

Department of Defense High Performance Computing Modernization Program FY16 HPCMP Annual Report



*A report by the Department of Defense
High Performance Computing Modernization Program*

Contact Information

US Army Engineer Research and Development Center

3909 Halls Ferry Road

Vicksburg, Mississippi 39180

Phone: 601-634-4692

Email: HPCMP@hpc.mil

Distribution A

Approved for Public Release; Distribution Unlimited.

Message from the Director



Dr. David A. Horner
Director, HPCMP

In FY 2016, the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) focused on several enhancements to our HPC ecosystem aimed at improving the usability, utility, and effectiveness of our investments in support of the Department's efforts in science and technology, test and evaluation, and acquisition engineering. This report highlights several of the efforts of the more than 2,400 scientists and engineers who utilized the over 7 billion core-hours of computing, high-speed networking, and advanced software provided by HPCMP.

To meet the Department's objectives in key research and engineering areas, the HPCMP initiated Frontier projects in 2014. Frontier projects receive approximately 30% of the Program's computing hours. These 11 efforts reflect the Services' and Department's highest priorities for high-performance computing. Investigations into stratified turbulence, hypersonic flows, railgun electromagnetics, and terminal ballistics are essential to the protection of warfighters and the development and deployment of the weapon systems needed to ensure battlefield dominance in the future.

This year saw a significant increase in computing capacity at our Defense Shared Computing Resource Centers (DSRC). The combined FY-14 and FY-15 Technology Insertion (TI) purchases placed 18.77 petaflops of computing capacity in the hands of DoD researchers at the start of FY-16. The FY-16 TI process will add new Knights Landing technology to the HPCMP footprint.

Likewise, the HPCMP network capacity significantly increased its capability. More than doubling its aggregate bandwidth from 2015, the Defense Research and Engineering Network (DREN) and its secret overlay serviced 163 active unclassified and 71 secret sites across the nation. The DREN network also provided connectivity for 26 major Joint Mission Environment Test Capability test and evaluation events.

The Computational Research and Engineering Acquisition Tools and Environments (CREATE) software development program provides critical modeling and digital engineering support to the acquisition community. By late 2016 more than 160 defense-related organizations from the Services, industry, academia, and other government agencies were using CREATE tools. These tools, with the Program's computers and networks, are now being used to support decisions in acquisition programs such as the Joint-Multirole Technology Demonstration program, future submarine procurements, and next-generation aircraft and ground vehicles.

In 2016, the HPCMP explored several strategic initiatives to better support DoD's missions. One of these efforts is the transition of the Maui High Performance Computing Center from a shared large-scale compute resource center to a Vanguard Center that explores new HPC technologies and algorithms for potential incorporation into the program. This effort will be key to keeping the HPCMP positioned to support the growing and diversifying demands for HPC resources.

As cyber threats continue to grow and evolve, the HPCMP is seeking to determine the best ways to

apply HPC to mitigate these threats. The HPC Architecture for Cyber Situational Awareness (HACSAW) initiative seeks to explore novel and innovative analytical capabilities using the program's very powerful computers to detect the anomalous activity that may indicate nefarious actions.

Service and Agency requirements indicate a growing demand for highly classified HPC resources. While the Program has traditionally supported these requirements with smaller, special purpose procurements, the maturity of program-supported codes that have been under development for many years is driving the need for larger machines. In FY 2016 the program initiated a study of the demand and options for above-secret shared HPC computing. This study informed supercomputer acquisitions in FY17 and beyond. We anticipate the first above-secret shared HPC resources coming on line in FY18, with additional procurements thereafter.


As we look to the future, the computational requirements needed to meet the demands of research and acquisition of future weapons systems such as hypersonic systems, directed energy systems, and autonomous systems, are continuing to grow. The HPCMP will continue to execute its mission to accelerate technology development and transition into superior defense capabilities through the strategic application of high performance computing, networking and computational expertise.

Table of Contents

1.0 HPCMP Overview	6
2.0 Program Synopsis.....	10
2.1 HPC Centers and Supercomputing Resources	10
DoD Supercomputing Resource Centers (DSRCs)	10
Dedicated HPC Project Investments (DHPIs).....	11
2.2 Networking and Cybersecurity.....	12
Defense Research and Engineering Network (DREN).....	12
Cybersecurity.....	14
2.3 Software Applications and User Experience	17
Computational Research and Engineering Acquisition Tools and Environments (CREATE)	17
HPCMP Applications Software Initiative (HASI).....	23
Frontier Projects.....	27
Data Analysis and Assessment Center (DAAC)	32
2.4 Workforce Development	32
3.0 HPCMP Contributions to Mission Success.....	36
4.0 Strategic Initiatives	56
4.1 Above-Secret Computing	56
4.2 HPC Architecture for Cyber Situational Awareness (HACSAW)	57
4.3 Technology Innovation: Maui Vanguard Center for HPC.....	58
4.4 Looking to the Future	59

1.0 HPCMP Overview


The Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) provides science and technology, test and evaluation (S&T, T&E), and acquisition engineering community customers with a full ecosystem to address their most challenging problems. The HPCMP ecosystem consists of world-class high performance computational resources, scientific software, and reliable networking supported by sound cybersecurity practices. HPCMP personnel provide supercomputing technology delivered with a portfolio of services and capabilities that the Department's scientists and engineers use every day to develop, test, and field the new technologies that keep America strong. The HPCMP is a technology-led, innovation-focused program committed to



extending high performance computing (HPC) to users across the Department to address the DoD's most significant challenges. The HPC ecosystem enables the development and use of physics-based modeling and simulation, and high performance data analytics to optimize the development,

acquisition, and sustainment of innovative weapons systems. We enable the Department of Defense to maintain the US strategic warfighting advantage over current and future adversaries.

To achieve our goals, the HPCMP is organized into three functional components: HPC Centers and Supercomputing Resources, Networking and Cybersecurity, and Software Applications—each discussed in detail in this report. The effort requires more than just the hardware and software of supercomputers and networks, the HPCMP also brings deep domain-specific knowledge to bear to enable our users to solve the challenging problems facing the DoD. We provide solutions that are cost-effective, flexible, and responsive. Our success is measured by our ability to positively impact our customers' achievement of their critical mission objectives in service to the nation.



The HPCMP responds to a diverse set of stakeholders which includes our governance and advisory boards, RDT&E and acquisition engineering users and customers, collaborators, and other interested parties. The Program is managed by government personnel from the Services and Agencies, and is supported by a large number of contractor staff essential for the day-to-day operations of the HPCMP supercomputing ecosystem.

The Program is a tri-Service effort managed by the Assistant Secretary of the Army, Acquisition, Logistics and Technology, ASA(ALT), and is executed by the US Army Engineer Research and Development Center (ERDC) in Vicksburg, Mississippi.

The Executive Steering Group (ESG) provides Senior Executive Service guidance and oversight for the HPCMP, and serves as the forum for coordinating and adjudicating enterprise policies, practices, responsibilities, and priorities. The ESG membership includes three chairs: Deputy Assistant Secretary of the Army (Research and Technology), Deputy Assistant Secretary of Defense (Research and Engineering), and Deputy Assistant Secretary of Defense (Command, Control, Communications, Cyber, and Business Systems); and three members: Deputy Assistant Secretary of the Navy (Research, Development, Test, and Evaluation), Deputy Assistant Secretary of the Air Force (Science, Technology and Engineering), and the Director of the US Army ERDC, *ex officio*.

The HPC Advisory Panel (HPCAP) advises the HPCMP on policy, procedures, and strategic focus. Members serve as primary Service/Agency representatives to the HPCMP, providing Service and Agency priorities, determining resource allocations, and vetting technical proposals. The HPCAP members frequently interact with HPCMP leadership providing a Service-/Agency-level perspective.

At a more technical level, the Computational Technology Area (CTA) Advisory Panel assesses the technical impacts of potential HPCMP actions on DoD science and engineering challenges. Working

from a subject-matter expert perspective, the CTA Leads offer insight into how algorithm development and changing computer architectures will affect HPCMP users. The User Advocacy Group (UAG) facilitates the exchange of information between the user community and the HPCMP at a tactical level. The UAG represents the different user communities within each of the Services and Agencies. They serve as points-of-contact for bringing the needs of a diverse community of users to the operational leadership of the HPCMP.




Figure 1. HPCMP Stakeholders

The HPCMP serves a large, diverse group of DoD RDT&E and acquisition engineering users, each supporting numerous customers, as shown in Figures 1 and 2. Additionally, the HPCMP enables and supports collaborations across academia, industry, federal agencies, and internationally with key partner nations. Lastly, interested parties include Congress, the White House Office of Science and Technology Policy (OSTP), the DoD industrial base, and supercomputer, networking, security, and software vendors.

	<u>Army HPCMP Participation</u> Presence in 25 Army labs Connected to 9 major Range and Test facilities ARL & ERDC DSRCs SMDC ARC 3 DHPis (FY13-FY16) 40 DREN Sites 13 PETTT Staff 741 Users 3 HASI Projects 3 Frontier Projects
	<u>Navy HPCMP Participation</u> Presence in 17 Navy labs Connected to 6 major Range and Test facilities Stennis DSRC NRL-DC/ARC Navy SSC-SD ARC 4 DHPis (FY13-FY16) 57 DREN Sites 10 PETTT Staff 671 Users 4 HASI Projects 3 Frontier Projects
	<u>Air Force HPCMP Participation</u> Presence in 11 Air Force labs Connected to 7 major Range and Test facilities AFRL & MHPCC DSRCs AFRL-RI ARC 5 DHPis (FY13-FY16) 29 DREN Sites 12 PETTT Staff 809 Users 5 HASI Projects 2 Frontier Projects
	<u>Defense Agencies Participation</u> DTRA, MDA, DARPA 3 DHPis (FY13-FY16) 24 DREN Sites 128 Users 2 HASI Projects

Figure 2. Summary of Service and Agency participation

As the demands and availability for computational resources have grown, the user base and level of utilization (size and complexity of jobs) have grown to take advantage of the opportunities afforded by increased resources. In 2016, there were approximately 2,400 active HPCMP users, with nearly 1,200 users utilizing at least 100,000 core-hours, and of those, 600 used at least 1 million (M) core-hours. Often it takes time for new users to mature into sophisticated HPC users and for their HPC requirements and allocations to grow. An important trend is that the number of users utilizing at least 1M core-hours has grown from 25 users in 2006 to almost 600 users in 2016. Large-scale HPC simulations have transformed from a niche activity to a mainstream activity.

The HPCMP made available to its user base 7.38 billion (B) core-hours on HPCMP systems during FY16 representing a substantial computational resource that would require individual research programs significant buy-in to procure equivalent resources via cloud services or individual HPC platforms, networks, and data storage. HPCMP users represent a wide spectrum across well-established S&T projects, some routinely use hundreds of thousands to millions of core-hours per year, while emerging communities of users from T&E and acquisition engineering generally have more modest use of these powerful resources. The HPCMP user demographic is shown in Figure 3.

2.0 Program Synopsis

The HPCMP delivers to its users an integrated ecosystem to enable the users to conduct research, to test systems, and to identify solutions for problems facing the Department of Defense. The Program provides high performance computing platforms, connected by secure high-speed networks that enable a distributed user base to develop and use advanced software to address the challenges of maintaining military superiority in an increasingly challenging world.

2.1 HPC Centers and Supercomputing Resources

The HPCMP delivers world-class high performance computational capabilities to DoD's S&T, T&E, and acquisition engineering communities, and supports five DoD Supercomputing Resource Centers (DSRCs). The DSRCs provide DoD scientists and engineers with the resources necessary to solve the most demanding computational problems. The DSRCs provide a robust complement of HPC capabilities which include: large-scale HPC systems, high-speed networking, multi-petabyte archival mass storage systems, and customer support services. Complementing the DSRCs are four Affiliated Resource Centers (ARCs), Dedicated HPC Project Investments (DHPIs), and an HPC Help Desk supporting the DoD HPC community. The figure below highlights the diversity of HPCMP supercomputing resources.




Figure 4. Distribution of HPCMP Supercomputing Resources

DoD Supercomputing Resource Centers (DSRCs)

The DSRCs strive to meet the needs of customers using a combination of synchronized resource, workload, and enterprise system management services. Collectively the DSRCs support some of the world's largest computational projects, address expanding DoD requirements for HPC capability, and delivered 7.38 billion core-hours of computing power in FY 2016. They assist their customers through collaboration, technology, and service support. A wide array of advanced HPC technologies, resources, and service support enables scientific productivity by HPCMP customers.

The five DSRCs are:

- Air Force Research Laboratory (AFRL) DSRC, Wright-Patterson AFB, Ohio
- Army Research Laboratory (ARL) DSRC, Aberdeen Proving Ground, Maryland
- Army Corps of Engineers, Engineer Research and Development Center (ERDC) DSRC, Vicksburg, Mississippi
- Maui High Performance Computing Center (MHPCC) DSRC, Kihei, Maui, Hawaii
- Navy DSRC, Stennis Space Center, Mississippi

Through an annual Technology Insertion (TI) process, the HPCMP delivers and updates supercomputing capability at its DSRCs. The diversity of supercomputing platforms mirrors the diverse needs of the DoD HPC user community. The figure below summarizes the FY16 shared systems within the HPCMP.




Figure 5. HPCMP supercomputers at the five DoD Supercomputing Resource Centers (DSRCs)

To access HPCMP systems, users are required to have, at a minimum, a National Agency Check with Inquiries (NACI). The HPCMP recognizes the value of contributions of individuals or organizations who may be engaged in basic research, but are lacking their NACI in order to conduct work on DoD supercomputing assets in the Program. To address this gap, the Program provides an Open Research System (ORS) consisting of a Cray supercomputer and a mass storage data archive to support computational work where the results are specifically intended for public release, and the software used to obtain those results is not export-controlled.

The ORS is co-located with the ERDC DSRC, and currently operates a 460-node Cray XE6m system, providing 14,270 cores.

Complementing the DSRCs, the Affiliated Resource Centers (ARCs) are comprised of DoD Laboratories and Centers that acquire and manage HPC resources as part of their local infrastructure, and share these resources with the broader HPCMP user community.

Dedicated HPC Project Investments (DHPis)

DHPis serve a unique and essential purpose in the HPCMP by providing moderate-size HPC systems that are used to solve problems with computational requirements that cannot be satisfied by resources provided by the DSRCs. Examples of such requirements include above-secret HPC capabilities, systems requiring near-real-time data analysis, hardware-in-the-loop projects, embedded HPC systems, and emerging HPC technologies. DHPis typically fall within two primary categories: addressing specialized functional requirements such as above-secret computing, and exploratory systems, such as emerging HPC systems that may mature into solutions that will become part of the DSRCs.

Proposals are endorsed and prioritized by the Service/Agency HPCAP members before further consideration is given. Once endorsed, proposals are evaluated based on the following criteria: 1) merit of the proposed project, 2) computational approach used to address the project's requirements, 3) potential for significant progress; 4) resource appropriateness and O&M/facilities support quality, and 5) ability of the host organization to meet financial execution requirements. Organizations selected to receive a DHPI are required to produce an annual report summarizing the impact of the work being performed with the DHPI. After the DHPI has been in successful production mode for two years, the HPCMP releases oversight to the home organization. At this point, the home organization takes full custodianship of the DHPI. The table below provides a summary of all active DHPIs.

Year Awarded	Service/Org	Project Title	Classification
FY16	Air Force, A9	Integrated Military Analysis Information Technology Suite (IMAITs)	TS/SCI
FY16	Defense Threat Reduction Agency	WMD Defeat Studies	TS/SCI
FY15	Army Research Laboratory	Program-Restricted Computational Fluid Dynamics/Structural Dynamics Simulations of Rotorcraft	Secret/SAR
FY15	Naval Postgraduate School	Data Analysis of Real-Time Streams (DARTS)	Unclassified
FY15	Defense Threat Reduction Agency	CBRN Threat Assessment Graph Analytics	TS/SCI
FY14	Army, ATEC (WSMR, Aberdeen Test Center)	Real-time Hadoop/Hbase HPC system	Unclassified
FY14	Air Force Lifecycle Management Center	Applications of the Digital Thread to Capability Planning and Analysis	TS
FY14	Naval Research Laboratory	On-Demand GPU Applications	Unclassified
FY13	Naval Surface Warfare Center, Carderock Division	Modeling and Simulation of Weapon Effects on Surrogate Test Items and Ship Structure	TS
FY13	Air Force Research Laboratory, Sensors Directorate	Countermeasures for Infrared Seeking Missiles	Secret/SAP
FY13	Air Force Arnold Engineering Development Center	Cluster for Enhanced Spacecraft Signatures	TS/SCI
FY13	Air Force Research Laboratory, Sensors Directorate	HPC System for Real-Time Gotcha Radar Surveillance and Autonomous Sensing	Unclassified
FY13	Army, ATEC	Real-Time Data Intensive High Performance Cluster	Unclassified

Table 1. Summary of active DHPIs

2.2 Networking and Cybersecurity

Defense Research and Engineering Network (DREN) is the HPCMP's information-assured wide-area network that minimizes the impact of distance on time-to solution to support the RDT&E and acquisition engineering communities anytime, anywhere with a versatile, low-latency, high-throughput communications network. In conjunction with DREN, the HPCMP Cybersecurity team executes its mission to apply security intelligently to ensure proactive protection and promote a productive computational environment.

