Fall 2016

DOD HPC INSIGHTS

A publication of the Department of Defense High Performance Computing Modernization Program

High-Fidelity Simulations of Primary Breakup and Impinging Jet Atomization

HPCMP "Hero" Awards

Highlights of the HPC Internship Program (HIP)

Contents

Cover & Inside Cover Photos:
Turbulent Fuel Jet Atomization in Army UAV Engine
courtesy of:
Luis Bravo, Army Research Laboratory, DoD HPC Data Analysis & Assessment Center, and ARL Vehicle Technology Directorate

First word

Dr. David Horner
Director, High Performance Computing Modernization Program

Welcome back to HPC Insights. There is significant work being done throughout the HPCMP program that directly impacts the Department of Defense and the warfighters that we support. The Program provides the ecosystem of high-performance computers, fast and reliable networks, and advanced software that powers efforts in science and technology, test and evaluation, and acquisition engineering necessary to maintain the US military's advantage over potential adversaries.

In this issue, we are showcasing work on high-fidelity simulations of jet atomization – an important capability that will assist the Army in their mission by providing reliable and efficient propulsion for direct injection fuel delivery systems to power military vehicles more efficiently and effectively. Similarly, the work on jet engine combustors will impact aircraft in all the services.

I congratulate our Hero Award winners. This year the program received a record high 49 nominations, and 13 winners. We added a team award, which was won by a group from the US Naval Academy for establishing a high-performance computing research center and recruiting faculty to use these resources. In addition to recognizing the lifetime achievements of distinguished program alumni Dr. Cliff Rhoades and Mr. Brad Comes, we are excited to recognize up and coming talent with four awards.

Again, with an eye toward the future, our HPC Internship Program (HIP) has been a great success. There is a great article showcasing some of our HIP project results where the interns have been approached to pursue a career in HPC by their mentor's organization or other DoD entities. There are some great young minds out there, and we've caught a few of them in our HIP program!

The coming year will see the installation of new HPC machines at the DSRCs in ERDC and ARL as part of the 2016 Technology Insertion (TI) and the requirements gathering for TI-17. The arrival of new computing architectures offers new capabilities, challenges, and opportunities. The HPCMP will continue to be a key resource in solving challenging problems for the Department of Defense.

DoD HPC Insights

High-Fidelity Simulations of Primary Breakup and Impinging Jet Atomization

Figure 1: *Army Gray Eagle MQ-1C, powered by turbo-charged diesel engine (160hp) on mission.*
(Photo courtesy of the US Army Acquisition Support Center)

Dr. Luis Bravo
Army Research Laboratory, Vehicle Technology Directorate

Background and Objective

To support the Warfighter, the Army needs to provide reliable and efficient propulsion for heavy fuel engine platforms that exclusively rely on direct injection fuel delivery systems. Combat vehicles, such as the Gray Eagle MQ-1C (Figure 1) and the Joint Light Tactical Vehicle (JLTV), are powered by diesel engines running on military JP-8, or F-24 fuels. Technology breakthroughs in engine and fuel conversion efficiencies require a fundamental understanding of key phenomena, including fuel/air mixture formation due to primary breakup and impinging jet atomization phenomena.

In diesel sprays, the process begins with ejection of an intact stream of liquid fuel from a nozzle into a high-pressure high-temperature charge of oxidizer (air). The stream experiences surface instabilities due to hydrodynamic effects leading to disintegration of the liquid core via primary breakup mechanics. The process is driven by shear, turbulence, and cavitation, and ultimately produces primary drops. Primary drops can break-up further, or collide and coalesce further downstream into the dilute region. In the dilute region, some of the fuel droplets may impact and interact with the cylinder piston surface. Interactions of fuel drops and the piston surface can result in various outcomes, including formation of wall film, drop rebound, drop breakup, and drop boiling on the surface. These outcomes are determined by the impingement regime that is characterized by the drop properties, flow conditions, and surface temperatures. The optimal control of resulting primary and secondary droplet distributions outcomes will impact evaporation and mixing rates; and is therefore a key to the quality of the combustion process.

Despite the relevance of the atomization process, its description is still among the weakest parts of practical engineering simulation models. The most common approach is to avoid a detailed description in favor of a semi-empirical model describing the sudden appearance of large droplets with specific momentum that then break up into finer droplets and vaporize. Such models rely upon experimental data to set adjustable model parameters. Their limitations are due to the lack of information and restricted experimental access in the dense and near-wall regions [1]. Although experimentalists have had success with modern methods such as ballistic imaging and x-ray techniques; extraction of the full four-dimensional information with sufficient spatial and temporal resolution for a detailed analysis is still unfeasible [2]. As a result, a comprehensive theory of turbulent atomization has remained elusive. Fully predictive modeling is, thus, not possible at this time with these approaches. However, with the recent advances in supercomputing power and numerical algorithms, first-principle simulations of the atomization processes are emerging today as a viable research tool to study and predict fuel/air mixture formation. See Figure 2 for an example showing in-cylinder atomization process in an Army engine.

At the Army Research Laboratory, a multi-institutional effort towards the prediction of multi-physics sprays is underway. It is a 5-year Frontier project that inaugurated at ARL in FY15, and is supported by the High Performance Computing Modernization Program (HPCMP). The focus of this study is to conduct scientific research to increase the knowledge and understanding through first-principle continuum simulations of atomization for two overarching scenarios including: primary breakup, and spray impingement atomization. The project addresses issues relating to the non-reacting behavior of atomizing sprays and impinging jets, including the role of perturbation-driven instabilities on breakup and droplet formation, complex evaporation, supercritical effects on fuel/vapor surface, and drop-wall interactions. The major impact of this effort is the creation of a suite of breakthrough computational tools with the ability to predict the microscale flow physics of atomizing and impinging jets using fundamental principles. The Army can apply these predictive models to the performance of any chemical-propulsion device that uses spray combustion, with the vetted models being particularly helpful in reducing the experimental steps necessary.

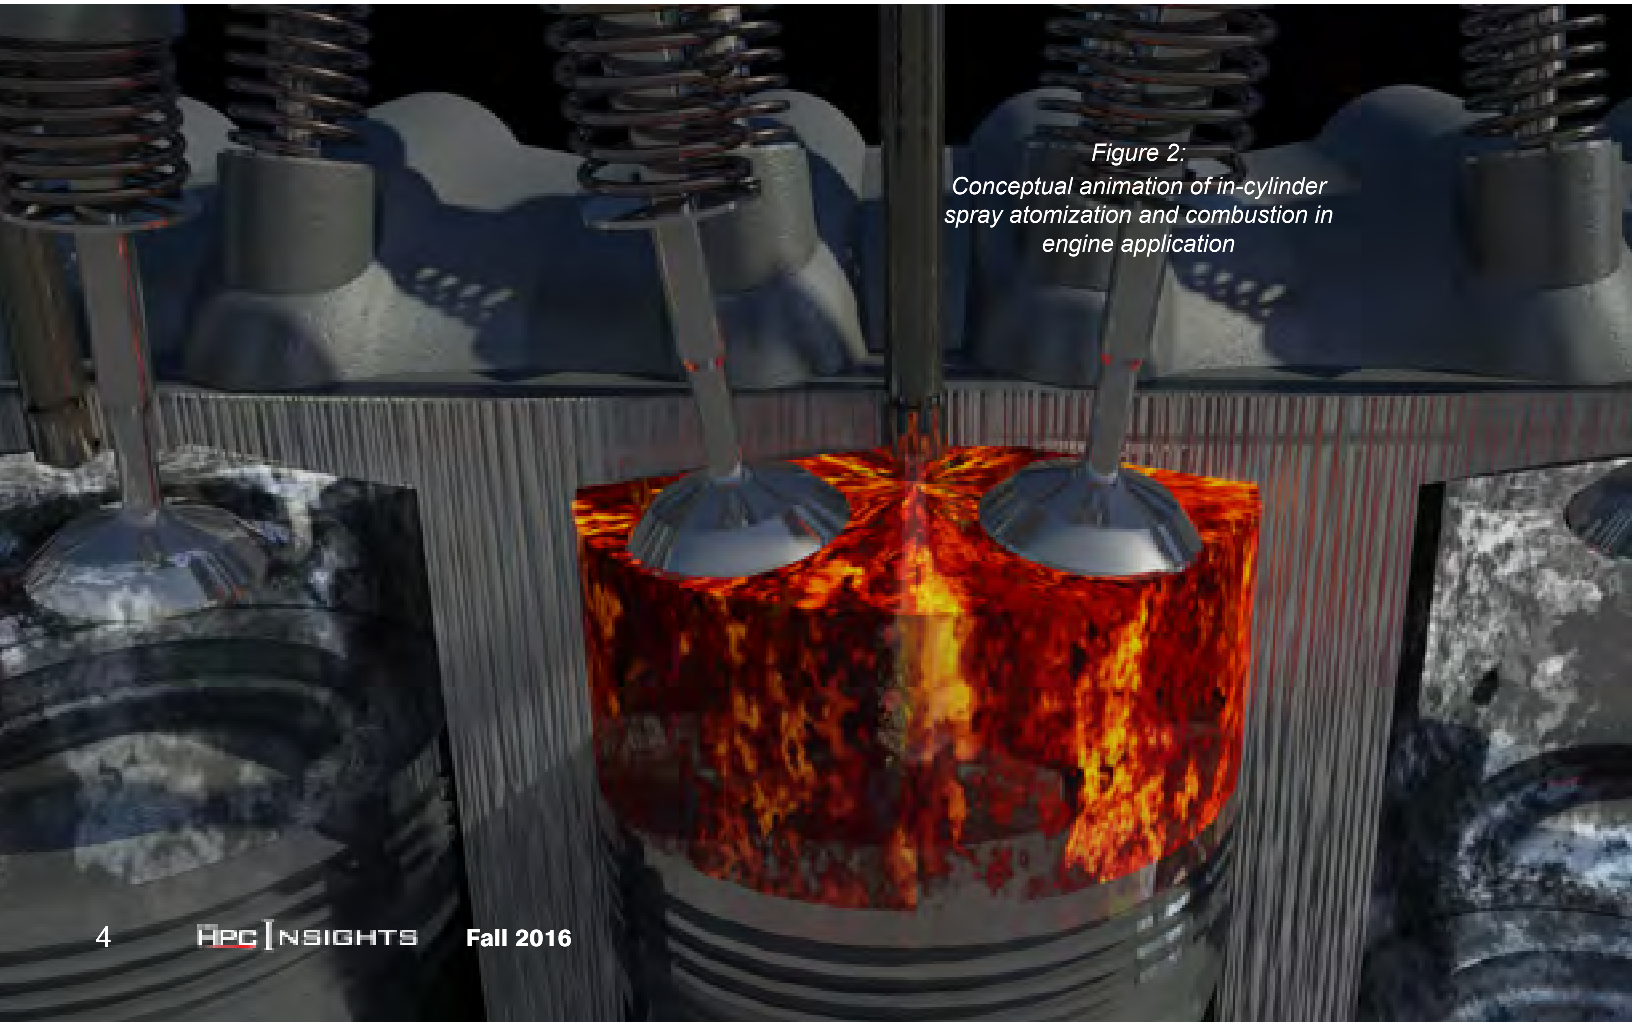

Figure 2: Conceptual animation of in-cylinder spray atomization and combustion in engine application

Approach

The methodology for simulating spray primary breakup is based on the solution of the Navier Stokes equations (NSE) coupled to a geometric unsplit interface-capturing method for immiscible fluids. The volume of fluid (VOF) method ensures discrete conservation and boundedness of the volume fraction, F, by utilizing non-overlapping flux polyhedra for donor volumes. The method uses piecewise linear interface calculation (PLIC) representation to resolve the liquid/gas interface. For consistency (and stability), mass and momentum are convected using the geometric VOF method. The simulation framework is designed for the distributed computations of unstructured mesh-based methods on HPC systems. The unstructured grids allow an accurate body-fitted representation of the complex curvature internal passageway regions, e.g., internal injectors. Several validation studies have been performed that tested the numerical accuracy and robustness of the solver [3].

In simulating the physics of an impinging jet, a methodology based on smoothed particle hydrodynamics (SPH) is employed. In SPH, a field function (e.g., fluid property) is described by the integral representation method, which is reformulated based on the use of computational particles. In SPH, the drop, surrounding gas, and the solid wall are discretized into free-moving and/or fixed particles. As a result, it becomes straightforward to track drop deformation and the interface of the liquid and gas. The governing equations to describe the fluid motion are discretized into the particle space, instead of grid space that is used in conventional computational fluid dynamics. As a result, SPH has the advantage of reproducing drop deformation and incorporating the drop properties and wall conditions easily.

Simulation Details

The high-pressure common-rail diesel injector geometry studied accounts for the complex internal features including the mini-sac region (0.2mm3), needle valve position, and converging nozzle with 90 µm exit orifice. The injector was scanned at Argonne National Laboratory using X-ray phase contrast imaging techniques with 5 µm spatial resolution.

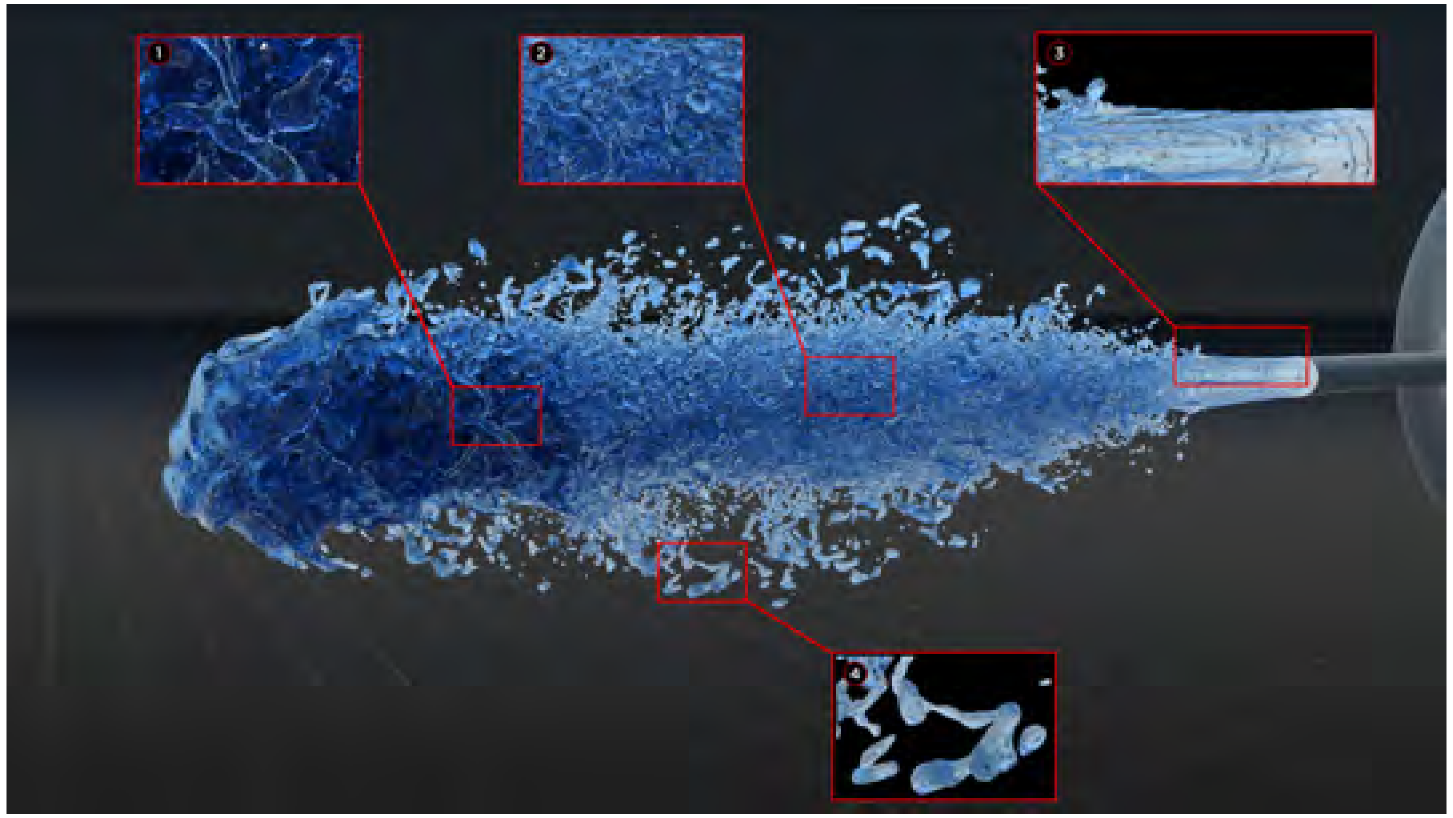

Figure 3: High-resolution visualization using ray-tracing of spray primary breakup phenomena from diesel injector.

The fuel pressure was specified at 150 bar, imposing Reynolds and Weber numbers Re = 25,573 and We = 125,806. The physical properties were based on a fuel temperature at 298K, as an approximation to the water-cooled injector jacket temperature in the laboratory. The liquid fuel properties employed are density ρ = 686 kg/m3, viscosity μ = 0.475 mPa.s, and surface tension σ = 18.6 mN/m. To specify diesel-type conditions, the chamber gas density is set to ρ = 22.8 kg/m3, by using 100% filled gaseous nitrogen at 303K, and a back-pressure at 20 bar. The simulation initializes with a liquid-filled injector and prescribes a rate-of-injection profile with bulk inflow velocities based from nozzle flow measurements. In addition, a stochastic turbulent inflow generation condition is employed to capture the transition to internal flow turbulence dynamics. The simulations provide a detailed diagnostics in the optically-dense region, 0 < x/d < 30 jet diameters, seamlessly calculating droplet clouds, and fuel/air mixture formation processes [4]. Figure 3, shows the spray formation process, and is enumerated to depict at points 1-2 the ligament structure in the dense region, at point -3 the onset of surface instabilities, and at point -4 a characteristic outer ligament and droplet length-scales for this injector. The results provide insights and new understanding of the breakup phenomena and droplet formation process for pulsed-injection diesel sprays.

The computational expense of resolving all the critical length scales at large Weber numbers is prohibitively high, so the number of detailed numerical simulations conducted at realistic diesel injector conditions has been limited. Note, a liquid jet moving with an O(100) m/s relative velocity with respect to the quiescent gas can generate droplets with diameters as small as a few microns. Predictions employing interface-capturing methods and Lagrangian particle representations will have a high computational demand. To model the range of spatial scales present, spanning over six orders-of-magnitude, requires computer codes that can exploit massively-parallel architectures, and with it millions of CPU-hours to describe the physics. The simulations are hence only possible with Frontier project resources. The simulations presented here were performed using a direct numerical simulation (DNS) approach exploiting large distributed memory parallel computers [1].

The numerical simulation of an impinging jet is accomplished by accurately predicting the details of drop-wall interactions. In a diesel engine, the size of a typical fuel drop can be 50~100 micro meters during wall impingement, with a velocity around 50~100 m/s. It is anticipated that a drop will be discretized by a few thousand SPH particles. The size of a numerical SPH particle can be 2.5~5 micro meters in diameter, which is the resolution of the computational domain inside and surrounding the drop. At such a small scale, the liquid-gas interface can be resolved in detail. Moreover, the surrounding gas phase can also be resolved with high levels of detail, similar to those employed in typical DNS for flow simulation.

Figure 4 shows a sequence of predicted images during a drop-wall interaction event. A series of diesel drops impact the piston surface at a 45-degree angle. The initial drop diameter is 100 micro meters, with a velocity of 50 m/s. The leading drop impinges on the piston surface and creates a liquid film. The subsequent drops impact the film, causing the film to spread and generate liquid ligaments and secondary droplets. These ligaments can further form droplets as time progresses. In a combustion engine, the gas flow will alter the trajectory of fuel drops, ligaments, and secondary droplets. The high-temperature gas and wall in the combustion chamber will also cause the liquid drops and wall films to vaporize and create combustible mixtures. These phenomena require further investigation by coupling the present numerical method with advanced physical models.

