White Paper

DoD Research and Development Agenda

For

High Productivity Computing Systems

(Prepared for)

Dr. Charles J. Holland
Principal Assistant Deputy Under
Secretary of Defense for Science and Technology

3080 Defense, Pentagon
Washington, DC 20301-3080

June 11, 2001

Background

On February 27, 2001, Cray Inc. announced an agreement with NEC of Japan to withdraw Cray Inc.'s anti-dumping complaint (in order to lift import tariffs on Japanese vector supercomputers) in exchange for $25 million cash investment and exclusive North American sales rights for NEC vector supercomputer products and support services. This historic agreement between these two intense competitors served to focus high-level attention on an issue critical to national security; i.e., the critically weak state of the U.S. industrial base in high-end computing, specifically in tightly-coupled, high-bandwidth "vector" computing systems. After considerable analysis, DoD leadership recognized the need to reinforce this important technology sector through the research and development program described in this white paper.

In late April 2001, Mr. David Oliver (Principal Deputy Under Secretary of Defense for Acquisition, Technology, and Logistics) tasked Dr. Charles J. Holland (Principal Assistant Deputy Under Secretary of Defense for Science and Technology) to develop a plan to address these national security needs. The purpose of this white paper is to outline a research and development plan to (1) revitalize high-end computer industry, providing options for high-end computing systems for the national security community and (2) develop improved tools to enhance software portability, performance, and ease-of-programming for a wide range of computer architectures (commodity clusters to high-end vector systems).

Dr. Holland assembled a team of technical experts in high performance computing to write and review this document. The writing team was led by Mr. John Grosh (Office of the Deputy Under Secretary of Defense for Science and Technology (ODUSD(S&T)) and included Mr. Robert Graybill (Defense Advanced Research Projects Agency (DARPA)), Dr. Bill Carlson (Institute for Defense Analysis Center for Computing Sciences), and Ms. Candace Culhane (National Security Agency (NSA)). The review team consisted of Dr. Frank Mello (DoD High Performance Computing Modernization Office), Dr. Richard Games (The MITRE Corporation), Dr. Roman Kaluzniacki (National Security Agency), Mr. Mark Norton (Office of the Assistant Secretary of Defense, Command, Control, Communications, and Intelligence), and Dr. Gary Hughes (National Security Agency).

Executive Summary

High performance computing is at a critical juncture. Over the past three decades, this important technology area has provided a superior computational capability for many important national security applications. Government research, including substantial DoD investments, has enabled major advances in computing, contributing to the U.S. dominance of the world computer market. Unfortunately, current trends in high performance computing are creating technology gaps that threaten the continued U.S. superiority for important national security applications. This white paper addresses these gaps with a focused research and development program, creating new generations of high-end programming environments, software tools, architectures, and hardware components in order to realize a new vision of high-end computing, *high productivity computing systems.*

Today, high performance computing research and development is focused on either short-term evolutionary commodity-based cluster and distributed computing or very long-term technologies such as quantum computing. Current research does not adequately address important medium-to-long-term (i.e., five to ten years) requirements encountered in high-end applications, such as cryptanalysis and weather modeling. Such applications require tightly coupled architectures, typically incorporating high-bandwidth/low-latency memory and storage subsystems. Without implementations of new computer architectures and improved software tools, high-end systems will become increasingly inefficient (e.g., computation, power, and size) and continue to be difficult to program. The need for a program to address these issues is urgent, due to the poor health of the U.S. industry in high-end computing and low-level of funded academic research in computer architectures and software tools.

This proposed agenda addresses these challenges through a focused research and development program specifically designed to produce innovative ideas and reinvigorate the academic and industry research community to develop and build new generations of high-end systems and software. Three program thrusts – technology development, concept demonstration, and industry adoption – are employed to ensure the development and manufacture of high productivity computing systems. Integrated into this effort is the development of software methodologies, tools, and standards to address issues of productivity, portability, and performance. Throughout all phases of this effort, DoD operational and research software applications will serve as the requirements driver for architecture and software research and systems assessment. Industry adoption is seen as a central strategy to ensure that cost-effective solutions are made available to the national security community. The end result will be integrated hardware and software systems tuned for high-end national security applications, rather than the loosely coupled solutions or clusters commonly available today.