Defense Research and Engineering Network (DREN)

DREN III is the third-generation instantiation of networking capability for HPCMP and includes Internet Protocol (IP) services, as well as next-generation transport technologies such as Ethernet and Optical Services. DREN III includes bandwidth capacities starting at 50 megabits per second (Mbps), and

increasing by four orders-of-magnitude up to 100 gigabits per second (Gbps). For comparison, the maximum bandwidth at sites on the DREN II contract was 2.4 Gbps.

The following highlights DREN use cases (Figure 6).

- **HPC Access:** The primary purpose for DREN, and for the secret-level counterpart, SDREN, is to connect the HPCMP's supercomputing resources to the users of those systems. Unclassified supercomputing resources are connected to DREN, while classified supercomputing resources are connected to SDREN. With DREN's connectivity to the Internet and peering with other research and education networks, the supercomputing users need not be located at military installations, but may remain at DoD contractor or academic institutions.
- **RDT&E Access:** In addition to HPC access, DREN provides connectivity to the DoD RDT&E and acquisition engineering communities, so they can communicate among their own geographically dispersed units, and others in these communities at their appropriate levels of security.
- **Test Events:** The DoD Test and Evaluation community operates numerous test events over the DREN backbone. Many of these test events are classified and conducted over SDREN. These events can include live test articles (such as aircraft or other sensors), human-in-the-loop participants, and virtual assets simulated via computers.
- **Closed Communities:** DREN supports several DoD Services and Agencies in creating a closed community of interest over the DREN backbone. The Service/Agency operates their own overlay connecting any number of DREN sites together, using NSA Type 1 encryptors and their own key material. These closed communities are accredited by applicable Service/Agency Authorizing Officials.




Figure 6. DREN Use Cases

At the end of FY16 there were:

- 163 active DREN sites and two sites ordered, but not yet active. DREN sites are located across the continental United States (CONUS), plus Alaska and Hawaii
- 71 active Secret DREN (SDREN) sites. SDREN is implemented as a network overlay on the DREN backbone for classified communications at the Secret-collateral level
- 6 DREN peering points, two Internet transit points and three DoD private peering points (for NIPRNet peering). DREN currently peers with 38 other national and international networks including research and education networks such as Internet2, DOE's ESnet, NASA's NREN, California's CENIC, and Ohio's OARnet
- 40 Gbps connectivity at each CONUS-based DSRC, plus SPAWAR-Pacific in San Diego, California
- 334.2 Gbps total aggregate bandwidth
- 2x increase in aggregate DREN bandwidth from January 2015 (163.6 Gbps) to October 2016 (334.2 Gbps), which continues to increase. (Most of this growth is attributable to the 40 Gbps bandwidth upgrades at the CONUS-based DSRCs.)

Additional contributions during FY16 were:

- The Joint Mission Environment Test Capability (JMETC) conducted 26 significant test and evaluation (T&E) events over DREN. In addition, 20 of 22 major Range and Test Facilities are connected to DREN.
- 1 Gbps optical service for the Army Research Laboratory's (ARL) Open Campus initiative at Aberdeen Proving Ground (APG) was delivered via the DREN III contract in early April 2016. ARL will use the connectivity to connect an open HPC system to the Internet via the University of Maryland's MAX-GigaPOP.

The following highlights some ongoing DREN activities:

Cybersecurity Enhancement Project (CSEP)

A Cybersecurity Laboratory has been established at SPAWAR-Pacific in San Diego to evaluate vendor cybersecurity products. Recommendations on products to acquire and deploy will be provided in early FY17. The laboratory is part of the Cybersecurity Enhancement Project that will upgrade the perimeter defenses of DREN at the Internet Access Points (IAPs) for Internet peering and transit. Similar capabilities will be implemented at the NIPRNet peering locations, following the roll-out at the IAPs. The major CSEP upgrade/improvement efforts include: Domain Name Services (DNS), Distributed Denial-of-Service (DDoS) attack prevention, full packet capture, malware detection and prevention, web content filtering, and security analytics.



DREN Joint Sensors (DJS)

The DJSs are undergoing a re-design to support the inclusion of next-generation Intrusion Detection Software (IDS) and 40 Gbps data flows. In FY16, a suitable hardware platform was identified, and 30 servers were procured to support this research initiative. A full technology refresh for all existing sensors at the other sites is scheduled to occur in FY17.

Software-Defined Networking (SDN) Pilot Project

The HPCMP has partnered with the DREN III service provider, CenturyLink, to explore the use of Software-Defined Networking (SDN) and Network Function Virtualization (NFV) technologies across the DREN backbone. DREN will deploy four SDN Micropods at two CenturyLink facilities in Virginia and Colorado, and at two government facilities in California and Ohio, for the evaluation of SDN capabilities. Three major project phases have been identified, and will explore aspects such as securing and operating the devices in a production-like environment. In addition, the pilot will investigate the deployment and use of virtual routers, firewalls, encryption, and performance tools at the four pilot sites.

SDREN Technology Refresh

In FY16, Networking began a project to perform a technology refresh on the equipment in SDREN. SDREN is a classified overlay on the DREN backbone using NSA Type 1 encryptors and a common key. The new architecture will take advantage of DREN III capabilities and build upon the lessons-learned from previous generations of SDREN. The architecture consists of a layer-2, full-mesh design between all SDREN sites. A next-generation layer-2 encryptor, was selected as the basis for the new architecture. New routers/switches are being deployed to complete the technology refresh. Deployment is scheduled for FY17.

Cybersecurity

As cyber threats continue to evolve and advance, so have the HPCMP's cybersecurity activities and associated research initiatives. In continuing its mission to apply security intelligently to ensure proactive protection and promote a productive environment for the RDT&E and acquisition engineering communities, the HPCMP Cybersecurity Team led two Program-wide efforts to increase protections against our cyber adversaries.

DoD Cybersecurity Campaign Plan

In response to the DoD Chief Information Officer (CIO), Undersecretary of Defense for Acquisition, Technology, and Logistics (USD AT&L), and US Cyber Command (USCYBERCOM)-issued plan, the HPCMP Cybersecurity Team managed a multifaceted response to enforce full cybersecurity compliance and accountability across the Program. From internal coordination amongst the HPCMP's Information Systems Security Manager (ISSM) panel, to external briefings delivered to the DoD Security/Cybersecurity Authorization Working Group (DSAWG) and Privileged User Working Group (PUWG), the HPCMP took critical actions to harden and protect HPCMP systems, and the DREN wide-area network (WAN). Given the nature of HPCMP's users throughout the S&T, T&E, and acquisition engineering communities, the HPCMP Cybersecurity Team implemented tailored solutions, such as YubiKey, to provide strong authentication mechanisms between HPCMP assets and its users. The HPCMP's implementation of YubiKey is currently under review and consideration by a Service-level CIO for a transferred technology to be used within their area of responsibility.

DoD Risk Management Framework (RMF) Transition

Under the direction and leadership of the HPCMP Authorizing Official (AO), the HPCMP continued its transition from the DoD Information Assurance Certification and Accreditation Process (DIACAP) to the Risk Management Framework (RMF) for DoD Information Technology (IT). In establishing an integrated enterprise-wide decision structure for risk management, the HPCMP has categorized HPCMP systems based on potential impacts resulting from loss of confidentiality, integrity, and availability. With the transition to RMF, HPCMP system authorizations are now being actively managed, tracked, and updated within the Enterprise Mission Assurance Support Service (eMASS) system. The Program-wide transition has made significant strides due to coordination and cooperation between Cybersecurity and IT personnel throughout the HPCMP, including ISSMs and Information System Security Officers (ISSOs).



Cybersecurity Metrics of Interest

As the HPCMP continued its commitment to provide secure, trustworthy and protected HPC resources, and an enterprise network to the RDT&E and acquisition engineering communities, several metrics reveal active engagement in defending the HPCMP ecosystem.

Cyber intrusions, either root-level (Category 1) or user-level (Category 2), did not occur on any HPC systems over the course of the year.

Strong authentication, designed to degrade the adversaries' ability to maneuver on a DoD network, increased by **35%** over the course of the year.

Device hardening, intended to prevent common exploitation techniques through proper configuration and vulnerability patching, increased by **57%** over the year.

Reducing the attack surface, which focused on minimizing external attack vectors into the HPCMP and DREN, increased by **17%** over the course of the year.

Research Activities

To effectively counter the cyber threats facing the HPCMP, the Cybersecurity Team initiated several research and development initiatives, investigating cybersecurity capabilities and techniques that add increased protection and defense of the HPCMP, DREN, and its customers.

Cybersecurity Environment for Detection, Analysis, and Reporting (CEDAR)

To assess and extend network security monitoring capabilities employed on the DREN, the HPCMP Cybersecurity Team commenced the CEDAR project by leveraging modernized technologies capable of detecting anomalous or malicious network traffic. With a focus on tailoring the significant advancements from industry and academia in networking and computational technologies, CEDAR enabled the HPCMP's collection, analysis, evaluation, and performance of detection capabilities on





Figure 7. CEDAR Dashboard

DREN. The unique characteristics of the DREN's robust, high-capacity, low-latency connectivity amongst the RDT&E and acquisition engineering communities requires a measure of effectiveness (MOE) framework to quantify the performance and utility of security monitoring capabilities. The HPCMP Cybersecurity Team successfully designed, built, and implemented twenty-one (21) CEDAR sensors, and an advanced processing pipeline for rapid collection, analysis, and visualization of high-value network security monitoring data.

Information Security Continuous Monitoring (ISCM) – Jigsaw

In continuing the HPCMP's progress to implement a risk-based approach to protecting networks and information systems, the HPCMP Cybersecurity Team worked closely with stakeholders within the HPCMP Cybersecurity Service Provider (CSSP), Comprehensive Security Assessment (CSA) Team, and DSRC Cybersecurity personnel to advance Jigsaw capabilities. Jigsaw was successfully transitioned to the HPCMP-wide Security Services Enclave (HWSSE) to provide its users throughout the HPCMP and DREN with streamlined access to the application. Additional capabilities were researched, developed, and implemented to provide features for data feed reporting, and plan of action and milestone (POA&M) tracking. The HPCMP Cybersecurity Team continues to employ, exercise, and advance Jigsaw as the HPCMP's information security continuous monitoring (ISCM) platform on the DREN.



Rapid Audit of Unix

System administrators and cybersecurity personnel throughout the HPCMP have used a host-based scanning tool, called Llama, for years as a lightweight application to perform security assessments of Unix-based systems. This year the HPCMP Cybersecurity Team finalized a complete refactoring of the tool, now called Rapid Audit of Unix (RADIX), providing a comprehensive update and increased functionality to the HPCMP and its community. In succeeding the Llama tool, RADIX has been

implemented as a customized scanning capability for HPC resources and Unix-based systems throughout the HPCMP.

As the HPCMP Cybersecurity Team looks to the future, a continued focus is on delivering a defense-in-depth approach to protect and defend the HPCMP, DREN, HPC resources, and the S&T, T&E, and acquisition engineering communities we serve. The HPCMP Cybersecurity Team will intelligently apply security from DREN Internet Access Points (IAPs), to Service Delivery Points (SDPs), to system-level and end-users to provide secure, productive access to HPCMP resources for the RDT&E and acquisition engineering communities.

2.3 Software Applications and User Experience

The HPCMP's Software Applications and user experience efforts provide a suite of software development and support services aimed at optimizing software capabilities to design, develop, test, and deploy superior DoD capabilities. These efforts and services include:

- Computational Research and Engineering Acquisition Tools and Environments (CREATE) – Develops key physics-based engineering software applications to enable DoD engineers to develop and test virtual prototypes of DoD weapon systems.
- HPCMP Applications Software Initiative (HASI) – Ensures continued effective use of future HPC capabilities through software performance improvements.
- Frontier Projects – High mission impact DoD RDT&E and acquisition engineering computational efforts that require taking full advantage of the extensive resources offered by the HPCMP to achieve success.
- User Productivity Enhancement, Technology Transfer and Training (PETTT) – Provides key intellectual services through university and industry partners.
- Data Analysis and Assessment Center (DAAC) – Transforms enormous amounts of computational and experimental data into scientific visualizations that foster insight and meaningful data interpretations.

Computational Research and Engineering Acquisition Tools and Environments (CREATE)

The CREATE program provides critical modeling support in the world of digital engineering prototyping, and is maintaining a steady pace of adoption and acceptance of its physics-based engineering software tools across an expanding customer base. CREATE is the HPCMP's premier vehicle for addressing the DoD's current and future design and analysis efforts for its major acquisition programs. CREATE provides innovative applications of its software tools for developing and optimizing aircraft, ship, ground vehicle, and radar antenna designs, and allows the acceleration of the acquisition timeline. As of November 2016, the CREATE program has expanded the acceptance, use, and adoption of its various physics-based software tools to over 160 defense-related organizations from the Services, industry, academia, and other government agencies with a base of more than 1,200 users. The success of the CREATE program increases with each new annual software release. Its products are now becoming an integral part of major defense acquisition programs for design space exploration, design analysis and performance prediction and testing across the weapon system lifecycle.




Figure 8. CREATE Digital Engineering Acquisition Paradigm Shift

A key element of the success of the CREATE software development process is its hybrid organizational structure that embeds its workforce with DoD Service clients. These government-owned and government-developed complex engineering software applications enable high-fidelity, physics-based virtual modeling and simulation using computational fluid dynamics (CFD), finite element analysis (FEA), and other physics-based mathematical algorithms for a wide variety of DoD Science and Technology (S&T), test and evaluation (T&E) and acquisition programs of record.

Present acquisition programs largely follow an empirical "design—build—test" iterative methodology. This results in late discovery of design flaws, immature technology issues, and system integration problems. Rework and redesign efforts contribute substantially to cost overruns and schedule delays. By employing a "model—test—build" paradigm, optimized engineering designs can be developed early in the acquisition process using CREATE tools. Costs can be substantially reduced, schedules shortened, and design and program flexibility, and agility increased. Above all, acquisition program performance will be improved through the reduction of design flaws, the quick and flexible development of sound engineering concepts and designs, and beginning the systems integration engineering process much earlier in the acquisition process (Figure 8).

Within five CREATE project areas (Figure 9), the CREATE program has developed a suite of 11 robust physics-based engineering software tools that perform critical design space exploration and performance assessments using the HPCMP DSRC high performance computing systems, and local high performance computing systems at industry sites. It is this marriage of engineering software tools and high performance supercomputers that provides significant modeling and simulation capabilities to the DoD acquisition engineering community that have not been previously available. This enables the use of an unprecedented capacity to accomplish virtual prototyping of major weapons systems programs in a fraction of the traditional time.




Figure 9. CREATE Project Areas

CREATE-AV (Aviation Vehicles): Kestrel – Fixed-Wing Aircraft Design and Analysis

Kestrel provides accurate predictions of DoD air vehicle performance with a specific focus on fixed-wing military aircraft. It integrates computational fluid dynamics, structural dynamics, propulsion and control for subsonic through supersonic aircraft operation. The capabilities available in Kestrel include aerodynamics (Navier-Stokes solvers and a full suite of boundary conditions and turbulence models), structural dynamics (modal models or finite-element analysis for aero-structure interactions), flight control systems (control surface movement), and propulsion (engine “cycle decks” for propulsion effects or direct engine simulation, including inlet and rotating machinery, nozzle and moving walls) (Figure 10). This set of capabilities is unique in the U.S. defense community, allowing users to develop major innovations in the design of next-generation aeronautical weapon systems.




Figure 10. Modeling UFY-219 Aircraft Engine Surface Currents

With Kestrel, engineers can verify designs prior to key decision points and fabrication of test articles or full-scale prototypes; plan and rehearse wind-tunnel and full-scale flight tests; evaluate planned or potential operational use scenarios; perform flight certifications (airworthiness, flight envelope expansion, mishap investigation, stores carriage and release, etc.); and generate response surfaces usable in DaVinci, flight simulators and other environments that require real-time access to performance data. Kestrel has been applied to the analysis of over 30 fixed-wing DoD aviation systems, including at least 10 unique unmanned aerial vehicles.

Recent impacts with Kestrel include assessment of the static forces and moments on the F/A-18E Super Hornet that enabled analysis of the effectiveness of proposed modifications to the flight control system.

Additional impacts include lifecycle 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.

CREATE-AV (Aviation Vehicles): Helios – Rotary-Wing Aircraft Design and High-Fidelity Analysis



Helios is a high-fidelity, full-vehicle, multi-physics analysis tool for rotary-wing aircraft. Helios can calculate the performance of a full-sized rotorcraft, including the fuselage and rotors. It can handle arbitrary rotor configurations, and analyze and predict prescribed maneuvers with tight coupling of rotor aero-structural dynamics. A highly accurate treatment of vortex shedding from rotor blade tips using adaptive mesh refinement gives Helios the unique capability to assess the interaction of these vortices with the fuselage and nearby rotor blades. Helios provides all the benefits for rotary-winged aircraft that Kestrel does for fixed-wing aircraft.




Figure 11. CH-47F Rotor Vortices in Helios

The Army Rotorcraft program, together with Boeing, used Helios to generate early-design stage predictions of helicopter performance for a proposed rotor blade upgrade of the CH-47F Chinook helicopter to achieve up to an estimated 2,000 pounds increase in hover thrust (~10%) with limited degradation of forward-flight performance. This added performance significantly enhances the combat capability of the 400+ Chinooks in the Army inventory (Figure 11). The Army Joint-Multi-Role Technology Demonstrator (JMR-TD) program also recently used Helios to provide decision data for proposals from four vendors. Helios enabled Army aviation engineers to conduct an independent analysis of contractor proposals. The Army Rotorcraft program is currently using Helios to assess H-60 tail rotor effectiveness for providing directional control of helicopters, in combination with increased engine and main rotor performance.

