

The Use of High Performance Computing (HPC) to Strengthen the Development of Army Systems

John W. Lyons

Richard Chait

Charles J. Nietubicz

Center for Technology and National Security Policy
National Defense University

November 2011

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the 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 of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 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 a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE NOV 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011	
4. TITLE AND SUBTITLE The Use of High Performance Computing (HPC) to Strengthen the Development of Army Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) National Defense University, Center for Technology and National Security Policy, 300 Fifth Avenue, Fort Lesley J. McNair, Washington, DC, 20319				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 25	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

The views expressed in this article are those of the authors and do not reflect the official policy or position of the National Defense University, the Department of Defense, or the U.S. Government. All information and sources for this paper were drawn from unclassified materials.

John W. Lyons is a Distinguished Research Fellow at the Center for Technology and National Security Policy (CTNSP), National Defense University. He was previously director of the Army Research Laboratory and director of the National Institute of Standards and Technology. Dr. Lyons received his PhD from Washington University. He holds a BA from Harvard.

Richard Chait is a Distinguished Research Fellow at CTNSP. He was previously Chief Scientist, Army Material Command, and Director, Army Research and Laboratory Management. Dr. Chait received his PhD in Solid State Science from Syracuse University and a BS degree from Rensselaer Polytechnic Institute.

Charles J. Nietubicz is a Research Leader in High Performance Computing at Battelle Memorial Institute. He was previously with the Army Research Laboratory (ARL) serving as the Chief of the Advanced Computing and Computational Sciences Division and Director of the ARL Department of Defense Supercomputing Resource Center. He holds BS and MME degrees in mechanical engineering from the University of Dayton.

Acknowledgements

The authors wish to acknowledge the support of Dr. Marilyn Freeman, Deputy Assistant Secretary of the Army for Research and Technology and Dr. Linton Wells, Director of the Center for Technology and National Security Policy at the National Defense University. We are grateful to the reviewers: Dr. Douglass Post of the High Performance Management Office of the Department of Defense, Mr. Loren Miller, DataMetric Innovations, LLC, and Dr. John Hurley, Information and Resource Management College of the National Defense University.

Table of Contents

I. Introduction	1
II. The Computer Revolution and HPC	4
III. Army Use of Computing in R&D and Systems Development	7
IV. Advancing the Industrial Use of HPC and Some Case Studies.....	12
V. Some Government Programs to Advance the Use of Computers in Design and Manufacturing.....	14
VI. Summary and Recommendations	19
Appendix A. More Details on the Computational Aerodynamics Branch at the Army Research Laboratory	22

I. Introduction

The long, drawn-out and expensive acquisition process in the military is a serious problem at any time, but especially so in these times of tight and declining budgets. Many times the Department of Defense (DOD) has been criticized for the manner by which it procures new weapons systems. Recent articles in the Washington Post^{1,2} and Defense News³ review the Army acquisition process, presenting an analysis of some \$38 billion spent on systems that were terminated for high costs and lengthy development times. Two cases in point that failed: the Crusader artillery system which foundered, in part, on the failure of a new propellant system, high costs, and a too-long development cycle, and the Comanche stealth helicopter, again terminated because of cost overruns and slippage of schedules. The available funding for the Comanche was redirected into buying and upgrading existing helicopter programs where the procurement risks were reduced. A question is how much of this problem is due to technical inefficiencies, and how much is due to other factors such as the requirements process, Defense acquisitions regulations, and the complexity of the lengthy budgeting process or changes in the governing political system. We seek to show how one factor, the expanded use of high performance computing, can contribute to improving the design and production of weapons systems. This would contribute to a more responsive, more economical acquisition process.

In the dictionary, the term “manufacturing” means the production process of converting raw materials, labor, and capital into finished products. However, in the discussions of the use of high performance computing, “manufacturing” has been applied more broadly to the part of the innovation process that includes both developing the design of the product in research and development and production on the factory floor. HPC has been used, on occasion, by the military and some large companies to speed the design process by using

“High performance computing, HPC” generally refers to the use of the latest, most powerful supercomputers or clusters of computers to solve the largest and most demanding computational problems. In the past, the use of HPC has been mostly in research at the frontiers of certain fields of physics, chemistry, materials, and engineering. More recently HPC has been used in many other aspects of technology, such as modeling for design, for managing complex production processes, for handling large data sets, and so on.

computer modeling and computer experimentation in place of the much longer and more expensive traditional process of building prototypes, testing, redesigning, building more prototypes, and more testing, until the design is final. While the Army has used HPC for some of its work, the use has historically been modeling the design of devices and components, for example, munitions (see the history provided in Chapter III). We note also that in the Army, the design process is often shared with contractors and production is usually done in the private sector. This requires close collaboration between the Army and its suppliers.

The wars in Iraq and Afghanistan focused the need for rapid fielding of systems not based on radically new technologies. (Even so, HPC was, and still is, useful in such cases as well as in the design of components and upgrades of existing designs.) Thus, the Stryker medium-armored combat vehicle and the mine-resistant, ambushed-protected (MRAP) vehicle were developed in this way – the Stryker to provide more survivability than the high-mobility, multipurpose

¹ Marjory Censer, “Go Big or Go to War with the Weapons You Have,” The Washington Post, May 30, 2011.

² Greg Jaffe, “A Soldier Who Stays Focused on Realities on the Ground,” The Washington Post, May 30, 2011.

³ Michael Hoffman, “U.S. Army’s New Mission: Fix Acquisition”, Defense News, 26, No. 38, October 17, 2011.

wheeled vehicle (HUMMV) but less than the main battle tanks and the MRAP to protect against the improvised explosive devices (IEDs) that were wreaking havoc in Iraq and in Afghanistan. The Stryker vehicle weighs about 18 tons and has a unit cost considerably less than the Abrams tank. The Stryker is based on an existing General Dynamics Canada vehicle. Initially called for by the Army Chief of Staff in 1999, the first fielding was in 2003. The MRAP is based on an earlier design from the late 1970s. The U.S. program was initiated in 2006 and first deliveries were taken in 2007. The MRAP has a specially designed V-shaped bottom to deflect the blast and fragments from IEDs. The vehicle is large (14-25 tons) and costs from \$500,000 to \$1M. It is difficult to handle on narrow dirt roads and has a tendency to roll over in certain situations (high center of gravity). The Pentagon has ordered more than 10,000. The MRAPs look like trucks and are basically replacement for the HMMWV and various military light trucks. Both the Stryker and the MRAP are built largely from off-the-shelf technology and are based on previously vetted designs.

Recent articles (see footnotes 1 and 2) of the views of the Secretary of Defense and the incoming Chairman of the Joint Chiefs of Staff suggest that emphasis in future procurements is likely to be on improvements to successful existing systems rather than trying to build new and more complex systems. There are several factors that contribute. One is the requirements process itself; it is subject to continuing changes in what the warfighter wants – in the middle of an acquisition cycle such changes create havoc in terms of delays, recycling of the research and development effort, and attendant cost increases. Another factor is the nature of the research and development (R&D) process as it has been traditionally conducted. As noted above, experimental research leads to building a prototype that is then subjected to developmental testing. These take time and money -- sometimes significant amounts of both. The bugs revealed in the testing are corrected in further R&D to produce a second prototype for testing, and so on through many cycles, each of which is time-consuming and expensive. Subsequently there may be design competitions for the private sector, competitions that may cause the military to alter the requirements. There may also be interventions by the Congress that change the basic assumptions of the program. The combination of changing requirements and repetitive cycles of laboratory experiments, prototype building, developmental testing, and external factors has led to very long acquisition times.