In order to ensure a focus on national security requirements, DARPA, NSA and ODUSD(S&T) will form the core management team. This management team will provide technical oversight for the program, facilitate interactions with the user community, and coordinate with complementary efforts within the DoD and other government agencies. Government, industry, and academia consortia will be established to ensure community and industry adoption. Through the efforts described in this white paper, a new class of computing systems will be realized that promotes user productivity and high-end system efficiency.

1. Introduction

Over the past three decades, high performance computing has enabled major advances in defense research and engineering and provided unmatched computational capability supporting the national security intelligence and operations communities. Unfortunately, immediate and long-term issues present a significant challenge to the continued U.S. superiority in this technology area. This white paper proposes a research and development (R&D) agenda to create a new generation of high *productivity* computing systems characterized by balanced system architecture including high effective bandwidth, robust implementation, and responsive software/hardware components. These new systems will address the inherent difficulties associated with the development and use of current high-end systems and applications, such as programming productivity, performance, portability, scalability, reliability, and tamper resistance.

As reported in recent DoD studies, there is a national security *requirement* for high productivity computing systems.[1,2] Without government R&D and participation, high-end computing will be available only through commodity manufacturers primarily focused on mass-market consumer and business needs - a solution that is ineffective for important national security applications. To provide the DoD with increased options for high-end computing, we propose a program with two primary goals. The first is to provide economically viable *high productivity computing systems* for the national security and industrial user community. The second is to reinvigorate the high-end hardware and software communities to develop a new generation of researchers, engineers, and leaders to drive the advancement and development of new high-end architectures and tools throughout the decade.

To achieve these goals, this program must address three overarching issues impeding the development and utilization of high-end systems:

- **Balanced System Performance:** The increasing imbalance among processor speed, communications performance, power consumption, and heat removal results in high-end systems that are chronically inefficient for large-scale applications. This effect increases the time to solution and the cost of programming, operation, and facilities acquisition (e.g., cooling, power, and floor space).
- **Improved Software Tools and Methodologies:** There exists a critical need for improved software tools, standards, and methodologies for effective utilization of multi-processor computers. As multi-processor systems become pervasive throughout the DoD, such tools will reduce software development and maintenance – a major cost driver for many Defense system acquisitions.
- **R&D Funding:** Near-elimination of R&D funding for high-end hardware architectures has resulted in a dramatic decrease in academic interest, new ideas, and people required to build the next generation high-end computing systems.

To address these issues, we propose a strategy that focuses on providing a pipeline of new ideas and people to strengthen the U.S. industrial base in high productivity computing systems. This program will involve research and development designed specifically to reinvigorate university research, gain industry and DoD participation, create a more effective

[1] "Task Force on DOD Supercomputing Needs," Defense Science Board Study, October 11, 2000.
[2] "Survey and Analysis of the National Security High Performance Computing Architectural Requirements," Presentation by Dr. Richard Games, MITRE, April 26, 2001.

industrial base, and respond to national security and dual use computing requirements. The program is comprised of three linked thrusts designed to create true high productivity computing system solutions throughout the decade: technology development, concept demonstrations, and industry adoption. The program incorporates the lessons learned from the recent history of high-end computing, economic and industrial base realities, and future innovative technologies.

2. Program Overview

The high productivity computing system (HPCS) R&D program is uniquely structured to address the complex programmatic, technical, and end user challenges. The HPCS program addresses the following broad topics across all of its activities:

- **Performance:** Improve the computational efficiency and performance of critical national security applications.
- **Productivity:** Reduce the cost of developing, operating, and maintaining HPCS application solutions.
- **Portability:** Insulate research and operational HPCS application software from system specifics.
- **Robustness:** Deliver improved reliability to HPCS users and reduce risk of malicious activities.
- **Academia and industry:** Reinvigorate interest and activity in HPCS architectures and software/hardware components.