CREATE-AV (Aviation Vehicles): DaVinci – Conceptual Aircraft Design

DaVinci enables integrated, multi-fidelity, physics-based digital engineering and analysis of aerospace (fixed-wing, rotary-wing, and space access) vehicles for their entire lifecycle, including vehicle performance, effectiveness, cost, schedule, thorough design space exploration, uncertainty quantification, risk assessments, and sensitivity studies. DaVinci includes a capability to allow conceptual designers to quickly and intuitively build parameterized, associative, attributed models (Outline Markup Language and corresponding layout of internal structure and sub-systems) necessary to feed higher-fidelity analysis tools (e.g., Kestrel and Helios). It also allows user contributions of additional capability not delivered with the product through knowledge management, model persistence, and collaboration across organization and Service boundaries.



CREATE-Ships (Maritime Vessels)

The overarching goal of the CREATE-Ships project is to develop software that enables comprehensive exploration of trade space design options for complex maritime systems; provides confidence in the results of computational predictions in required disciplines across all phases of the acquisition process; and meets required acquisition timelines for DoD clients, primarily the US Navy.

CREATE-Ships (Maritime Vessels) RSDE-Rapid Ship Design Environment

RSDE is a concept design tool that allows engineers and naval architects to assess the tradeoffs inherent in designing ships to meet a spectrum of competing performance parameters (Figure 12). RSDE provides the following critical capabilities: 1) Generates a wide range of ship concept options, and allows selection of optimum solutions; 2) Develops physics-based response surfaces to act as surrogates for full-physics analyses in the concept design process; and 3) Performs hierarchical, multidisciplinary design optimization, in which a top-level ship-based function, e.g., cost, is optimized subject to constraints at the individual discipline level.



Figure 12. RSDE Interface

Employing the concept of design space exploration, engineers and naval architects can provide data for decision makers on the impact of tradeoffs in range, speed, armament, and aviation support on the size—and, in large measure, the cost—of a proposed ship concept. RSDE can generate tens of thousands of candidate ship designs with varying hull forms, including compartment and machinery arrangements. To date, RSDE has been used to enable set-based design on 10 different Navy acquisition programs, allowing down-selection of a ship design to occur earlier in the requirements decision process.

CREATE-Ships (Maritime Vessels) IHDE-Hydrodynamics

IHDE provides a highly-automated environment that allows naval architects and hydrodynamicists to build computational models, submit “jobs” to HPC centers, and retrieve and validate results for all required areas of hydrodynamics, and at all required levels of fidelity. The Leading-Edge Architecture for Prototyping Systems (LEAPS) is used to manage geometry and analysis information, and allows data sharing between several knowledge domains. LEAPS provides a warehouse for ship model information, which provides several advantages over existing prototyping systems: 1) Interoperability for different activities which already use LEAPS; 2) Creation of a foundation of common terms across different disciplines for ship geometries and characteristics, thereby preventing translation errors; and 3) Synergy in software development and integration using a common product model that can be shared with other activities, eliminating inefficient and time-consuming resetting of the same ship modeling inputs when using different analysis methods.

CREATE-Ships (Maritime Vessels) NavyFOAM—High-Fidelity Hydrodynamics

NavyFOAM is a high-end, full-physics code that is Navy-developed and Navy-maintained. It is based on the OpenFOAM libraries and code architecture. NavyFOAM is a fully parallelized, multi-physics, computational fluid dynamics framework developed using modern, object-oriented programming. The code enables high-fidelity hydrodynamic analysis and prediction of ship performance, including resistance, propulsion, maneuvering, seakeeping, and seaway loads. It has demonstrated accuracy against experimental data for several target applications, such as resistance, propeller characteristics, hull/propulsor interaction, and six-degrees-of-freedom ship motion for underwater vehicles and




Figure 13. Submarine Propulsor Flow Modeling

surface ships (Figure 13). Offering a suite of Navier-Stokes-based flow solvers tailored to specific applications, including single- and multi-phase solvers. NavyFOAM allows assessment of alternative hull and propulsor designs. With NavyFOAM, users can evaluate a ship's performance in a wide array of operating conditions, including subsea and surface operations. It has been applied to many naval systems including assessment of the safe operating envelope of the DDG-1000 Zumwalt Class guided-missile destroyer, propeller designs, and the US Marine Corps Amphibious Combat Vehicle. Most recently, the NavyFOAM program has supported the new Columbia Class SSBN program, using a custom physical model for flow predictions, and rotating-arm simulations to better understand the underlying physics used in design decisions.

CREATE-Ships (Maritime Vessels) NESM-Shock/Damage Analysis

NESM builds on the Department of Energy Sandia National Laboratory's shock analysis tool, Sierra Mechanics, to provide a means to assess ship and component response to external shock and blast using accurate HPC tools (Figure 14). NESM can reduce the time and expense required for physical shock testing of ship classes. It also improves the initial ship design process by assessing planned component installations for shock performance prior to final arrangement and installation decisions. Capabilities include structural dynamics (implicit linear-elastic solvers), solid mechanics (explicit plasticity solvers), fluid dynamics (Euler solvers), and fluid-structure interaction. The solution algorithms in NESM can exploit massively parallel computers and scale to thousands of cores, enabling efficient computer use and the ability to address full-sized naval vessels up to, and including, next-generation aircraft carriers and submarines.



NESM is being used to support Live-Fire Test and Evaluation requirements for the new Ford Class nuclear powered aircraft carriers. All ship classes are required to be tested for shock damage resistance. Historically, this has been done exclusively with external explosive charges. Effects were recorded and necessary alterations were designed and implemented for the ship class. This not only requires a ship to be dedicated to the lengthy test process, but also requires any deficiencies in shock resistance to be retrofitted into the ship design. NESM allows for a computational assessment of shock on the ship and equipment while still in the design stage, when fixes can be incorporated prior to construction. NESM has been approved by the Navy to supplement and reduce costly and damaging full-ship shock trials,



Figure 14. Live Fire Test

CREATE SENTRi (RF Antennas and Radars)

As integral components of almost all DoD weapon systems, antennas and radars perform critical communication, identification, and navigation functions necessary for warfighter effectiveness and survivability. Computational Electromagnetics (CEM) tools have matured in recent years and are widely used for antenna concept exploration and design by the DoD and its contractors.



SENTRi is a set of highly accurate, maintainable, extensible, validated, and productive software tools to support rapid antenna system design and integration into various structures of interest including air, sea, land, and space-based platforms. The project's development of state-of-the-art government-owned antenna modeling capabilities is critical to our nation's defense, and represents a national asset.

SENTRI provides customers with the best available schemes of rapid design space exploration and high-fidelity solutions to quickly find both optimal and robust performing designs (Figure 15). SENTRI is a full-wave electromagnetic prediction code for RF modeling of antennas, microwave circuits, and radar cross-section predictions. It is designed to model complex structures, including highly heterogeneous material structures with multi-scaled features. A key goal is the calculation of the simultaneous performance of multiple-antenna systems embedded on a platform. Its electromagnetic features are based on solutions of Maxwell's equations with advanced hybrid finite-element boundary-integral techniques, which provide high accuracy with the ability to solve large, complex problems. SENTRI is continually validated with DoD measurements, and is being used for antenna design, antenna in situ analysis, RF signature prediction, electromagnetic interference, electromagnetic compatibility, material modeling, microwave device analysis, phased array antenna systems, and apertures. SENTRI is being used by over 60 organizations in government and industry.

CREATE-GV (Ground Vehicles) Mercury & MAT

The CREATE-GV (ground vehicles) project was recently established to develop physics-based, high performance computing (HPC) tools to: 1) enhance ground vehicle concept design; 2) inform requirements development; and 3) provide requisite data for trade space analysis to positively impact cost, schedule, and performance with significant reduction in risk for the acquisition community. The CREATE-GV project is developing two software tools: 1) Mercury, a physics-based HPC tool for tactical mobility with co-simulation of terrain mechanics, such as soil modeling, with vehicle systems including suspension, tire and track options, and powertrain simulation (Figure 16); and 2) MAT

(Mobility Analysis Tool), a model-based analysis software tool used for predicting tactical mobility for a wide variety of terrain conditions. MAT incorporates many factors such as soil condition, vehicle performance and configuration, vegetation density, average surface roughness, and average slope.



Figure 17. Conceptual Light Vehicle Model Used in GV Tools for Trade space Analysis

The 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 methods. In collaboration with Engineered Resilient Systems (ERS) program, Mercury software has been used in a trade space analysis pilot project of the Army's Light Reconnaissance Vehicle (LRV). 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 (Figure 17).




Figure 15. Electric field on CV-22 calculated by SENTRI

Figure 16. Engine Thermal Analysis Using GV Tools

CREATE-MG: (Mesh & Geometry) Capstone-Geometry and Mesh Generation

Capstone is a computer-aided design (CAD)-neutral application that provides two distinct capabilities. The first is the ability to develop numerical representations of a




Figure18. CAPSTONE algorithms work seamlessly in SMLib™ database and discrete geometry.

DoD weapon system (i.e. a Non-uniform Rational B-Splines (NURBS)-based digital product model

consisting of the platform geometry with the associated attributes). The second is the capability to generate a mesh from the geometry. Valid and easily produced meshes with the required accuracy are the essential starting point for the other CREATE (solver) tools for detailed analysis. In addition, a number of non-CREATE groups use Capstone for its geometry and mesh generation capability for their applications (Figure 18).



CREATE Going Forward:

The future of the 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 CREATE tools coupled with supercomputing capability are powerful examples of leading technology that is starting to expand their usage and acceptance throughout the defense industry. New technology areas such as hypersonics, space and missile systems, directed-energy, future submarine design, multihull ship design, and unmanned air, land, and sea, and underwater systems, and autonomous systems design and analysis are all prime candidates for future CREATE tool applications. Of special note, Kestrel is being upgraded for design and analysis of hypersonic systems.

HPCMP Applications Software Initiative (HASI)

The HPCMP initiated the HPCMP Applications Software Initiative (HASI) in mid-FY15 to address the RDT&E and acquisition engineering communities' continuing need for modern applications software that incorporates the latest understanding of physical phenomena, and which executes effectively and efficiently on next-generation high performance computers. Investments in HASI projects were made to address needs identified due to the rapidly emerging RDT&E software gaps identified in algorithms, application codes, and frameworks. The emergence of new hardware accelerators and many-core processors, and the need for software that can take full advantage of these also served as a strong HASI driver. HASI projects are a key component of the HPCMP's overall approach to applications software, and will help ensure continued effective use of HPC capabilities to address current and emerging DoD challenges.

A call for proposals was released in FY15 for HASI projects. Full proposals were subjected to extensive peer review by a technical review panel composed of technical experts within and outside of the Department of Defense, and examined closely by the Services/Agencies and the Office of the Secretary of Defense for their mission impact. Fourteen HASI projects selected were ultimately awarded for a two-year period initiating in FY15 and FY16.

The HASI project portfolio is shown in the table below. Several of these are also highlighted as key mission successes.

Service/Org	Project Title	Name
Air Force/AFIT	Efficient Predictions of Structure and Spectra for Nanomaterials Using ACES III and IV	Rod Bartlett
Air Force/AFRL	Enabling Exascale Calculations for Electronic Theory	Mark Gordon
Air Force/AFRL	Directed-Energy High Performance Computing	Nathaniel Lockwood
Air Force/AFRL	SPACE –a Scalable Physics-based Advanced Computational Engineering Platform for Liquid Rocket Combustion Simulations	Venkateswaran Sankaran
Air Force/AFRL	Multiphysics Simulations of Multicomponent, Off-design Aircraft Engine Operation Using Dynamic Hybrid RANS/LES	Parthiv Shah
Army/ERDC	Higher-Order Computational Methods for Penetrating Weapons Effect Modeling	Kent Danielson
Army/ERDC	The Integration of ERDC Geotechnical and Structural Laboratory Geomaterial Models Into a Modern Eulerian-Lagrangian Coupled Solver for Use in Blast-load and Impact Calculations	Ramon Moral
Army/AMRMC	Software Capabilities to Simulate Non-impact, Blast-induced Traumatic Brain Injury	Jacques Reifman
Navy/NRL	Quantifying Uncertainty in the Battlespace Environment	Jim Hansen
Navy/NRL	High Performance Bioinformatics Workflow for Integrative “-omics” Data Analytics	Judson Hervey
Navy/NRL	Achieving Scalability for the MICHELLE Charged-particle Beam Optics Code on Heterogeneous HPC Architectures	John Petillo
Navy/NRL	Optimizing Global and Regional Earth Prediction	Timothy Whitcomb
DTRA	FEMAP : Complex Multiphysics of Future Exascale Systems	Joseph Baum
DTRA	Refactoring Advanced CFD Applications to Accelerate Optimally on Next-Generation DOD HPCMP Hardware	Donald Kenzakowski

Table 3. FY15-FY16 HPCMP Applications Software Initiative (HASI) Projects

Efficient Predictions of Structure and Spectra for Nanomaterials Using ACES III and IV

PI Rod Bartlett (Air Force, Air Force Institute of Technology): This HASI project is a joint effort between the University of Florida, the Air Force Institute of Technology, and the Materials Directorate of the Air Force Research Laboratory to create first-principles computational chemistry methods to predict excited states for nanomaterials. This research is essential to realize practical design of materials to enable future technologies. The impact on HPC "materials by design" efforts in DoD for nano-optics and nano-electronics are wide-ranging. Given AFIT's interest in nano-excitonics, the potential is huge for creating robust, efficient photovoltaics and electro-optics materials to replace fragile organic materials for use in demanding environments of high temperature and radiation fields. Already the project has enabled reliable prediction of properties of inorganic charge-transfer molecules superior to those currently available. For the longer term, the work will facilitate design of materials for excitonics quantum computers that will not only provide advantage for information and communication security technologies, but ultimately will make artificial intelligence a practical reality. The tools created by this HASI project put the DoD on the forefront of materials-intense photonics and excitonics technologies.

Enabling Exascale Calculations for Electronic Structure Theory

PI: Mark Gordon (Air Force, Air Force Research Laboratory): This HASI project is developing a suite of quantum chemical methods (GAMESS, NWChem, and PSI4) to greatly expand the size of chemical systems that can be addressed with highly accurate electronic structure theory methods,

analyzing methods that can potentially reduce power consumption while maintaining or improving computational throughput and fault resilience to pave the way to exascale computing, and increasing the number of potential users and capability of these methods via enhanced interoperability. Quantum chemistry methods have become an essential partner to experiments in such DOD-critical domains as the development of new high energy fuels, the design of new materials with desired properties such as selective heterogeneous catalysts and resistance to extreme environments, the elucidation of atmospheric phenomena, development of polymeric materials with non-linear optical properties, and the study of intense laser fields. In addition to making significant progress on improving the capabilities of all three major application codes, a detailed study of the relative merits of ARM and Intel Haswell computer architectures was completed and published.

Directed-Energy High Performance Computing Applications Software Initiative (DE HASI)

PI Nathaniel Lockwood (Air Force, Air Force Research Laboratory): Software developed as part of the Directed-Energy (DE) HASI project has assured a strategic advantage to the DoD by enabling high-fidelity virtual prototyping and optimization of DE systems and components to maximize their effectiveness while reducing size, weight, and power to maximize their military utility to the warfighter. The DE HASI project has also supported the evaluation of a hybrid DE air defense system's ability to defeat swarms of unmanned aerial systems in all weather conditions, which will provide the DoD a strategic advantage over our adversaries.

SPACE - a Scalable Physics-based Advanced Computational Engineering Platform for Liquid Rocket Combustion Simulations

PI Venkateswaran Sankaran (Air Force, Air Force Research Laboratory): The SPACE HASI project enables next-generation DoD propulsion systems by adding internal combustion capability to the CREATE repertoire of codes. Applications include gas turbines, augmentors, liquid rockets, solid rockets, and hypersonic propulsion systems. Beyond DoD, SPACE impacts US national interests such as environmentally compatible land-based and marine-based power.

Multiphysics Simulations of Multicomponent, Off-design Aircraft Engine Operation Using Dynamic Reynolds-averaged Navier-Stokes/Large-Eddy Simulation (RANS/LES)

PI Parthiv Shah (Air Force, Air Force Research Laboratory): This HASI project provides strategic advantage to the DoD's efforts to deploy future hypersonic aircraft by addressing the physics, computing, and engineering challenges associated with massively parallelized simulation of high-speed, air-breathing propulsion systems such as scramjet engines. The team is developing and testing new multiscale, multiphysics algorithms that more accurately describe unsteady, turbulent scramjet flows that include reacting chemistry in a computationally efficient manner on current and future DoD supercomputers. The HASI work also aims to better explain outcomes from recent flight vehicle demonstrations that are hypothesized to originate from a strong fluid-thermal-structural coupling between the engine and aircraft.

Higher-Order Computational Methods for Penetrating Weapons Effect Modeling

PI Kent Danielson (Army, Engineer Research and Development Center): This HASI project provides advanced analysis capabilities, in the form of commercial-grade DoD production HPC software, for simulating deformable high-speed weapon impact into traditional and new construction material types. Model creation (meshing) of new warhead design concepts and threat projectiles for protective structural and lethality analyses can now be more easily created in hours or minutes, instead of the weeks or days frequently required with previous methods, while significantly reducing model sizes, computing requirements, and analysis turnaround time, as well as extending their applicability to new impact conditions and material types.