Scientific Data Visualization

A suite of high-resolution visualizations and animations were created with support of the Data Analysis and Assessment Center (DAAC) scientific visualization team [5]. The visualizations are helping us uncover several complex three-dimensional features of surface tension-dominated flows and spray atomization dynamics. The visualizations describe in detail the spray formation process that is inaccessible, and remains extremely challenging to laboratory measurements. The application of ray-tracing algorithms to the visualization further enhances the realism of the spray rendering, and enables us to make an assessment of the spray structures and detailed droplet features.

ARL's state-of-the-art high-resolution Cinemassive tiled display unit (Figure 5) is also playing a crucial role in accelerating our discovery process. The 24-tiled display is in a 6x4-curved configuration with a total of 49.76 megapixel available for display, and uses ParaSAGE to enable the scientific visualization application [6]. Because of the high-resolution environments we're able to seamlessly visualize and analyze Terabytes of data generated from our massively-parallel HPC simulations of spray atomization. The simulations resolve the droplet particle down to 1micron in diameter and, as such, require a high-resolution visualization solution and

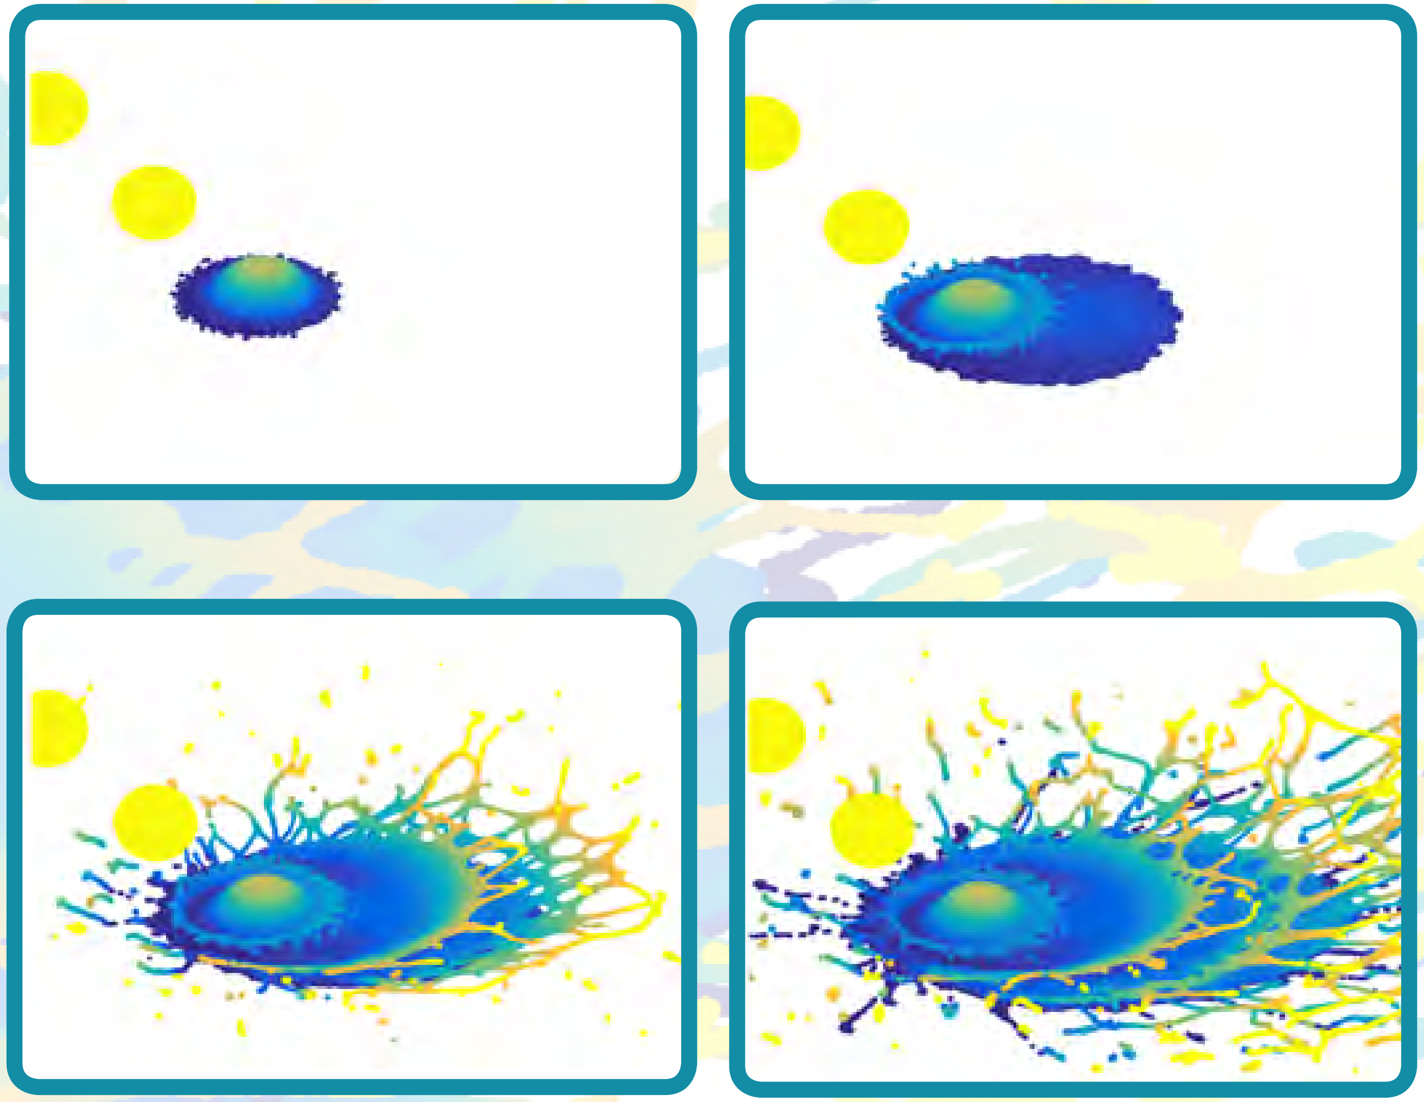

Figure 4
Sequence of simulated images describing the interactions of fuel droplets with piston wall.

display unit to extract critical surface morphology features. The visualizations enhance our understanding of the spray primary breakup quantifying the complex 3D interface topology of coherent liquid structure emerging during atomization, such as ligaments, lobes, and bi-modal droplet distributions. Scientific visualization played an important role in identifying the micron-sized liquid structures and droplets in 3D. Understanding this process is crucial for discovering new physical mechanisms and theories for multi-physics sprays.

In summary, the visualizations are effective in providing insights of the fuel/gas interface behavior and morphology including the formation of surface corrugations and rupture of the liquid core into the finer structures and droplets. Without this tool, we would not be able to identify and capture the various complex and interrelated features occurring during start of injection. Predicting the cloud of fuel droplets accurately and efficiently can have a transformative impact in improving design of military and industrial chemical propulsion applications and can lead to quantum leaps in the performance of Army combustors.

Acknowledgements
Financial support for this work was provided by the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) Frontier Project titled, "High-Fidelity Simulation of Spray Atomization and Spray Wall Interactions at High Pressure and Temperature Conditions". Support is also gratefully acknowledged from the Army Research Laboratory, Vehicle Technology Directorate (VTD), 6.1 base program. Computing resources were provided through the DoD High Performance Computing Modernization Program. We also thank Mr. Vu Tran, Rick Angelini, and Simon Su, from the Data Analysis and Assessment Center (DAAC) for their assistance with data visualization.

References
1. Fansler, T.D., Parrish, S.E., "Spray measurement technology: a review", Measurement Science and Technology, 26 (1), 2015.

2. Bravo, L., Kweon, C.B., "A review on liquid spray models for diesel engine computational analysis", ARL Technical Report, TR-6932, May 2014.

3. Bravo, L., Ivey, C.B., Kim, D., Bose, T., "High-fidelity simulation of atomization in diesel engine sprays", Center for Turbulence Research Proceedings of the Summer Program, 2014.

4. Bravo, L., Kim, D., Tess, M., Kurman, M., Ham, F., Kweon, C., High-Resolution Numerical Simulations of Primary Atomization in Diesel Sprays with Single Component Reference Fuels, 27th Annual Conference on Liquid Atomization and Spray Systems, Rayleigh, NC, 2015.

5. Bravo, L., "Simulations of High-Pressure Primary Atomization using a Complex Diesel Injector", DSRC HPC Data Analysis and Assessment Center, https://daac.hpc.mil/gallery, 2016.

6. Kobayashi, D., Su, S., Bravo, L., Leigh, J., Shire, D., "ParaViewWebSAGE2: A Community-Based, High-Resolution, Universal Scientific Visualization Application", IEEE Vis 2016, Oct 23-28, Baltimore, MD, 2016

Figure 5
High-resolution visualization using ParaSAGE on a 49.76 Megapixel screen [6].

DREN
Bandwidth Upgrades

Ralph A. McEldowney,
Associate Director for Networking

Imagine enough network bandwidth to transfer the Library of Congress' entire print collection (which was 208 TeraBytes in 2011) in less than 12 hours from Aberdeen, MD to San Diego, CA. Imagine enough bandwidth to move 3,600 DVD-quality high-definition (HD) movies (at an average of 5 GigaBytes per movie) or 720 BluRay-quality HD movies (at an average of 25 Gigabytes per movie) in 1 hour from Dayton, OH to Vicksburg, MS. Imagine the possibilities with large bandwidth network connections…

Fortunately, DoD High Performance Computing Modernization Program (HPCMP) users no longer have to imagine large bandwidth network connections. Large bandwidth is now a reality with the recent network upgrades at the DoD Supercomputing Resource Centers (DSRCs) on the Defense Research and Engineering Network (DREN).

In the fall of 2014, the HPCMP issued DREN bandwidth upgrade orders, from 10 Gigabits per second (Gbps) to 40 Gbps, for the four CONUS-based DSRCs and SPAWAR-San Diego, CA. These upgrades were anticipated to support greater file transfers between the Centers, enable next-generation network technology and protocol research, and permit coast-to-coast testing at some of the highest data rates available on the DREN.

By contract, the DREN service provider (CenturyLink) had 180 days to deliver the new service. On May 28-29, 2015, the government conditionally accepted the 40 Gbps service at the following DSRC sites: Navy DSRC at Stennis Space Center, MS; Engineer Research and Development Center (ERDC) DSRC in Vicksburg, MS; and the Air Force Research Laboratory (AFRL) DSRC at Wright-Patterson AFB (WPAFB), OH. These sites were conditionally accepted, as CenturyLink worked with the HPCMP and industry partners to develop a test suite capable of demonstrating an 80 percent single- stream (i.e., 32 Gbps) data transfer test. A demonstration of this type is not easily achieved, and pushes the technical boundary of compute servers to generate the required amount of traffic. As an alternate test for conditional acceptance, CenturyLink demonstrated 34 Gbps of traffic across three parallel data streams at each DSRC. The three data streams consisted of one 22 Gbps stream and two 6 Gbps streams. Although not a single-stream test, the alternate test was also exceptionally difficult to demonstrate.

Following the conditional acceptance tests, on June 9, 2015, the Navy DSRC successfully transitioned its site enclaves from the 10 Gbps service to the upgraded 40 Gbps service, becoming the first DSRC to complete the transition. On June 18, 2015, the ERDC DSRC transitioned its site enclaves; and shortly thereafter on June 29, 2015, the AFRL DSRC transitioned its site enclaves. On July 8, 2015, the HPCMP conditionally accepted the upgraded 40 Gbps service at the Army Research Laboratory (ARL) DSRC at Aberdeen Proving Ground (APG), MD. The ARL DSRC completed the CONUS-based DSRC transitions by moving its site enclaves to the upgraded 40 Gbps service on July 17, 2015.

The upgraded service at SPAWAR-San Diego was significantly delayed due to a fiber build-out construction project in the San Diego area. However, on April 25, 2016, the government accepted the 40 Gbps service at SPAWAR-San Diego and shortly thereafter, the site transitioned their site enclaves to the new service.

Throughout the fall of 2015 and winter of 2016, CenturyLink continued to work with the HPCMP and industry partners to develop a test suite capable of demonstrating a 32 Gbps single-stream data transfer test. On March 31, 2016, CenturyLink successfully demonstrated the single-stream data transfer test, and full acceptance of the 40 Gbps service upgrades at the DSRC sites was achieved.

The upgraded bandwidth at the DSRCs represents a 4X increase over the previous bandwidth, and a 16X increase in bandwidth over DREN II. This is a significant increase in capacity that will greatly benefit the HPCMP research and test communities; resulting in improved network performance and data transfer capability. Additionally, it will allow the HPCMP to perform next-generation network technology and protocol research for high-speed encryption, Software-Defined Networking (SDN), and high-speed Information Assurance (IA). It will also support new 40/100 Gbps testbeds located at WPAFB and SPAWAR-San Diego. As of June 2016, the number of active or ordered DREN locations is 166, and the number of SDREN sites is 71. The current total aggregate bandwidth for all DREN locations is 333.9 Gbps–a nearly 8x increase over the DREN II bandwidth. With large bandwidth connections available at the DSRCs and SPAWAR-San Diego, HPCMP users can now imagine all the possibilities.

Highlights of the HPC Internship Program (HIP)

Cynthia Dahl
HPC Insighs Editorial Board

Veterans of the specialized high-performance computing (HPC) world understand that the use of HPC resources requires unique knowledge, skills and abilities which are rarely introduced or obtained in school. Most HPC users relate that their initial supercomputing experience was obtained on the job. By pairing incoming talent (interns) with experts in a specialized field (mentors), the HPC Internship Program (HIP) is closing this gap, and better supporting the DoD's future Warfighter needs by providing future workforce candidates with the computational skills and experience necessary to become knowledgeable HPC users before even entering the workforce. Returning HIP mentor Dr. Mel Roquemore (AFRL) said, "For many interns, this is their first experience with computations and supercomputing. This [HPC internship] has provided a basis for their entire career." Dr. Alexei Poludnenko (NRL) concurs. "Nobody trains students for high-performance computing. They learn programming, but not supercomputing—things like how to set up a job, etc. That's why these internships are so important."

In addition to closing the educational technology gap in the future DoD workforce, HPC internships also help DoD organizations and mentors identify the best prospective DoD employees. "HIP internships allow you to really get to know the interns—you get a different, more in-depth knowledge of people," says Poludnenko. They allow mentors to "review the talent," adds Roquemore, "and see if there is a possibility to bring them to work on our program long-term."

This year, the HPCMP awarded 20 HIP projects with a total of 34 internship slots in DoD laboratories and organizations across the United States, some of which are highlighted here.

Air Force Research Laboratory – Wright-Patterson AFB, OH

Computational Insights into Extreme Light Phenomena

Mentors: Mel Roquemore and Chris Orban
Featured Intern: Scott Feister

The Extreme Light group at AFRL is investigating a next-frontier area of fundamental research involving ultra-intense laser-matter interactions. These interactions are important to the DoD, because they can produce very short pulses of x-rays and fundamental particles (electrons, protons and neutrons, etc.) that can be used as sources for the next generation of diagnostic systems for medical and materials research, and inspection of weapon systems components.

The use of high-resolution hydrodynamic and electromagnetic simulations is essential for understanding the complexities of laser-matter interactions. For the last three years, the HPC Internship Program (HIP) has provided student support to AFRL's extreme light simulation studies. Thanks to this support, seven different interns developed expertise with six different computer codes written in multiple languages to work on various aspects of the dynamics of the extreme light interaction processes. The value of HPC interns to AFRL's research is evident, since they are authors or co-authors to several peer-reviewed scientific papers.

HIP intern Scott Feister is the lead author of one of these papers. With an experimental background, but little experience in HPC, Feister was one of the first HIP interns who participated in the 2014 program. His research focused on how the expansion of plasma affects ultra-intense laser interactions. The group had studied plasma expansion experimentally, but this wasn't sufficient to understand the problem. Feister adapted and utilized the alpha version of a C- and Python-based open-source code to gain insight into the expansion process. Using this tool, he determined the optimal conditions for x-ray and particle acceleration. Feister's work turned into a scientific paper that was published in Review of Scientific Instruments, and he was hired by an on-site contractor to continue the HPC research he had started. This research also became the basis for his PhD thesis, which Scott received in May 2016 from the Ohio State University.

Dr. Scott Feister returned as an HIP intern in summer 2016. He performed massively-parallel simulations on DoD HPCMP resources to understand why certain plasma conditions accelerate electrons so efficiently using the hydrodynamic code FLASH (see figure below) and a particle-in-cell code in this project. He became so skilled at running 3D simulations with the FLASH code that the Flash Center for Computational Science at the University of Chicago offered him a postdoc position, which he accepted. His research there will involve using the 3D FLASH code on the Mira supercomputer at Argonne National Lab.

Several other HIP interns from the Extreme Light Group at AFRL have also had major successes. Gregory Ngirmang's 2014 summer research resulted in him being the lead author in a peer-reviewed paper published in Physics of Plasmas. This research will be an important part of his PhD thesis. Summer 2015 intern Abe Handler made major contributions to a scientific paper from the group, and recently accepted a permanent staff position at Lawrence Livermore National Laboratory. The HPCMP's internship program has provided crucial support to these students as they integrate HPC methods into their skillset and pursue scientific careers.

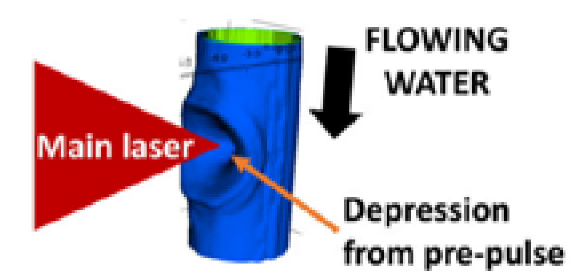

3D Flash simulation

Naval Research Laboratory – Washington, D.C.

Turbulence-Chemistry Interactions and Turbulent Compressibility Effects in High-Speed Reacting Flows

Mentor: Alexei Y. Poludnenko

Featured Interns: Colin A.Z. Towery and Ryan Darragh

The US Naval Research Laboratory, Washington, DC (NRL) is currently investigating the interactions between turbulence and flames—work that is at the heart of a wide variety of energy-generation and propulsion systems of critical importance to the DoD. These range from diesel power plants to conventional jet engines to more novel designs, such as scramjets for hypersonic flight. During the course of their HIP internships, Ryan Darragh and Colin Towery from the University of Colorado, Boulder have contributed to a broader research effort at NRL aimed at gaining the fundamental understanding of turbulence-flame interactions from two novel perspectives: (i) Lagrangian analysis of the thermochemical evolution of the fluid passing through premixed flames, and (ii) analysis of compressible turbulence in the context of the combustion dynamics in high-speed flows.

