



ARL-CR-0861 • JUNE 2021



The History of the Department of Defense High-Performance Computing Modernization Program

by J Michael Barton

Parsons

6210 Guardian Gateway, Suite 200

Aberdeen Proving Ground, MD 21005

under contract W911QX-18-D-0001

Approved for public release; distribution is unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



The History of the Department of Defense High-Performance Computing Modernization Program

by J Michael Barton

Parsons

6210 Guardian Gateway, Suite 200

Aberdeen Proving Ground, MD

under contract W911QX-18-D-0001

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188		
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) June 2021		2. REPORT TYPE Contractor Report		3. DATES COVERED (From - To) January 2020–March 2021	
4. TITLE AND SUBTITLE The History of the Department of Defense High-Performance Computing Modernization Program			5a. CONTRACT NUMBER W911QX-18-D-0001		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) J Michael Barton			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Parsons 6210 Guardian Gateway, Suite 200 Aberdeen Proving Ground, MD 21005			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) DEVCOM Army Research Laboratory ATTN: FCDD-RLC-S Aberdeen Proving Ground, MD 21005			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) ARL-CR-0861		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Computing has changed science and engineering since the introduction of the first general-purpose scientific computer in 1946, the Electronic Numerical Integrator and Computer (ENIAC), produced by the University of Pennsylvania Moore School of Electrical Engineering for the Army. The ENIAC was created for studying the science of ballistics and for the development, testing, and fielding of Army ordnance and creation of associated firing tables. Since that time, computing has provided access to larger-parameter spaces for design, analyzed more cases in a given unit of time or budget, provided access to phenomena not measurable or to conditions not accessible to experimentation, completed production and analysis of ever-larger data sets, and even contributed to mathematical proofs (e.g., the four-color theorem). In turn, science and engineering continue to change computing, creating demand for ever-larger and faster platforms and illuminating new phenomena that lead to new materials, components, fabrication techniques, architectures, and entirely new types of computing. We describe the Department of Defense High-Performance Computing Modernization Program, the environment in which it was created, the people who helped bring it into existence, and the impact it has had on science and engineering, problem solving, and ultimately, on our primary customer, the Warfighter.					
15. SUBJECT TERMS high-performance computing, history of computing, modeling and simulation, ENIAC, Linux, supercomputers, HPCMP					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 64	19a. NAME OF RESPONSIBLE PERSON J Michael Barton
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code) 410-278-3722

Contents

List of Figures	v
Acknowledgments	vi
1. Introduction	1
2. The World of HPC	1
2.1 Setting the Stage	1
2.2 Identifying the Players	3
2.3 Assembling the Team	3
2.4 Action	5
3. The Road to the HPCMP	7
4. Adventures in Computing	11
4.1 The Modernization Plan	11
4.2 The Implementation Plan	12
4.3 Organizing the Program	14
4.4 Beginning Operation	15
4.5 Test and Evaluation Join the Team	19
5. The Network	22
6. PET, PETT, PETTT, or What?	26
7. Program Management	28
8. Impact	34
8.1 Space Shuttle Columbia Accident Investigation	34
8.2 Army CH-47 Advanced Chinook Rotor Blade	35
8.3 Aircraft Stores Carriage and Certification	36
8.4 Armor	37

9. Modernizing the HPCMP	38
10. Conclusions	40
11. References	42
List of Symbols, Abbreviations, and Acronyms	50
Distribution List	54

List of Figures

Fig. 1	Structure of the HPCWG in April 1993 ⁵⁴	12
Fig. 2	Time line for the creation of the HPCMP	21
Fig. 3	The I-DREN circa late calendar 1994.....	23
Fig. 4	Key events leading to the DREN	25
Fig. 5	Appreciation from theater	38

Acknowledgments

Rather than saving the best for last, I want to recognize and thank Dr Larry Davis first. If you look up the definitions of patient and untiring, you may find Larry's image. Through a *mélange* of in-person meetings, telephone calls, video calls, and e-mail exchanges, Larry has provided a true wealth of information, careful reviews with attention to detail, historical documents, advice, meeting notes, Figure 1, and encouragement to complete this work. He is a true colleague and friend and much of the value of this review derives from his support. It is my privilege to have worked with Larry since I joined the High-Performance Computing Modernization Program (HPCMP), via Army test and evaluation, in late 2001. Ms Cathy McDonald was the HPCMP requirements maven at that time. She is one of the most gracious people I have ever known, and any time I needed anything even remotely related to requirements, Cathy was helpful. She was there before the beginning, along with Larry, and knows the story. It was a joy to reconnect with her for this history.

Mr Robert Reschly provided extensive background on how the Defense Research and Engineering Network (DREN) came into existence and the personalities that made it happen, forming the basis for the material in Section 4.6, and he provided Fig. 3. Bob seems to have instant recall as soon as I mention a date, person, or incident and has been a delight to interview. He also has a true eagle eye when it comes to reviewing and editing. Mr Rodger Johnson described I-DREN and DREN developments, and how the high-speed encryption devices came to be. I thoroughly enjoyed sitting on the DREN Technical Advisory Panel while Rodger was Program Manager (PM) DREN. Dr Roy Campbell succeeded Rodger as PM DREN. Roy has true passion for anything he does and he is fun to work with. Though a US Army Combat Capabilities Development Command Army Research Laboratory, now DEVCOM ARL, employee for many years, he is detailed to the HPCMP, currently as its Chief Strategist. Every conversation with Roy is a learning experience, especially if we are talking about emerging technology, and every time I wonder where he gets all of his energy.

Dr Aram Kevorkian provided background and documents on Navy High-Performance Computing (HPC), on early organization of the HPCMP, and insight into the first Major Shared Resource Center (MSRC) source selections. Mr Cray Henry educated me on early functions of the program, the annual budget process, congressional and DOD oversight, the programmatic path to Computational Research and Engineering Acquisition Tools and Environments (CREATE), and activities that led to transitioning of the program to the Army. I had the good pleasure to sit on the HPC Advisory Panel (HPCAP) for 10 years while Cray was

the Program Director and Larry was Deputy Director. Mr John West educated me on activities that led to transitioning of the program to the Army and execution of the program's joint service mission by the Army. While John was director, I worked with him in HPCMP outreach, especially to the DOD test and evaluation community. Dr Valerie Miller provided background on the Major Automated Information System (MAIS) and Major Defense Acquisition Program (MDAP) oversight processes and documentation, their criticality to the success of the HPCMP, and insight into the program matriarchs and patriarchs. She also helped me understand the importance of the Service/Agency Approval Authorities (S/AAAs) and power users throughout the program. Ms Deborah Schwartz corrected my misunderstanding of the Return on Investment (ROI) study and provided the material and background on the workforce development efforts. She was key to both while working in the HPCMP.

Dr Doug Post schooled me on CREATE, providing background on the motivation, environment, prime movers, and value. Dr Edward Kraft shared his passion for and experience in integrated test and evaluation, HPC in test and evaluation (T&E), digital engineering, stores certification, material for Section 8.3 and the origins of CREATE. Ed is a true visionary and one of the prime movers behind digital engineering. Dr Clifford Rhoades walked me through the early years of HPC at Kirtland Air Force Base, Air Force involvement in early supercomputing, and the origin of the Maui High-Performance Computing Center (MHPCC). Cliff is meticulous in his attention to detail and has extensive knowledge of HPC.

Mr Harold Breaux provided detailed recollections on the background of the Army in HPC and the HPCMP, and notes from meetings with Army leadership and with Dr Jones. Harold made it happen for DEVCOM Army Research Laboratory and led the way for the Army and the HPCMP. Mr Charles Nietubicz gave me background on Ballistic Research Laboratory (BRL) activities that led up to the ARL MSRC proposal. Charlie was there helping make it happen, and then became its first director. Mr Robert Sheroke offered up his wealth of historical documents and personal knowledge from BRL and ARL and the formative days of the HPCMP. Mr Matt Goss and Ms Lee Ann Brainard, Director and Deputy Director, respectively, of the ARL DOD Supercomputing Resource Center (DSRC), have been tireless in their support, encouragement, and patience. Lee Ann helped me understand BRL, and later ARL, operations during the time of delivery of the Cray 2 and KSR-1. I have a delightful photo of four young computer scientists proudly gathered around a newly delivered Cray 2 in 1987, one of whom might be Lee Ann, but she will not let me use the photo.

Ms Virginia To is a colleague and friend of longstanding with a 25-year history at ARL. She offered invaluable insight into the workings of ARL, personalities involved, and the role of High Performance Technologies Inc. (HPTi) in the Programming Environments and Training (PET) contracts. Ginny managed the HPTi PET contract at ARL. Dr Bharat Soni and I began working together in 1981 at Arnold Engineering Development Center (AEDC), and he is the one who introduced me to Army T&E in 2001, when I became an Intergovernmental Personnel Act (IPA). Bharat was by that time at Mississippi State University (MSU) and part of PET 1 and PET 2. I have availed myself of his knowledge repeatedly, especially concerning the MOS Consortium. Professors Stan Ahalt and Ashok Krishnamurthy were at The Ohio State University (OSU) when I began working with the HPCMP and I was an advisor to PET. They helped me understand the Ohio Supercomputer Center (OSC) role in PET 2, the value of PET training, and the signal and image processing (SIP) function. They also provided colorful anecdotes.

Dr Robert Lucas has extensive knowledge and experience with Defense Advanced Research Projects Agency (DARPA), DOD, and Department of Energy (DOE) HPC programs. I met him in the program and his insights, knowledge, and encouragement to undertake this task have been invaluable. Mr Tom Kendall has helped me navigate the stovepipes, personalities, and technology of the HPCMP. I met Tom when I was on the DREN Technical Advisory Panel about 15 years ago and for the past 5+ years have had an office next to his at ARL. We still have adjacent offices, but they are mostly unvisited. Mr Jeff Graham spent 2 h with me prior to his retirement and later exchanged extensive e-mail relating how the Air Force Research Laboratory (AFRL) DSRC came to be and educated me on Air Force involvement in the HPCMP. I am not sure any DSRC director has had a greater heart for the users than Jeff did.

Dr Robert Meakin graciously provided background material that constitutes much of Section 8.1, and carefully reviewed my summary of his work. Mr Jeff Highland provided background on Army T&E involvement in the HPCMP and early distributed centers awarded to the test centers. I worked with Jeff in the Army Developmental Test Command when I first became involved with the HPCMP. Jeff's early guidance and advice greatly smoothed my introduction to the HPCMP. Dr Andrew Wissink provided material for Section 8.2. Mr David Kleponis provided material for Section 8.4. Mr Brian Simmonds produced all of the final graphics and located many historical documents for me.

Thanks to every one of you, and to those I have inadvertently omitted. I appreciate your help more than I am able to express.

1. Introduction

The Department of Defense (DOD) High-Performance Computing (HPC) Modernization Program (HPCMP) was created on 5 December 1991 when President George H W Bush signed the National Defense Authorization Act for Fiscal Years 1992 and 1993.¹ It was established under the Office of the Director, Defense Research and Engineering and in 2011 was transferred to the US Army, which assigned it to the Engineer Research and Development Center (ERDC) in Vicksburg, Mississippi.

After nearly 30 years, the program still executes its mission to modernize the supercomputer capability of DOD laboratories and test centers with large-scale computing resources; high-bandwidth, low-latency networks; commercial and government-owned software as well as open source tools; and resident subject matter expertise and training in the above. Demand for high-performance computers, advanced networks, scalable software, and algorithmic and programming expertise continues to increase monotonically.

The program was born amid rapid technological change nationally, during a time when many organizations—government, industry and university—were advocating for widespread access to leading edge computers as a requirement for the advancement of scientific discovery, corporate innovation, American international competitiveness, and for education in science and engineering. It has adapted to changing requirements, responded to national emergencies, and enabled current and future generations of DOD scientists and engineers.

2. The World of HPC

2.1 Setting the Stage

Before we extol the contributions of the HPCMP to DOD science and engineering, let us look at the environment in which it arose. The legislation was passed in December 1991. That same month the High-Performance Computing Act of 1991² was passed, establishing the High-Performance Computing and Communications Initiative (HPCCI). The Act was known as the Gore Bill, after its author and champion, then Tennessee Senator Al Gore. Its purpose was to “ensure the continued leadership of the United States in high-performance computing and its applications. . . .”² It established a National High-Performance Computing Program; directed the President to establish a High-Performance Computing

Advisory Committee; mandated the creation of a one gigabit per second or greater National Research and Education Network (NREN); and established Grand Challenges as driving applications.

Grand Challenges were fundamental problems in science and engineering with broad application whose solutions required HPC resources.^{3,4} A criticism of the HPCCI (later renamed the Network and Information Technology Research and Development [NITRD] Program) was that Grand Challenges were focused on science and engineering but with minimal impact on the average citizen. In response, the initiative was extended to include National Challenges addressing critical needs of society.⁵

The High-Performance Computing Act directed the Defense Advanced Research Projects Agency (DARPA)* to conduct research and development of advanced fiber optics technology, switches, and protocols needed to develop the NREN. The National Science Foundation (NSF) was made responsible for pulling the NREN together based on previous NSF-developed network capabilities (NSFNET) and agreements with DARPA. There is a host of other provisions in the Act, with citations for the National Aeronautics and Space Administration (NASA), the Department of Commerce, Department of Energy (DOE), and other agencies, but the above are key for our purpose.

This was a truly foundational piece of legislation. It was extended and modified several times. For example, Public Law 105-305, Next Generation Internet Research Act of 1998,⁶ aimed to remove economic and geographic impediments for network access. Public Law 110-69, America COMPETES (Creating Opportunities to Meaningfully Promote Excellence in Technology, Education, and Science) Act of 2007⁷ removed the NREN section (since the Internet was commercial at this point); reinforced developments in networks, security, software, computational science and engineering, re-emphasized Grand Challenges, and included funding for science, technology, engineering and mathematics. Public Law 114-329, Sec. 105, Networking and Information Technology Research and Development Modernization Act of 2016,⁸ expanded the scope by replacing HPC with networking and information technology.

* Note that the Advanced Research Projects Agency (ARPA) became DARPA, which became ARPA, which again became DARPA. Rather than follow the time-appropriate acronym in this article, we always refer to DARPA.