The Integration of ERDC Geotechnical and Structural Laboratory Geomaterial Models into Modern Eulerian-Lagrangian Coupled Solver for Use in Blast-load and Impact Calculations

PI Ramon Moral (Army, Engineer Research and Development Center): This HASI project is transferring ERDC's knowledge about modeling geomaterials to a large- scale multiphysics simulation package maintained at Sandia. The impact is that the models will be available to DoD scientists for little to no cost to the HPCMP DSRCs, regardless of the number of processors used to solve the problem. Our goal is to provide high-fidelity constitutive models to the DoD to increase the accuracy of simulations that include weapons effects into geomaterial-based targets.

Software Capabilities to Simulate Non-impact, Blast-induced Traumatic Brain Injury

PI Jacques Reifman (Army, Army Medical Research and Materiel Command): Model development in this HASI project will ultimately allow the DoD to determine the mechanism by which blast waves can interact with the human body and cause mild traumatic brain injury. Without knowing the underlying mechanism for such an injury, the DoD cannot develop protective gear.

Quantifying Uncertainty in the Battlefield Environment (QUBE)

PI James Hansen (Navy, Naval Research Laboratory): Progress to date on quantifying uncertainty in the battlefield environment has made a substantial impact to the Army Program of Record (PoR) Distributed Common Ground System - Army (DCGS-A). This is the Army's primary operational system for intelligence on the battlefield. The weather component of DCGS-A (DCGS-A Wx) received the latest version of Automated Impacts Routing (AIR) technology, which was developed by the Army Research Laboratory. This new version of AIR contains significant documented enhancements that were developed as part of the QUBE-HASI-funded effort through collaboration with NRL Marine Meteorology Division, Monterey (NRL-MRY). Specifically, the QUBE project enabled enhancements to data structures resulting in speed improvements for aircraft path optimization calculations in AIR on the order of 300x. The latest version of AIR (v2.0-b3), is set to be integrated and deployed as part of DCGS-A Wx release 3.2.6 in the FY17-FY18 timeframe.

High Performance Bioinformatics Workflow for Integrative “-omics” Data Analytics

PI Judson Hervey (Navy, Naval Research Laboratory): This HASI project aligns with the Office of the Secretary of Defense's S&T Priority area of Synthetic Biology, a field which generates large amounts of '-omics' -level biomolecular data. Adapting bioinformatics applications for emerging and future HPC architectures bolsters computational capabilities in biological data mining, pathway analysis, modeling and simulation, and genetic circuit design in Synthetic Biology. Further, these software tools enable acceleration of development and predictive power of gene network function in "non-traditional" chassis organisms directly applicable to military environments (relative to model chassis such as E. coli). As such, this effort assures that the DoD maintains strategic advantages in both HPC and Synthetic Biology.

Achieving Scalability for the MICHELLE Charged-Particle Beam Optics Code on Heterogeneous HPC Architectures

PI John Petillo (Navy, Naval Research Laboratory): This Navy HASI project is developing a predictive capability to model electron sources, transport, and collection for current and next-generation high-power RF amplifiers and oscillators critical to DoD, using the MICHELLE application code, developed with DoD funding, which is the leading electrostatic beam optics code in terms of capability and broad scope of applications that it can address. This development takes advantage of new HPC architectures, such as GPUs and Intel Xeon Phi, and HPCMP software infrastructure, such as Galaxy Simulation Builder and the CREATE mesh construction capabilities. Significant re-engineering of the code has now been accomplished that has greatly increased the scalability of the code, allowing much larger and more accurate modeling and simulation of these critical RF devices.

Optimizing Global and Regional Earth Prediction

PI Timothy Whitcomb (Navy, Naval Research Laboratory): This HASI project is dedicated to improving the accuracy and availability of timely environmental information to provide a strategic advantage for DoD. This work will ensure the continued capability to provide earth system characterization and forecasts at ever-increasing spatial and temporal scales at operational tempo on current and future computing platforms, even as timing requirements and parallel architectures continue to evolve.

FEMAP: Complex Multiphysics on Future Exascale System

PI Joseph Baum (DTRA): This HASI project supports multiscale and multiphysics improvements and code enhancements applied to the Massive Ordnance Penetrator (MOP) weapon development effort. FEMAP couples CFD and CFD methodology to evaluate weapon performance against deeply-buried high-value targets. The program has evolved into both weapon improvements and timely targeting of large, complex underground facilities.

Refactoring Advanced CFD Applications to Accelerate Optimally on Next-Generation DoD HPCMP Hardware

PI Donald Kenzakowski (DTRA): High-fidelity CFD provides insight into spatial and temporal scales impacting DTRA support of warfighter decisions for chemical/biological threat mitigation. The interacting multiphase, chemical kinetic, thermodynamics, and geometric modeling aspects to the problem make such simulations computationally expensive to complete with current runtime methodologies and hardware platform usage; these bottlenecks are being addressed through this HASI effort to significantly accelerate solution run-times and permit rapid parametric studies and probabilistic assessments for improving Agent Defeat effectiveness.

Frontier Projects

The HPCMP established DoD Frontier Projects in FY13 to enable the exploration of RDT&E and acquisition engineering outcomes that would not be achievable using typically available HPCMP resources. These projects are DoD high-impact RDT&E computational efforts that are selected through a rigorous evaluation process that includes both OSD and Service/Agency mission relevance and technical excellence. The Frontier Project portfolio represents the Program's most computationally-demanding, resource-intensive set of projects that require sustained and extensive assistance from the entire HPCMP ecosystem (DSRCs, user support, software development, PETTT, and networking) to fully succeed. Since inception in July 2013, Frontier has grown to eleven Frontier Projects, with the selection of 3 new projects in FY16, representing all three Services. Looking forward, a steady-state of 12-15 Frontier Projects is expected to be achieved after selection of the FY18 projects. At steady-state, the Frontier portfolio will utilize approximately 30% of HPCMP total program computational resources, and represent a considerable resource investment on potential breakthrough concepts.

The Frontier Project portfolio spans the entire RDT&E and acquisition engineering spectrum, from basic research through direct support, to acquisition programs at the decision-making stage. Many of these projects are interdisciplinary in nature, focusing on the boundaries among two or more technical disciplines. DoD organizations executing these projects include each Service's primary research laboratories and basic research granting agencies, as well as other Service laboratories. Principal investigators are from government and academia. Over 100 scientists and engineers are active Frontier users on the HPCMP systems. Collectively, these Frontier Projects used 26% of total HPCMP computing resources in FY2016, accounting for approximately 1.6 billion core-hours on HPCMP systems. Specific projects include leading-edge research in hypersonic systems, combustion, chemistry and turbulence, and development of weather and ocean models to support DoD operations. In addition, there are projects that are critical components of the Navy's Electromagnetic Railgun development, Army rotorcraft acquisition programs, and Army armor development. The Frontier Project portfolio is summarized in chronological order in the table below.

Service/Org	Project Title	Name	Duration
-------------	---------------	------	----------


Navy/ONR	Multi-scale Interactions in Stratified Turbulence	Stephen de Bruyn Kops	FY14-18
Air Force/AFRL	Unsteady Pressure & Heating Environment on High-Speed Vehicles	Ryan Gosse	FY14-18
Air Force/AFOSR	Dynamics and Properties of High-Speed Turbulent Reacting Flows	Alexei Poludnenko	FY15-18
Army/ARL	Simulation of Atomization and Spray/Wall Interactions	Luis Bravo	FY15-19
Navy/NSWC	Navy Electromagnetic Railgun	Joel Mejeur	FY15-17
Navy/NPS	Advancing DoD Modeling and Prediction Capabilities in the Arctic	Wieslaw Maslowski	FY16-19
Army/ARL	Terminal Ballistics for Lethality and Protection Sciences	Robert Doney	FY16-20
Air Force/AFRL	Development of Multi-scale Models for Materials Design	Mark Gordon	FY16-20
Army/AMRDEC	Modeling and Simulation to Support DoD Rotorcraft Acquisition Programs	Andrew Wissink	FY17-19
Air Force/AFRL	Validation of Turbulence and Turbulent Combustion Models	Venkateswaran Sankaran	FY17-21
Navy/NRL	Earth System Prediction Capability	Alan Wallcraft	FY17-19

Table 4. Active Frontier Projects

A brief description of each Frontier Project is provided below. Several of these have also been chosen to be highlighted as key mission successes in a later section of this document.


Multi-scale Interactions in Stratified Turbulence

PI: Stephen de Bruyn Kops (Navy, Office of Naval Research): The technical goal of this project is to generate research simulations at Reynolds numbers that are relevant to engineering problems in the ocean or atmosphere. In addition, turbulence subjected to stabilizing buoyancy forces will be modeled so that fast simulations can accurately predict the performance of vehicles, sensors, and weapons in the deep ocean and atmosphere. Very large direct numerical simulations of simple flow configurations will be performed in order to understand the dynamics of stratified turbulence and its potential military applications.



Unsteady Pressure and Heating Environment on High-Speed Vehicles with Responding Structures

PI: Ryan Gosse (Air Force Research Laboratory): This Frontier Project seeks to determine the fully-coupled high-fidelity fluid-thermal-structural response of high-speed vehicles at full-scale that resolves all relevant physics of the fluid and structure domains. Specific coupled computational fluid dynamics (CFD) and structural application codes will be optimized, scaled, and tested for their ability to address important flow features in boundary layers. These techniques will be applied to hypersonic Air Force strike weapons to impact and control costs for test programs of these vehicles.



Dynamics and Properties of High-Speed Turbulent Reacting Flows

PI: Alexei, Poludnenko (sponsored by the Air Force Office of Scientific Research and the Air Force Research Laboratory):

This project will carry out a systematic, first-principles study of turbulent flames in a range of fuels from hydrogen and methane to realistic jet fuels for a broad range of turbulent intensities and system sizes under Air Force-relevant conditions. The results will provide key insights into the fundamental physical processes controlling combustion dynamics in extreme regimes characterized by high pressures, significant flow compressibility, and large turbulent intensities. The overall goal of such a study is to inform the development of the next generation of physics-driven combustion models that will impact the design of a range of novel systems of interest to DoD, including scramjet engines for hypersonic flight, and detonation-based engines for efficient on-board power generation and propulsion.



Petascale High-Fidelity Simulation of Atomization and Spray/Wall Interactions at High Temperature and Pressure Conditions

PI: Luis Bravo, (Army Research Laboratory): This project seeks to perform extensive three-dimensional simulations of the mixture formation process in direct injection engines, with a particular emphasis on the spray atomization process and spray/wall interactions. An understanding of the complex physics involved in these processes from these simulations will be used to enhance current engineering models for predicting these processes. The outcomes of this project will include both new high-fidelity direct Navier-Stokes simulation codes, as well as the detailed understanding of these processes, which can then be used to improve several Army direct-injection engine technologies by increasing power density and engine efficiency, while reducing battlefield signatures and fuel consumption in tactical situations.




Navy Electromagnetic Railgun

PI: Joel Mejeur, (Naval Surface Warfare Center Dahlgren Division): This Frontier Project supports basic science and applied studies on development of the electromagnetic railgun launcher technologies. These studies have improved understanding of the phenomena that occur at the dynamic, molten aluminum interface between the armature and the rail, provided boundary conditions for modeling the interface layer, and developed detailed multi-physics models of the railgun launcher to assist with detailed analysis of the launcher performance. This detailed understanding of processes within the railgun will directly improve the efficiency of designing and deploying such systems.


Advancing DoD Modeling and Prediction Capabilities in the Arctic

PI: Wieslaw Maslowski (Naval Postgraduate School): This Frontier Project has the overall goal of advancing the understanding of the Arctic Ocean and sea ice systems, through improved representation of processes and feedbacks controlling their operation, to reduce uncertainty in modeling those systems, and ultimately to advance prediction of Arctic sea ice and climate using state-of-the-art modeling applications. The results of this work will aid Arctic stewardship and US strategic interests in line with the DoD and national strategies for the Arctic region, and improve the accuracy of projections of regional climate change.



Development of Multi-scale Models for Material Design


PI: Mark Gordon (sponsored by the Air Force Office of Scientific Research and the Air Force Research Laboratory): This goal of this Frontier Project is the development and application of a multiscale method that seamlessly integrates electronic structure theory methods, parameter-free coarse-graining methods, and molecular dynamics/Monte Carlo simulation methods to provide accurate and efficient predictions of bulk properties of advanced materials, without the need for empirically fitted parameters. The



new multiscale simulation methods will be implemented in the widely-used GAMESS code, thus making these new capabilities available to a broad user base in government, industry, and academia. The methodologies will be applied to accurate computation of the properties of ionic liquids, which, as potential new propellants, may have significant impacts to DoD, including (a) reliable and cost-effective access to space, (b) improved satellite maneuverability and increased on-orbit lifetime, (c) reduction of environmental and toxicological hazards, and (d) improved safety associated with propellant storage and handling.

Terminal Ballistics for Lethality and Protection Sciences


PI: Robert Doney, (Army Research Laboratory): This Frontier Project seeks to advance the state-of-the-art in terminal ballistics by focusing on three themes: rigorous uncertainty quantification of ballistic events, understanding human response to ballistic loading, and developing breakthrough *capability* for materials modeling. Advances in materials modeling, in particular, require a multiscale approach as many physical processes happen at length scales shorter than what can be resolved in



the continuum. Successful investigation of these three themes will further develop the quality of shock physics codes which provide critical modeling capability to enable significant DoD advances in armor, lethality, and hypervelocity impact across multiple military platforms.


HPC Modeling and Simulation to Support DoD Rotorcraft Acquisition Programs

PI: Andrew Wissink, (Army Aviation Development Directorate, AMRDEC): The overall goal of this Frontier Project is to demonstrate, for the first time, that high-fidelity modeling and simulation tools, such as the CREATE-AV Helios flow solver, can have a direct significant impact in the acquisition of major defense systems by reducing their cost, development time, and risk. This project will perform high-fidelity multidisciplinary computational modeling and simulation for the JMR-TD and CH-47 DoD acquisition program vehicles. These computations will be used to evaluate performance, loads, vibration, noise, and safety of flight in the proposed new designs prior to and during their upcoming flight tests.




Validation of Turbulence and Turbulent Combustion Models for Air Force Propulsion Systems

PI: Venkateswaran Sankaran (Air Force Research Laboratory): The goal of this Frontier Project is the development, validation, and application of advanced turbulence and turbulent combustion models designed specifically for Air Force propulsion applications, including gas turbines, scramjets, and rockets. Reacting direct numerical simulations (DNS) and large-eddy simulations (LES) coupled with existing and new turbulence, combustion, and turbulent combustion models will be evaluated using a hierarchy of unit physics, canonical, and grand challenge problems in gas turbines, augmentors, rockets, and scramjets.



Earth System Prediction Capability

PI: Alan Wallcraft (Naval Research Laboratory): The overall goal of this Frontier Project is to perform the R&D necessary to produce the Navy's contribution to the national Earth System Prediction Capability (ESPC). Specifically, this will be our first operational global long-range coupled forecast system for the atmosphere, ocean, sea ice, and waves that extends beyond a week to a month or more. The core components of this ESPC system are the Navy's current global prediction models for seven-day forecasts. Data assimilation will also initially use the Navy's current separate atmosphere and ocean products loosely coupled via the coupled forecast model as a first approximation. Multi-year re-analyses and re-forecasts are used to test and



understand the system. The target for initial operational capability is a 30-day ensemble forecast, but much of the testing will be with 45- or 60-day re-forecasts, since the expectation is to extend the range for full operational capability.

User Productivity Enhancement, Technology Transfer, and Training (PETTT)

The user Productivity Enhancement, Technology Transfer, and Training (PETTT) program's mission is to enhance the capabilities and productivity of the HPCMP's user community through HPC-related science and technology support, training, collaboration, tool development, technical expertise supporting software development, technology tracking and transfer, and outreach to users. PETTT has numerous industry and university partners.

PETTT accomplishes its mission through the following activities: planned and special projects, an HPC training program and outreach, and reactionary assistance to users.

Planned and special projects

In 2016, PETTT developed 10 new pre-planned projects through a competitive review process, in the following technical focus areas: algorithms for scalable parallelism, accelerators, parallel mesh generation, parallel I/O, and multiscale methods. A few highlighted projects include (1) Adaptation and Modernization of Dissipative Particle Dynamics Methods into the LAMMPS Software, (2) Characterization and Optimization of Data Movement, (3) Strategies for Large-scale I/O for Application Performance Improvement, and (4) Scalable Software Framework and Algorithms for Parallel Mesh Generation and Adaptation. Another major goal of PETTT is to provide HPCMP users with state-of-the-art software tools and environments. As part of the Performance Enhancement Thrust Area (PETA), PETTT installed two key performance tools, OpenSpeedShop and Tau Commander, on all major DSRC systems, creating a uniform environment for evaluating code performance. In 2016, PETTT initiated the Data Science Analytics Thrust Area to identify application areas for High Performance Data Analytics (HPDA), and help enable Big Data(BD)/Data Analytics(DA) research within the HPCMP.

HPC Training Program and Outreach

PETTT also delivers a broad range of HPC-related training to the HPCMP user community. In 2016, PETTT offered 57 training events attended by 828 students. Thirty-one training events (62%) were webcast; all 31 of these events were posted online for future viewing. Some examples of training events include: (1) *In Situ* Analysis and Visualization, (2) Advanced Heterogeneous Computing Systems for High Performance Multidisciplinary Simulations, (3) Code Performance Enhancement Analysis Tools and Practices, (4) Enhancing Many-core Performance with Kokkos, and (5) SGI User Training. PETTT also supported 21 additional training events attended by 247 users (195 in person and 52 webcast).

Working closely with the CREATE team, ten of these events directly supported training in CREATE software (SENTRI, Kestrel, and Capstone). PETTT also organized a Technical Interchange Meeting on High Performance Signal Processing with Accelerators with 64 participants, which included tutorials from leading experts in academia and industry.

