Summer 2014

DOD HPC INSIGHTS

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

DREN III Rollout Ramp Up

Optimizing Daylight Imaging on the 3.6-Meter AEOS Telescope

Massively Parallel Simulations of Supersonic Jet Noise from Complex Nozzles

HPC Insights is a semiannual publication of the Department of Defense Supercomputing Resource Centers under the auspices of the High Performance Computing Modernization Program.

Publication Team

AFRL DSRC, Wright-Patterson Air Force Base, OH
Gregg Anderson
Chuck Abruzzino

ARL DSRC, Aberdeen Proving Ground, MD
Debbie Thompson
Brian Simmonds

MHPCC DSRC, Maui, HI
Betty Duncan

Navy DSRC, Stennis Space Center, MS
Bryan Comstock
Lynn Yott

HPCMPO, Lorton, VA
Denise O'Donnell
Leah Glick

ASSOCIATE DIRECTOR FOR HPC CENTERS
Christine E. Cuicchi

MANAGING EDITOR
Rachelle Hintson

TECHNICAL EDITOR
Denise O'Donnell, HPCMPO

DESIGN/LAYOUT
Betty Watson, ACE-IT

COVER DESIGN
Data Analysis and Assessment Center

Contents

About the Cover: The roll out on the new DREN III contract is making great strides. Awarded to CenturyLink in December 2012, the $750M wide-area network services contract has over 260 sites worldwide. Over 91% of sites identified to transfer to DREN III have successfully completed acceptance testing, and nearly 85% have already fully transitioned to the new network. Users of HPCMP resources at the DSRCs will see immediate benefits – bandwidth will increase by four times the capacity available under DREN II, while bandwidth to the Internet will increase by eight times its previously available capacity! HPCMP users should be well-served by the new network for at least the next 10 years.

First Word

By Christine E. Cuicchi, Associate Director for HPC Centers

As your new Associate Director for HPC Centers, let me first thank our director, John West, for this opportunity to serve the Program, and for my predecessor, Bobby Hunter, for his service to the HPCMP during his tenure. Bobby's stewardship of the Centers was exemplary and he will no doubt continue that excellence as Director of the ERDC DSRC. I am honored and excited to begin this role after 17 years of service to the Program, and to work with our users, Centers, and the many components of the Program in accomplishing the mission and goals of the HPCMP.

HPC system procurements and deployments are a complex business, and we have begun a new process to ensure that our users receive the best and most up-to-date supercomputers as quickly as possible. This new process was used for the Technology Insertion (TI) process, and TI-13 resulted in the award of three new Cray XC30 systems at the AFRL and Navy DSRCs for a total of 2.7 petaFLOPS capability:

- The Air Force Research Laboratory DSRC at Wright-Patterson Air Force Base, Ohio, received a Cray XC30 system built upon the 2.7 GHz Intel Xeon E5-2697 v2 ("Ivy Bridge EP") processor. This system consists of 56,112 compute cores and 150 terabytes of memory, and has 1.2 petaFLOPS peak computing capability. The system will be named *Lightning* after the U.S. Air Force F-35 aircraft.
- The Navy DSRC, of the Naval Meteorology and Oceanography Command located at Stennis Space Center, Miss., received two Cray XC30 systems built upon the 2.7 GHz Intel Xeon E5-2697 v2 processor and the 1.05 GHz Intel Xeon Phi Coprocessor 5120D. These two systems are identical; each consisting of 29,304 compute cores, 7,440 coprocessor cores, 78 terabytes of memory, and 0.75 petaFLOPS of peak computing capability. The systems, which provide 1.5 peak petaFLOPS combined, are designed as sister-systems to provide continuous service during maintenance outages. They will bear the names *Armstrong* and *Shepard*, honoring naval aviators and NASA astronauts Neil Armstrong and Alan Shepard.

The systems are expected to enter production in the second half of calendar year 2014 and will provide an increase of approximately 50% in the HPCMP's computing capability.

In this edition of *HPC Insights*, the HPC Centers team is pleased to feature the new Center's website, centers.hpc.mil, and the co-located User Dashboard, which provide an array of information regarding systems, Program resources, and user specific information to our users, staff, and other HPCMP stakeholders. We strongly encourage you to visit these sites, and very much appreciate your feedback on the sites or any other Center's issues you wish to weigh in on at feedback@centers.hpc.mil.

We hope that this informal feedback avenue, in addition to the user satisfaction surveys and the well-respected User Advocacy Group, will strengthen the communications between users and our Centers. Effective communications are crucial for the success of our users and our Program.

DREN III Rollout Ramp Up

by Ralph McEldowney, Associate Director for Networking, HPCMPO, Lorton, VA

The Defense Research and Engineering Network (DREN) III is on a roll and will soon be coming to a site near you. After nearly a year of planning for transition and successfully achieving an initial operational capability with the first 10 sites in the fall of 2013, the DREN III network rollout continues on a brisk pace with sites transitioning to the network every week. The entire transition is expected to be complete by June 2014.

DREN III is a 10-year, firm-fixed price, indefinite-delivery/indefinite-quantity, wide-area network services contract for more than 260 sites worldwide, with a value up to $750M. Through a full and open competition, the Defense Information Technology Contracting Organization (DITCO) awarded the DREN III Contract to CenturyLink, Inc. in December 2012.

The HPCMP is transitioning 149 sites from DREN II to the new DREN III network. Once CenturyLink installs the DREN III equipment and wide-area circuit at a site, the service undergoes an extensive acceptance test to verify that it performs as specified in the contract. Once acceptance testing is complete, CenturyLink and the DREN team transition all of the enclaves at a site to the new network. Depending on the complexity of a site, this transition can require a few hours for simple sites, to several weeks for the most complex sites. In either case, the transition of every site will be accomplished professionally and successfully before the June 2014 deadline.

Progress to-date on site acceptance tests and site transitions has been impressive. Since the fall of 2013, over 91% of the 149 sites scheduled to transition have successfully completed acceptance testing; while nearly 85% have either fully or partially transitioned to the new network. Aside from DOD user sites, transitioned sites include two Internet transit locations, three Internet peering locations, and two DREN II to DREN III gateways. The two gateways ensure sites on each network can continue to communicate with each other during the transition without requiring this traffic to traverse the Internet. Users should not notice any difference in service when accessing the Internet from DREN II or DREN III.

Non-classified Internet Protocol Router Network (NIPRNet) peering with DREN, on the contrary, presents some challenges. Though a final DREN-NIPRNet peering solution will result in increased bandwidth and more reliable connectivity, this solution won't be fully implemented until the summer of 2014. In the interim, some HPCMP users may experience additional latency and network congestion as data flows between the two networks. This is necessary because the west coast DREN-NIPRNet peer has been dedicated to support only the DREN III sites, while the east coast DREN-NIPRNet peer has been dedicated to support only the DREN II sites. While the DREN's engineering team is examining all options to reduce the current latency issue, patience will be required until the final solution is implemented.

DREN connectivity across the continental U.S. and Alaska/Hawaii.

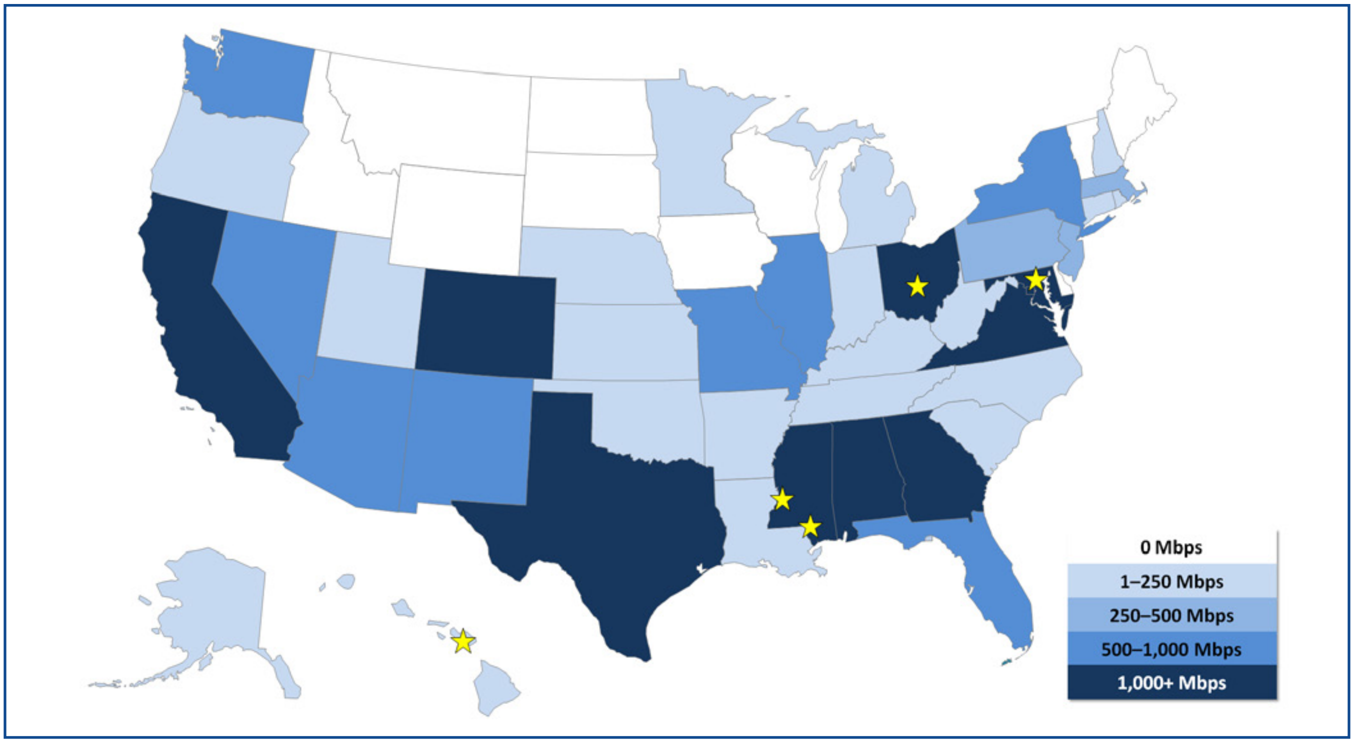

Aggregate DREN III bandwidth implemented per State. White represents no bandwidth deployed, while dark blue represents the highest bandwidth deployed. States with DSRCs are typically dark blue indicating the highest DREN III bandwidth deployed.

Users of HPCMP resources at DOD Supercomputing Resource Centers (DSRCs) will see immediate benefits from the new network. It has been well-known for several years, that the HPCMP research and test communities have strained the capacity of DREN II, not only at the DSRCs but to the Internet as well. As a result of the new DREN III network, bandwidth at the DSRCs will increase by four times the capacity available under DREN II. Likewise, bandwidth to the Internet on DREN III will increase by an impressive eight times its previously available capacity. Increasing the bandwidth on DREN II was simply not cost-effective under a firm-fixed price contract that established bandwidth prices over 10 years ago. Now with the DREN III contract in place, the HPCMP can take advantage of the decrease in bandwidth prices and order capacities not previously affordable or imagined a decade ago.

Comparing the total aggregate bandwidth between DREN II and DREN III is equally impressive. Under DREN II, the total aggregate bandwidth ordered for all sites was 33.7 Gigabits per second, while for DREN III the total aggregate bandwidth ordered for all sites is 212 Gigabits per second – more than a six fold increase!

Along with noticeable increases in bandwidth capacity, there are many new technical capabilities that will likely go unnoticed by most HPCMP users. These technical capabilities include new data transport options between sites using Ethernet and Optical Services. Both options provide network engineers and administrators with new ways of connecting the HPCMP sites together. For example, the HPCMP's Data Recovery network was redesigned under DREN III to take advantage of the new Ethernet option. This approach provides a more isolated and secure network, which is less visible and, thus, less vulnerable to the Internet at-large. While the Optical Services solution is not currently implemented on DREN III, it is available for future uses, including use as a platform for network research.

To prepare DREN Site POCs at each of the 149 locations for the DREN III transition, the HPCMP sponsored the DREN III Users' Forum in September 2013. To preserve valuable travel funding, the Users' Forum was held via Video Teleconference (VTC) and conference calls. Two identical four-hour sessions were presented on back-to-back days. Topics included the DREN III network design, operations, security, transition plans and the transition schedule. Although the use of VTC and conference calls was more cumbersome than a face-to-face meeting, the Users' Forum was a resounding success. Over 270 participants gained significant insight into the new DREN III network, as well as information on the transition plans and schedule. As a result, site POCs are better able to prepare their site for the new network.

In summary, the new DREN III network rollout is making great progress, with new sites transitioning every week. DREN III brings not only new technical capabilities, such as Ethernet and Optical Services, but greatly increased bandwidths to the DSRCs and the Internet; benefiting the entire HPCMP user community. Over 91% of sites identified to transition to DREN III have successfully completed acceptance testing, while nearly 85% of the sites have already transitioned to the new network. The total aggregate bandwidth ordered on DREN III represents a six fold increase over the total aggregate bandwidth ordered on DREN II. HPCMP users should be well-served by the new DREN III network for the next 10 years.

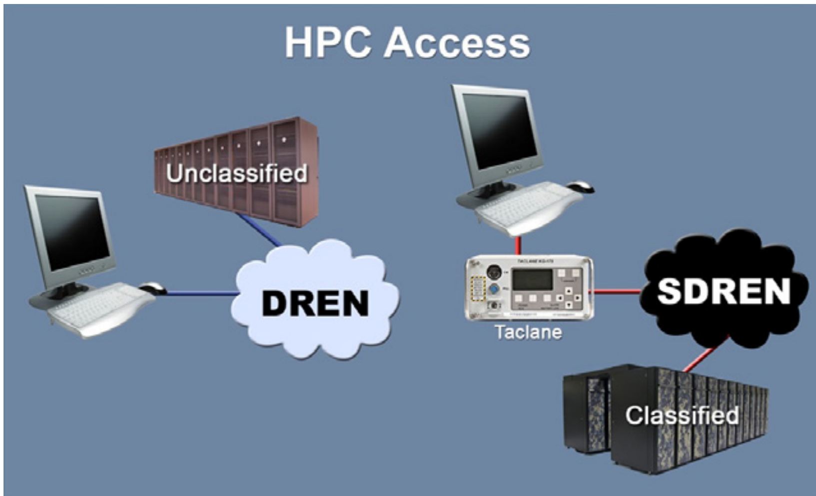

Traditional HPCMP Connectivity. Users are connected to unclassified HPC resources over DREN, and to classified HPC resources over Secret DREN (SDREN).

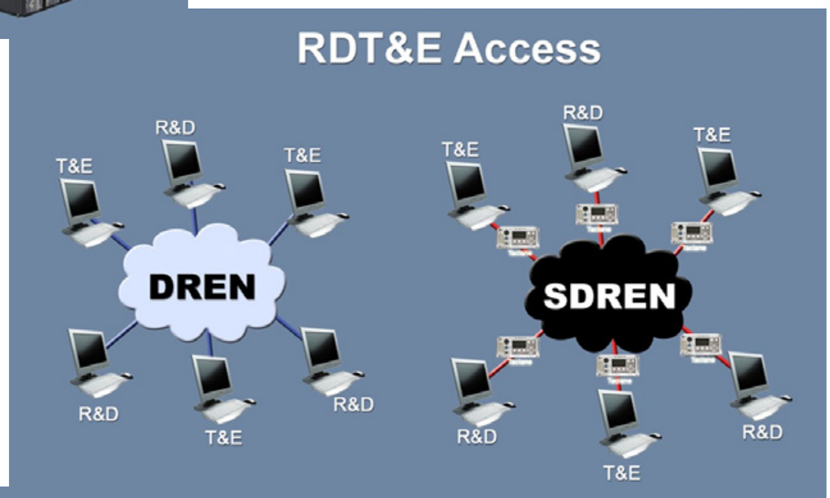

Typical RDT&E Connectivity. Unclassified RDT&E users are connected over DREN, and classified users are connected over SDREN.

