the first milestone. Although it was forecast to have been met in 2017, there still remains much work to be done, including establishing a format for the data that can represent the necessary range of integrated and nonintegrated data, along with its geometry provenance and data uncertainty. Cambridge University [72], Intelligent Light [73], ANSYS [74], and Sandia National Laboratory [75] have all made significant progress on developing web server-based tools for data processing and analysis.

The second timeline in this domain focuses on large-scale visualization with a technology demonstration slated for 2020 to provide "on-demand analysis/visualization of a 10B point unsteady CFD simulation." At the SC19 conference, NASA and NVIDIA demonstrated an interactive visualization of a 6 billion node, time-dependent simulation of a Mars lander [23]. This user-controlled animation of 150Tb of time-series data in real time required four dedicated NVIDIA DGX-2™ systems each with 16 NVIDIA Tesla™ V100™s and coupled with 16 SSDs to hold the data using NVIDIA's GPUDirect™ to move the data directly across the network to the GPU memory. While this level of hardware is not yet generally available, it provided a powerful demonstration of the near-term potential.

To achieve the levels of performance required, particularly with the larger increases in CFD throughput compared to the increases in storage speed, new strategies are required to meet the visualization demands. These include development of in situ/in transit methods that perform at least part of the data reduction while the CFD simulation is being performed. With in situ processing, the flow solver is instrumented with a data processing/visualization library, such as VisIt-Libsim [76] or Paraview-Catalyst [77], which shares the memory space of the flow solver. This capability often requires that the simulation code pauses while the data are processed. With in transit processing, the simulation data are transferred to a separate set of compute nodes, which processes the data and allows the simulation to resume after the data have been transferred.

Data reduction is often needed to make visualization and knowledge extraction tractable. Data reduction can be performed via either data extracts or by directly creating an image. With data extracts, the visualization graphic objects are created in situ, in transit or via traditional post-processing tools and then written to disk. The user can then load the extract into a conventional post-processor for further analysis or render an image. Since the extract is just a subset of the volume data (i.e., a coordinate cut plane or line extract), it can be many orders of magnitude smaller than the volume data, reducing the amount of data stored and reducing the processing time during user interaction. This approach was demonstrated by Kirby [78] in 2018 to extract knowledge about the wake breakdown of the vortical wake structure in a wind turbine simulation and apply Proper Orthogonal Decomposition (POD) methods to correlate wake breakdown features to POD modes for a 1.2B degree of freedom time-dependent simulation. Another data extract approach is to save off a series of images. The Cinema technology [79] developed at Los Alamos National Laboratory provides mechanisms to interactively visualize large collections of images.

As an indication of the present limits of visualization technology, Tecplot has used their product to visualize data sets from a 1 trillion cell canonical dataset on desktop-sized systems [80]. Cielo et al, [81] visualized the "World's Largest Turbulence Simulation" of an actual simulation — interstellar turbulence that focused on the formation of stars and the associated magnetic fields. Using a custom build of VisIt with OSPRay running on a parallel computer system (512 nodes, 24 cores per node), the researchers were able to perform volume renderings of a turbulent dataset sized at $10{,}048^3$ (1 trillion) cells.

Steady progress is being made on both timelines within this domain, but the Integrated Databases tasks are falling behind the forecast in the original Study. In addition to establishing a flexible and robust schema for representing CFD data, data fusion approaches need to be matured and accepted in the aerospace community to facilitate realizing the large scale databases necessary for achieving the challenge problems and vision laid out in the Study.

## F. Multidisciplinary Analysis and Optimization

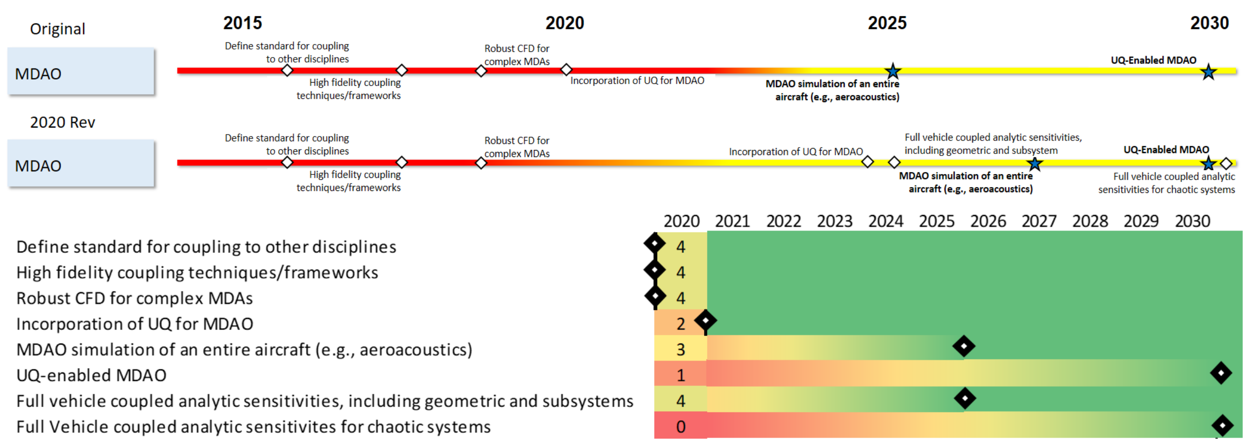


There has been considerable progress made in many areas of MDAO since the release of the Roadmap including increased use of high fidelity CFD in optimization on a more routine basis with a primary multidisciplinary interaction focusing on combining external aerodynamics and structural analysis. To accomplish a full vehicle MDAO, there are typically 100-200 critical conditions for a loads survey that must be assessed, leading to $O(10^4) - O(10^5)$ constraints. Similarly, the number of design variables are large for a system-level, full vehicle MDAO. For example, the number of aerodynamic design parameters such as wing span, chord, and thickness are $O(10^2)$, structural sizing parameters such as skin thickness, spar thicknesses, and moments of inertia are $O(10^3)$, control effector parameters such as number, size, and location of control effectors is $O(10^2)$, system level propulsion parameters are $O(10^2)$, and system thermal management parameters are $O(10^2)$. Hence, a large scale system level full vehicle MDAO problem consists of potentially 10 multifidelity disciplines, 10 couplings, $O(10^4)$ design variables, and $O(10^5)$ constraints. This poses a daunting problem for integrating high-fidelity methods like CFD, but there have been a number of advancements that increase the state of the art.

The MDAO Domain of the Roadmap called for development of coupling standards among disciplines, coupling techniques and frameworks, robust CFD processes, and inclusion of UQ for MDAO by 2020. Significant advancements have been made in these areas, but complete success has not yet been attained.

Primary approaches for multidisciplinary coupling involve either solving the coupled systems of equations representing different disciplines simultaneously or using a partitioned approach where conditions are exchanged at discipline boundaries. The latter approach is more modular and provides more options for readily adding additional disciplines as well as utilizing independent methods for each discipline, but adds complexity with the coupling that must occur at the discipline boundaries, including the geometry. Considering methods for transferring loads and displacements between CFD and computational structural dynamics (CSD) solvers, Kiviaho and Kennedy [82] classified load and displacement transfer techniques into projection-based methods and interpolation-based methods and developed the method of matching-based extrapolation (MELD) that achieves localization similar to projection-based methods for computational efficiency with nonintrusiveness similar to interpolation-based methods. This method also incorporates a consistent approach for developing the necessary sensitivity information required to perform MDO with coupled analysis. In addition to providing a coupling mechanism, the functionality for a CFD code is also required to be defined. While there is no formal standard that has been accepted, there has been progress in developing standard interfaces. Mader et al. [83] developed an approach with the solver as a compiled library with direct memory access and defined both a minimal API and an advanced API for interfacing to other disciplines. This approach has been incorporated into openMDAO [84]. Commercial software vendors such as Dassault Systèmes' SIMULIA® [85], ANSYS [86] and COMSOL® [87] have implemented chained analysis capabilities, but at present none compute fully coupled sensitivities for steady and transient responses that can be used in gradient-based optimization.