2.2 Identifying the Players

Supercomputing began in the mid-1970s with the introduction of vector computers, though the Control Data Corporation (CDC) 6600, introduced in 1964 and designed by Seymour Cray, is often cited as the first supercomputer. Texas Instruments had their Advanced Scientific Computer,⁹ introduced in 1973 with a total of seven machines built. CDC, with a long history of scientific computers, built the STAR-100¹⁰ and delivered the first of only three machines in 1974. It was superseded by the Cyber 203, then by the Cyber 205, which was spun off to the ETA Corporation as the ETA-10, then reabsorbed back into CDC around 1987 and discontinued, not long before CDC itself ceased to exist.

Cray Research, Inc., founded by Seymour Cray in 1972, delivered its first Cray-1¹¹ in 1976, eventually selling about 80 of them before producing, under different corporate entities, the Cray-2, Y-MP, T3D, XT5, and a host of others, continuing to this day but now as part of Hewlett Packard Enterprise.

Burroughs tried to get into the market with the Burroughs Scientific Processor (BSP),¹² building on experience from producing the ILLIAC IV¹³ parallel computer for the University of Illinois at Urbana-Champaign. The ILLIAC IV eventually went to the NASA Ames Research Center in California and became operational in 1975, seeing extensive use for computational fluid dynamics, and was the beginning of a long history of NASA in supercomputing. But the ILLIAC IV was a one-off computer and the BSP never came to light. In very short order, the vector computer market slimmed down to a single supplier, Cray, at least domestically. For a time, supercomputer was synonymous with Cray computer.

Owning, operating, and maintaining a supercomputer was expensive; consequently, not everyone had one or access to one. The limitation prompted studies, meetings, and legislation. Congress, industry, and academia recognized its implications for education, industry, and national security. Congress took actions to rectify the situation.

2.3 Assembling the Team

The Army, and later the Los Alamos National Laboratory (LANL), Princeton Institute for Advanced Study, DARPA, and the National Bureau of Standards, built custom computers to meet the demands of weapons development for computational science and research to explore new computer architectures, and for other requirements.¹⁴ Custom computers then gave way to commercial purchases as the computer industry matured. In 1976 the Army Ballistic Research Laboratory (BRL) took delivery of a CDC 7600, the same year that LANL accepted Cray-1 serial 1.

Lawrence Livermore National Laboratory received two STAR-100s, with the third going to the NASA Langley Research Center. The Naval Research Laboratory (NRL) took delivery of a Texas Instruments Advanced Scientific Computer. Other government and industry organizations followed, creating pockets of advanced computing capability. The academic community worked with many of these organizations, gaining access to a Cray or other vector computer, but the number was severely limited and depended largely on personal connections.

The NSF¹⁵ published the results of a study in 1981, recognizing computational science as the third pillar of science, complementing experiment and theory, and warned of a crisis in computational physics. The study highlighted the deficiency of NSF funding for computational facilities, for networks to access the facilities, and for ability to manage and visualize large amounts of data. The study cited the importance of access to both state-of-the-art supercomputers as well as regionally distributed complementary computing capacity.

In April of 1982, the NSF/DOD Coordinating Committee requested that a workshop be organized to explore the needs and opportunities in large scale computing. The workshop was sponsored by the NSF and DOD with cooperation of the DOE and NASA. The panel assessed the role of supercomputing in scientific and engineering research; surveyed the current use, availability, and adequacy of supercomputers; and considered near-and long-term needs. The Lax report,¹⁶ named after its chairman, Peter Lax, outlined the results of the workshop and subsequent discussions and contributions. NSF, DOE, and DOD HPC programs and subsequent initiatives all trace their genesis to this report, its findings, and its recommendations.

The major finding was that in the 1970s the US government diminished its support for supercomputers while other countries increased theirs. Two problems stood out: 1) segments of the research and defense communities lacked access to supercomputers (e.g., students were not familiar with their capabilities or use) and 2) the capacity of supercomputers was several orders of magnitude too small for problems of urgency.

Government, industry, and academia were of one accord. The Lax Report contained white papers from a spectrum of experts substantiating the need for supercomputing in research, estimating the required resources for scientific discovery, design, and defense applications, and heralding the risk of letting Japan gain parity with the United States in supercomputers. At this time DARPA was funding research and development for ARPA Network (ARPANET),¹⁷ which connected defense and military sites and a few universities performing DOD work. The NSF created CSNET (Computer Science Research Network)^{18,19} to connect university computer

science departments in an effort to mitigate the growing negative effects for schools without ARPANET access. The NSF and DARPA developed plans and moved out; NASA and the DOE were already moving.

2.4 Action

Prior to the Lax Report, NASA had taken delivery of the ILLIAC IV, funded by DARPA, and had applied it for computational aerodynamics. Subsequent developments led to the Numerical Aerodynamic Simulation Program,²⁰ initially planned as custom-designed computers to achieve gigaFLOPS (10^9 floating-point operations per second) performance, but switching to commercial platforms. Cray 2 serial 2 with four processors was the first computer installed, followed by SGI clusters and other Cray computers. NASA upgraded network connectivity to make the resources more easily accessible and expanded the program scope. With creation of the HPCCI, NASA was immediately involved in Grand Challenges.²¹ NASA Advanced Supercomputing is still flourishing.²²

The LANL was the first user of the Electronic Numerical Integrator and Computer (ENIAC) in 1946, developed for the Army by the University of Pennsylvania Moore School of Electrical Engineering. LANL continued with custom computer designs and later with purchasing commercial computers.²³ DOE has a long tradition of innovation in advanced computing, in hardware, software, visualization, cluster computing, networking, languages, applications software, and much more.²⁴ The Accelerated Strategic Computing Initiative (ASCI) did for the DOE, on a grander scale, what the HPCMP did for the DOD; it gave the labs a common purpose,²⁵ and led to the Office of Advanced Scientific Computing Research within the DOE Office of Science.²⁶ The DOE is now at the leading edge of exascale (10^{18} FLOPS) computing and is pioneering post-exascale research.²⁷

NSF conducted an internal study in 1983, referred to as the Bardon-Curtis Report,²⁸ to review the Lax report,¹⁶ validating its conclusions and developing requirements and a plan. In response, by 1986 the NSF created five supercomputer centers, rolled out the NSFNET (replacing CSNET) to connect researchers to the centers, and created the Office of Advanced Scientific Computing to provide access by researchers to the supercomputers, supercomputer services, and to encourage growth and development of advanced scientific computing.

In 1983, DARPA started its 10-year Strategic Computing Initiative,²⁹ the explicit goal being to “provide the United States with a broad line of machine intelligence technology and demonstrate applications of the technology to critical problems in defense.”³⁰ It sought autonomous vehicles, a Pilot Associate that could understand plain English, and an artificial intelligence-based Battle Management System to

coordinate everything. The lofty goals were not achieved after more than \$1 billion invested, but many technological advancements were, chief among them parallel computing.

Congress took related action in 1986, passing the National Science Foundation Authorization Act for Fiscal Year 1987.³¹ The Act directed the federal Office of Science and Technology Policy (OSTP) to “undertake a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and Federal research facilities in the United States.”³¹ The study was to look forward 15 years; consider file transfers, electronic mail, remote computer access, and communications; and investigate advantages and disadvantages of fiber-optic systems.

The OSTP response is contained in a three-volume report.^{32–34} Volume I contains recommendations on developing computer networks to support research; Volume II contains summaries of February 1987 workshop discussions; and Volume III contains invited white papers from industry, government, and academia. The three recommendations were to 1) interconnect the current Internet system developed by DARPA and the networks supported by agencies for researchers; 2) expand and upgrade existing computer networks with 1.5 million bits per second capabilities; and 3) provide network service to every research institution in the United States with transmission speeds of three billion bits per second. The gigabit network mentioned in 3) was envisioned to be achievable as a staged program of research and development (i.e., by increasing degrees of capability, not all at once).

The director of the OSTP also chaired the Federal Coordinating Council for Science, Engineering and Technology (FCCSET). The OSTP and the FCCSET were advisory bodies to the president on science and technology issues, and in coordinating federal science and technology efforts. Subsequent to the Lax report, the FCCSET created a panel on supercomputers in 1983 and initiated a series of studies. The goals were to ensure that the United States maintained its lead in supercomputing, strived to make supercomputers available to more researchers, particularly in universities, and increased coordination among government funding contributing to the technology base.³⁵

The next year, 1984, the National Research Council³⁶ addressed computer architectures, novel and conventional, looking out 5–10 years. It considered chips and mass storage technologies; the role of special purpose computers; algorithmic and software technology for large-scale parallel computing; communication technology to make supercomputers available to the US research community; efforts abroad that might undercut the dominant US role in HPC; the relationship of supercomputing to the broader field of artificial intelligence; and policy

questions of how to get the US research and industrial communities to move forward in rapid and effective cooperation.

Threats from Japan were itemized in terms of Japanese semiconductor production, on which US supercomputers depended, in their competitive vector computers, and in two new Japanese initiatives: a National Super Speed Computer Project to produce a teraFLOPS (10^{12} FLOPS) machine, and the Fifth Generation Computer Project aimed at artificial intelligence applications. The culmination of US government, academic, and industry pressure was the research and development strategy³⁷ in 1987 and subsequent program³⁸ for HPC in 1989.

Contemporary with these reports was an independent study by the National Research Council³⁹ in 1988 that came to the following conclusions:

- A strong domestic high-performance computer industry is essential for maintaining US leadership in critical national security areas and in broad sectors of the civilian economy.
- Research progress and technology transfer in software and applications must keep pace with advances in computing architecture and microelectronics to more fully exploit parallel systems.
- US government, industry, and universities should coordinate research and development for a research network to provide a distributed computing capability that links the government, industry, and higher education communities.
- Long-term support for basic research in computer science should be increased within available resources. Industry, universities, and government should work together to improve the training and utilization of personnel to expand the base of research and development in computational science and technology.

Congress followed in 1988 with its own study,⁴⁰ reaching harmonious conclusions. The culmination of all the activity was the High-Performance Computing Act of 1991, and creation of the High-Performance Computing Modernization Program by the National Defense Authorization Act for Fiscal Years 1992 and 1993.¹

3. The Road to the HPCMP

The services met their computational requirements in several ways prior to establishment of the HPCMP. The Army BRL designed custom computers beginning with the ENIAC in 1946 and continued into the 1960s. The last in-house designed computer was the all-digital BRL Electronic Scientific Computer II

(BRLESC II), which entered service in 1967 and was retired in 1978. The first commercial computer purchased by BRL was a CDC 7600 in 1976, followed by a Heterogeneous Element Processor (HEP) in 1983 (one of the first parallel computers), a Cray X-MP, a Cray 2, and others.⁴¹

Harold Breaux was chief of the BRL Advanced Computer Projects Branch, which later became the High-Performance Computing Division. Harold's team included Tony Pressley, Bob Reschly, Phil Dykstra, Charles Nietubicz, Rodger Johnson, Valerie Miller (nee Thomas), and others you will meet later in this report. The Army joined the ARPANET team, working on protocols and adapting the ideas to create campus area networks. As the BRL effort unfolded, the Army deemed the objective essential for other key Army activities.

Bob Reschly recalls one experience that reinforces the adage "Be careful what you ask for."⁴² In late 1984, BRL advocated to leadership for a supercomputer, specifically a Cray 2, for delivery in 1987. The commander of the Army Materiel Command became so convinced of the need that he decided the Army could not wait on a Cray 2 and needed an interim solution. A three-phase plan was approved in April 1985 to purchase commercial supercomputer time, then purchase a Cray-X-MP, then a Cray 2. A Cray X-MP 4/8 was delivered in 1986, actually taking funds away from the Cray 2 effort, but BRL did eventually get the Cray 2 in 1987.⁴³

In October 1987, the Army created the Army Supercomputer Program and a Department of the Army-level Supercomputer Functional Coordinating Group chaired by the Deputy Assistant Secretary of the Army for Research and Technology, Mr George Singley. Mr Singley appointed Harold Breaux as the Executive Secretary with the task of identifying issues and objectives and coordinating Army meetings.⁴⁴ The Army program was very successful, expanding the earlier BRL success. Besides procuring the BRL supercomputers, it acquired and placed a Cray 2 at the Tank-automotive and Armaments Command (TACOM), Warren, Michigan (1988), and a Cray Y-MP at the Corps of Engineers Waterways Experiment Station (CEWES), Vicksburg, Mississippi (1988). The plan also established an academic center, the Army HPC Research Center (AHPCRC), at the University of Minnesota (1989). The AHPCRC offered the first broadly available parallel computers, and AHPCRC staff scientists were located at both the university and Army sites and were an important bridge between university and Army researchers.⁴⁵

In August of 1989, Harold Breaux's network staff created the unclassified Army Supercomputer Network (ASNet), connecting five Army laboratories and AHPCRC, and the classified CASNET connecting four Army laboratories; the architect of both was Bob Reschly. The Army established a funding line for the

acquisition and operational costs of the total program. With this as background, in 1992 Mr Singley named Mr Breaux as the Army representative on the DOD HPC Working Group (HPCWG).

In the early 1980s, Kirtland Air Force Base (AFB) purchased commercial time on a Cray 1 from United Computer Services and made it available to other Air Force organizations (e.g., to the Arnold Engineering Development Center [AEDC] to support aircraft engine and wind tunnel testing). From the 1980s to early 1990s, the Air Force Systems Command (AFSC) purchased high-performance computers. The Air Force Weapons Laboratory at Kirtland AFB originally had a Cray 1 and then upgraded to a Cray 2. Lieutenant Clifford Rhoades was one of three authors (Major, later Colonel, Edmund Nawrocki and Mr Denzil Rogers were the others) of a report which justified the Cray 1. After the HPCMP was created, Dr Clifford Rhoades took an IPA position to work at the Phillips Laboratory at Kirtland AFB to help the Laser and Imaging Directorate stand up the Maui-High Performance Computing Center (MHPCC).⁴⁶ Dr Rhoades was later hired into Dr Charles Holland's vacated position at the Air Force Office of Scientific Research as Director, Mathematics and Geosciences, and assumed Dr Holland's duties as Air Force Science & Technology Principal on the HPCAP. Dr Rhoades has served as the HPC Software Institute Program Manager, Technical Director of MHPCC, and as Defense Research and Engineering Network (DREN) III and DREN 4 Acquisition Lead. He remains active in the program today.