The basic technical challenges underlying these broad HPCS research objectives are outlined below and are discussed in more detail in the *technical challenge areas* section.

- **High effective bandwidth:** Focus the attention of academia and industry on high bandwidth/low latency hierarchical memory systems based on future CMOS (complementary metal-oxide semiconductor) and other emerging technologies.
- **Balanced system architecture:** Scalable computer systems should be designed from an overall system perspective balancing the performance of processors, memory systems, interconnects, system software, and programming environments.
- **Robustness strategy:** Address system brittleness and susceptibility of large complex computing systems by exploring balanced system hardware and software reliability/fault tolerance capabilities, active application software bug tolerance, and intrusion identification and resistance techniques.
- **Performance measurement and prediction:** Develop a new class of metrics and benchmarks to measure and predict performance of system architecture and applications software.
- **System Tailorability:** To improve system efficiency for a broader class of user applications, hardware and software characteristics must adapt and optimize to changing workload and user requirements. Examples of tailorable features include support for multiple programming models, selectable machine abstractions, and configurable software/hardware architectures .

The proposed HPCS R&D program, Figure 1, is comprised of three linked programmatic thrusts designed to provide usable HPCS solutions throughout the decade. Such solutions will span a wide range of technologies and products: from software methodologies and tools to pilot high productivity computing systems. The primary programmatic thrusts are technology development, concept demonstration, and industry adoption as described below.

Figure 1 - HPCS Program Plan

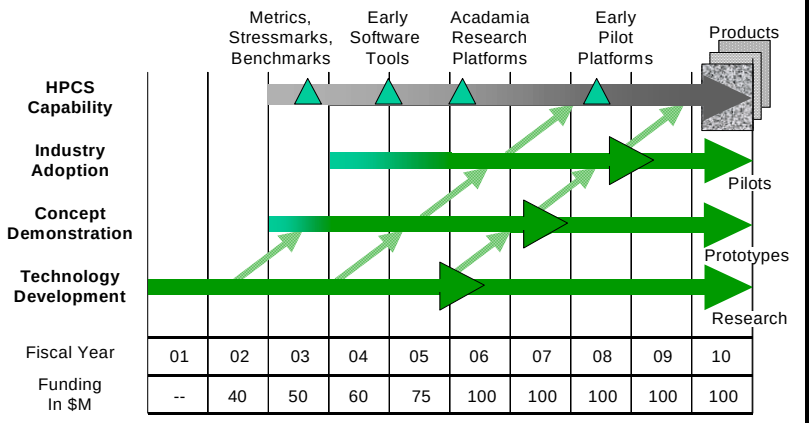

The technology development thrust will emphasize research in strategic technologies in high-end computing system architectures, micro-architectures, packaging, interconnects, runtime software, programming environments and methods, productivity profiling tools, runtime performance predictors, and HPCS specific benchmarks/stressmarks and metrics. Research of this type is best carried out by universities and research laboratories where scientists can focus on medium-to-long-term research. Also, the collocation of small high-end platforms with software researchers will be considered. End user and industry advisory forums will be used to provide top-level requirements and suggestions in order to ease technology transition in the later stages.

The concept demonstration thrust will incorporate incremental technology developments from the technology development thrust into multiple HPCS subsystem level developments and demonstrations. Developments of this type are best carried out by balanced industry, university, and research laboratory consortia. One goal is to replicate a number of these early prototypes to be placed in strategic universities and research centers for hands-on-research as a means to encourage and create a new generation of researchers.

During the industry adoption thrust, full-scale pilot systems will be developed with substantial investment by both government and industry. By building on the industry ties developed in the first two thrusts, these pilot systems will represent a synergy between industry and government interests. The user community will evaluate these emerging HPCS systems using "real world" applications, guiding further system development by industry. The new HPCS performance predictors and benchmarks/stressmarks developed earlier in the program will aid in long-term computing product procurements by providing an accurate measure of emerging commercial industry HPCS products and effectiveness for intended DoD applications.

Integrated into each of the thrust areas is end user involvement, where research and operational software applications will serve as the forcing functions for the development, prototyping, and testing of new HPCS hardware and software solutions. This is an important aspect of the program, tightly coupling the end use with the research and development.

