

## WHITE PAPER

---

### **Bridging the Capability Gap: Cray Pursues "Adaptive Supercomputing" Vision**

Sponsored by: Cray Inc.

---

Addison Snell

Christopher G. Willard, Ph.D.

Jie Wu

Earl Joseph, Ph.D.

February 2007

### **IDC OPINION**

High-end HPC users are looking for solutions that provide higher levels of performance on their specific applications. Standard processor speed improvements have slowed to a snail's pace, so users need to look to other architectural approaches to provide performance improvements. This has created a gap between standards-based products and the needs of high-end HPC end users. To help users bridge this gap, Cray Inc. is pursuing the following strategy:

- ☒ Cray is positioned to deliver on an "Adaptive Supercomputing" vision that delivers multiple HPC capabilities in a single integrated architecture. This approach is also referred to as heterogeneous processor-based computers or hybrid-computers. Adaptive supercomputers can improve users' productivity by providing technology designed to address a wider range of application requirements.
- ☒ Cray has experience and intellectual property in four processing paradigms — scalar, vector, multithreading, and attached coprocessors — that excel at different types of workloads.
- ☒ Cray's vision for Adaptive Supercomputing is targeted to users whose range of HPC applications requires varying dimensions of performance and scalability. IDC expects the heterogeneous nature of application requirements to remain a common problem in the industry.
- ☒ Cray is planning to use the Linux operating system, with the goal of delivering an adaptive supercomputer that can be integrated into most enterprises in a straightforward manner. Cray is also pursuing compiler technology that would hide the complexity of programming for different processing models, with the goal of increasing user productivity.
- ☒ The U.S. government agency DARPA recently selected the Cray Adaptive Supercomputing design as one of the approaches in which to invest \$250 million in R&D in order to have the type of computers it needs by 2010.

## IN THIS WHITE PAPER

This IDC White Paper analyzes Cray's vision for Adaptive Supercomputing and the role it can play in addressing current and future needs in HPC. By integrating its current disparate product architectures into a cohesive product with multiple capabilities, Cray is aiming to address different dimensions of performance with a single supercomputer that can adapt to different applications' needs. This white paper discusses the advantages of Adaptive Supercomputing and some of the challenges and opportunities Cray will encounter as it moves forward with its vision.

## SITUATION OVERVIEW

Cray's Adaptive Supercomputing vision capitalizes on current HPC market trends and current product expertise within Cray. Adaptive supercomputing is a powerful concept in the industry today, and Cray's current product lines put the company in position to move forward with the concept.

---

### Market Trends

In 2006, the HPC market grew to over \$10 billion in revenue, recording its fourth year of strong growth. One of the primary contributors to this growth has been the incorporation of commodity components (i.e., Intel x-86 based processors, Ethernet and InfiniBand interconnects, etc.) into systems primarily designed for non-HPC markets, that are able to capitalize on the economies of scale to establish new price/performance points in the industry.

One effect of the move to standard processors is that system performance increases are driven almost entirely by Moore's law improvements in general-purpose processor performance and by increasing the number of processors in a given system. Although these approaches have led to significant increases in performance over time, they can be limited by such factors as:

- ☒ Trade-offs between clock speed and power consumption
- ☒ Imbalances between processor performance and memory performance
- ☒ Trade-offs between system density and heat dissipation
- ☒ Application scalability

In addition, the move to leveraged products has tended to channel efforts to increase system performance along dimensions of increased clock speed and higher processor counts, leaving other paths unexplored.

IDC discussions with end users indicate that many organizations, particularly at the high end of the market, believe that the current approaches to increasing system performance will not meet the future requirements. These discussions have led IDC to believe that a "capability gap" between next-generation HPC requirements and system capabilities may be developing.

## ***Four Paths to Processor Performance***

The HPC industry has seen several processing paradigms, each of which is still pertinent today. Today's computational engines fall into four main categories: general-purpose scalar processors, vector processors, multithreaded architectures, and attached coprocessors.

**General-purpose scalar processors**, which initiate one discrete command at a time, are the most familiar engines. Scalar processors are designed to perform general workloads efficiently and affordably, and they are found at the heart of PCs, workstations, servers, and the majority of today's HPC systems. Scalable HPC systems with multiple scalar processors first entered the market in the late 1980s, and improvements in symmetric multiprocessing (SMP) architectures fueled their growth through the 1990s. Over time scalar processors incorporated many of the techniques used in vector processors, such as pipelined functional units and overlapped instruction execution. Such advances, combined with continued increases in clock speed, led to very capable systems by the end of the decade.