Often individual discipline solvers are interfaced through a scripting language "glue" to provide a best in class approach (BCA) for defining a process and solving a selected problem. Frameworks that support the BCA can be found in commercial offerings such as Phoenix Integration's ModelCenter® and Analysis Server® [88], ESTECO's modeFRONTIER and VOLTA [89], VR&D Visual Doc [90], and Dassault Systèmes' SIMULIA Isight Simulation Engine [91]. There are also some U.S. government or open source efforts such as OpenMDAO [84], Dakota [92], GEMS [93], and MSTC Engineering [94] that can support BCA MDAO. While there are multiple benefits of this approach including selecting preferred analysis techniques for individual disciplines, challenges exist due to different data structures and formats for the different tools, often leading to a fragile connection among the disciplines and difficulties in identifying failures. Furthermore, many of these approaches are unable to accept user-supplied sensitivities when performing gradient-based optimization, but this limitation is evolving. For example, OpenMDAO excels in solving design problems involving coupled numerical models of high-fidelity, complex engineering systems. It also has a framework for the computation of the derivatives of these coupled models for use with gradient-based optimization algorithms. This enables the solution of high-fidelity, large-scale ($O(10^3) - O(10^4)$ design variables), gradient-based design optimization. Dakota supports BCA integration with user supplied sensitivities but its primary strengths are in UQ, optimization algorithms, and surrogate modeling. Dakota's UQ methods primarily focus on the forward propagation of uncertainty where probabilistic or interval information on parametric inputs are mapped through the computational model to assess statistics or intervals on outputs. Further, Dakota supports optimization under uncertainty (OUU) and reliability-based design optimization (RBDO). GEMS focus is on developing an automatic programming of MDO processes along with distributed and multilevel MDO formulations (or MDO architectures). Finally, MSTC-Engineering is primarily a research code exploring efficient ways of performing multifidelity large scale geographically distributed MDAO.

Clark [95] gives a comprehensive review of recent work done in including uncertainties and expensive-to-evaluate simulation models for optimization and classifies optimization under uncertainty techniques into three primary areas; reliability-based design optimization (RBDO), robust optimization, and distribution matching techniques. All of these techniques become computationally challenging as the complexity of the modeled physics increases. Recently a large research effort has been funded in the area of UQ by DARPA to address many of the challenges associated with incorporating uncertainty in design optimization. From 2015–2018, DARPA funded EQUiPS (Enabling Quantification of Uncertainty in Physical Systems) [96] acknowledging the importance in characterizing model, parametric, environment, and measure uncertainty when designing complex physical systems, devices and processes. Under this effort, Scalable Environment for Quantification of Uncertainty and Optimization in Industrial Applications (SEQUOIA) developed a strategy for UQ and Design Under Uncertainty (DUU) that advances the scale and complexity of problems [58]. The three thrusts of the program were: 1) scalable algorithms to accommodate a large number of design variables and random parameters, 2) model-form uncertainty using multiple approaches, and 3) design and decision making under uncertainty. Researchers from the Massachusetts Institute of Technology and the University of Texas at Austin have advanced the state of the art in multifidelity MDAO under uncertainty through new methods that leverage Monte Carlo variance reduction techniques and machine learning along with effective reuse of information from past optimization iterations in multifidelity Efficient Global Reliability Analysis (mfEGRA) [97–99].

For progressing toward full-system MDAO, there appears to be two primary vectors that are being followed in the aerospace community. The first is focused on incorporating more disciplines into the vehicle level design process and the second is to bring higher fidelity, coupled analysis into the process. Both are aimed at capturing physics that are being missed during early design. In time, it will be necessary that these two vectors merge to solve the system level full vehicle MDAO problem. The first vector is heavily dependent on two technology areas: robust parametric associative geometric modeling of the inner and outer mold line of the vehicle and analytic sensitivities.

Parametric geometry has been somewhat discussed in the Geometry Modeling and Mesh Generation Domain. Analytic sensitivities from CFD analyses often use adjoint solutions because there is typically a large design space with few constraints and objective functions. Over the last five years, there have been numerous examples (see [100–102] among others) of single-discipline gradient-based optimization using analytic sensitivities. There have also been numerous two-discipline efforts including aerostructural [103–105], aeropropulsive configurations [106, 107], and rotorcraft aerostructural problems [108, 109]. The team of AFRL, Northrop Grumman, and Stanford is executing the quantification of the utility of aerospace derivatives (QUAD) [110] program to accelerate the transition of this capability to industry. The second vector focuses on bringing more disciplines into the design environment. Efforts that focused on including subsystem models into the MDAO environment include the works by Chakraborty and Mavris [111, 112] and the Optimized Integrated Multidisciplinary Systems (OPTIMUS) program [113, 114]. Recently, the Expanded MDO for Effectiveness-Based Design Technologies (EXPEDITE) [115, 116] program has focused on early conceptual MDAO with an important part of the program goals being to advance the state of the art for effectiveness-based design (EBD) [117]. EBD uses mission scenario measures of effectiveness (MoEs) and shifts away from the traditional paradigm of performance-based design, which uses aircraft performance metrics (weight, range, drag, etc.) as the design objective. EBD on the other hand uses mission effectiveness objectives such as cost per available seat mile for commercial air travel. In order to achieve full vehicle MDAO, it is critical that these two primary vectors must merge. In order for that to occur, continued advances must be made in the area of modeling and sensitivity analysis. Efforts need to be initiated to develop analytic sensitivities for many of the subsystems being modeled such as generators, actuators, electrical systems and thermal management systems along with the continued development of high-fidelity physics sensitivities for coupled transient analysis.

While the Study does discuss the importance of sensitivities for MDAO and UQ, it did not create a milestone for sensitivities. It could be reasonably stated that high-fidelity, CFD-based MDAO and UQ will only be possible with analytic derivatives, as every other nongradient-based technique is too expensive if the number of free parameters exceeds a few dozen. Hence, additional specific milestones have been added for "Full vehicle coupled analytic sensitivities, including geometry and subsystems" and "Full vehicle coupled analytic sensitivities for chaotic systems". Although the focus in recent years has been on developing the sensitivities of coupled high-fidelity physics, efforts have to be initiated to compute analytic sensitivities of geometry and subsystem models as well. Examples of this are the open source ESP/CAPS [118–120] effort for solid geometry and PyCycle [121, 122] for chemical equilibrium used in propulsion models. Martins and Hwang [123] provide the unifying chain rule approach to derive sensitivities; this approach has been implemented in the modular analysis and unified derivatives (MAUD) architecture [124]. Transient or unsteady sensitivity analysis is immature at this time due to the challenge of developing efficient and accurate coupled sensitivities. There has been some progress in this approach by numerous researchers using different approaches to avoid the chaotic behavior of solutions [125–129]. Unless efficient and accurate sensitivities are available for all responses (objectives and constraints) participating in the optimization, it will not be possible to perform full gradient-based design, which is currently the only realistic way to handle the number of design variables ($O(10^3) - O(10^4)$) that will be present in full vehicle MDO problems.