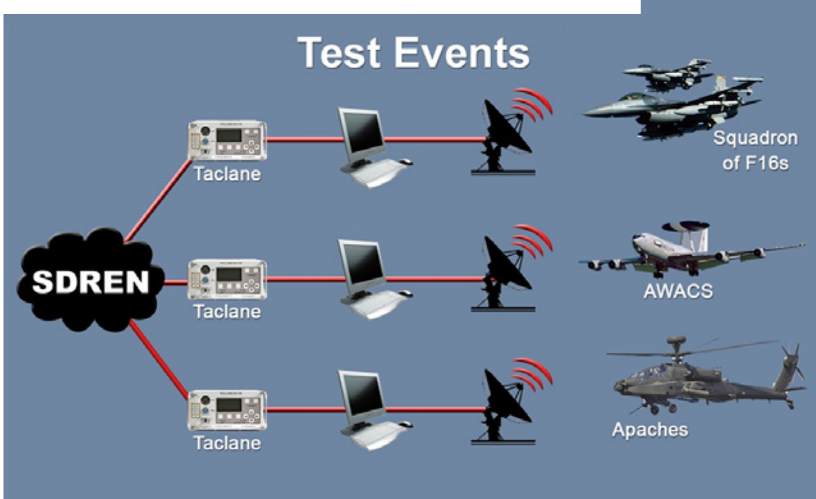

Traditional Test Event Connectivity. Deployed sensors on aircraft, ships, and so on are connected with users over SDREN.

Closed Community Connectivity. Members of closed communities of interest are connected to each other over closed, private networks.

Optimizing Daylight Imaging on the 3.6-meter AEOS Telescope

By Kanoe Hardin, undergraduate Intern from the University of Hawaii Maui College, Mentored By Dr. Stacie Williams, AFRL/DE Physical Scientist, Air Force Research Laboratory

HPC Resources: IBM iDataPlex (*Riptide*), MHPCC DSRC

The Maui Space Surveillance Complex (MSSC), located on the summit of Haleakala, has an array of telescopes that provide space situational awareness in the strategically relevant Pacific arena. One of the critical missions of the MSSC is imaging high priority sun synchronous objects. For more than a decade, MSSC has collected images during daylight hours with a 1.6-meter telescope. To improve resolution of collected imagery, MSSC plans to use the largest optical telescope in the Department of Defense, the 3.6-meter Advanced Electro-Optical System (AEOS), to collect daylight imagery. Based on imaging aperture alone, this 3.6-meter upgrade would provide a 2.25x the resolution of the current 1.6-meter capability.

Upgrading AEOS to a daylight imaging asset requires sun mitigation to prevent thermal damage to the telescope, to reduce sky background, and to prevent stray light degradation of collected imagery. AEOS is an open-truss telescope enclosed in a unique dome that completely lowers during operations. This design maximizes imagery under dark conditions by alleviating the issue of dome-induced turbulence which degrades imagery, but provides no sun protection during daylight (Figure 1). The Air Force Research Laboratory Directed Energy Directorate will be designing a sun mitigation system to increase daytime use and significantly improve MSSC's daylight imaging resolution capability.

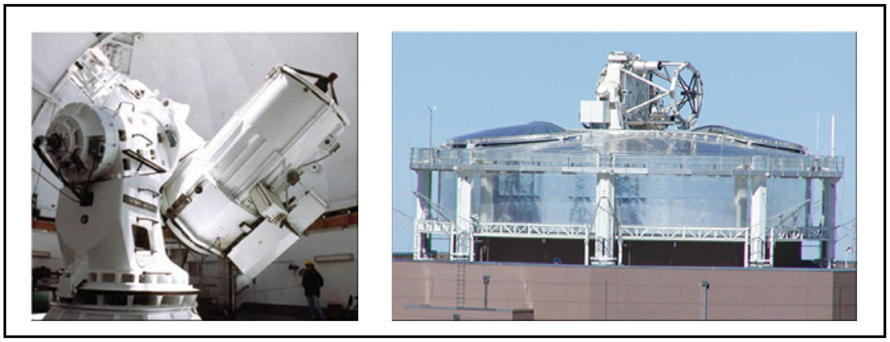

Figure 1. Photos of the 1.6-meter telescope (left) and the 3.6-meter telescope (right) at the MSSC. The 1.6-m telescope is the traditional tube style enclosed in a dome providing protection from the sun. The 3.6-m telescope is an open truss design with a dome that completely lowers during data collection. This design overcomes anticipated turbulence-induced image degradation generated by artificial structures.

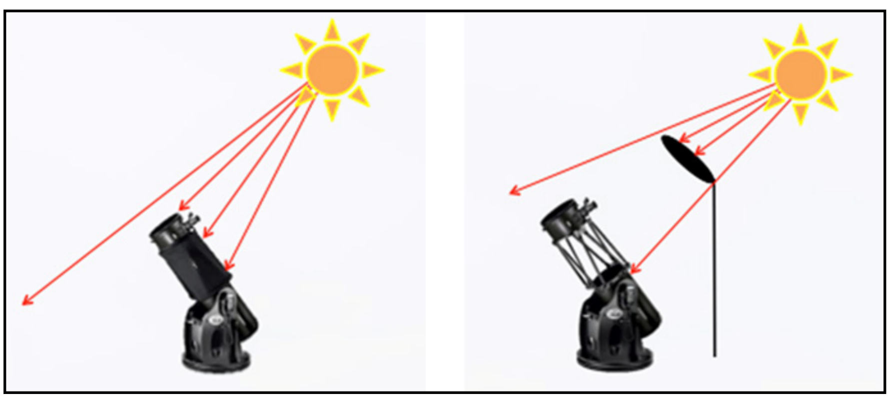

Figure 2. The image on the left shows a cloth baffle wrapped around the truss tubes. The image on the right shows a spherical disk decoupled from the telescope structure placed between the sun and the telescope. Given the short duration of the author's internship, the parasol was not optimized for this optical system, but was over-sized to ensure complete shading of the telescope structure providing consistency in comparative measurements.

Two sun mitigation options were investigated in this study to determine which method best improved imagery by reducing stray and background light (Figure 2). The first option is the traditional approach for mitigating stray light, which is to retrofit AEOS with a baffle that covers the truss structure. Unfortunately, this solution requires a large sun exclusion zone, and would not protect the telescope from solar damage in the event of inadvertent sun exposure on the telescope mirror system, which has the potential to significantly damage the telescope. Another approach is to decouple the sun mitigation from the telescope and construct a shading mechanism, referred to as a parasol, that serves as a physical barrier between the sun and the telescope.

While this solution decreases the sun exclusion angle and protects the system in cases of inadvertent sun exposure, it is more complex and more costly than the baffle. To compare these two sun mitigation options, a series of tests were conducted on a commercial small-scale telescope (Figure 3) with a similar open truss design to AEOS in order to determine which solution best lowered sky background and best improved daylight imagery.

Images were collected under identical conditions with the telescope in three configurations; telescope with no sun mitigation, telescope with a baffle, and telescope without the baffle but shading the telescope from direct sunlight with a disk referred to as a parasol. The images shown in Figures 4A-4C reveal significant improvement with the baffle and the parasol; however, the parasol (Figure 4C) shows the most significant improvement. Sky background was also measured as a function of angle from the sun (Table 1). In all data sets collected, the parasol lowered the sky background by a factor of two more than the traditional baffle. Furthermore, it is anticipated that the parasol method will allow measurements very close to the sun by providing a solar exclusion zone that approaches zero.

In conclusion, the sun mitigation that best improved imagery and lowered sky background was the parasol method. Design of an optional parasol will be a complex and costly undertaking. To determine whether this sun mitigation option will be implemented at MSSC will require additional studies that include statistical analysis of high-priority objects to determine required solar exclusion angle and engineering design solutions for both the baffle and parasol implementation. A cost-benefit analysis will be used to determine which sun mitigation meets the needs of the Air Force Space Command.

Acknowledgements: Ms. Hardin completed this work as an Akamai intern for the Air Force Research Laboratory Directed Energy Directorate under the direction of Dr. Stacie Williams. The Akamai Internship Program offers community college students and undergraduates that are attending college in Hawaii, or that are from Hawaii but studying on the mainland, an opportunity to get involved in high-tech research and industry. Support for the 2013 Akamai Internship Program and the Akamai Workforce Initiative is provided by: the National Science Foundation (AST#0836053); Air Force Office of Scientific Research (FA9550-10-1-0044); University of Hawaii; Thirty Meter Telescope Corp.; University of California, Santa Cruz; and Kamehameha Schools. Ms. Hardin would also like to thank the members of the MSSC Daylight Imaging Team for all of their help particularly John Valiant, Dr. Michael Werth, Daniel Thompson, Dr. Brandoch Calef, Dr. Steven Griffin, Dr. Ed Walker, and Andrew Whiting.

Table 1. A graph of normalized sky background counts vs. angle from the sun is shown below. The parasol sun mitigation system utilized for this study was not optimized for this optical system, and prevented data collection below 40 degrees. The dashed line is an extrapolation of the data. The parasol curve is not anticipated to follow the same trend as the other two curves at angles below 15 degrees. The Maui sky above Haleakala has very low aerosol levels and scattering is anticipated to be low, thus sky background is low when the parasol prevents direct sunlight from reaching the optical system (Barry Labonte. Sky brightness measurements at Haleakala, 1955–2002. *Solar Physics*, 217:367–381, 2003.).

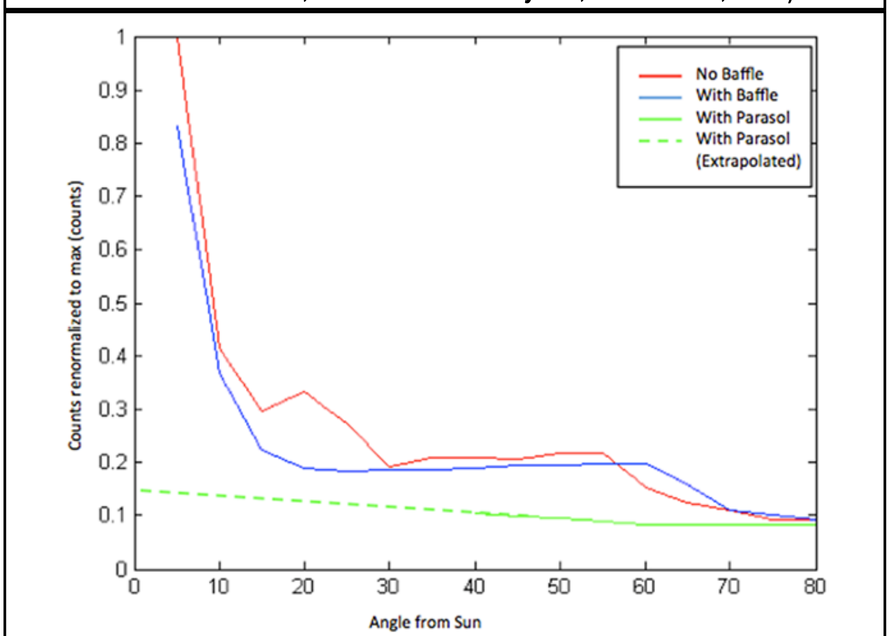

Figure 3. This is the commercially-available Orion SkyQuest XX12g GoTo Dobsonian telescope.

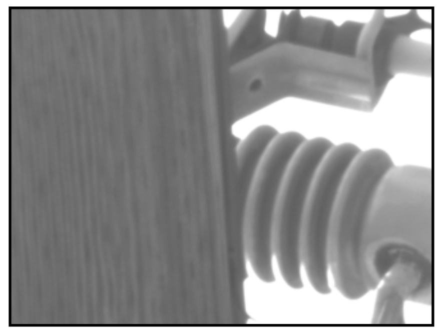

Figure 4A. Image of power transformer, taken with no sun mitigation.

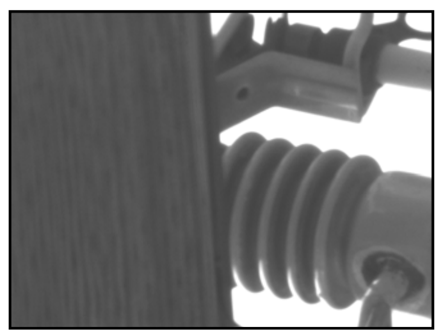

Figure 4B. Image of power transformer, taken with baffle.

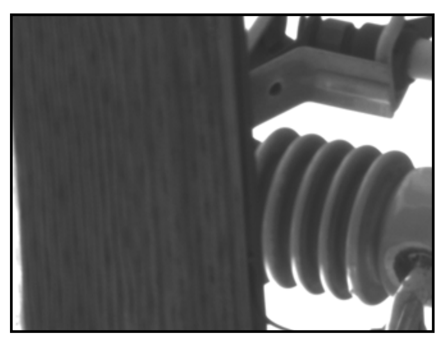

Figure 4C. Image of power transformer taken with parasol.

Massively Parallel Simulations of Supersonic Jet Noise from Complex Nozzles

by Guillaume A. Brès, Frank E. Ham, Cascade Technologies Inc., Palo Alto, CA; Joseph W. Nichols, Sanjiva K. Lele, Stanford University, Stanford, CA; and John T. Spyropoulos, Naval Air Systems Command, Patuxent River, MD

HPC Resources: Cray XE6 (*Raptor*), SGI ICE X (*Spirit*), Cray XE6 (*Garnet*), ERDC DSRC

Importance and Challenges of Jet Noise Predictions for Tactical Aircraft

The operation of modern tactical aircraft, in particular from aircraft carrier decks, exposes military personnel to extreme noise environments that can result in significant and permanent hearing loss. Unlike commercial airplanes, the propulsion systems of combat aircraft are based on low bypass ratio turbojet engines. The exhaust jets are supersonic and hot, especially with the augmenter/afterburner in operation. In some cases, this can lead to crackle, which is an especially irritating component of supersonic jet noise because it is associated with N-shaped acoustic waveforms with sudden strong compressions followed by more gradual expansion. In addition, the static pressure at the exit of the nozzle is generally not matched with the ambient pressure during high-powered takeoffs and landings. The pressure mismatch generates a shock-cell system in the jet exhaust plume that interacts with the jet turbulence. These interactions generate additional jet noise components, namely broadband shock-associated noise (BBSN) and tonal screech noise.

It has been estimated that, in the US military, the disability claims associated with hearing loss exceed $1B each year. Finding effective strategies to mitigate the exposure to the severe noise without severely impacting propulsive performance are therefore critical. Until recently, the development of such strategies has relied largely on laboratory and full-scale testing, as well as some numerical prediction tools typically based on empirical databases, modeling and Reynolds–Averaged Navier–Stokes (RANS) approaches. However, cost constraints and the high complexity of the flow often limit the range of the parametric investigation and the success of the design optimization. Likewise, while semi-empirical numerical methods provide useful first estimates, they tend to lack the sensitivity to design/configuration changes needed in a design tool, and are largely untested for complex settings such as multi-stream nozzle or carrier deck takeoff conditions.

In this context, high fidelity, unstructured, large eddy simulation (LES) is emerging as an accurate yet cost-effective computational tool for prediction of turbulent jets and their acoustic fields for multiple reasons. First, it is a fully predictive approach based on first principles, free of user-defined calibration or empirical constants. Second, the complex nozzle geometries such as chevrons, faceted military-style nozzle, and multi-stream architecture, can be explicitly included in the simulations thanks to the unstructured mesh capabilities enabling the direct study of the variation of flow and emitted noise by different nozzle configurations. Third, LES provides access to the complete transient flow field, allowing in-depth probing of the physics of jet noise production. Finally, advancement in high performance computing (HPC) has enabled large-scale, high fidelity LES to become a practical engineering solution for design studies. Access to current (and future) supercomputer facilities have brought forward a new engineering paradigm for massively parallel methods. These methods contrast with past schools of thought and existing approaches, which rely on low-cost methods, modeling, and empirical databases.

The last point is particularly relevant for the DOD High Performance Computing Modernization Program. As previously mentioned, the RANS approach is currently an industry standard in computational fluid dynamics (CFD) for aerospace applications, but significant modeling and calibration are required to account for flow turbulence and to ultimately predict noise. The other important limitation of the RANS approach for future design studies is the absence of a path to higher accuracy with increasing computer power. That is, the predictive accuracy of the RANS method is not necessarily improved by increasing the spatial and temporal resolution of the computational grid. It is limited by the irreducible errors of the turbulence closure models. In contrast, LES has a well-defined path to increased predictive capability through increased computer power. Via increased spatial and temporal resolution, afforded by advances in computer technology, the LES equations approach the true flow equations derived from first principles as the effect of phenomenological sub-grid scale models diminish. The current industry standard engineering design tools have plateaued in their ability to harvest increasing computer power, and to resolve the critical technical challenges that arise as the complexity of the nozzle geometry and operating conditions are increased. The next generation of prediction tools for engineering design of future tactical aircraft nozzles should be developed with first principles approaches, such as LES, and anticipate the continued advancements in high-performance computing.