The program consists of multiple waves of R&D over the next decade to provide a continuous stream of HPCS technologies addressing medium and long-term national security computational needs. The steady-state funding levels required for this program are based on previous high-end computing technology programs in the early 1990's and take into consideration the high cost of entry in this specialized field today. The management plan is designed to accommodate broad DoD and commercial community interests, yet be lean enough to rapidly make decisions on a day-to-day basis. Senior members of DARPA, NSA, and ODUSD(S&T) will make up the core HPCS management team with DARPA providing the lead, day-to-day, execution during the early years of the program. In addition, we foresee active involvement by NSA's high-end computing programs and the DoD High Performance Computing Modernization Program. Every three years, milestone reviews will be conducted to ensure that this program is meeting its goals and objective, tracking strategic national security requirements, and appropriately balancing medium and long-term efforts across the R&D portfolio. A larger more comprehensive stakeholder team representing DoD, academia, and industry interests will provide end user input and technical oversight on a yearly basis. The goal is to encourage the active participation of a broad spectrum of researchers representing academia, research laboratories, industry, and DoD laboratories and agencies across all three thrusts.

3. Technical Challenge Areas

This section describes some of the key technical elements necessary to research and develop high productivity computing systems. These new systems will address the current technical challenges that confront both development and use of current high-end systems and applications, such as programming productivity, performance, portability, scalability, reliability, and tamper resistance. A balanced technology and system development effort will be required across all technical elements and programmatic thrusts. Early performance characterization, metric development, and prediction activities along with end user and industry involvement will provide active feedback required to meet the challenging R&D program goals. The following key technical elements will be discussed: performance characterization and prediction, programming models, system architecture, and hardware and software components.

Performance Characterization and Prediction: Key to the research and development of new high productivity computing systems is the ability to measure and understand critical performance characteristics for the entire system – both hardware and software. The best known benchmark for high-end systems is Linpack. Unfortunately, this benchmark does not provide insight into critical high-end performance characteristics such as bandwidth, communications latencies, and I/O. To address this need, benchmarks and performance measurement methodologies must be developed that are comprehensive, widely implemented, and correlate to user software applications performance. Such an effort would include continual analysis of user applications codes to develop and refine lightweight/low-cost synthetic benchmarks used by the acquisition, development, and vendor communities to characterize performance. As performance measurement technologies progress, tools will be developed to predict the performance of codes on a given HPCS architecture. The ability to characterize and predict performance will provide a clearer picture on future hardware and software requirements and serves as a critical basis for evaluation and development of high-end systems.

Programming Models: A critical aspect of productivity is the ability to provide an abstraction of actual computing systems that allows programmers to concentrate on a simple and portable target for their efforts. For many years there have been two dominant programming models used in high performance computing systems: message passing and shared memory. The message passing model is attractive because of its performance and widespread adoption throughout industry, while the shared memory model is generally considered easier to use. Each of these models is relatively low-level and disconnected from the domain of most scientists and engineers but at the same time has real performance implications. Such fixed-point compromises in software/middleware/hardware design are common in today's high-end computing systems. As a result, a key attribute of the HPCS is the ability to support more than one programming model, transparently to the user, while adapting the underlying system architecture to efficiently select a different virtual machine and corresponding hardware/software micro-architecture. Research efforts in the programming model area should focus on obtaining new and/or enhanced models that ease use while not sacrificing the overall performance of the system.

Systems Architecture: Many current high-end systems adopt a design philosophy of "build the hardware, worry about the rest later." This has resulted in unbalanced systems for which many complex simulation codes demonstrate poor computational efficiency. This is a very real issue because the physical size of our current teraflop systems is starting to exceed our capacity to facilitate and maintain these large machines. Fundamental to the problem is the imbalance between processor speed and system latency and bandwidth. For complex simulations, this imbalance can result in applications software with poor scalability (utilizing less than a dozen processors for optimized codes) and poor single processor performance (less than 10% of peak processor performance). In addition, industry is just starting to acknowledge that there is a serious concern that Moore's law, the doubling of microprocessor performance every 18 months, will not continue to hold through this decade. Challenging chip power density issues plus the expected saturation of arithmetic capacity threaten to slow annual processors performance increases from approximately 60% to under 20%. New and innovative ideas in computer architecture will be required to maintain the current pace of CPU performance improvement. Fortunately, recent research indicates that significant gains may still be achieved by increasing the "effective" productivity of current high-end solutions.