An example is the Abrams main battle tank.⁴ In 1963, the Army mounted a joint program with the Federal Republic of Germany to create a new main battle tank, but the initial estimates of cost were deemed too expensive and the work was terminated. In 1971, a new set of requirements was established. Only in 1980 was the first Abrams, the M1, produced. It took nine years from the reset and seventeen years from the original first efforts. Such long time lags cannot be tolerated – the product, whether an entirely new system or a major upgrade to an existing one, will be obsolete before it is fielded. A better way must be found. One way is to make full use of computational science in the design of the system and in computer-controlled production or factory automation.

⁴ Richard Chait, John Lyons, and Duncan Long, *Critical Technology Events in the Development of the Abrams Tank*, Defense & Technology Paper 22 (Washington, D.C.: Center for Technology and National Security Policy, National Defense University, December 2005).

The history of computing in modeling and simulation of products and systems parallels the evolution of the modern computer. There are many kinds of models and simulations. In this paper, we focus on computer models based on the physics, chemistry, behavior of materials, and engineering, taken together in the finished product or system. Early models could simulate the behavior of components, but could not model a complete system. This was because the computers could not run such a program in any acceptable length of time. However, with developments in HPC, such modeling is now possible. With a successful computer model of the proposed product or system, one can then build prototypes on the computer and study the effects on the performance of changing the many variables. Thus, the cycle of experiment, prototype building and developmental testing can be run on the computer. Many such cycles can be run very quickly until the optimum is located. Then, this optimal design can be built and tested to verify the computed performance. In examples discussed below, this part of the acquisition cycle has been reduced roughly by a factor of three.

The military has been doing scientific modeling and simulation (M&S) for many years, largely in research and engineering as opposed to production processes. The Army has used HCP in developing weapons platforms such as the modeling of a projectile striking an Abrams tank⁵, the interaction between the rotor wash of the Apache attack helicopter and its horizontal stabilator, and various aspects of the Stinger and Javelin missile systems including flight simulation.⁶ These are just a few examples.

We took note of an Army forecasting effort, known as STAR 21 (Strategic Technology for the Army in the Twenty-first Century), published in the early 1990s.⁷ The STAR 21 study complemented a newly-issued Army Tech Base Master Plan to bring focus to Army science and technology programs. The goal for the National Academy of Sciences' STAR 21 study was "to assist the Army in improving its ability to incorporate advanced technologies into weapons, equipment, and doctrine." Key among the objectives was to identify the advanced technologies most likely to be important to ground warfare in the next century.

The technology forecast on manufacturing contained in the STAR 21 study did not predict the potential value of advanced computing in the Army acquisition process. Since the time of that study, modeling and simulation in designing new weapons and their performance has made many advances. Recent National Defense University (NDU) publications have discussed an anticipated new Army technology forecasting effort.^{8,9,10} Our reviewers felt that the predictions

⁵ Paul H. Dietz, Harry L. Reed, Jr., J. Terrence Klopick, and James N. Walbert, *Fundamentals of Ground Combat System Ballistic Vulnerability/Lethality*, American Institute of Aeronautics and Astronautics, Inc., Reston, VA, 2009.

⁶ John Lyons, Duncan Long, and Richard Chait, *Critical Technology Events in the Development of the Stinger and Javelin Missile Systems*, Defense & Technology Paper 33 (Washington, D.C.: Center for Technology and National Security Policy, National Defense University, July, 2006).

⁷ *STAR21 – Strategic Technologies for the Army of the Twenty-First Century*, Board on Army Science and Technology, (Washington, D.C.: National Academies Press, 1992).

⁸ John Lyons, Richard Chait, and Jordan Willcox, *An Assessment of the Science and Technology Predictions in the Army's STAR21 Report*, Defense & Technology Paper 50 (Washington, D.C.: Center for Technology and National Security Policy, National Defense University, July 2008).

⁹ John W. Lyons, Richard Chait, and James J. Valdes, *Forecasting Science and Technology for the Department of Defense*, Defense & Technology Paper 71 (Washington, D.C.: Center for Technology and National Security Policy, National Defense University, December 2009).

rated only a “D”, the lowest in the assessment. One intent of this paper is to re-evaluate the advanced manufacturing area in light of the many recent developments in computational capability and availability. These topics are addressed in the next chapter.

With the above in mind, we begin with a discussion of the growth and trends in HPC. We then discuss how the Army has utilized HPC in research. Included here is the Army’s exploration of the possibilities for design and production of new systems. Turning to manufacturing-related topics, we then discuss how industry has advanced the use of HPC and include several case histories. Next we present some contributions to HPC advancement in several current programs being conducted by other government agencies. The paper concludes with closing comments and recommendations for strengthening Army acquisition using advanced computer technology.

II. The Computer Revolution and HPC

Among the underlying building blocks of science and engineering are the computational tools that are needed to develop an understanding of scientific principles, to solve problems, and to make advances. While we take many capabilities for granted today, they all encountered some suspicion, trepidation, and many times required a long learning curve. Looking for the first use of a computing tool takes us to the abacus, a tool found in Babylon in 2400BC. This can be considered as the first computer. From this beginning, there has been one revolution after another, each of which advanced our computational capability. Moving from the abacus, to the early Antikythera (150-100BC) used to calculate astronomical positions, to mechanical analog devices, to the difference engine, to the analytical engine, to devices such as planimeters and nomographs, the quest has been to develop capabilities which would help with our understanding of science, physics, and engineering. If we fast-forward to the 20th and 21st centuries, the tools are more easily understood and recognizable. Many of us today can remember our days at the university when the only devices used for solving homework or test problems were analog devices, such as the slide rule, a mechanical analog computer, which is now rarely found except perhaps in museums. We have moved on to calculators, digital mainframe computers, mini-computers, personal computers, workstations, and multiple hand-held devices, but only a relatively small number of scientists and engineers, mostly in research, have realized that high performance computers are the modern-day tool of choice.

The development and use of HPC has an important and impressive legacy, especially within the Army, going back to the ENIAC (Electronic Numerical Integrator and Calculator) in 1946. Its need was born, as with many new discoveries, from DOD requirements for the need for speed in producing artillery firing tables. Continuing into today, we still have the need for speed and are fortunate to have the increasing processing power being made available to us by the likes of Intel, AMD, and supercomputer companies such as Cray, IBM, SGI, etc. Given the current trends of moving to exa-flop (10^{18} floating point operations per second) computing, which will be available around 2017 (see Figure 1), and the natural trend to have slightly less powerful systems readily available to the engineering community, the capability of physics-based

¹⁰ John W. Lyons, Richard Chait, and Simone Erchov, Editors, *Improving the Army’s Next Effort in Technology Forecasting*, Defense & Technology Paper 73 (Washington, D.C.: Center for Technology and National Security Policy, National Defense University, September 2010).

modeling and simulation is very encouraging. Basically, today's supercomputers will be tomorrow's laptops.