Reactionary Assistance (RA)

PETTT reactionary assistance is Tier 3 technical support for HPCMP users. Users in the HPCMP community submit real-time requests for expert advice and assistance in developing and using HPC applications. In 2016, 457 RA tickets were opened and 444 RA tickets were closed. The average time-to-close was 29 days. Among the projects this year, the HPCMP benchmarking team requested PETTT conduct a survey of current HPC benchmarks across US Agencies with significant national HPC involvement (DOE, NASA, NOAA, etc.) and the PRACE consortium representing supercomputing in the European Union. This assistance was needed to inform the HPCMP's Technology Insertion benchmarking efforts for FY18 and beyond.


PETTT supports HPCMP goals and priorities including: benchmark development for Technical Insertion, scalable codes for hypersonic applications code refactoring for increasingly complex heterogeneous supercomputing architectures, and data-intensive computing and decision analytics.

Data Analysis and Assessment Center (DAAC)

While supercomputers generate enormous quantities of data, transforming those data into *insight* through scientific visualization is the mission of the HPCMP Data Analysis and Assessment Center (DAAC). Scientific visualization is a branch of data science that is concerned with visually analyzing and interpreting data. The data being analyzed may be the product of an experiment, observation (such as information collected from sensors), or computational simulation. The goal of analyzing and interpreting data, and hence scientific visualization, is to gain insight and understanding into the processes being studied.

Scientific visualization also includes the placement of scientific phenomena into visual context to better understand or communicate the data under study. For example, laminar fluid flow around a waterborne hull form in isolation may be insightful, but an animation of an actual vessel moving on or through the sea, combined with the fluid flow data, now takes on a contextual realism that is significantly more visually compelling and communicative. In FY16, DAAC, working with Army Research Laboratory researchers, produced a video entitled "High Resolution Visualization and

Simulation of High-Pressure Primary Atomization using a Complex Diesel Injector" that was one of six finalists for Best Visualization at SC16, the premier international supercomputing conference. With 13,000 people attending this conference, making it as a finalist is a prestigious accomplishment.



Whether by simply making visualization software resources available to HPCMP users, or by working closely with them on a custom project that includes animation, narration, and video production, DAAC provides the tools and talent to help DoD scientists and engineers gain insight into problems of interest to the Department and the nation.

2.4 Workforce Development

Through its Workforce Development initiatives, the HPCMP supports students and faculty at the Department of Defense Service academies, and supports US students at colleges and universities around the country through the HPC Intern Program by providing funding for summer interns to work with DoD researchers on projects using HPC. Since 2001, the Workforce Development effort has supported STEM education objectives, and the development of the next generation of science and technology leaderships necessary to sustain the nation's military technological superiority and develop the systems vital to the Third Offset Strategy.

Specific Highlights from 2016:

US Naval Academy

The HPCMP initiative at the US Naval Academy (USNA) has focused on encouraging the awareness of and fostering the use of HPC technology in faculty research, thereby providing midshipmen with the tools, techniques, and talents to become leaders in the DoD science and research communities. This year the USNA made a huge leap forward in HPC capabilities with the acquisition and installation of a Cray XC30-AC, the architecture of which was chosen to match the architecture of larger DoD HPC systems. The machine will provide significant computational power for researchers and students in the years to come.



Midshipmen at Grace Ribbon Cutting

Furthermore, HPCMP funding provided for five HPC-related internships with nine students and five professors—research that resulted in five papers, including a Best Undergraduate Presentation Award: MIDN 1/C Mark Schnabel. Lattice Boltzmann

Modeling of Turbine Tip Gap Leakage. American Nuclear Society Student Conference, 2016.

The USNA's efforts are focused through the "Center for HPC Education and Research" which includes 19 faculty members from five different departments across two divisions. The Center received the 2016 HPCMP Hero Award for Team Achievement for its efforts at USNA.

US Military Academy

High performance computing at the US Military Academy (USMA) spans multiple disciplines and academic departments. In FY16, cadets and faculty from USMA used more than 7 million core-hours on DoD HPC resources to study computational fluid dynamics for turbomachinery applications, structural mechanics for blast response, parallel computing algorithms, atomic molecular and optical physics, and computational chemistry. More than 25 cadets were involved in this research that led to seven publications and thirteen presentations or posters at technical conferences and symposia.



Cadets at Tapia Conference


Several recent graduates have continued their work on HPC as they transition to their new roles as junior officers.

A new course offered by the Math department, "A Survey in Computational Engineering Mathematics," was offered for the first time in spring 2016. This course taught cadets about modeling and simulation, and featured guest lecturers from multiple academic departments. Finally, a large-enrollment introductory engineering course, "Fundamentals of Engineering Mechanics and Design," underwent an extensive redesign over the summer, and now prominently features digital simulation-based engineering leveraging the HPCMP Portal.

US Air Force Academy

In FY16, HPCMP support funded summer research for seven US Air Force Academy (USAFA) cadets. USAFA has cadets from the departments of Computer Science and Electrical Engineering, and Aeronautics and Astronautics participating in this program. More than 120 cadets have used DoD HPCMP supercomputers for their academic courses over the past two academic years. All of the Academy's Aeronautical Engineering majors take at least one course employing HPC-enabled modeling and simulation.

HPCMP funding has helped support faculty and cadet research in the areas of computational structural mechanics, computational aerodynamics, weather modeling, computational modeling of austere airstrips, aero optics, and modeling of new optics for data communications and lasers. The newest areas of support are hypersonic gas dynamics and parallel optimization of legacy codes for next-generation computational hardware (such as the Intel Xeon Phi and NVIDIA GPGPUs).



During the past two years, HPCMP funding supported research that has been published in at least seven journal articles (including one from the Department of Civil Engineering, a new area for HPC support), and more than 17 conference papers. This research includes investigation of aircraft structural life extension, operations from unimproved airfields, weather modeling, air-to-air refueling, and space situational awareness in support of the Air Force mission.

HPC Internships

The HPC Internship Program (HIP) provides an introduction to the knowledge and skills needed to employ high performance computing during a career in science and engineering that many students would otherwise not obtain. The HPC Intern Program provides funds to pair a promising undergraduate or graduate student with a Department of Defense scientist or engineer in a DoD laboratory or test center to conduct DoD research. The research must use HPC tools, resources, and methods.

In FY16, 37 HIP proposals were submitted, which made the selection process extremely competitive. Due to the exceptional response and the quality of the proposals submitted, additional funding was identified which allowed funding for 20 total projects (instead of only 12 originally) and 31 total interns, plus two interns funded by matching organizational funds.

Deliverables for each HIP project include a research paper, a presentation (attended by members of their organization and HPC Workforce Development personnel), and a poster submitted by each intern, as well as a quad chart from the mentor summarizing the project efforts and results.

Service/Org	Project Title	School
AF/AFRL	HIP Studies of Extreme Light/Target Interactions	Ohio State University
AF/USAFA	Improved Aerodynamic Modeling and Simulation of Aerospace Vehicles Using High Performance Computing	University of Colorado Springs
AF/AFRL	High Throughput Bioinformatics Methods for Analyzing Nucleotide Polymers Demonstrating Human Physiological Effects of Interest to the United States Air Force	Miami University
AF/AFRL	High Performance Computing (HPC) for Real-time Automatic Target Recognition in Contested Environments	University of Florida University at Buffalo
AF/AFRL	Parallel profiling and optimization of advanced plasma simulation framework	University of California Los Angeles (UCLA)
AF/AFIT	4D Weather Effects Visualization for Future Fires Decision Aids	Cedarville University Johns Hopkins University
ARMY/ARL	Atomistic Understanding of the Role of Grain Boundary Structure in Lightweight Armor Systems	Drexel University

ARMY/ARL	High Fidelity Simulation of Engine Jet Combustion and Application to Turbomachinery	Kansas State University Ohio State University
ARMY/TARDEC	Amphibious Vehicle Hydrodynamics Simulation (CFD)	City College of New York
ARMY/ARL	Tiltrotor Whirl Flutter Analysis with CFD/CSD Coupling	University of Alabama
ARMY/ARL	Molecular simulation of thermal transport in mechanically-stressed energetic materials under pressure: Polycrystalline nitromethane with complex interfaces	University of South Florida
ARMY/ATEC	White Sands Missile Range Hadoop Parallel and Reliable Sensor Data Collection	University of Texas at El Paso (UTEP)
ARMY/TARDEC	Transition and initial utilization of RAMDO software for RBDO simulation of military ground vehicles	University of Iowa
NAVY/NUWC	Water impact of deformable elastic spheres	University of South Florida
NAVY/NSWC	Computing Buried Blast Effects on Military Vehicles and Test Structures	University of Florida
NAVY/NSWC	Technology Insertion Set-Based Design Study	University of Michigan SUNY Maritime
NAVY/NRL	Parametric Sensitivity and Reliability Analysis of Vacuum Electron RF Device Models Using DoD's High Performance Computing Resources	University of Maryland, College Park
NAVY/NRL	Toward the Next Generation of High-Fidelity, Large-Eddy-Simulation Models for High-Speed Turbulent Combustion	University of Colorado Boulder
NAVY/NSWC	Ship Hull Form Optimization Tool Implementation	Virginia Commonwealth University, University of Michigan
NAVY/NSWC	Heterogeneous Computing for Unstructured Engineering Simulations	Case Western Reserve University

Table 5. 2016 HIP Interns


3.0 HPCMP Contributions to Mission Success

These stories are Distribution C until cleared

AFRL Gotcha Radar Program

PI: Linda Moore, Air Force Research Laboratory

The vision of the Air Force Research Laboratory Sensors Directorate (AFRL/RY) Gotcha Radar Program is to provide persistent, continuous, all-hour, all-weather, forensics-capable, radar-based surveillance of wide areas. Using synthetic aperture radar (SAR) technology, the Gotcha Radar surveillance system will provide the ability to generate a radio frequency (RF) hologram of a ground spot, e.g., a city that can be processed in different ways to simultaneously perform what would otherwise be the functions of several radars on several platforms. In other words, a single data stream will be employed for all exploitation - a departure from the philosophy of traditional radar systems, which perform specific tasks through distinct radar modes. The program anticipates that some exploitation products will be created in real-time using on-board processing, while other products may be created in near-real-time, or several hours or days later, i.e., forensics, using an RF data link and a ground-based processor.



The success of the Gotcha Radar Program depends largely on high performance computing (HPC) technology and the support of the DoD HPC Modernization Program (HPCMP). Technical challenges for the Gotcha Radar Program include real-time computation, on-board and/or ground-based processing, RF data communications, and radar hardware design for wide-area surveillance. The Gotcha Radar Program first leveraged HPC computational resources at the AFRL and ARL DSRCs. The program continually engages the HPCMP PETTT program for advice and support of employing HPC for its data-intensive SAR applications. The program's Dedicated HPC Project Investment (DHPI) system, "Desch" addressed real-time computation requirements. Desch is comprised of an SGI Altix ICE 8200 LX system, which is a distributed memory cluster, and serves as the main computational component for Video SAR processing; and a 32-core SGI Altix 450 system, which is a smaller shared memory parallel system, and serves as an auxiliary system for other specialized processing, testing, and data ingesting.

A future deployable Gotcha Radar system will not only leverage previous research performed on HPC systems, but will indeed include one or more HPC systems to perform real-time processing of surveillance SAR data in the field. The Gotcha Radar Program's successful advancement of wide-area surveillance technology is a direct result of both the availability of many different HPC systems, and the accessibility of leading DoD computing experts, through HPCMP DSRC and PETTT personnel.

Delivery of Accurate Environmental Models to the Naval Warfighter

PI: Alan Wallcraft, Naval Research Laboratory

PI: Wieslaw Maslowski, Naval Postgraduate School

The Naval Research Laboratory (NRL) has a long history of developing, testing, and making operational a number of increasingly accurate atmospheric and ocean models to ensure that operational naval forces have the best possible understanding of environmental conditions under which they operate around the world. In addition, NRL plays a major role in national and international programs to conduct research and development of future predictive models. Recently, a new US law, “An Act to Improve the National Oceanic and Atmospheric Administration’s Weather Research,” specifically requires NOAA to “build upon existing forecasting and assessment programs and partnerships,” one of those being the interagency Earth System Prediction Capability (ESCP), which is the new coupled atmospheric/ocean/sea ice prediction system under development by NRL and sister federal agencies.

The Naval Postgraduate School (NPS) has also been involved in the model development, evaluation, and prediction of environmental conditions, with specific emphasis on the pan-Arctic region. NPS research targets some of the key requirements, including those published in the Navy Arctic Roadmap 2014-2030 (2014), and in the Arctic Research Plan: FY2017-2021, regarding the need for advanced modeling capabilities for operational forecasting, and strategic climate predictions at sub-seasonal to decadal time-scales.

Two DoD Frontier Projects address this important capability. The NRL Frontier Project specifically targets verification and validation of the global Earth System Prediction Capability model as it meets its planned initial operating capability (IOC) for Navy forces in 2018. This activity is the latest in a series of improvements to the Navy’s capability to make weather/ocean modeling predictions for the operational fleet that has been accomplished on HPCMP systems over the past 25 years. The NPS Frontier Project has developed and used very high-resolution Arctic Ocean and sea ice models within the Regional Arctic System Model (RASIM), which is at the forefront of Arctic ice and climate simulation to reduce uncertainty and advance predictive capability for the Navy beyond the operational (i.e. 7-14 days) time-scales.

The NRL Frontier Project “Earth System Prediction Capability” is performing validation of the Navy’s fully-coupled atmospheric/ocean/sea ice global model for weather/ocean predictions out to at least one month. This work builds upon previous NRL work that has developed and tested the individual atmospheric, ocean, and sea ice modeling capability; the coupling of these models will allow more accurate environmental forecasts for a longer time-span. This Frontier Project builds on several years of initial development of ESCP. Accomplishments in 2016 include re-tuning of the NAVGEM global atmospheric model to run longer than the 7 days typical of uncoupled prediction systems. The Madden-Julian Oscillation was a particular focus because it is the largest element of the 30-90 day variability in the tropics, and depends on air-sea interaction. The figure below shows the improvement in rainfall characteristics between V1 and V2 NAVGEM physics.

The NPS Frontier Project “Advancing DoD Modeling and Prediction Capabilities in the Arctic” also builds upon considerable prior work in developing and testing a regional capability to predict Arctic sea ice

over several decades. The project made significant progress in 2016, including its first seasonal prediction of the minimum Arctic sea ice extent in 2016, as a contribution to the community Sea Ice Prediction Network (SIPN), which is an activity supported by the Interagency Arctic Research Policy Committee (IARPC). This project has also developed procedures for systematically quantifying climate model skill by using a combined metric that includes both variance and correlation between modeled and observed quantities. Use of these techniques consistently showed that the Regional Arctic System Model (RASAM), used in this project, offers improvements in representation of the mean states and variability of the Arctic Ocean, compared to global climate models offering similar predictive capability. Comparisons between RASAM and several global climate models that participated in the Fifth Coupled Model Inter-comparison Project (CMIP5), observational climatology (PHC), and recent hydrographic data from ice-tethered profilers proved its overall improved accuracy. In addition, these comparisons demonstrated the need for better representation of the upper-ocean hydrography in the Arctic, which can store heat absorbed from the atmosphere during the summer, and has been hypothesized as a potentially important heat source accelerating sea ice melt above. These detailed studies and ongoing model improvements set the stage for more accurate predictions of the extent and thickness of sea ice in the Arctic over coming years.

Future Tactical Truck System (FTTS): R&D to Acquisition

David Lamb, TARDEC

The US Army is preparing for the future by providing state-of-the-art technology for the warfighter. Ground vehicles are critical for both soldier and marine warfighters, and replacements for the aging fleet of High Mobility Multi-Purpose Wheeled Vehicles (HMMWVs) are a pressing need. This spurs the Joint Light Tactical Vehicle (JLTV) program, which is a key enabler for a substantial increase in effectiveness of ground forces for both the US Army and the US Marine Corps. An important technology input to this program is the Future Tactical Truck System (FTTS) Advanced Concept Technology Demonstration (ACTD). As the primary defense laboratory for military ground vehicle systems, the US Army Tank-Automotive Research, Development, and Engineering Center (TARDEC) in Warren, MI, is supporting the Army and the Marine Corps by providing advanced technical solutions to improve the ground fleet.

TARDEC was the technology developer for FTTS. During the modeling and simulation (M&S) phase, vehicle concepts were developed with significantly increased capabilities in survivability, deployability, mobility, connectivity, and capacity using HPCMP computational resources. Three prototypes were built. M&S was required for the JLTV to:

- Model propulsion, duty cycles, and fuel economy for traditional and hybrid electric powertrains
- Simulate vehicle signatures, protection levels, lightweight materials, adaptive variable suspensions, and underbody blast threats
- Develop numerous advanced concepts as requirements for transportability, payload, mobility, and sustainability evolved as part of a Whole Systems Trade Study

The importance of this support to FTTS was recognized by the Deputy Assistant Secretary of the Army for Research and Technology:

“The Future Tactical Truck System (FTTS) is a good example of how Army S&T seed corn enabled the Joint Lightweight Tactical Vehicle (JLTV) acquisition program...TARDEC led the S&T technology development effort that demonstrated mature technologies which ultimately led to realistic and achievable requirements.” —Ms. Mary Miller, then Dep. Asst. Sec. Army for R&T, Feb. 2016

HPCMP resources were critical for performing the required M&S rapidly enough to influence the design cycle. If it had not been for the ability to use HPC for these simulations, FTTS would either have not used M&S, or would have experienced significant delays. Rapid simulation allowed for design iterations that improved the FTTS, and ultimately JLTV, by increasing all the different “-ilities” that are best assessed using M&S during the design stage, prior to any prototypes. HPC enabled this, making a significant impact to the acquisition.