Charles: The High Fidelity Massively Parallel Compressible Flow Solver

Through several projects funded by NASA, Air Force, NAVAIR, and the Office of Naval Research (ONR), and in collaboration with Stanford University, there has been a significant and continuous effort over the past few years to improve understanding and develop predictive capabilities of propulsive jet aeroacoustics, through high fidelity, physics-based simulations with the unstructured LES framework developed at Cascade Technologies. For aeroacoustics, the framework is composed of preprocessing tools (e.g., for unstructured mesh refinement and adaptation capabilities, see next section), a compressible flow solver, Charles, and post-processing tools (e.g., for far-field noise prediction). The solver, Charles, features all the capabilities required to efficiently perform large-scale, high fidelity aeroacoustics simulations of turbulent jet flows from complex geometries (i.e., unstructured grid capabilities, low-dissipative numerical schemes, shock capturing scheme, sub-grid scale modeling, highly parallel computing performances, etc.). For the far-field noise predictions, an efficient frequency-domain permeable formulation (Lockard 2000) of the Ffowcs Williams–Hawkings (FW-H) equation (Ffowcs and Hawkings 1969) has been implemented in Cascade's massively parallel unstructured LES framework (Brès et al. 2012). As an example of the typical LES setup for aeroacoustics, Figure 1 shows the computational domain for a hot supersonic jet issued from a rectangular nozzle (Nichols, Ham, and Lele 2011), including the FW-H surface used to propagate the near-field LES data to the far-field and compute the radiated noise. The instantaneous temperature and pressure fields are also shown, to visualize the shock cells, screech noise, and broadband jet noise.

One of the unique features of Charles is the solver performance and scalability on massively parallel environments. The solver is designed and implemented using the Message Passing Interface (MPI) to function well in a massively parallel, distributed memory environment. Calculations are carried out routinely on several thousand processors at DOD supercomputer facilities at ERDC and AFRL. Scalability studies performed during the Cray XE6 testing period of 2010 (CAP and CAP-2 early access programs) demonstrated that the code exhibits perfect scaling at the 20,000 core level. The algorithm efficiency, defined as the speedup divided by the number of participating processors, remains at 97% for 20,000 cores compared to 100% at the 2,048 core baseline. More recent scalability tests on different systems showed that the code exhibits good scaling for up to 160,000 cores. Figure 2 shows scaling statistics (speedup vs. number of cores) for a jet noise simulation (Nichols et al. 2012) on a mesh with 528M control volume, computed on the *Intrepid* system at Argonne National Labs (ANL). The calculation on 163,840 cores used the full capabilities of the system at 80% parallel efficiency. Testing on these different systems (IBM BlueGene P&Q, Cray XE6, SGI, commodity Linux clusters, etc.) also demonstrated great cross platform portability.

Moreover, the Charles solver was recently run on over one million cores, reaching a new milestone in high performance computing. This breakthrough was realized in January 2013 during early science testing of the newly installed Blue Gene/Q supercomputer at Lawrence Livermore National Laboratory (LLNL), in collaboration with Stanford University and LLNL computing staff. The jet noise calculation performed during that breakthrough challenged all parts of these supercomputers because waves propagating throughout the tightly-coupled simulation required a well orchestrated balance between computation, memory, communication, and I/O. The data from this simulation is being used to understand how such jets emit crackle noise, a significant and especially irritating part of the overall acoustic field (see section below).

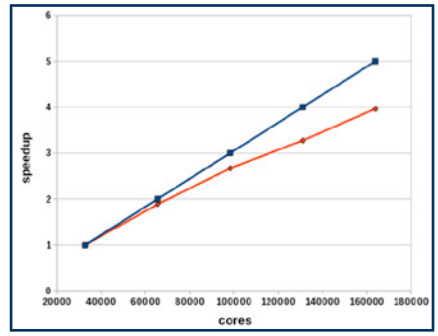

Figure 2. Scalability of the LES code Charles on the Intrepid system at Argonne National Labs (up to 163,840 cores).

From Simple Jet Geometries to Chevrons and Military-Style Faceted Nozzles

In past studies, Charles has been used to investigate a wide range of high-speed high-Reynolds number, unsteady flow processes for various jet configurations, including impinging flows (Khalighi et al. 2011, Brès et al. 2011), circular (Khalighi et al. 2010, Khalighi et al. 2011, Brès et al. 2012) and rectangular (Nichols, Ham, and Lele 2011, Nichols et al. 2012) jets, chevrons (Brès et al. 2012, Nichols et al. 2012), faceted military style nozzle (Nichols et al. 2013, Nichols, Lele, and Spyropoulos 2013), and internally mixed dual-stream jet (Brès et al. 2013). Access to DOD supercomputers and "Challenge" status granted by the HPC programs were critical to the success of these research projects.

Simulations have progressed from simple geometries to more complex configurations in recent years with studies showing a commonality

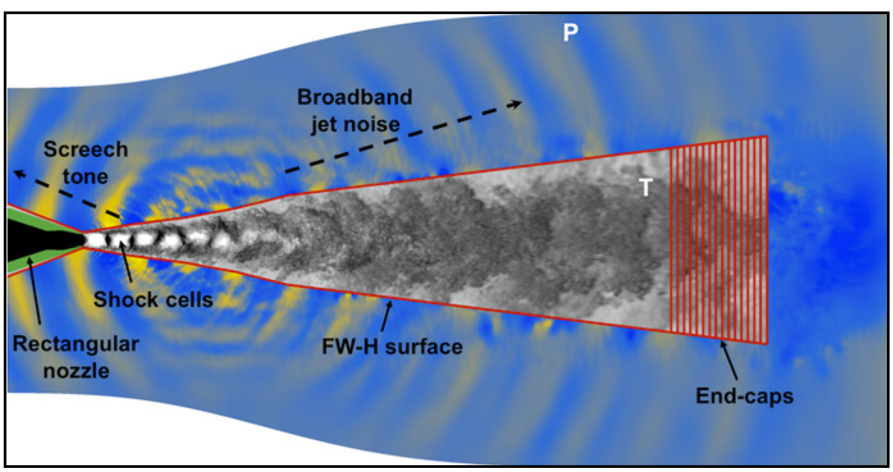

Figure 1. Typical LES setup for aeroacoutics: example of a hot supersonic jet issued from a rectangular nozzle (Nichols, Ham, and Lele 2012). Inside the FW-H surface, the shock cells and turbulence of the supersonic jet flow are visualized by grayscale contours of temperature T, with black and white corresponding to hot and cold, respectively. The direct near-field sound outside of the FW-H surface is displayed by color contours of pressure P.

that the complex nozzle is explicitly included in the computational domain using unstructured body-fitted mesh and adaptive grid refinement. Here, it is important to point out that the generation of a high-quality mesh for a complex geometry remains a pacing issue in high fidelity flow simulation. LES places strict requirements on mesh resolution, element quality, and even the level of allowed mesh anisotropy. To address these challenges, a massively parallel tool has been developed in Cascade's solver infrastructure to give the user detailed control over the local mesh resolution in their grid. This mesh adaptation module, Adapt, features localized adaptive refinement capabilities and produces high-quality yet economical unstructured grids suitable for capturing turbulence dynamics.

An example of the use of Adapt is shown in Figure 3 for a 12-count chevron appended on a round converging-diverging nozzle. The starting point is a very coarse grid containing about 0.16 million control volumes. Several embedded zones of refinement can then be simply defined by the user and automatically enforced by the adaptation tool Adapt. The code includes a smoothing algorithm to avoid sharp grid transitions between different refinement zones. A surface projection method is also applied to respect non-planar mesh boundaries during refinement, ensuring accurate representation of the underlying geometry. The use of very coarse grids as a starting point is a convenient approach, as it not only greatly simplifies the meshing process and reduces the burden on the users, but it also allows for complete control of the location and length-scale of the refinement.

For this project, a grid resolution study was performed by increasing the mesh size from 55 × 106 to 118 × 106 and 386 × 106 cells, with the bulk of the mesh in the jet potential core and enclosing the FW-H surface. The grid refinement analysis showed that while additional resolution is always beneficial to the numerical predictions, the grid size range of 50 to 60M cells seems to be a good compromise between accuracy and runtime/data storage for current computational resources. Figure 4 shows the flow field visualization for a heated supersonic over expanded jet issued from the chevron nozzle presented in Figure 3, on the 386M mesh. The development of fine turbulent structures from the chevron's nozzle exit to the end of the potential core, as well as the shock cell pattern, is clearly visible in the top figure. This example highlights some of the multi physics capabilities of the solver as well as the low dissipative aspect of the numerical scheme.

Figure 5 shows the noise spectra comparison between LES predictions and experimental measurements carried out at United Technologies Research Center (UTRC) for the same geometry and operating conditions. Both near and far field noise predictions match the experimental results accurately over a wide range of frequencies. It should be underscored that the results presented here correspond to blind comparisons, in the sense that all the simulations and postprocessing were performed without prior knowledge of the experimental data. This level of agreement is the state-of-the-art in the field of computational aeroacoustics and demonstrates the fully predictive

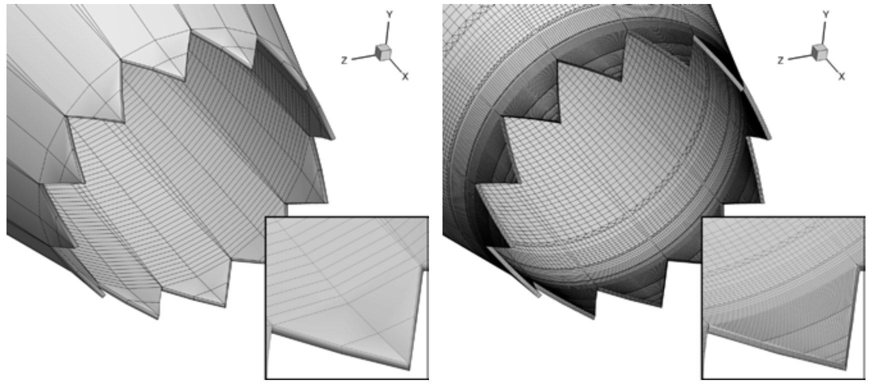

(a) Initial coarse "skeleton" mesh (0.16M cells) (b) Adapted mesh (55M cells)

Figure 3. Chevron nozzle surface mesh and zoomed-in view of chevron tip.

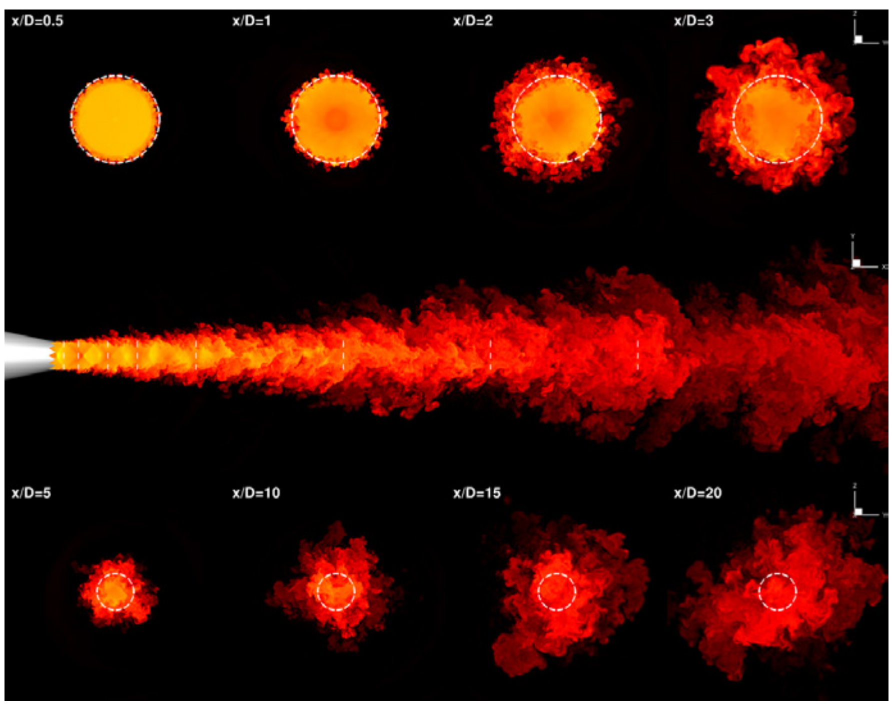

Figure 4. Instantaneous temperature field for a heated supersonic over expanded jet issued from the chevron nozzle. The cross flow cuts of the temperature at different locations downstream of the nozzle are shown in top and bottom subfigures. The dashed circular line corresponds to the nozzle exit outline.

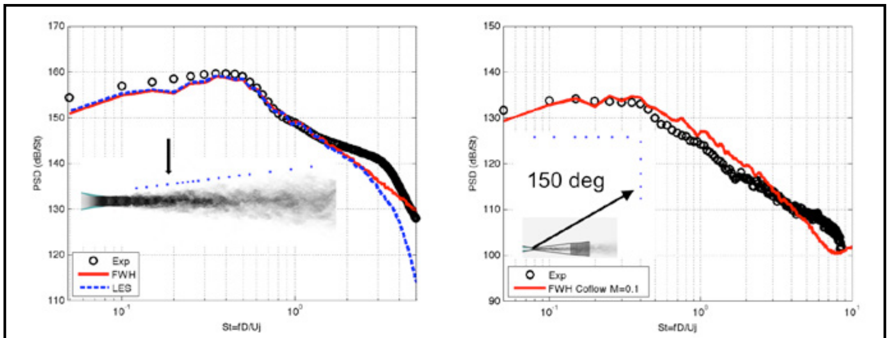

(a) Near-field probe (b) Far-field microphone

Figure 5. Blind comparison of the sound spectra at different locations between experimental (symbols) and LES (lines). See (Brès et al. 2013) for details.

capabilities of the present high fidelity LES. The large database generated during these studies is currently being used to evaluate chevron design, not only for noise reduction, but also in terms of performance and thrust loss.

Crackle: The Most Annoying Component of Supersonic Jet Noise

While blind comparisons to aero-acoustic laboratory experiments have demonstrated the predictive power of LES, the potential of advanced simulation methods for scientific understanding is greater still. As mentioned above, LES provides access to complete three-dimensional transient flow fields, in a fashion unparalleled by experimental measurements. In essence, LES functions as a computational microscope which scientists can use to understand the root causes of complex phenomena such as jet noise.

In this spirit of computation based discovery, part of our research effort has been directed towards applying high fidelity LES to understand the mystery of crackle noise that sometimes occurs in hot supersonic jets, such as those produced by the exhausts of high-powered tactical aircraft. Crackle noise is characterized by N-shaped, acoustic waveforms that result in intermittent sudden compressions in the pressure field at observer locations (Ffowcs Williams, Simson, and Virchis 1975). When crackle occurs, it is extremely irritating to human perception. Moreover, even though crackle is an intermittent phenomenon, it can account for a significant portion of the overall sound pressure level in the direction of the peak noise radiation (Krothapalli, Venkatakrishnan, and Lourenco 2000). Its elimination would therefore have the added benefit of decreasing the peak jet noise, an important factor for noise regulations, ensuring safety of carrier deck personnel, and maintaining community standards close to military bases.

Although mitigating crackle would have clear practical advantages, the fundamental mechanism by which it is produced has been the subject of recent debate. One possible explanation is that nonlinear effects cause acoustic waves of sufficiently high amplitude to steepen as they propagate, eventually leading to N-shaped waves interpreted as crackle. The propagation distance required for nonlinear effects to become significant, however, depends on the wave's initial amplitude, and in most cases is longer than distances relevant to an aircraft carrier deck (Ffowcs Williams, Simson, and Virchis 1975). Instead, our simulations have allowed us to observe N-shaped waves emerging directly from a spatially developing turbulent round supersonic jet issuing from a faceted military-style nozzle [Nichols *et al.* 2013, Nichols, Lele, and Spyropoulos 2013). An interesting aspect of our simulations is that once an N-shaped wave is detected at an observer location, the simulation, through checkpointing, is able to revert back to a precise earlier time, reconstructing the complete flow field responsible for producing that particular wave. Figure 6 shows an average of several such crackle producing flow fields. Because time flows in one direction in the real world, such time-lagged phase-averaging is difficult to measure experimentally, except through the storage of tremendous amounts of data, which by necessity are of reduced spatial and temporal accuracy.

