EDITORS: Barry I. Schneider, bis@nist.gov
Gabriel A. Wainer, gwainer@sce.carleton.ca

DEPARTMENT: COMPUTER SIMULATIONS

Realizing the Vision of CFD in 2030

Andrew Cary, *Boeing Research and Technology, St. Louis, MO, 63042, USA*
John Chawner, *Cadence Design Systems, Fort Worth, TX, USA*
Earl Duque, *Intelligent Light, Rutherford, NJ, USA*
William Gropp, *University of Illinois in Urbana-Champaign, Urbana, IL, USA*
Bil Kleb, *Langley Research Center, NASA, Hampton, VA, USA*
Ray Kolonay, *Air Force Research Laboratory, Wright-Patterson Air Force Base, OH, USA*
Eric Nielsen, *Langley Research Center, NASA, Hampton, VA, USA*
Brian Smith, *Lockheed Martin, Dallas-Fort Worth, TX, USA*

In 2014, NASA released a report outlining a future state for aerospace computational fluid dynamics (CFD), the CFD Vision 2030 Study (the Study).[1] Developed by experts from industry, government, and academia, the Study provided a forecast of CFD capabilities required for turbulent, transitional, and reacting flow simulations across a broad Mach number regime. In addition, the Study provided an aspirational role for future CFD as part of a routine, efficient, and physics-based aerospace design and development process. This future role of CFD was summarized in the Study as follows: "A single engineer/scientist must be able to conceive, create, analyze, and interpret a large ensemble of related simulations in a time-critical period (e.g., 24 hours), without individually managing each simulation, to a prespecified level of accuracy."

The Study assessed industrial requirements, identified impediments, and formulated Grand Challenge (GC) problems to highlight capabilities needed to achieve the Vision as well as a development plan (the Roadmap) to achieve these goals. The Study became a focus for NASA CFD efforts, and other organizations and universities used it to identify development needs. Although the Study predated much of the current emphasis on digital engineering and model-based engineering, the Vision is highly aligned with many of the objectives that these focal areas seek to achieve.

While the Roadmap defined a series of technical milestones and demonstrations needed to advance the aerospace community toward the Vision, an advocate was required to maintain momentum. Original sponsors and authors of the Vision, along with other industry leaders, formed the CFD Vision 2030 Integration Committee (the IC) in 2018 within the American Institute of Aeronautics and Astronautics (AIAA) to foster the Vision and promote accomplishments toward its realization. This effort has led to the establishment of a website,[a] a five-year anniversary special session at AIAA Aviation 2019 with top experts from across CFD technologies, and the recent refinement of three GC problems in commercial air transport, full engine simulation, and space vehicle design. Another key outcome of the IC in 2021 has been a careful review of the advancement of technologies identified on the Roadmap since its release[2] and a corresponding update to the Roadmap to reflect current status and trends.[3]

THE ROADMAP TO CFD IN 2030

The updated Roadmap (see Figure 1) has a timeline divided into six domains: high-performance computing, physical modeling, algorithms, geometry modeling and mesh generation, knowledge extraction, and multidisciplinary analysis and optimization. Each of these domains has one or more key technology areas identified with a colored bar representing the technology readiness level evolution. Critical technology milestones and demonstrations are distributed along

1521-9615 © 2022 IEEE
Digital Object Identifier 10.1109/MCSE.2021.3133677
Date of current version 14 March 2022.