## IV. Next Steps

The completely updated Roadmap image is depicted in Figure 2 and illustrates some of the key areas of continued CFD focus toward the Vision. The AIAA CFD Vision 2030 Integration Committee will continue to monitor developments and annually review progress by updating the TRL for each milestone based on public knowledge. These updates will be available on the website [130]. Additional participation in collecting data and evaluating progress is welcome with more details available on the website. Another detailed update of the Roadmap is planned for 2025.

## Appendix

This appendix provides an overview of all of the milestones appearing on the Roadmap organized by domains with brief descriptions to provide more details of the item than is possible in the graphic itself.

### A. High Performance Computing

- Demonstrate implementation of CFD algorithms for extreme parallelism in NASA CFD codes (e.g., FUN3D): Intended to encourage modification of NASA and related CFD codes to efficiently execute on hierarchical memory (GPU/coprocessor) systems by 2020.

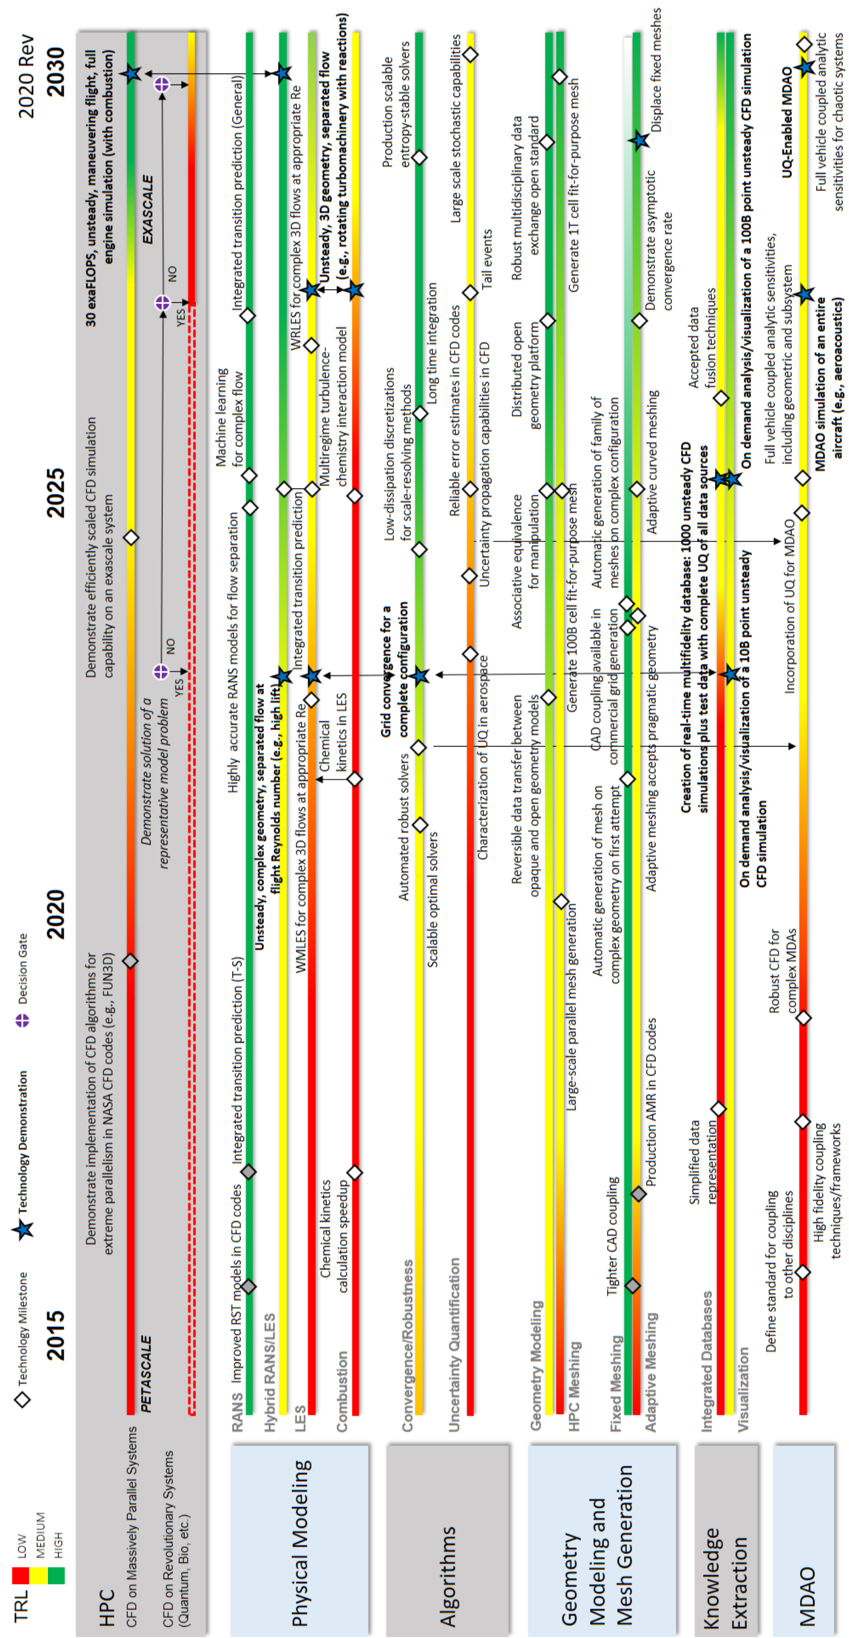


**Fig. 2 Roadmap graphic updated to 2020.**

- Demonstrate efficiently scaled CFD simulation capability on an exascale system: An initial evaluation of exascale performance on a representative CFD problem should be performed.
- 30 exaFLOPS, unsteady, maneuvering flight, full engine simulation (with combustion): A mature exascale implementation should be used to efficiently simulate one or more of the relevant Grand Challenge problems as proposed by the CFD 2030 Integration Committee and the broader AIAA community.
- Demonstrate solution of a representative model problem: Intended to establish the ability of these systems to address problems of interest to the aerospace community (i.e., typical of a Poisson problem for PDE-based simulations). If, at any of these decision points, the technology clearly shows its expected potential, increased investment to accelerate the use of these machines for CFD applications is recommended.