Figure 1. Projected Supercomputer Performance Development 2009¹¹

At the time that individual workstations and, later, personal or desktop computers were being introduced, the most advanced computers were central machines such as the Control Data Corporation's 6600 (in 1964) and 7600 (1969-78). Additionally, the creation of a packet-switched network, Advanced Research Projects Agency Network (ARPANET), in the 1960s by DARPA, changed the way computer interactions occurred. The Department of Energy created a network of T-1 bandwidth (1.5 megabits per second) to connect a few of its machines. At about this time, local area networks began to appear, mostly based on information packets circulating with collision avoidance techniques. Computers in the laboratory began as laboratory automation, for example, mass spectrometers with databases included so that the spectra would be read by the machine and the most likely candidates printed out. A little bit later the first personal computers (PCs) came on the market, mostly as curiosities. The operating systems were either ms-dos or cp/m. The machines were programmed in a simple language called "basic" and were generally difficult to use. But soon, IBM (the IBM PC came on the market in 1981) and, later, Apple offered friendlier PCs beginning with Apple's LISA (based on the GUI or graphical interface). With the advent of Windows 95, the use of PCs increased very rapidly and eventually PCs became the ubiquitous workstation we see today. Office secretaries turned in their IBM Selectrics and similar machines for PCs. With the local area networks, typed documents could be circulated for critique and review. Bench scientists began to draft papers and reports on their PCs rather than writing them by hand.

¹¹ http://www.top500.org/lists/2009/06/performance_development

Then Control Data's designer Seymour Cray formed his own company and built the CRAY series of supercomputers (the CRAY 1 appeared in 1976). Previously insoluble problems could be computed quickly (in terms of machine time) and economically. As the network spread, scientists who previously had to mail their programs and data to the few advanced computing centers, particularly at Department of Energy (DOE) sites, now could accomplish this electronically. Meanwhile, back at the laboratory very capable mini-computers such as the DEC VAX were being installed to serve one or a few labs. These machines were sufficiently user-friendly that there appeared to be no need to hire computer specialists to operate and maintain them. (This was a time when many duties now reserved to people trained in computer science and engineering were carried out by research scientists and engineers.) Now there developed a dichotomy between the need for compute power and the need for accessibility. Soon a tendency developed for research staff to hold out for a dedicated (under their control) mini-computer rather than to use a central computer (sometimes a supercomputer) either at the laboratory or on the network. The supercomputers were not always available – one had to wait in a queue. Software was difficult to transition from the local minis to the supers; the software wasn't always scalable and the languages might be different and, as an aside, it was fun to “fiddle” with your “own” system. Nonetheless, the most difficult problems justified the need for the continued development of ever more powerful supercomputers. As each generation of machines appeared, the Federal government not only funded their development and early testing, but also purchased a number of machines and made them available to all comers. DOE, the National Aeronautics and Space Administration, the National Security Agency, the National Science Foundation, and DOD were some of the major players in the development and use of these HPC systems.

The DOE has a compelling rationale for obtaining the most powerful computers, namely the need to model the design and behavior of nuclear weapons. Thus the DOE National Laboratories nearly always have had the most advanced systems. Similarly NASA's requirements for increased capability with the space race provided a need for the most advanced computer systems and computational software. Computer codes in, for example, computational fluid dynamics (CFD), were adopted by the Army research community for use in modeling various phenomena for arms and armaments systems. Collaborations between Army labs, Livermore, Los Alamos, NASA Ames and others have been very fruitful. Each new generation of supercomputers was quickly adopted by researchers as they realized that problems they needed to address, but couldn't be run on earlier machines, could now be handled. Thus new opportunities opened up in many fields. The cost and small number of such machines in the past kept the use of these systems to only a relatively small subset of researchers

DOD had also recognized in the early 1990s that high performance computing was an essential enabling technology that was facing a critical juncture in the United States. The current systems on hand were aging fast (life time of a supercomputer is < 4 years); European and Asian interest was increasing; new technology was around the corner – parallel computing and the development of the World Wide Web for collaboration. However, the Services could not, within existing budgets, independently afford the latest supercomputers. Out of this need, a Working Group with members from each Service, and directed by the DOD Director for Defense Research and Engineering (DDR&E), developed the HPC Modernization Plan in March 1992¹²

¹² Director, Defense Research & Engineering, *DoD High Performance Computing Modernization Plan* (Washington, D.C.: Office of the Secretary of Defense, March 1992).

which was then submitted to Congress. This plan called for the organization of the High Performance Computing Modernization Program with the vision that an HPC capability comprised of the most advanced systems, scientific and visualization workstations, special purposed systems and high-speed networking be built for direct support to the DOD Research, Development, Test and Evaluation (RDT&E) community. The program has made great progress¹³ and continues today.

A recent report from the National Research Council¹⁴ discusses the use of HPC (termed high-end capability computing (HECC) in the report to refer to the very latest HPC machines) at the frontiers of four scientific and engineering disciplines, pushing the limits of what can be done. The definition used in the committee's report is:

“...[HECC] means advanced computing that pushes the bounds of what is computationally feasible...whatever sort of advanced non-routine computing system is needed to push the computational science or engineering of a given field.”

By this definition, the power of the HPC computer needed will vary depending on the field of use. For some problems, such as meteorology or astrophysics, the most powerful systems may still fall short of what is needed.

The NRC report discusses the requirements for effective use of such computing power. One needs “models, algorithms, software, hardware, facilities, education and training, and a community of researchers attuned to its special needs.” In other words, it is not enough simply to have available a supercomputer, the area of science must be “ready.” Each of the fields discussed in the NRC report generates vast amounts of data, much of which must be added to large existing files to update the models in use. For example, in atmospheric science and meteorology, data are collected daily or hourly and fed into worldwide weather models to update forecasts. Doing research on the weather requires being able to include all past data. This in turn means the ability to warehouse, search, and manipulate many years worth of information. The same is true in astrophysics, which keeps adding volumes of new data from the many telescopes searching the sky. The size of these data sets requires the power of the latest in high performance computers.

III. Army Use of Computing in R&D and Systems Development

The use of HPC has been recognized by the Army's Ballistic Research Laboratory (BRL) as an absolute need and enabler for getting its mission completed since the construction of the ENIAC in 1946. The need to develop firing tables for Army artillery during World War II led to the birth of the computer age with its insatiable need for greater and greater computing power.

¹³ Office of the DoD High Performance Computing Modernization Program, *Determining the Value to the Warfighter – A Three-Year Return on Investment Study* (Washington, D.C.: Office of the Secretary of Defense, February 2009).

¹⁴ *The Potential Impact of High-End Capability Computing on Four Illustrative Fields of Science and Engineering*, National Research Council, National Academies Press, Washington, D.C., 2008.