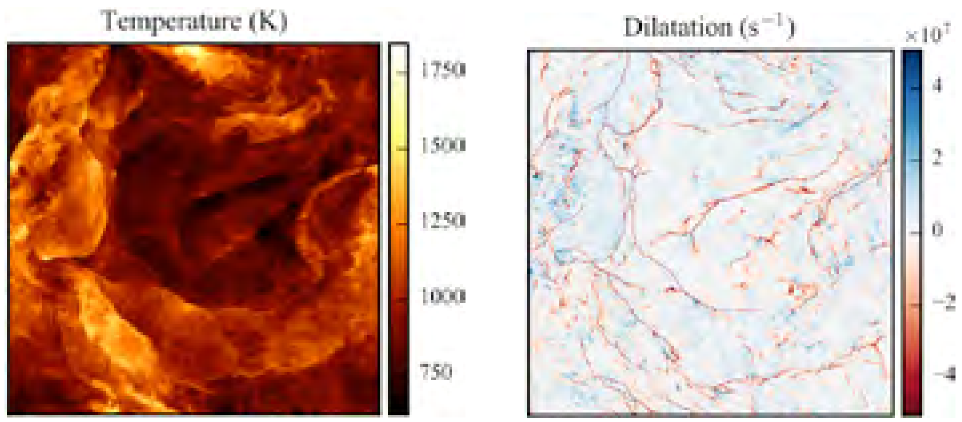

Sample structure of the flow-field formed by the highly-compressible turbulence (Mach 1.0). Left panel shows the distribution of temperature, while the right panel shows the dilatation of the velocity vector field.

Note the presence of extremely large temperature fluctuations in the range of ~700–1800 K in a flow with an average temperature of ~1000 K. The right panel illustrates the presence of pronounced shocklets dominating the flow.

Lagrangian analysis has long been used to analyze non-reacting turbulent flows, and has recently gained attention in the combustion community. An accurate calculation of the Lagrangian trajectories can, however, be rather difficult due to the chaotic nature of turbulent flows and the added complexity of reactions. Over the course of his 2016 summer internship, first-time intern Ryan Darragh ran extensive tests to accumulate data of the trajectories, then analyzed results. Darragh's work showed a surprising phenomenon, namely that fluid parcels were often found to undergo not simply gradual heating as fuel was consumed by reactions, but rather a sequence of alternating episodes of heating and cooling. This result shows that intense turbulence often found in modern propulsion systems can have a significant effect on the chemical evolution, which is typically not accounted for in modern turbulent combustion models used for the design of practical combustion applications.

In the second research effort, returning intern Colin Towery focused on the effect of fluctuations in the thermodynamic state of the fluid imprinted by high-speed turbulence on the combustion dynamics. In the course of this project, Towery used DoD HPCMP resources to perform a detailed verification study of the Athena-RFX code, developed in part at NRL, in the context of its ability to model high-speed, compressible turbulence. Results of this verification study are currently being used to develop a suite of novel, high-resolution DNS models of compressible reacting turbulence (see above Figure) intended to provide insights into the effects of compressible turbulence on ignition, flame propagation, and the probability of catastrophic flow transitions, such as the spontaneous detonation formation.

Both projects required over 30 million CPU-hours, and were only made possible by the resources available on DoD HPCMP's supercomputers. In the course of their HIP work, both students gained first-hand experience in carrying out a large-scale computational research effort on the DoD HPCMP resources, and had an opportunity to experience the research environment at a major DoD research lab. After finishing his PhD work at the University of Colorado, Boulder, Ryan Darragh intends to pursue a career with a national lab. Colin Towery plans on seeking a postdoctoral fellowship in uncertainty quantification (UQ), assimilation, and reduced-order modeling (ROM), after which he intends to pursue a career in research and development in aerospace propulsion and power systems.

Army Research Laboratory – Aberdeen Proving Ground, MD

HPC Simulations of N-dodecane Jet Spray at Diesel Conditions

Mentors: Luis Bravo, Anindya Ghoshal

Featured Intern: Louis Wonnell

The engine fuel injection process plays an important role in determining fuel-air mixing, combustion characteristics, and emissions formation. Due to the DoD's ever-rising need for improved fuel economy and increased power density, a fundamental understanding of its process is essential and remains a challenging task. At ARL's Vehicle Technology Directorate, intern Louis Wonnell employed high-fidelity simulations and high-performance computing to analyze the turbulent jet spray that is produced by the high-pressure injection of these complex fuels. Specifically, Wonnell implemented the LESLIE Multiphysics solver on DoD HPCMP resources. In this study, the diesel spray fuel/air mixture formation process was modeled with particular emphasis of the effects of chamber temperature (700K-1000K) on the mixing mechanisms controlling the penetration jet lengths. The results demonstrate that the existing breakup models can capture chamber temperature sensitivities that significantly influence the vaporization processes and control the penetration lengths.

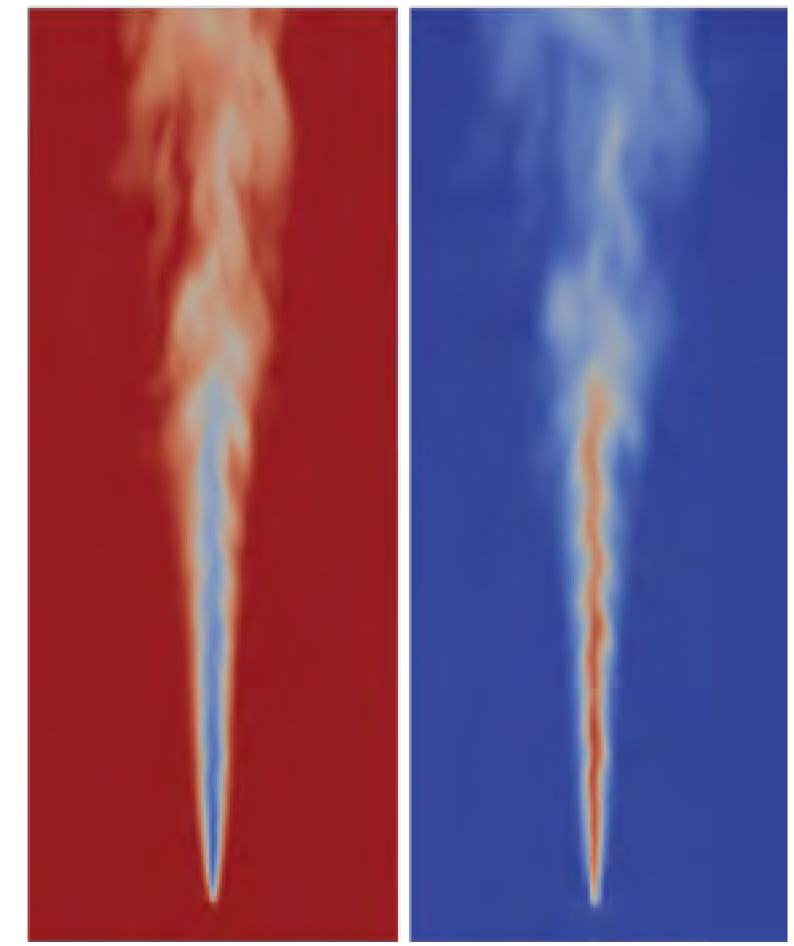

Above: Contours of n-dodecane Jet Spray shown at 0.6ms after start-of-injection (SOI) (a) temperature variations with min (blue) and max (red) values 360K-720K, and (b) velocity variations with min (blue) and max (red) values 0-370 m/s.

As an example, the figure shows a contour of temperature and velocity at the 700K and 60bar surrounding conditions. The figures show the fuel spray formation process, and the turbulent structure as depicted by the presence of eddies. During this study the liquid and vapor global penetration profiles were extracted and compared to experimentally measured values. With more accurate fluid models that can describe the thermodynamic behavior, the jet spray mixing can drastically change with respect to the diffusive mixing behavior, and can have an impact on the penetration jet lengths.

This assessment reveals that LESLIE can capture the essential physics of subcritical atomizing flows, and will serve as a foundation towards the development of more advanced models that can include supercritical combustion framework. The new tool will provide insights into the diffusive mixing process, and provide a predictive framework to evaluate futuristic engine designs and concepts.

Louis Wonnell's future academic goals include conducting PhD research at the Multi-scale Computational Physics Lab at Kansas State University using Morphing Continuum Theory (MCT). His work as a HIP intern has helped provide an Army mission experience with access to DoD leadership HPC platforms, conducting basic research in fluid mechanics and combustion science directly related to his career objectives as a researcher.

Army Research Laboratory – Aberdeen Proving Ground, MD

HPC-Enabled Computational Study on the Feasibility of Using Shape Memory Alloys for Gas Turbine Blade Actuation and on the Aerodynamic Performance of an Articulating Gas Turbine Blade Cascades

Mentors: Luis Bravo, Anindya Ghoshal

Featured Interns: Kathryn Esham and Richard Blocher

Two additional HIP internships at ARL's Vehicle Technology Directorate are part of a long-term project focused on increasing the efficiency of jet turbine engines and extending aircraft flight ranges by changing the shape (articulation) of the turbine blades in a jet engine through the use of shape memory alloys (SMAs). Existing gas turbine (GT) engines feature turbine blades that are fixed at an optimal design angle for a single flight condition. However, changes in flight conditions (velocity, altitude, etc.) can affect the optimal blade pitch angle. Articulating turbine rotors and stators, which maintain the optimum fluid-structure interaction for different flight conditions, would significantly increase the efficiency of aircraft engines and extend aircraft flight ranges. SMAs are a good candidate material for this application, because they offer an actuation mechanism which does not require the use of motors or hydraulic fluid.

Intern Kathryn Esham, using DoD HPCMP resources, focused on the material properties of shape memory alloys (SMAs), which provide a low-cost, compact, lightweight and high-power density option to change the shape of turbine blades via thermal actuation. Issues arise from high-temperature regimes required for operation within the engine, which can reach 900C, but the development of new high-temperature SMAs looks to push operating temperatures from the standard 100C for NiTi into these higher regimes. One way to prepare SMAs for high temperature applications is to add a third element (often Pt, Pd, Zn, Au, or Hf) to the NiTi composition, which creates nickel-rich nanoscale precipitates. How these nanoprecipitates work chemically and mechanically is not well-understood. This work uses phase field simulations on HPC systems to explore the microstructure in the local environment of a precipitate. Initial simulations investigated the effect of precipitates on structural variations. The figure shows two structural variations, depending on whether a third element is present. Qualitatively, the third element has changed the microstructure. Further analysis into the local stress around the precipitate will provide a quantitative measure of the variant disruption.

Future studies will investigate transformation temperatures and actuation properties through thermal cycling. Once the material properties are understood, microstructural finite element simulations will provide predictions for the stress development in the SMA during actuation. By investigating a novel material, high temperature SMAs, this work provides the basis for creating an actuating turbine blade system, which will increase efficiency and reliability of jet turbine engines.

While Esham's work focused primarily on calculating properties of the turbine blade materials that could be used to reshape them, Richard Blocher calculated the flow properties through the turbine, using DoD HPCMP resources to see what effect the changing blade shapes has on engine performance. In order to design an SMA actuator for turbine blades, the fluid flow across turbine rotors during actuation must be understood. This study presents computational fluid dynamics (CFD) results for simulated flow over turbine blades during actuation about the leading-edge. The CFD tool chosen for this study is CONVERGE multi-physics software. CONVERGE's features include a fully compressible Navier-Stokes equation formulation, rigid-body fluid-structure interaction modeling for immersed objects such as the rotors and stators in this study, and adaptive mesh refinement (AMR) to increase fidelity in areas of high local intensity. For a sinusoidal actuation amplitude of 10 degrees of rotation about the leading-edge of the airfoil, there was no significant difference in airflow behavior between an actuation period of 0.5 seconds and an actuation period of 0.25 seconds. Temporal snapshots of the velocity fields around the airfoils during actuation are shown in the figure below. In the near future, the use of different points of actuation, such as the aerodynamic center, geometric centroid, and center of gravity of the airfoil, will be explored. Future work will also examine the forces on turbine rotors and stators during actuation, and the effect of those forces on the phase transformation behavior of the SMA actuator through a finite-element (FE) simulation.

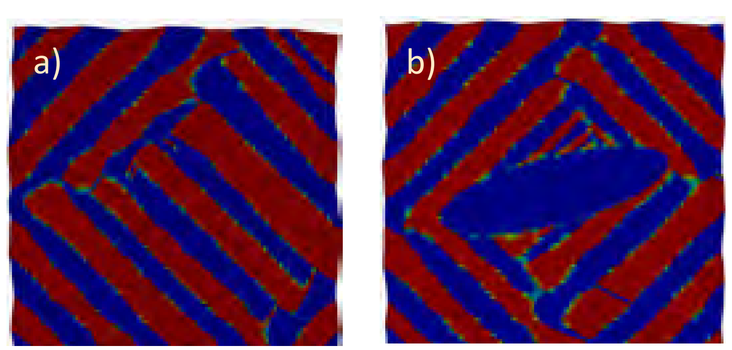

Martensite variant distribution in (a) precipitate-free NiTi and (b) NiTi with precipitate at 30°

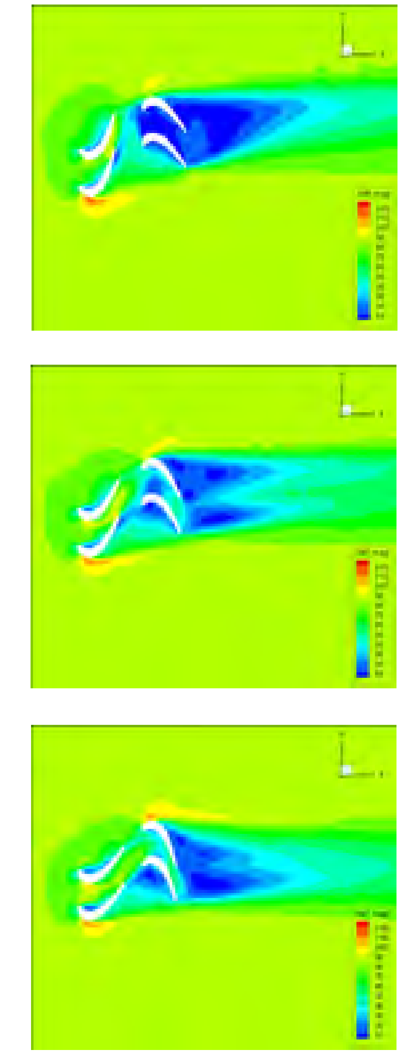

Velocity fields around rotor and stator blades at +10 degrees of articulation (top), zero articulation (center), and -10 degrees of articulation (bottom).

After completion of her summer internship Kathryn Esham will return to Ohio State University this fall to continue her graduate work focusing on precipitates in shape memory alloys. Having the opportunity to perform computational studies in a professional environment outside of academia has expanded her understanding of how high-performance computing can open new doors to explore novel materials, and has cemented her intentions to continue computational materials science research. Richard Blocher will continue work on this project remotely, while beginning his graduate studies in Materials Science at the Ohio State University this fall with the eventual goal of a PhD in Materials Science. The familiarity with HPC resources gained during this summer's internship, and the access to those resources, will significantly facilitate his research work.

Mentors across the Services agree that the work these HIP interns are doing is credible and valuable—both to their individual projects and to the DoD in general. But perhaps more important than the amount of data produced, methods developed and documented, or even papers published by these interns, is the interns' understanding that they have contributed to a much larger effort. Mentor Chris Orban (AFRL) says that he defines a successful internship as one in which "the interns come away from the experience feeling like they're a part of a team, and realizing that the problems they are working on are highly relevant to a large and important outstanding question." He likens this effort to the parable of the blind men and the elephant, "where each intern," he explains, "can eventually see their place in the bigger picture." The HPC Internship Program is helping mentors and interns work together to find solutions to real-world problems and to discover answers to these "important outstanding questions," while also helping prospective DoD employees find a place in the state-of-the-art world of DoD high-performance computing.

Hero Award Winners Announced

Denise O'Donnell,
High Performance Computing Modernization Program Office

Since their inception in 2003, the annual High Performance Computing Moderization Program Hero Awards recognize outstanding contributions by individuals who support the overall mission of the Department of Defense (DoD) and the HPCMP. The initial Hero Awards categories were Technical Excellence, Innovative Practices and Long Term Sustained Performance. In the years following, several additional categories were added to include: Up and Coming within the HPCMP, which recognizes those newcomers who are making excellent contributions to the program; and the HPCMP Interconnector, recognizing those whose support brings people together to solve one or more of the program's challenges. This year, we added the HPCMP Team Achievement award to recognize a group effort in support of the program and the DoD.

Traditionally, the Hero Awards were presented by the HPCMP Director at a ceremony that took place at the program's annual Users' Group Conference (UGC). Unfortunately, 2011 was our last UGC, due to the Department of Defense curtailing the hosting of conferences. However, the Hero Awards were deemed important to the program, so the tradition continues via teleconference or ceremony at the recipients' work locations. This year and last, Dr. David Horner, HPCMP Director, along with members of the program's senior leadership have personally presented and congratulated the Hero Award winners for their outstanding contributions to the program.

Lifetime Achievement Award:

In 2015, we were pleased to present a special Lifetime Achievement award to:

Dr. Jean Osburn

Dr. Osburn (Jeanie) has been an instrumental force in the continued modernization of the DoD's Research, Development and Engineering computational capabilities through advocacy, hard work and a passion for excellence. Her expertise and influence spans many computational domains, rapidly evolving computer technologies, and customer-focused computer center operations. She's one in a select community that understands the full potential for HPC's impact on the future of science and technology. Jeanie is a leader in the HPC community. The HPCMP extends our personal thanks and appreciation to Jeanie for the many years of service which she has given to our country.

In 2016 , we are pleased to present two outstanding individuals with Lifetime Achievement awards:

Mr. Brad Comes

Dr. Clifford Rhoades

Both have each been dedicated, hard-working members of the program for numerous years, making many exceptional contributions

While working for the program, Mr. Comes provided eight years of leadership to the highly visible, HPC Centers component. He also was instrumental in our technological insertion (TI) process, acquiring some of our greatest assets. Mr. Comes went on to become the Deputy Director of HPCMP, during such time he advocated and implemented a new Frontier Project activity. He ensured that the most dynamic computational projects received strong support from all program activities. Mr. Comes continued to make strides for HPCMP as an integral part of the program office, serving as an employee of the Intergovernmental Personnel Act (IPA).

Dr. Rhoades also wore several hats while with the program. He was an integral part of establishing the long-lasting structure of the HPCMP while serving as a crucial member of the selection committee for HPCMP shared resource centers. Dr. Rhoades has served as the S & T principal on the HPC Advisory Panel (HPCAP) serving as an instrumental player in the formation of the strategic plans, goals, and direction of the HPCMP. One of his greatest achievements, for which he will always be remembered for, is his role as our DREN III Transition Manager, where he interacted closely with CenturyLink, the DREN III service provider, and successfully transitioned 141 sites from DREN II to DREN III in less than 18 months!

It is with great pride and enthusiasm that the program celebrates all of our Hero Award winners. We encourage them to keep up the good work, and we wish our Lifetime Achievement Award winners luck in their future endeavors.

The Hero Awards

Long Term Sustained Performance:

Awarded to an individual whose support provides a superior contribution to the Research, Development, Technology and Evaluation (RDT&E) community for the last five years or longer.

Stan Carver, HPCMP DREN, Secure Mission Solutions (2015)

"For superb customer service and outstanding technical expertise in the implementation, operation and maintenance of the DREN and SDREN, supporting the High Performance Computing Modernization Program (HPCMP) community and the DoD RDT&E mission."

Sharon Amerg, ARL (2016)

"For demonstrated expertise and professionalism in supporting Army HPCMP users. Providing over 20 years support to the HPCMP, beginning as a systems administrator at the inception of the program, and then providing outstanding support as an Army Service Account Approval Authority. Her efforts in bringing United Kingdom researchers into the HPCMP community are particularly noteworthy."

Aram Kevorkian, SPAWAR Pacific (2016)

"For his dedication and contributions over the past 20 years to multiple aspects of the DoD HPC Modernization Program including establishing and leading the Baseline Configuration Team, participating on past Source Selection Evaluation Board Technical Panels, and serving as a Navy S&T representative on the User Advocacy Group."

The HPCMP Interconnector:

Awarded to an individual whose support brings a group of people together to effectively solve HPCMP challenges.

Bryan Comstock, Navy DSRC (2015)

"For demonstrated leadership, technical expertise, and professionalism in leading the technical team to consensus for the High Performance Computing Modernization Program (HPCMP) Technology Insertion FY14 and FY15 acquisitions."