## B. Physical Modeling

- RANS
  - Improved RST models in CFD codes: Reynolds stress transport models that are in general use and provide significant prediction improvement over other RANS models for at least some flow phenomena.
  - Integrated transition prediction (T-S): A transition prediction model that can be applied in a highly automated fashion with a RANS flow field simulation to a full scale transport configuration. The transition location for these applications is often dominated by Tollmein-Schlicting instabilities. At the high Reynolds numbers for these applications, adverse pressure gradient flow typically triggers transition.
  - Highly accurate RANS models for flow separation: Improved Reynolds stress closure models that significantly improve predictions of separation and reattachment relative to present RANS models.
  - Machine learning for complex flow: Demonstration of the application of machine learning methods to improve prediction of a RANS turbulence model for a complex, 3D flow regime. This should be a critical portion of an aircraft related flowfield where standard RANS approaches do not predict the flow well.
  - Integrated transition prediction (General): As detailed in this report, there is are many applications where extended runs of laminar flow and weak pressure gradients, crossflow, environmental factors, hybrid laminar flow techniques, and supersonic flow complicate transition prediction and required more sophisticated transition prediction methods than can be captured by methods available at this time.
- Hybrid RANS/LES
  - Unsteady, complex geometry, separated flow at flight Reynolds number (e.g., high lift): Ability to generate consistent, improved predictions with hybrid RANS/LES models that are predictive. Predictive models switch from RANS to LES without a priori knowledge of transition locations for boundary layer separation.
  - Integrated transition prediction: A transition prediction method than is applicable in a region where flow separation occurs in a region of laminar or transitional flow and the hybrid model switches to LES mode in this region.
- LES
  - WMLES for complex 3D flows at appropriate Re: A single or unified method that can accurately predict complex flows such as realistic configuration aircraft at lift break with good accuracy in both attached and separated flow regions of the flow field at flight Reynolds numbers.
  - Integrated transition prediction: An accurate automated transition prediction method that is embedded within an WMLES or WRLES simulation method
  - WRLES for complex 3D flows at appropriate Re: A single or unified method that can accurately predict complex flows such as realistic configuration aircraft at lift break with good accuracy in both attached and separated flow regions of the flow field at flight Reynolds numbers.
  - Unsteady, 3D geometry, separated flow (e.g., rotating turbomachinery with reactions): A single method that can accurately predict complex flows including rotation, chemical reactions, separation for realistic flow conditions. Flow withing a compressor, combustor and turbine would be an example.
- Combustion
  - Chemical kinetics calculation speedup:
  - Chemical kinetics in LES
  - Multiregime turbulence-chemistry interaction model

## C. Algorithms

- Convergence / Robustness

  - Scalable optimal solvers: Flow solver algorithms that scale well with increased problem size and number of processors to reach toward exascale-level performance
  - Automated robust solvers: Flow solvers that routinely and without user intervention or customization can reliably converge to machine zero, including when there is additional stiffness associated with multidisciplinary simulations.
  - Grid convergence for a complete configuration: Demonstrated solution independence of grid for relevant quantities of interest on a complete aerospace vehicle.
  - Low-dissipation discretizations for scale-resolving methods: Demonstration of methods that have small enough numerical dissipation compared to the physically modeled dissipation at the desired LES cut-off frequency on relevant grid topologies.
  - Long time integration: Demonstrate algorithms that are efficient at reaching long times (10s of seconds for complex events) while appropriately representing necessary high frequencies.
  - Production scalable entropy-stable solvers: Extension of algorithms that are inherently mathematically stable to production environment whileremaining applicable to exascale class machines.
- Uncertainty Quantification
  - Characterization of UQ in aerospace: Accepted approaches and representations for describing uncertainty elements present in aerospace applications.
  - Uncertainty propagation capabilities in CFD: Development and acceptance of efficient methods for assessing the impact of uncertainty in output quantities of interest predicted using CFD. Uncertainty sources include input uncertainty, model form uncertainty, and numerical/discretization errors.
  - Reliable error estimates in CFD codes: Demonstrable methods of discretization error estimation appropriate for modern solvers on representative production CFD problems.
  - Tail events: Development of methods that will identify infrequent events and characterize the likelihood of occurrence.
  - Large-scale stochastic capabilities: Demonstration of approach that provides appropriate uncertainty distributions for a full configuration across a range of operating conditions with representative number of uncertain quantities.

### D. Geometry Modeling and Mesh Generation

- Geometry Modeling
  - Reversible data transfer between opaque and open geometry models: Quantified, reversible data transfer demonstrated between opaque and open geometry model representations.
  - Associative equivalence for manipulation: Associative equivalence demonstrated for OML manipulation schemes.
  - Distributed open geometry representation platform: Established a distributed (multimachine) geometry representation that is openly available.
  - Robust multidisciplinary data exchange open standard: Accepted robust, quantifiable multidisciplinary data exchange supported by open data standard.
- HPC Meshing
  - Large scale parallel mesh generation: Creation of a mesh on a complete configuration using $O(100+)$ cores with scalable efficiency
  - Generate a 100 billion cell, fit-for-purpose volume mesh.
  - Generate a 1 trillion cell, fit-for-purpose volume mesh.
- Fixed Meshing
  - Tighter CAD Coupling: CFD solvers, regardless of the platform on which they run including HPC and regardless of the format in which geometry models are provide, will have efficient access to geometry definitions for operations such as mesh adaptation and shape changes.
  - Automatic generation of mesh on complex geometry on first attempt: Automatic generation of a suitable mesh on the first attempt such as for the 4th CFD High-Lift Prediction Workshop.
  - CAD coupling available in commercial grid generation
  - Automatic generation of family of meshes about a complex configuration: This family of meshes could represent a grid refinement study or parametric variations of a component of the design.
- Adaptive Meshing

- Production AMR in CFD codes:
- Adaptive meshing accepts pragmatic geometry: Adaptive meshing techniques will accept typical assembly tolerance levels and unfavorable B-Rep topologies to accept a pragmatic interpretation of geometry.
- Adaptive curved meshing: Support for higher-order solvers will be available from multiple implementations.
- Demonstrate asymtotic convergence rate: Accurate CFD solutions are verified by asymptotic convergence rate demonstration or low variation between independent implementations.
- Displace fixed meshes: Adaptive mesh computations displace fixed meshes as the default and practitioners will rarely visualize the mesh directly.

### E. Knowledge Extraction

- Integrated Databases
  - Simplified data representations: Development of a standard schema for representing different types of CFD data, including meta data, scalar results, and field data.
  - Creation of real-time multifidelity database: Development of a database with 1000 unsteady CFD simulations plus test data with complete UQ of all data sources representative of the type of information necessary for performing a flight profile analysis.
  - Accepted data fusion techniques: Established and accepted methods for combining multiple types of data, including different CFD approaches and test data, to leverage the cost and accuracy of each approach.
- Visualization
  - On demand analysis/visualization of a 10B point unsteady CFD simulation
  - On demand analysis/visualization of a 100B point unsteady CFD simulation

### F. Multidisciplinary Analysis and Optimization

- Define standard for coupling to other disciplines: Accepted definition of API or data required to be provided to/from a flow solver to interface with other disciplines for generating multidisciplinary simulations.
- High-fidelity coupling techniques/frameworks: Demonstrated accurate and robust approaches for interfacing high-fidelity CFD codes to exchange both discipline interface information and associated sensitivities.
- Robust CFD for complex MDAs: Demonstration of CFD processes, including both mesh deformation and solution evaluation, that is automated and robust for a range of multidisciplinary analyses for a complex configuration representative of a full vehicle or subsystem.
- Incorporation of UQ for MDAO: Demonstration of the computation of uncertainties in relevant performance metrics including the effect of multiple disciplines.
- Full vehicle coupled analytic sensitivities, including geometric and subsystems: Computation of multidisciplinary analytic sensitivities to appropriate quantities of interest for a range of design variables representative of a full vehicle analysis.
- MDAO simulation of an entire aircraft (e.g., aeroacoustics)
- UQ enabled MDAO: Appropriate inclusion of uncertainty quantification in a MDAO analysis to provide a design that is robust to both physical and modeling unknowns.
- Full vehicle coupled analytic sensitivities for chaotic systems: Demonstration of analytic sensitivities generated from a chaotic flow field such as at high angles of attack with massive separation (e.g., commercial high lift).