[a] www.cfd2030.com

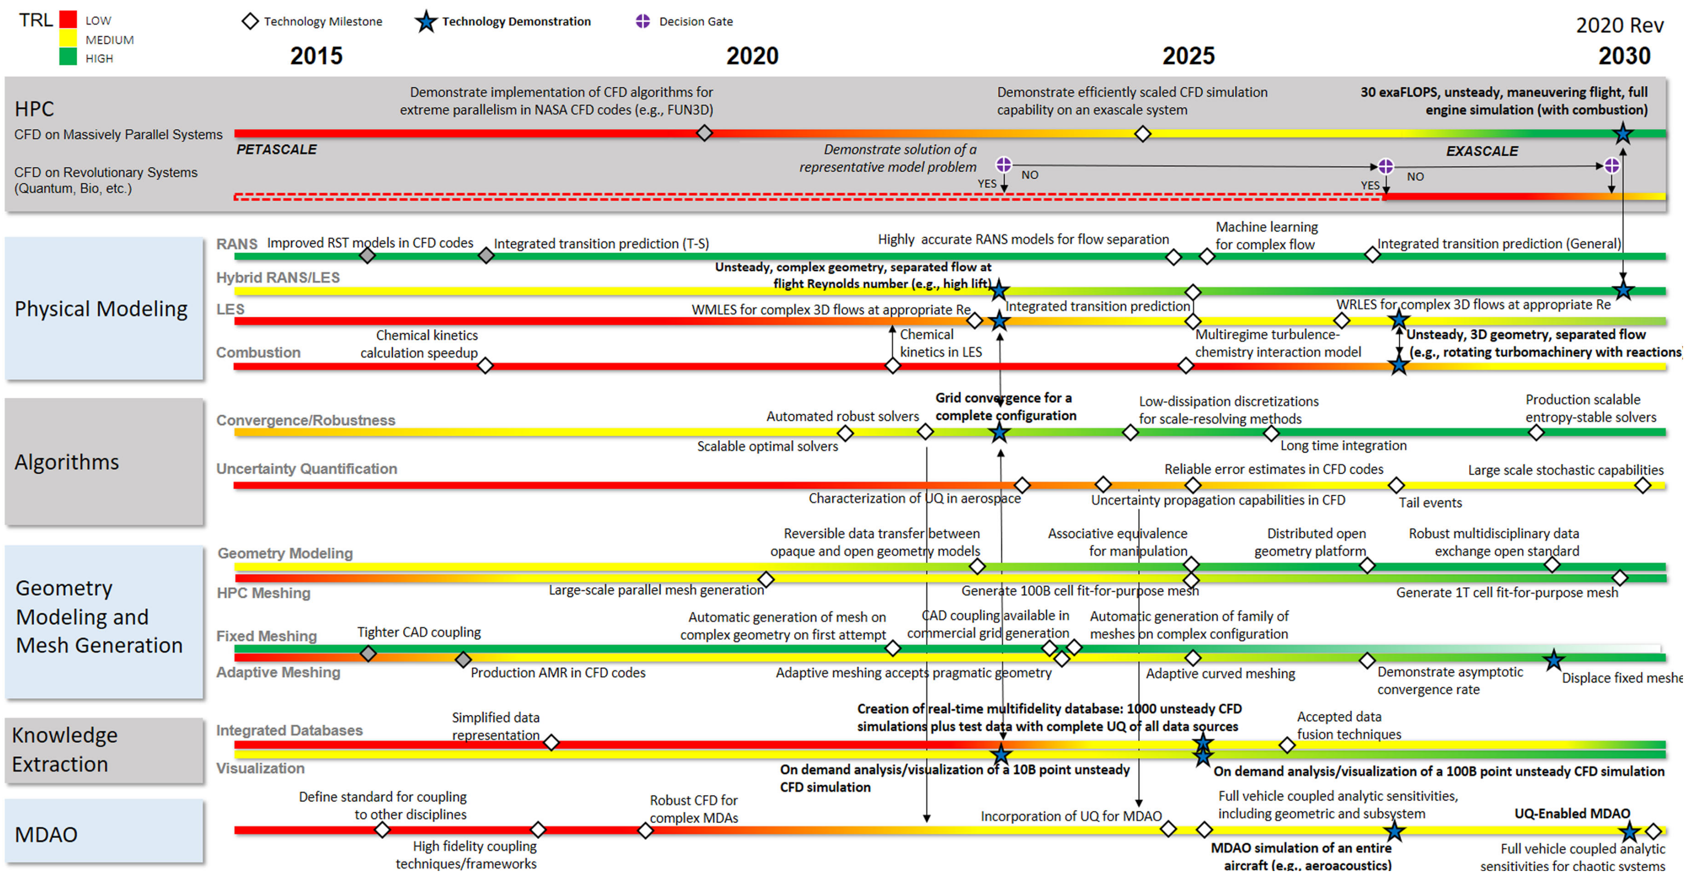

FIGURE 1. CFD Vision 2030 Roadmap, updated 2020.

these bars. The 2020 Roadmap update included adjusting the dates for these milestones to reflect accomplishments as well as addition of items to reflect learning that has taken place since 2014.