Figure 6 specifically shows color contours of the phase-averaged pressure on an axial cross section through a military style nozzle (green, at left). The magenta circle in the upper right marks the observer location. In addition to the shock cells present in the core of the jet, a region of high pressure is observed to form on the upper edge of the jet, spanning the turbulent shear layer at that location. This suggests that crackle is generated by large-scale flow features embedded in the turbulence,

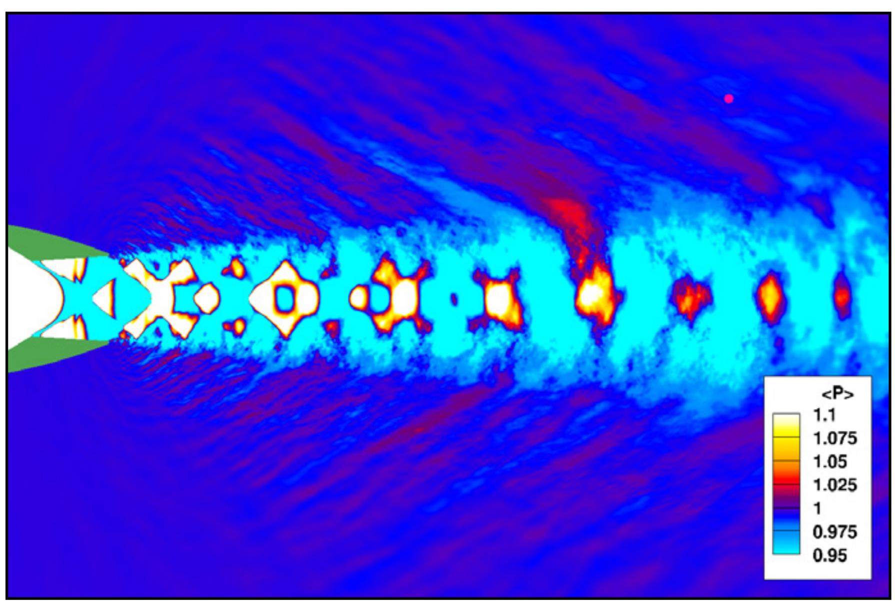

Figure 6. Contours of phase-averaged pressure on an axial cross section through a high fidelity LES of a crackling supersonic jet.

and explains why recent full-scale tests involving military-style nozzles with chevrons (which tend to breakup large-scale flow features) have shown a reduction in crackle noise (Martens, Spyropoulos, and Nagel 2011).

Conclusions and Outlook

The paper is intended as an overview of the recent progress made with massively parallel large eddy simulation for jet noise predictions of tactical aircraft. The compressible flow solver Charles used in the present jet noise studies is part of the multiphysics unstructured LES framework developed at Cascade Technologies, which extends to multiphase and combustion applications. Over the past year, nearly 20 million computing hours from the DOD HPC Modernization Program were used to simulate heated supersonic jets and aero engine injectors. These resources enabled us not only to investigate chevron effects on round nozzles and multiphase flows in complex combustor/augmentor geometry, but also to explore other applications relevant to tactical aircraft, such as twin nozzle architecture (see Figure 7) and crackle noise.

A continued effort is being made to conduct further analysis of the large LES database collected over the course of the current HPC Challenge Project. Experience gained from these studies is currently used for the mesh design, numerical setup, and acoustic post-processing of ongoing work, to continue advancing existing methodologies towards best practices for jet noise predictions with unstructured LES.

Acknowledgments

The chevron nozzle study is supported by NAVAIR Small Business Innovation Research/Small Business Technology Transfer program contracts N68335-11-C-0026 under the supervision of Dr. John T. Spyropoulos. The twin nozzle study is supported by ONR grant N00014-12-C-0010, under the supervision of Dr. Brenda Henderson. Most of the simulations were performed as part of the HPC Challenge Project "Large Eddy Simulations of Supersonic Jet Noise and Combustor & Augmentor Spray Atomization." The majority of calculations were carried out on *Garnet*, *Raptor*, and *Spirit*, at DOD supercomputer facilities at ERDC and AFRL.

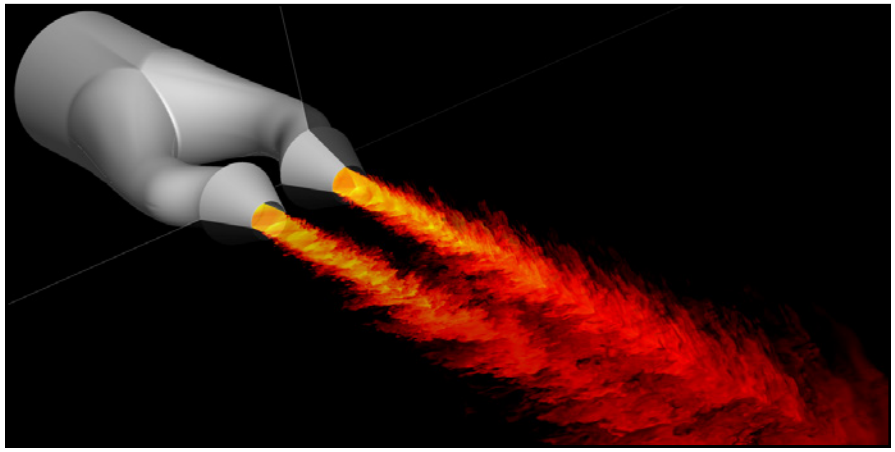

Figure 7. Temperature field for heated supersonic jets issued from a twin nozzle.

References

Brès, G. A., Ham, F. E., Nichols, J. W., and Lele, S. K. 2013. "Nozzle Wall Modeling in Unstructured Large-Eddy Simulations for Hot Supersonic Jet Predictions," AIAA paper 2013-2142.

Brès, G. A., Khalighi, Y., Ham, F., and Lele, S. K. 2011. "Unstructured Large-Eddy Simulation Technology for Aeroacoustics of Complex Jet Flows," *Proceedings of Inter-Noise 2011*, Sept. 2011, Osaka, Japan.

Brès, G. A., Nichols, J. W., Lele, S. K., Ham, F. E., Schlinker, R. H., Reba, R. A., and Simonich, J. 2012. "Unstructured Large-Eddy Simulation of a Hot Supersonic Over-Expanded Jet with Chevrons," AIAA paper 2012-2213.

Brès, G. A., Nichols, J. W., Lele, S. K., and Ham, F. E. 2012. "Towards Best-Practices for Jet Noise Predictions with Unstructured Large-Eddy Simulations," AIAA paper 2012-2965.

Ffowcs Williams, J. E., and Hawkings, D. L. 1969. "Sound generation by turbulence and surfaces in arbitrary motion," *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, Vol. 264, pp. 321–342.

Ffowcs Williams, J. E., Simpson, J., and Virchis, V. J. 1975. "Crackle: an annoying component of jet noise," *J. Fluid Mech.*, Vol. 71(2), pp. 251–271.

Khalighi, Y., Ham, F., Moin, P., Lele, S. K., Colonius, T., Schlinker, R. H., Reba, R. A., and Simonich, J. 2010. "Unstructured Large-Eddy Simulation Technology for Prediction and Control of Jet Noise," *Proceedings of ASME Turbo Expo 2010*, Glasgow, UK, GT2010-22306.

Khalighi, Y., Ham, F., Moin, P., Lele, S. K., Schlinker, R. H., Reba, R. A., and Simonich, J. 2011. "Noise Prediction of Pressure-Mismatched Jets Using Unstructured Large-Eddy Simulation," *Proceedings of ASME Turbo Expo 2011*, Vancouver, Canada, GT2011-46548.

Khalighi, Y., Nichols, J. W., Ham, F., Lele, S. K., and Moin, P. 2011. "Unstructured Large-Eddy Simulation for Prediction of Noise Issued from Turbulent Jets in Various Configurations," AIAA paper 2011-2886.

Krothapalli, A., Venkatakrishnan, L., and Lourenco, L. 2000. "Crackle: A dominant component of supersonic jet mixing noise," AIAA Paper 2000-2024.

Lockard, D. P. 2000. "An efficient, two-dimensional implementation of the Ffowcs Williams and Hawkings equation," *J. Sound Vib.*, Vol. 229, pp. 897–911.

Martens, S., Spyropoulos, J. T., and Nagel, Z. 2011. "The effect of chevrons on crackle engine and scale model results," *Proceedings of the ASME Turbo Expo*, GT2011-46417.

Nichols, J. W., Ham, F. E., and Lele, S. K. 2011. "High-Fidelity Large-Eddy Simulation for Supersonic Rectangular Jet Noise Prediction," AIAA paper 2011-2919.

Nichols, J. W., Ham, F. E., Lele, S. K., and Bridges, J. 2012. "Aeroacoustics of a supersonic rectangular jet: Experiments and LES predictions," AIAA paper, 2012-678.

Nichols, J. W., Lele, S. K., Moin, P., Ham, F. E., and Bridges, J. E. 2012. "Large-Eddy simulation for supersonic rectangular jet noise prediction: effects of chevrons," AIAA paper 2012-2212.

Nichols, J. W., Lele, S. K., and Spyropoulos, J. T. 2013. "The source of crackle noise in heated supersonic jets," AIAA paper 2013-2197.

Nichols, J. W., Lele, S. K., Ham, F. E., Martens, S., and Spyropoulos, J. T. 2013. "Crackle Noise in Heated Supersonic Jets," J. of Eng. for Gas Turbines and Power.

Data Analytics and Visualization for Large Field Test Datasets

by Ken Renard and Jerry Clarke, US Army Research Laboratory

HPC Resources: Appro Xtreme (*Abutil*, *Cabutil*), IBM iDataPlex (*Hercules*), ARL DSRC

Systems-of-systems evaluation is critical for mission success as individually complex systems are combined in even more complex environments. Analysis of test data can yield essential information for enabling interoperability, improving performance, and engineering resilient combat systems. In the field of digital communications, large amounts of valuable test data can be generated, often with significant latent information. The processing and analysis of these large datasets becomes a high value endeavor. Our growing dependence on digital communications and the availability of highly sophisticated communications systems requires more scalable methods to process test data into a form that analysts can use to make critical decisions.

Data collection from a “small” test of 20-30 networked radio systems can yield on the order of 1 Terabyte of data for a single 12-hour test period. Processing this data into a data model for use by analysts can take up to 2-3 days. With larger, more complex test scenarios becoming necessary, the current processing approach will soon become overwhelmed. The Army Research Laboratory (ARL) and Aberdeen Test Center (ATC) have partnered to utilize high performance computing (HPC) to address the growing requirements for data processing and analysis. Using the HPC resources at ARL’s DOD Supercomputing Resource Center (DSRC), a framework has been designed for distributed processing and analysis of test data that can scale to large data sets. A series of map-reduce operations distribute the processing load of data among hundreds (up to thousands) of processing cores. Each map-reduce cycle can decompose the problem and data in different domains (such as time, data source, packet attributes, etc.) to achieve higher levels of parallelism. The framework uses pluggable processing modules that can perform various types of predetermined data analysis functions, as well as building a reduced data model that can be studied in depth by analysts.

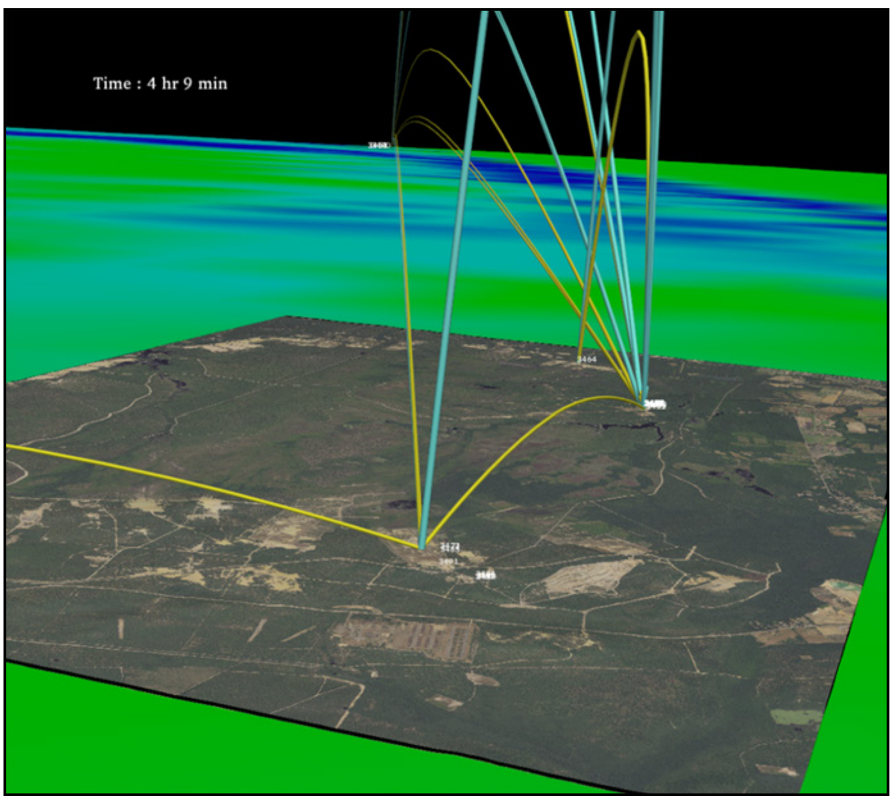

Figure 1. Geographic test data visualization depicting conversations between platforms.

The existing code was built with Java to process data serially, creating many intermediate database tables. The input/output (I/O) requirements of this method were easily overwhelming the database with the sheer quantity of operations. The code was rewritten from scratch using Python and Message Passing Interface (MPI) to distribute the load and I/O requirements among many processing cores. The result is a framework with several packet processing modules that match multiple instances of packet captures, allowing for completion percentages and latency values to be determined. Higher layer analysis measures the performance of the Transmission Control Protocol (TCP) over these *ad hoc* networks in the presence of a dynamic routing environment and packet loss. Additionally, application-layer metrics are produced, providing insight into the effectiveness of command and control messaging over the tested network, and performance of protocols such as Voice-over-IP in these environments.

Between map-reduce phases, intermediate data must be exchanged among processors. The memory requirements for exchange are too great to avoid accessing the file system. However, I/O performance is enhanced using PyTables which employ Hierarchical Data Format (HDF) and compression for efficiency. This also allows for the storage of intermediate results which support post mortem data analysis and verification of results. Reading of the PyTable data is efficiently filtered using Structured Query Language (SQL)-like syntax. The final results are read from PyTables and ingested into the existing database schema, preserving the interface that analysts have become accustomed to.

Runs have been executed on the ARL DSRC’s utility server and IBM cluster. While the swap space available on the Utility Server compute nodes

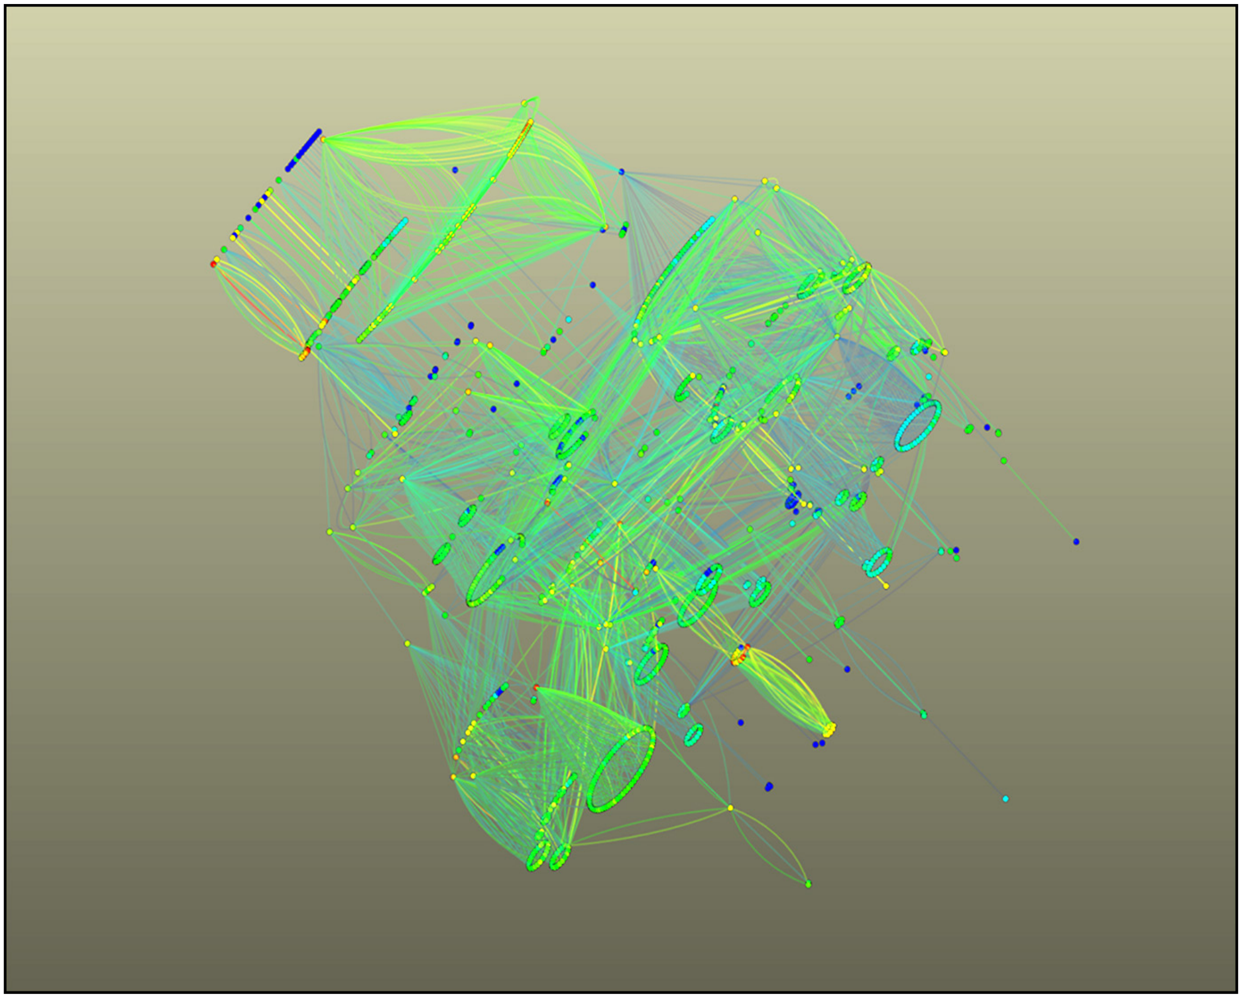

Figure 2. Spanning tree layout strategy.

prevents widely varying data sets from running out of memory, the file system performance of the IBM cluster has proven such an advantage that codes were modified to reduce the memory footprint enough to run there. Data is collected from various tests, including the Army Test and Evaluation Command's (ATEC) Network Integration Evaluation (NIE) tests at White Sands Missile Range, as well as other Army test centers. It is sent back via Secure Defense Research and Engineering Network (SDREN)/Defense Research and Engineering Network (DREN) to Aberdeen for processing.