The Air Force issued a moratorium on supercomputer purchases in 1987, until a comprehensive master plan was developed. Jeff Graham from Wright-Patterson AFB led a working group from the labs to find a path forward for centralized computing rather than dispersing it. The plan his group created was approved in January 1989 and one key element was development of the Air Force Supercomputer Network (AFSNet). Jeff Graham recalls⁴⁷:

The AF was also running supercomputers as part of the AF Supercomputer Masterplan. Resources at Kirtland (Capt Roie Black), Eglin (Calvin George) and WPAFB (Joe Dowdell) – along with the onset of the Maui Center that was operated by staff at the University of New Mexico. Unfortunately, it was fee for service – where customers had to pay 10% of the cost via taxing algorithms – a plan that didn't work well.

In the early 1990s, Major General Robert Rankine Jr, Deputy Chief of Staff for Technology, Headquarters, AFSC, had a commitment to HPC and taxed the labs (~\$800,000 per lab annually) to pay for the existing systems. These funds paid for most of existing systems operating costs but General Rankine believed the users' home organization should share a small portion of the cost; hence the 10% direct

cost to users that Jeff references. The program received mixed reviews from the labs, some of which claimed to have little in the way of HPC requirements.

NRL invested early in supercomputing, beginning with delivery of Texas Instruments ASC serial 7 in 1976. In 1990, the Navy produced a strategic plan⁴⁸ identifying HPC as essential to the success of the Navy's research and development efforts. The Navy took an evolutionary approach focused on local needs, locally managed facilities, addressing grand challenges, and creating a nationwide network of supercomputing resources. The report decomposed Navy research and development into 10 computational sciences interest groups, reminiscent of what would become the computational technology areas of the HPCMP.

There are three names in the list of authors of the Navy strategic plan that feature prominently in the creation of the HPCMP: Mr Don Endicott, Dr Leland Williams, and Dr Aram Kevorkian. Two user requirements representatives from the report, Mr Myles Hurwitz and Mr Stephen Schneller, became mainstays in the HPCMP.⁴⁹

The House of Representatives Committee on Appropriations directed the DOD to submit a 5-year master plan for the acquisition and use of supercomputers. The response was the DOD Supercomputer Acquisition Master Plan published in October 1990.⁵⁰ The purpose was to establish a strategy and mechanism for coordinating the requirements, plans, policies, acquisition and use of supercomputers for all DOD components. The Army had already published its Supercomputing Roadmap, which served as its master plan. The Air Force completed its master plan in 1989 and the Navy in 1990.

The Supercomputer Acquisition Master Plan enumerated the then available supercomputers in the DOD:

- The Army had four: two at BRL (Cray X-MP 4/8 and Cray 2), one each at TACOM (Cray 2) and CEWES (Cray Y-MP 8/6128); and a fifth one due to AHPCRC, a Thinking Machines CM-2. The Army had created the ASNet connecting BRL, other Army research facilities, and some universities.
- The Navy had supercomputers at Fleet Numerical Oceanography Center (Cyber 205), NRL (Cray X-MP 2/4), Naval Weapons Center, China Lake (Cray X-MP 1/16), Naval Underwater Systems Center (Cray X-MP 2/8), and David Taylor Research Center (Cray X-MP 2/4).
- The Air Force had supercomputers at the Global Weather Center (Cray X-MP 2/4), AEDC (Cray X-MP 1/2 and Cray 1s), Kirtland AFB (Cray 2 and Cray 1s), and Aeronautical Systems Division (Cray X-MP 1/2). The Air Force had also created the AFSNet. (Note that Aeronautical Systems Division became the Aeronautical Systems Center [ASC] in 1992.)

- The Defense Nuclear Agency (DNA, now known as the Defense Threat Reduction Agency, DTRA) had a Cray X-MP 4/16 located at and operated by LANL.

Conditions were right for a concerted DOD-wide initiative for shared supercomputing resources.

The National Defense Authorization Act for Fiscal Years 1992 and 1993¹ brought the services together under a single HPC umbrella. It was established without prescribed form, tasking the Director, Defense Research and Engineering (DDR&E) to submit a plan to Congress for the new program by 1 April 1992. The same month that the legislation was enacted, a new DDR&E was appointed, Dr Victor Reis. Dr Reis, formerly the director of DARPA, pulled together service representatives, the Institute for Defense Analyses (IDA), and his former DARPA colleague, Dr Gil Weigand, to create and deliver the plan in less than 4 months.

4. Adventures in Computing

4.1 The Modernization Plan

On 7 January 1992, the DOD Supercomputing Advisory Group met, including representatives from the DDR&E, Office of the Secretary of Defense (OSD), DARPA, DNA, Army, Navy, and Air Force. The purpose of the meeting was to organize to write the plan for Congress. At the meeting, Joe Batz of OSD discussed the Supercomputing Master Plan. Nine days later, a subgroup of the DOD Supercomputing Advisory Group met to select DARPA as lead for a working group to develop the plan for Congress. The first Supercomputer Master Plan Working Group meeting took place on 4 February 1992 at DARPA, led by Gil Weigand.

The main topics of discussion were supercomputer centers and networking. A decision was made that funding should be centralized with decentralized execution. Dr Weigand, via DARPA, funded early access systems (parallel architectures). He pushed for parallel systems over the more traditional vector systems, a position in which DARPA was heavily invested but which was not uniformly accepted as the approach the program should take.⁵¹

Ms Cathy McDonald was a member of the working group and was employed by IDA, where weekly working group meetings were held. Cathy McDonald recalls that during lunch breaks, she would print the revised documents and distribute them to members for the after-lunch session. The final version was printed on a Sunday and hand delivered to Gil Weigand's home.⁵² The HPC Modernization Plan⁵³ was approved by the DDR&E and submitted to Congress in May 1992, though dated 31

March 1992. Notably absent from the plan was any mention of software or subject matter expertise. The next challenge was to write the Implementation Plan.

4.2 The Implementation Plan

By September, the Supercomputer Master Plan Working Group re-chartered itself as the HPCWG. Key participants included Lt Col Larry Davis, PhD (AF), Paul Lewis (AF), Gil Weigand (DARPA), Harold Breaux (US Army Research Laboratory [ARL]), Leland Williams (NRL), Pam Fields (OSD), Tony Pressley (ARL), Denise Lenois, and Ken Hong-Fong (DNA). Their primary purpose was to write the Implementation Plan. Gil Weigand led the HPCWG with Joe Batz as co-chair. Joe led program charter development and was tasked to define the authority of the STWG, a DDR&E advisory group, relative to the HPCWG. Figure 1 shows the HPCWG organization.⁵⁴ (In 1992, BRL and other Army laboratories were reorganized into the Army Research Laboratory, ARL. It is now known as the US Army Combat Capabilities Development Command Army Research Laboratory, DEVCOM ARL.)




Fig. 1 Structure of the HPCWG in April 1993⁵⁴

Congress had appropriated \$42 million for procurement and an additional \$45 million was expected for FY93. IDA was hired for ongoing support and future meetings were held at IDA. A draft call was prepared soliciting service/agency interest, requirements, and proposals that could be incorporated into the Implementation Plan. During the fall of 1992, Tony Pressley took the lead for developing criteria for evaluating proposals and the HPCWG elected to use outside

reviewers, such as from NASA and the National Oceanic and Atmospheric Administration. They had to figure out how to spend their funding before it was rescinded.

They developed a Program Objective Memorandum (POM) initiative, with DDR&E support, documenting requirements and proposed management structure. This was the crucial step to get the program into the DOD budget instead of it being a year-to-year congressional appropriation. An Early Access Systems Committee was formed, led by Harold Breaux; a Networking Committee, chaired by Don Endicott; and a Stable Systems Committee, chaired by Lt Col Larry Davis. They agreed that the executive director of the program (later renamed the director) would be a rotating slot among the services with a 2-year tenure. Discussion ensued on how to divide the resources among classified versus unclassified systems, large versus small systems, stable versus newest systems; handling operations and maintenance, and allocations.

Meetings were held with the Defense Information Systems Agency (DISA) regarding the network (buy-in from DISA was necessary to have an independent network for the program). In November 1992, the HPCWG was notified that the POM initiative was successful and the program would be put into the budget. They prepared a call for early access systems proposals that would be capped at \$5 million per system. Example early access systems were the Thinking Machines CM-5 provided to NRL and AHPCRC, a Kendall Square Research-1 (KSR-1) provided to ARL, and an Intel Paragon for the Air Force ASC at Wright-Patterson AFB.

On 17 December 1992, the HPCWG briefed the Implementation Plan to the DDR&E and the STWG; the DDR&E approved the Plan.⁵⁴ There was discussion of HPC processes and an HPC consortium; \$9 million was committed to a consortium with a possibility of that being increased to \$12 million. The consortium became the National Consortium for HPC (NCHPC),⁵ a DARPA-NSF collaboration and a collection of academic institutions with expertise in HPC and parallel computing. HPC expertise and software were components missing from the Implementation Plan.

The approved plan stipulated that hardware consist of stable HPC systems and early access to newest systems. Stable HPC systems were intended to reduce science and technology computational deficiencies immediately. The criteria required that they be commercial off-the-shelf systems with mature (high degree of reliability and ease of use) operating systems, software development tools, applications software, and comprehensive user support. Early access systems were intended to accelerate transition to scalable systems to solve next generation DOD problems, and to gain

knowledge necessary when selecting systems for acquisition as future stable HPC systems. The criterion was simple—they were to be the newest HPC systems.⁵⁵

4.3 Organizing the Program

The new year, 1993, kicked off with a defining meeting for the future HPCMP, 27–29 January.⁵⁴ In attendance were Pam Fields, Don Endicott, Harold Breaux, Joe Batz, Gil Weigand, Larry Davis, and Ken Hong-Fong. There was a discussion of contracting options for the hardware and networking and how to handle requirements for massively parallel processor systems. A brief was prepared on the Program Management Office to be presented to the DDR&E. It was decided that three plans would be prepared annually: 1) a Forecast Plan each December; 2) a Modernization Plan each February, and 3) an Implementation Plan each September.

A process was developed for the next round of proposals. A suggestion was made to combine early access systems with applications development and to possibly define grand challenges (which later became Challenge Projects and later still morphed into the Frontier Projects we have today). The allocation to the services and agencies of 30%-30%-30%-10% was adopted, with the services given authority to set priorities within their allocation. There was debate over who would be allowed to use the computers, civilians or also contractors. It was concluded that contractors would be allowed access as long as they had a DOD sponsor. They also agreed to recommend that the program not be fee-for-service.

Don Endicott briefed the schedule for the Interim Defense Research and Engineering Network (I-DREN) and DREN. It was concluded that the stable systems would be purchased all at once for all centers. The HPCWG discussed preparing a brochure to be ready in time for the user group meeting planned for CEWES. There were concerns that if the systems were considered infrastructure then the HPCMP might have to charge for their use under (then) DOD policy. The plan was to have a contract in place by mid-April to buy the first stable systems. Early access systems were scheduled to be delivered by the end of March.

The following March meetings were held to plan a \$128 million budget for 1994 for the program. The HPCMP was designated a Major Automated Information System (MAIS) Program, subject to MAIS Review Council (MAISRC) oversight. As a consequence of the size of the upcoming procurements, and special interest from Congress and OSD, it was also designated a Major Defense Acquisition Program (MDAP) and subject to that oversight as well. Valerie Miller from ARL established integrated product teams and led the group that produced all of the required documentation and obtained OSD approval in record time, allowing the procurements to proceed.⁵⁶ Valerie eventually moved to the Pentagon, where she

was responsible for organizing and coordinating the process for the research directorate to review and comment on all MDAP and MAIS major milestone reviews, and worked closely with OSD leadership in developing the Reliance 21 Communities of Interest.

4.4 Beginning Operation

In the summer of 1993 the first stable systems were to be installed. Early access systems were approved for Maui and Alaska; at least 30% of these machines were reserved for environmental research and development.

The first User Group Conference was held in Vicksburg, Mississippi, at CEWES, 18–19 May 1993, hosted by Steve Adamec. The Roadmap, Implementation Plan, Networking, and DREN concept were presented. The I-DREN initial operating capability was scheduled for 15 August 1993. The first six technical emphasis areas had been selected and were announced: computational mechanics, computational fluid flow, computational chemistry, computational electromagnetics, climate-weather-ocean modeling, and signal/image processing. These six areas reflected the NCHPC technical areas and represented an important early collaboration.

A discussion arose on how to structure onboarding of users, assigning resources, managing this process, and so forth. It was agreed that each service/agency organization would be responsible for allocating, prioritizing, and managing their respective allocation and providing a close connection to users. A position was created, entitled the Service/Agency Approval Authority (S/AAA), and continues to this day, now led by a Service/Agency Allocations Officer, with subordinate roles delegated throughout each organization. The Allocations Officer manages the accounts, allocates and tracks core-hours, and helps onboard new users. Today, the HPC Advisory Panel (HPCAP) principals are the senior officials validating both requirements and allocations to the HPCMP.

The HPCWG met on 19 May and the Stable Systems Committee chartered the Mature Systems Advisory Panel (MSAP), which later became the Shared Resource Center (SRC) Advisory Panel and later still became the User Advocacy Group, which exists today. Kay Howell from the Navy chaired the MSAP. The HPCWG decided that each site should use their own security procedures and policies. (This later became an issue in establishing a common login and common user experience across the HPCMP.) The MSAP selected NASTRAN, GAUSSIAN, and ABAQUS as three of the first applications to be installed. The Stable Systems Committee tasked the MSAP with management of software acquisition. Jeff Graham remembers: “It seemed that the first ‘mature system’ had no application software

whatsoever. I worked on Kay's team and all services brought their requirements, and we broke them into 4 categories in terms of priority."⁴⁷

In July 1993, Vic Reis and Gil Weigand left the DOD and moved to DOE, creating and managing the landmark Accelerated Strategic Computing Initiative.²⁵ Dr Anita Jones, from the University of Virginia, was appointed the new DDR&E in August 1993, bringing a strong background in computer science and software.

During the fall of 1993, the nomination and planning processes and draft schedule for SRCs were discussed, and the term Major Shared Resource Center (MSRC) was used for the first time. Dr Jones issued a call for proposals to the services and agencies for MSRC candidate sites. The agenda was planned for a 7 December user group meeting at Wright-Patterson AFB in Ohio. The MSAP now assumed responsibility for user group meetings.