David Lamb, TARDEC (2016)

"For demonstrated leadership, technical expertise, and professionalism in guiding the TARDEC Analytics Directorate in leveraging HPCMP resources to achieve significant impact on multiple programs of record such as Advanced Amphibious Assault and Light Armored Vehicles."

Timothy Madden, AFRL (2016)

"For demonstrated leadership, technical expertise, and professionalism in training researchers from the AFRL Directed Energy Directorate to access, leverage, and optimize HPCMP computing resources to better achieve the directorate's mission goals and reach program milestones earlier than anticipated."

HPCMP Team Achievement :

Awarded to a team whose support has resulted in specific exemplary performance of importance to the HPCMP and the Department of Defense over the past year.

US Naval Academy Center for HPC Education and Research – Carl Albing, Nate Chambers, Gavin Taylor (2016)

"For efforts in establishing the US Naval Academy Center for HPC Education and Research, encouraging 18 new faculty to use HPC resources, creating new HPC internships, and acquiring a Cray system as a testbed for education and research at the US Naval Academy."

Look for the Call for Nominations for the Hero Awards on our website in spring 2017.

www.hpc.mil.

Innovative Practices:

Awarded to an individual whose support demonstrates creative business practices to improve the overall HPCMP business model.

Robert Lunceford, Navy DSRC, Lockheed-Martin (2015)

"For developing a creative business practice for upgrading tape technology at the High Performance Computing Modernization Program (HPCMP) DoD Supercomputing Resource Centers, ensuring that users have the storage capabilities needed while enabling cost-effective budget solutions for storage upgrades within the HPCMP."

Robert Thornhill, Navy DSRC (2016)

"For demonstrated technical expertise, leadership, and professionalism in providing oversight of the design and implementation of facility modifications to two separate facilities required to support FY15 DoD HPCMP Technology Insertion. Additionally, he implemented creative approaches to reducing energy consumption."

Technical Excellence:

Awarded to an individual whose support demonstrates scientific or engineering excellence using HPCMP resources to advance creative and effective technology.

Thomas Kendall, Army Research Laboratory (2015)

"For technical expertise in leading the effort to engineer and oversee the construction effort to create two new HPC facilities with state-of-the-art energy-efficient power and cooling infrastructure. The on-time completion of this major effort enabled an early delivery, acceptance, and transition to production of the 101,312-core Excalibur system."

Joel Mejeur, NSWC Dahlgren (2016)

"For demonstrated technical expertise, leadership, and professionalism in spearheading the computational modeling and simulation effort of the Navy's electromagnetic railgun program. Mr. Mejeur's efforts have significantly impacted launcher design to increase performance and reduce risk. His work has encouraged and improved collaboration between the government and contractors for the design and development of the railgun."

Up and Coming within the HPCMP:

Awarded to an individual who has been with the program for two years or less whose support provides a distinct contribution to the HPCMP community.

Robert Knapke, Arnold Engineering Development Center, Aerospace Testing Alliance (2015)

"For rapidly becoming proficient at executing and validating the CREATE-AV Kestrel product in support of the Air Force F-22 Digital Thread effort, demonstrating the exact skills needed by program offices to accelerate vehicle design processes and minimize costly redesign – proving that significant ground test cost savings are possible."

Kelly Dalton, Air Force Research Laboratory DSRC (2015)

"For leading a comprehensive outreach effort to expand the use of high-performance computing as a true enabling technology within the Air Force. These outreach efforts will right-size resource allocations in the Air Force to the mission areas that will have the highest impact to the DoD Warfighter."

Andrew Barnes, ERDC (2016)

"For being an instrumental part of the analysis and benchmarking of first-principle codes used for blast events to validate and improve their predictive capabilities. His dual role as an experimentalist and numerical modeler is highly valued within the DoD HPCMP community, and contributes to better communication between computational and experimental researchers."

Ben Parsons, ERDC-ITL (2016)

"For being instrumental to the HPCMP in increasing the understanding of data analytics tools and how data analytics tools can be effectively utilized on HPCMP's supercomputers. Additionally, for having established himself as the "go-to" guy for new HPC research endeavors."

John Gilbert, NSWC Carderock (2016)

"For demonstrated leadership and significant technical contributions to the Navy Energetic Modeling Oracle (NEMO) software development. Dr. Gilbert's efforts greatly facilitated the initial program release in April 2016. His expertise in computational fluid dynamics and interest in big data analytics are of true benefit to the HPCMP and Navy."

Stephen Guimond, AEDC (2016)

"For participation in designing replacement rotor blades for the AEDC 16T compressor using HPCMP CREATE™ software and HPCMP resources. His efforts providing critical aerodynamic load information for the structural design of one of the nation's premier wind tunnels were essential to maintaining the project's schedule and design goals".

First-Principles Modeling of Modern Jet-Engine Combustors with Petascale Computations:

Alexei Y. Poludnenko
Naval Research Laboratory, Washington, D.C.

Christopher Lewis, Vu Tran, Miguel Valenciano, Michael Wissmann
HPCMP DAAC, Vicksburg, MS

Few physical processes are as fundamental to the functioning of modern society as combustion – energy extraction from a fuel through rapid, highly exothermic chemical reactions. This can involve fuels in all phase states – gaseous, liquid, solid, and even plasma (in the case of thermonuclear reactions); however, in this discussion we will primarily focus on combustion of gaseous and liquid fuels. Virtually all of us rely daily on turbulent flames inside the automotive and aircraft engines for transportation, in gas-turbine power plants for energy production, or in countless industrial settings for chemical processing, manufacturing, etc.

In recent years, the growing awareness of the environmental dangers associated with hydrocarbon combustion has created a powerful momentum to replace fossil fuels with alternative energy sources in the automotive and power-generation industries. At the same time, for aircraft and rocket propulsion, presently there is virtually no alternative to hydrocarbon (and hydrogen) fuels. This does not mean that alternative energy sources are not being explored. At the time of the writing of this article, we have been witnessing a remarkable aviation milestone as the aircraft Solar Impulse 2 is preparing for the last leg of the bid to circumnavigate the Earth for the first time by a fixed-wing solar-powered aircraft.

Despite this achievement, major scientific and engineering breakthroughs will be required in the future before alternative energy sources become capable of providing performance comparable to modern jet and rocket engines in terms of thrust, specific impulse, etc. This aspect is particularly crucial for defense applications, in which gains in environmental impact cannot come at the expense of combat characteristics. As a result, in the foreseeable future, qualitative improvements in the design of conventional aircraft and rocket propulsion systems, as well the development of the novel ones, will depend on the advances in our ability to understand and model combustion processes in liquid and gaseous fuels.

Historically, development of combustion systems proceeded through an empirical "trial-and-error" approach, without necessarily an in-depth understanding of all aspects of the highly turbulent, reacting flow inside a combustor. Such an "agnostic" approach, while costly, still allowed conventional aircraft and rocket engines to reach a remarkable level of perfection. In recent years, both the advances in traditional combustion systems, e.g., turbine jet engines, as well as the emergence of more novel designs, e.g., scramjet engines for hypersonic flight, have been pushing reacting flows into ever more extreme conditions characterized by higher pressures, faster flow speeds, and leaner, more premixed mixtures. Such novel combustion regimes are far less understood, they are more difficult to diagnose experimentally, and it is far more costly to design engines for such regimes through simple trial-and-error. This has led to the emerging recognition of the potential provided by various computational modeling approaches. Numerical models can be used to probe the flame dynamics and understand the behavior of a reacting flow both during stable operation and, more importantly, in off-design regimes. Such insights often cannot be obtained through modern experimental techniques, and they are invaluable for guiding the development of stable, reliable systems providing major savings both in time and cost.

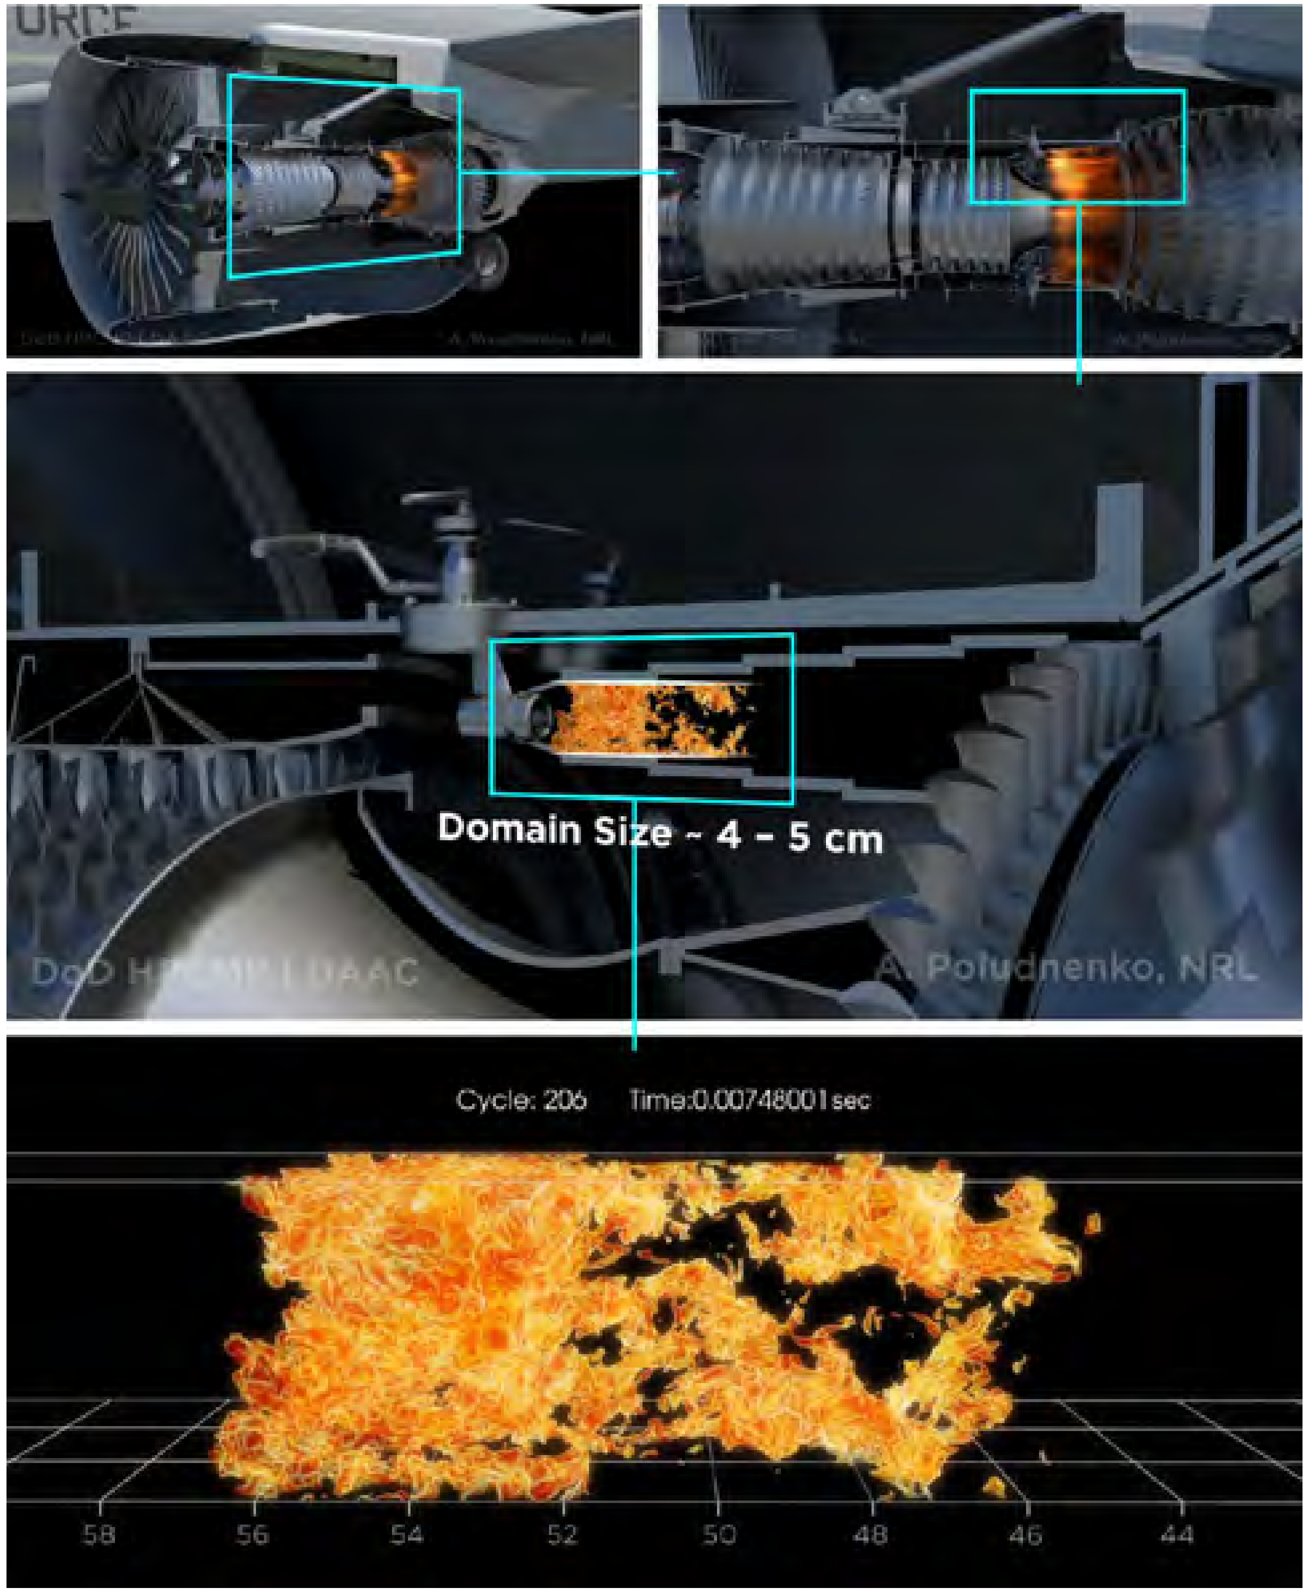

Figure 1:
Location and characteristic scales of a combustor in a typical jet engine. For comparison, the lower two panels show on the same scale a turbulent flame in the largest direct numerical simulation (DNS) performed at the Naval Research Laboratory to-date, containing 32 billion grid cells. In particular, shown is the isosurface of the fuel mass fraction corresponding to the peak reaction rate. Note that the turbulent-flame DNS was not performed in an actual combustor geometry, but rather in an idealized rectangular domain shown in the lower panel. Axis scale in the lower panel is in cm.

Virtually all existing computational approaches face one major hurdle, which is often referred to as the "tyranny of scales". This is illustrated in Fig. 1. A typical modern high-bypass turbofan jet engine has a fan diameter of under 3 m and an overall length of around 5 m (upper-left panel), while typical combustors have sizes on the order of 10 cm (upper-right and middle panels). On the other hand, characteristic scales of physical processes inside a combustor can be vastly smaller. For instance, the typical width of a flame which is folded inside a complex turbulent-flame structure (lower panel) varies with pressure, and can range from hundreds of microns under atmospheric conditions, to just a few tens of microns at high pressures present in modern turbine engines.

Such vast dynamical range of spatial (and, consequently, temporal) scales is further complicated by the rich, multi-physics nature of the reacting flow, which can involve complex chemical reactions, molecular transport processes, radiation, multiple phase states, etc., all evolving in a highly-dynamic, turbulent flow-field. All computational approaches face the challenge of capturing that complexity. In general, this cannot be done entirely from first-principles without a multitude of simplifying assumptions, which ultimately leads to an obvious question: How accurate are the resulting computational models, and how robust is their predictive capability?

The main compromise that a modeler must typically make is the decision to focus on either large or small scales of the flow. In particular, when the temporal dynamics of the flow cannot be ignored, this gives rise to the two classes of approaches known as Large-Eddy (LES) and Direct Numerical (DNS) simulations. LES are intended to model the entire flow-field inside a combustor taking proper account of a realistic combustor geometry, inlet and outlet conditions, boundaries, etc., while not resolving small scales explicitly. DNS, on the other hand, aim to provide the highest possible physical fidelity on small scales at the expense of not capturing the large-scale flow dynamics. Ultimately, these two complementary approaches are intended to form a synergy: DNS allows one to probe in detail the small-scale flow physics, which can then be distilled into accurate and efficient models that can subsequently be used in LES calculations to capture the effects of the unresolved small scales.

The field of turbulent combustion has largely inherited this paradigm from its counterpart focused on studies of non-reacting turbulent flows. Relative simplicity of the underlying physical picture in the latter case has led to a great success of this approach, even though work still continues actively in the community aimed at improving the LES small-scale (also known as "subgrid-scale") models.

One of the key premises of this paradigm, and the resulting LES-DNS duality, arises from the classical hypotheses regarding the nature of non-reacting turbulence, which were first introduced by A. Kolmogorov and colleagues, and which have since been examined and verified in many experimental and numerical studies. In particular, this paradigm is based on two, effectively empirical, observations: 1) the overall flow dynamics is predominantly governed by large-scale motions with small scales providing essentially an energy sink for the dissipation of the kinetic energy of the flow; 2) small scales have a universal nature, which is largely independent of the details of motions on large scales. Such scale separation is what allows one to study small-scale dynamics in idealized DNS settings, and distill the obtained insights into a compact set of "universal" subgrid-scale models, which can then be used in LES to model the infinite diversity of realistic flow-field configurations.

This picture becomes significantly more complicated once we introduce highly-energetic exothermic processes into the turbulent flow. At its core, combustion in a turbulent flow is a "two-way street". On one hand, turbulence controls combustion by folding and wrinkling the flame, and by providing a powerful mechanism to transport heat and various reacting species. On the other, combustion itself can modify turbulence, thus fundamentally changing the coupling between the two. Therefore, two key questions must be answered in this respect in order to gain confidence in the numerical modeling of reacting flows:

1. How does energy injection by the flame, which occurs primarily on small scales, impact the resulting turbulent flow-field? In particular, do key assumptions of the modern theory of nonreacting turbulence concerning small-scale universality and the separation of small- and large-scale motions still hold?
2. If and when the dynamics of turbulence change due to the presence of combustion, how does this affect the resulting flame properties and the rate of energy release in the system?

Recent studies by us, as well as other groups, of a broad range of turbulent combustion regimes are starting to form a picture of reacting turbulence, which is profoundly different from its nonreacting counterpart. In high-speed regimes, such as found in scramjet engines or in modern high-pressure jet-engine combustors, turbulence is no longer a minor perturbation; but rather is an important, and potentially dominant, factor shaping the resulting flame dynamics. We found that high-speed flames can be highly unstable, exhibiting large variability of the burning speed even under the most idealized upstream conditions. Somewhat paradoxically, highly subsonic reacting turbulence can also spontaneously produce shocks and even transition to a detonation without any external assistance of walls or obstacles – a phenomenon which has no analog in nonreacting turbulence[1]. Furthermore, combustion can also generate significant amounts of turbulence, far in excess of what can be ascribed to the upstream turbulent flow. Such energy is injected on small scales and is redistributed to larger scales, which can lead to the acceleration of the overall turbulent flame producing, in a sense, a self-propelling system.