## Acknowledgments

As part of gathering the progress and refining assessments and additional needs, the authors received feedback from a wide variety of individuals, including M.S. Anand (Rolls-Royce), Anirban Chaudhuri (University of Texas at Austin), Daniel Clark (AFRL), Clifton Davies (Lockheed Martin), Boris Diskin (NIA), Edwin Forster (AFRL), John Gallman (Northrop Grumman), Ramana Grandhi (Air Force Institute of Technology), Justin Gray (NASA), Jim Guglielmo (The Boeing Company), George Huang (Wright-State University), Kenneth Jansen (University of Colorado), Dmitry Kamenetskiy (The Boeing company), Stephen de Bruyn Kops (University of Massachusetts), Steve M. Legensky (Intelligent Light), Joaquim Martins (University of Michigan), Mory Mani (The Boeing Company, retired), Dimitri Mavriplis (University of Wyoming), David McLaurin (Siemens Digital Industries Software), Jeffrey Moder (NASA Glenn Research Center), Hiro Nishikawa (NIA), Michael A. Park (Computational AeroSciences Branch, NASA Langley Research Center), Chris Rumsey (Computational AeroSciences Branch, NASA Langley Research Center), Scott Sellers

(The Boeing Company), Jeff Slotnick (The Boeing Company), Nagendra Somanath (Pratt & Whitney), Philippe Spalart (The Boeing Company, retired), Bret Stanford (NASA), Marc Stelmack (Lockheed Martin), Nigel Taylor (MBDA UK Ltd), Nick Wyman (Pointwise), Karen Wilcox (University of Texas at Austin), and PK Yeung (Georgia Institute of Technology). Your contributions and considerations are very much appreciated.

# References

[1] Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D., "CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences," Tech. Rep. NASA/CR-2014-218178, NASA, March 2014.

[2] Cavolowsky, J., "NASA Aeronautics and CFD 2030," *Oral Presentation*, June 2019.

[3] Slotnick, J., "Progress Towards CFD Vision 2030 - An Industrial Perspective for Air and Space Vehicle Applications," *Oral Presentation*, June 2019.

[4] Medic, G., "Impact of Vision 2030 on CFD Practices in Propulsion Industry," *AIAA Paper 2019–2943*, June 2019.

[5] Nielsen, E., "High Performance Computing Towards CFD Vision 2030," *Oral Presentation*, June 2019.

[6] Mavriplis, D., "Progress in CFD Discretizations, Algorithms and Solvers for Aerodynamic Flows," *AIAA Paper 2019–2944*, June 2019.

[7] Chawner, J., and Taylor, N., "Progress in Geometry Modeling and Mesh Generation Toward the CFD Vision 2030," *AIAA Paper 2019–2945*, June 2019.

[8] Spalart, P., and Strelets, M., "Turbulence Prediction in Aerospace CFD: Reality and the Vision 2030 Roadmap," *Oral Presentation*, June 2019.

[9] Cary, A. W., Chawner, J. R., amd W. Gropp, E. P. N. D., Kleb, B., Kolonay, R., Nielsen, E., and Smith, B., "The CFD Vision 2030 Roadmap: 2020 Status, Progress, and Challenges," Tech. rep., 2021.

[10] NASA, "Technology Readiness Level Definitions," `https://www.nasa.gov/pdf/458490main_TRL_Definitions.pdf`, 2019. Accessed: 2020-10-15.

[11] "Technology Readiness Assessment (TRA) Deskbook," Tech. rep., Department of Defense, July 2009.

[12] Rumsey, C., "3rd AIAA CFD High Lift Prediction Workshop (HiLiftPW-3)," `https://hiliftpw.larc.nasa.gov/index-workshop3.html`, 2016. Accessed: 2020-10-15.

[13] Cary, A., Yousuf, M., Li, P., and Mani, M., "Current Practice Unstructured Grid CFD results for 3rd AIAA High Lift Prediction Workshop," *AIAA Paper 2018–1037*, January 2018.

[14] Konig, B., Fares, E., Murayama, M., and Ito, Y., "PowerFLOW Simulations for the Thurd AIAA High-Lift Prediction Workshop," *AIAA Paper 2018–1255*, January 2018.

[15] Gok, K., Bose, S., and Moin, P., "Wall-Modeled Large Eddy Simulation of an Aircraft in Landing Configuration," *AIAA Paper 2020-3002*, June 2020.

[16] Park, M., Barral, N., Ibanez, D., Kamenetskiy, D., Krakos, J., Michal, T., and Loseille, A., "Unstructured Grid Adaptation and Solver Technology for Turbulent Flows," *AIAA Paper 2018–1103*, January 2018.

[17] Fidkowski, K., "Three-Dimensional Benchmark RANS Computations Using Discontinuous Finite Elements on Solution-Adapted Meshes," *AIAA Paper 2018–1104*, January 2018.

[18] Balan, A., Park, M., Anderson, W., Kamenetskiy, D., Krakos, J., Michal, T., and Alauzet, F., "Verification of Anisotropic Mesh Adaptation for Turbulent Simulations over ONERA M6 Wing," *AIAA Journal*, Vol. 58, April 2020.

[19] Kamenetskiy, D., Bussoletti, J., Hilmes, C., Venkatakrishnan, V., Wigton, L., and Johnson, F., "Numerical Evidence of Multiple Solutions for the Reynolds-Averaged Navier-Stokes Equations," *AIAA Journal*, Vol. 52, 2014.

[20] Salian, I., "How NASA is Helping Humans Reach the Red Planet, Using GPUS and 3D Visualization," `https://blogs.nvidia.com/blog/2019/11/18/nasa-mars-landing-simulation-gpus`, 2019. Accessed: 2020-10-15.

[21] Strohmaier, E., Dongarra, J., Simon, H., Meuer, M., and Meuer, H., "The Top 500 List," `https://www.top500.org`, 2020. Accessed: 1 January 2021.

[22] Korzun, A., Nielsen, E., Walden, A., Jones, W., Carlson, J.-R., Moran, P., Henze, C., and Sandstrom, T., "Effects of Spatial Resolution on Retropropulsion Aerodynamics in an Atmospheric Environment," *AIAA Paper 2020–1749*, January 2020.

[23] Moran, P., Sandstrom, T., Korzun, A., Nielsen, E., and Walden, A., "Visualizing Supersonic Retropropulsion for Mars: The Summit Early Science Results," *Scientific Vizualization and Data Analytics Showcase at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Denver, CO, 2019.

[24] "The OpenMP Website," `http://www.openmp.org`, 2021. Last accessed 8/1/20.

[25] `https://www.mpi-forum.org`, 2021. Last accessed 8/1/20.

[26] `https://www.european-processor-initiative.eu`, 2021. Last accessed 8/1/20.

[27] Steijl, R., and Barakos, G., "Parallel Evaluation of Quantum Algorithms for Computational Fluid Dynamics," *Computers and Fluids*, 2018, pp. 173:22–28.

[28] Morsbach, C., Franke, M., and di Mare, F., *Differential Reynolds Stress Modeling for Separating Flows in Industrial Aerodynamics*, Springer International Publishing, 2015, Chap. Application of a Low Reynolds Differential Reynolds Stress, pp. 1–18.

[29] Eisfeld, B., and Rumsey, C. L., "Length-Scale Correction for Reynolds-Stress Modeling," *AIAA Journal*, Vol. 58, No. 4, 2020, pp. 1518–1528.

[30] Wang, S., Dong, Y., Deng, X., Wang, G., and Wang, J., "High-Order Simulation of Aeronautical Separated Flows with a Reynold Stress Model," *Journal of Aircraft*, Vol. 55, No. 3, 2018, pp. 1177–1190.