In November, a meeting was held to brief Dr Jones. Attending were COL Al Sullivan (military assistant to Dr Jones), Larry Davis, Harold Breaux, Joe Batz, and Don Endicott. Software was identified as a bottleneck in effective system utilization and Dr Jones directed the team to add software as a new component of the program. A decision was made to leverage software tools from other sources (e.g., DOE and NSF). Dr Jones stated emphatically that "if our only goal was to put hardware at the DOD Labs, I would end this office's role in the Program and give the money directly to the Labs."⁵⁷ She also clarified that in evaluating MSRC proposals that centers should be located at sites with a concentration of critical research exploiting HPC. Finally, Dr Jones' behind-the-scenes support helped streamline the MAISRC process, which enabled Valerie Miller to rapidly meet approval milestones.⁵⁷

On 7–8 December 1993, the second user group meeting was held, at Wright-Patterson AFB, Ohio. The program was officially named the High-Performance Computing Modernization Program.

The first quarter of 1994 brought many developments.⁵⁸ It kicked off with the HPCWG addressing milestone requirements for the MAISRC. A recommendation was made to hire Cathy McDonald full time for the requirements analysis, which was due by 7 June. Tony Pressley was named the first HPCMP director. The MSAP met to discuss software requirements and benchmarks; and the full Stable Systems Committee met to discuss the acquisition plan, which could not be finalized until completion of the requirements analysis and the analysis of alternatives. Tony Pressley had met with Dr Jones and debriefed the HPCWG. Tony recommended that the working group become an advisory body to the HPCMP and recommended deferring the software initiative to FY95. (The HPCWG became the HPC Advisory Panel that exists today.)

Proposals for MSRCs were received from the services and agencies in response to Dr Jones' call. The Air Force nominated the Aeronautical Systems Center at Wright-Patterson AFB. The Army actually nominated two sites, contrary to the data call, ARL and the Tank-Automotive and Armaments Research, Development and Engineering Center. The DNA proposed LANL to host their center. The Navy selected the Naval Oceanographic Office (NAVO) with a unique arrangement. The Navy proposed sharing 15% of the HPC resources with Navy operational environmental forecasts, with the Navy offsetting the associated costs with in-kind funding. This has been very successful for the Navy and continues today.

In making the final selection, LANL was judged to be highly qualified and offered strong credentials for hosting an MSRC; however, their proposed costs far exceeded those of the other proposals. In the end, the Air Force ASC, ARL, and NAVO were selected. CEWES had been operating the largest stable systems and was made a de facto MSRC. They did, however, submit a proposal that was evaluated comparatively with the other MSRC proposals.

In a 2 March 1994 HPCWG meeting at IDA, Tony Pressley gave a program overview and stated that requirements should be collected that reflect more than just current user needs. Steve Adamec discussed the CEWES system; Leland Williams briefed the Scalable Software Initiative; Don Endicott briefed DREN; and Hank Dardy briefed the NRL system. Dr Jones arrived and stated that hardware is the least important part of the program, and that software is crucial to the DOD. She sought long-term relationships with consultants and encouraged attendees to keep their focus on applications and not on communications (the network). She noted that scalability is important and that the program should closely follow emerging standards. (Recall that MPI 1.0 was released in May 1994.⁵⁹)

By April 1994, the requirements analysis⁶⁰ was completed, conducted by a team of eight people led by Dr Larry Davis and included Mr Jeff Graham from the Air Force and Ms Cathy McDonald. It began with a questionnaire and follow-up in person interviews. An early estimate placed the number of potential users at 2,000, much too large for individual surveys. The services and agencies were requested to organize their functional requirements into technical areas, and to consolidate them by user groups, rather than by individual. In other words, requirements represented aggregated requirements of an organization. The requirements populated the original six technical areas, which had come from collaboration with the NCHPC, and revealed three new ones: Forces Modeling and Simulation/C⁴I; Environmental Quality Modeling and Simulation; and Computational Electronics and Nanoelectronics. The conclusion was that a complete HPC environment was needed, not just hardware: a balance of hardware, memory, primary storage, archival storage, network, software, programming environments and tools,

expertise, and training. It identified the need for teraFLOPS systems over the succeeding 5 years.

Larry also led development of a companion requirements addendum⁶¹ and an analysis of alternatives,⁶² both published in May 1994. The addendum addressed the requirement for system integration contracts for the MSRCs and support for parallelization for a range of computational architectures. The analysis of alternatives focused on program management, distribution of HPC resources, and system integration, concluding that the most advantageous option was a centralized joint program with a negotiated systems integration service contractor to provide complete HPC environments at four MSRCs. Smaller distributed sites would be created to meet additional local requirements.

The HPCWG met at CEWES on 16–17 May 1994. Larry briefed the Advanced Planning Briefing for Industry (APBI) that would be presented in preparation for a stable systems request for proposal (RFP). Dr Charles Holland, from the Air Force Office of Scientific Research and senior Air Force representative to the HPCMP, briefed DARPA programs.

On 13 June 1994, the HPCMP briefed service executives. It was recommended that the HPCMP develop closer coordination with the basic and applied research community. There was discussion of a recent US Government Accountability Office (GAO) review⁶³ of the HPCC Program requested by the House Armed Services Committee. The GAO recommended that the HPC Program explicitly delineate its overall goals, objectives, and development strategy; set priorities and measures for specific technology areas; take additional steps to promote industry participation, especially in the program planning process; and provide greater support for standards setting. These recommendations were also considered by the HPCMP.

A User Group Meeting was held in San Diego on 19–20 July 1994 at the Naval Command, Control and Ocean Surveillance Center.⁵⁸ Tony briefed a possible expansion of the program, the need for expertise in scalable software and discussed the Scalable Software Initiative. Phil Dykstra briefed the network. Tony encouraged users to organize autonomously to share their experience. Denise Brown of ARL briefed experience with the KSR-1; thumbs down for scientific computing. AHPCRC, CEWES, and the Air Force ASC were reviewed. The MSAP decided that initially, CEWES would be unclassified only, ARL would be classified only, and NAVO and ASC would offer both.

After the user group meeting, the Statement of Work Team met and Harold Breaux argued for a 4-year hardware cycle, still used today. The team created a multilevel performance contract for stable systems. It would contain three performance levels

and a research and education component, which became PET (Programming Environments and Training) and part of the integration contracts.

An HPCAP meeting was held to discuss the Modernization Plan for the following March. (Transition from the HPCWG to the HPCAP was now complete.) Distributed centers (which replaced early access systems) were being treated like mini-centers—they were expected to serve a broad local user community and the HPCMP paid operations and maintenance costs. There was a discussion of what the critical applications were and of the underfunding of software. DOE platforms were considered for additional hours on a “spot market” basis.

4.5 Test and Evaluation Join the Team

The autumn of 1994 saw much activity. An HPCAP meeting was held with a major topic being inclusion in the program of developmental test and evaluation (T&E). Dr Charles Holland suggested rolling T&E requirements into the overall service allocation (instead of having a separate T&E requirements document). It was concluded that there would be no effect on the MSRCs but there could possibly be on the DREN, depending on the need to add new sites. Knowledge of this would have to await requirements gathering from T&E organizations.

The SRC Advisory Panel (no longer referring to itself as the MSAP) made software recommendations organized by CTA (Computational Technology Area, the first official mention of this term), and made a recommendation to form a security committee. The Scalable Applications Software Initiative (SASI) was discussed and \$30 million were allocated with a data call to be sent to services. SASI later became the Common HPC Software Support Initiative (CHSSI) and the budget was eventually reduced to \$20 million annually. Dr Jones insisted that the program have an alignment with the Grand Challenges. Larry Davis participated in responding to her request: “We designed our Challenge Project selection process based on what NSF was doing to sponsor large grand challenge projects. I remember observing their selection process before we established our process.”⁶⁴ In FY1997 Challenge Projects were created, later becoming Frontier Projects in FY2015. (Frontier Projects are intended to explore research, development, test and evaluation and acquisition engineering outcomes that would not be achievable using typically available HPCMP resources.) Sandia National Laboratories later submitted a proposal to the HPCMP to team with them on a teraFLOPS system initiative, one facet of the HPCC Program; Dr Jones declined.

On 25 October 1994, the HPCAP met. Phil Webster of the Air Force came on board to take over the requirements effort; Larry Davis and Cathy McDonald worked for him. The issue of how to distribute the workload among the centers relative to the

CTAs arose. PET and the CTAs were already aligned. It was agreed to assign specific CTAs to specific MSRCs, depending on MSRC host organization research. Computational Fluid Dynamics (CFD) was assigned to every center since it was so large and pervasive. It was decided that PET would handle the common software tools and that benchmarks would be created by CTA.

On 29 November 1994, Tony Pressley, Joe Batz, Cathy McDonald, Phil Webster, and Larry Davis visited John Bolino, T&E liaison to the joint staff, and his team to discuss T&E. This was a foundational meeting for HPCMP support to the T&E community. They discussed the need for a broadband network connecting T&E centers and for distributed centers. Arnold Engineering Development Center and Eglin AFB were suggested as distributed center sites. Dr David Brown and Jeff Highland of the Army Test and Evaluation Command visited Tony to discuss the DREN and a distributed center for White Sands Missile Range (WSMR). (A CM-5 was later awarded to WSMR in 1996 and upgraded in 1997.⁶⁵)

Several T&E site visits were scheduled. At AEDC, an Air Force developmental T&E center, personnel presented requirements for engine and wind tunnel testing, the process of test planning through test analysis, and the unique test capability offered by AEDC. A senior Air Force official stated that a long-term goal of modeling and simulation (M&S) was to reduce flight testing by 80% by 2015. (This did not happen.)

The next visit was made to Eglin AFB, an Air Force developmental T&E and flight test center. Eglin personnel discussed their hardware-in-the-loop Guided Weapons Evaluation Facility, Seek Eagle Office, time-space-position measurements, and store certification. Later trips were made to the Army Combat Systems Test Activity (now Aberdeen Test Center), Pax River, China Lake, Edwards AFB, Navy Pt Mugu, WSMR, Dugway Proving Ground, Yuma Proving Ground, and the Joint Interoperability Test Command at Fort Huachuca, Arizona. Test and evaluation was formally incorporated into the HPCMP by Congress on 1 December 1995.⁶⁶ A tenth CTA was created, Integrated Modeling and Test Environments. Dr Danny Weddle of NAVAIR at Pax River, a strong proponent of T&E in the HPCMP, recommended Dr Andy Mark of ARL as the CTA lead. Figure 2 is a time line for the first 3 years of the program, highlighting major milestones and developments.




Fig. 2 Time line for the creation of the HPCMP

On 23 January 1995, an HPCAP meeting was held. Leland Williams discussed SASI and success stories, CTAs, inclusion of T&E in the program, and invited T&E to join the CTAs. Tony began his presentation by stressing that Congress had directed the program to view T&E as a partner in the HPCMP and that T&E management should be considered as a functional proponent for the program equivalent to the DDR&E. The members discussed the urgency of SASI in terms of when Dr Jones would require them to implement it. It was noted that commercial software companies were not making money on scalable software, since not many businesses were using parallel computers. Tony briefed the status of the MSRC RFP, discussed upcoming evaluations, and questioned if the HPCMP had enough evaluators. It was decided that two awards would be made in FY95 and two in FY96. Bob Reschly was confirmed as the Program Manager (PM) for DREN. A joint network Program Office was established with DISA, and DISA assigned a representative to work with Bob.⁴²

On 24 October 1995, a user group meeting was held at NRL. Kay Howell welcomed everyone; the Shared Resource Center Advisory Panel (SRCAP) was now an important part of planning the meetings. Steve Schneller provided an SRCAP overview and Jeannie Osborn of NRL gave administrative details. Dr Jones

announced that Tony would be leaving the program and that Kay would succeed him as the next director. Jeannie replaced Kay on the SRCAP and remained there until it became the User Advocacy Group, which she supported for approximately 20 years, several of those as chair. She was a long-time S/AAA for NRL and designed and implemented the Navy's process for apportioning Navy allocations across systems. She was a tireless advocate for both NRL and program users, prodding program leadership to always consider the users' point of view. She conducted training for S/AAAs for most of her tenure. She was one of the most influential people in the program for the entire time she was associated with it.

Tony Pressley returned to ARL leaving a strong program legacy: first director, standing up the program office and staff, getting hardware and network procurements moving, standing up the I-DREN, standing up MSRCs, and getting MAISRC and MDAP documentation in place.

Calendar year 1996 brought a new director, completion of the four MSRCs (CEWES, ARL, ASC, and NAVO), as well as centers in Maui and at the University of Alaska at Fairbanks. Steve Adamec became the first CEWES MSRC director. Charles Nietubicz was the first director at ARL; Tom Dunn was the first director at NAVO; and Paul Shahady was the first director at ASC. The technology refresh process, termed Technology Insertion, was developed, and the DREN contract was awarded, replacing the I-DREN.

Contracts of 8-years' duration were placed for integration contractors at each of the four MSRCs, including operations, maintenance, help desk support, and subject matter and parallel programming expertise (via the PET component). Nichols Research Corporation was awarded two of the contracts, for CEWES and ASC in March and May 1996, respectively. (Nichols Research was later acquired by CSC.) Grumman Data Systems was awarded the NAVO contract in May 1996, and E-Systems, Inc. was awarded the ARL contract in August 1996. (E-Systems was subsequently acquired by Raytheon.)

5. The Network

The program was under a great deal of pressure to spend its funding exclusively on hardware to support US supercomputer manufacturers. Tony spent a large percentage of his time advocating with congressional staffers for a balanced program. Don Endicott, Harold Breaux, Larry Davis, and Cliff Rhoades recognized the need for networks and were vocal proponents, with Don leading the Network Management Committee of the HPCWG. The DOD lagged the DOE, which had established programs in computational hardware, and the NSF, which had a strong program in networks and collaboration, and HPC centers.