Recent studies by us, as well as other groups, of a broad range of turbulent combustion regimes are starting to form a picture of reacting turbulence, which is profoundly different from its nonreacting counterpart. In high-speed regimes, such as found in scramjet engines or in modern high-pressure jet-engine combustors, turbulence is no longer a minor perturbation, but rather is an important, and potentially dominant, factor shaping the resulting flame dynamics. We found that high-speed flames can be highly unstable, exhibiting large variability of the burning speed even under the most idealized upstream conditions. Somewhat paradoxically, highly subsonic reacting turbulence can also spontaneously produce shocks and even transition to a detonation without any external assistance of walls or obstacles – a phenomenon, which has no analog in nonreacting turbulence[1]. Furthermore, combustion can also generate significant amounts of turbulence, far in excess of what can be ascribed to the upstream turbulent flow. Such energy is injected on small scales and is redistributed to larger scales, which can lead to the acceleration of the overall turbulent flame producing, in a sense, a self-propelling system.

An important consequence of these findings is that they point to a much more complex coupling between small and large scales than a traditional LES-DNS paradigm can account for. Indeed, if small scales can energize the entire turbulent flow-field, they become an important, or sometimes even the main, driving force controlling the flow dynamics and turbulence-flame coupling. This is contrary to the traditional logic of the LES approach, which assigns such role to large scales. This does not necessarily imply that we need to discard the LES approach, but rather we need to re-examine it more carefully in the context of high-speed turbulent reacting flows; and in particular we need to understand in detail the nature of the interplay between small scales and all other scales of the flow. From a computational point of view, this means that all relevant scales need to be captured in a calculation. In other words, one would need to extend the DNS beyond its traditional realm to model a realistic combustion system in its entirety. Is this possible?

Due to the "tyranny of scales", as well as the physical complexity of a reacting flow discussed above, a traditional answer would be – No! As will be shown below, in a general sense, that is still indeed true. At the same time, in certain cases, recent advances in physical models, numerical tools, and high-performance computing platforms place this seemingly impossible goal within reach.

Consider a flame in a mixture of vaporized jet fuel and air. Under atmospheric conditions, a typical flame thickness is ~0.4 mm. In our code, Athena-RFX, grid resolution of ~16 cells per flame width is generally sufficient to capture the flame dynamics accurately. This means that a computational grid with 2,048 × 2,048 x 8,192 cells (32 billion cells) would represent a domain with a size of 5 × 5 × 20 cm, which starts to approach a size of a realistic combustor. Figure 1 (lower panel) shows an example of a DNS of a similar size (32 billion cells in a 1,024 × 1,024 x 16,384 domain) performed by us on 65,536 computational cores on Cray XE6 (Garnet) at the ERDC DSRC.

That calculation used simplified single-step reaction kinetics. It is important to emphasize, however, that it is not sufficient to focus only on achieving a realistic range of scales representative of a typical jet-engine combustor. It is equally important to provide also a comparable level of realism of a physical model and, in particular, of the description of chemical reactions for a realistic jet fuel. Recent advances in the understanding of chemical kinetics of heavy hydrocarbons have resulted in a new generation of reduced reaction-kinetics models, which are sufficiently compact (< 30 species) and efficient to allow their use in DNS calculations.

The emergence of such reaction models has allowed us in the last two years, in the course of the ongoing Frontier project, to carry out the first systematic study of turbulent flame properties for realistic jet fuels (dodecane and JP-8/Jet-A) under flow conditions representative of realistic jet- engine combustors, i.e., high pressures (30 atm) and high temperatures (700 K), for a broad range of turbulent intensities – from relatively low to ultra-high. Figure 2 shows an example of such turbulent flame in a jet-fuel/air mixture.

From a physical standpoint, this survey has uncovered a number of surprising effects in terms of the flame structure and dynamics, which are characteristic of higher hydrocarbon fuels undergoing pyrolysis, and which are not present in lighter fuels, such as methane. On the other hand, from a practical point-of-view, it has demonstrated that it is computationally feasible to model realistic jet-fuel chemistry in large-scale DNS. For instance, this survey included 30 DNS calculations with the largest domain size of 1 billion cells, and it required a total of ~ 100 million CPU-hours on the HPCMP Excalibur and Thunder platforms at the ARL and AFRL DSRCs. In particular, the observed code performance with a stiff reaction-kinetics solver and mixture averaged molecular transport model was ~30,000 cells/core/s (or 33 µs/cell) both on Cray XC40 and SGI ICE-X platforms. We currently estimate that ongoing code modifications aimed at better vectorization, as well as improvements in the hybrid parallelization will increase performance on the HPCMP platforms by a factor of 2-3 on conventional CPUs without any assistance from external accelerators pushing a 100,000 cells/core/s (10 µs/cell) mark.

Let us finally revisit the question that we asked above: is DNS of an entire, realistic combustor possible? Figure 2 shows that the largest scales accessible today in a DNS at high-pressure conditions do not exceed 1 mm. Therefore, at high pressures, a DNS of a system with a size of ~10 cm would require a domain ~10^6 times bigger than shown in Fig. 2, or approximately 1000 trillion cells. This means that in our quest for a full realism – realistic combustor scales, realistic jet fuels, and realistic flow conditions, i.e., pressure, – all three conditions cannot be satisfied simultaneously in a DNS. At the same time, the first two criteria are certainly within reach of DNS on the petascale-class platforms available in the DoD HPCMP program. For instance, a system with a size of 10 × 10 × 20 cm, would require a domain with ~128 billion cells. With a notional code performance of 10 µs/cell, discussed above, this will translate into ~14 s per computational time-step on a 100,000-core-class platform, which is suitable for practical computations.

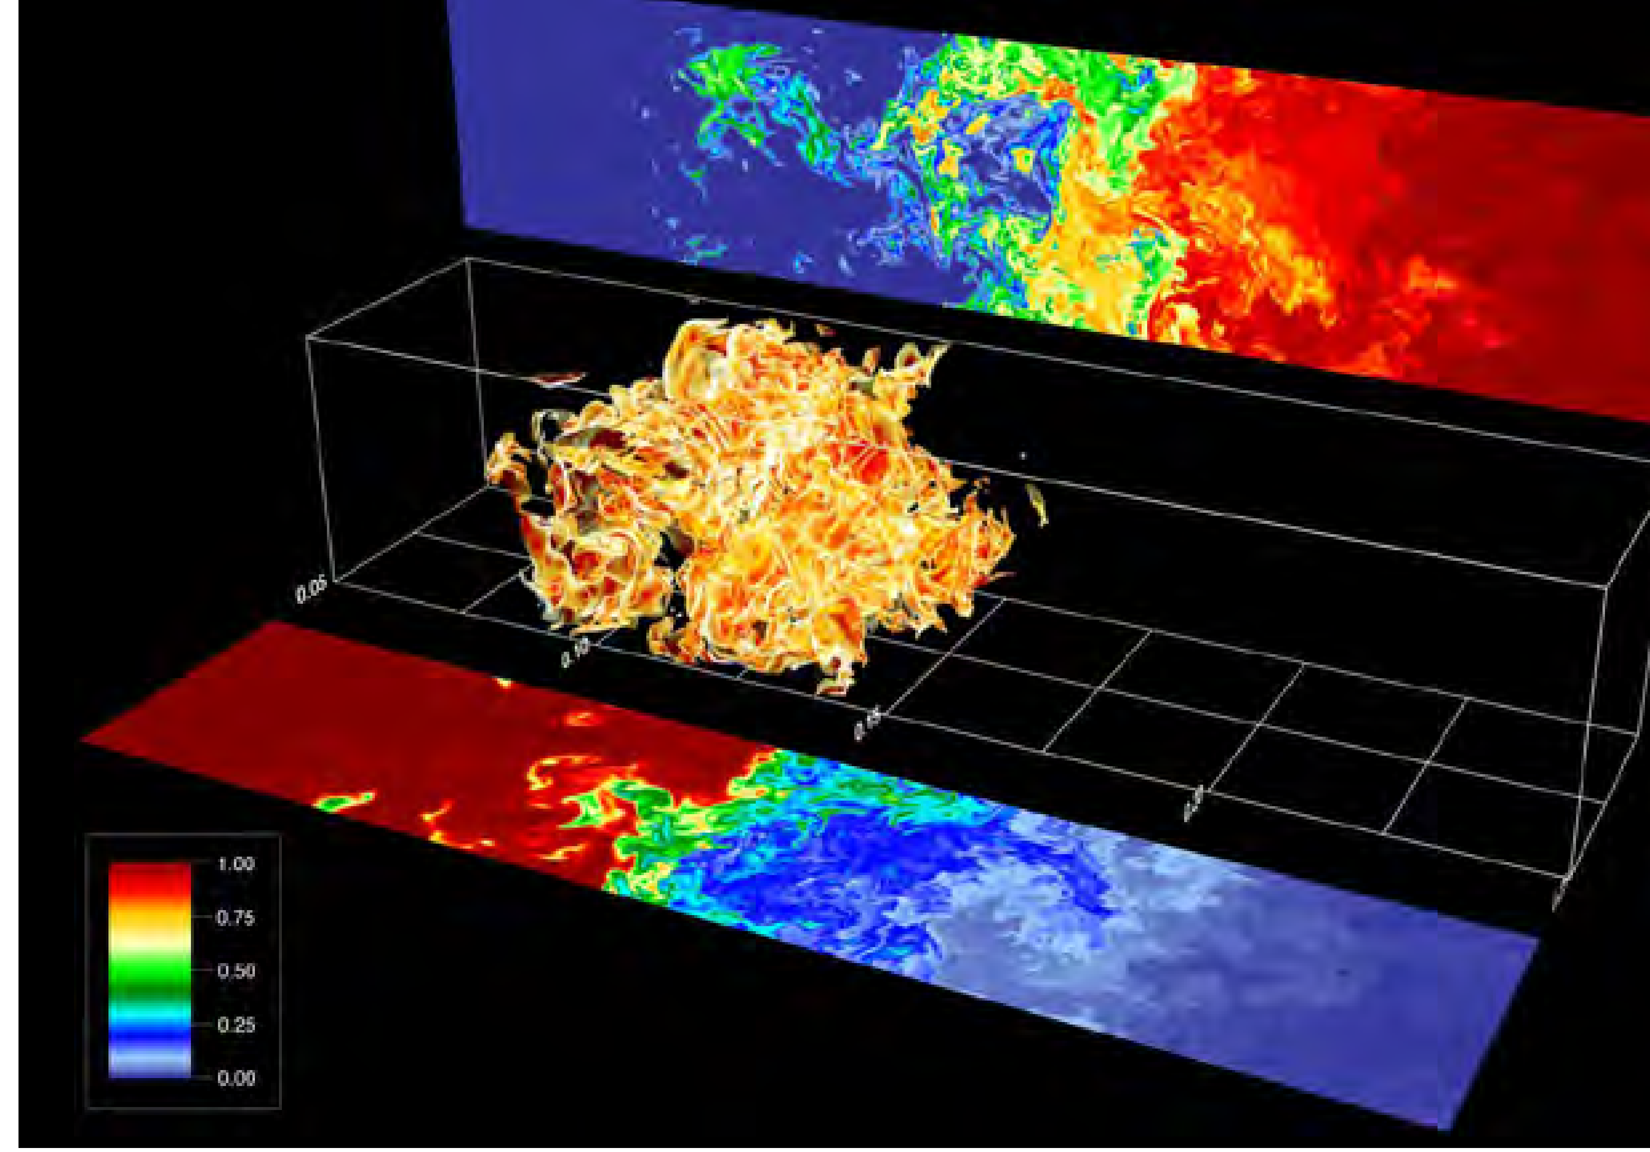

Figure 2:
Complex structure of the turbulent chemical flame in a lean mixture ($\phi = 0.7$) of jet fuel (n- dodecane) and air. Calculation is performed under high-temperature (700 K), high-pressure (30 atm) conditions characteristic of jet engine combustors in a flow with high-turbulent intensity (Karlovitz number 103). Uniform computational grid has dimensions 0.42 × 0.42 × 3.37 mm and resolution of ~0.82µm, shown in the center is the iso surface of temperature (1839 K) corresponding to the maximum rate of heat release (axis scale is in cm). Upper and lower planes are the corresponding 2D cuts through the domain showing the normalized mass fractions of fuel (upper-plane) and one of the burning products, namely water (lower-plane). In this direct numerical simulation, all characteristic scales of burning and turbulence are fully-resolved.

Such large-scale DNS of an entire realistic combustor, while obviously computationally expensive, would allow one for the first time to explore the dynamics of reacting turbulence over a realistic range of scales from first-principles with minimal model assumptions. As was discussed above, reacting turbulence can be profoundly different from its nonreacting counterpart, exhibiting a remarkable wealth of unexpected physical phenomena. Increasing the level of realism of DNS calculations will open a path to deepen our understanding of how those physical effects manifest themselves in realistic systems, to compare those findings directly with experiments, and ultimately to enable the development of a new generation of more accurate and reliable LES subgrid-scale models.

Acknowledgments
Financial support for this work provided by the Air Force Office of Scientific Research is gratefully acknowledged. Computing resources were provided by the DoD High Performance Computing Modernization Program under the Frontier project award.

Using HPC to Build a More Complete Space Object Catalog

Chris Sabol, *US Air Force Research Laboratory, Directed Energy Directorate, Kihei, HI*

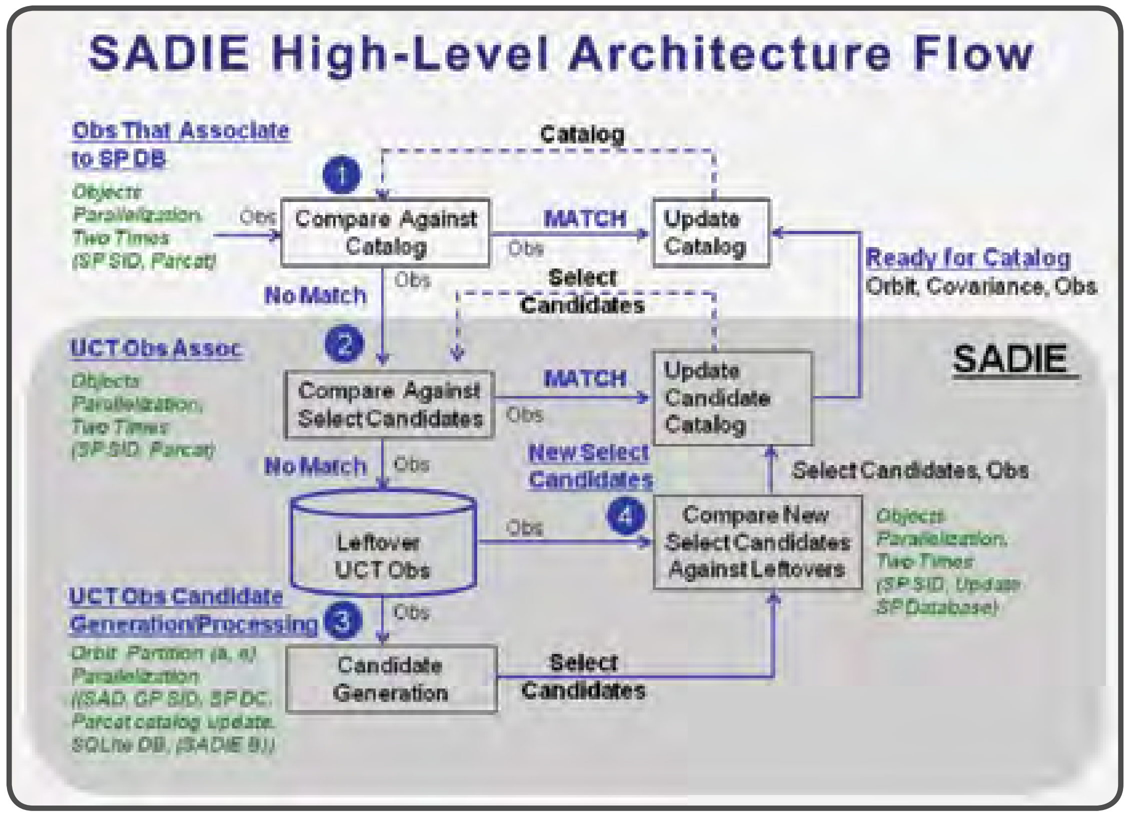

Figure 1: SADIE High-Level Architecture and Flow

Introduction

Making a list of things you know you have to do is easy. Preparing for all of the potential crises that may pop up, and quickly recognizing and effectively dealing with those that do is far more difficult. That is essentially situational awareness, and for the Department of Defense's (DoD's) interest in Earth-orbiting objects, going from maintaining a catalog of resident space objects, essentially the known to-do list, to predictive Space Situational Awareness (SSA) presents multiple challenges.

No other nation relies more on space than the United States for military operations, not to mention the large role it plays in our civil lives (television, GPS, Earth observation, etc., all rely on free access to space). Thus, maintaining the ability to use space, or maintaining space superiority, is of critical importance. SSA includes the location, movement, status, capabilities and characteristics of forces and objects in, from, to, and through space; the knowledge of man-made or natural threats; and the rapid analysis, organization and presentation of this information to operational decision makers. Space surveillance is a key component of SSA, and is the means by which we maintain custody of known objects and identify new objects in space. Much like keeping awareness of the cars around you in traffic enables you to drive defensively, maintaining awareness of the space object population, allows the DoD to operate in space.

The essential problem of space surveillance is data association. Once a sensor detects an object, do those observations correlate to a known object maintained in the space object catalog? And if not, what is it? Observation tracks that do not correlate to known objects are called Uncorrelated Tracks (UCTs), and associating sufficient UCTs together to form a catalog-ready orbit has been a challenge. Consider a hypothetical example of a collision in space, where two satellites suddenly become two debris clouds. Sensor 1 views scores of UCTs and labels them 1,2,3… Then another sensor, Sensor 2, later observes scores of UCTs and labels them a,b,c… How do we know which numbered-track from Sensor 1 is the same object as any lettered-track from Sensor 2? This example illustrates the space-track data association problem, except that the data association problem extends beyond collisions, it extends to all UCTs.

The DoD tracks more than 22,000 satellites and pieces of space debris on orbit, and those numbers are anticipated to increase greatly in the years ahead as improved sensors, such as the Space Surveillance Telescope (SST) and Space Fence, are expected to observe an order-of-magnitude more objects than are cataloged today. This will result in a substantial increase in the number of UCTs, and an unprecedented space-track data association challenge that requires high-performance computing.

Background

Data association for space is inherently different than the air, ground, and sea regimes. In space surveillance, we are data-poor. Surveillance sensors only get brief glimpses of objects, and observation tracks can be separated by hours or days. Furthermore, since the sensors and space object viewing geometries can be significantly different with each data collection, reliable and repeatable signature features are not available to aid discrimination. The good news is that gravity is the predominant acceleration driving the trajectories of orbiting objects, and dynamics can be used to help link hypothesized orbits with sparse data sets. This is different from other regimes where reasonably predictable dynamics are not available, and considerable effort is applied to collect observations that support kinematic and feature-aided tracking.

Building a more complete space object catalog was a challenge well-known to the High Performance Computing Software Applications Institute for Space Situational Awareness (HSAI-SSA), and developing a solution became a priority at the urging of the Joint Forces Combatant Commander, General Willie Shelton. The Search-and-Determine Integrated Environment (SADIE) project was initiated as a joint effort by Air Force Research Laboratory (AFRL) and the Naval Research Laboratory (NRL) under the HSAI-SSA to fully address this emerging need. SADIE has been developed to automatically build and maintain a database of orbits using radar and optical data, addressing Joint Space Operations Center (JSpOC) mission requirements.

HPC is a critical enabler of SADIE for a variety of reasons. First, SADIE was not simply an integration project, and considerable effort went into algorithm development and tuning. Nine different radar processing implementations were investigated, and testing involved dynamic closed-loop simulation of up to 100,000 space objects over multiple days. For the optical processing, computationally-expensive six-degree-of-freedom simulations were incorporated to better capture the dynamics and detectability of thousands of objects. HPC enabled modeling and simulation to support timely algorithm development and testing. Second, SADIE was developed using Special Perturbations (SP) orbit propagation code. The term "Special Perturbations" means that numerical integration is used to propagate satellite orbits. The related term "General Perturbations" (GP) signifies that the orbits are propagated with approximate analytical formulas that typically are much less accurate than the SP solutions. SP orbit propagation generally takes significantly more compute-time than GP solutions. Furthermore, the optical association approach scales as N3, and realistic data processing is not possible without HPC. Simply put, SADIE could not be developed, tested, or executed operationally without HPC.