Beginning in the late 1990s, two distinct but interrelated trends boosted scalar processing even further. Rather than purchasing larger, monolithic systems, users began building clusters of smaller systems that communicated over a network to solve HPC problems. For programmers who could figure out how to exploit coarse-grained parallelism in their applications, clustering reduced the cost inherent to system infrastructure. The other trend involved the use of commodity microprocessors that are not proprietary to any system vendor, which again lowered the system cost and provided regular clock-speed increases. More than half of all system revenue could be attributed to clusters in 2006, and commodity microprocessors from Intel or AMD made up over 90% of the microprocessor volume.

Scalar processing, driven by trends in clustering and commoditization, has established new levels of price/performance in the industry, but the lack of specialization has left holes in the market for users who demand greater capabilities. Other processing paradigms offer users options for increased performance.

**Vector processors** formed the basis of traditional supercomputers. Twenty years ago, vector processors were at the heart of all major supercomputers, and some of the world's most significant supercomputing sites still rely on vector processing today. The advantage of vector processors is that they can stream sets of similar data items (i.e., "vectors") from memory to arithmetic units of the processor and return results to memory at rates close to the peak speed of the processor. Thus, complicated calculations can be completed quickly, with relatively few processors. Vector processors effectively exploit fine-grained parallelism within codes. Applications that involve repeated calculations over large data sets, such as global climate modeling or virtual stress testing of a complete aircraft model, can benefit significantly from vectors. The disadvantages to vectors are their relatively high cost per processor and, in recent years, their lack of commonality with the other systems in an IT enterprise.

**Multithreaded architectures** switch between threads very quickly in order to hide memory latencies. Multithreaded architectures also have the potential to run independent threads in parallel. Thus, whereas vectors process a single command across multiple pieces of data, multithreaded processors execute multiple

independent calculations. They therefore operate as scalar processors, but they can execute many scalar threads concurrently. Their primary advantage is latency tolerance due to the abundance of available threads to execute, especially for codes with irregular memory reference patterns, where individual threads spend most of their time waiting for memory responses.

**Attached coprocessors**, including field-programmable gate arrays (FPGAs), work in conjunction with a system's regular processors (usually scalar) to accelerate computationally demanding kernels via direct hardware execution. FPGAs can be reconfigured at runtime to perform a particular operation or calculation in dedicated hardware circuits. For example, an application that is highly dependent on fast Fourier transforms (FFTs), such as a signal processing application, could benefit from FPGAs. The FPGAs would be connected to the system, either through I/O or directly to memory. The system microprocessors would handle all other tasks, sending the FFTs to the FPGAs as required.

Although coprocessor elements can offer obvious performance benefits, sometimes by multiple orders of magnitude, their adoption has been slow because they require a change to program structure and usually a new programming language. Nevertheless, Cray and other vendors have introduced products with integrated FPGAs, and FPGA developers are working on productivity enhancements that make coprocessors easier to adopt. A key to successful system design is placing the FPGAs close to the other processors and main memory and not on a remote bus.

### ***Striving for Productivity***

Productivity is perceived to transcend performance to encompass the complete experience, including programmability, optimization, portability, manageability, support, and system robustness. The commoditization of hardware components and the resulting capability gap have many users in the HPC community seeking measurements of performance that look beyond the capabilities of the machine to the overall productivity of the users. This process involves considering the multiple dimensions of scalability required by the application, as well as the user's ability to utilize the system's full capabilities in a transparent fashion (i.e., without needing expertise in programming to the computer's individual architecture).

As discussed in the previous section, different workloads can benefit from different processing techniques. In addition, applications may have different requirements in terms of I/O scalability and data management. Most users have a range of applications that in many cases have different scalability profiles. Furthermore, the proliferation of clusters has provided a solid base for high-performance computing at a strong price/performance point, but there has been a resulting increased burden placed on users and application developers as the responsibility for scalability and optimization has passed from hardware to software.

The capability gap is strongly recognized in the user and vendor communities. Both first- and second-tier vendors offer middleware solutions for cluster optimization. In 2004 and 2005 the U.S. Council on Competitiveness commissioned IDC to study barriers to scalability among the HPC ISV community. These studies found that HPC is critical to industrial competitiveness and that most users have a business need for increased capability.

In addition, the Defense Advanced Research Projects Agency (DARPA) embarked on an initiative called the High Productivity Computing Systems (HPCS) program, with the goal of funding the design of a new generation of high-productivity supercomputers that will offer higher performance and ease of use. In late 2006, Cray was selected as one of two vendors to receive HPCS Phase III funding to bring its designs to market. Cray's successful proposal was centered around its Adaptive Supercomputing vision.