The HPCS program embodies a holistic approach to high-end system architecture and design, promoting adaptability and scalability to provide a balanced system for a broad application space. Adaptable systems involve either software adapting to a given hardware environment or a hardware environment adapting to a given software load. Current efforts on software adaptability are beginning to pay dividends in the area of automatically choosing good implementations of algorithms for a given architecture and available computing resources. ATLAS[3] is an example of this: it automatically chooses the best implementation of basic linear algebra functions for any hardware/compiler combination. Examples of early research in adaptive micro-architectures are the DARPA's Data Intensive Systems (DIS), and Polymorphic Computing Architectures (PCA) programs. DIS primarily focuses on the processor memory subsystem bottleneck by exploring processor-in-memory, adaptive caches, and adaptive algorithms. The HPCS R&D program will leverage this early research and extend it by creating fully adaptable high-end systems that mutually seek the best operating point for a given problem.

[3] Automatically Tuned Linear Algebra Software (ATLAS), http://www.netlib.org/atlas/.

Hardware Components: In developing the required hardware components necessary for future high productivity computing systems, the system architects are encouraged to research and evaluate a range of hardware technologies in order to achieve a balanced system level architecture and programming environment. Architects should make distinctions between three classes of components: 1) the class of components where commodity COTS technology makes sense; 2) the class of components where modified COTS or IP (intellectual property) core technology can be leveraged; and 3) the class of components where technology gaps are identified and new micro-architectures, chips, network elements, memory devices, and packaging techniques need to be developed. Current examples of specialized systems that combine specialized processors and commodity memory components include the NEC SX-5 and Cray SV-2. It is clearly acknowledged that technology gaps do currently exist. Suggested technology areas for investigation and evaluation include, but are not limited to, photonics technology, storage and memory technology (e.g., magnetic memory, processing in memory, holographic memory), communications infrastructure, power conversion and power efficiency, and energy/thermal management techniques.

Software Components: To both fully evaluate and use new programming models and architectures, academic and industry partnerships should be funded to investigate and implement software tools for program production and understanding. Program production tools allow programmers to rapidly produce new programs and adapt old ones to new problems. Such tools would include traditional techniques such as compilers, development environments, and library definitions. In addition, the adaptation of new model abstractions, middleware, and tools into domain-specific frameworks will also merit investigation. Program understanding tools allow both programmers and users of programs to quickly gain insight into both correctness and performance. Research is required in offline and online tool suites to predict, measure, and profile target application characteristics for a specified programming model and associated virtual machine architecture instantiation. This capability enables the understanding of new programming models and virtual machine architecture abstractions through simulation before hardware delivery. Another promising research topic is derived from the ability to perform online profiling correctness to support software bug tolerance and intrusion resistance capability. Finally, enhancements to operating and runtime systems will be required to support these objectives.

4. Interactions and Community Involvement

To realize the goals of this agenda, DoD must engage the national high performance computing community on many levels. From a programmatic perspective, DoD will coordinate with complementary programs conducted by DOE, NASA, NSF, and other government agencies to maximize resources and minimize duplication of effort. From a strategic perspective, high level DoD management attention will be required to ensure buy-in from industry leadership. From a bench-level perspective, DoD computational scientists from a wide range of organizations (research to operations) will collaborate with academia and industry on setting the technical agenda for the program. In the initial phases of this program, we envision active collaboration with the high performance computing R&D community in such areas as software tools, benchmarks, and device technology. Through programmatic coordination, strategic alliances with academia and industry, and active technical participation, DoD will ensure that this agenda tracks national security requirements and that *high productivity computing systems* solutions are delivered throughout the decade.