SADIE Architecture

The SADIE architecture, shown in Fig. 1, separates the catalog maintenance and building processes into four major stages:

1) observation correlation against the catalog and catalog update;
2) UCT correlation against previously-generated candidates and updates of a catalog of candidate orbits;
3) generation of new orbit candidates (association hypotheses) from the remaining UCTs; and
4) observation correlation against the newly created candidates.

While critical to the process, the first processing stage, correlation against the catalog, is not considered part of the SADIE tool. From a functional and software perspective, stages 1, 2, and 4 are very similar, but can operate on different data sets with different correlation criteria and gating. The architecture was developed to allow alternate approaches to correlation and observation or track associations to be implemented and enable continued evolution and improvement of capabilities.

SADIE builds upon the integration of Covariance-Based Track Association (CBTA) and optical candidate generation modules, Satellite Identification (SID), and Parallel Catalog (ParCat). SID and ParCat are based on operational legacy codes, and are used for correlating observations with known orbits and orbit estimation, respectively. Various utilities and other software modules have been integrated and developed to enable SADIE to execute in a closed-loop fashion. SADIE is designed to utilize shell and perl scripts for building and for run-time execution, with most of the source code written in Fortran and C. Parallel execution is supported by either MPI or PVM. All codes were developed and are hosted on the Riptide machine at Maui High Performance Computing Center (MHPCC).

Observation Correlation with SP SID

SID is a fixed-gate observation correlation tool originally developed for use at Naval Space Command, now DSC2-Dahlgren. It uses fixed correlation gates projected into the orbit radial, transverse and normal directions. As part of the SADIE project, the ability to use SP vectors and ephemeris tables was added to SID, which is now referred to as SP SID. For deep space satellites, we expected that the SP predictions would allow for correlation gates an order-of-magnitude smaller than gates used for GP predictions. A similar reduction would be expected for low Earth orbits, mainly due to better drag and gravity modeling in the SP predictions. The utilization of smaller correlation gates decreases the likelihood of false-positive correlations (e.g., falsely identifying object B as object A), a consideration that is especially important in densely populated orbit regimes. New baseline SP-based tolerances were empirically derived for SP SID.

SP SID performance results using simulated Space-Based Space Surveillance (SBSS) and SST data test cases are presented here. Over 4000 deep space objects were chosen from the predicted NASA Debris Catalog and propagated for 10 days using an SP propagator. Some of the objects used full 6-degree-of-freedom SP propagation, while others had either prescribed attitude profiles or simple spin rates around a fixed axis. We incorporated attitude motion to model radiation pressure effects for the simulation truth-reference orbits, which are more realistic than those for orbits estimated with the spherical-satellite model used for the catalog maintenance. Canonical models, such as box and wing, cylinders, flat-plates, and random combinations of plates were used to model the space object population. Random body-component variations were added to the objects to provide a diverse population. The size distribution was kept consistent with the debris catalog. Right ascension and declination were generated using the geometric line of sight from the observer to the satellite, accounting for light time. Probability of detection was calculated based upon our available understanding of sensor performance. Gaussian random noise with 1 arcsecond standard deviation was added to the data. The simulations were based on our best understanding of the measurement characteristics and operating modes of these sensors.

An initial SP catalog was created by performing differential corrections to the first nine days of data, assuming perfect correlation for 2000 of the objects. For the SP catalog, the objects were assumed to be spherical, and a solar radiation pressure coefficient was estimated. The catalogs were then fed into SP SID for correlating all of the observations on the tenth day, using various gate sizes. This resulted in 92976 SBSS and 21807 SST observation pairs, respectively. SP SID was run using 0.255 x 0.510 x 2.232 km (radial, transverse, and normal, respectively) gates, scaled by factors of 10 to 50. By this means, we were able to study the frequency of occurrence of both false-positive and false-negative associations as a function of the correlation gate sizes. Fig. 2 summarizes the results.

The basic trend is easy to understand. As we increase the nominal gate size, the number of false-positives (observations associated with the wrong cataloged orbit) must increase, on average, and the number of false-negatives (observations falsely labeled as UCTs when they belong to a cataloged orbit) must decrease, again on average. These opposing effects necessarily result in a local minimum in the number of incorrectly associated observations. In this case, the optimum occurs close to a gate-scale factor of 30, and at this value we can correctly correlate approximately 98% of the observation data for this dense population.

In order to process large numbers of observations against a large catalog in a timely manner, SP SID has been parallelized to match the associated ParCat catalog update step. Fig. 3 plots the scaling results (the run-time vs. the number of processors) for a combined SP SID and ParCat run for data set described above. In these results, ten days of data are processed in one-day batches: one day of data is correlated and then the catalog is updated, repeat with the next day's data. The results were generated on the MHPCC Riptide system. As we go from 3 to 12 compute-nodes, the associated run-time drops from approximately 54 to 16 minutes. This is a very acceptable factor of 3.4, compared to a factor of 4 in an ideal speed-up. In operational practice when only real-time processing is required, SID will not require HPC; however, in the R&D stages, the ability to process large data sets much faster than real-time was very important.

The above results did not include any multi-tag resolution capability. Multi-tags can occur when multiple simultaneous observations from the same sensor correlate against a single object; a likely occurrence in dense populations such as that observed by SST, or in the less common scenario of a single observation correlating against multiple objects. Initial testing has shown that the multi-tag resolution capability can reduce false correlations below the 1% level.

SADIE Candidate Generation from Optical Data

Six pieces of information are needed to describe a simple orbit, be it orbital elements, position and velocity, or some other form. A single optical line of sight only provides two pieces of information. In optical track association, one has three choices: 1) hypothesize the additional four pieces of information to form an orbit candidate, and then correlate additional data to the hypothesized candidate; 2) hypothesize additional pieces of information (like a range) for each of the optical data, and then form hypothesized orbit candidates by combining pairs of observations; or 3) form combinations of three pairs of optical observations to form an initial orbit candidate. Each results in a significant number of hypotheses to be evaluated. After exploring option 2, SADIE ultimately incorporated the third approach. If a fourth temporally-separated optical observation is combined with the three previously identified pairs, a simple orbit determination can test whether that hypothesis is a viable orbit.

This angles-hypothesis approach is depicted in Fig. 4. Consider a sensor that surveys a region of space four times. A hypothesis is generated for every combination of detections from each of those four surveillance sweeps. For example, in Fig. 4, Hypothesis 1111 takes the first detection from each data set, and Hypothesis 12n2 takes detection 1 from the first set, 2 from the second, nth from the third, and 2 from the fourth. Orbit determination can be performed for each hypothesis and if all observations are not utilized, or fit residuals are inconsistent with the sensor performance, that hypothesis can be pruned. If the fit to the data is good, then the hypothesis is retained as a candidate orbit. This approach "connects the dots" in angles space. While the approach initially utilizes a very large number of hypotheses, the hypotheses can be quickly pruned.

In practice, optical data usually comes with more than one observation per track and, therefore, even a three-track hypothesis is over-constrained enough to prune a large number of hypotheses. If sufficient angle-rate information is available, one can use rate values to eliminate hypotheses, but that is not incorporated into SADIE at present. The current approach forms three-track hypotheses, and then searches for a fourth track for only the successful three-track combinations. The computational burden is still high: N3 scaling with respect to observations is not computationally attractive, even with supercomputing resources available. There are ways to mitigate this load, such as binning the observations by time and geometry, but it is still a computationally-challenging approach.

The angles-hypothesis approach was tested with the same large SST simulated data set that was described above in the SID discussion. Only one day of the simulated data was processed, but the data set consisted of 3493 tracks produced by 760 uncataloged objects. Of those objects, 405 had at least four tracks, meeting the current SADIE criteria for generating a candidate from angles-only data. Given those observations, and restricting candidate generation to 50-degree hour-angle bins, the angles-only hypothesis approach produced 391 correct 4+ track candidates with an additional 69 incorrect candidates. These results are summarized in Figure 5.

Inspection of the results revealed that most of the false candidates resulted from multiple combinations of one erroneous track with the three tracks that belonged to a single object. In this case, it happens that the one erroneous track was within a handful of arcseconds of the true trajectory.

Next, four consecutive nights of real SST data provided by Lincoln Laboratory for testing were processed. Restricting the data set to objects believed to be free from mis-tags reduced the number of objects in the data set to 51 and the

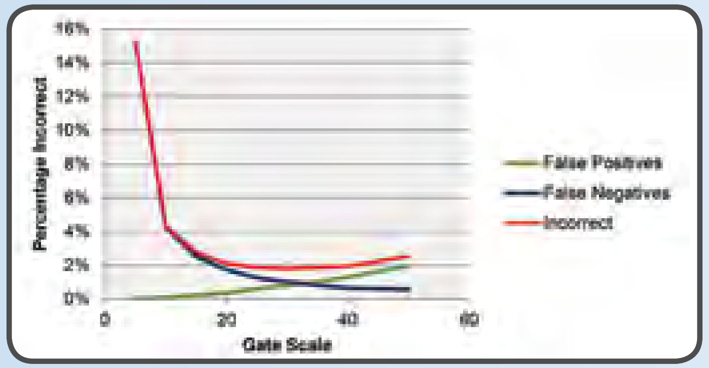

Figure 2: SP-SID Results with Simulated SST and SBSS Data

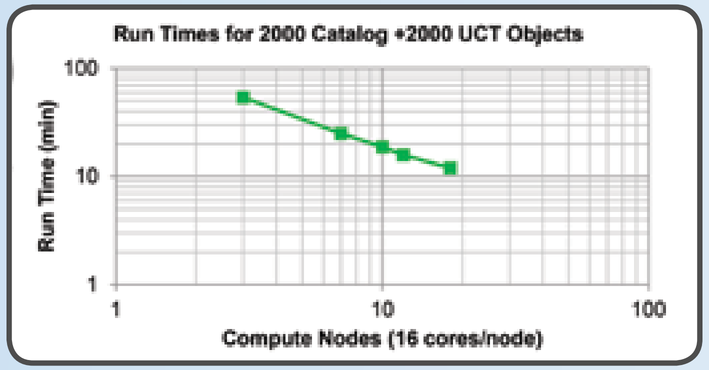

Figure 3: SP SID and ParCat Scalability Results for Simulated SBSS and SST Data

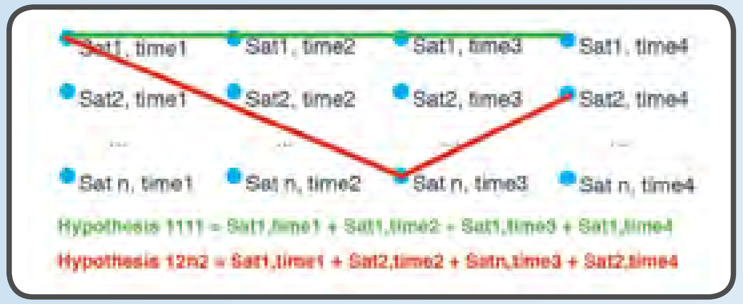

Figure 4: SADIE Candidate Generation Approach for Optical Processing

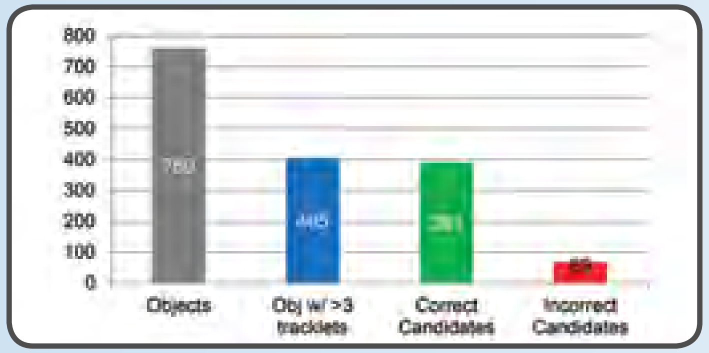

Figure 5: Angles-Only Hypothesis Approach Simulated Data Results Summary

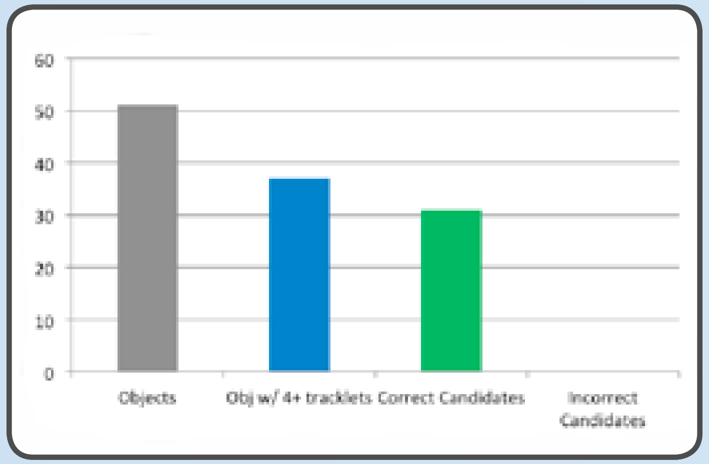

Figure 6: SADIE Real-Data Results Summary

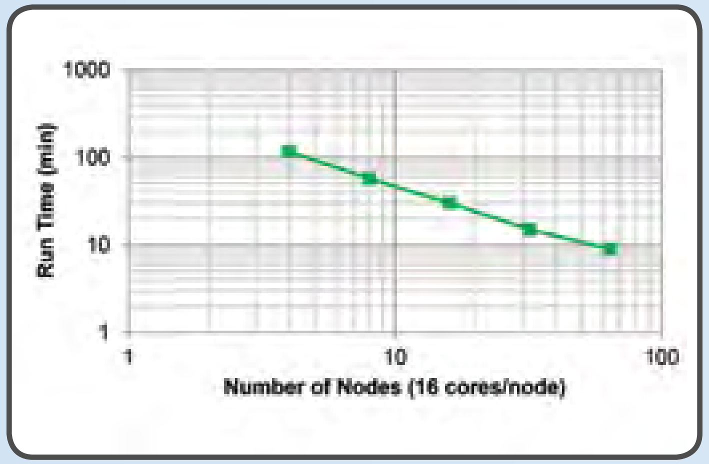

Figure 7: Scalability of the SADIE Angles-Only Candidate Generation Prototype

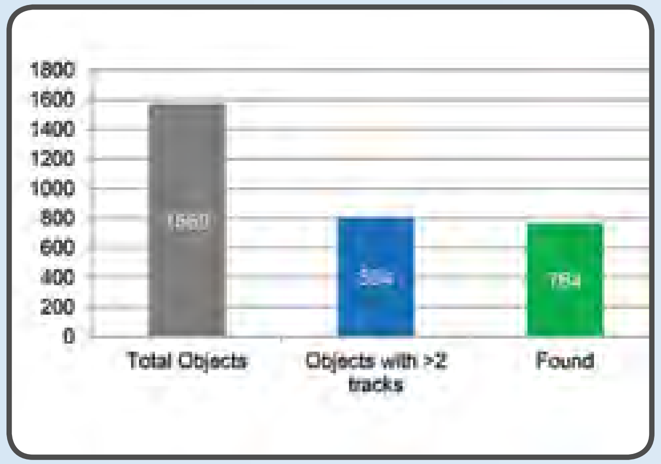

Figure 8: SADIE CBTA Real-Radar-Data Results

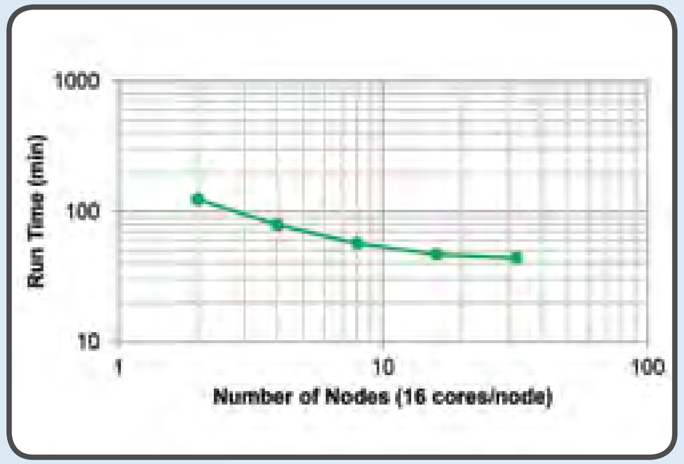

Figure 9: Scalability of the SADIE CBTA Module

number having 4 or more tracks during the data span to 37. The SADIE angles-only hypothesis approach recovered 31 of the 37 4+ track objects with no false candidates. The lack of false candidates is not surprising since this is a very sparse data set in comparison to the simulated data described above. These results are summarized in Fig. 6. Analysis is needed to determine why some of the candidates were missed but, as with the simulated data results, the angles-hypothesis approach turns out to be very effective.

The results presented here have been generated with a prototype version of the SADIE angles-only candidate generation approach, and the scalability of that code has been tested with a subset of the simulated SST data set. The results are plotted in Fig. 7 and reveal near-ideal scaling. This means that as we double the number of processors, we halve the run-time. Since the formation and testing of one hypothesis is independent of the handling of any other hypothesis, the processing can be efficiently split among many processors. The SADIE implementation of this approach utilizes processors in a similar way and, while ideal speed-up is not expected, scaling similar to SID and ParCat, as shown in Fig. 3, is expected.

SADIE Candidate Generation from RADAR Tracks

If the error probability distributions are Gaussian, it has been proven that the probability of correct association in the single-target case is constant in time when a Mahalanobis-type distance (chi-squared statistic) is used as the association metric. Although Gaussian assumptions do not hold rigorously for orbital motion or typical space surveillance sensor systems, they are often close enough to allow for an effective CBTA approach. As a slight improvement to Mahalanobis distance, a likelihood ratio developed by the Numerica Corporation was introduced into the SADIE CBTA module.

SADIE's radar track association was tested with a reasonably large set of real-space surveillance network radar tracks from 2004. This data has been especially well scrutinized, so it is as close as we can get at present to real-data "truth-reference". Fig. 8 plots results in terms of objects and tracks for the total number of objects, the number of objects with three or more tracks (our candidate generation criterion), and the number of objects that CBTA "found".

Here, the number of objects "found" by CBTA includes only those that both satisfied the association criterion and resulted in a successful ParCat SP orbit solution with the specified number of tracks. Occasionally, an object is "found" as 2 or more distinct orbit solutions, and the number-of-tracks criterion is not sufficient since the tracks are split among two candidates. For instance, if an object had four tracks, CBTA may have recovered two two-track candidates, neither of which count as a "found" 3+ track object. In other words, the track cohesion for the candidate orbits is less than ideal. Of the 38 objects with more than two tracks missed by CBTA, 18 had one or more two-track combinations.

CBTA relies heavily upon ParCat, which is parallelized, and has implemented an orbit element clique-based approach to separate incoming tracks into a scalable track association architecture. However, there are still several serial steps in the process, which limits scalability. Figure 9 plots the execution time results as a function of processing nodes for the 2004 data set on the MHPCC Riptide system. From the plot, one can see that going from two to four nodes results in a 1.5x decrease in run-time, which is about 75% of ideal, and going from two to eight nodes results in a 2.2x reduction in run-time, which is only 54% of ideal. (Eight processing nodes is 128 processing cores.) Despite non-ideal scaling, the run-times are still on the order of an hour on our machine for this fairly large data set. Scaling improvements could surely be realized, but have not been judged necessary for SADIE's projected use.