As Army technology evolved and research and development faced even greater challenges, the use of HPC was becoming an indispensable tool. While there was limited and difficult access to supercomputers of the day (the ARPANET was in its infancy), an important missing piece was the application software. An extremely beneficial exchange program was established between the BRL and NASA Ames Research Center in Mountain View, California. While the collaboration with NASA Ames continues to today, the most active timeframe was from the late 1970s into the mid-1980s. Through this program, scientists from BRL spent from 4



Figure 2. Programming the ENIAC by moving connecting cables by hand

In 1946, using the ENIAC, it would take skilled technicians weeks to change programs, moving large cables from one area of the machine to the other. Figure 2 shows some of the first programmers, who were mostly women, getting the ENIAC ready for one of its runs. While we do not have to move cables today, we still need to know about parallel programming, Message Passing Interface, command line interfaces, scientific visualization, graphical processing units, etc. So the trend continues even today. Each new generation of supercomputers requires learning new software and techniques.

months to 1 year at NASA and worked with some of the early pioneers in Computational Fluid Dynamics. This early connection to one of the world's premier research centers jump-started the computational efforts and in-house capability by at least five years. The association lasted well into the 2000s, with the last major connection being

the NASA partnership with Stanford University in the latter's selection as the new home for the Army High Performance Computing Research Center.

Armed with developed application software and access to some of the most capable machines of the time - the CDC 7600, Cray 2 and Cray XMP - the BRL exploited this capability on projects like the transonic and supersonic magnus coefficients, base flow research, the XM829A1 silver bullet, and Sense and Destroy Armor Munitions (SADARM). Early work in developing the advanced numerical techniques (it was not yet called CFD) was begun to develop both the experimental and computational capabilities to help understand the unexpected flight stability of longer projectiles. This work was the beginning of a




Figure 3. Experimental Range at ARL

long and successful effort in bringing the capability of HPC to bear on the understanding of complex projectile aerodynamics. In fact, the name of the Branch was changed from the "Wind

Tunnel Branch” to the “Computational Aerodynamics Branch,” signifying the acceptance by the community that computational work was a force to be reckoned with. Work by this team led to some significant accomplishments and represented some pioneering efforts in the ability to compute projectile aerodynamic behavior that had not been done before. More information on the accomplishments of the Computational Aerodynamics Branch and the cost savings from the computational approach to a severe problem in the development of SADARM is given in Appendix A.

A more recent example in the successful application of HPC in projectile aerodynamics is the development of the Digital Virtual Aerodynamics Range (DVAR). Prior to having a fully-integrated computational capability, the aerodynamic characteristics of new projectile configurations had to be determined from actual range tests (Figure 3). This meant that after a concept was initialized, the design would have to be completed, fabrication of the prototype would have to take place, and finally the tests would be completed. If




Figure 4. Virtual fly-out of a 40mm grenade concept in the DVAR

there was a failure or if slight changes in the concept needed to be made, this process would have to be repeated from scratch. With the advent of the DVAR, the design, development, and production cycle of new projectiles can be dramatically reduced. The DVAR is the result of a multi-discipline approach to simulating the flight of a projectile using a supercomputer and physics-based modeling and simulation techniques. Through the integrated solution of the Navier-Stokes equations coupled with the rigid body, six-degree of freedom equations, a projectile concept (Figure 4) can be virtually flown in the computer and one can accurately determine the flight trajectory and characteristics. Additionally, with the computational solution there is a complete numerical and visual history of the entire flight pattern, which is not available through traditional range testing. While a simulation of this type for 100 meters still takes a significant time to complete computationally (three days, running 24 hours, on multiple processors), it is still much cheaper and faster than the traditional approach. Also, with the advent of multiple-core electronic chips, the time to solution will significantly decrease. Multiple-core technologies reduce supercomputing process costs and will improve access from users’ desktops.

A similar account of the introduction and value of advanced computing within the Corps of Engineers can be found in the history of work in hydraulics engineering at the Corps’ Engineering Research and Development Center (ERDC).¹⁵ In the mid 1960s, the physical modeling of rivers, bays, and estuaries was reaching a peak, while at the same time, the digital computer was making its presence known. Early discussion of the potential application of computational techniques and HPC of the time for modeling flood routing and reservoir releases was met with comments such as “It is not considered that electronic computers as such will ever be an alternate or substitute for hydraulic models except in special cases,” and that “The computer is merely a means of solving more readily analytical problems that might otherwise be too laborious on account of the volume of computations.” Nonetheless, the Corps’ Waterways

¹⁵ B.H.Fatherree, *The First 75 Years. History of Hydraulics Engineering at Waterways Experiment Station*, Engineering Research and Development Center, Vicksburg, MS, 2006.

Experiment Station (now ERDC) began to use computational models to calculate currents in the Narragansett Bay and the effects of hurricane barriers in the Galveston Bay. Efforts to recruit and train engineers in the new computer technology and a strong management commitment were the keys to success. The ERDC continues its long history and tradition today in developing, applying, and advocating the use of HPC as a key ingredient of fundamental research and development.

We have also witnessed a very successful buildup of a computational capability where little existed before in the Medical Research and Materiel Command of the Army Medical Command, where they are now using HPC in computational biology. This field is a mixture of the physical sciences, computer science, and the life sciences of biology, physiology and medicine. The start of this capability came from funding from the DOD High Performance Computing Modernization Program (HPCMP) to create a DOD Biotechnology HPC Software Applications Institute. The institute began six years ago with the vision and dedication of a few people who realized that computational biology would be a significant research capability within the DOD. With the support of the HPCMP, the institute, which started with just 4 people, now has a staff of 40, 75% of which hold PhDs. The team has developed 12 codes that run on HPCMP assets and take advantage of the multi-core systems. This has allowed researchers to save countless hours when searching for DNA matches unique to an organism. Additionally, it has allowed for calculations that would have taken nine years to complete on serial systems to be completed in nine weeks using the newly-developed computational techniques. The HPC Institute resides within the Army Medical Command's Telemedicine and Advanced Technology Center. Its creation and continued success is a direct result of a strong management commitment, acceptance of HPC as a necessary tool for research, and development of a strong interdisciplinary team. This same success can be achieved in many areas within the Army RDT&E community.

These and similar stories illustrate that some significant DOD work could not have been done without the capabilities of high performance computers. These projects also depended on a dedicated and knowledgeable staff that was willing to travel to where the computers were located and to remain there for extended periods. Supportive upper management, funding, and training were critical factors. To move forward, we need to make HPC systems more accessible and more user-friendly. The advent of the World Wide Web and fiber optic links is making desktop HPC access possible. The Army should leverage the HPC developments in industry and other government agencies as discussed in the next two sections.

There are currently a number of research centers with a mission related to the design and development of ground vehicles.

-- The US Army Automotive Research Center (ARC) - a Center of Excellence for Modeling and Simulation of Ground Vehicles, is led by the University of Michigan in collaboration with universities in Ohio, California, Iowa, Virginia, South Carolina, and Alaska. Established in 1994, ARC's research thrusts include dynamics and control of vehicles; human-centered modeling and simulation; high-performance structures and materials; advanced and hybrid power trains; vehicle system integration, optimization, and robustness.

-- Established in 1993 and chartered by the Secretary of the Army, the National Automotive Center (NAC) was created to be the DOD and Army focal point for collaboration in ground vehicle research and development. The NAC, a part of the Tank and Automotive Research, Development, and Engineering Center in Warren, Michigan, focuses on automotive technology, collaboration mechanisms and partnerships, international cooperative R&D and special initiatives. The collaboration between government, academia and industry provides opportunities for exchange of information and technology.