[31] Yang, R., guo Wang, Z., xin Zhao, Y., cheng Wang, Q., and heng Feng, W., "Numerical investigation on spatial development of the secondary flow in a supersonic turbulent square duct," *Aerospace Science and Technology*, Vol. 100, 2020.

[32] Pereira, F. S., Eça, L., and Vaz, G., "Simulation of Wingtip Vortex Flows with Reynolds-Averaged Navier–Stokes and Scale-Resolving Simulation Methods," *AIAA Journal*, Vol. 57, No. 3, 2019, pp. 932–948.

[33] Kegerise, M. A., Neuhart, D., Hannon, J., and Rumsey, C. L., "An Experimental Investigation of a Wing-Fuselage Junction Model in the NASA Langley 14- by 22-Foot Subsonic Wind Tunnel," *AIAA Paper 2019-0077*, January 2019.

[34] Lowe, T., Beardsley, C., Borgolz, A., Devenport, W. J., Duetsch-Patel, J. E., Fritsch, D. J., Gargiulo, A., Roy, C. J., Szoke, M., and Vishwanathan, V., "Status of the NASA/Virginia Tech Benchmark Experiments for CFD Validation," *AIAA Paper 2020-1584*, January 2020.

[35] Slotnick, J., "Integrated CFD Validation Experiments for Prediction of Turbulent Separated Flows for Subsonic Transport Aircraft," *NATO Science and Technology Organization, Meeting Proceedings RDP, STO-MP-AVT-307*, 2019.

[36] Lynch, K. P., Lance, B., Lee, G. S., Naughton, J. W., Miller, N. E., Barone, M. F., Beresh, S. J., Spillers, R., and Soehnel, M., "A CFD Validation Challenge for Transonic, Shock-Induced Separated Flow: Experimental Characterization," *AIAA Paper 2020-1309*, January 2020.

[37] Pandya, M. J., Diskin, B., Thomas, J., and Frink, N., "Improved Convergence and Robustness of USM3D Solutions on Mixed Element Grids," *AIAA Journal*, Vol. 54, No. 9, 2016, pp. 2589–2610.

[38] Pandya, M. J., Diskin, B., Thomas, J., and Frink, N., "Assessment of USM3D Hierarchical Adaptive Nonlinear Iteration Method Preconditioners for Three-Dimensional Cases," *AIAA Journal*, Vol. 55, No. 10, 2017, pp. 3409–3424.

[39] Nielsen, E., and Diskin, B., "High-Performance Aerodynamic Computations for Aerospace Applications," *Parallel Computing*, Vol. 64, 2017, pp. 20–32.

[40] Galbraith, M., Carmen-Ioana Ursachi, C.-I., Chandel, D., Allmaras, S., Darmofal, D., Glasby, R., Stefanski, D., Erwin, J., Holst, K., Hereth, E., Mukhopadhaya, J., and Alonso, J., "Comparisons of HPCMP CREATE-AV Kestrel-COFFE, SU2, and MIT SANS RANS Solutions using Output-Based Adapted Meshes for a Multi-Element Airfoil," *AIAA Paper 2021-1080*, January 2021.

[41] Wood, S., and Anderson, W., "Finite-Element and Finite-Volume Solutions for the Juncture-Flow Experiment," *AIAA Paper 2020-2751*, June 2020.

[42] Ursachi, C.-I., Galbraith, M., Allmaras, S., and Darmofal, D., "Output-based adaptive RANS solutions using higher-order FEM on a multi-element airfoil," *AIAA Paper 2020-3220*, June 2020.

[43] Mehta, U., Romero, V., andJ. Pearce, D. E., and Keim, N., "The JANNAF Simulation Credibility Guide on Verification, Uncertainty Propagation and Quantification, and Validation," *AIAA Paper 2015-0502*, June 2015.

[44] Mehta, U., Eklund, D., Romero, V., Pearce, J., and Keim, N., "Simulation Credibility, Advances in Verification, Validation, and Uncertainty Quantification," Tech. rep., NASA/TP 2016-219422, JANNAF/GL 2016-0001, 2016.

[45] Barth, T., "An Overview of Combined Uncertainty and A-Posteriori Error Bound Estimates for CFD Calculations," *AIAA Paper 2016-1062*, January 2016.

[46] Quagliarella, D., Serani, A., Diea, M., Pisaroni, M., Leyland, P., Montagliani, L., Iemma, U., Gaul, N., Shin, J., Wunsch, D., Hirsch, C., Choi, K., and Stern, F., "Benchmarking Uncertainty Quantification Methods Using the NACA 2412 Airfoil with Geometrical and Operational Uncertainties," *AIAA Paper 2019-3555*, 2019.

[47] Le, H., Ghia, U., Bayyuk, S., Oberkampf, W., Roy, C., Benek, J., Rumsey, C., Powers, J., Bush, R., and Mani, M., "Development and Use of Engineering Standards for Computational Fluid Dynamics for Complex Aerospace Systems," *AIAA Paper 2016-3811*, 2016.

[48] Schaefer, J., Hosder, S., West, T., Rumsey, C. L., Carlson, J.-R., and Kleb, W., "Uncertainty Quantification of Turbulence Model Closure Coefficients for Transonic Wall-Bounded Flows," *AIAA Journal*, Vol. 55, No. 1, 2016.

[49] Schaefer, J., Cary, A., Mani, M., Krakos, J., and Hosder, S., "Grid Influence on Turbulence Model Coefficient Uncertainties in Transonic Wall-Bounded Flows," *AIAA Journal*, Vol. 56, No. 8, 2018.

[50] Stefano, M. D., Hosder, S., and Baurle, R., "Effect of Turbulence Model Uncertainty on Scramjet Isolator Flowfield Analysis," *Journal of Propulsion and Power*, Vol. 36, No. 1, 2019.

[51] Erb, M., and Hosder, S., "Analysis of Turbulence Model Uncertainty for Shock-Wave/Boundary-Layer Interaction Simulations," *Journal of Spacecraft and Rockets*, 2020.

[52] Schaefer, J., Cary, A., Mani, M., and Spalart, P., "Uncertainty Quantification and Sensitivity Analysis of Spalart-Allmaras Turbulence Model Coefficients in Two and Three Dimensions," *AIAA Paper 2017-1710*, 2017.

[53] Cary, A., Schaefer, J., Duque, E., and Lawrence, S., "Application of a CFD Uncertainty Quantification Framework for Industrial Scale Aerodynamic Analysis," *AIAA Paper 2019-1492*, 2019.

[54] Mishra, A., Mukhopadhaya, J., Iaccarino, G., and Alonso, J., "Uncertainty Estimation Module for Turbulence Model Prediction in SU2," *AIAA Journal*, Vol. 57, No. 3, 2018.

[55] Schaefer, J., Romero, V., Schafer, S., Leyde, B., and Denham, C., "Approaches for Quantifying Uncertainties in Computational Modeling for Aerospace Applications," *AIAA Paper 2020-1520*, 2020.

[56] West, T., and Phillips, B., "Multifidelity Uncertainty Quantification of a Commercial Supersonic Transport," *Journal of Aircraft*, Vol. 57, No. 3, 2020.

[57] Wendorff, A., Alonso, J., and Bieniawski, S., "Using Multiple Information Sources to Construct Stochastic Databases to Quantify Uncertainty in Certification Maneuvers," *AIAA Paper 2016-1420*, 2016.