During the spring 2013 tests, the HPC codes began taking over primary responsibility for key parts of the data reduction. For the fall 2013 tests, the codes assumed primary responsibility for the entire data reduction process. This approach has reduced the turnaround time for data from 50 hours to less than five. As test complexity and data reduction requirements increase, we expect to maintain or improve these tutnaround times.

In addition to producing the existing data model products with a quicker turn-around time, ARL engineers will be able to apply data analytics and visualization algorithms to the data sets that can provide better insights into trends and correlations in the network metrics. By looking at "conversations" between nodes at different layers of the network stack, trends become apparent that can help analysts decipher bottlenecks, explain behaviors apparent in other metrics, and identify opportunities for optimization. With large datasets such as these, the organization of a visual representation becomes critical to gaining useful insight. For example, a three-dimensional spanning tree algorithm is used to display IP address conversations for networks with thousands of IP endpoints. Trends and associations become immediately apparent with these types of graphical layouts. Another visualization of particular use for analysts is a geographic representation of nodes with an overlay of satellite imagery that depicts data exchanges over time. GPS data collected from nodes enables visual correlation of events that have happened in the exercise with their associated network effects. We provide the ability to filter and plot data to see effects on specific applications over time and location. HPC enables the use of these algorithms on such large data sets and the rendering of analytics to help analysts make more effective decisions on Army communications platforms.

Fatigue Crack Growth in Aircraft Structures

by James M. Greer, Jr., CAStLE Technical Director

HPC Resources: Cray XE6 (*Garnet*), ERDC DSRC; IBM POWER6 (*Davinci*), Navy DSRC

Abstract

The US Air Force Academy's Center for Aircraft Structural Life Extension (CAStLE) continues to advantageously use HPCMP assets for enhancing the safe and economic operation of the Air Force's aging fleets of aircraft. With some fleets consisting of hundreds of aircraft with more than four decades of service behind them, fleet managers need accurate data and capable tools to help them continue operations. By generating accurate stress-intensity solutions for a wide range of damage geometries, and high resolution finite element models, CAStLE helps fleet managers realistically assess the structural health of their aircraft.

Scope and Principal Objectives

Recent work under this project has focused on the development of high-accuracy crack stress-intensity solutions via the finite element method and large-scale structural analysis of damage propagation in complete aircraft sections with all parts modeled in detail (Giga-DOF-sized problems). *CAStLE* has a portfolio of projects that leverage, and are leveraged by, HPC-centric work. The overarching objective of the HPC work is to provide tools and data for physics-based modeling of crack growth behavior, residual strength, residual life and, in the long-term, material degradation associated with corrosion kinetics. Realistically predicting the advancement of corrosion in the spatial and time domains is an extremely complex and challenging task, due primarily to the scarcity of accurate inputs (with environmental exposure parameters being chief among these).

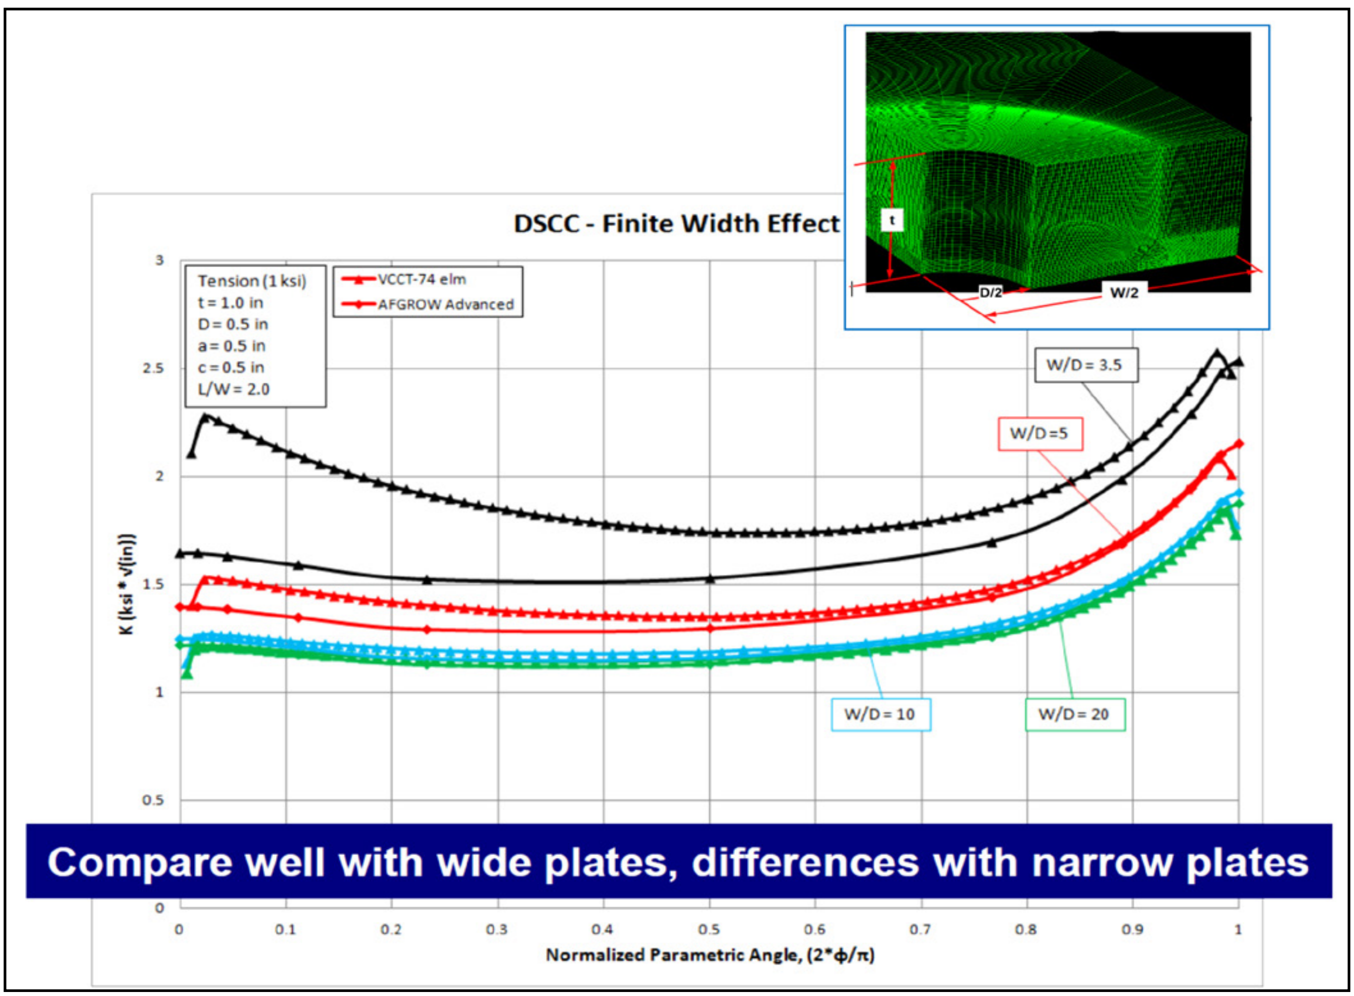

Figure 1. Differences in K as a function of W/D [3]. Inset shows typical mesh for quarter symmetric model where t, W, and D are specimen thickness, specimen width, and hole diameter.

Methods and Approach

The methods for the prediction of structural degradation center on the finite element (FE) method: both *h*-version (more elements for better convergence) and *p*-version (higher order elements for better convergence). HPC assets provide platforms that make solving these very complex and mathematically-intensive FE problems tenable. Work by Fawaz and Hammond [1–3] has focused on extracting stress-intensity (*K*) solutions in cases where the flaw length is significant in comparison to the width dimension of the structure—a seldom explored part of the crack geometry solution space. This work was motivated by the discovery that currently used *K* solutions were inaccurate for these geometries. Using the Virtual Crack Closure Technique (VCCT) and the FE method, they found significant differences between their results and results obtained from the AFGROW crack growth software tool. As seen in Figure 1, the agreement with legacy solutions is very good for larger ratios of *W/D* (wider plates), but not for narrower ones. This is a very important result, as small variations in *K* can lead to very large variations in part life.

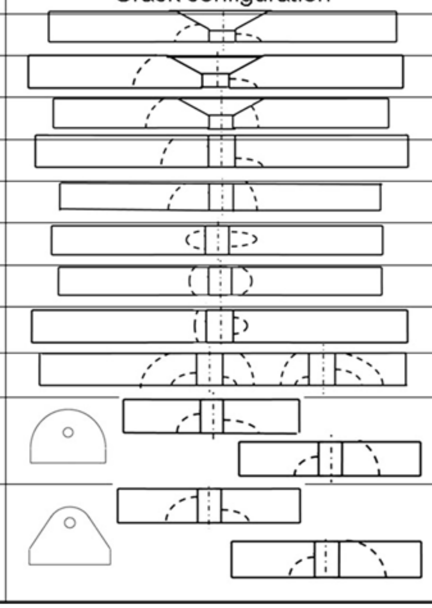

Case	Crack configuration	Cracks	Solut.(M)	Solut.(M)
1		1-2	7.9	7.9
3		1-2	11.0	11.0
5		1-2	7.7	7.7
2		1-2	0.7	0.7
4		1-2	0.3	0.3
6		1-2	1.0	1.0
7		1-2	0.3	0.3
8		1-2	0.7	0.7
9-32		1-2(4)	38.4	1788.0
33-34		1-2	0.9 0.7	0.9 0.7
35-36		1-2	1.8 1.6	1.8 1.6

Figure 2. Hole and lug crack geometries of interest in HPC analyses.

In parallel with this work is an effort to find *K* solutions for a wide variety of scenarios involving cracks at holes in plate structure and lugs (very common in aerospace structures). Using an *hp*-version code called "STRIPE"; Andersson is continuing HPC work done with Fawaz [4–9]. Hole/lug/crack geometries of interest include those shown in Figure 2. To date, about two billion solutions in the "Case 9-32" family of geometries have been generated, and work continues.

The STRIPE code is also very well-suited for solving structural FE models having billions of degrees-of-freedom. Work continues on development of an extremely detailed model of a C-130 aircraft center wing box [10]. The

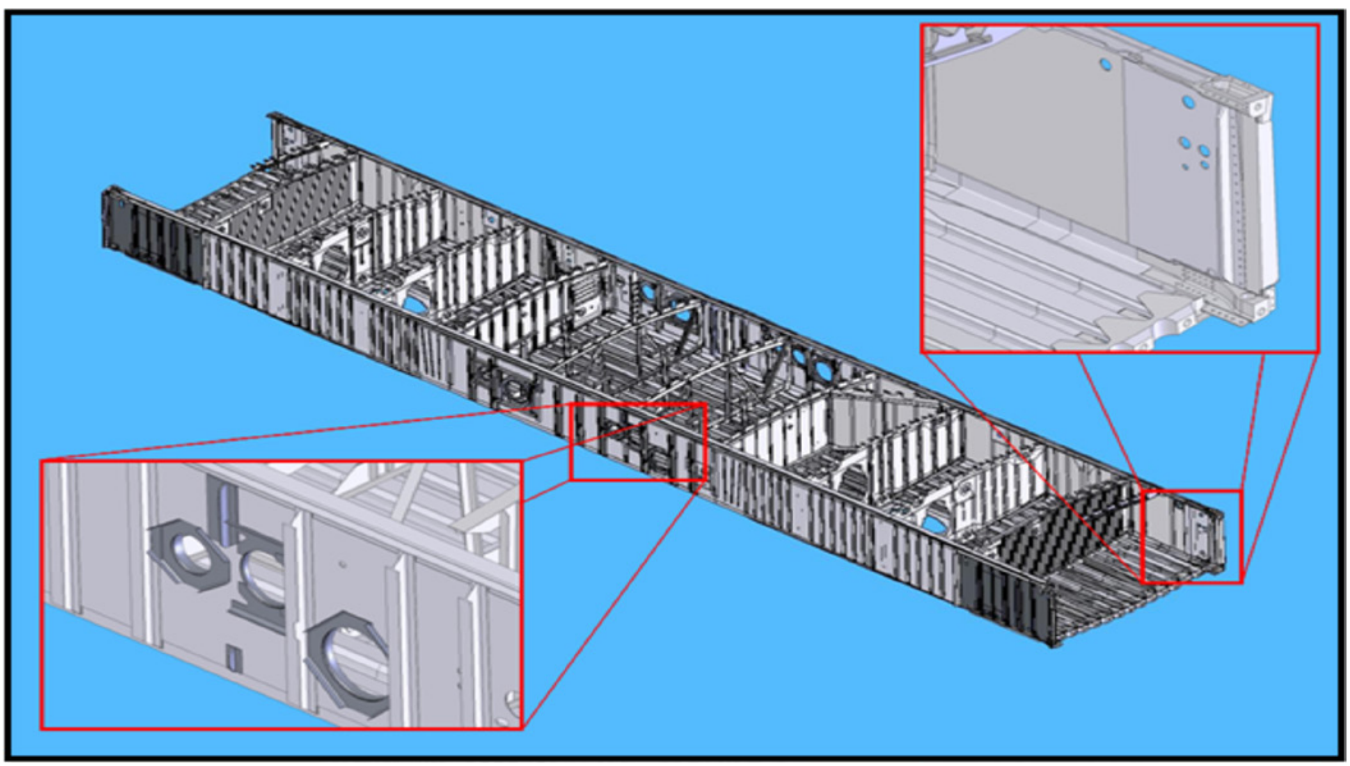

Figure 3. Finite element model of C-130 center wing box structure.

model, when complete, will include approximately 50,000 explicitly modeled fasteners in addition to the hundreds of parts, with everything modeled using solid elements (Figure 3).

The goal of this effort is to give fleet managers a means of assessing damage as it relates to residual life and strength. It may provide a means, for example, of clearing a battle damaged aircraft for a one time ferry flight, or for otherwise determining the impact of well characterized damage.

Summary and Conclusions

This vital HPC effort continues to provide fleet managers with tools and data to more effectively manage aging fleets of aircraft safely and economically. Accurate stress-intensity solutions are the linchpin in accurate crack growth prediction, and crack growth prediction tools are relied upon heavily by fleet managers. They are used to quantify risk, establish inspection intervals, determine economic life, and ultimately safely accomplish the missions generated by military leadership.

Value and Impact to DOD

As accurate *K* solutions continue to be generated, crack growth simulation tools, like AFGROW, are able to provide more accurate structural analyses for fleet managers. Overly conservative analyses result in too frequent inspections and aircraft downtime, with the accompanying loss of mission readiness and increased cost. Non-conservative analyses put lives and valuable hardware at risk. The current work aims to narrow the band of uncertainty, improving economic performance while enhancing safety.