-- In 2003 the Center for Advanced Vehicular Systems (CAVS), led by W.L. Giles Distinguished Professor Roger King, was established at Mississippi State University in part due to the decision by Nissan Motor Company to locate a manufacturing plant nearby. This is an interdisciplinary center with a research, development, and technology transfer mission. Research is focused on material science, computational structural and fluid dynamics, manufacturing process and multi-scale modeling, and vehicular systems engineering. A significant note is that the CAVS is directly connected with and serves as a focal point for the Mississippi State College of Engineering High Performance Computing Center. They have significant experience in HPC as Mississippi State was one of the early recipients of a National Science Foundation award to fund an Engineering Research Center focused on Computational Field Simulations.

The work ongoing at these Centers is key to advancing vehicle design and development and, coupled with an increase use of HPC, can help transform the current work within the Army. While there is direct involvement by the Army in the ARC and the NAC, the Army would benefit from a close working relationship with the CAVS.

The HPCMP today continues an aggressive approach in providing HPC services to the DOD community. There are currently five DOD Supercomputing Resource Centers (DSRCs): the Air Force Research Laboratory (AFRL) DSRC in Dayton, Ohio; the Army Research Laboratory (ARL) DSRC at Aberdeen Proving Ground, MD; the ERDC DSRC in Vicksburg, MS; the Maui DSRC in Kihei, Maui, HI; the Navy DSRC at the Stennis Space Center, MS.

The five DOD centers are shared with users across the Defense community and are connected by the high-speed Defense Research and Engineering Network. They would thus be accessible to Army R&D staff and their contractors to launch a program in HPC for Army-related design and manufacturing. Initially the High Performance Computing Management Office (HPCMO) was placed in the Office of the Secretary of Defense, but in 2011 the HPCMO was transferred to the Assistant Secretary of the Army for Acquisitions, Logistics, and Technology (ASAALT). ASAALT has decided to have the ERDC take on the management of the HPCMO for all users across DOD.¹⁶

¹⁶ See the testimony of Dr. Marilyn Freeman, Deputy Assistant Secretary of the Army for Research and Technology before the Subcommittee on Emerging Threats and Capabilities, Committee on Armed Services, U.S. House of Representatives on the U.S. Army's Science and Technology (S&T) Program for Fiscal Year 2012, March 1, 2011.

IV. Advancing the Industrial Use of HPC and Some Case Studies

There is renewed interest in manufacturing as a key factor in economic competitiveness. This is made the more urgent given the current poor performance of the U.S. economy. The Council on Competitiveness, the Manufacturing Institute (of the National Association of Manufacturers), the National Center for Manufacturing Sciences, the Alliance for High Performance Digital Manufacturing, the National Academy of Engineering, DARPA, and the HPCMO are all working on the challenge. The use of HPC in the design and production of parts, components, and finished systems holds promise to speed up the process and reduce the cost of going from research to ultimate fielding. New programs have been emerging over the past few years to engage the mid-sized subcontractors that make up the “missing middle” described below. One example comes from a leading DOE HPC Center at Oak Ridge National Laboratory where they have developed an industrial partnership program providing American companies with access to some of the most powerful supercomputers in the world. The program also provides access to the expertise necessary to take full advantage of the supercomputer’s power. A significant problem that has been identified lies with small businesses that may lack the resources or expertise to utilize advanced technology. This has been called the “missing middle.”

The concept of the “missing middle” is the gap between what large sophisticated companies can do with supercomputers and what small businesses can do. (See Figure 5.)




Figure 5: The “missing middle” showing the vast majority of computer users do not use high performance machines and how much productivity may be lost.

On the one hand, you have a firm with a cadre of computer engineers with access to a supercomputer and a well-equipped factory floor with highly capable production engineers. On the other hand, you have small businesses in the supply chain with access, at best, to PCs and workstations in the office and a factory floor with some computer-aided machines perhaps managed with a local area network. The question then is how to bridge this gap, how to fill the “missing middle” so as to enable the use of high performance computing all up and down the supply chain. Looking back at reference 1 we see what is needed: hardware, software, and well-trained engineers capable of dealing with the results of designs built on supercomputers.

As we will see in the next section on government programs, an extension program has been established in the Department of Commerce to provide technical assistance in the use of computer-based systems to small companies. This program could be enhanced to include the opportunities presented by HPC.

Currently only a few companies, mostly very large but some mid-sized ones, have the ability to buy and staff high performance computers. Some of these Fortune 20 and 200 firms are using HPC in the manner indicated above. The Council on Competitiveness has published several case studies of which four are discussed briefly below.

AAI. AAI develops aerospace and defense technologies for the Army and the Marine Corps, especially in unmanned aircraft such as the Shadow. The company uses HPC in applying computational fluid dynamics (CFD) in the design of the landing gear and in fuselage design both to reduce drag and to funnel cooling air to the rear-mounted engines. They also design the propeller blades for optimum performance using HPC CFD. If the customer changes the type and weight of the payload, the HPC system can quickly alter the design to accommodate the shift in weight. Michael Guterres, AAI's Director of Engineering, says "the use of high performance computing and advanced CFD software allows us to keep our customers happy, be more competitive and maintain our position as a premier supplier of tactical unmanned aircraft systems."¹⁷

Ford. Ford has long used supercomputers in its engineering work. In a recent case study, Ford described work on a new EcoBoost engine using turbo charging and direct fuel injection. The new engine is said to be 15-20% better in fuel economy. The case study also discussed development of several safety features, work on noise and vibration, and the use of computational fluid dynamics to improve the external body shape. Ford's executives say "The flexibility and speed made possible by HPC lets us simulate a wider range of scenarios, component combinations and associated trade-offs than would have been possible with physical testing."¹⁸ They used computer simulation for some of their road testing and for wind tunnel experiments. The computer work reduced the amount of physical prototyping.

Whirlpool. Tom Gielda, Whirlpool's Global Director of Mechanical Structures and Systems, says HPC "allows the Whirlpool engineers to run more complex simulations and do them faster. In addition, engineers are now beginning to run stochastic models in order to better predict how a product or material will function."¹⁹ Whirlpool models appliances such as washing machines or refrigerators with all the components and functions in one HPC model. It even models the packaging used in shipping. They experiment with the model on the HPC system, thereby reducing the number of prototypes needed for validation. They save both time and money on the path to market.

¹⁷Michael Guterres, *Builds Competitiveness with HPC and the "Black Art" of Computational Fluid Dynamics*, Case Study for the Council on Competitiveness See <http://www.compete.org/publications/detail/1680/hpc-case-study-council-showcases>, Washington, D.C., 2010. -power-of-high-performance-computing-at-aa/.

¹⁸Nand K. Kochhar and Alex Akkerman, *From Safety to EcoBoost: HPC Enables Innovation and Productivity at Ford Motor Company*, Case Study for the Council on Competitiveness, Washington, D.C., 2011; See <http://www.compete.org/publications/detail/1664/case-study-council-showcases-power-of-high-performance-computing-at-ford-motor-company/>.