Oshkosh JLTV



FTTS Prototype

MQ-4C Triton – Maritime Patrol UAV

Gary Estep, NAVAIR



The MQ-4C Triton is a DoD maritime surveillance unmanned aerial system (UAS), under developmental test and evaluation by the US Navy and Northrop Grumman, with an early operational capability (EOC) date planned for 2018.

The MQ-4C Triton provides real-time intelligence, surveillance, and reconnaissance (ISR) including vessel detection, tracking, and classification over vast ocean and coastal regions. The aircraft is an integral part of the US Navy (USN) Maritime Patrol and Reconnaissance Force's (MPRF) core capabilities, and fills a vital role for delivering persistent maritime ISR by enhancing situational awareness of the operational environment and shortening the sensor-to-shooter kill chain. The MQ-4C Triton will bring significant awareness of the maritime environment with the capability to maintain five continuous operational sites (orbits) around the globe. Collected data from the Triton will be posted to the Global Information Grid (GIG) and support a variety of intelligence activities and nodes. In addition, the MQ-4C will be used to respond to theatre-level operational or national strategic tasking.

The MQ-4C Triton system consists of multiple capabilities. Its multifunctional active sensor suite provides unprecedented 360-degree maritime domain awareness for the US Navy. Triton can fly up to 24 hours at altitudes over 50,000 feet with a range up to 8,200 nautical miles, and can monitor 1 million square miles of ocean in a single flight, using high-resolution images and full motion video. The aircraft incorporates a reinforced airframe for increased internal payload and wing for hail, bird strike, and gust load protection, along with de-icing and lightning protection systems. These features allow the aircraft to descend and ascend through harsh maritime weather environments to gain a closer view of targets at sea.

MQ-4C Triton was flown cross-country for the first time from Palmdale, CA to Patuxent River, MD in 2014 to test the platform's ability to: (1) remotely receive and execute long-distance flight plans, and (2) communicate with a simulated operational battlefield network during flight. The MQ-4C Triton program achieved its 100th test flight in October 2016, and began flight tests with integrated functional capability (IFC) in March 2017. It is now preparing for EOC. Flight and ground tests between the UAS and ground support stations are conducted daily, and rely on the HPCMP's DREN/SDREN for connectivity.



Hypersonic Aircraft and Missiles


Ryan Gosse, AFRL

Alexei Poludnenko, AFRL

A recently published study by the Mitchell Institute ("Hypersonic Weapons and US National Security: a 21st Century Breakthrough," January 2016) discussed the breakthrough capabilities of hypersonic weapons: projection of striking power at longer range while minimizing exposure to increasingly sophisticated defenses; compression of the shooter-to-target window; providing the opportunity for numerous types of strikes, and enhancement of future joint and combined operations. The study concluded that a focused commitment to sustained hypersonic research, development, test and evaluation (RDT&E) will produce usable hypersonic weaponry within the next decade.

HPCMP provides the necessary HPC resources for two Frontier Projects to address complementary aspects of hypersonic flight, both currently sponsored by the Air Force. One, "Unsteady Pressure and Heating Environment on High-Speed Vehicles with Responding Structures," aims to demonstrate fully-coupled high-fidelity fluid-thermal-structural response of high-speed vehicles at full-scale that resolves all relevant physics of the fluid (minus combustion chemistry) and detailed structure domains. As a practical matter, this goal translates to large-eddy simulations (LES) of the HiFiRE 6 hypersonic vehicle at full-scale by the end of the project in FY18. The other, "Dynamics and Properties of High-speed Turbulent Reacting Flows," seeks to develop a fundamental understanding of the interplay of combustion and turbulence in high-speed jet engines. Both made considerable progress toward these goals in 2016.

During 2016, the full-scale hypersonic vehicle modeling project responded quickly to a change in the Hi-FIRE 6 program, and strongly influenced the design and execution of wind tunnel tests of the vehicle in AFRL and NASA wind tunnels. The project used a variety of computational fluid dynamics (CFD) codes to calculate several points along a hypersonic flight path for HiFiRE 6 and the results from these cases included variations of angle-of-attack, presence of appendages, and simplifications of the geometry. Comparisons with wind tunnel results validated the major features of the CFD calculations, and showed that these calculations must be done using LES, instead of commonly used RANS techniques, in order to correctly simulate key features associated with the engine inlet and the flow through the engine. In addition, the project was able to demonstrate a framework for conducting fluid-structure interactions for this vehicle. Over FY17-18, the team plans to focus on analysis with this capability. These results have been used to plan and define the most useful wind tunnel tests to further study HiFiRE 6 vehicle performance throughout its hypersonic flight envelope.



Computational Fluid Dynamics (CFD) simulations of HiFiRE-6 Adaptive Flight Control Experiment

For the jet engine simulation project, one of the primary results is a better understanding of the coupling between high-speed turbulence present in jet engines and supersonic/hypersonic combustors and the complex chemical kinetics of realistic jet fuels. A systematic study was completed with advanced direct numerical simulation (DNS) techniques surveying the properties of a range of hydrocarbon fuels from methane to heavy hydrocarbons, primarily, single-component (dodecane) and multi-component jet fuels (Jet A2). The primary focus was on high-pressure, high-



Structure of a hydrogen-air turbulent flame formed in the decaying upstream turbulence.

temperature conditions and a broad range of turbulent intensities, from relatively low to very high, characterized by large compressibility effects such as present inside SCRAM jet combustors. The results of this study showed that under virtually all turbulent conditions, all heavy hydrocarbon fuels exhibited remarkably similar behavior, which, however, was pronouncedly different from that of light hydrocarbons. A separate line of inquiry in this study focused on the nature of the turbulent energy cascade in reacting turbulence and, in particular, on the so-called backscatter phenomenon, which involves the ability of small-scale combustion processes to energize large-scale turbulent motions in the system. These results are crucial for the accurate modeling of high-speed reacting flows. Armed with this detailed understanding of turbulent combustion in extreme regimes, this project is now ready to move to modeling real combustors with a realistic range of length- and time-scales, which will be the focus of this project for its two remaining years.

Pacific Missile Range Facility (PMRF) Ensures Safety of MDA Missile Defense System Tests

Danielle Franklin, PMRF

When the Missile Defense Agency (MDA) conducts flight tests of missile defense systems, safety of the surrounding area is essential for successful conduct of these tests. PMRF is intimately involved in ensuring the safety of these tests, and utilizes the significant resources of the HPCMP in support of that goal. Recently, on 17 May 2016, MDA, the US Pacific Command, and US Navy sailors aboard USS John Paul Jones (DDG 53) successfully conducted a flight test involving the launch of a medium-range ballistic missile target from PMRF located on Kauai, Hawaii. The flight test, designated Flight Test Other-21 (FTX-21), successfully demonstrated the ability of an Aegis Baseline 9.C1 (BMD 5.0 Capability Upgrade) configured ship to detect and track a medium-range ballistic missile (MRBM) target within the Earth's atmosphere. In addition, MDA and Navy sailors aboard the USS Hopper (DDG 70) successfully conducted two developmental flight tests of the Standard Missile-3 (SM-3) Block IB Threat Upgrade guided missile on 25 and 26 May 2016 off the west coast of Hawaii. The flight tests, designated Controlled Test Vehicle (CTV)-01a and CTV-02, demonstrated the successful performance of design modifications to the SM-3 third-stage rocket motor (TSRM) nozzle. The results of these flight tests will support a future SM-3 Block IB production authorization request. For both of these crucial tests PMRF range safety experts provided required safety evaluations before and during the event.

Specifically, PMRF range safety support begins with safety design analysis, using HPC modeling and simulation on HPCMP supercomputers, with access provided via SDREN, that determines risk to civilian population centers, the firing Naval vessel, test barges, commercial shipping, fishing craft, and commercial and participant aircraft. Then, a design test concept is created, which determines, models, predicts, and simulates worst-case failure maneuvering of targets and missiles by providing model impact debris fields on a test range of 42,000+ square miles of controlled airspace and a temporary operating area of over two million square miles. Access to classified HPC capabilities allows runs from six-degree-of-freedom (6-DoF) debris caused by breakups from a structural failure or flight termination event, and these types of analysis runs may take on the order of two weeks of HPC utilization time. This analysis was previously contracted out with lead-times ranging from three to six weeks. In order to structure events concepts of operation, the range safety analyst must first model the predicted failures, then create flight termination limits and propagate the debris to impact. In the event the data slightly changes, which regularly occurs, the impacts to the DoD and the range can highly impact cost, schedule and performance risks. With the use of HPC, not only is PMRF Range Safety able to bring the capability in-house, but the analysis results are turned around in record times. Day-of-launch wind analysis is another crucial part of this process. Due to lack of resources and limited time, the ranges need to use computational techniques to predict worst-case winds. The HPCMP, with its ability to process this data more quickly has allowed PMRF Range Safety to adopt new standards to improve mission safety and success. This analysis requires 250,000 core-hours on classified HPCMP systems per mission. As mission complexity and safer, smarter modelling becomes available, additional high



SM-6 Launch
Photo Credit US Navy

performance computing time will be required to support critical national defense tests. For missions such as FTX-21, use of HPCMP assets made it possible for Range Safety to be intimately involved with the issues and concerns to ensure a safe and successful test. Estimated cost savings for each major test event like FTX-21 are \$500K. For 2016, there were at least 4 major test events with estimated cost savings to DoD of approximately \$2M. This work is making a critical difference in the development of viable missile defense systems.



A Standard Missile-3 (SM-3) Block 1B interceptor missile is launched at PMRF
Photo Credit US Navy


Multi-scale Interactions in Stratified Turbulence

Stephen de Bruyn Kops, ONR

Turbulence is present under many fluid flow conditions, and has been very difficult to model accurately. Stratified turbulence, for which lighter fluids sit atop heavier fluids and can stabilize turbulence, is even more difficult to model because the range of eddy sizes and time-scales that need to be modeled is even larger than unstratified turbulence. Stratified turbulent conditions prevail in many environments that are important to DoD, including the deep ocean, the stratosphere, the atmospheric boundary layer over land at night, and the boundary-layer over ice. This Frontier Project is developing a combined Direct Numerical Simulation (DNS) – Implicit Large-Eddy Simulation (ILES) simulation of stratified turbulent wakes at high Reynolds numbers to predict turbulence signatures in the deep ocean, and high-fidelity idealized wakes from a submerged object to understand turbulence development, radiated internal wave fields, and their interaction. ILES-simulated internal wave strain field analysis indicates that wake-emitted waves, at least in the context of an idealized stratification, can produce a remotely detectable surfactant pattern at operationally relevant wake parameter values. For a given depth of the submerged wake, specific wavelengths arrive at a predictable time at the surface.

High Reynolds number stratified wake datasets analyzed in this study are the result of a combination of high-accuracy, high-resolution ILES and DNS. ILES ran the full evolution of a stratified wake at body-based Reynolds numbers up to 400,000, which is on the order of 100 times larger than its laboratory counterpart, and at three different body-based Froude numbers. By comparing these simulations with simulations matched to the laboratory experiments, the flow regimes of the simulations with high and low Reynolds numbers are shown to be markedly different. Snapshots in time of ILES full-wake velocity and density fields were interpolated onto a DNS grid and advanced in time until small-scale motions established quasi-equilibrium with respect to large-scale motions. By this method, exact solutions of the flow at all length- and time-scales are computed, in particular in the regions of the wake in which the turbulence-generated waves are produced.

A dynamic range of turbulent scales up to nearly two decades wide was found inside the wake core, which progressively diminishes as one approaches the turbulent/non-turbulent interface. Both turbulence and internal gravity waves impart motion to the fluid, but the mechanisms are much different. It is important to distinguish the two, both for understanding the flow physics and for developing models for fast simulations. The overall flux of energy released into these internal waves is found to be comparable to the overall rate of energy loss due to viscous kinetic energy dissipation within the wake. Wavelet analysis of the surface-incident radiated internal waves in the ILES data revealed a universal power law for the time evolution of the characteristic length-scale of the surface-incident internal waves. This universal power law allows an extrapolation of these results to high Reynolds number above that accessible by research simulations. To verify this extrapolation, a large simulation at higher Reynolds number will be run. Finally, at the higher Froude and Reynolds numbers considered, the wave-induced strain field is capable of producing remotely observable slicks at the surface.




Simulation of stratified turbulence.

End-to-End Modeling of Laser Directed Energy Devices


Timothy Madden, AFRL

In the Directed Energy (DE) Directorate of the Air Force Research Laboratory at Kirtland AFB, NM, the DE High Performance Computing Software Applications Institute (HSAI) developed end-to-end models of high power electromagnetic (HPEM) devices that have been used to directly impact the design of weapon systems. The HPCMP is supporting a follow-on HPCMP application software initiative (HASI) project to develop the same types of end-to-end models of laser production devices. These devices include fiber lasers, diode-pumped alkali lasers (DPALs), chemical oxygen-iodine lasers (COILs), and bulk solid-state lasers. The modeling capability will also include propagation of the laser light from the platform through the atmosphere to the target and into the interaction of the laser light with the material in the target.

In FY16, the DE High Performance Computing Software Applications Institute (HSAI) made great strides in the pursuit of this laser end-to-end modeling effort. First, models for both DPAL and fiber lasers have been integrated successfully into the DE HASI suite of software and run in optimization loops within Galaxy, software produced by the institute to enable end-to-end simulations combined with advanced analysis, HPC execution, and collaborative data-basing all in one tool. A setup for the Galaxy simulation of the DPAL source is shown in the figure below. This simulation allowed researchers to determine parameter values for the DPAL that were not experimentally measurable. Another critical component of laser end-to-end modeling is the propagation through the atmosphere. To model this, the High Energy Laser End-to-End Operational Simulation (HELEEOS) code was also integrated into Galaxy. This enabled researchers to quickly process 170,000 weather data points on DoD supercomputers to produce beam on target irradiance at each weather data point and resulted in a 1000 times speed up of atmospheric laser propagation modeling by moving the execution from a single PC. Finally, the Institute has developed the Galaxy Coupling Manager (GCM) to provide a time-dependent architecture for coupling of codes that represent various parts of the laser end-to-end model that require time-dependent coupling during the simulations.



Aero-Optic effect phase screen



Beam on target Irradiance modeled with HELEEOS

Assessing Molecular Contamination of Spacecraft


Michael Weaver, SMC

The deposition of molecular contaminants onto a space vehicle (SV) can reduce the performance or the mission life of the SV. Maximum allowable molecular contamination during pre-launch processing and launch of an SV is typically limited to tens or hundreds of Angstroms. One source of contamination during launch is the impingement of any plumes from the launch vehicle (LV) upper stage onto the SV. For instance, plumes from upper-stage engines that burn fuel containing hydrocarbon molecules may result in molecular contaminants that adhere to an SV surface. For this reason, pre-flight predictions of contaminant levels on the SV are required in order to assess if the requirements can be met.


A challenge is that the relevant flow regimes vary from continuum conditions for the plumes, to rarefied conditions surrounding the SV, and free-molecular conditions in the far field. With independent research and development funding from The Aerospace Corporation, a methodology was developed to span the flow regimes by coupling a Reynolds-averaged Navier-Stokes (RANS) model to a Direct Simulation Monte Carlo (DSMC) model, giving an integrated solution for the entire domain. Timely utilization of this numerically intensive methodology is enabled through utilization of DoD supercomputing resources.

Illustrating one application, the deposition of molecular contaminants generated by an upper stage engine onto a representative SV was numerically modeled. The first phase of the analysis addressed the convection and deposition of products of incomplete combustion during steady state flight. In this phase, an axisymmetric RANS solution of the nozzle flow-field was coupled to a DSMC simulation of flow around the upper stage and SV. The second phase of the analysis modeled the evaporation, convection, and deposition of unburned fuel droplets released during engine shutdown. In this phase, 1-D droplet evaporation and convection models were employed in conjunction with the RANS and DSMC modeling tools used in the first phase to estimate contaminant deposition. Combining the results of the first and second phase of this analysis with a typical mission profile allowed for conservative estimates of total molecular contamination due to upper-stage engine operation to be made.

This coupled methodology has also been utilized to assess contamination environments and plume-impingement heating, supporting missions and launch vehicles for National Security Space.



(A) Contours of density and streamlines for nozzle flow-fields. Black dots represent location of continuum breakdown surface.




(B) Number density contours in DSMC flow-field. Contour scale is logarithmic and spans eight orders-of-magnitude.

Optimizing the Performance of Diesel Engines for Army Systems


Luis Bravo, ARL

Fuel injection sprays are a critical part of today's military propulsion platforms powering US Army and DoD tactical and combat vehicles, such as the Gray Eagle Unmanned Aerial Vehicle (UAV) or the Joint Light Tactical Vehicle (JLTV). These vehicles are powered by diesel engines running primarily on JP-8 or F-24 fuel. To meet and exceed the range and operational requirements, these engines need to provide significant advancements in fuel flexibility and combustion efficiencies leading to highly optimized, high performance combustion engines.