Summary

In summary, SADIE has been developed jointly by the Air Force and Naval Research Laboratories under the HSAI-SSA to resolve uncorrelated tracks automatically, and to build a more complete space object catalog. Enabled by HPC, SADIE has been shown to recover almost all potential orbit candidates from real and simulated data, and is complemented by SID's improved ability to correlate observations against known objects. Together, they have overcome one of the long-standing challenges to space surveillance.

Significant contributions to the SADIE program have come from Jason Addison, Terry Alfriend, Kathy Borelli, Shannon Coffey, Bruce Duncan, Keric Hill, Aaron Hoskins, Kevin Roe, Alan Segarman, and Paul Schumacher. Vicki Soo Hoo and Adam Mallo helped provide simulation data and SST data were provided by Lincoln Laboratory with the permission of DARPA. Sponsorship of the SADIE development effort was provided by the AFRL Directed Energy Directorate, DoD High Performance Computing Modernization Program office, DARPA, and the AFRL Rapid Innovations Program. Betty Duncan of MHPCC was instrumental in the preparation of this article. We acknowledge and thank the Air Force Research Laboratory and the Naval Center for Space Technology at the Naval Research Laboratory for their continuing support of this work. We greatly appreciate the computational support provided by the Maui High Performance Computing Center and the DoD High Performance Computing Modernization Program office.

Naval Research Laboratory Meteorologists Receive Naval Technology Achievement Award

Daniel Parry
US Naval Research Laboratory Public Affairs

"This work was performed on HPC systems at the Navy DSRC"

(from the left)
Rear Adm. Mat Winter, Chief of Naval Research, presents the Dr. Arthur E. Bisson Award to, Dr. Patrick Reinecke, Dr. Melinda Peng, Dr. Yi Jin, Dr. Hao Jin, Dr. Eric Hendricks, Dr. James Doyle, and Dr. Sue Chen,

Photo: Courtesy of US Naval Research Laboratory, photo by John F. Williams

A team comprised of eleven research meteorologists from the US Naval Research Laboratory (NRL), Marine Meteorology Division, received the Dr. Arthur E. Bisson Prize for Naval Technology Achievement at a ceremony hosted by the Office of Naval Research (ONR), on 26 August 2015.

The award honors the team for expertise and innovative scientific work resulting in the rapid development, from basic research to transition to operations, of an innovative and versatile Numerical Weather Prediction (NWP) system, significantly improving the prediction of tropical cyclones (TC) - one of the most significant threats to DoD operations in the tropical and mid-latitude ocean areas around the world.

The team consisting of Dr. James D. Doyle, team lead and Mesoscale Modeling Section Head at NRL, and meteorologists Drs. Sue Chen, Eric Hendricks, Richard Hodur, Teddy Holt, Hao Jin, Yi Jin, Jonathan Moskaitis, Melinda Peng, Patrick Reinecke, and Shouping Wang were named for their achievements in contributing to the improved safety of Navy personnel, DoD assets, and the broader civilian population in coastal regions through their development of the Coupled Ocean/Atmosphere Mesoscale Prediction System for Tropical Cyclones (COAMPS-TCTM).

Doyle and his entire COAMPS-TC team vigorously engaged the Navy's Meteorology and Oceanography (METOC) and warfighting leadership, as well as the broader civilian meteorological community to inform the development and advancement of COAMPS-TC into a leading tropical cyclone model. Their coordination with a variety of Navy, DoD and interagency collaborators, including the National Oceanic and Atmospheric Administration (NOAA), National Science Foundation (NSF), and NASA resulted in the expanded use of COAMPS-TC in the research and operational forecasting communities.

Currently, the Joint Typhoon Warning Center (JTWC) and National Hurricane Center (NHC) use COAMPS-TC for forecast guidance for the wind speed intensity of tropical cyclones (TC), and to forecast their speed and direction of movement. Because COAMPS-TC showed significant promise in predictive skill, both JTWC and NHC incorporated its products into their official 'Consensus' forecasts in 2012, well before the official model transition date. Improvements are continuing to be made to COAMPS-TC that will ultimately provide more accurate guidance for DoD and US government forecasters.

Increasingly sophisticated developmental versions of COAMPS-TC will continue to be transitioned to Navy operations in support of the Joint Typhoon Warning Center and the National Hurricane Center. A key additional enhancement will be a fully-coupled ocean-atmosphere version in which the NRL Costal Ocean Model (NCOM) and the Wave Watch III (WWIII) model will provide the ocean circulation and wave components, respectively.

Rigorous testing using the Navy DoD Supercomputing Resource Center (DSRC) has demonstrated the modeling system to be among the finest and most accurate tools available to DoD and US civilian forecasters. COAMPS-TC has garnered numerous awards and accolades for its predictive skill and technical features during its development and since its transition to operations in June 2013.

In June 2011, COAMPS-TC was one of nine worldwide winners of the inaugural High Performance Computing (HPC) Excellence Award presented at the ISC-11 International Supercomputing Conference in Hamburg, Germany - an award presented annually to recognize noteworthy achievements by users of HPC technologies. As a result, COAMPS-TC was recognized for achieving a significantly improved model for tropical cyclone forecasting. COAMPS-TC development benefited significantly from the DoD HPCMP computational assets at the Navy DoD Supercomputing Resource Center (DSRC) at Mississippi's Stennis Space Center.

The COAMPS-TC project received numerous DoD HPCMP Challenge Awards during its development, due to its innovative technology and overwhelmingly potential benefit to the Navy and DoD. Real-time development of COAMPS-TC and its support of Navy exercises and operations were the impetus behind the Marine Meteorology Division being awarded a Cray XE6m supercomputer through an HPCMP Dedicated HPC Project Investment (DHPI) grant in 2012, and the inaugural HPCMP 'Pathfinder Project' in 2014.

The advancement in TC intensity forecasts with COAMPS-TC are based on the long-term science and technology (S&T) investment in mesoscale processes, and model development from the NRL base program and the ONR Marine Meteorology Program. The understanding of tropical cyclone dynamics has been accelerated in recent years through several ONR-supported field observation campaigns that include Coupled Boundary-Layers Air-Sea Transfer (CBLAST), TCS-08, ITOP-10 and Tropical Cyclone Intensity (TCI). The final technical push of COAMPS-TC model development came from a Rapid Transition Program project jointly supported by ONR and the Oceanographer of the Navy through PEO C4I&Space PMW-120. Advancements of COAMPS-TC and real-time demonstrations have also been supported through NOAA's Hurricane Forecast Improvement Project (HFIP).

The Bisson Prize is named in honor of the late Dr. Arthur E. Bisson (1940-1996), who provided a model of principled, effective leadership in transitioning S&T to naval capabilities. In his last assignment, Dr. Bisson was Director of Science and Technology for ONR. He was a prime mover in the integration of all naval S&T in a single command, capable of managing new technology from earliest scientific concepts through prototyping and manufacture. His integrated program provided the Navy with a wholly new paradigm for faster, better coordinated, and more predictable S&T transition to acquisition and operations.

For more news from Naval Research Laboratory, visit www.nrl.navy.mil or www.navy.mil/local/nrl/.

HPCMP CREATE™

Tools Continue to Impact DoD Acquistion Programs

Scott M. Sundt - Senior Operations Director, HPCMP CREATE™ Program
Dr. Douglass Post - Associate Director, HPCMP CREATE™ Program
Dr. Robert Meakin - Program Manager, CREATE-AV
Dr. Richard Vogelsong - Program Manager, CREATE-Ships
Dr. John D'Angelo - Program Manager, CREATE-RF
Dr. Larry Lynch - Program Manager, CREATE-GV
Dr. Saikat Dey - Program Manager, CREATE-MG

INTRODUCTION — MAINTAINING OUR TECHNICAL EDGE IN A RAPIDLY CHANGING WORLD

We live in a globalized, highly interconnected world that is in the process of a major technological revolution in warfare arenas such as air, maritime, land, space, cyber, and the electromagnetic spectrum. Our technological superiority is a key element to our national security, both at home as well as abroad; and our ability to design, procure, and field new equipment is critical to that security. How does the DoD maintain that security in the face of shrinking budgets and closing gaps in our technological superiority while meeting the challenges of a dynamically changing world? How can the DoD speed new and innovative equipment to our Armed Forces that will provide "game-changer" capabilities on the battlefield against potential adversaries? Things like hypersonic aircraft directed-energy weapons, autonomous vehicles, new ship and submarine designs, improved ground combat vehicles—where can the DoD conduct the research and development, testing and evaluation, and modeling and simulation needed to field these new technologies in a timely and cost-effective manner?

This is primary focus of the Computational Research Engineering Acquisition Tools and Environments (CREATE) program, which is under the oversight of the High Performance Computing Modernization Program (HPCMP). The HPCMP CREATE™ program provides critical modeling support in the world of computational prototyping, and is maintaining a steady pace of adoption and acceptance of its physics-based engineering software tools across an expanding customer/client base. It is HPCMP's premier vehicle for addressing the DoD's current and future design and analysis efforts for its major acquisition programs. HPCMP CREATE™ provides innovative applications of its software tools for optimizing aircraft, ship, ground vehicle, and radar antenna designs.

As of May 2016, the HPCMP CREATE™ program has expanded the acceptance, use and adoption of its various physics-based software tools to 129 organizations comprised of the DoD, industry, academia, and other government agencies. The success of the HPCMP CREATE™ program continues to grow with each new software release, and its products are now becoming an integral part of major defense acquisition programs in their design space exploration and validation efforts.

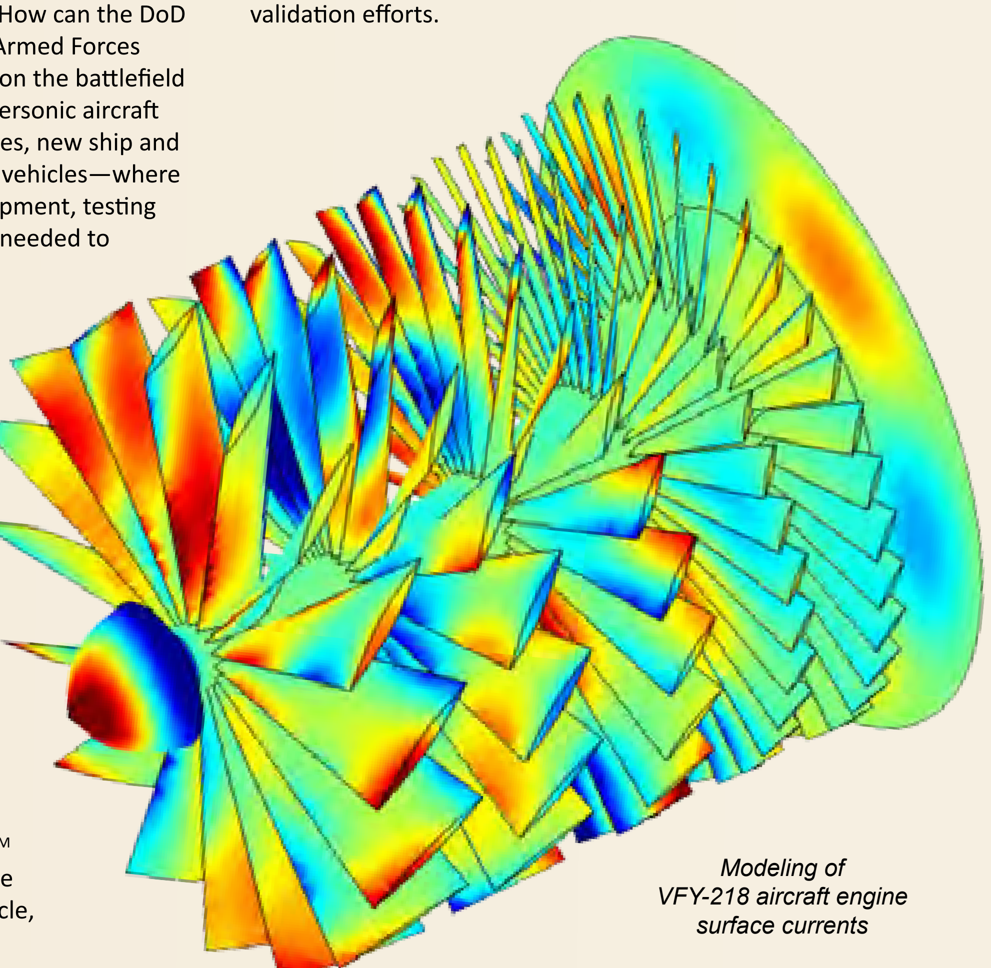

Modeling of VFY-218 aircraft engine surface currents

HPCMP CREATE™ — TOOLS CONTINUE TO DEVELOP — CREATE-GV JOINS THE GROUP

The HPCMP CREATE™ program is comprised of five distinct major defense weapons systems procurement project areas; CREATE-AV (aviation), CREATE-Ships (shipbuilding), CREATE-RF (antenna design), CREATE-MG (meshing and geometry), and the newest addition to the HPCMP CREATE™ family of products, CREATE-GV which incorporates physics-based engineering software tools for design and analysis of ground vehicles for the Army and US Marine Corps. Within each of these project areas, the CREATE program has developed a total of 11 robust physics-based engineering software tools that perform critical design space exploration and assessments, using the high-speed computing systems that are found at the five HPCMP DoD Supercomputing Resource Centers (DSRCs) located around the country at various government research laboratories. The critical success of the HPCMP CREATE™ program is the result of a solid business plan, an innovative hybrid management system that embeds the software developers within the Service research and development organizations, and a steady annual delivery of product development and releases that are specifically focused on meeting the needs of the Service R&D clients.

In early 2016, the CREATE-GV (ground vehicles) program was established to "develop physics-based, high-performance computing (HPC) tools to enhance ground vehicle concept development, inform requirements development and provide requisite data for trade-space analysis to positively impact cost, schedule and performance with significant reduction in design risk for the acquisition community"[1]. As its initial project charter, the CREATE-GV program is developing two software tools: (1) Mercury, which is an HPC physics-based tool for tactical mobility with co-simulation of terrain mechanics with vehicle systems and components including suspension, tire and track, soil modeling and powertrain simulation, and (2) a Mobility Analysis Tool (MAT), which is a model-based analysis software tool used for predicting tactical mobility for a wide variety of terrain conditions, incorporating many factors such as soil condition, vehicle performance and configuration, vegetation density, average surface roughness, average slope, and other factors.

These CREATE-GV tools have already had a significant impact in reducing design flaws while allowing for multiple design options to be tested at a fraction of the cost of traditional design, build, and test development systems. To date, the Mercury program, under the overarching Engineering Resilient Systems (ERS) concept, is in the process of completing a trade-space analysis pilot project of the Army's Light Reconnaissance Vehicle (LRV), in which 65,000 unique configurations were analyzed for five key mobility performance parameters. Likewise, the MAT program has completed a baseline version that supports optimized trade-space analysis by allowing thousands of design iterations to be tested and synchronized with Mercury program inputs.

CREATE-GV : Modeling lane changing analysis of Light Reconnaissance Vechicle (LRV)

HPCMP CREATE™ — TOOLS ADDING VALUE, GROWING ADOPTION

While the CREATE-GV program has had its early successes, the more established HPCMP CREATE™ projects have continued their steady progress of greater exposure and adoption of their tool sets in the design and analysis of major weapons systems. During the recent CREATE Developers' Review 2016 that was held in May, each CREATE project provided a series of success stories that highlighted the application and impact of their physics-based software tools in major Defense acquisition programs over the past year.

CREATE-AV highlighted the successes of their three engineering software tools—Helios (which is a full [rotary-wing] aircraft design analysis and testing tool that uses high-fidelity, multi-disciplinary, physics-based simulation), Kestrel (which is a full [fixed-wing] design and analysis testing tool similar to Helios), and DaVinci (which is a model-centric conceptual design tool). These highly successful engineering software tools have now become an integral part of design application analysis for Navy and Air Force multi-billion dollar acquisition programs. Recent impacts with Kestrel include assessing the effectiveness of static forces and moments on the F/A-18E Super Hornet, which enabled the ability to assess effectiveness of proposed modifications to the flight control system. Other impacts with Kestrel include life cycle support for the Air Force A-10 Warthog, and database development for the Navy's E-2D Advanced Hawkeye all-weather, carrier-capable tactical airborne early warning aircraft for improved flight simulator performance assessments. Helios has had similar impact success with the Army's CH-47F Chinook heavy-lift helicopter early design stage predictions, and the H-60 Sikorsky tail rotor effectiveness testing.

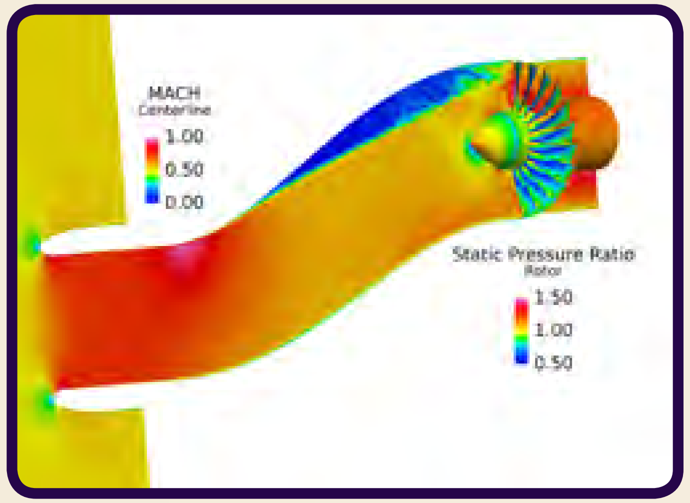

CREATE-AV
Mach wave pressure modeling of military jet engine fan

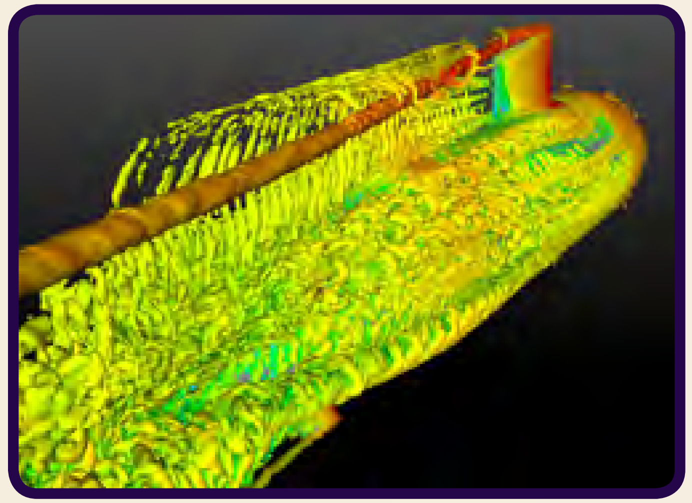

CREATE-Ships
Simulation of submarine at drift angle”

The CREATE-Ships project is currently fielding four specific physics-based engineering software programs; RSDE (Rapid Ship Design Environment), NESM (Navy Enhanced Sierra Mechanics), NavyFOAM, and IHDE (Integrated Hydrodynamics Design Environment). Each software tool has been heavily involved in recent major Navy T&E assessments, as well as major shipbuilding acquisition program design analysis efforts.

The Rapid Ship Design Environment, or RSDE, tool is a rapid development assessment and integration of candidate ship designs to avoid costs vs. capability mismatches. It allows users to comprehensively explore alternate design solutions while there is still a maximum range of options available through detailed, physics-based and HPC-based analysis early on in the design cycle. Most recently, RSDE supported the LX(R) Analysis-of-Alternatives study for the Navy’s next amphibious assault ship through design space exploration of 22,000 concept designs in only three months, as well as identifying major cost vs. capability trades.

The NESM program provides a massively-parallel, enhanced, physics-based modeling and simulation suite for prediction of ship shock response and damage due to weapon engagements. This critical analysis tool has been extensively used as a Full Ship Shock Trial (FSST) alternative and Live-Fire Test and Evaluation (LFT&E) support for design analysis on the Navy’s Littoral Combat Ship (LCS) and the Navy’s CVN-78/79 nuclear aircraft carrier programs.