¹⁹Tom Gielda, *Whirlpool's Home Appliance Rocket Science: Design to Delivery with High Performance Computing*, Case Study for the Council on Competitiveness, Washington, D.C.; See <http://www.compete.org/publication/detail/682/whirlpools-home-appliance-rocket-science-design-to-delivery-with-high-performance-computing/>.

Goodyear. Goodyear technical staff was challenged to develop an entirely new tire to help pull the company out of a financial slump. The competitive situation was severe and time was of the essence. The objective was “optimizing tire handling on dry, wet, icy or snowy surfaces, we also had to make major improvements in tread wear, noise reduction and vehicle handling. So, in a sense, we were going beyond optimization, beyond creating a tire that was quiet, rode well and handled well. We were chartered with creating something new; a tire that would perform better than any other tire in history. And we were under severe time pressures to come up with the new design.”²⁰ The tire had at least 18 different components/materials that had to be factored into the physics- and materials-based computer model to provide all the desired characteristics and to meet requirements for production at affordable costs. The work was modeled on Goodyear’s own HPC system with collaborative advice from experts in HPC at Sandia National Laboratory. They modeled the product and then conducted experiments on the HPC, making changes in the variables and evaluating the performance, all in computer simulation. The project was successful. The new award-winning tire, the Assurance, was on the market in just one-third the time Goodyear traditionally needed.

V. Some Government Programs to Advance the Use of Computers in Design and Manufacturing

An early government research effort on using computers in manufacturing was the Automated Manufacturing Research Facility (AMRF) at the National Bureau of Standards. Beginning in 1979 and 1980, the Bureau began assembling hardware, software and highly-skilled research scientists and engineers to build what became, according to American Machinist magazine,²¹ “the free world’s largest and most advanced public research facility for the study of automated manufacturing.” This project grew out of research on computer correction of errors in coordinate measuring machines. The technique was soon applied to correcting errors in machine tools which led to attempting to use computers, sensors, new software and architectures, and robots to produce, automatically, batches of machined parts of unparalleled precision and accuracy. The AMRF consisted of several workstations for horizontal and vertical machining, turning, deburring, and post-production inspection. These were supported by robot carts carrying materials and feeding the metal blanks into the machine tools. All were connected by a computer network. Control of the system was done by computer by decomposing tasks into a hierarchy of layers, each describing a portion of the work. This approach became known as the RCS – the Real-Time Control System – now the basis for much of automated control of computer-driven systems.

A key aspect of the AMRF was a set of standards for the many factors in computer control. The AMRF was constructed from machine tools, robots, and computers from disparate

²⁰ Loren Miller in *Goodyear Puts the Rubber to the Road with High Performance Computing*, Case Study for the Council on Competitiveness, Washington, D.C., 2010; See <http://www.compete.org/publications/detail/685/Goodyear-puts-the-rubber-to-the-road-with-high-performance-computing>. For more details see Loren K. Miller, *Simulation-Based Engineering for Industrial Competitive Advantage*, Computing in Science & Engineering, May/June, 2010.

²¹ Quoted by Raymond Kammer in the Foreword to Joan M. Kazen, *Automating the Future, A History of the Automated Manufacturing Research Facility, 1980-1995*, NIST Special Publication 967, National Institute of Standards and Technology, Gaithersburg, MD, March 2001.

manufacturers, each with its own proprietary software. So a neutral data interface standard was necessary so these machines could talk to each other and to the central controller. Standards for describing parts, for factory architecture, and for describing various machines were developed and proposed to standards bodies for adoption.

After about six years of research and development, the AMRF was demonstrated to a large audience of people from industry and universities as well as senior government officials. The single operator sat at a computer terminal deciding what parts to make.

The AMRF program was staffed by NIST employees augmented by about 50 visiting researchers from industry and people from nearly 40 universities. Much of the financial support came from the Navy's Manufacturing Technology program. In some cases, a company could conduct its own research at the facility as long as it shared non-proprietary data.

An important point is the need, in using high performance computers, to have skilled computer engineers working with subject matter experts to carry out the design development and then to transfer it to a receiving group in the manufacturing department. The design model will have to be converted to specific tasks on the manufacturing floor, much as was done in the early applications of CAD/CAM (computer-aided design/ computer-aided manufacturing) for making simple parts. Large companies should be able to develop this expertise fairly easily, but small and medium- sized companies will have greater difficulty. When the Commerce Department wanted to transfer to industry new technology on automated manufacturing in the late 1980s, it found that smaller firms could not handle the technology for lack of computer expertise. As a result Congress approved a program called the Manufacturing Extension Partnership (MEP)²² to work with the states to provide information and assistance to small businesses (up to 500 employees). Modeled after the agricultural extension services to farmers, the MEP quickly learned that many of the smaller companies had no computer knowledge nor did they even have a PC in the office. So the early task was to educate the clients in the fundamentals. Usually the MEP centers had access to demonstration facilities where small businesses could try out some of the computer-based technologies. Starting with three centers in 1991 the MEP program now operates 59 centers spread over some 400 sites covering every state in the country. A similar exercise would be needed to introduce high performance computing to small businesses. Since many of these companies are supplying large companies, much of the work of technology transfer would logically be done by the large companies. Perhaps the MEP would work on this as a logical extension of the current program. The National Center for Manufacturing Sciences²³ is proposing a national network of Predictive Information Centers that seems similar to the MEPs, but focused on HPC.

The military has long needed to decrease the time it takes to move new technology from the laboratory to the warfighter. Part of the time is spent on developing designs, building prototypes, running development testing, and then repeating this cycle many times (owing in part to unforeseen design flaws) until the desired result is achieved. In the early use of computers to design products, the designs were primarily two- and three-dimensional drawings of solid parts,

²² See the NIST Manufacturing Extension Partnership at <http://www.nist.gov/mep/>.

²³ *Enabling Digital Manufacturing: A Strategy to develop a National Innovation Network*, National Center for Manufacturing Sciences, Wayne, PA, September 2010.

typically for machining. This did not involve evaluating, within the computer models, performance of the design, such as heat transfer, shock resistance, flow around the object, or energy efficiency. The advent of very fast and very capable supercomputers, together with interdisciplinary computational techniques, has made possible the formulation of models based on physics, chemistry, and material properties and experimenting with all the variables on the computer, leaving actual field testing as final validation of the resulting models. One can evaluate the effect of changing variables on ultimate performance of the product by use of the computer, thereby saving large blocks of development time – factors of three and more. This is a long step beyond the computer-aided designs of mechanical drawings that are in use in many small machine shops today.

We turn now to current DOD programs designed to promote the use of HPC in manufacturing.