The ARL Frontier project, "Petascale High-Fidelity Simulation of Atomization and Spray/Wall Interactions at High Temperature and Pressure Conditions," has provided significant contributions towards the understanding of fuel injection sprays and spray/wall interactions regimes. This addresses a significant knowledge gap in liquid fuel sprays including spray formation, droplet generation, and droplet impingement, which is crucial to accurately predict the fuel/air mixture formation process to improve combustion efficiencies. An important 2016 research milestone achieved was the quantification via Direct Numerical Simulation (DNS) of the pulsed spray formation process issued from a complex diesel injector subject to rapid transients. The model captured, for the first time, the complex liquid structures including: helical waves, ligaments, surface holes that lead to droplet generation, and turbulence interactions in an engine environment. Further, the DNS was validated against X-ray radiography measurements of spray density fields for two fuels, including an average properties jet-propellant (A2) and a high-viscosity alternative jet fuel (C3). The high-viscosity C3 fuel showed improved atomization quality and faster penetration speeds by 5%. The unique findings are of high relevance to the Army mission. The results reveal the underpinning physics of fuel injection sprays in engines such as Gray Eagle UAV, and is helping the Army develop the next-generation models to design high-efficiency visionary concepts for 2040+ vehicle propulsion platforms.



Frontier Simulation of Turbulent Dense Spray in UAV Engine



Experimental validation of spray penetration of A2 and C3 fuels

Redesign of the V-22 Engine Exhaust System

Allan Aubert, NAVAIR

The V-22 Osprey is a critical air platform for DoD. The MV-22 entered service with the Marine Corps in 2007, and provides troop and supply transport capabilities that are significantly improved over legacy platforms. Components in the engine exhaust airflow path can experience high failure rates. The design of these components was performed long before recent advances in computational flow modeling tools emerged. The NAVAIR Propulsion Internal Flow Modeling Team (NIF-T) has recently developed in-house capability to model and design engine inlet and exhaust airflow paths, quickly and accurately, predicting both inefficiencies from air flow separation and dynamic pressure fluctuations from turbulent and separated flow. A team has been assembled including NIF-T and contractors to propose an extended-life, low-cost, limited redesign of the subject parts. This team is competing against other, higher cost redesign proposals being developed by other contractor teams. Early indications are that the dynamic pressure fluctuations seen in the NIF-T model results were not included in the anticipated loads for current parts. The team is moving quickly to assess this redesign for possible production, which would result in an expected large savings in annual maintenance, and the return of this crucial air platform to greater availability to support the marines, airmen, sailors, and soldiers who use it.




NIF-T is using SBIR-developed, Combustion Research Technology, Inc. (Craft Tech) Large-Eddy Simulation (LES) time-accurate turbulent air flow modeling tools to compute flow-induced forcing on existing parts. Such tools were, until recently, considered to be impractical for these purposes because of the large computer time required to solve cases. However, the availability of HPCMP computational resources is now making the use of these enhanced accuracy, time-accurate tools practical. The final LES model utilized over 100 million cells to model the very complex flow paths in the engine

exhaust. The LES study has identified areas of the flow path that may be redesigned to reduce turbulent forcing on components. This forcing reduction, together with structural redesign considering the newly understood turbulent forcing is supporting a better focused redesign process. NIF-T produced surface pressure time history predictions covering the key flight regimes were delivered in early 2017, and are being input to the contractor Finite Element Analysis (FEA) dynamic models to support redesign for increased life of the parts. Future models will explore the use of a fully-coupled fluid structure interaction model to determine if it will significantly affect the result. The proposed redesign is expected to be finalized by fall 2017, with flight test in 2018.

Atomistic Simulation of an Explosive Formulation

DeCarlos Taylor, ARL



To enable a computational capability for design of next-generation energetic materials (EMs), ARL has made a large investment in the development of a multiscale modeling and simulation (MM&S) paradigm that incorporates all of the length scales (atomistic through continuum) that have to be addressed in order to computationally characterize an explosive. Although MM&S is not new, the novelty of the ARL approach is strict adherence to the principle that all models used in the MM&S hierarchy must be derivable purely from the physical and mathematical considerations underlying the problem of interest. This can be achieved by development of models at the atomistic level using accurate first-principles quantum mechanical approaches with subsequent homogenization and “upscaling” of the first-principles-derived reference data into models that operate at the nanoscale, mesoscale, and continuum levels.

Currently, our target modeling material is the real formulation Composition A3 Type II, which is a two-component formulation consisting of cyclotrimethylene trinitramine (RDX) energetic and polyethylene (PE) plasticizer (see figure above). The overarching goal of this research effort is to demonstrate a completely coupled, multi-scale simulation of this test system with experimental validation of all findings. However, before this can be accomplished, material models being applied in this study have to be developed and validated.


The goal of the Multiscale Response of Energetic Materials program, under which this work is being executed, is the development of advanced energetic materials with increased energy output and decreased sensitivity to initiation. Although advanced experimentation can be used to develop and characterize novel EMs, the experimental development cycle is a costly process that historically requires a minimum of twenty years from initial concept to fielding. However, the overall development time can be significantly reduced via application of advanced modeling and simulation techniques which enable virtual screening, virtual optimization, and virtual exploration of novel materials; thus expediting the design and development of new explosives for use by the Department of Defense.

Composition A3 Type II consists of RDX (91%) and PE plasticizer. As such, atomistic modeling of this system requires not only a proper description of the intra-component (RDX-RDX and PE-PE) interactions, but also the inter-component (RDX-PE) interactions. The RDX-RDX and PE-PE interactions were treated using quantum mechanically-derived potentials available in the literature. The cross-terms (RDX-PE) were developed at ARL by fitting a classical interaction potential of exponential-6 to form to interaction energies of a polyethylene chain on an RDX surface. The resulting potential was validated via comparison to experimental interaction energies obtained using force-distance spectroscopy.

Following validation, the potential was used to perform large-scale shock simulations of Comp A3 Type II in order to determine the effect of the polymer on the response of RDX to shock loading. All simulations were performed using the LAMMPS molecular dynamics software package on 3,200 cores of a Cray XC40 (Excalibur) at the ARL DSRC. The simulation cells contained 2 million atoms, and the shock was initiated by driving a piston with a particle velocity of 1 km/s into the RDX edge of the cell. The material was shocked along multiple orientations including the 100, 010, 001 and 210 directions.

A snapshot of shock-compressed Comp A3 Type II, for a shock propagating along the 100 direction is shown in the figure below. In both panels, only the center of mass of each RDX molecule and PE oligomer is represented, and the bottom panel is color-coded by the internal temperature of each molecule. As shown at the given compression rate, shear bands nucleate in the RDX layer; however, RDX remains crystalline near the polymer interface. This results from stress relaxation that occurs in RDX due to the higher compressibility of the PE at the leading-edge of the shock. It can also be observed that there is significant heating of the sample on the polymeric side of the interface.

The models developed in this work, through derived use of quantum mechanical methods, are non-reactive potentials. Given the significant heating that is evident in the shock-compressed polymer, it can be expected that chemical reactions will occur. At present, the models we have developed cannot account for any degree of chemical transformation beyond molecular conformation; therefore, development of reactive models is currently underway.




Composition A3 Type II following shock compression. Shear banding is evident in RDX layer (upper-panel) and heating of polymer is also observed (lower-panel).

High Performance Computing for Geospatial, Network, and Vehicular Data Analysis

Jim Feight, U.S. Army Test Center

Growing dependence on digital communications and the availability of highly sophisticated communications systems require more complex and detailed test events, producing terabytes of data per day. When aggregated over time, the processing requirements are beyond the scope of commodity computational resources or even small clusters. Scalable, parallel methods are required to transfer and process data into a form that analysts can use in a timely manner.



Geospatial visualization of communication links in mobile tactical network


Prior to using high performance computing (HPC) resources to process test data, a Java-based application was executed on several large (64 core) Linux servers to process collected data. This could involve several days of processing for a single day's collection set.

The design goal for a new, HPC-based process was to complete all data processing and provide the results within several hours. To accomplish this order-of-magnitude reduction in time, the Aberdeen Test Center (ATC) partnered with ARL to develop a prototype HPC framework and created a new automated data transfer and job control pipeline.

Tactical networks support the activities of a military unit during operations, with high assurance and minimal delays, as the unit maneuvers to accomplish its mission. Vehicles are also becoming more complex with Controller Area Networks (CAN) buses that support vehicle operation and status monitoring.

Large-scale tests are required to provide senior leaders with data necessary for system evaluation and milestone decisions. The software developed provides scalable, parallel methods, which transfer data from test ranges, and process it into a form that analysts and decision makers can use in a timely manner. An example of this data is shown in the figure above, which shows temporally controlled geospatially-rendered satellite (blue) and terrestrial (green) communications links between tactical vehicles.

An automated event-driven pipeline (depicted in the figure below) was created to transfer data, marshal HPC jobs, and load results around the clock for multiple data sets simultaneously; it enabled decoupling distributed HPC processing from the results database.



Automated data collection, reduction, and analysis pipeline

A framework was designed to distribute tasks in multiple phases over an arbitrary number of processing cores. Input files are read once per phase, and multiple modules “subscribe” to data record types and receive a copy of the data record for processing.

The framework consists of many data-processing modules, grouped and executed in separate sequential phases based on data dependencies. Early phases are typically dependent on only raw data, while later phases typically rely on only the output of the early phases. Data exchanged between phases is typically performed via files. While this approach increases input-output (I/O) requirements, it is essential for the orthogonal redistribution of data between phases.

The HPC framework produces results with an order-of-magnitude reduction in time compared to previous methods. The results are automatically loaded in a relational database schema that enables analysts and other customers to combine data from multiple sources for determining system reliability and performance.

Future research will replace the current HPC code and database with an open-source Apache Hadoop ecosystem. Hadoop provides a distributed file system and query engine that will alleviate I/O and data movement issues; the processing and querying both run on the same distributed HPC platform. Queries will also gain the speed-up of this distributed system. Web front-ends and interfaces will enhance analysis and expedite decision making.

Turbulent Minority Carrier Species Transport


Neal Blackwell, CERDEC

Several hypotheses exist for the origin of $1/f$ noise (flicker noise) in electronics. While electronics exhibit several types of noise, $1/f$ noise is typically the dominant noise in low frequencies. One hypothesis is that the origin of $1/f$ noise is turbulence in carrier (electrons and holes) transport. The US Army Communications-Electronics Research, Development and Engineering Center (CERDEC) Night Vision and Electronic Sensors Directorate (NVESD), along with consulting expertise from Dr. Paul Norton, PETTT personnel, DAAC personnel and NVESD-Signal Processing personnel, are numerically investigating this hypothesis. Experiments at NVESD are buttressed by state-of-the-art nanoscale experiments by the University of California - Santa Barbara (UCSB).


Discovery of the origin of $1/f$ noise will yield indicators for mitigation. For night vision, mitigation yields force multipliers of 1) improved night vision in extreme low-light conditions, and 2) extended range. Since $1/f$ noise is common to electronic devices, discovering the origin and mitigation of $1/f$ noise will provide improvements for nearly all DoD electronic devices including sensors, communication systems, lasers, displays, computers, and guidance systems (missiles, unmanned aerial systems, unmanned ground systems, manned air/sea/ground vehicles), to name a few. In the same way, the discovery and mitigation of $1/f$ noise would produce similar widespread advancements in civilian electronics.

The HPCMP project entitled “Turbulent Minority Carrier Species Transport” numerically supports this investigation into the origin of $1/f$ noise. While there are many journal papers on the subject of turbulent carrier flow, electron gas viscosity, and the application of the Navier-Stokes equations to carrier flow, this work is unique. Technically, the challenge is to numerically discover the turbulent spectrum of electron gas flow down to the limits of continuity using the instantaneous Navier-Stokes equations. Specifically, the turbulent spectrum of current fluctuations is sought after for a range of currents, represented by Reynolds Numbers (Re), and over the full spectrum of frequencies down to 1 Hz. Statistically-significant averages of the current fluctuation magnitudes at 1 Hz are required. Computationally, the challenge is the combination of the small time-steps required to resolve the smallest turbulent structures, and the long integration periods required to conduct numerical experiments that will capture the current fluctuations across the spectrum, down to 1 Hz with statistical significance. This involves Direct Numerical Simulations (DNS) and Large-Eddy Simulations (LES) to capture the range from the Kolmogorov frequency to 1Hz. Another computational challenge is the visualization and data mining of raw data spanning hundreds of millions of grid points over millions of time-steps.


Numerical results are being compared to measurements on micro-channel (NVESD) and nano-channel (UCSB) carrier transport test structures. Numerical results are valued for the transient 3D visualization of the turbulent structures of interest, and the unobstructed measurement of the transient current flow, spatially-averaged over sampling planes. The numerical experiments cover free jet flows and flows through micro-channels, representing the smallest NVESD pixels. Preliminary computations have yielded excellent agreement with the hypothesized $1/f$ spectrum, as shown in the figures below. Comparison with preliminary measured results has also been encouraging.



Preliminary results



Partially complete production results



Dashed sampling planes along free jet

4.0 Strategic Initiatives

The HPCMP supports the Department of Defense in maintaining and building our strategic advantage over current and emerging adversaries. While near-term military budgets are focused on restoring readiness, following more than a decade of combat, it remains necessary to invest in the future systems that will be needed on the battlefield of tomorrow. The HPCMP is positioned to support the DoD goals of ensuring the continued lethality of the joint forces.

Among the initiatives underway, the HPCMP is increasing its investment in classified and above-secret computing. As the unclassified research in areas such as armor technology, hypersonic flight, and advanced submarine and ship design has matured, the move to design and analyze future systems necessitates the increased classification of the products and test results. This is discussed in greater detail below.


The HPCMP is uniquely positioned to undertake an effort to apply high-end computing to the cybersecurity challenge. With its multiple data streams from the DSRCs and the DREN and SDREN, the HPCMP collects and analyzes significant cyber event data. In 2016, the Program initiated the HPC Architecture for Cyber Situation Awareness (HACSAW) project to better understand what HPC can contribute to the cybersecurity effort.

With the retirement of the Riptide supercomputer at the Maui DSRC, the HPCMP is refocusing efforts at the Maui DSRC to address the urgent need to explore new HPC architectures and algorithms. This will position HPCMP to be able to meet customers' needs in an evolving hardware environment. This effort is referred to as the "Vanguard Center for HPC", and is detailed below.

4.1 Above-Secret Computing

The HPCMP identified an initial set of DoD organizations and users needing above-secret high performance computing capabilities. Above-secret is defined as classification levels ranging from collateral secret with restricted data, up to and including a mix of top secret with special access program (SAP), and sensitive compartmented information (SCI) components.

The initial set of users do not represent all potential Service and Agency above-secret HPC requirements; however, it is expected that interest will increase after an above-secret supercomputing capability is deployed. Consequently, it is anticipated that these requirements will expand as the capability is matured. It should be noted that not all above-secret HPC requirements can be met with shared resources, due to technical requirements such as a need for specific architectures, and/or security-based policy issues.



Multiple service delivery options were considered to attempt to meet the Services' and Agencies' above-secret HPC requirements. The options assessed were Dedicated Processing, Shared Period Processing, and Multi-Level Security (MLS) Shared Processing. Dedicated processing means that an individual HPC system is provided for a specific project for the duration of the project or life of the system. Shared period processing allocates a machine to one project for a limited period of time, typically on the order of weeks to a few months. The machine is then sanitized and turned over to another project for a limited time period, and so on. The third option is Multi-Level Security Shared Processing. With MLS, multiple jobs simultaneously run at potentially different levels of security (TS, TS/SCI, SAP, etc.) with specialized software to maintain data segregation. This option poses the most risk from a security policy perspective.

An in-depth study was conducted to articulate the DoD requirements for shared above-secret HPC resources, and provide a comprehensive assessment of trade-offs, costs, and impacts of incorporating shared above-secret HPC resources and services into the portfolio of capabilities the HPCMP acquires, provisions, and delivers to the DoD research, development, testing and evaluation (RDT&E), and acquisition engineering communities. Based on the Services' and Agencies' above-secret HPC requirements, along with technical, budget, and policy factors, the following recommendations were made:

Policy and Studies

- Implement shared HPC above-secret capabilities starting with Technology Insertion 17 (TI-17) acquisition to meet known user requirements.
- Promote high performance network connectivity solutions for above-secret traffic, with the DREN being the preferred option. The DREN backbone is approved to transport appropriately encrypted above-secret traffic in a closed community of interest between any number of sites, i.e., two or more. In instances where DREN is not feasible, for technical or other reasons, alternate network connectivity solutions will be considered.
- Commission a Multi-Level Security (MLS) study to determine efficacy of an MLS HPC system deployed with the intent to be shared across DoD Services and Agencies. MLS has the potential to be more cost-effective and responsive in addressing above-secret requirements, as compared to using dedicated systems or periods-processing approaches alone.

Hardware

- Purchase three above-secret systems for the AFRL DSRC as part of the TI-17 procurement process. Recommend two 6,000 core systems, using periods-processing, and one 10,000 core system running as a shared Top Secret/Sensitive Compartmented Information (TS/SCI) system.
- As part of the TI-17 procurement, a large unclassified HPC system should be purchased that has the same architecture as the above-secret systems, to enable above-secret applications to be tested at the unclassified level to the maximum extent possible.
- To meet growing requirements, ARL DSRC will expand Secret/RD/CNWDI capabilities by extending the lifecycle and repurposing the existing collateral secret resources (TI-12) to support Secret/RD/CNWDI applications.
- In parallel with the study, a small-scale MLS system should be designed and deployed as a "proof-of-concept" demonstration system for the DoD HPCMP in order to investigate applicability and acceptance in the HPCMP community, and viability of achieving policy/security agreement reciprocity across Services and Agencies.

The HPCMP is implementing these recommendations, with the exception of the MLS proof-of-concept. This may be pursued in the future.