In a 1992 meeting at IDA, the interim network idea was hatched.⁴² The team recognized that the immediate solution might differ from the long term, permanent solution. The Army approach had always been to develop custom solutions in house, test them, implement them, and develop lessons learned, then move to an off-the-shelf solution for the next generation. This approach was chosen for the HPCMP network as well. The Army had developed ASNet based on leased T-1 circuits, bought its own routers and managed everything directly. For an immediate solution to the HPCMP network requirement, ASNet, along with the Air Force AFSNet, became the initial wide-area network (WAN) and initial nodes. T-1 circuits were leased from DISA but the HPCMP bought and managed the end routers. The Navy had some leased T-1 circuits and dedicated WAN connectivity for the needs of specific projects but did not yet have general research-oriented WAN capability. The Navy T-1 circuits were aggregated with ASNet and AFSNet to form the I-DREN in August 1993. Figure 3 shows a layout of the I-DREN as of late calendar 1994.




Fig. 3 The I-DREN circa late calendar 1994

The network was formally named the DREN based on the NREN, which utilized commercial services. The I-DREN used leased lines and purchased and managed hardware; the DREN would be commercial services, reflecting the NREN/NSF approach. Bob Reschly recalls that Ron Broersma of the Navy may have been the first person to point out that DREN is NERD spelled backward. Bob led the RFP development, along with Carlos Fernandez from Kirtland AFB and Ron Broersma

from the Navy in San Diego. Joe Batz of OSD and Tony Pressley were a big help in getting the DREN moved beyond DISA.

Meetings were held with Lt Gen Albert Edmonds, Director of DISA, and his senior staff, and with LTG Emmett Paige Jr, Assistant Secretary of Defense for Command, Control, Communications and Intelligence, and his senior staff. The primary argument was that, while all agreed that DISA has the mission to provide DOD networking, the small (in comparison to overall DOD requirements) HPCMP environment needed cutting edge and nimble services that would be difficult for DISA to provide and manage among all the huge efforts DISA already managed. It was further argued that by being a small, focused, and nimble effort, the DREN could serve as a testbed for advanced services that would eventually be incorporated into mainstream DISA networking services. (Excellent examples of this include production use of Asynchronous Transfer Mode (ATM) at award of the first DREN contract, and DREN being designated a DOD pilot network for IPv6 services.) It was also acknowledged that DISA may someday absorb the DREN into mainstream services.

DISA assigned Ed Schoenborn to the HPCMP network team as Bob's deputy. Bob recalls Ed's essential role⁶⁷:

Ed was instrumental in ensuring we maintained appropriate linkages within the DISA technical and management structure. The interchange of information and ideas between the DISA and HPCMP groups improved the design and resulting solicitation, and meaningfully improved the odds the contract would be successfully awarded. Ed (with input from DISA staff) was able to identify some gaps and a number of elements that had caused problems with other DISA awards, allowing us to make changes to reduce the likelihood of trouble.

The network team also had a crew assigned from MITRE, one of whom was Bill Pittinger. Bill was former military and had conducted large military procurements. He was a very strong program manager and had the ability to look at a scenario and tell you the consequences. If you did not like the consequences, you revised the scenario. This was an invaluable skill in moving the DREN procurement forward. As the DREN effort grew, and specifications and the RFP had to be written, Bob stepped away from I-DREN and Rodger Johnson took over.

A DREN Industry Day was held on Aberdeen Proving Ground, Maryland, in 1995 and roughly 80 people attended. The network team described the scope and requirements. The DREN contract would treat the end product as a service and put the onus of managing it on the contractor (e.g., service delivery points [SDPs], metrics for service, and service level agreements). Bob recalls the AT&T

representative informing him that this was all well and good, but that AT&T would get the contract.

AT&T was in fact awarded the first DREN contract in the summer of 1996, a 5-year effort valued at somewhat under \$500 million. (This was no accident. In network team discussions it was felt that if the procurement were \$500 million, people would perceive it as one-half billion dollars, which would mentally produce much greater expectations than if the value was four hundred-some-odd million dollars.⁴²) The first SDPs were installed in December of 1996 and SDP0 testing took place until February 1997.

Rodger Johnson invested about \$2 million for development of high-performance encryption devices; everyone agreed they were needed. The investment was approved and he served as the focal point. This arose from ARL involvement with the Defense Simulation Internet Joint Program Office. The effort was successful; the General Dynamics KG-175 was the result and became a mainstay of the HPCMP and other programs.

Murray Huffman from CEWES succeeded Bob as PM DREN in 1996; Fig. 4 shows the early DREN time line. The DREN 1 contract ran from 1996–2002 and upgraded the T-1/T-3 service of the I-DREN to packet-based Ethernet and cell-based ATM. The 10-year DREN 2 contract, 2002–2012, was awarded to Verizon and provided IPv4 over Multi-Protocol Label Switching. The current contract, DREN 3, awarded to CenturyLink, runs from 2012–2022 and provides IPv6 over advanced community-of-interest collectives.⁶⁸ The DREN bandwidth, number of service delivery points, availability, quality of service, and capabilities continue to grow.




Fig. 4 Key events leading to the DREN

6. PET, PETT, PETTT, or What?

Vic Reis, Gil Weigand, Anita Jones, Charles Holland, and many other senior leaders involved during the formative period of the HPCMP understood the importance of scalable computers, software, and the need for expertise in parallel programming. The services had at best niches of parallel computing experience. DARPA had funded parallel computer development and addressed software, but it had yet to make an impact on the DOD. Several universities were growing resident capability, and creation of the NCHPC was intended to provide access to it. The scalable software initiative was slow in arriving, but leadership still needed to develop subject matter expertise native to the HPCMP. The PET initiative, incorporated into the first MSRC integration contracts, was the result. It provided MS- and PhD-level domain and parallel computing experts, many co-located at the MSRCs and other DOD laboratories. Identifying algorithms and codes for parallel implementation, training, and porting vector codes to parallel platforms consumed a large part of their time early in the contracts.

The first integration contracts, from the fall of 1996 to the fall of 2004, included software and programming expertise and each contractor had to propose how they would provide this. They all chose to subcontract that part of the work. The successful teams were E-Systems (later Raytheon) at ARL with High Performance Technologies, Inc. (HPTi) and the National Center for Supercomputing Applications for PET; Nichols Research (later CSC) at CEWES and Mississippi State University (MSU) for PET, and at ASC with the Ohio Supercomputer Center (OSC) for PET; and Grumman Data Systems at NAVO with the San Diego Supercomputer Center for PET.

The first PET contract was more geocentric than later contracts (i.e., the support was mostly local to the host organization, such as ARL) and the PET expertise was aligned with local needs. For example, OSC provided signal and image processing (SIP) support with requirements concentrated at ARL. As a result, nearly all OSC meetings were held at Aberdeen Proving Ground, Maryland.⁶⁹

Professor Stan Ahalt, currently with the University of North Carolina at Chapel Hill and Director of the Renaissance Computing Institute (RENCI), was with The Ohio State University (OSU) for the first PET contract and was the SIP lead. Since most of the OSC PET support would be provided by graduate students, Stan remembers being in meetings to work out exactly how they would handle government review of publications, especially thesis material, to ensure nothing restricted was put into the public domain but respecting the graduate student time line. They all agreed that reviews would be completed within a maximum of 30 days.

The PET academic leads advocated to the HPCMP office for independent PET contracts. Stan recalls being contacted by Professor Joe Thompson from MSU regarding interest in forming a consortium to compete for the follow-on PET contracts, resulting in the MOS consortium. The MOS consortium consisted of MSU (Prime Contractor), OSC, University of Texas, University of Tennessee-Knoxville, University of Alabama at Birmingham, and University of Hawaii, plus industry partners. The PET component of the integration contracts was removed after 5 years and competed as independent support. The MOS consortium successfully bid on three of the four contracts, at ERDC (stood up in 1999, formerly CEWES), ASC, and NAVO; HPTi won the contract at ARL. The result was the 8-year User Productivity Enhancement and Technology Transfer (PETT) contracts that began in the fall of 2001, affectionately referred to as PET 2. (The “U” is silent and invisible.) The domain expertise was chosen to align with the 10 CTAs.

For PET 2, Professor Ashok Krishnamurthy, also of the University of North Carolina at Chapel Hill and RENCi, was with OSU and supported SIP. He remembers wondering how the training portion of the contract would actually work. He credits PET 2 for teaching him how to take a large amount of material and, in 2–4 days, teach it in such a way that the students became productive with it right away. Students would come to PET training with their respective problems to solve. Stan recalls Ashok leaving class, going to his hotel, staying up much of the night, and returning to class the next morning with solutions for the students’ problems. PET training became a valuable resource.⁶⁹

During PET 2, the author worked with the Army Test and Evaluation Command and represented Army T&E on the HPCAP. PET 2 support for the Army test centers was excellent, especially the training. PET leadership actually created training for us when none existed. Test centers also took advantage of quick-reaction support, where a problem could be presented to a PET subject matter expert and, if he or she was available, could immediately begin working on it without further bureaucracy. We derived substantial benefit from PET 2 support.

Ashok relates an anecdote from PET 2. He and PET colleague, Dr John Nehrbass, were to teach a MATLAB course at NRL Monterey. They caught a flight out of Columbus, Ohio, and landed in the Bay area. John had two huge suitcases for at most a 4-day course. Ashok asked John if he was planning to extend his trip and take vacation. John said no but declined to disclose what was in the suitcases. After they checked into their hotel, John invited Ashok to dinner, knowing Ashok is a vegetarian. Instead of heading to a restaurant, John led them to his room where he prepared a complete vegetarian dinner, salad and all, from his suitcases. Ashok was surprised, the mystery was solved, and nourishment was had. Future travels with John no longer caused suspicion.⁶⁹

At the end of the PET 2 contracts, the HPCMP decided to consolidate all such support into a single contract; a comparable decision was made with regard to consolidating the MSRC integration contracts. The result was the User Productivity Enhancement, Technology Transfer and Training (PETTT) contract awarded to HPTi in August 2009 for 10 years. And, of course, it was referred to as PET 3 and the “U” was still silent and invisible. PET 3 consisted of multiple levels of support, from quick-reaction and short-term projects, to projects selected competitively and awarded annually, to projects funded directly by DOD customers to the PETTT contractor. It brought increased oversight compared to PET 2 and more inertia for obtaining project support and training. For example, even quick reaction support required review and approval up and down the PET 3 management chain.

The conclusion of PET 3 brought anticipation, what would the fourth “T” represent in PETTTT? Alas, the contract is named User Productivity Enhancement and Training and goes by PET once again. HPCMP leadership changed emphasis in this contract, awarded to General Dynamics Information Technology in July 2019 for 5 years. The new contract uses, almost exclusively, structured proposals and competitive evaluations for customers to obtain support; Mission Projects are the vehicle. Once approved, a project is assigned a PET expert for the duration, typically not longer than 6 months and less than 50% level of effort. As lessons are learned during project execution, a follow-on proposal can be submitted. Transition to the current PET contract also changed the expertise landscape. Many of the highly experienced and PhD-level staff moved on to other jobs and were replaced with less experienced BS- and MS-level staff. The loss of much PET corporate knowledge and quick-reaction assistance is an adjustment, but thus far the approach is working and providing expertise to the services and agencies.

7. Program Management

Tony Pressley presided over the HPCMP from its beginnings as a plan through its existence as a program, an organization, and a physical office. He managed the transition, established many of the elements and functions that still exist, initiated the effort to move the program beyond vector computing into parallel computing, and set the program on its course. Kay Howell took the fledgling program, completed the procurements and put the major contracted elements in place for hardware, the DREN, and for the MSRCs, created the scalable software initiative, CHSSI, brought T&E into the community, created Challenge Projects, implemented secure login via Kerberos, created the PET initiative, and negotiated with the DOD Comptroller’s Office to avoid fee-for-service for MSRC resources.

Kay Howell left the program in 1997 to take over the National Coordination Office (NCO) for Computing, Information, and Communications, which is the forerunner of the NITRD Program. Kay was tasked to develop and deliver a plan to put the NITRD Program in place. Cathy McDonald also moved to the NCO and helped write the plan, then returned to the HPCMP in 2000 when Kay retired.

Tom Dunn from NAVO succeeded Kay Howell as the director. Tom was the first MSRC director at NAVO, and later held positions in the HPCMP as PET project manager, MSRC project manager, HPCMP deputy director to Kay, and HPCMP technical director. Tom implemented the annual Technology Insertion process. He also initiated HPC success stories and championed the program for its impact on the research scientist and the Warfighter.

Dr Charles Holland preceded and succeeded Tom Dunn, and was the only member of the Senior Executive Service to be the program director. Dr Holland was the senior Air Force representative to the program during its formation and early recognized that data-intensive problems and near-real time data analytics would become future drivers. Dr Holland recommended that the position of director be changed from rotating among the services to a permanent position. He went on to serve as the Deputy Under Secretary of Defense for Science and Technology (the HPCMP reported directly to him), and as the director for the DARPA High Productivity Computing Systems Program.

Cray Henry became the first permanent director (i.e., not detailed to the HPCMP) in 2000, having worked as its Financial Manager and DREN PM. Cray led the program through change, adapting to emerging user needs and responding to new OSD oversight requirements until 2011, when the Army was given responsibility for the HPCMP. Cray oversaw the transition from traditional HPC platforms running custom versions of Unix, to the era of large-scale Linux clusters, which account for every platform in the program and the TOP500 today*. The first Linux cluster in the HPCMP was installed at ARL in 2003.

In 2005, Cray hired Dr Douglass Post from the LANL as Chief Scientist. Doug brought an extensive background in large-scale scientific software development and high-performance computing. In 2006, the GAO presented testimony before the House Armed Services Committee regarding DOD weapons acquisition.⁶³

In the past 5 years, DOD has doubled its planned investments in weapons systems, but this huge increase has not been accompanied by more stability, better outcomes, or more buying power for the acquisition dollar. Rather than showing

* www.top500.org

appreciable improvement, programs are experiencing recurring problems with cost overruns, missed deadlines, and performance shortfalls.

Later in the testimony we find, “At this time, however, DOD is simply not positioned to deliver high quality products in a timely and cost-efficient fashion.”⁶³ Doug and Dr Edward Kraft of AEDC proposed a solution.