The HPCMO conducts a program called CREATE (Computational Research and Engineering Acquisition Tools and Environments) to focus on improved design as opposed to production. The objective is “developing and deploying scalable multiphysics-based computational engineering tools to design and analyze DOD weapons system performance.”²⁴ The program’s goals include the following²⁵:

- Prevent defects and design flaws early in the acquisition process
- Reduce rework thereby enabling faster system deployment
- Reduce experimental testing time and effort through analysis of virtual prototypes

The CREATE program consists of four project areas: Ships, Aircraft, RF Antennas (integrated with platforms), and Mesh & Geometry. The CREATE Air Vehicle tools are focused on aerodynamics, structural mechanics, control, and propulsion. The Ships tools address shock hydrodynamics to assess vulnerability, hydrodynamics to assess resistance, seaway loads and sea keeping, and ship configuration optimization including hull-form optimization and compartment lay-out. The RF Antennas tools are focused on the electromagnetic performance of antenna systems, including co-site interference effects from adjacent antennas and platform features. The Mesh and Geometry tools will enable designers to generate and repair the numerical representations of the weapons systems. The CREATE effort consists of teams of governmental technical experts and defense contractors and academicians who are developing and testing software for use by DOD and Defense industry engineers in developing weapons systems. The selected topics resulted from a call to the services requesting large problems where a computational engineering approach could make a significant difference.

²⁴ Douglass Post, *The DOD Computational Research and Engineering Acquisition Tools and Environments (CREATE) Program*, National Defense Industry Association Workshop, 1 Feb. 2011, Arlington, VA.

²⁵ Robert Rassa, *CREATE Program Observations*, National Defense Industry Association, SE Division, April 13, 2011.

The intent is to significantly reduce acquisition costs early in the development cycle to minimize the number of expensive re-designs inherent in most large DOD acquisition systems. In CREATE software models, the computational designs would include not only the physical shape and type of materials, but also calculate the projected performance of the result. The simulated design is then used, in lieu of building prototypes, for experimentation on the computer. Actual experiments are used only for validation of the models. The models should also be useful in preparing and updating training and maintenance manuals. For the program to be successful, industry must be engaged since it also consumes time and expense in perfecting the new product design and preparation for production.

It is unclear how these software models might be useful in designing implementation of the computer models on the factory floor. The leader of the CREATE program says that beginning studies on the application to the factory floor are a year or two away. (However, the product design models from CREATE should enable the construction of computer-aided design models that can form the basis of instructions for machine tools and other manufacturing processes.) There is some discussion in the literature of the use of HPC on the factory floor. An early report²⁶ describes the use of supercomputing in chemical processes. Through studying various reaction pathways and changing reactants on the supercomputer, pilot plants and interim production facilities were no longer needed. This would be a good saving in time and money. There is some use of HPC in managing factories in areas such as production planning, design of tooling and of how the tools are used, inventory management, product assembly, energy management, data management, and quality control. Many of these functions can be handled by workstations and would only require HPC for a very complex system

There are ten individual CREATE projects. Two are of special interest here. One is Project Helios that is modeling the fluid dynamics around rotorcraft rotors and the frame. This is being done with and for the Aeroflightdynamics Directorate of the Aviation and Missile Research and Development Center co-located with NASA at Ames, California. The software is now in its second release. The Army's Chief Design Engineer at Ames told us during a July 13, 2011 teleconference that "the HPC model in Helios makes it practical for the first time to model the rotor wake back to the tail of the helicopter. Use of the model reduces, but probably does not eliminate, problems that will show up in flight tests. But it is a significant reduction in problems."

In Project SENTRI, HPC will make it easier to calculate multiple antenna arrays for small- and large-scale systems. This is a long-standing problem for any platform that has many antennas. The Communications and Electronics Research, Development, and Engineering Center of RDECOM is interested in, and is monitoring, this project. SENTRI's focus, however, is on the Navy and Air Force.

ARPA has several programs that relate to this discussion. One, iFAB, proposes to perform manufacturing in a central foundry-like facility where customers submit design software and the iFAB produces the desired product. The iFAB vision²⁷ is

"to move away from wrapping a capital-intensive manufacturing facility around a single defense product, and toward the creation of a flexible, programmable, potentially distributed production capability able to accommodate a wide range of systems and system variants with extremely rapid reconfiguration timescales."

²⁶ _____, *21st Century Manufacturing Enterprise Strategy*, The Iacocca Institute, Lehigh University, U.S. Dept. of Defense, Office of Manufacturing Technology, Diane Publishing, 1991, p21.

²⁷ See iFAB at http://www.darpa.mil/our_work/TTO/programs/adaptive_vehicle_make_AVM.aspx.

The proposal is modeled on the silicon chip foundries in the electronics industry today. This approach, for complex systems, likely will need HPC in some form.

A second DARPA effort is the ubiquitous high-performance computing program or UHPC. The objective of the UHPC is:

“to provide the revolutionary technology needed to meet the steadily increasing demands of DOD applications – from embedded to command center and expandable to high performance computing systems. This may be accomplished by developing highly parallel processing systems with significantly increased power efficiency enabling ease of programming application development for the user; resilient execution through all modes of system failure.”

This program holds the promise of facilitating the use of HPC systems and extending their use from the research arena to more applied subjects and making HPC use easier for the non-specialist.

The program goals are in the referenced solicitation.²⁸

A new DARPA program, Adaptive Vehicle Make (AVM), is directly aligned with the comments and recommendations being made in this report, i.e. the need to shorten the acquisition process and bring new weapon systems to our armed forces faster, cheaper, and fully capable. It is clear that HPC will make a significant impact when used in the AVM program. AVM is a portfolio of programs that is looking to revolutionize the design, development and acquisition of complex DOD systems including ground vehicles. The AVM portfolio consists of three primary programs: 1. META, a \$60M program to build a high-level meta-language for the development, verification and validation of designs using a component model library; 2. iFAB, see above, and 3. FANG (the Fast, Adaptive, Next-Generation Ground Combat Vehicle) that seeks to develop an infrastructure, conduct design challenges based on open sources, and generate the next infantry fighting vehicle. The use of Vehicleforge and Manufacturing Experimentation and Outreach (MENTOR) are also part of the FANG program. Contracts have already been awarded to Vanderbilt University to develop and operate vehicleforge.mil for development of an open source environment and the University of Pennsylvania for a credentialing and verification scheme. A previous award was made to Carnegie Mellon for providing drive-train cooling system models using META. These major DARPA initiatives must be fully considered by the Army as new tools that can be coupled with the existing HPC capability to modernize the vehicle development process. If successful, they can provide a fundamental change to the way ground vehicles are designed, prototyped, built and tested.

A related program, this time in the Office of the Assistant Secretary of Defense for Research and Engineering, is called Engineered Resilient Systems²⁹ This program is working on design or

²⁸ *Request for Information (RFI) Ubiquitous High Performance Computing*, DARPA-SN-09-46, Information Processing Technology Office, DARPA, Arlington, VA.

²⁹ See a briefing paper by Robert Neches, “DoD’s Engineered Resilient Systems (ERS), S&T Priority, Office of the Assistant Secretary of Defense for Research and Engineering, June 2011.

redesign of warfighting systems to make them more adaptable to changing threats, more trustworthy, and affordable. More modeling and simulation is anticipated. HPC is not explicitly mentioned, but it can certainly play a significant role.

In contrast to the industry case histories, the military does not do its own manufacturing. Design is begun in the Army's research laboratories, is formalized at the Army Research, Engineering, and Development Centers and refined by the Army Acquisition System's Program Executive Officers and Program Managers. Potential contractors for manufacturing may offer new design details as they prepare their bids. Therefore, design is a joint exercise between the Army and its contractors. Production is the province of the contractors. In industry, this whole process is usually conducted solely by an individual company.