This new effort will put the HPCMP on a trajectory to support a substantial portion of the known near-term above-secret HPC requirements starting in FY18, and position the HPCMP to effectively support Service and Agency future requirements.

4.2 HPC Architecture for Cyber Situational Awareness (HACSAW)

In December 2015, the Executive Steering Group (ESG) conducted their annual oversight meeting of the HPCMP. As a result of this meeting, the HPCMP was tasked with developing an initiative to leverage HPC to advance emerging challenges for cyber Situational Awareness (SA). Executing this task to "explore the applicability of HPC resources to cyber situational awareness," the HPCMP conducted a

formal review of applicable capabilities, technologies, and future trends. Based on the analysis, the HPCMP identified significant potential benefits from employing resources to advance the Department's cyber situational awareness capabilities. The rapidly increasing volume of structured and unstructured cyber data sets, the infusion of HPC resources and big data technologies, such as machine learning, data mining, and data analytics, together present potentially game-changing means to radically enhance cyber situational awareness within the HPCMP.

At the October 2016 annual ESG meeting, the HPCMP received endorsement for its plan and was directed to proceed with a fast-track effort. This effort, HPC Architecture for Cyber Situational Awareness or HACSAW, is intended to drastically increase the HPCMP's current and predictive understanding of cyberspace on the DREN, which entails the collection of unclassified data sources from the edge of the network, such as Internet Access Points (IAP), down to the host level. Through the application of HPC resources, HACSAW will explore novel and innovative analytical capabilities based on a comprehensive data set.

The integration of HPC within the cyber workflow provides an opportunity for fusion and assessments of disparate data streams and real-time analysis using data science algorithms and machine learning (both structured and unstructured data). This approach also provides automated, dynamic response mechanisms to significantly reduce the response time to threats (days to minutes).

The HACSAW platform has a formal Authority to Operate (ATO) and is fully accredited under the Risk Management Framework (RMF).

Building blocks of cyber SA can be decomposed into a variety of activities that are assigned to organizations as Mission-Essential Tasks (METs). METs build on each other to ensure that decision makers have the understanding necessary to make effective decisions. Researchers will develop analytics that align with one or more MET(s). There are six MET categories: Monitoring, Detection, Alerting, Cyber Threat Analysis, Cyber Risk and Event Analysis, and Share/Collaborate.

As part of this effort, the HPCMP issued a Call for Proposals to targeted collaborators, which solicits contributions from DoD, federally-funded research and development centers, and academia in the areas of cyber SA and HPC. The collaborator efforts over the course of the next year are intended to yield results that demonstrate potential integration into the cyber SA operational environment and aligns with defined METs.

4.3 Technology Innovation: Maui Vanguard Center for HPC

The Maui High Performance Computing Center (MHPCC) DSRC, established in 1993, is an Air Force Research Laboratory (AFRL) center managed by the University of Hawaii under contract to the Directed Energy Directorate at Kirtland Air Force Base in New Mexico. After serving the HPCMP as an all-purpose shared resource for 23 years, the MHPCC DSRC is transitioning to a Vanguard Center for High Performance Computing. The Vanguard Center's mission is to evaluate and optimize early production HPC technology and provide breakthrough software solutions to significantly increase the breadth and impact of the HPCMP. Maui received a new unclassified system with 2-3 times the computation power of the machine it is replacing. The new machine has a more energy-efficient design. Maui is also fielding two identical shared memory system computers for classified computing (secret and above-secret). The classified systems will support data-intensive computing and special applications related to the US Pacific Command area of responsibility.

Maui's Vanguard Center mission will be accomplished via a cadre of experts immersed in a culture of innovation and collaboration, working together as HPCMP pathfinders to explore, evaluate, and apply emerging HPC technology, develop HPC-backed solutions for non-traditional users, and enhance the

user productivity through a modern HPC ecosystem. Initial hardware investments include an IBM Power 8+ supercomputer and two SGI UV300 supercomputers.

4.4 Looking to the Future

With its Frontier projects, CREATE high-fidelity physics and digital engineering tools, and advanced software and computing platforms, the HPCMP is poised to provide the computational power to drive future weapon system innovations. The programs and technology development being supported include:

Hypersonic Systems: The HPCMP has dedicated over 400 million core-hours to support advanced hypersonic research in support of Army and Air Force Research Lab projects. Additionally, the CREATE-AV tools are being improved to enable assessment of hypersonic weapons and platforms. Discussed in greater detail in Section 3, hypersonic research is fundamental to support the development of both US and allied weapons for 2020 and beyond, as well as developing the knowledge necessary to effectively counter the weapon systems being developed by potential competitors. China has openly discussed its efforts and investments in supercomputing to underpin hypersonic research.



Hypersonic Strike Weapon

The HPCMP is currently supporting research on hypersonic propulsion systems, airflow over and through potential hypersonic platforms, and computational assessments of hypersonic test events. This research is directly contributing to the Services' efforts to field hypersonic weapons platforms. The planned investments in above-secret computing are necessary to support the transition of hypersonic platforms from the research phase to the prototyping and deployment of actual weapon systems. The complexities involved in these calculations require the most advanced supercomputers in the HPCMP ecosystem to reach solutions in a timely manner.

Autonomous Systems: Autonomous sensor and shooter platforms offer future joint force commanders combat multipliers. HPC offers decision makers greater and more timely insights into the physical performance capabilities of proposed unmanned vehicles and the potential operational environments into which they will be employed.

Digital prototyping allows for high-fidelity simulation of the platforms and sensors proposed for future autonomous vehicles. From these high-fidelity models, multiple reduced-order models can be run simultaneously to model distributed platform operations for air, land, and surface/subsurface platforms.



HPC-driven artificial intelligence modeling may provide additional insights for weapon systems designers to enable the development of complex communication and control architectures necessary for the employment of distributed platforms operating under teaming concepts of manned and unmanned platforms.



Future Vertical lift TDs

Future Vertical Lift: Frontier and CREATE-AV Helios software directly supports the Joint Multi-Role Technology Demonstration (JMR-TD) project. The CREATE-AV Helios software was used in the project demonstrator down-select, and is continuing to be used in the development and assessment of test vehicles.

One of the FY17 Frontier projects is dedicating 100 million core-hours to support the developmental testing of the JMR-TD platforms. With the software and computing power provided by the HPCMP, Army engineers are able to assess contractor designs and work to promptly identify and assist in the mitigation of potential test program problems. As more powerful engines are being developed and proposed to offset the increases in platform weight and evolving operational employment patterns of the fleet of legacy helicopters, the Improved Turbine Engine Program (ITEP) is using CREATE-AV Helios and Kestrel tools to analyze integration of advanced engines into current and future rotary-wing aircraft. These calculations will ensure that proposed changes to main engines do not adversely affect the flight characteristics of current airframes.

Advanced Electromagnetic Design and Directed-Energy Systems: Design and analysis of the next generations of photonics, directed-energy, and electromagnetic pulse protection systems depend on the advanced computing and algorithm development currently funded by the HPCMP.



Directed Energy

The installation of lasers or other directed energy systems on current and future airframes and surface platforms offers significant advances in combat effectiveness. To achieve this goal, it is necessary to understand not only the waveforms and energy generation problems, but how the dynamic motion of the platform impacts the ability to form and sustain the energy beams. The coupled physics problems of fluid and structural dynamics, plasma formation, and environmental effects require HPC to achieve timely and accurate solutions.

These and other projects in space, rocket propulsion, ship design, and service life prediction depend on advanced computing to both design and digitally test future weapons systems. The need for more rapid innovation and fielding of systems points to new manufacturing and testing paradigms incorporating digital design and evaluation to reduce the acquisition timeline and more effectively incorporate innovation. The complexity of the interactions in hypersonics, hydrodynamics, and electromagnetics require powerful supercomputers to achieve results in a timely manner. The HPCMP is working now to make innovation practical.

Acronyms

ACTD – advanced concept technology demonstration

AFIT – Air Force Institute of Technology

AFOSR – Air Force Office of Scientific Research

AFRL – Air Force Research Laboratory

AFRL-RI – Air Force Research Laboratory, Information Directorate

AI – artificial intelligence

AIR – automated impacts routing

AMRMC – Army Medical Research and Materiel Command

AO – authorization official

APG – Aberdeen Proving Ground

ARC – Affiliated Resource Center

ARL – Army Research Laboratory

ASA – Assistant Secretary of the Army

ASA(ALT) – Assistant Secretary of the Army for Acquisition, Logistics and Technology

ATC – Aberdeen Test Center

ATEC – Army Test and Evaluation Command

ATO – authority to operate

BD – big data

CAD – computer-aided design

CAN – Controller Area Network

CBRN – chemical, biological, radiological and nuclear

CEDAR – cybersecurity environment for detection, analysis, and reporting

CEM – computational electromagnetics

CENIC – Corporation for Education Network Initiatives in California

CERDEC – US Army Communications-Electronics Research, Development and Engineering Center

CFD – computational fluid dynamics

CIO – chief information officer

CMIP5 – Fifth Coupled Model Inter-comparison Project

CNS – computational numeric simulation

CNWDI – critical nuclear weapons design information

COIL – chemical oxygen-iodine lasers

CONUS – continental United States

COR – contracting officer's representative
CORAL – Collaboration of Oak Ridge, Argonne, and Livermore program
CREATE – Computational Research and Engineering Acquisition Tools and Environments
CSA – comprehensive security assessment
CSEP – cyber security enhancement project
CSSP – cybersecurity service provider
CTA – computational technology area
CTAAP – Computational Technology Area Advisory Panel
CTV – controlled test vehicle

DA – data analytics
DAAC – Data Analysis and Assessment Center
DARPA – Defense Advanced Research Projects Agency
DARTS – Data Analysis of Real-time Streams
DASA – Deputy Assistant Secretary of the Army
DASA(ALT) – Deputy Assistant Secretary of the Army for Acquisition, Logistics and Technology
DASA(R&T) – Deputy Assistant Secretary of the Army for Research and Technology
DASAF – Deputy Assistant Secretary of the Air Force
DASAF(STE) – Deputy Assistant Secretary of the Airforce for Systems, Test and Evaluation
DASD – Deputy Assistant Secretary of Defense
DASD(C³ & Cyber), Deputy Assistant Secretary of Defense for Command, Control, Communications and Cybersecurity
DASD(C3CB) – Deputy Assistant Secretary of Defense for Command, Control, Communications, Cybersecurity and Business Systems
DASN – Deputy Assistant Secretary of the Navy
DASN(RDTE) – Deputy Assistant Secretary of the Navy for Research, Development, Testing and Evaluation
DCGS-A – distributed common ground system - Army
DCGS-Wx – distributed common ground system - weather
DDoS – distributed denial-of-service
DE – directed energy
DHPI – dedicated HPC project investments
DIACAP – DoD Information Assurance Certification and Accreditation Process
DJS – DREN joint sensors
DNS – direct numerical simulations
DoD – Department of Defense
DOE – Department of Energy
DoF – degrees-of-freedom

DPAL – diode-pumped alkali laser
DREN – Defense Research and Engineering Network
DSAWG – DoD Security/Cybersecurity Authorization Working Group
DSMC – Direct Simulation Monte Carlo
DSRC – DoD Supercomputing Resource Center
DTRA – Defense Threat Reduction Agency

eMASS – Enterprise Mission Assurance Support Service
EM – energetic material
EOC – early operational capability
ERDC – USACE Engineer Research and Development Center
ERS – Engineered Resilient Systems
ESG – Executive Steering Group
ESnet – Energy Sciences Network
ESPC – Earth System Prediction Capability
ExaFLOPS - 10^{18} floating-point operations per second

FEA – finite element analysis
FEMP – finite element modeling and post-processing
FTTS – future tactical truck system
FY – fiscal year

GB – gigabytes (10^9 bytes)
Gbps – gigabits per second (10^9 bits per second)
GCM – Galaxy Coupling Manager
GIG – global information grid
GPGPU – general-purpose graphics processing unit
GPU – graphics processing unit
GV – ground vehicles

HACSAW – HPC architecture for cyber situational awareness
HASI – HPCMP Applications Software Initiative
HELEEOS – high energy laser end-to-end operational simulation
HIP – HPC Internship Program
HMMWV – high mobility multi-purpose wheeled vehicle
HPC – high performance computing
HPCAP – HPC Advisory Panel

HPCMP – High Performance Computing Modernization Program
HPDA – high performance data analytics
HPEM – high power electromagnetic
HSAI – High Performance Computing Software Applications Institute
HWSSE –HPCMP-wide security services enclave

IAP – internet access point
IARPC – Interagency Arctic Research Policy Committee
IDS – intrusion detection software
IFC – integrated functional capability
IHDE – Integrated Hydrodynamics Design Environment
ILES – implicit large-eddy simulation
IMATS – integrated military analysis information technology suite
INP – innovative naval prototype
I/O – input/output
IP – internet protocol
ISCM – information security continuous monitoring
ISSM – information systems security manager
ISSO – information system security officer
IT – information technology
ITEP – Improved Turbine Engine Program

JLTV – Joint Light Tactical Vehicle
JMETC – Joint Mission Environment Test Capability
JMR-TD – Joint Multi-role Technology Demonstrator
JSF – joint strike fighter

KNL – knights landing processor

LEAPS – Leading-edge Architecture for Prototyping Systems
LES – large-eddy simulations
LRV – Light Reconnaissance Vehicle
LV – launch vehicle

M&S – modeling and simulation
MAT – Mobility Analysis Tool
MAX GigaPOP – Mid-Atlantic Crossroads (MAX) Gigabit Point of Presence

Mbps – megabits per second (10^6 bits per second)
MDA – Missile Defense Agency
MET – mission-essential task
MHPCC – Maui High Performance Computing Center
MLS – multi-level security
MM&S – multiscale modeling and simulation
MOE – measure of effectiveness
MOP – massive ordnance penetrator
MPRF Maritime Patrol and Reconnaissance Force
MRBM – medium-range ballistic missile

NACI – National Agency Check with Inquiries
NASA – National Aeronautics and Space Administration
NATO – North Atlantic Treaty Organization
NESM – Navy Enhanced Sierra Mechanics
NFV – network function virtualization
NIF-T – NAVAIR Propulsion Internal Flow Modeling Team
NIPRNet – Non-secure Internet Protocol Router Network
NOAA – National Oceanic and Atmospheric Administration
NPS – Naval Pos
NREN – NASA Research and Engineering Network
NRL – Naval Research Laboratory
NRL MRY – Naval Research Laboratory Monterey
NSA – National Security Agency
NURBS – Non-Uniform Rational B-Splines
NVESD – Night Vision and Electronic Sensors Directorate

O&M – operations and maintenance
OARnet – Ohio Academic Resources Network
OASD – Office of the Assistant Secretary of Defense
OASD A – Office of the Assistant Secretary of Defense for Acquisition
OASD(R&E) – Office of the Assistant Secretary of Defense for Research and Evaluation
OLCF – Oak Ridge Leadership Computing Facility
ONR – Office of Naval Research
OpenMP – open multi-processing
ORS – Open Research System
OSD – Office of the Secretary of Defense

OSTP – Office of Science and Technology Policy

PE - polyethylene

petaBytes - 10^{15} bytes

PETA – Performance Enhancement Thrust Area

PETTT – Productivity Enhancement, Technology Transfer and Training

POA&M – plan of action and milestones

PoR – Program of Record

PF - petaFLOPS (10^{15} floating-point operations per second)

PM – program manager

PMRF – Pacific Missile Range Facility

PUWG – Privileged User Working Group

QUBE – quantifying uncertainty in the battlefield environment

R&T- research and technology

RA – reactionary assistance

RADIX – Rapid Audit of Unix

RANS – Reynolds-average Navier-Stokes

RANS/LES – Reynolds-averaged Navier-Stokes/Large-Eddy Simulation

RASM – Regional Arctic System Model

RD – restricted data

RDT&E – research, development, test and evaluation

RDX – cyclotrimethylene trinitramine

Re – Reynolds number

RF – radio frequency

RMF – risk management framework

RSDE – Rapid Ship Design Environment

SA – situational awareness

SAR – synthetic aperture radar

SDB – small diameter bomb

SDN – software-defined networking

SDP – service delivery point

SDREN – Secret DREN

Secret/SAP – Secret/Special Access Program

Secret/SAR – Secret/Special Access Required

SES – Senior Executive Service
SIPN – Sea Ice Prediction Network
SM – standard missile
SMDC – Space and Missile Defense Command
SPAWAR – Space and Naval Warfare Systems Command
SPH – smoothed-particle hydrodynamics
SSC-SD – Space and Naval Warfare Systems Center, San Diego
SSL – secure sockets layer
STEM – science, technology, engineering, and mathematics
S&T – science and technology
SV – space vehicle

T&E – test and evaluation
TARDEC – US Army Tank-Automotive Research, Development and Engineering Center
teraFLOPS - 10^{12} floating-point operations per second
TFA – Technical Focus Area
TI – technology insertion
TSRM – third-stage rocket motor
TS/SCI – Top Secret/Sensitive Compartmented Information

UAG – User Advisory Group
UAS – unmanned aerial system
UAV – unmanned aerial vehicle
UCSB – University of California – Santa Barbara
USACE – US Army Corps of Engineers
USAFA – US Air Force Academy
USCYBERCOM – US Cyber Command
USD AT&L – Undersecretary of Defense for Acquisition, Technology, and Logistics
USMA – US Military Academy
USNA – US Naval Academy

VoF – volume of fluid

WAN – wide-area network
WMD – weapons of mass destruction
WSMR – White Sands Missile Range

DOD
HPC

MODERNIZATION PROGRAM