[58] Alonso, J., Eldred, M., Constantine, P., Duraisamy, K., Farhat, C., Iaccarino, G., and Jakeman, J., "Scalable Environment for Quantification of Uncertainty and Optimization in Industrial Applications," *AIAA Paper 2017-1327*, 2017.

[59] Eldred, M., Geraci, G., Gorodetsky, A., and Jakeman, J., "Multilevel-Multidelity Approaches for Forward UQ in the DARPA SEQUOIA project," *AIAA Non-Deterministic Approaches Conference*, 2018.

[60] Geraci, G., Eldred, M., Gorodetsky, A., and Jakeman, J., "Recent advancements in Multilevel-Multifidelity techniques for forward UQ in the DARPA Sequoia project," *AIAA Paper 2019-0722*, 2019.

[61] Haimes, R., and Dannenhoffer, J., "EGADSlite: a lightweight geometry kernel for HPC," *AIAA Paper 2018-1401*, 2018.

[62] Dawes, W., and et al, "Digital Geometry to Support a Gas Turbine Digital Twin," *AIAA Paper 2019-1715*, 2019.

[63] Bryson, D., Haimes, R., and Dannenhoffer, J., "Towards the Realization of a Highly Integrated, Multidisciplinary, Multifidelity Design Environment," *AIAA Paper 2019-2225*, 2019.

[64] Mestreau, E., Aubry, R., Dey, S., and Richardson, M., "CREATE-MG Capstone Architecture and Feature Modeling Improvements," *AIAA Paper 2019-1716*, 2019.

[65] Andersson, J., and et al, "Parametrically Controlled Subdivision Surfaces for Conceptual Design," *AIAA Paper 2017-4005*, 2017.

[66] Payot, A. D. J., and et al, "Parametric Surfaces with Volume of Solid Control for Optimisation of Three Dimensional Aerodynamic Topologies," *AIAA Paper 2019-1473*, 2019.

[67] Chawner, J., Michal, T., Slotnick, J., and Rumsey, C., "Summary of the 1st AIAA Geometry and Mesh Generation Workshop (GMGW-1) and Future Plans," *AIAA Paper 2018-0128*, 2018.

[68] Taylor, N., and Gammon, M., "The Wind Tunnel Model of the NASA High Lift Common Research Model: a Geometry-Handling Perspecitve," *AIAA Paper 2019–3460*, June 2019.

[69] Ibanez, D., Barral, N., Krakos, J., Loseille, A., Michal, T., and Park, M., "First Benchmark of the Unstructured Grid Adaptation Working Group," *Procedia Engineering*, 2017, pp. 1877–2058.

[70] `https://www.github.com/UGAWG`, May 2019.

[71] Park, M., Kleb, B., Jones, W. T., Loseille, A., Krakos, J., and Michal, T., "Geometry Modeling for Unstructured Mesh Adaptation," *AIAA Paper 2019-2946*, 2019.

[72] Pullan, G., "A Web-Based Database Approach to CFD Post-Processing," *AIAA Paper 2017-0814*, 2017.

[73] Duque, E., and Lawrence, S., "Spectre: A Computational Environment for Managing Total Uncertainty Quantification of CFD Studies," *AIAA Paper 2019-2221*, 2019.

[74] "Introduction, Nexus," , May 26, 2021. URL `https://nexusdemo.ensight.com/docs/html/Introduction.html`.

[75] Friedman-Hill, E. K., Hoffman, E. L., Gibson, M. J., Clay, R. L., and Olson, K. H., "Incorporating Workflow for V&V/UQ in the Sandia Analysis Workbench," Tech. rep., SAND2015-0425C, 2015.

[76] Whitlock, B., Favre, J. M., and Meredith, J. S., "Parallel In Situ Coupling of Simulation with a Fully Featured Visualization System," *Eurographics Symposium on Parallel Graphics and Visualization*, 2011. Lawrence Livermore National Labs LLNL-SM-446033.

[77] Ayachit, U., Bauer, A., Geveci, B., O'Leary, P., Moreland, K., Fabian, N., and Maudlin, J., "Paraview catalyst: Enabling in situ data analysis and visualization," *Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*, 2015.

[78] Kirby, A. C., Yang, Z., Mavriplis, D., Duque, E. P., and Whitlock, B. J., "Visualization and data analytics challenges of large-scale high-fidelity numerical simulations of wind energy applications," *AIAA Paper 2018-1171*, 2018.

[79] Turton, T. L., Banesh, D., Overmyer, T., Sims, B. H., and Rogers, D. H., "Enabling Domain Expertise in Scientific Visualization With CinemaScience," *IEEE Computer Graphics and Applications*, Vol. 40, 2020, pp. 90–98.

[80] Imlay, S. T., Mackey, C. A., and Taflin, D. E., "Visualizing a Trillion-Cell Simulated CFD Solution on an Engineering Workstation," *AIAA Paper 2018-xxx*, 2018.

[81] Cielo, S., Iapichino, L., Gunther, J., Federrath, C., Mayer, E., and Wiedemann, M., "Visualizing the world's largest turbulence simulation," *arXiv preprint arXiv:1910.07850.*, 2019.

[82] Kiviaho, J. F., and Kennedy, G. J., "Efficient and Robust Load and Displacement Transfer Scheme Using Weighted Least Squares," *AIAA Journal*, Vol. 57, No. 5, 2019.

[83] Mader, C. A., Kenway, G. K. W., Yildrim, A., and Martins, J. R. R. A., "ADflow: An Open-Source Computational Fluid Dynamics Solver for Aerodynamic and Multidisciplinary Optimization," *Journal of Aerospace Information Systems*, Vol. 17, No. 9, 2020.

[84] Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A., "OpenMDAO: an open-source framework for multidisciplinary design, analysis, and optimization," *Structural and Multidisciplinary Optimization*, Vol. 59, 2019, pp. 1075–1104.

[85] `https://www.3ds.com/products-services/simulia/disciplines/multiphysics-simulation/`, 2020.

[86] `https://www.ansys.com/products/platform/multiphysics-simulation/fluid-structure-interaction`, 2020.

[87] `https://www.comsol.com/products`, 2020.

[88] `https://www.phoenix-int.com/`, 2020.

[89] `https://www.esteco.com/modefrontier`, 2020.

[90] `http://www.vrand.com/products/visualdoc/`, 2020.

[91] `https://www.3ds.com/products-services/simulia/products/isight-simulia-execution-engine/latest-release/`, 2020.

[92] Adams, B. M., Bohnhoff, W. J., Dalby, K. R., Ebeida, M. S., Eddy, J. P., Eldred, M. S., Hooper, R. W., HOUGH, P. D., Hu, K. T., Jakeman, J. D., Kahill, M., Maupin, K. A., Monschke, J. A., Ridgway, E. M., Rushdi, A. A., Seidl, D. T., Stephens, J. A., SSwiler, L., and Winokur, J. G., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.12 User's Manual," Tech. rep., Sandia National Laboratory, SAND2020-5001, 2020.

[93] Gallard, F., Vanaret, C., Gu'enot, D., Gachelin, V., Lafage, R., Pauwels, B., Barjhoux, P.-J., and Gazaix, A., "GEMS: A Python Library for Automation of Multidisciplinary Design Optimization Process Generation," *AIAA Paper 2019-2991*, 2019.

[94] Kolonay, R. M., "MSTC Engineering – A Distributed and Adaptive Collaborative Design Computational Environment for Military Aerospace Vehicle Development and Technology Assessment," *AIAA Paper 2019-2992*, 2019.