Working across institutional lines is now possible by using software that allows various players to work on a design package simultaneously, thereby avoiding the need to send the package back and forth repeatedly. To reap the benefits, design engineers must be able to access the high performance computer from their desktops. Past experience shows that designers will be reluctant to use HPC if it means going somewhere other than their own offices and if it is cumbersome to access and utilize.

VI. Summary and Recommendations

High performance computing has enabled the use of physics-based models for design and manufacturing of very complex systems such as military platforms. The evolution in computer technology is transforming supercomputing from a very complex, expensive, and limited tool to one that now provides HPC capacity, available remotely via the Internet, at the desktop. This means a program of continual upgrades in machines, software, and personnel capability must be pursued if one wishes to be in a position of leadership in materiel acquisition. This evolution must be pressed forward by management at all levels.

Most of the HPC work in the Army has been primarily by users in the research and development laboratories. This is only the front end of the acquisition cycle. To reap the advantages fully, HPC must become a routine factor in the entire process. This will involve close collaboration among the research, development, and engineering laboratories, the program executive officers (PEOs) and program managers (PMs), and the contractors. Only when the commitment is made by the Army to involve all participants in acquisition will the full benefits of physics-based modeling on high performance computers be realized.

Given that the Army's manufacture of materiel is carried out by industry contractors, it is important that these industry partners participate fully in using the latest in HPC. Some firms will be ahead of us in this regard but most, especially those smaller firms in the supply chain, will need help. One way is to encourage these firms to use the DOD's HPC centers. Existing policy provides for private sector use of government experimental facilities under certain conditions. Government expertise can be made available as needed. This idea has been advanced

independently by the National Center for Manufacturing Sciences.³⁰ They propose a national network of Predictive Information Centers that would assist companies in using HPC systems.

To help ensure that the Army gains the benefits of using HPC in all aspects of acquisition, we present a set of recommendations for actions by the ASAALT and the Army S&T community at large to establish an Army program on the use of HPC in the acquisition community for design and manufacturing:

For the Deputy Assistant Secretary for Research and Technology

1. *Conduct a technology forecast for High Performance Computing for design, development, and manufacturing.* The Army should update the forecast in manufacturing made in 1992 (and judged by us to be poorly done) and should focus on the use of advanced computational modeling and simulation as well as on the factory floor.
2. *The ASAALT should seek support from top Army leaders by making the case for shortening the acquisition cycle through the use of HPC.* The technology should be of value in upgrading existing systems as well as in planning and fielding new systems.
3. *The ASAALT should designate a lead person to promote the use of HPC in design, development, and manufacturing, and in the supply chain.* The Assistant Secretary should appoint representatives from RDECOM and from the PEO/PM community to develop a program on HPC usage.
4. *The ASAALT should request that periodic workshops be held with participants from the RDECs, PEOs and PMs, industry, HPCMO, CREATE, and DARPA.* This workshop would likely parallel a workshop held by the NDIA, but would have as its theme on how to bring HPC into design, development, and manufacturing to the Army.
5. *The S&T Office and the HPCMO should identify and implement ways to make HPC more available on the desktop.* The use of HPC will be enhanced by streamlining accessibility and ease of use.

For the Army S&T Community

1. *The Army laboratories should learn about the Air Vehicle, Ships, and Antennae components within the HPCMO's CREATE program and see what applies to the Army.* The Army is already involved with the Air Vehicle portion of the CREATE programs. The Army should leverage the CREATE program and work to become more involved in developing a Ground Vehicle component.
2. *Managers should learn more about the DARPA programs in manufacturing.* DARPA's programs are focused on improving the manufacturing component of military acquisition. Some of the effort is on HPC.

³⁰ *Enabling Digital Manufacturing: A Strategy to develop a National Innovation Network*, National Center for Manufacturing Sciences, Wayne, PA, September 2010.

3. *Managers and subject matter experts should become familiar with the activities in the private sector on HPC in design and manufacturing.* There are several groups in the commercial sector that are promoting HPC use; for example:

- The National Center for Manufacturing Sciences
- The Alliance for High Performance Computing Digital Manufacturing
- Intel High Performance Computing Program Office
- Council on Competitiveness, HPC Initiative
- U.S. Manufacturing Round Table
- National Defense Industry Association
- International Data Corporation reports on HPC

Appendix A. More Details on the Computational Aerodynamics Branch at the Army Research Laboratory

One of the goals of the Computational Aerodynamics Branch was to develop a computational capability that would help augment the excellent experimental facilities at BRL. While exterior ballistic research had been a fundamental area since the beginning of the BRL, most of the work up to this time has been primarily experimental and resulted in some gaps in the technology. As an example, the ability to measure aerodynamic drag was ideally suited for experimental ranges, but the ability to determine both supersonic and transonic magnus coefficients was tenuous at best. The early computational efforts were thus targeted to address these shortfalls. One of the first papers in this area reported on the development of a parabolized Navier-Stokes code for accurately predicting the supersonic magus effect about spinning cones, ogive-cylinders, and boat-tailed afterbodies. This work led to the successful characterization of the aerodynamic coefficients for projectile configurations. This was then followed by the more difficult predictions of transonic magnus effects. The difficulty was attributed to the lack of solid-state memory on the supercomputers of the day. Further advances were made for projectiles with domed bases, with base bleed, with rotating bands, configurations with fins, fuzes, etc. The trend during the development was to obtain the capability to address all physical characteristics of a projectile configuration that would then lead to a truly full predictive capability with the goal of reducing the number of prototypes and thus reducing the development costs.

An example of the potential savings from computational predictions can be seen with the SADARM projectile (Sense and Destroy Armor Munition), XM898. The SADARM was a concept for top armor attack. It consisted of two sub-munitions within a 155mm caliber round to be fired from a howitzer or up to six sub- munitions from the Multiple Launch Rocket System. The projectile nose was fitted with an electronic time fuze, preset to detonate at approximately 1000 meters above the target area. Upon the ejection of the sub-munition from the carrier and after separation had occurred, the sub-munitions contained a glide parachute that allows the sub-munitions to hang at an angle of 30 degrees and, during its slow rotation, to search the area for targets. Once detected and within range (approx <152m), the sub-munition would fire an explosively-formed penetrator (EFP). While full-scale production was scheduled for FY94, a major problem was occurring during testing which showed that there was a dud rate of almost 80%. A SADARM red team was formed and a number of potential causes and fixes were developed. During this time frame, a new computational capability was being looked at for the prediction of bodies in relative motion. It was called the Chimera technique and would be applicable to the SADARM analysis. The thought was that the SADARM sub-munitions that were contained in the spinning carrier projectile, when ejected from the base, would initially clear the carrier. The trailing sub-munition, however, would be in the draft of the first sub-munition and they would eventually collide with each other. Because of the very high spin rate, there would be a cookie cutter effect that would destroy the sensitive guidance and firing electronics of the sub-munitions. Navier-Stokes computations clearly showed that, indeed, the munitions would collide with no intervention, but that with the addition of the fins, separation of the sub-munitions would occur, thus collisions and duds could be avoided. This example is one of the early indications that the correct use of physics-based modeling and simulation, used early in the design process, could potentially lead to cost and time savings and enable major weapon system programs to be completed in time and under budget.