Value and Impact of HPCMP Resources to the Project

The models used by this project to generate these important tools and data could not be practically solved without HPCMP resources. They inherently require large amounts of shared memory and fast execution of millions and millions of floating point operations. These resources speed deployment of critical solutions to the fleet managers who so vitally need it. The value in these resources comes primarily from their sharing. It obviates the need for disparate organizations to purchase their own very high end workstations or supercomputers.

References

[1] Fawaz, S.A. and Hammond, M.J., "Stress-Intensity Factors in Finite-Width Plates," presented at the 2012 AFGROW Workshop, 11 September 2012, Layton UT (www.afgrow.net/workshop/documents/2012/Scott_Fawaz_AFGROW_Workshop_2012.pdf).

[2] Hammond, M.J. and Fawaz, S.A., "Test Method for Pit-to-Fatigue Crack Transition: Stress-Intensity Values for Narrow Finite-Width Plates," SAFE, Inc. Contractor Presentation, 25 July 2012.

[3] Hammond, M.J. and Fawaz, S.A., "Stress-Intensity Factors for Finite-Width Plates," *Proceedings of the 2012 Aircraft Airworthiness & Sustainment Conference,* Baltimore MD, 5 April 2012.

[4] Fawaz, S.A. and Andersson B., "Accurate Stress-Intensity Factor Solutions for Corner Cracks at a Hole," *Engineering Fracture Mechanics 71*, Number 9, June 2004, pp. 1235-1254(20).

[5] Andersson, B. and S. A. Fawaz, "Statistical Fatigue and Residual Strength Analysis of New/Aging Aircraft Structures," *The 2004 High Performance Computing Modernization Program Users Group Conference*, Williamsburg, VA, 2004.

[6] Andersson, B. and S. A. Fawaz, "Statistical Fatigue and Residual Strength Analysis of New/Aging Aircraft Structures," *The 2006 High Performance Computing Modernization Program Users Group Conference*, Denver, CO, 2006.

[7] Andersson, B. and S. A. Fawaz, "Statistical Fatigue and Residual Strength Analysis of New/Aging Aircraft Structures," *The 2007 High Performance Computing Modernization Program Users Group Conference*, Pittsburgh, PA, 2007.

[8] Fawaz, Scott and Börje Andersson, "Fast and Reliable Solution of GDOF-Problems on NAVO/BABBAGE and AFRL/HAWK Systems," *The 2008 High Performance Computing Modernization Program Users Group Conference*, Seattle, WA, 2008.

[9] Andersson, Börje and Scott Fawaz, "Solution of Ultra-Large Structural Mechanics Problems during CAP-I 2008 on the IBM POWER6 System DAVINCI," *The 2009 High Performance Computing Modernization Program Users Group Conference*, San Diego, CA, 2009.

[10] M. Hammond, B. Andersson, S. Fawaz, J. Greer, Jr., and R. Rainsberger, "Detailed Three-Dimensional Modeling of the C-130 Center Wing Box for Damage Tolerance Analyses," *Proceedings of the 2010 Aircraft Airworthiness and Sustainment Conference*, Austin, TX (2010).

HPC Centers Website and User Dashboard Offer Informative, Tailored Information to HPCMP Community

by Christine E. Cuicchi, Associate Director for HPC Centers

DOD HPCMP users, staff, and visitors now have a new, centralized resource for HPC Centers-related information they most often need and use, including information tailored to their everyday workflow on our HPC systems.

HPC Centers Website

The HPC Centers team, which oversees the five DOD Supercomputing Resource Centers, launched centers.hpc.mil in the fall of 2013. Users, staff, and other visitors are able to quickly find information and documentation about the HPC systems hosted at the five DSRCs, and have fast access to the broad array of tools, services, and expertise available to users within the HPCMP, including accounts information and information on the Consolidated Customer Assistance Center (CCAC). The site replaces the CCAC web page, which was retired in September 2013.

The new HPC Centers site features:

- A central place to find information about all unclassified HPC systems in the HPCMP
- A systems summary page which shows the up, down, or degraded status of all unclassified HPC systems
- Information and quick links to HPC Domain Expertise, Tools, and Services available to our users
- News and publications about HPC systems and activities within the HPCMP

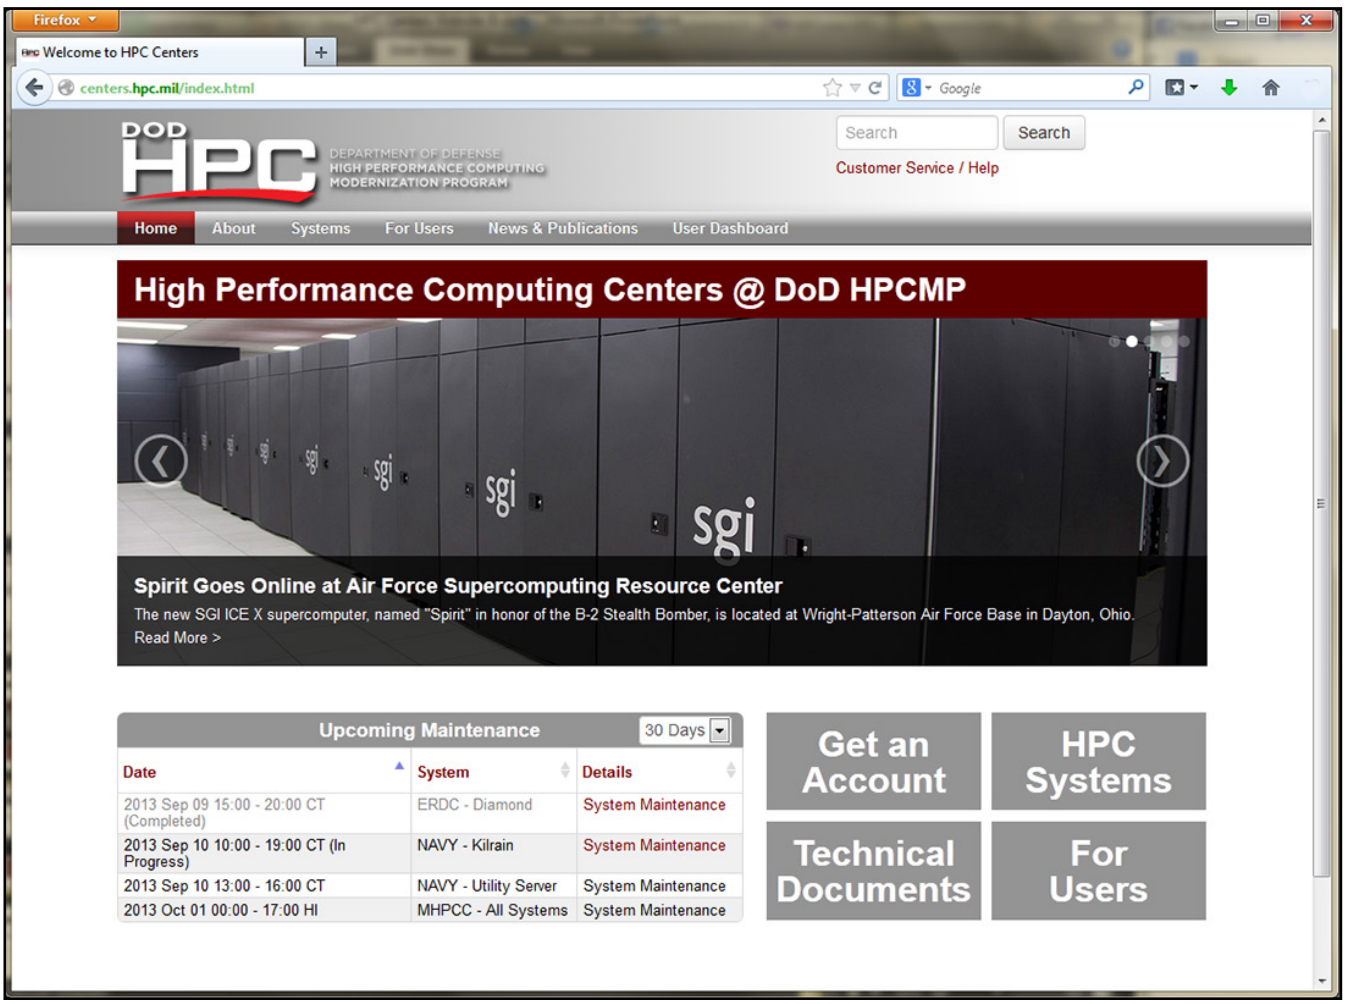

Figure 1. Home page of HPC Centers website features news stories, system maintenance, and quick links.

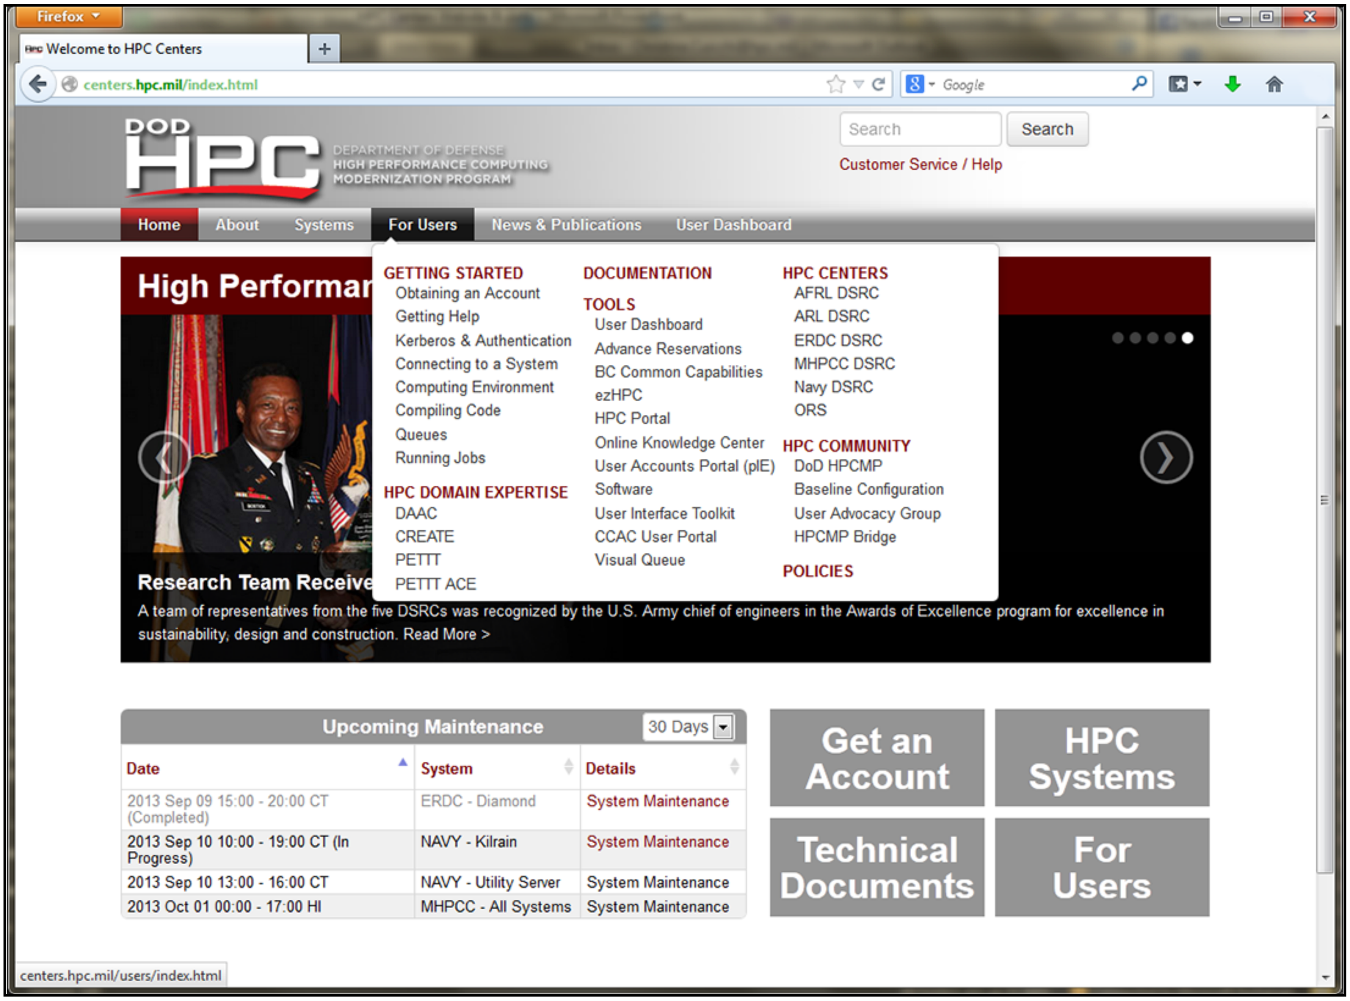

Figure 2. New menu designs help users quickly find the information they are looking for. The Systems menu provides information about our HPC systems. The For Users menu provides information on Documentation, Tools, Getting Started, Domain Expertise, HPC Centers and HPC Community, and Policies.

User Dashboard

The User Dashboard, https://dashboard.hpc.mil, provides users with access to real time information about their accounts and jobs—typically accessed via command line interfaces—in a single place. The Dashboard also houses information that requires authentication to access, and is accessible to anyone in the DOD with an active portal to the Information Environment (pIE) account, even those who do not currently have an active account on HPC systems.

Users enter the Dashboard using the OpenID authentication system, which works with your CAC, YubiKey, or SecurID card. Once in the Dashboard, users see modules with information tailored to them:

- Utilization of the systems on which they have accounts
- A comprehensive listing and status of the jobs they have queued, are running, or completed on those systems
- The number of COTS software licenses available to them via the Software License Broker
- Listing and status of any Advance Reservations they have made
- Maintenance for their HPC systems

The Dashboard is a great start, but it is still a work in progress—we need your input and ideas for new information modules and improvements. Please send your ideas, suggestions, and feedback to feedback@centers.hpc.mil.

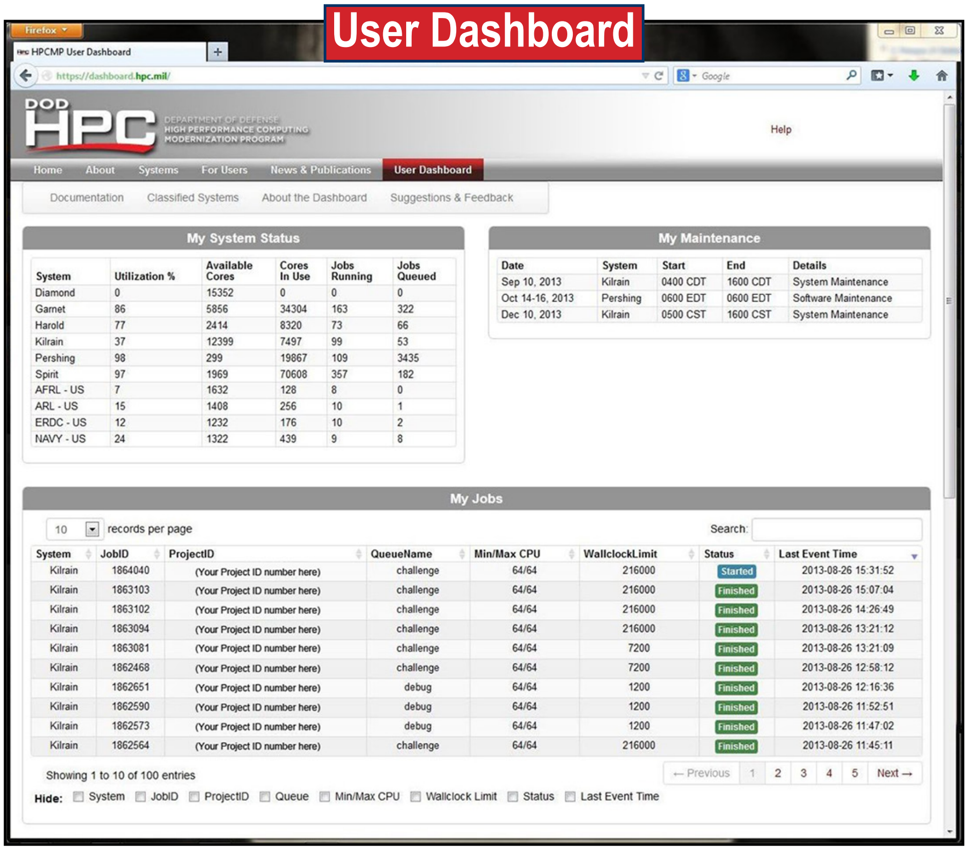

Figure 3. The User Dashboard as it appears to a user with active accounts on several HPC systems across the HPCMP. The My Jobs module lists all the jobs a user has submitted on HPC systems, and is searchable. An actual user will see the project ID numbers affiliated with his or her jobs.