### **Cray's Vision for Adaptive Supercomputing**

To address the industry's need for higher productivity supercomputing, Cray is promoting a new vision it calls adaptive supercomputing. In its full vision for adaptive supercomputing, Cray plans to incorporate all of its currently disparate computing architectures into a single platform that can adapt to different application workloads in a user-transparent fashion.

Cray is unique in that it has products in each of the four processing paradigms discussed above. The Cray X1E continues Cray's legacy of vector processing. Cray has sold large X1E configurations to the Korea Meteorological Administration and to Oak Ridge National Laboratory, among others. The Cray XT3 and XT4, which occupy four of the top 20 spots on the November 2006 Top500 supercomputers ranking, use x86 64-bit AMD Opteron single- or dual-core microprocessors in a massively parallel processing (MPP) architecture. These scalar microprocessors employ HyperTransport technology for high-speed communication between nodes. The Cray XD1 allows the combination of Opteron processors with FPGA accelerators, and the Cray MTA product line uses multithreaded processors (see Table 1).

**TABLE 1**

Current Cray Systems		
System	Processor/Architecture	Notable Customers
Cray X1E	Vector	Oak Ridge National Laboratory; Korea Meteorological Administration
Cray XT3, XT4	Scalar MPP	Sandia National Laboratories; Atomic Weapons Establishment
Cray XD1	Scalar plus FPGA	Sony; Rice University
Cray MTA-2	Multithreaded	Naval Research Laboratory; Electronic Navigation Research Institute (Japan)

Source: Cray Inc., 2006

In the first phase of its Adaptive Supercomputing vision — code-named Rainier — Cray will provide a unified user environment across all its platforms, using Opteron processors and the Linux operating system on top of the XT system infrastructure. This environment will provide a necessary first step in bringing the different systems together by providing a single log-in, a global file system spanning the different compute platforms, and the ability to launch applications on different types of nodes.

In the next phase, Cascade plans to build a single integrated system with heterogeneous processing blades. This system will allow programmers to leverage different types of processing capabilities within a single application, thus bringing multiple computing technologies to bear on a single problem. It will also provide benefits for organizations that run different types of applications; the single heterogeneous architecture would allow each application to be optimized according to its own needs.

## **CHALLENGES/OPPORTUNITIES**

Cray's full vision for Adaptive Supercomputing is extremely bold and requires a sizable amount of R&D. As such, it is subject to significant challenges and risks in bringing it to market. Cray has already crossed one major hurdle in acquiring \$250 million in HPCS funding. The Phase III selection was important not only for the influx of R&D dollars, but also because it provides an external validation of plan for increased productivity.

With its DARPA HPCS Phase III win, Cray is pursuing both a common heterogeneous architecture — code-named Cascade — and the compiler technologies that would optimize applications to take advantage of the multiple system capabilities in a user-transparent fashion. For example, the compiler would automatically recognize portions of the code that are well-suited to vector processing and create vectorized binaries for them. The goal is to create a combined executable with different binary formats tightly integrated together, as if the application were composed of a single binary format.

Cray's ultimate plan is to deliver a system that adapts to the application, rather than forcing the application to adapt to the system. The incorporation of multiple processing capabilities, with automatic, user-transparent optimization, in a familiar Linux user environment would provide benefits that transcend performance to better user productivity.

Until Cascade is eventually delivered, Cray will continue to be subject to market forces. The continued maturation of technologies such as FPGAs would benefit Cray. At the same time, Cray should hope that the vector market will remain relatively stable until adaptive supercomputers come to fruition.

Cray's main opportunity lies in the potential for helping to initiate and participate in a new technology cycle within the high-end technical computing markets. We believe that technologies that provide a discontinuous jump in delivered performance can spark new buying cycles within the high-end segments of technical computing markets and that these technologies can then move down-market to replace older alternatives. If Cray can capture a technology leadership position in adaptive supercomputing, it can then leverage that leadership to capture share, first in the high end of the market and then within larger-capacity markets.

## **CONCLUSION**

There is a clear market need for the performance advances that Cray is pursuing through its Adaptive Supercomputing strategy. However, the plan is bold, and Cray will need to overcome many technological barriers to deliver on the full vision.

The potential risk is great, and so is the potential reward. Many users could benefit from the performance improvements alone, and the additional benefits to productivity would be a boon to the industry. As long as the costs are controlled, Cray should find no shortage of interest in the user community.

---

## **Copyright Notice**

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2007 IDC. Reproduction without written permission is completely forbidden.