The result is the Computational Research and Engineering Acquisition Tools and Environments (CREATE) Program, approved by the Under Secretary of Defense for Acquisition, Technology and Logistics in 2006, with funding programmed to start in 2008.⁷⁰ Ed provides insight into its origins⁷¹:

[T]he incentive to pursue a software initiative that resulted in CREATE was birthed at the HPC Users Group meeting in June 2005 at the Opryland Hotel in Nashville and took its initial shaping at a meeting at the Gossick Leadership Center in September of that year. By January 2006 the initial CREATE topics were selected and supporting letters from the services were prepared. Doug did a phenomenal job running the traps in DoD and got it put into the POM cycle.

Doug offers additional background⁷²:

I spent a few months in the Pentagon after I joined the HPCMP in June 2005 trying to identify something that the HPCMP could do that was immediately useful for the Military Services. The DoD was floundering with acquisition. Most of the DoD acquisition programs were behind schedule, over budget and building systems that didn't work very well.

Virtual prototyping using physics-based high performance computing codes was a win-win for the Military Services and the HPCMP. Our business model was to fund the services to build the codes. For instance, we sponsored Carderock staff to build codes to design ships. Since we were the sponsor, we were able to get them to use sound software engineering practices that Richard Kendall and I had picked in our careers at LANL and LLNL. We had many very helpful stakeholders from senior people like Ed Kraft (AF AEDC), Kueichien Hill (AFRL), and Scott Littlefield (Navy).

The purpose of CREATE is to reduce the cost, time, and risks of DOD acquisition programs by developing and deploying multi-physics-based HPC software applications for the design and analysis of military aircraft, ships, and radio frequency antenna systems. CREATE improves acquisition program performance by reducing design flaws, developing sound engineering designs quickly and flexibly, and enables systems integration earlier in the acquisition process. Today, CREATE tools are being used by more than 110 DOD organizations and there are more than 600 active user software licenses.⁷³

Some needs cannot be met by centralized batch computing and testing is a prime example. Once a test is under way, computing must be available on demand, not waiting in a batch queue. Some T&E applications require embedded and real-time computers, such as hardware-in-the-loop sensor/seeker test facilities and cockpit simulators. Some applications require above Secret computing resources not available at a DSRC. (The MSRCs were renamed DOD Supercomputing Resource Centers, DSRCs, in 2008.) None of these requirements could be met by the original structure and facilities of the HPCMP, but the program adapted to accommodate the needs.

Distributed centers, which had replaced early access systems, were treated as mini-shared resource centers, intended to serve a large local community of users. As user requirements evolved, especially T&E, it became expedient to repurpose distributed centers as dedicated resources belonging to the host organization. This allowed the program to meet an immediate requirement as well as reduce some of its oversight responsibilities and costs inherent in distributed centers. Dedicated HPC Project Investments (DHPis) were created and are competed annually to address specialized computing needs.⁶⁸

Cray Henry established workforce development outreach, which consisted of three programs: the midshipmen/cadet HPC summer program, HPC-focused fellowships in the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program, and an intern program funded and managed by each of the MSRCs. In 2008 the intern programs were consolidated under the HPCMP's Oversight and Outreach component, led by Valerie Miller, and recreated as the Joint Educational Opportunities for Minorities (JEOM) Program.

The midshipmen/cadet HPC summer program provided funding for faculty/student teams to work on a project during the summer at an HPC center or a laboratory, typically lasting 3–4 weeks and amounting to \$20,000–\$40,000 per service academy. The NDSEG Fellowship is a congressionally directed program sponsored by the Army, Navy, and Air Force to increase the number of US citizens in science and engineering disciplines; the services split this funding. The HPCMP directly funded 10 students per year over and above the NDSEG funding, or about 100 graduate students between 2001 and 2011. The JEOM program concentrated on eligible US citizen minority students (from freshman through postdoctoral persons of color, female, or persons with a disability) maintaining a 3.0 GPA or higher.

Cray initiated a return on investment (ROI) study that ran from 2006 to 2009 to quantify the value to the warfighter of investments made by the program.⁷⁴ It was led by Deborah Schwartz, who was detailed to the HPCMP from the Navy at Pax River. The study was conducted in three parts, focused in year 1 on the armor/anti-

armor portfolio of projects, in year 2 on the climate, weather, ocean portfolio, and concluding with the air vehicles portfolio. The result was that ROI across the three portfolios was between \$6.78 and \$12.92 for every dollar invested by the HPCMP in these portfolios, the variation resulting from considering multiple variants of a weapon system and uncertainty in estimating some categories of cost avoidance. In 2011, the HPC User Forum initiated the innovation award program with one category for ROI. To maximize its use, the HPCMP submitted sections of the ROI report as nominations over 2011, 2012, and two in 2013.*

In 2011, Cray Henry oversaw the smooth handoff of the program to the Army as one of his final duties as director. The Office of the Secretary of Defense had revisited the fundamental question of where acquisition programs belong within the DOD, having first taken it up in 2004. After a great deal of internal political debate and strong community involvement, a plan was formulated to move the HPCMP to the Army, which assigned it to ERDC, while protecting the budget and providing a strong oversight role for the DDR&E and key service principles.⁷⁵

John West of ERDC succeeded Cray as director, having been part of the source selection evaluation board for the original MSRC contracts as a new CEWES engineer. He led during transition of the program to the Army and was part of the team that made it happen.⁷⁶

The program was being shopped to at least one other service, but it had just suffered a \$30M cut by DDR&E, and no one wanted to take the program with a hole in its budget. However, Cray Henry patched that hole, and I knew the program was not only safe to take, but an excellent strategic opportunity.

The devolvment from OSD was directed by Congress as an effort to get DDR&E out of the business of program execution (devolvment in general; the HPCMP wasn't specified as far as I know, but as a \$300M program it was a very large budget piece that could help them meet their Congressional goal), but one of the few things it had that actually belonged in OSD was the HPCMP. The transition leadership that continued through the first several years, including Mary Miller and Jeff Holland and Ron Jost, provided strong leadership and had a very strong commitment to the program remaining purple. The risk now with the program in the Army is that as those leaders have left or retired the program is at risk of slowly being absorbed into the Army, and losing its DOD focus. This will be a real loss for the Department if it happens.

Besides bringing the program to a new steady state under Army leadership, John emphasized outreach and consolidated the three workforce development outreach

* <https://hpcuserforum.com/innovationaward/winners-previous.html>

components under a new directorate called Workforce Development, with Deborah Schwartz as its associate director. In 2014, John left the DOD to accept the position of Director for Strategic Initiatives for the Texas Advanced Computing Center (TACC) at the University of Texas-Austin. TACC currently operates the Frontera platform, which is number 9 on the TOP500 list and is the fastest computer at any university in the world.

Deborah Schwartz recalls⁷⁷:

In 2012 a DOD study was published suggesting that there were hundreds of “intern” programs and that they should be disbanded. The study resulted in the reevaluation and restructuring of the HPCMP’s workforce development efforts. During the restructuring, John West, consolidated the three components under a new Directorate called Workforce Development. Our own analysis during this time showed that only 2 or 3 of our 100 NDSEG graduates became DOD scientists or engineers. It was also determined that our support for the service academies was minimal in comparison to the NDSEG support. Our thesis was that the cadets and midshipmen were the future leaders in the DOD and exposing them to high performance computing and to the HPCMP now could benefit the DOD and programs well into the future as these young career military men and women become Project/Program Managers in Program Executive Offices. As a result, we reprogrammed the NDSEG funds directly into the service academies to enhance their HPC efforts, expanding their participation from a summer student program to include enhancing their computer science and engineering curricula as well.

As part of the 2012 reevaluation, we restructured the JEOM Program into the High Performance Computing Internship Program (HIP). HIP is a mentor driven process where mentors submit proposals to the HIP team and successful mentors then select their interns. Interns are expected to write a paper, prepare and present a brief and prepare a poster. The program created a Mentor’s and Intern’s Briefcases for guidance for both the mentors and interns. At the time of my retirement (2015), we exposed more than 600 interns and cadets/midshipmen to HPC and the HPCMP. While I was tracking the statistics, we successfully captured greater than 20% of the JEOM/HIP interns into DOD civilian and contractor positions.

Dr David Horner of ERDC followed John as director in 2015, leading the HPCMP toward new horizons. David was a researcher at ERDC and a long-standing member of the Army research and development community. He was very familiar with the program, having been an early HPC user and project leader, including leading one of the first Challenge Projects. As big data grew in importance, a Decision and Data Analytics CTA was created. The DOD also saw growing demand for computing above the level of Secret. These needs were addressed previously through DHPIs.

It became expedient to consider providing above-Secret shared resources; these are hosted today by AFRL and ERDC, and can accommodate up to Top Secret computing needs. Dr Horner was selected for senior executive service in 2018 and became director of the ERDC Information Technology Laboratory.

Dr Will McMahon of ERDC succeeded Dr Horner as HPCMP director. Will has continued to expand program capabilities, not just through the annual technology insertion process, but through new roles for the program. Cloud computing has been around for more than a decade but now is being assessed as a source of infrastructure-as-a-service, software-as-a-service, and platform-as-a-service for the HPCMP. Pushing computation to the data is also a mantra of the day and one response of the HPCMP was to create a mobile computing capability, named SCOUT (SuperComputing OUTpost), specially designed for machine learning and data intensive workloads. It is almost completely self-contained, requiring only external power, providing in excess of 6 petaFLOPS (10^{15} FLOPS) of single precision performance and 1.3 petabytes of storage in a 53-ft trailer.

Large-scale data analytics, visualization, and machine learning present a new workflow compared to compute-intensive work, and different architectures are required. Under Dr McMahon's leadership, in 2020 the program announced its first large-scale reconfigurable platforms designed specifically for data analytics and machine learning. They are provided by Liquid Computing and located at ARL and ERDC.

8. Impact

Many examples exist of the impact of the HPCMP for the DOD and outside the Department over the past 30 years. We present a few memorable ones.

8.1 Space Shuttle Columbia Accident Investigation

The Space Shuttle Columbia and its seven crew members on STS-107 were lost 1 February 2003 during reentry, the result of an unnoticed incident as the vehicle ascended on 16 January. A piece of foam insulation was shed from the external tank and struck the leading edge of the left wing of the Orbiter. NASA and Army personnel and codes, and NASA and HPCMP computational resources, were applied to support the Columbia Accident Investigation Board. Analyses included steady-state and unsteady calculations performed with the Overflow and Cart3D flow solvers. Unsteady calculations included moving body, 6 degree-of-freedom simulations of foam debris shed from the region of the left bipod-ramp of the vehicle. The analysis provided an estimate of the speed at which a piece of debris

would strike the wing leading edge of the Shuttle Orbiter, results which were supplied to the Columbia Accident Investigation Board.

Dr Robert Meakin, then an Army Civil Servant with the US Army Aeroflightdynamics Directorate at the NASA Ames Research Center, was working under an Army/NASA joint agreement. He supported the investigation using the Overflow code, which he helped develop and to which the HPCMP had contributed. Within one day, the ARL MSRC provided Bob with 40,000 processor hours and a dedicated queue on two ARL platforms. The NASA Advanced Supercomputer Facility at Ames Research Center contributed 400,000 additional processor hours, and the NASA Ames CFD Group provided a team of experts using Overflow and Cart3D.^{78,79} The MSRC staff helped port Bob's code to MSRC computers, made special provisions for large storage space, and assisted in flow visualization and 3-D animations for communicating results.

The primary data provided by the simulations to NASA Johnson Space Center included debris trajectories as functions of debris mass, shape, and mechanism of separation; debris velocity at impact; debris rotational velocity at impact; and animations. The data was used by NASA to establish the range of conditions for ground-based foam-firing tests conducted by the Southwest Research Institute. The tests demonstrated that it was possible for a piece of foam debris to cause massive damage to the Shuttle Orbiter wing reinforced-carbon-carbon panels and T-seals, creating a breach where hot gases could enter the wing structure during reentry.^{80,81} The analyses were made possible by a dedicated team of NASA, Army, university, and industry specialists provided with state-of-the-art high-performance computing resources and experimental facilities.

8.2 Army CH-47 Advanced Chinook Rotor Blade

A major part of the Cargo Helicopter Project Office (PMO) CH-47F Block II Chinook upgrade program is the complex Advanced Chinook Rotor Blade (ACRB), designed to improve hover performance in high/hot conditions. While flight tests demonstrated the desired hover payload performance, it also revealed larger than expected flight control system loads during high-speed forward flight. The PMO formed a team with the US Army Combat Capabilities Development Command Aviation and Missile Center (formerly the Aviation and Missile Research, Development and Engineering Center) and Boeing to address the problem. They applied the CREATE-AV Helios code,⁸² which couples fluid dynamics and structural dynamics for complex unsteady aerodynamics phenomena.

Helios reproduced the observed aerodynamic problems with the ACRB and was applied to explore mitigation strategies for the high control system loads. The

ability to simulate large numbers of potential rotor and control system designs with Helios and the HPCMP computer systems resulted in significant cost savings compared to flight testing various rotor and control system designs. The 150 full-vehicle simulations and 20 million DSRC core hours enabled engineering decisions and Helios continued to provide critical support for the CH-47F Block II program in 2018 and 2019. The hours were provided as part of a Frontier Project for Future Vertical Lift. COL Gregory S Fortier, Program Executive Office, Aviation, Cargo Helicopters Project Office stated, “Without Helios, the PMO would have spent at least \$50M and months of building and testing each possible design iteration using legacy processes.”⁸³

8.3 Aircraft Stores Carriage and Certification

Every service maintains a process and tools for ensuring aircraft-store compatibility. The process establishes safe and acceptable carriage and release limits, loading and unloading procedures, safe escape parameters, and ballistic accuracy. The process is also applied after modifications that alter the aircraft or store aerodynamic or structural characteristics, or the ejection characteristics of suspension equipment. Examples of stores include weapons, deployable countermeasures, suspension equipment, tanks, and pods carried internally or externally.

Historically, store certification was accomplished through a combination of ground test and flight test, with little feedback from flight test to ground test to improve models and scaling uncertainty. It was an expensive and time-consuming process that put personnel and equipment at risk. Computational modeling coupled with HPC have improved all aspects of store certification, from performing virtual fit checks for the store to obtaining the data needed to verify safe and acceptable aircraft-store compatibility. The HPCMP early recognized the potential to improve this process and committed resources for its support through HPC projects, a CHSSI portfolio of projects, DHPIs for dedicated resources to meet test deadlines, and the Institute for HPC Applications to Air Armament to help coordinate Air Force, Navy and Army databases, software development, models, and processes. HPC plus M&S are used today to provide knowledge of aircraft and store loads, aircraft flutter, vibration, aircraft stability and control, and store separation analysis, and are applied for ground and flight test support.