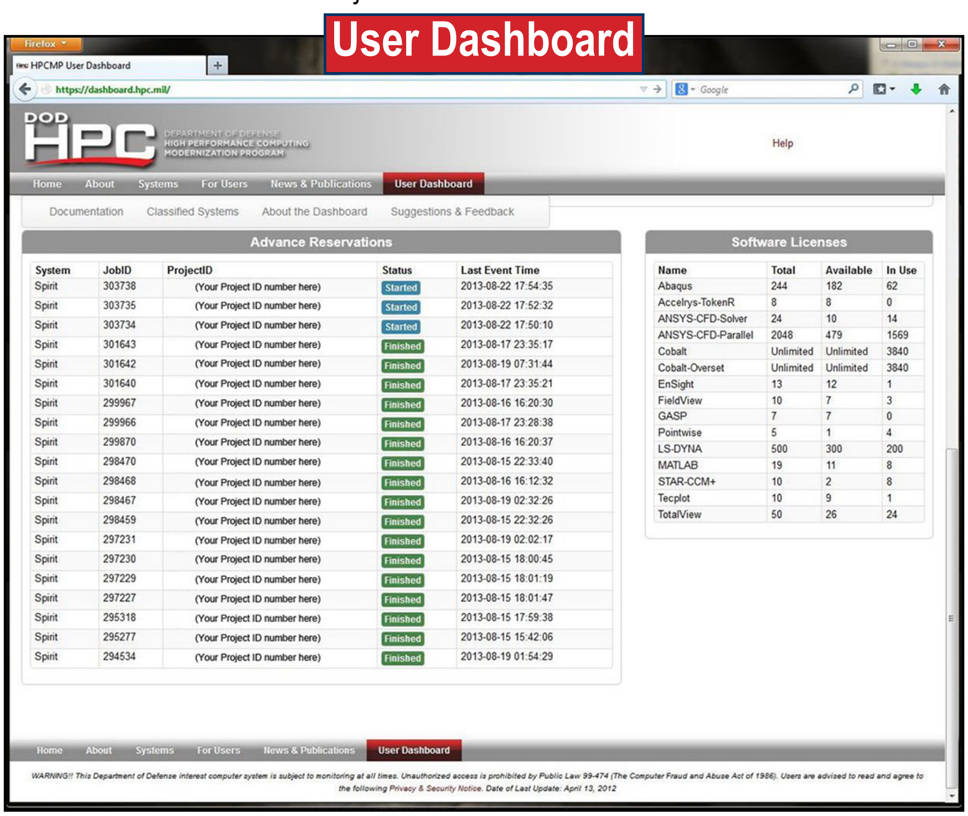

Figure 4. Additional modules available on the User Dashboard show a user's Advance Reservations and the number of commercial software licenses available to the user.

ARL Marks Opening of New Facility

by Dr. Raju Namburu, Director ARL DSRC

On Monday June 10th, 2013, the Army Research Laboratory's (ARL's) Computational and Information Sciences Directorate, Computational Sciences Division, hosted an open house of the ARL Supercomputing Research Center facility. The event marked the formal opening of ARL's newly acquired supercomputing center and featured exhibits of ARL's supercomputing history and demonstrations of the current role of supercomputing and computational science research in the DOD Army RDT&E community for the design, development, testing and rapid fielding of emerging military technologies and soldier systems.

Formal dedication and ribbon cutting of the ARL Supercomputing Research Center From left to right: Mr. Robert Carter, U.S. Army ATEC; Dr. Thomas Russell, ARL Director; Dr. Patricia Falcone, Office of Science and Technology Policy, White House; BG Daniel Hughes, CG U.S. Army RDECOM; Sen. Ben Cardin (D-MD); COL Gregory McClinton, U.S. Army CECOM; Lt. COL Richard Moyers, ARL SLAD; Dr. Raju Namburu, Chief, ARL Computational Sciences Division and Director, ARL DSRC, Mr. John West, Director, DOD HPC Modernization Program; Dr. John Pellegrino, Director, ARL Computational and Information Sciences Directorate.

The ARL Supercomputing Research Center facility was provisioned to ARL through the proactive efforts of senior laboratory management and the strong support of top civilian and military leadership at Aberdeen Proving Ground. One of the most appealing features about our newly acquired facility is that it provides the footprint to design and plan for a scalable HPC infrastructure for the ARL DSRC as well as ARL's advanced computing research programs. In addition to providing a scalable supercomputing environment, a key benefit of the facility is the ability to host the entirety of ARL's HPC resources in one facility, rather than three non contiguous facilities. By FY16, we expect to realize notable cost savings and energy efficiencies of operations and maintenance for our HPC resources.

The open house featured a keynote talk by Dr. Patricia Falcone, Assistant to the President for Science and Technology, Director of the White House Office of Science and Technology Policy; and speeches by Sen. Ben Cardin (D-MD); and ARL Director, Dr. Thomas Russell; as well as several other senior military and civilian stakeholders in the ARL HPC Center. The logistics of the event were directed by the ARL Public Affairs Office, and were conducted with all of the pomp and circumstance of an official Army ceremony complete with The Army Band, Color Guard and ribbon cutting.

After the ribbon cutting, the distinguished visitors and guests were provided tours of the facility to include exhibits and demonstrations of past and current examples of ARL's role in high performance computing and advanced networking technologies. One of the most well received exhibits was the Army Research Laboratory

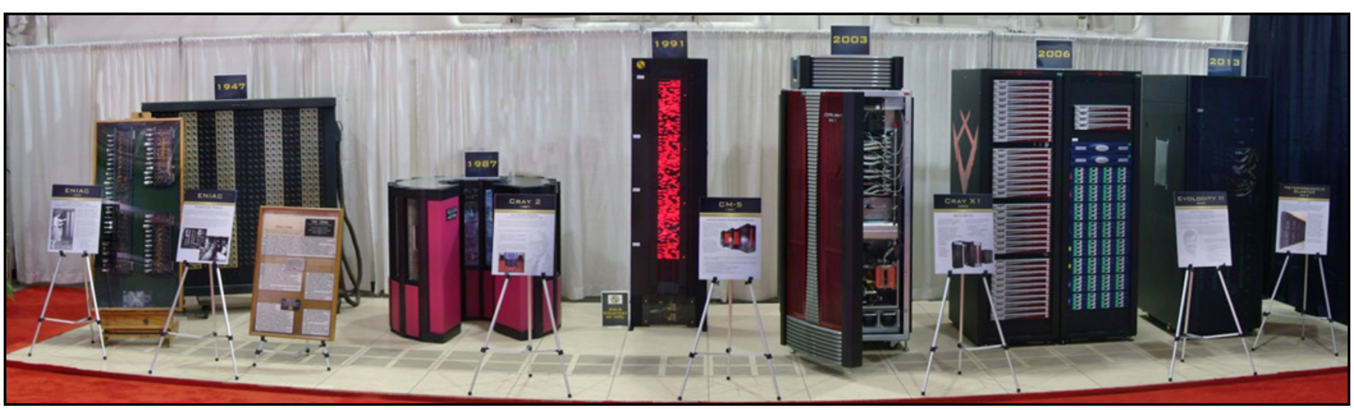

ARL Supercomputing Research Center Open House Computing History Exhibit. From left to right: ENIAC Decade counter (1946), ENIAC Function Table (1946), Cray-2 (1987), Thinking Machines CM-5 panel (courtesy of NRL), Cray X1 (2003), Linux Networx Evolocity (2007), IBM Heterogeneous Cluster (2013)

Computing and Networking History Exhibit. The exhibit featured the timeline of ARL's 70 year history of scientific computing beginning with the ENIAC in 1946, through today's current state-of-the-art supercomputing systems. The display area featured several artifacts from the world's first general-purpose computer, the ENIAC, and other historically significant scientific computing systems in ARL's HPC lineage. For each computer system, examples of the breakthrough scientific advances enabled for various military applications by these computational systems were highlighted. The exhibit also contained a timeline which chronicled the pioneering efforts of ARL researchers in the evolution and development of modern networking (from ARPANET to the Army Supercomputing Network (ASNET) to the DREN), as well as ARL's current research in next generation advanced networking.

Other exhibits included the DOD High Performance Modernization Program, ARL's Next Generation Advanced Networking Research Initiatives, ARL's Battlefield Network Modeling Research, Novel Materials Modeling, ARL's 'Green Computing' Initiatives, and the Center's robust Outreach and Student Programs in High Performance Computing.

As you might imagine, it was a very proud day for the ARL, the HPCMP, the ARL DSRC, and the computational science research community who have the privilege, honor and commitment to support our soldiers and our national defense.

Energy Aware Scheduler Demonstrates Innovative Power Savings

by Kathy Smith and Tom Kendall, ARL DSRC; Michael Knowles, Lockheed Martin

In FY12, the HPCMP's Energy Aware Scheduler (EAS) project was initiated as a one year project under the HPCMP's Green HPC Initiative in collaboration with the U.S. Army Research Laboratory, Lockheed Martin, Altair and Instrumental. The goal of the project was to evaluate the feasibility of reducing power and cooling costs, and the associated environmental impacts that resulted from the energy used by idle compute nodes. The HPC systems across the DSRCs run on an average 80-90% busy, and the DSRC Utility Servers (US) run on the average 20% busy; making it attractive to reduce the power requirements on the portion of the systems that are idle. Scheduling of user reservations, system maintenance, and scheduling of large node count jobs all create opportunities to power off idle nodes. The EAS capability is architecture-specific and works with the current Program wide scheduler, Altair PBS Professional™ Scheduler, by intelligently powering off resources (nodes) that are idle and not actively being used by running jobs or reserved for near future jobs. The team tested the resiliency and feasibility before and after EAS was installed on the SGI Altix ICE, Cray XE6 and Appro Xtreme (Utility Server architecture) by running benchmarks, with the results showing no significant performance degradation in job execution times.

Typically idle nodes consume about 50% as much energy as active nodes, and it is attractive to power off resources that are idle to save energy. Altair's PBS Professional™ Scheduler already incorporated a capability that was oriented toward power monitoring and control. We integrated the EAS with Altair PBS Professional™ Scheduler to control and minimize the power required by resources that are not in use. After performance testing the original Perl version, Altair ported the exiting EAS code to a Python implementation that was better at delivering the requirements of the HPC systems and provided improved performance on what the PBS scheduler anticipates will happen in the future, rather than trying to make the determination itself.

EAS has been operational at the ARL DSRC on the SGI Altix ICE and Appro Xtreme architecture since early in 2012. What we have typically seen is there is good opportunity at both the early and late stage of the system life cycle to save energy. During the middle of the system lifecycle, HPC systems tend to run at least 80% busy and there is less opportunity to power-off nodes. Scheduling of user reservations, system maintenance and scheduling of large node-count jobs all create additional opportunities to power off idle nodes. Since 2012 the ARL DSRC has saved a total of 349190 kWh and a dollar savings of approximately $35k at $0.10/kWh across the systems.

EAS was configured for a three month trial period from May-July 2013 on five of the DSRCs Utility Servers at AFRL, ARL, ERDC, and NAVY to control and minimize the power of resources that are not in use, testing the effectiveness and reliability of the system with the scheduler. The Utility Servers are fairly new to the Program and as such not highly utilized. There is excellent opportunity to power down idle nodes to save energy. EAS has configuration parameters that are easily tunable to allow the scheduler to meet site specific needs. Some of the configuration parameters include how long idle nodes stay on before powering-down, maximum number of nodes to power down in each cycle, how far in advance a currently powered off node can be powered on if a job will need it, etc. Since the Utility Servers are geared to always be up and available to run user jobs with little or no waiting, EAS is configured with a keep alive feature that keeps a minimum number of nodes always up and available. Also, given that new jobs are typically not submitted during weekend and evening hours suggests that we have additional idle nodes that are candidates for power down and additional savings. On the ARL Utility Server there is an average of 53 nodes powered off daily. The daily usage across the DSRC Utility Servers during the three-month trial period was 1 May – 31 July 2013. During the three month trial period, EAS saved approximately $14,500 across the particpating DSRCs. The annual savings is estimated at $58k. EAS usage and savings on the five Utility Servers are illustrated in Figures 1-5.

In this article we provided a brief overview of the efforts to date to deploy the EAS on the DSRC Utility Servers. Implementing the Energy Aware Scheduler on these systems has thus far saved thousands of kWh and has the potential to save in reducing associated cooling requirements, with very little fallout from hardware failures. With this enhanced power efficiency, we can lower the energy costs at the DSRCs. Early results have proven there is potential to save a projected millions of kWh throughout the HPCMP. Clearly, if the HPCMP is to continue to successfully meet the majority of requirements in the Department of Defense, efforts focused on reducing energy cost need to be sustained.

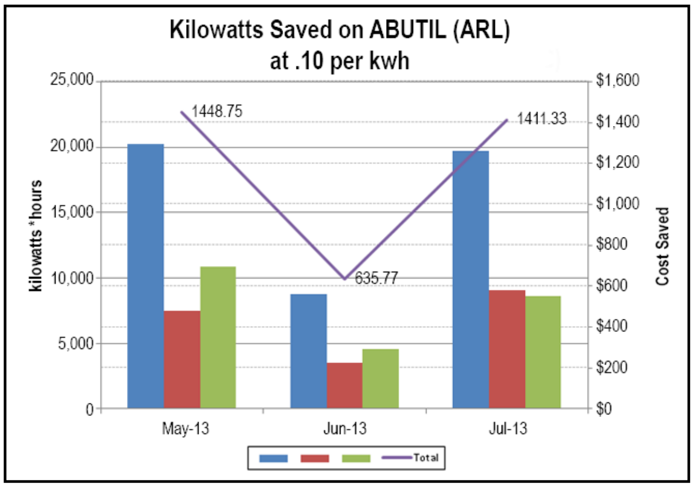

Figure 1. ABUTIL (ARL)

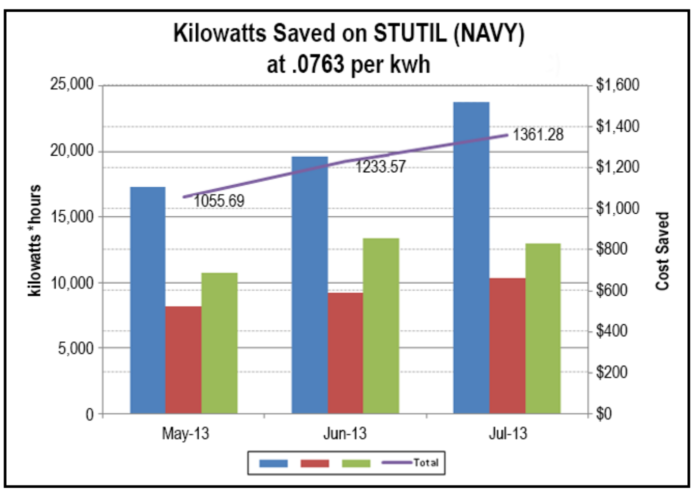

Figure 2. STUTIL (NAVY)

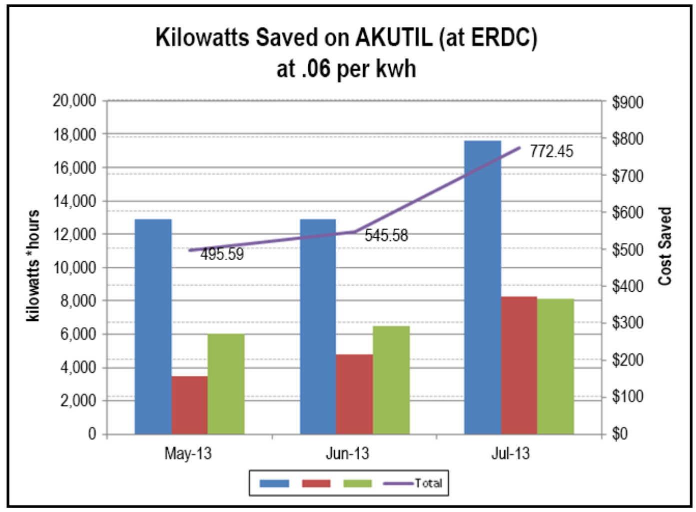

Figure 3. AKUTIL (ERDC)

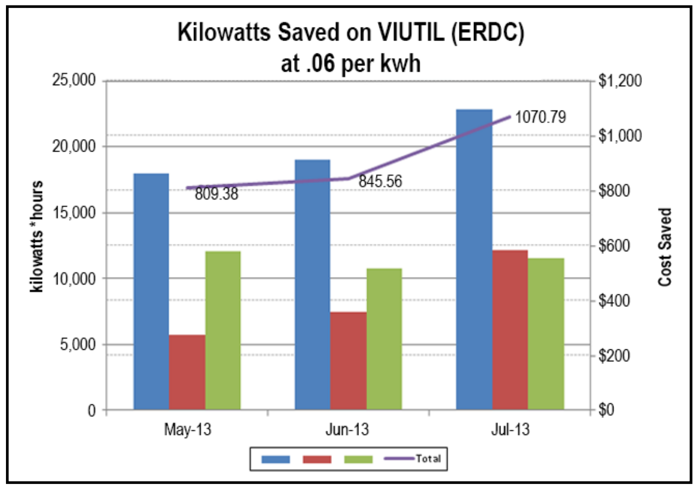

Figure 4. VIUTIL (ERDC)

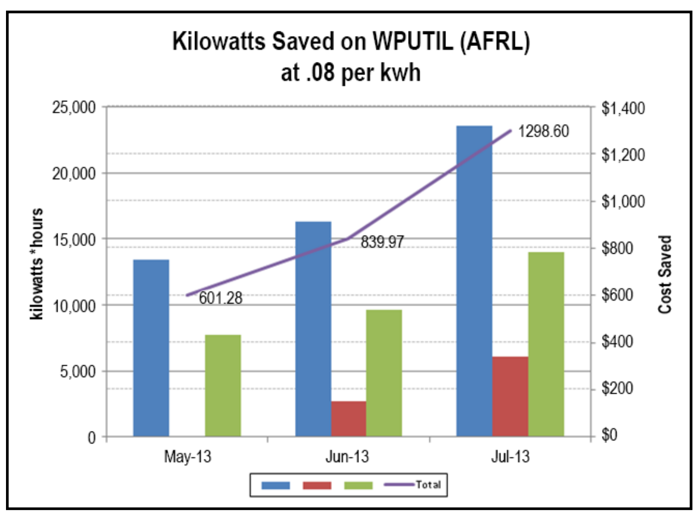

Figure 5. WPUTIL (AFRL)

Methods for Improving Software on *Garnet*

By Dr. Megan L. Cason, ERDC DSRC Computer Scientist

New documentation has been added to the ERDC DSRC website (www.erdc.hpc.mil) as part of a larger effort to transition the usage model of the shared supercomputing resource Cray XE6, *Garnet*, toward capability class computing. That process involves changes in the management of the system and changes in user practices to enable next-generation supercomputing.