[95] Daniel L. Clark, J., "Non-Deterministic Metamodeling for Multidisciplinary Design Optimization of Aircraft Systems Under Uncertainty," Ph.D. thesis, Wright State University, 2019.

[96] Vandenbrande, J., "Enabling Quantification of Uncertainty in Physical Systems (EQUiPS)," `https://www.darpa.mil/program/equips`, 2020.

[97] Chaudahuri, A., Marques, A. N., and Willcox, K. E., "mfEGRA: Multifidelity Efficient Global Reliability Analysis," *arXiv preprint*, 2019.

[98] Chaudahuri, A., Marques, A. N., Lam, R., and Willcox, K. E., "Reusing Information for Multifidelity Active Learning in Reliability-Based Design Optimization," *AIAA Paper 2019-1222*, 2019.

[99] Chaudahuri, A., Kramer, B., and Willcox, K. E., "Information Reuse for Importance Sampling in Reliability-Based Design Optimization," *Reliability Engineering & System Safety*, 2020.

[100] Mangano, M., and Martins, J. R. R. A., "Multipoint Aerodynamic Shape Optimization for Subsonic and Supersonic Regimes," *AIAA Paper 2019-xxxx*, 2019.

[101] Munu, B., Economon, T. D., and Alonso, J. J., "Discrete Adjoint Framework for Low-Boom Supersonic Aircraft Shape Optimization," *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2017.

[102] Rallabhandi, S. K., Nielsen, E. J., and Diskin, B., "Sonic-Boom Mitigation Through Aircraft Design and Adjoint Methodology," *Journal of Aircraft*, 2014.

[103] Kenway, G., and Martins, J. R. R. A., "Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations," *Journal of Aircraft*, Vol. 52, No. 5, 2015, pp. 935–951.

[104] Jacobson, K., "Adjoint-based Aeroelastic Optimization with High-Fidelity Time-Accurate Analysis," Ph.D. thesis, Georgia Institute of Technology, 2018.

[105] Ghazlane, I., Carrier, G., Dumont, A., Marcelet, M., and Desideri, J., "Aerostructural Optimization with the Adjoint Method," *Eurogen Conference*, Capua, Italy, 2011.

[106] Gray, J., Mader, C., Kenway, G., and Martins, J. R. R. A., "Coupled Aeropropulsive Optimization of a Three-Dimensional Boundary-Layer Ingestion Propulsor Considering Inlet Distortion," *Journal of Aircraft*, Vol. 57, No. 6, 2020.

[107] Chauhan, S. S., and Martins, J. R. R. A., "RANS-Based Aerodynamic Shape Optimization of a Wing Considering Propeller–Wing Interaction," *AIAA Paper 2020-1764*, 2020.

[108] Mishra, A., Mani, K., Mavriplis, D., and Sitaraman, J., "Time Dependent Adjoint-Based Optimization for Coupled Fluid-Structure Problems," *Journal of Computational Physics*, Vol. 292, No. 1, 2015, pp. 253–271.

[109] Wang, L., Diskin, B., Biedron, R., and Bauchau, O., "Evaluation of High-Fidelity Multidisciplinary Sensitivity Analysis Framework for Multipoint Rotorcraft Optimization," *Journal of Aircraft*, Vol. 57, No. 5, 2020, pp. 830–842.

[110] Deaton, J. D., Gallman, J., and Alonso, J. J., "Introduction to AFRL-NG-SU QUAD: Quantifying the Utility of Coupled Sensitivity Analysis in Industry Design Processes," *AIAA Paper 2020-xxxx*, 2020.

[111] Chakraorty, I., and Mavris, D. N., "Heuristic Definition, Evaluation, and Impact Decomposition of Aircraft Subsystem Architectures," *AIAA Paper 2020-xxxx*, 2016.

[112] Chakraorty, I., and Mavris, D. N., "Integrated Assessment of Aircraft and Novel Subsystem Architectures in Early Design," *Journal of Aircraft*, Vol. 54, No. 4, 2017.

[113] Carrere, A., Sellers, P. S., and Allison, D. L., "Building an Expanded MDAO Process with Subsystem Analysis for Next Generation Air Dominance Platforms," *AIAA Paper 2019-xxxx*, 2019.

[114] Sellers, P. S., Carrere, A., and Allison, D. L., "Results, Conclusions and Lessons Learned from the OPTimized Integrated MUltidisciplinary Systems (OPTIMUS) Program," *AIAA Paper 2019-xxxx*, 2019.

[115] Allison, D. L., and Kolonay, R., "Expanded MDO for Effectiveness Based Design Technologies: EXPEDITE Program Introduction," *AIAA Paper 2018-3419*, 2018.

[116] Davies, C., "Lockheed Martin Overview of the AFRL EXPEDITE Program," *AIAA Paper 2019-0173*, 2019.

[117] Harper, D. J., "Operations Analysis Integration for Effectiveness-Based Design in the AFRL EXPEDITE Program," *AIAA Paper 2020-1130*, 2020.

[118] Dannenhoffer, J. F., and Haimes, R., "Generation of Multi-fidelity, Multi-discipline Air Vehicle Models with the Engineering Sketch Pad," *AIAA Paper 2016-1925*, 2016.

[119] Haimes, R., Dannenhoffer, J., Bhagat, N. D., and Allison, D. L., "Multi-fidelity Geometry-centric Multi-disciplinary Analysis for Design," *AIAA Paper 2016-4007*, 2016.

[120] Dannenhoffer, J. F., and Haimes, R., "Using Design-Parameter Sensitivities in Adjoint-Based Design Environments," *AIAA Paper 2017-0139*, 2017.

[121] Hendricks, E. S., and Gray, J. S., "pyCycle: A Tool for Efficient Optimization of Gas Turbine Engine Cycles," *Aerospace*, Vol. 6, No. 8, 2019, p. 87.

[122] Hearn, T. A., Hendrcks, E., Chin, J., and Gray, J. S., "Optimization of Turbine Engine Cycle Analysis with Analytic Derivatives," *AIAA Paper 2016-4297*, 2016.

[123] Martins, J. R. R. A., and Hwang, J. T., "Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models," *AIAA Journal*, Vol. 51, No. 11, 2013.

[124] Hwang, J. T., and Martins, J. R. R. A., "A Computational Architecture for Coupling Heterogeneous Numerical Models and Computing Coupled Derivatives," *ACM Trans. Math. Softw.*, Vol. 44, No. 4 Article 37, 2018.

[125] Mani, K., and Mavriplis, D., "Geometry Optimization in Three-Dimensional Unsteady Flow Problems using the Discrete Adjoint," *AIAA Paper 2013-0662*, 2013.

[126] Mani, K., and Mavriplis, D., "A Discrete Adjoint Framework for Unsteady Aerodynamic and Aeroacoustic Optimization," *AIAA Paper 2015-3355*, 2015.

[127] Prasad, R., and Choi, S., "Aerodynamic Shape Optimization for Flutter/LCO based design using Coupled Adjoint," *AIAA Paper 2020-0405*, 2020.

[128] Makhija, D. S., and Beran, P. S., "Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models," *Structural and Multidisciplinary Optimization*, Vol. in review, 2018.

[129] Makhija, D. S., and Beran, P. S., "Time Scale Effects in Topology Optimization of the Interior Material Distribution of a Body Subject to Transient Conjugate Heat Transfer," *AIAA paper 2019-1468*, 2019.

[130] "CFD Vision 2030 Report Updates," `http://www.cfd2030.com/report`, 2021.