While NESM is providing shock and damage analysis on Navy ship designs, the NavyFOAM program is providing high-fidelity hydrodynamics to accelerate and improve all stages of ship hydrodynamic design for surface ships and submarines. These design elements include the impact of seaway loads, seakeeping, resistance, and powering loads on various ship design models. Most recently, the NavyFOAM program has supported the Ohio Replacement Program, using a custom physical model for flow predictions and rotating-arm simulations to better understand the underlying physics used in design decisions. The DDG-1000 Zumwalt guided missile destroyer used NavyFOAM to determine hull forces, and provide related hull maneuvering coefficients to support safe operating envelope design decisions. Additionally, NavyFOAM supported both the Marine Corps Amphibious Combat Vehicle programs, as well as the R&D efforts for future high-speed multi-hull vessel optimization design.

The CREATE-Ships Integrated Hydrodynamics Design Environment, or IHDE, is a user interface that allows naval architects to access Navy hydro design codes and the Navy model database (LEAPS—Leading-Edge Architecture Prototyping System). The IHDE system provides the Navy with a suite of analysis methods that can be used to impact ship designs for fast turnaround needs, as well as providing an integrated user design environment for using different levels of fidelity methods, in both the design and analysis domains, for simultaneously evaluating different ship elements (e.g., resistance, powering, maneuvering, seakeeping, etc.). Most recent success impacts for IHDE have involved the DDG 51 Flight III bow bulb design assessments, and support to the Navy’s Small Surface Combatant Task Force (SSCTF) in evaluating multiple ship designs in a very short timeframe.

The CREATE-RF project revolves around a single physics-based engineering tool known as SENTRi. SENTRi is an electromagnetic analysis and design tool that uses high-fidelity, first-principle methods and physics-based simulation for a variety of classified radar and electronic warfare systems. While the majority of the recent stories for SENTRi reside in classified defense programs, areas such as High-Power Microwave modeling and Red Target high-fidelity modeling for potential adversary (red) targets have made important impacts with a host of government and industry organizations.

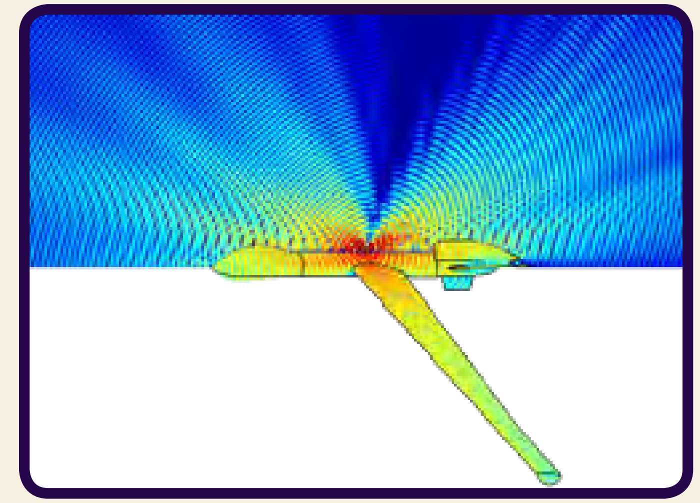

CREATE-RF
RF Gain Modeling of an Unmanned Aerial Vehicle (UAV)

CREATE-MG
Boundary-layer meshes for flow around model aircraft

The final CREATE project is the CAPSTONE tool, which resides in the CREATE Meshing and Geometry (MG) project. CAPSTONE is a physics-based interface tool that improves the ease, speed, flexibility, and quality of geometry and mesh generation, and enables the generation of CAD-neutral digital representations and product models of weapons systems and platforms, as well as operational terrains and environments. As such, CAPSTONE is used by all of the other CREATE tools to generate the visual and physical design attributes of an aircraft, ship, vehicle, or antenna. CAPSTONE has very robust anisotropic (having a physical property that has a different value when measured in different directions) unstructured surface and volume meshing capabilities for complex geometries. Because of the wide usage of CAPSTONE in every CREATE engineering tool, it has had significant impact on air vehicle, ground vehicle, and ship design applications. Notable success stories for CAPSTONE include support for the Aerostar, F/A-18E/F, T-45 and P-8 aircraft design analysis, submarine mesh design analysis for the Ohio Replacement Program, and GPS antenna integration evaluations.

Although the impacts and successes of the HPCMP CREATE™ program are grounded in the support these physics-based engineering tools have provided to the DoD Services on major defense acquisition programs and R&D efforts, their growing adoption and acceptance within the wider computational prototyping community has been acknowledged through significant professional recognition. In January of 2016, 26 papers from the HPCMP CREATE™ program were accepted at the American Institute of Aeronautics and Astronautics (AIAA) annual Science and Technology Conference. At the recent American Society of Naval Engineers/Office of Naval Research Conference in March of 2016, the HPCMP CREATE™ program was asked to give five presentations on various CREATE computational prototyping software development projects, and one of its engineers won the ASNE 2015 Solberg Engineering Award for sustained engineering excellence. Additionally, the HPCMP CREATE™ program recently published six articles in the IEEE/AIP publication “Computing in Science and Engineering”. The HPCMP CREATE™ program continues to gain recognition participation in major professional conferences such as the upcoming Supercomputing 2016 conference this fall, as well as several National Defense Industry Association (NDIA) conferences on systems engineering and modeling and simulation.

CONCLUSION — THE WAY AHEAD

The future of the HPCMP CREATE™ program is focused on building upon its steady success of high-fidelity, physics-based engineering tools that meet the needs of the Defense R&D community and the warfighter in all warfare domain battlespaces in a timely and cost effective manner. The necessity of being able to speed-up the acquisition cycle, while providing state-of-the-art weapons systems at a reduced cost, is critical to our national security and protecting our allies and interests abroad. The combination of physics-based computational prototyping provided by HPCMP CREATE™ tools, coupled with high-speed supercomputing capability, are powerful examples of leading technology that is starting to expand their usage and acceptance throughout the defense and civilian industry.

New technology areas such as hypersonics, directed-energy, future submarine design, multi-hull ship design, and unmanned vehicle design and analysis are all prime candidates for future HPCMP CREATE™ tool applications. As the capability and capacity of these hardware and software systems progress from the current PetaFLOP domain to the ExaFLOP domain in the next 30 years, computational prototyping, such as the products being developed by the HPCMP CREATE™ program and those adopted by the OSD’s Engineering Resilient Systems (ERS) and Air Force’s Digital Thread/Digital Twin programs, will become an integral part of all major DoD acquisition programs. The DoD’s technological future has much to look forward to.

PETTT Team Wins Best Poster Award at SC15

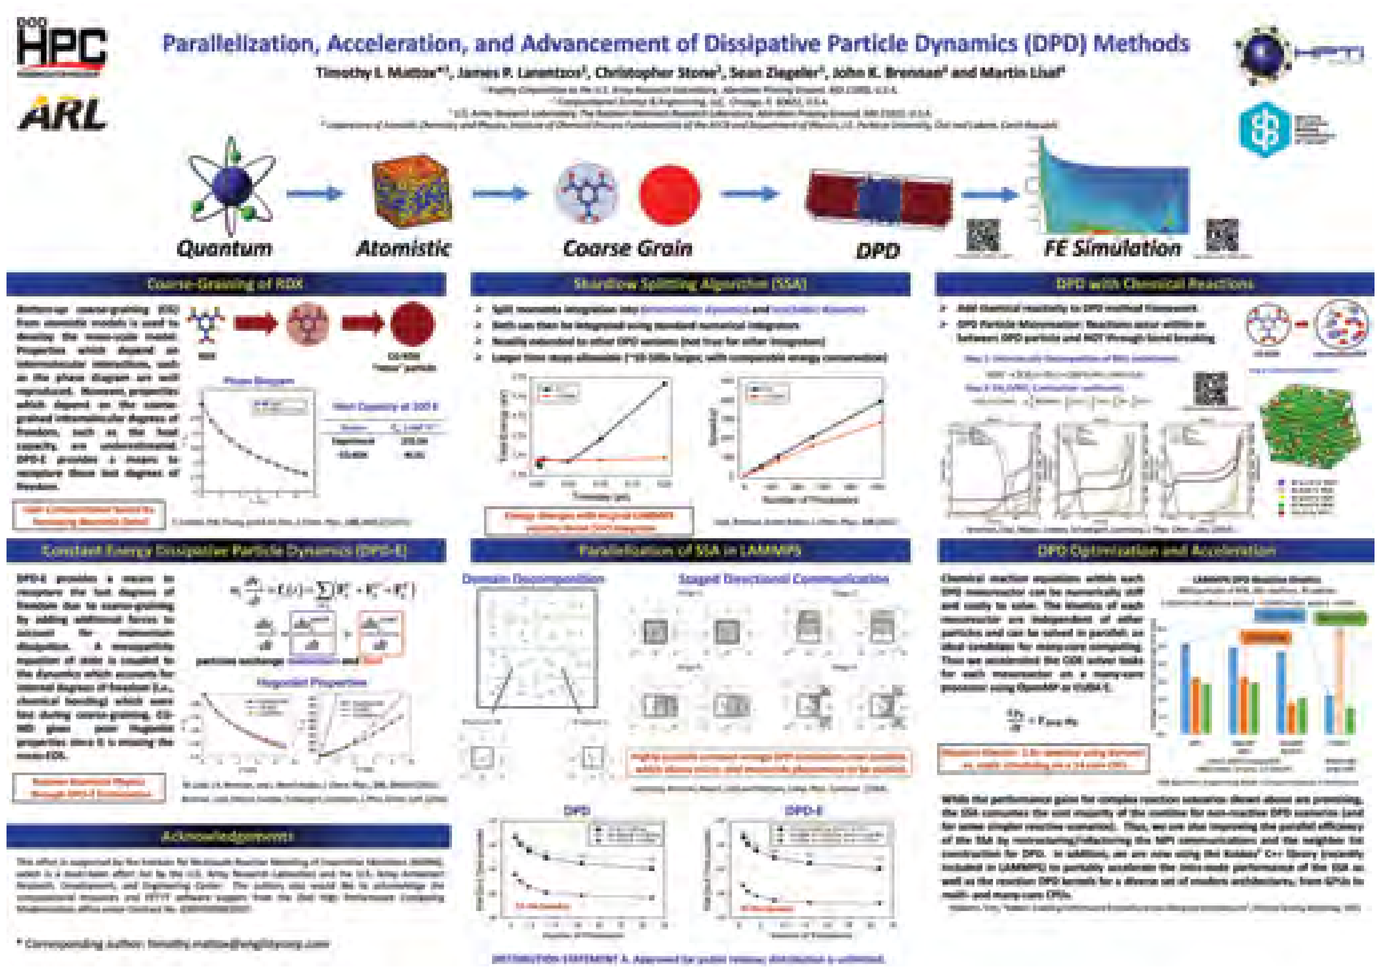

To see the winning poster, go to:
http://tinyurl.com/petttsc15, or visit our website at www.hpc.mil

The HPCMP congratulates the HPCMP PETTT Team whose poster, "Parallelization, Acceleration, and Advancement of Dissipative Particle Dynamics (DPD) Methods" won the Best Poster award at SC15, the 27th annual International Conference for High Performance Computing, Networking, Storage and Analysis. Each year, the SC Posters Committee identifies one poster as the best poster presented as part of the conference's technical program.

This year's winning HPCMP Team is comprised of Drs. Timothy Mattox, James Larentzos, and Sean Ziegler (DoD High Performance Modernization Program); Dr. Christopher Stone (Computational Science & Engineering, LLC); Dr. John Brennan (US Army Research Laboratory); and Dr. Martin Lísal (Institute of Chemical Process Fundamentals, J.E. Purkinje University). The award, accepted by PETTT senior computational scientist Dr. Timothy Mattox, was announced at the SC15 awards ceremony in Austin, TX, on November 19.

The award-winning poster highlighted ongoing efforts within the ARL Multiscale Response of Energetic Materials (MREM) research program, and continued PETTT support in transitioning novel DPD methods within the LAMMPS (Large-scale Atomic/Molecular Massively-Parallel Simulator) software package for efficient use of modern and next-generation HPC architectures. The team implemented micro- and mesoscale modeling capabilities that use coarse-grain models that are up-scaled from quantum-based reactive models. These novel modeling tools are routinely used to accurately simulate reactive chemistry of an energetic material under shock or thermal insult, and investigate the effect of material microstructure on the multiscale response of crystalline energetic materials, such as RDX , the current target system within the MREM research program. Through efficient use of HPC, the DPD simulations are now conducted at significantly longer length-and-time scales than previously possible, enabling microscale simulations that were heretofore inaccessible. This, in turn, has allowed for a novel concurrent multiscale simulation methodology in which first-principles-based microscale simulations drive a continuum scale simulation. Before this effort, such continuum scale simulations were driven by conventional, empirically-derived continuum-level material models. Moreover, these advanced DPD methods have been shown to be applicable to a wide set of materials, and are now being transitioned for use in other ARL computational research programs. Within ARL's MREM program, the development of these DPD computational tools is a multi-disciplinary and multi-person effort involving ARL scientists Dr. Sergey Izvekov, Dr. Michael Sellers, Dr. Brian Barnes, and Dr. Betsy Rice; Naval Research Laboratory scientist Dr. Igor Schweigert; and previous ARL scientist Dr. Joshua Moore.

DoD HPC INSIGHTS
is a semiannual publication of the
High Performance Computing Modernization Program.

Questions or Comments
feedback@centers.hpc.mil

DOD HPC MODERNIZATION PROGRAM SYSTEMS

The High Performance Computing Modernization Program (HPCMP) provides over 900,000 cores and 26 petaFLOPS to its users in the pursuit of improved scientific research, weapons design, force protection, and software development for the Department of Defense

Army Research Laboratory DSRC
Aberdeen Proving Ground, Maryland

EXCALIBUR
Cray XC40
Service Date 04/13/2015

Processor Cores	101,184
Memory (GB)	404,736
Disk (TB)	6,055
GPGPUs	32
teraFLOPS	3,769

PERSHING
IBM iDataPlex
Service Date: 01/07/2013

Processor Cores	20,160
Memory (GB)	45,696
Disk (TB)	2,370
teraFLOPS	419

Engineer Research and Development Center DSRC
Vicksburg, Mississippi

TOPAZ
SGI ICE X
Service Date: 07/01/2015

Processor Cores	125,440
Memory (GB)	443,584
Disk (TB)	17,496
GPGPUs	32
teraFLOPS	4,662

GARNET
Cray XE6
Service Date: 07/18/2013

Processor Cores	150,912
Memory (GB)	309,360
Disk (TB)	3,125
teraFLOPS	1,509

COPPER
Cray XE6m
Service Date: 11/19/2012

Processor Cores	14,975
Memory (GB)	29,250
Disk (TB)	443
teraFLOPS	139

Maui High Performance Computing Center DSRC
Kihei, Maui, Hawaii

RIPTIDE
IBM iDataPlex
Service Date: 05/07/2013

Processor Cores	12,160
Memory (GB)	25,984
Disk (TB)	1,500
teraFLOPS	253

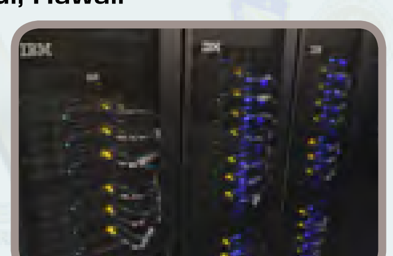

Air Force Research Laboratory DSRC
Wright-Patterson Air Force Base, Ohio

SPIRIT
SGI ICE X
Service Date: 03/25/2013

Processor Cores	73,440
Memory (GB)	146,880
Disk (TB)	2,458
teraFLOPS	1,528

THUNDER
SGI ICE X
Service Date: 09/08/2015

Processor Cores	125,888
Memory (GB)	460,288
Disk (TB)	17,568
GPGPUs	356
Coprocessors	356
teraFLOPS	5,620

LIGHTNING
Cray XC30
Service Date: 09/01/2014

Processor Cores	57,200
Memory (GB)	153,088
Disk (TB)	4,480
GPGPUs	32
teraFLOPS	1,281

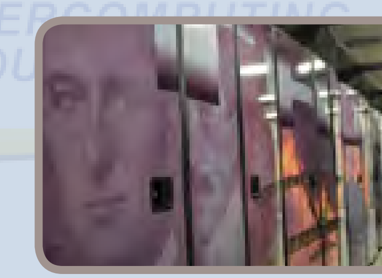

Navy DSRC
Stennis Space Center, Mississippi

SHEPARD
Cray XC30
Service Date 09/01/2014

Processor Cores	30,144
Memory (GB)	82,112
Disk (TB)	20,080
Coprocessors	124
GPGPUs	32
teraFLOPS	822

ARMSTRONG
Cray XC30
Service Date 09/01/2014

Processor Cores	30,592
Memory (GB)	82,752
Disk (TB)	2,080
Coprocessors	124
teraFLOPS	786

CONRAD & GORDON
Cray XC40
Service Date 06/19/2015

Processor Cores	61,088
Memory (GB)	204,544
Disk (TB)	1,560
Coprocessors	168
teraFLOPS	2011

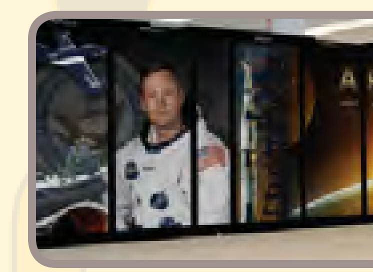

KILRAIN
IBM iDataPlex
Service Date: 01/07/2013

Processor Cores	19,584
Memory (GB)	42,240
Disk (TB)	2,304
teraFLOPS	407

HAISE
IBM iDataPlex
Service Date: 01/07/2013

Processor Cores	19,584
Memory (GB)	42,240
Disk (TB)	2,304
teraFLOPS	407

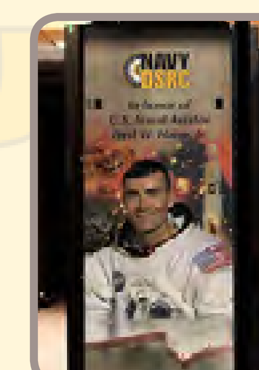

POWERFUL RESOURCES

The High Performance Computing Modernization Program (HPCMP) provides the Department of Defense supercomputing capabilities, high-speed network communications, and computational science expertise that enable DoD scientists and engineers to conduct a wide range of focused research, development, and test activities. This partnership puts advanced technology in the hands of US forces more quickly, less expensively, and with greater certainty of success.

Today, the HPCMP provides a complete advanced computing environment for the DoD that includes unique expertise in software development and system design, powerful high performance computing systems, and a premier wide-area research network. The HPCMP is managed on behalf of the Department of Defense by the US Army Engineer Research and Development Center.

To learn more about the HPCMP, please visit our website at: hpc.mil.

To learn more about our High Performance Computing (HPC) centers and HPC Systems, please visit: centers.hpc.mil.

HPCMP User Help Desk
1.877.222.2039
centers.hpc.mil

AFRL DSRC

US Air Force Research Laboratory
Wright-Patterson AFB, Ohio
http://www.afrl.hpc.mil

ARL DSRC

US Army Research Laboratory
Aberdeen Proving Ground, Maryland
http://www.arl.hpc.mil

ERDC DSRC

US Army Engineer Research & Development Center
Vicksburg, Mississippi
http://www.erdc.hpc.mil

NAVY DSRC

Navy DoD Supercomputing Resource Center
Stennis Space Center, Mississippi
http://www.navo.hpc.mil

MHPCC DSRC

Maui High Performance Computing Center
Kihei, Maui, Hawaii
http://www.mhpcc.hpc.mil