New documents can be found on the ERDC DSRC website under the Tips & Tricks section (www.erdc.hpc.mil/docs). This article presents the purpose of two of these new documents.

"Introduction to Parallel I/O, Lustre, and ADIOS"

The first document addresses users who need more advanced I/O behavior from their application to scale to larger size jobs. In this context, more advanced means more parallel, or more intelligent, or both.

The most common reasons for using more advanced parallel I/O in applications on *Garnet* are as follows:

- Speed up I/O during execution. For example, a simulation that outputs data for visualization and analysis at regular intervals would benefit from speeding up the write phase.
- Control file system access globally, instead of on a per core basis. For example, some parallel I/O libraries allow a user to tune how many simultaneous file accesses occur, independent of number processes.
- Make code more maintainable. Using parallel I/O libraries instead of writing out I/O logic allows the library implementation to change without affecting the main code.

The ERDC DSRC website Tips & Tricks section contains a document presenting an introduction to parallel I/O concepts in the context of high performance computing. "Introduction to Parallel I/O, Lustre, and ADIOS" focuses on the Cray XE6 system, *Garnet*, which uses the Lustre file system to support fast parallel data access by user jobs.

In addition to basic parallel I/O concepts, this document includes a brief introduction to the ADIOS parallel I/O library. ADIOS is a continually evolving tool to improve parallel I/O performance in HPC applications, reduce code complexity required for parallel I/O tasks, and allow applications to reconfigure I/O tasks without recompiling (www.olcf.ornl.gov/center-projects/adios/).

"Introduction to Checkpointing for Capability Computing"

This second document addresses the need for warm restart functionality in software running on *Garnet*.

The primary reasons for adding restart functionality to parallel applications are as follows:

- Provide fault tolerance in the face of node failures or other system level errors, which become more likely as jobs run longer or on more nodes.
- Strengthen the software for future increases in parallel width of jobs.
- Provide rollback points that may help in deducing sources of computational errors.

A secondary reason for warm restarts applies only to a batch scheduled share resource. Batch scheduling systems may require lower maximum job runtimes to fairly schedule very large jobs. In that case, the maximum job runtime may be lower than the time needed to complete computation. Software must checkpoint its state at least once toward the end of the time allotment, and then restart in a new job.

The ERDC DSRC website Tips & Tricks section contains a document presenting an introduction to checkpointing in the context of high performance computing. "Introduction to Checkpointing for Capability Computing" provides methods to calculate how often to checkpoint running jobs, describes how checkpointing overheads affect a job and the system, and outlines a procedure to add checkpointing to an application by hand.

Detailed information on both of the above tips is located at www.erdc.hpc.mil. After reaching the website, click on Documentation located in the heading, and scroll down to Tips & Tricks.

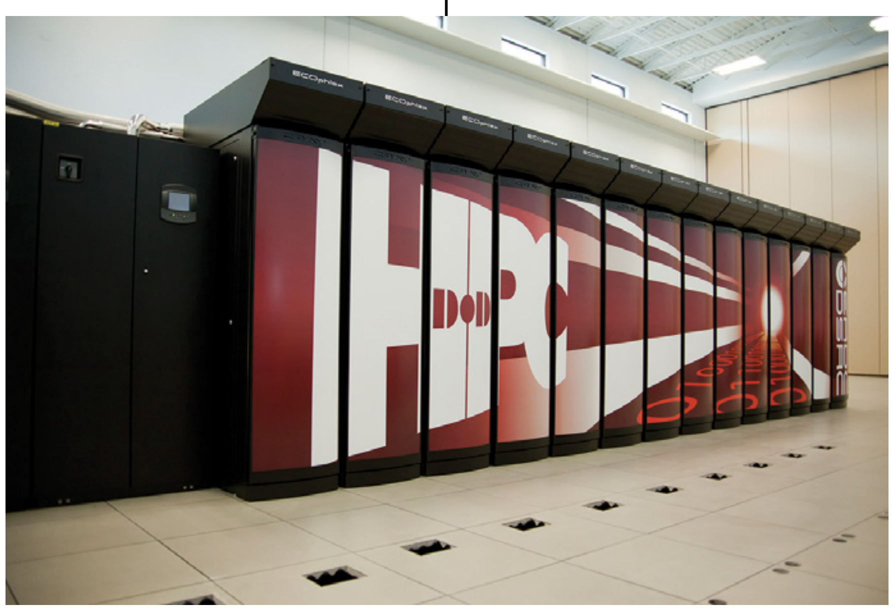

Cray XE6 (Garnet), ERDC DSRC

Technology Enhancements at MHPCC

by David Morton, Director MHPCC

In June 2013, MHPCC brought our new IBM supercomputer *Riptide* online and retired the *Mana* supercomputer. This was a significant step forward, not only from a performance perspective (peak FLOPS went up 2.4X) but also from a power perspective (power required to operate went down by 50%). The new system uses several technologies to reduce power needs, including direct warm water cooling of the compute nodes. Power costs are becoming an increasingly larger part of the overall cost of ownership of HPC systems in general, and in the current DOD financial climate it is becoming even more important to the HPCMP.

MHPCC is the primary DSRC supplying a new supercomputer delivery model called Dedicated System Partition (DSP) that is continuing to grow. This delivery model is geared toward users whose requirements don't fit in the normal batch computing environment, and delivers a dedicated cluster to those users for long periods of time. This new delivery model expands the capability of HPC delivered by the HPCMP to a new set of users and proves to be very beneficial to the DOD. You can contact Ms. Marie Greene at marie.greene@maui.afmc.af.mil for more information.

The HPCMP Supercomputing Software-as-a-Service (SaaS) infrastructure effort called HPCMP Portal is also continuing to grow and mature. The Portal infrastructure will be rolled out to all the DSRCs over the next few months. This effort has the ability to transform how users interact with the HPCMP's supercomputing resources. If you think about it, the way that users interact with supercomputers has been relatively unchanged since I logged onto my first Cray system back in 1984. In the rapidly changing tech world this is pretty remarkable. By making access easier, broader use of the DOD's supercomputing resources will be enabled, which will increase the capabilities of these incredible systems to solve additional problems for the Department.

The Portal effort is designed to reduce barriers of access and training, particularly for new users who may not want to learn about the computer science details of running their analysis efforts on supercomputers. The success of the effort will ultimately rest on the availability of applications, so considerable effort has been spent on providing tools and methodologies to port existing and new applications to this framework. A software development kit (SDK) has been developed and is available for use. The MHPCC development team has ported applications in a number of ways and are looking to consult and help other developers throughout the DOD port their applications. If you have an application that you think would benefit from easy web-based SaaS interface please contact Jeff Brown at Jeff.Brown.ctr@mhpcc.hpc.mil.

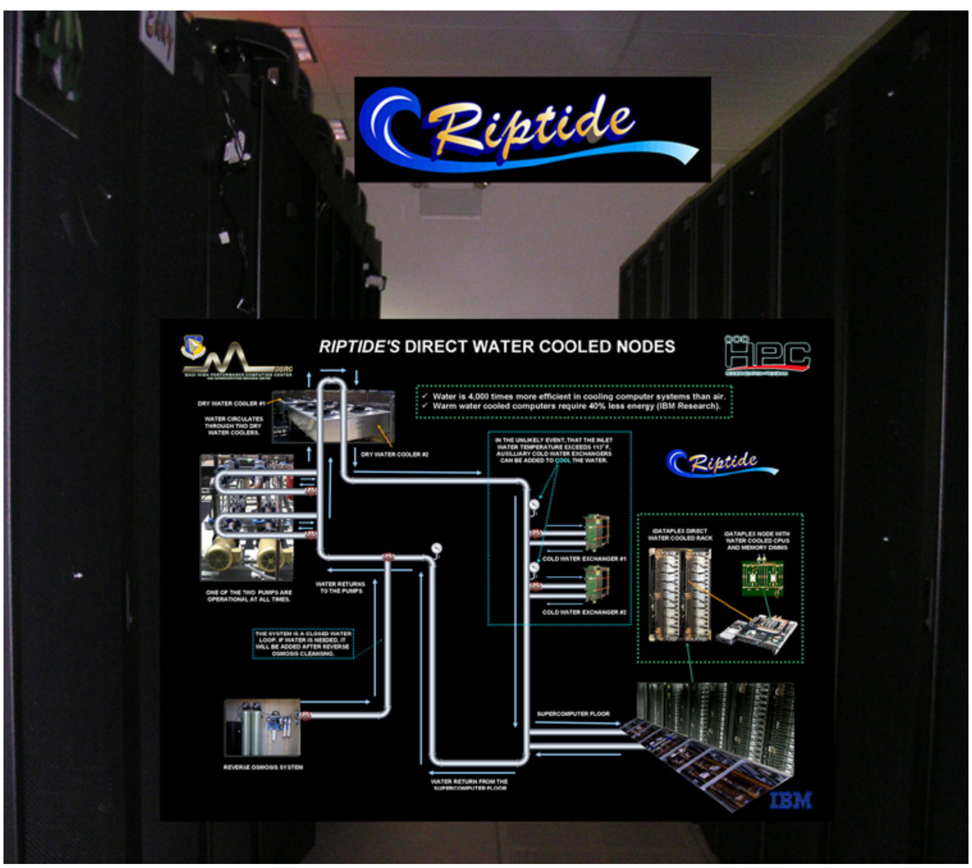

IBM iDataPlex (Riptide), MHPCC DSRC

HPCMP Portal Rollout

At the direction of the HPCMP, the MHPCC DSRC continues to take a leadership role in expanding HPC support to DOD science and engineering organizations through the use of a web-enabled “zero footprint” Portal. MHPCC is working with the USACE Engineer Research and Development Center (ERDC) DSRC, the Computational Research and Engineering Acquisition Tools and Environments (CREATE) Team, and HPCMP leadership to integrate ongoing activities. A second production version of the Portal is being deployed on the MHPCC *Riptide* System, and a development path was approved to provide portal capabilities across all DSRC’s to users in a timely and cost-effective manner. This release features the transition of the Portal front end to the Utility Server environment, with the ability to submit jobs to user selectable HPCs. An updated suite of virtual and web-based applications is provided, including CREATE-AV Kestrel versions 3 and 4, CREATE-AV Helios version 3, and the JSpOC Mission System (JMS) HPC development and test platform, *Arcade*. This Portal update also switches over from local authentication to a centralized OpenID authentication server. For more information on the HPCMP Portal roll out, please visit us at https://portal.mhpcc.hpc.mil.

TI13 Cray XC30 Machines Arrive at Two DSRCs

Workers unload the first cabinet and roll it into the AFRL DSRC.

A telehandler forklift unloading a carefully packaged cabinet for the Navy DSRC Cray XC30 system Armstrong.

In the foreground, one of four tractor trailer trucks is positioned to unload the Cray XC30 system Armstrong at the Navy DSRC.

Workers install the miles of wire connections that network the thirty cabinets of the AFRL Cray XC30.

DOD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM

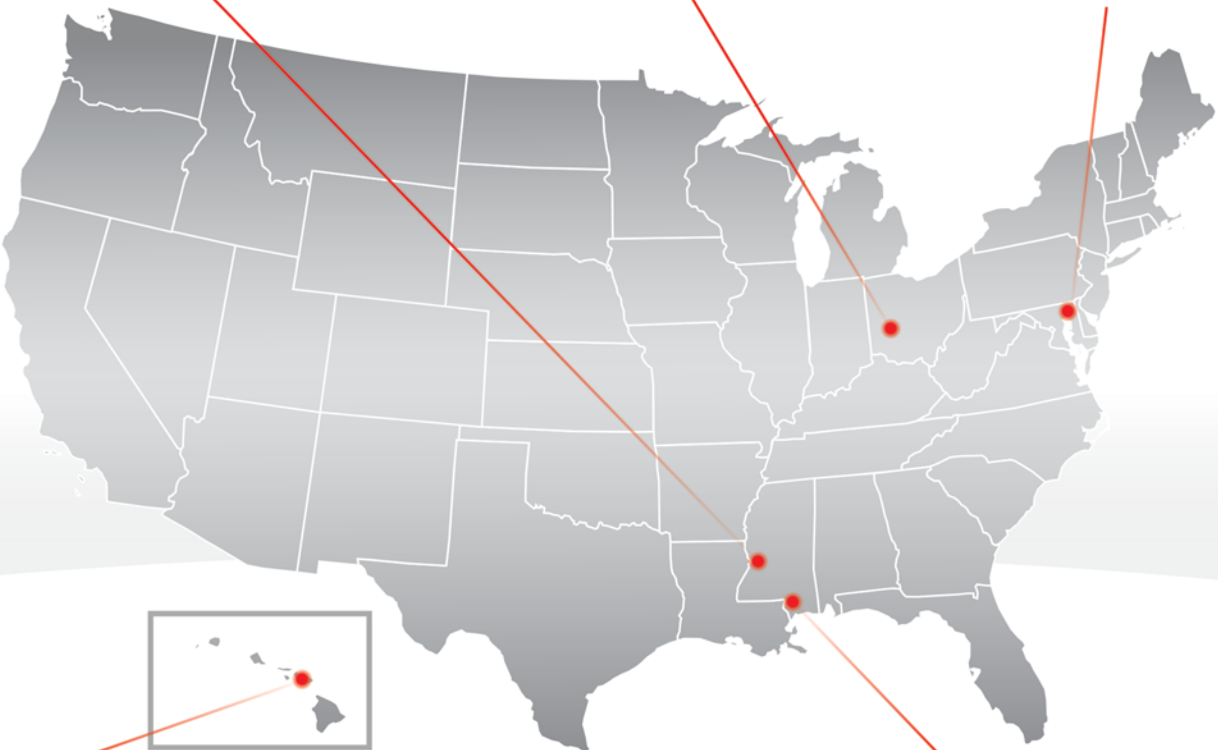

Scientists and engineers throughout the U.S. leverage the capabilities of the High Performance Computing Modernization Program (HPCMP) to solve the most time-consuming computational problems. They know that the HPCMP's supercomputing centers continually architect, deploy, and sustain their equipment to deliver world-class network and supercomputing capabilities, resulting in simulations with greater resolution and faster results than achievable on conventional workstations. You too can gain access to these powerful resources by calling to register for an account!

CCAC Accounts Center
1-877-222-2039

http://www.hpc.mil