Dr Edward Kraft, formerly of the Arnold Engineering Development Complex at Arnold AFB, Tennessee, recalls⁸⁴:

The early application of HPC resources to weapons integration by integrating CFD and testing was, in my opinion, a transformational change in the industry enabled

by HPC resources. In the mid 90's we were able to save the F-22 \$8M in wind tunnel testing to certify the safe carriage and release of external fuel tanks and ferry missiles.

Dr Kraft describes the Air Force integrated T&E process in Kraft.⁸⁵ The current state of the practice as embodied in CREATE-AV enables simulation of a maneuvering aircraft releasing stores, requiring coupled fluid mechanics, structural mechanics and flight controls with 6 degree-of-freedom calculations for the relative motion of the stores.

8.4 Armor

During Operation Iraqi Freedom, improvised explosive device (IED) explosively formed penetrators quickly became a very lethal threat to Army platforms and soldiers. Between 2003 and 2007 there was an average of 27.3 fatalities per month caused by IED attacks in Iraq. A solution to this threat had to be developed very rapidly and ARL worked closely with forward operational assessment team scientists to understand the problem. Working cooperatively with multiple organizations, ARL rapidly developed a solution known as Interim Frag Kit 6 (IFK6) to provide protection from this threat for the High-Mobility Multipurpose Wheeled Vehicle (HMMWV) platform. Critical to the process was the use of the HPC capability at the ARL DSRC.

The nature of the threat required a minimum of 10 experiments to fully understand the capability of a potential armor solution. Instead of 10 time-consuming and expensive experiments, HPC allowed scientists to conduct a single simulation to screen a candidate configuration, saving valuable time and money and conserving resources (materials, threats, range time, etc.). HPC enabled scientists at ARL to conceptualize candidate armor designs and rapidly examine armor mechanisms, screen design options, and down-select a final armor configuration. Subsequent experimentation validated the armor configuration and also validated the methodology of closely coupling HPC with experimentation to enable the rapid development of certain armor technologies.

IFK6 went from HPC conceptualization to fielding in nominally 4 months, reducing fatalities by 20%, Fig. 5. This work spawned explosively formed penetrator armor development for Mine-Resistant Ambush-Protected (MRAP) vehicles and the subsequent MRAP Armor Weight Reduction Spiral Program. HPC was heavily leveraged in these programs to rapidly guide development of armor solutions based on IFK6 within very tight program fielding time lines. HPC enabled the fielding of explosively formed penetrator armor on more than 8,500 MRAP vehicles along with numerous route clearance platforms.



Fig. 1 Appreciation from theater

9. Modernizing the HPCMP

The Greek philosopher, Heraclitus, who lived around 500 BC, is quoted as saying “change is the only constant in life” and the HPCMP is no exception. Technology changes, personnel change, requirements change, threats change, organizations change, and the HPCMP responds accordingly, adapting to current needs. Part of change often entails losing things near and dear to some to make way for modernization. The HPCMP is transforming itself to meet ongoing requirements for compute-intensive resources while accommodating the demand for data analytics, machine learning and other artificial intelligence tasks, and integrating cloud computing into the overall portfolio.

The DOD is challenged with balancing design, development, delivery, and sustainment of complex systems with rapidly changing operational and threat environments, restricted budgets, and aggressive schedules. Years of acquisition reform have not solved the problems. Being more responsive and agile in the software and hardware acquisition lifecycle is a requirement, especially with the time lines related to artificial intelligence and machine learning, which can be measured in days or weeks, not months and years. The DevSecOps process is intended to address the issues, from cradle to grave, fully integrating security testing into the continuous integration and continuous delivery pipelines.

The HPCMP does not perform operational functions and its networks do not transport operational traffic. The HPCMP does, however, directly support technology development that can be deployed operationally and as a consequence

plays a role in the development and security stages of the DevSecOps process. Specifically, it can provide 1) a data environment, for ingesting and storing raw, reduced, derived and cleaned data products; 2) a development environment consisting of a collection of development tools (machine learning, statistical analysis, graph analysis, visualization, etc.), large-scale computing resources, processes for managing/orchestrating development, and a cyber and security toolbox; and 3) processes for continuous integration and testing with interfaces to the data environment and operational environment.

Digital engineering is the integrated DOD digital approach that uses authoritative sources of system data and models as a continuum across disciplines to support lifecycle activities from concept through disposal. The DOD approach securely and safely connects people, processes, data, and capabilities across an end-to-end digital enterprise, enabling the use of models throughout the lifecycle to digitally represent the system of interest.⁸⁶ Digital engineering is in principle a means for implementing the DevSecOps process using the continuum of data sources and models, and using the physics-based models to inform DevSecOps. The HPCMP large-scale computing resources, software tools, expertise and connectivity, and its roots in science and technology, test and evaluation, and acquisition engineering, ideally position the HPCMP for being a key element of DOD digital engineering.

Physics-based M&S is the foundation of the HPCMP and has remained its core; however, recent developments are evidence of change even here. For example, machine learning is now applied to physics-based M&S and uncertainty quantification⁸⁷; for shape optimization coupled with high fidelity CFD codes⁸⁸; for using governing field equations to guide discovery in sparse and noisy fluid flow data⁸⁹; and for improving turbulence modeling in Reynolds-averaged Navier-Stokes calculations.⁹⁰ The skills and software stack for these applications are different than for traditional compute-intensive applications, but the program is responding—for example, the Liquid Computing reconfigurable platforms delivered in late calendar 2020.

The cloud is rapidly becoming a globally pervasive resource offering scalable infrastructure-as-a-service and software-as-a-service, more recently holding promise as an alternative, or complement, to on premise HPC capabilities. The HPCMP has studied cloud computing for 9 years, examining the types of jobs suitable relative to the HPCMP workload, cost, performance, response time, and sustainment. Already some of the workload is offloaded to the cloud, replacing a previously maintained open system accessible to academic and other researchers who do not have security clearances. The challenge for expanded use will be scheduling jobs between local and cloud resources while balancing cost and schedule for a given level of performance.⁹¹

Containers are a way of packaging executable code and the corresponding source code, as well as any dependencies including software libraries, supporting data, a web server, virtual networking, virtual file systems, and even a virtual operating system including system libraries and environment variables needed to configure the software or execution environment. Containers virtualize the operating system but share the Linux kernel that is running on the Linux host. Containers provide portability between operating systems and host platforms, especially significant for integrating on premise and cloud resources.

These are some of the directions in which the HPCMP is moving, representing a very different program than created in December 1991. Yet its mission and dedication to the Warfighter remain the same: accelerate technology development and transition into superior defense capabilities through the strategic application of HPC, networking and computational expertise.

10. Conclusions

Aram Kevorkian, coauthor of the Navy HPC strategic plan, member of the original MSRC source selection team, and long-time chair of the Baseline Configuration Team, reminisces⁹²:

The annual HPCMP User Group Conferences (UGCs) were by far one of the best organized and exceptionally valuable HPCMP events of the year. HPCMP UGCs offered high quality plenary sessions, technical sessions, birds-of-a-feather sessions, user forums, and most importantly the UGCs brought together DoD users from all three services to talk about and share invaluable experiences in the use of supercomputing for the benefit of DoD. In the same spirit, the annual PET Technical Reviews were also very valuable in the way they brought different CTA leads together to talk about the technologies and software driving each of the CTAs and thus creating opportunities for collaborations among the CTA leads and their respective communities. All these are alas gone. We need these venues to create the passion for high end computing within the DoD.

The UGCs and PET Technical Reviews fell victim to the government-wide travel and conference attendance restrictions in 2012. HPCMP leadership is working to recreate the UGCs, having supported individual service user group meetings over the past few years.

Another loss around 2012 was the requirements site visits, which had been made annually for many years. The site visits enabled HPCMP leadership to meet existing users face-to-face and to understand, firsthand, current and upcoming requirements. The visits also provided a direct personal connection between the HPCMP and all

levels of the user organizations. HPCAP members were invited on the visits enabling us additional contact with our dispersed organizations, and insight into the requirements of other services and agencies. The annual requirements collection and validation process was greatly enhanced by these site visits.

As noted in the Introduction, after nearly 30 years, the program still executes its mission to modernize the supercomputer capability of DOD laboratories and test centers. As of the end of September 2020, there are nearly 2700 unique users executing about 800 projects. The top 20 projects, as measured by utilization, consume about one-half of the total available core-hours. Of these 20 projects, 9 include Frontier Projects accounting for 100 of the users. Frontier Project users are recognized as power users. The top 20 projects are aligned with 7 of the now 12 CTAs (space and astrophysical science and data and decision analytics were added more recently), a trend dating to near the beginning of the program, with CFD and computational structural mechanics consistently accounting for the largest utilization among the CTAs.

I will conclude with one final quote from French writer Jean-Baptiste Alphonse Karr: “The more things change, the more they remain the same.” This statement is vividly embodied in an observation from Bob Reschly⁴²:

I have email saved back to the mid-80's. I measure my storage requirements in ‘spindles’ rather than bytes. I usually claim I need 10 spindles worth, and when you look at the cost of storage, 10 spindles worth is typically less than \$3000 regardless of timeframe (since 2000 AD), and the storage capacity per spindle just keeps increasing. So these days, 60TB of storage (6TB/spindle for midrange capacity) is trivial and my entire email history is just a drop in that bucket. Of course convincing others to think in these terms is always a challenge.

11. References

1. Public Law 102-190. National defense authorization act for fiscal years 1992 and 1993. 102nd Congress, 1991 Dec 5.
2. Public Law 102-194. High-performance computing act of 1991. 102nd Congress, 1991 Dec 9.
3. Office of Science and Technology Policy. Grand challenges: high performance computing and communications. The FY 1992 U.S. Research and Development Program. Executive Office of the President; 1991.
4. Office of Science and Technology Policy. Grand challenges 1993: high performance computing and communications. A report by the committee on physical, mathematical, and engineering sciences. The FY 1993 U.S. Research and Development Program. Executive Office of the President; 1992.
5. Office of Science and Technology Policy. High performance computing and communications: toward a national information infrastructure. A report by the committee on physical, mathematical, and engineering sciences, Federal Coordinating Council for Science, Engineering, and Technology; 1994.
6. Public Law 105-305. Next generation internet research act of 1998. 105th Congress; 1998 Oct 28 [accessed 2020 Apr 6]. <https://www.govtrack.us/congress/bills/105/hr3332>.
7. Public Law 110-69. America COMPETES act. 110th Congress; 2007 Aug 9 [accessed 2020 Apr 7]. <https://www.congress.gov/110/plaws/publ69/PLAW-110publ69.pdf>.
8. Public Law 114-329, Sec. 105. Networking and information technology research and development modernization act of 2016. 114th Congress; 2017 Jan 6 [accessed 2019 Apr 24]. <https://www.congress.gov/114/plaws/publ329/PLAW-114publ329.pdf>.
9. Cragon HG, Watson W. The TI advanced scientific computer. Computer. 1989 Jan;22(2):55-64.
10. Purcell CJ. The control data STAR-100-Performance measurements. In: 1974 Proceeding National Computer Conference and Exposition, Association for Computing Machinery; 1974 [accessed 2020 Mar 16]. <https://pdfs.semanticscholar.org/e575/4e59c16fc448802d454d608d17b55bfb3007.pdf>.

11. Russell RM. The CRAY- 1 computer system. *Communications of the ACM*. 1978;21(1):63–72.
12. Kuck DJ, Stokes RA. The Burroughs scientific processor (BSP). *IEEE Transactions on Computers*. 1982;C-31(5):363–376.
13. Barnes GH, Brown RM, Kato M, Kuck DJ, Slotnick DL, Stokes RA. The ILLIAC IV computer. *IEEE Transactions on Computers*. 1968;C-17(8):758–770.
14. Barton JM, Namburu R. 70 years of scientific computing in the Army. *The ITEA Journal*. 2017;38(3):196–1207.
15. National Science Foundation. Prospectus for computational physics. Report by the subcommittee on computational facilities for theoretical research. 1981 Mar 15.
16. National Science Foundation. Report of the panel on large-scale computing in science and engineering. 1982 Dec [accessed 2019 Feb 5]. http://www.pnl.gov/scales/docs/lax_report1982.pdf.
17. Leiner BM, Cerf VG, Clark DD, Kahn RE, Kleinrock L, Lynch DC, Postel J, Roberts LG, Wolff S. Brief history of the internet. 1997 [accessed 2020 Mar 31]. <https://www.internetsociety.org/internet/history-internet/brief-history-internet/>.
18. Comer D. The computer science research network CSNET: a history and status report. *Communications of the ACM*. 1983;26(10):747–753.
19. Denning PJ, Hearn A, Kern CW. History and overview of CSNET. *Proceedings of the Symposium on Communications Architectures & Protocols*. 13. SIGCOMM, Association for Computing Machinery; 1983 Apr.
20. National Research Council. Review of NASA's numerical aerodynamic simulation program. Committee on NASA Scientific and Technological Program Reviews, Commission on Engineering and Technical Systems, National Academy Press; 1984 Mar. Report No.: PB84-194919.
21. Bischof C, Simon HD. The implementation of the U.S. high performance computing and communications program. In: Meuer HW, editor. *Supercomputer '92*. Informatik aktuell. Springer; 1992.