There are three related technology demonstrations that serve as intermediate indications of future capabilities slated for 2023:

1) unsteady, complex geometry, separated flow at high Reynolds number;
2) grid convergence for a complete configuration;
3) on-demand analysis/visualization of an unsteady CFD simulation with on the order of 10 billion mesh points.

There have been several recent key indicators of progress on these topics. The AIAA CFD High Lift Prediction Workshop is highly aligned with the first of these. At the third workshop held in 2017, two participants performed time-dependent, scale-resolving simulations that provided encouraging results compared to Reynolds-averaged simulations. The fourth workshop (scheduled for 2022) includes over ten teams contributing scale-resolving solutions and generating preliminary results in better agreement with experiment than previous efforts. Less progress has been made on the grid convergence demonstration. In 2018, several teams were able to demonstrate grid convergence on an ONERA M6 wing using mesh adaptation and strong finite element solvers. Increasing the wing's angle of attack led to challenges with multiple solutions being obtained, which creates issues for achieving grid-converged results. This observation has resulted in additional effort being focused on controlling which solution is attained. At SC19, NVIDIA, NASA, and ORNL provided an interactive visualization of a lander entering the Mars atmosphere using six billion elements and leveraging state-of-the-art resources including computation of the results on Summit, the world's most powerful supercomputer at that time, and storage of 150 TB with visualization driven by four DGX-2™ servers. While each of these are significant accomplishments, they fall short of the intended goals on the Roadmap's timeline.

HIGH-PERFORMANCE COMPUTING (HPC)

The HPC domain made an accurate assessment of the increases in computing capacity since 2014 and identified some of the challenges associated with the transition to heterogeneous computing platforms such as GPUs. On the journey to exascale computing, there are several potential disruptive technologies that are identified, but these have yet to mature to a point for significant industrial CFD usage. The key milestones for the Vision identified in this domain include demonstration of algorithms for extreme parallelism in NASA CFD codes, demonstration of efficient scaling of CFD on exascale systems, and demonstration of a challenge problem consisting of a 30 exaFLOP simulation of a maneuvering flight with a full engine simulation. While we are still several years away from the latter two, NASA's FUN3D CFD code has been demonstrated to scale effectively on large systems such as Summit, running on over 16,000 GPUs with an equivalent performance of several million Xeon cores.

Reaching for the improved performance of exascale computing is necessitating a paradigm shift from the traditional computing advances. In the recent past, increases in clock speed and memory performance, addition of cores, and vectorization support have facilitated improvement in CFD without large-scale application changes using node-level parallelism on the order of a few hundred processes. Manufacturing constraints and power requirements have forced computer vendors to seek continued improvements through vastly higher levels of parallelism and developing more complicated memory hierarchies. These changes are necessitating a shift in not only the programming of CFD software, but also leading to CFD algorithms that are tailored to specific hardware classes. These changes stem from both the typical memory-bound nature of traditional CFD routines and the increased level of parallelism that must be exploited. There has also been progress on performing CFD on a range of revolutionary hardware systems including quantum computing, but these efforts are presently immature for industrial CFD applications.

PHYSICAL MODELING

The physical modeling domain of the Roadmap includes both turbulence model and chemical reaction enhancements to support increased prediction accuracy. One timeline focuses on Reynolds-averaged Navier–Stokes (RANS) modeling. RANS approaches have helped industry to provide much of the CFD success in modeling over the last 25 years, but have limitations as simulation objectives extend past cruise conditions (where the boundary layers tend to be attached) toward more separated flows at more extreme conditions. The dominant turbulent models for industrial applications were introduced in the early 1990s and have received several extensions but have not been replaced. At the time of the original Study, there was considerable hope that Reynolds Stress Transport models could provide much of the additional accuracy, but these have not yet gained

FURTHER READING:

- AIAA CFD Vision 2030 website (www.cfd2030.com) contains information on the original report, progress updates, activities, and GC problems.
- 2019 AIAA Aviation Conference Special Session
  - G. Medic, "Impact of Vision 2030 on CFD practices in propulsion industry," AIAA Aviation 2019 Forum, 2019, AIAA-2019-2943.
  - D. Mavriplis, "Progress in CFD discretization, algorithms, and solvers for aerodynamic flows," AIAA Aviation 2019 Forum, 2019, AIAA-2019-2944.
  - J. Chawner and N. Taylor, "Progress in geometry modeling and mesh generation toward the CFD Vision 2030," AIAA Aviation 2019 Forum, 2019, AIAA-2019-2945.
- CFD Grand Challenge Problems
  - J. Slotnick and D. Mavriplis, "A grand challenge for the advancement of numerical prediction of high lift aerodynamics," AIAA Scitech 2021 Forum, 2021, AIAA-2021-0955.
  - M. Anand, G. Laskowski, K. Suder, G. Medic, U. Paliath, and M. Malik, "Vision 2030 aircraft propulsion grand challenge problem: Full-engine CFD simulations with geometric fidelity and physics accuracy," AIAA Scitech 2021 Forum, 2021, AIAA-2021-0956.
  - D. Schuster, "CFD 2030 Grand Challenge: CFD-in-the-loop Monte Carlo flight simulation for space vehicle design," AIAA Scitech 2021 Forum, 2021, AIAA-2021-0957.

general acceptance. Instead, the scale-resolving approaches identified in the next two timelines (hybrid RANS/LES and LES) have become more prevalent. These approaches represent the dominant unsteady turbulent scales to improve the predictive accuracy; they differ primarily in the representation of turbulent flows near the wall. These methods have seen faster growth than originally predicted. Considerable questions need to be resolved for both hybrid RANS/LES and LES methods to realize their potential. The final timeline in this domain focuses on combustion modeling and reliably representing detailed chemical reactions, particularly in time-dependent flows. Key challenges in this area include increasing the computational efficiency for predicting complex chemical kinetics and improving the modeling of chemical kinetics in turbulent flows. This capability builds toward modeling time-dependent separated flows such as rotating turbomachinery including chemical reactions.

Additional key challenges identified in the physical modeling domain include better modeling of the transition to turbulence as well as improvements to models through machine learning. Turbulence transition is important for a wide range of applications from improving drag through increased laminar flow to hypersonic heat transfer predictions. Transition methods typically used in industrial applications are based on stability theory or, more recently, on empirical corrections included as a partial differential equation. These have focused on modeling transition based on Tollmien–Schlichting waves, but there are additional important transition mechanisms that need to be included. Some of the scale-resolving methods described earlier suffer from inconsistent generation of turbulence that stems from numerically induced transition. Machine learning for turbulence is an example of a CFD technology that did not appear on the original roadmap because of its infancy at the time. There are now numerous efforts to improve turbulence modeling through the use of machine learning to adjust turbulence models to better represent flow behavior.

ALGORITHMS

The algorithms domain recognizes that there are enhancements in numerical methods that are necessary to continue to make efficient use of new HPC hardware and to support the simulation requirements of industrial CFD applications. For example, scale-resolving simulations require low-dissipation numerical algorithms and the ability to maintain accuracy for long-time integrations. On top of these new challenges, there are ongoing issues associated with achieving grid-converged solutions as traditional CFD numerical methods have challenges reaching convergence. This has been improved upon in some cases, but has further highlighted the challenges of multiple solutions mentioned previously. Developing appropriate solution methods to ensure that iteration error is sufficiently small is an additional challenge. Both high-order methods and entropy-stable solvers are seen as options for reaching some of these goals, but there remains considerable development necessary to mature the scalability and robustness of these technologies for industrial applications.

A second timeline in the algorithms domain is uncertainty quantification. This discipline has been gaining increased awareness in its applications to CFD because of the increased reliance on CFD for emerging industrial analysis and design applications. Previously, the near-linear behavior for many applications did not require a detailed determination of the solution confidence as the solutions are heavily supported by experimental evidence. To reach the desired accelerated timelines for future aerospace applications, testing will need to be used more for validation of CFD results than performance estimation as used presently. This necessitates better characterization of the CFD results and understanding of the numerical, modeling, and condition uncertainty. Techniques for doing this are being developed, but reliable error estimation for CFD tools does not presently exist for a broad range of applications.

TO REACH THE DESIRED ACCELERATED TIMELINES FOR FUTURE AEROSPACE APPLICATIONS, TESTING WILL NEED TO BE USED MORE FOR VALIDATION OF CFD RESULTS THAN PERFORMANCE ESTIMATION AS USED PRESENTLY.

GEOMETRY MODELING AND MESH GENERATION

The Roadmap update expanded the elements associated with geometry modeling and mesh generation by the incorporation of timelines for both geometry modeling and HPC-based meshing methods. These new revisions are driven by increased recognition that significantly more automation will be necessary to provide the discretization needed for future engineering simulations. The techniques to robustly and accurately manipulate geometry representations and efficiently generate large (billions of cells) meshes require considerable development and coordination across a range of different organizations. Incorporation of multidisciplinary analysis and association to geometry will continue to be required for increased simulation complexity.

Fixed-grid meshing and adaptive mesh refinement have made progress but have not reached the expectations suggested by the original timelines, partially due to the increased simulation emphasis on time-dependent simulations. First-time quality of automated grid generation has improved significantly over the last decade, but user intervention is presently required for a number of key scenarios. Similarly, adaptive meshing has demonstrated increased capability for a range of flow domains and flow solvers, but is not yet an industry standard, despite having multiple options available to accomplish this. A significant portion of the challenges faced by mesh generation involve the use of geometry models from industrial computer-aided design applications that all too often include nonphysical gaps, overlaps, and other geometry problems.

KNOWLEDGE EXTRACTION

The knowledge extraction domain stems from the understanding that the increased level of simulation will necessitate better methods for gaining understanding, not only from a single simulation but also from complete simulation databases. The level of manual inspection of each simulation performed in the past is not feasible at present CFD throughput levels let alone the future levels. This domain focuses on extracting the key results from simulations, whether it be visualization of a single solution or understanding the fluid-dynamic sources that impact engineering quantities of interest across an entire database. Engineering applications require thousands of simulations to represent aerospace problems, and it is likely that multifidelity data fusion techniques will be required to represent different levels of approximation and test data. These techniques are important to achieve a digital thread design concept, but many present approaches cannot consume complete CFD solutions.

MULTIDISCIPLINARY ANALYSIS AND OPTIMIZATION

Because aerospace products are increasingly integrated across disciplines, a key element of the Vision is extending CFD analysis to interface with additional disciplines including structures, control, aeroacoustics, and others. To facilitate this, it is expected that standards need to be defined to enable both analysis and the determination of sensitivities across the interfaces of these domains. In some cases, such as structures, it might be straightforward to envision a cosimulation at the highest levels of fidelity, but the disparate time scales and solution costs for different domains will introduce challenges. Coupling of other disciplines is less direct at the high-fidelity level necessary for an integrated domain, but other approaches may suffice if sufficient sensitivity information can be obtained.

Combining this with the very large number of design points (hundreds for an aircraft loads survey), design variables (potentially tens of thousands), and design constraints (up to hundreds of thousands) highlights how daunting of a task this is.

While there are not yet accepted standards for coupling disciplines, approaches are emerging and some systems have demonstrated initial abilities for performing integrated multidisciplinary optimization. Other approaches use a scripting language to interface multiple analysis tools together, but this often limits the determination of fully coupled sensitivities.

THE PATH FORWARD

The CFD Vision 2030 is a comprehensive look at computational simulation technologies with a practical focus on overcoming today's challenges in order to deliver the necessary tools for achieving the digital engineering capabilities required for advanced aerospace applications in less than a decade. The challenges faced are not easily overcome. Several challenges are in domains that are largely technological: effective use of HPC platforms to drastically reduce turn-around time, numerical algorithms that exploit HPC architectures while also delivering the desired accuracy, and improved physical modeling of turbulence, separation, and combustion that will expand CFD's applicability near the boundaries of the flight envelope. Other challenges are in domains that are largely process-oriented: practical automation of mesh generation that accounts for both the complexity of the geometry model and the simulation needs of the CFD solver, tools for engineering knowledge extraction that provide insight from hundreds or thousands of simulations, and ecosystems of software that are truly interoperative and enable multidisciplinary simulations.

The CFD Vision 2030 Integration Committee is continuing to champion the objectives defined in the 2014 Study to support a stronger digital engineering capability in CFD. This effort includes not only monitoring technical progress for identification of new approaches and technology maturation, but also advocacy for technology risks associated with slowed development. Through both annual reports on the state of the Roadmap and industry-wide advocacy for specific GC problems representing enduring challenges, progress across the industry can be achieved. While this effort focuses on CFD, it parallels and complements industry-wide trends in digital engineering.