22. National Aeronautics and Space Administration. NAS Division 2008: 25 Years of Innovation. 2008. [accessed 2020 Mar 16]. https://history.arc.nasa.gov/hist_pdfs/riacs_25anniversary_bklt-tmln_2008.pdf.
23. Forslund, DW, Slocomb CA, Agins IA. Windows on computing: new initiatives at Los Alamos. Los Alamos Science. 1994;22.
24. Carr AB. Punched cards to petaflops. National Security Science. 2013 Apr;34–41.
25. Larzelere II AR. Delivering insight: the history of the accelerated strategic computing initiative (ASCI). Lawrence Livermore National Laboratory, UCRL-TR-231286; 2009 Sep [accessed 2019 Feb 28]. https://asc.llnl.gov/asc_history/.
26. Bashor J, Conner T, editors. Department of Energy, ASCR@40: Highlights and Impacts of ASCR's Programs. Department of Energy, Office of Advanced Scientific Computing Research; 2020.
27. U.S. Department of Energy. Future high performance computing capabilities. Summary report of the advanced scientific computing advisory committee (ASCAC) subcommittee; 2019 Mar 26 [accessed 2020 Jan 6]. <https://www.osti.gov/servlets/purl/1570693>.
28. Bardon M, Curtis K. A national computing environment for academic research. National Science Foundation; 1983 July.
29. Defense Advanced Research Projects Agency. Strategic computing. New-generation computing technology: a strategic plan for its development and application to critical problems in defense. 1983 Oct 28.
30. Roland A, Shiman P. Strategic computing: DARPA and the quest for machine intelligence, 1983-1993. MIT Press; 2002.
31. Public Law 99-383. National science foundation authorization act for fiscal year 1987. 114th Congress; 1986 Aug 21 [accessed 2019 Mar 14]. <https://www.govtrack.us/congress/bills/99/hr4184>.
32. Federal Coordinating Council for Science, Engineering and Technology. A report to the office of science and technology policy on computer networks to support research in the United States, a study of critical problems and future options. Volume I, recommendations. 1987 Nov 20.

33. Federal Coordinating Council for Science, Engineering and Technology. A report to the office of science and technology policy on computer networks to support research in the United States, a study of critical problems and future options. Volume II, workshop summaries. 1987 Nov 20.
34. Federal Coordinating Council for Science, Engineering and Technology. A report to the office of science and technology policy on computer networks to support research in the United States, a study of critical problems and future options. Volume III, a compendium of supporting technical data. Preliminary Issue, 1987 June.
35. Federal Coordinating Council for Science, Engineering and Technology. Report to the Federal coordinating council for science, engineering, and technology supercomputer panel on recommended government action to retain U.S. leadership in supercomputers. Office of Science and Technology Policy; 1983 Jan.
36. National Research Council. Report of the research briefing panel on computer architecture. In Research Briefings 1984. National Academies Press; 1984.
37. Executive Office of the President. A research and development strategy for high performance computing. Office of Science and Technology Policy; 1987 Nov 20.
38. Executive Office of the President. The federal high performance computing program. Office of Science and Technology Policy; 1989 Sep 8.
39. National Research Council. The national challenge in computer science and technology. National Academy Press; 1988 [accessed 2019 Feb 28]. <https://www.nap.edu/catalog/10331/the-national-challenge-in-computer-science-and-technology>.
40. U.S. Congress. High performance computing and networking for science – background paper. Office of Technology Assessment, OTA-BP-CIT-59, US Government Printing Office; 1989 Sep.
41. Bergin TJ, editor. 50 years of army computing, from ENIAC to MSRC. Army Research Laboratory (US); 2000 Sep. Report No.: ARL-SR-93.
42. Reschly B. In-person interview and follow-up e-mail correspondence, 2019 Sep 18.
43. Frasier JT. Army supercomputer program chronology and fiscal history. BRL Director's Point Paper to BRL; 1991 May 20.
44. Breaux H. Response to interview questionnaire, 2019 Mar.

45. Tezduyar T. Army high performance computing research center, final research program report for the period August 25, 1989 - January 8, 1995. University of Minnesota/AHPCRC; 1995. #ARO 27747.504-MACOE.
46. Rhoades C. In-person interview, response to interview questionnaire, and follow-up e-mail correspondence, 2018 Aug 29.
47. Graham J. In-person interview and follow-up e-mail correspondence, 2018 Sep.
48. Department of the Navy. Supercomputing in the 90's, a Navy Strategic Plan. Information Resource Management, Supercomputer Working Group; 1990 May.
49. Kevorkian A, Parnell L. DoD high performance computing modernization program nurtures skills at SSC San Diego. Outlook: Space and Naval Warfare Systems Center San Diego; 2000 Oct 27. p. 1–3.
50. Department of Defense. Supercomputer acquisition master plan. 1990 Oct 1.
51. Davis L. In-person interview and follow-up e-mail correspondence, 2018 Aug 29.
52. McDonald C. Video interview, 2018 Aug 29.
53. Department of Defense. DOD high performance computing modernization plan. 1992 March 31.
54. Davis L. In-person interview and follow-up e-mail correspondence, 2018 Oct 2.
55. Department of Defense. DOD implementation plan for high performance computing modernization plan. 1992 Dec.
56. Miller V. Telephone interview and follow-up e-mail correspondence, 2018 Oct 18.
57. Breaux H. 18 November DoD high performance computing working group briefing to Dr Anita Jones, DDR&E. Information Paper to George T Singley, III, Deputy Assistant Secretary of the Army for Research and Technology. 1993 Nov 23.
58. Davis L. In-person interview and follow-up e-mail correspondence, 2018 Oct 11.
59. MPI Forum. [accessed May 11, 2020]. <https://www.mpi-forum.org/docs/>.

60. Department of Defense. High performance computing modernization program requirements analysis. 1994 Apr.
61. Department of Defense. Performance computing modernization program addendum to the requirements analysis of April 1994, to support the acquisition of systems integration services for the DoD high performance computing major shared resource centers. 1994 May.
62. Department of Defense. High performance computing modernization program analysis of alternatives. 1994 May.
63. United States General Accounting Office. High performance computing and communications, new program direction would benefit from a more focused effort. House of Representatives, GAO Report to the Chairman, Committee on Armed Services, GAO/AIMD-95-6, 1994 Nov [accessed 2019 Mar 12]. <https://www.gao.gov/archive/1995/ai95006.pdf>.
64. Davis L. E-mail correspondence, 2020 Oct 28.
65. Highland J. Telephone interview and follow-up e-mail correspondence, 2018 Oct 9.
66. Public Law 104-61. Department of Defense Appropriations Act, 1996. 104th Congress; 1995 Dec 1.
67. Reschly B. E-mail correspondence, 2020 Oct 8.
68. Barton JM, Brainard LA, Campbell RL. Modernizing test and evaluation with high performance computing and advanced networks. *The ITEA Journal*. 2018;39(3):162–176.
69. Ahalt S, Krishnamurthy A, To V. Telephone interview and follow-up e-mail correspondence, 2020 Sep 3.
70. Henry C. E-mail correspondence, 2020 Oct 28.
71. Kraft EM. E-mail correspondence, 2020 Oct 27.
72. Post DE. E-mail correspondence, 2020 Oct 27.
73. Post DE, Atwood CA, Newmeyer KP, Meakin RL, Hurwitz MM, Dey S, D'Angelo JN, Vogel song RL, Hariharan NS, Kendall RP, et al. CREATE: Software engineering applications for the design and analysis of air vehicles, naval vessels, and radio frequency antennas. *Computing in Science & Engineering*. 2016 January/February;18(1)14–24.

74. Department of Defense. Determining the value to the warfighter, a three-year return on investment (ROI) study. High Performance Computing Modernization Program; 2011.
75. Henry C. Response to interview questionnaire and follow-up e-mail correspondence, 2018 Oct 23.
76. West J. Response to interview questionnaire and follow-up e-mail correspondence, 2019 Feb 6.
77. Schwartz D. E-mail correspondence, 2021 Mar 18.
78. Gomez RJ, Vicker D, Rogers SE, Aftosmis MJ, Chan WM, Meakin R, Murman S. STS-107 investigation ascent CFD support. 34th AIAA Fluid Dynamics Conference; 2004 28 June–1 July. AIAA Paper 2004-2226.
79. Wagner BJ. Answering the questions: ARL MSRC provides support the Columbia space shuttle accident investigation. ARL MSRC Link; Spring 2003.
80. National Aeronautics and Space Administration. Columbia accident investigation board report vol. 1. U.S. Government Printing Office; 2003 Aug.
81. National Aeronautics and Space Administration. Columbia accident investigation board report vol. 2. U.S. Government Printing Office; 2003 Oct.
82. Hariharan N, Wissink A, Potsdam M, Strawn R. First-principles physics-based rotorcraft flowfield simulation using HPCMP CREATE-AV Helios. *Computing in Science & Engineering*. 2016;18(6):19–26.
83. Department of the Army. Endorsement of Aviation & Missile Research, Development and Engineering Center (AMRDEC) Helicopter Modeling Software (HELIOS). Letter from COL Gregory S Fortier, Program Executive Office, Aviation, Cargo Helicopters Project Office to Aviation & Missile Research, Development And Engineering Center (AMRDEC), 2018 Jan 22.
84. Kraft EM. E-mail correspondence, 2020 June 18.
85. Kraft EM. Integrated test and evaluation - a knowledge-based approach to system development. 1st AIAA Aircraft Engineering, Technology, and Operations Congress; 1995 Sep 19–21. AIAA Paper 95-3982.
86. Department of Defense. Digital engineering strategy. Office of the Deputy Assistant Secretary of Defense for Systems Engineering; 2018 June.
87. Torrea E, Marelli S, Embrechts P, Sudret B. Data-driven polynomial chaos expansion for machine learning regression. *Journal of Computational Physics*. 2019;388:601–623.

88. Lyea KO, Mishraa S, Ray D. Deep learning observables in computational fluid dynamics. *Journal of Computational Physics*. 2020;410:109339.
89. Raissi M; Karniadakis GE. Hidden physics models: machine learning of nonlinear partial differential equations. *Journal of Computational Physics*. 2018;357:125–141.
90. Zhao Y, Akolekar HD, Weatheritt J, Michelassi V, Sandberg RD. RANS turbulence model development using CFD-driven machine learning. *Journal of Computational Physics*. 2020;411:109413.
91. von Alfthan S. High-performance computing in the cloud? Computer Science Corporation; 2017.
92. Kevorkian A. In-person interview and follow-up e-mail correspondence, 2018 Sep 27.

List of Symbols, Abbreviations, and Acronyms

ACRB	Advanced Chinook Rotor Blade
AEDC	Arnold Engineering Development Center
AFB	Air Force Base
AFRL	Air Force Research Laboratory
AFSC	Air Force Systems Command
AFSNet	Air Force Supercomputer Network
AHPCRC	Army HPC Research Center
APBI	Advanced Planning Briefing for Industry
ARL	US Army Research Laboratory
ARPA	Advanced Research Projects Agency
ARPANET	ARPA Network
ASC	Aeronautical Systems Center
ASCI	Accelerated Strategic Computing Initiative
ASNet	Army Supercomputer Network
BRL	Ballistic Research Laboratory
BRLESC II	BRL Electronic Scientific Computer II
BSP	Burroughs Scientific Processor
CDC	Control Data Corporation
CEWES	Corps of Engineers Waterways Experiment Station
CFD	Computational Fluid Dynamics
CHSSI	Common HPC Software Support Initiative
COMPETES	Creating Opportunities to Meaningfully Promote Excellence in Technology, Education, and Science
CREATE	Computational Research and Engineering Acquisition Tools and Environments
CSNET	Computer Science Research Network

CTA	Computational Technology Area
DARPA	Defense Advanced Research Projects Agency
DDR&E	Director, Defense Research and Engineering
DEVCOM	US Army Combat Capabilities Development Command
DHPIs	Dedicated HPC Project Investments
DISA	Defense Information Systems Agency
DNA	Defense Nuclear Agency
DOD	Department of Defense
DOE	Department of Energy
DREN	Defense Research and Engineering Network
DSRC	DOD Supercomputing Resource Center
DTRA	Defense Threat Reduction Agency
ENIAC	Electronic Numerical Integrator and Computer
ERDC	Engineer Research and Development Center
FCCSET	Federal Coordinating Council for Science, Engineering and Technology
FY	fiscal year
GAO	US Government Accountability Office
HEP	Heterogeneous Element Processor
HIP	High-Performance Computing Internship Program
HMMWV	High-Mobility Multipurpose Wheeled Vehicle
HPC	high-performance computing
HPCAP	HPC Advisory Panel
HPCCI	High-Performance Computing and Communications Initiative
HPCMP	High-Performance Computing Modernization Program
HPCWG	HPC Working Group
HPTi	High Performance Technologies, Inc.

IDA	Institute for Defense Analyses
I-DREN	Interim Defense Research and Engineering Network
IED	improvised explosive device
IFK6	Interim Frag Kit 6
IPA	Intergovernmental Personnel Act
JEOM	Joint Educational Opportunities for Minorities
LANL	Los Alamos National Laboratory
M&S	modeling and simulation
MAIS	Major Automated Information System
MAISRC	MAIS Review Council
MDAP	Major Defense Acquisition Program
MHPCC	Maui High Performance Computing Center
MRAP	Mine Resistant Ambush Protected
MSAP	Mature Systems Advisory Panel
MSRC	Major Shared Resource Center
MSU	Mississippi State University
NASA	National Aeronautics and Space Administration
NAVO	Naval Oceanographic Office
NCHPC	National Consortium for HPC
NCO	National Coordination Office
NDSEG	National Defense Science and Engineering Graduate
NITRD	Network and Information Technology Research and Development
NREN	National Research and Education Network
NRL	Naval Research Laboratory
NSF	National Science Foundation
NSFNET	NSF-developed network capabilities
OSC	Ohio Supercomputer Center

OSD	Office of the Secretary of Defense
OSTP	Office of Science and Technology Policy
OSU	Ohio State University
PET	Programming Environments and Training and User Productivity Enhancement and Training
PETT	User Productivity Enhancement and Technology Transfer
PETTT	User Productivity Enhancement, Technology Transfer and Training
PM	Program Manager
PMO	Project Office
POM	Program Objective Memorandum
RENCI	Renaissance Computing Institute
RFP	request for proposal
ROI	return on investment
S/AAA	Service/Agency Approval Authority
SASI	Scalable Applications Software Initiative
SCOUT	SuperComputing OUTpost
SDP	service delivery point
SIP	signal and image processing
SRC	Shared Resource Center
SRCAP	Shared Resource Center Advisory Panel
STWG	Science and Technology Working Group
TACC	Texas Advanced Computing Center
TACOM	Tank-automotive & Armaments Command
T&E	test and evaluation
UGC	User Group Conference
WAN	wide-area network
WSMR	White Sands Missile Range

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 DEVCOM ARL
(PDF) FCDD RLD DCI
TECH LIB

1 DEVCOM ARL
(PDF) FCDD RLC S
JM BARTON