The next set of Technology Demonstrations for 2023 include achieving accurate results for unsteady simulations on complex geometries, demonstration of grid convergence on a full configuration, and interactive visualization of a 10 billion mesh point unsteady simulation. None of these will be easy to achieve, but solid progress has been made across all Roadmap domains in the last five years, which will support reaching these goals.

REFERENCES

1. J. Slotnick *et al.*, "CFD vision 2030 study: A path to revolutionary computational aerosciences," NASA, Washington, DC, USA, Tech. Rep. NASA/CR-2014-218178, 2014.
2. A. W. Cary *et al.*, "The CFD vision 2030 roadmap: 2020 status, progress, and challenges," 2021. [Online]. Available: http://www.cfd2030.com/report
3. A. Cary *et al.*, "CFD vision 2030 road map: Progress and perspectives," AIAA, Reston, VA, USA, Tech. Rep. AIAA 2021-2726, 2021.

ANDREW CARY is currently a Technical Fellow with the Aerosciences Team, Boeing Research and Technology, St. Louis, MO, USA. Cary received the Ph.D. degree in aerospace engineering from the University of Michigan, Ann Arbor, MI, USA. Contact him at Andrew.W.Cary@boeing.com.

JOHN CHAWNER was the Co-Founder and President of Pointwise (CFD mesh generation), which was acquired by Cadence Design Systems, San Jose, CA, USA, in early 2021 where he is a Senior Group Director R&D. Chawner received the B.S. degree from Syracuse University, Syracuse, NY, USA, and the M.S. degree in aerospace engineering from the University of Texas at Arlington, Arlington, TX, USA. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics. Contact him at jchawner@cadence.com.

EARL DUQUE is currently the Vice-President of research and development at Intelligent Light, Rutherford, NJ, USA. His current research interests include CFD visualization and computational environments for uncertainty quantification and data analytics. Duque received the Ph.D. degree in mechanical engineering from the University of California, Davis, CA, USA. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics. Contact him at epd@ilight.com.

WILLIAM GROPP is the Director of the National Center for Supercomputing Applications and holds the Thomas M. Siebel Chair, Department of Computer Science, University of Illinois in Urbana-Champaign, Champaign, IL, USA. He was with Yale University and Argonne National Laboratory. His research interests include parallel computing, software for scientific computing, and numerical methods for partial

differential equations. Gropp received the Ph.D. degree in computer science from Stanford University, Stanford, CA, USA. He is a Fellow of AAAS, ACM, IEEE, and SIAM, and a Member of the National Academy of Engineering. Contact him at wgropp@illinois.edu.

BIL KLEB is currently the Senior Technologist for computational aerothermodynamics at NASA Langley Research Center, Hampton, VA, USA. Kleb received the B.S. and M.S. degrees from Purdue University, West Lafayette, IN, USA, the MBA degree from The College of William & Mary, Williamsburg, VA, USA, and the Ph.D. degree in aerospace engineering from the University of Michigan, Ann Arbor, MI, USA. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics. Contact him at Bil.Kleb@nasa.gov.

RAY KOLONAY is currently the Director of the Multidisciplinary Science and Technology Center, Aerospace Systems Directorate, Air Force Research Laboratory, Wright-Patterson AFB, OH, USA. Kolonay received the Ph.D. degree in aerospace engineering from Purdue University, West Lafayette, IN, USA. He is an ASME Fellow, AFRL Fellow, and an Associate Fellow of the American Institute of Aeronautics and Astronautics. Contact him at Raymond.Kolonay@us.af.mil.

ERIC NIELSEN is currently a Senior Research Scientist with the Computational AeroSciences Branch, NASA Langley Research Center, Hampton, VA, USA. Nielson received the Ph.D. degree in aerospace engineering from Virginia Tech, Blacksburg, VA. Contact him at Eric.J.Nielsen@nasa.gov.

BRIAN SMITH has been with Lockheed Martin, Dallas-Fort Worth, TX, USA, and its predecessor companies, since 1986, where he is currently a Lockheed Martin Senior Fellow. His research interests include computational fluid dynamics, turbulence modeling, and aero-optics. Smith received the Ph.D. degree in mechanical and aerospace engineering from Princeton University, Princeton, NJ, USA. He is a Fellow of the American Institute of Aeronautics and Astronautics. Contact him at Brian.R.Smith@lmco.com.