FRONTIERS OF EXTREME COMPUTING 2007

SAND2008-0231C

Nanoscale to Zettascale

Chaminade Executive Conference Center,
October 21-25, 2007, Santa Cruz, California

http://www.zettaflops.org

ZETTA FLOPS

Experts representing computer and computational science will meet to create a comprehensive picture of future computing.

TECHNOLOGIES

Multi-core
Nanotechnology
Quantum Computing
Advanced Algorithms
Advanced Architectures
Systems Software

APPLICATIONS AREAS

Scientific Research
Environment Science
Medical Applications
Defense Problems
Robotics
Spacecraft

Organizing Committee:

William Camp, Intel
Candace Culhane, Department of Defense
Mike Foster, National Science Foundation
Jag Shah, DARPA
Horst Simon, Lawrence Berkeley National Laboratory
Rick Stevens, Argonne National Laboratory
Thomas Theis, IBM
Stanley Williams, Hewlett-Packard
David Womble, Sandia National Laboratories

General Conference Chair: Erik DeBenedictis, Sandia National Laboratories

Deputy Conference Chair: Thomas Sterling, Louisiana State University

Organizing Institutions:
Argonne National Laboratory, DARPA, Department of Defense, Hewlett-Packard, IBM, Intel, Lawrence Berkeley National Laboratory, Louisiana State University, National Science Foundation, Sandia National Laboratories

Financial Support:

Sandia National Laboratories · CSRI Computer Science Research Institute · intel Leap ahead · hp invent · DARPA · Berkeley Lab · CRAY

Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Frontiers of Extreme Computing 2007
Zettaflops Workshop
Nanoscale to Zettascale

1. **Date:** October 21-25, 2007
2. **Location:** Chaminade Conference center, Santa Cruz, CA
3. **Participation:** By gracious invitation only
4. **Sponsors:** Sandia, Intel, Hewlett-Packard, DOE Office of Science, DARPA, Lawrence Berkeley National Laboratory/NERSC, Cray
5. **Steering Committee:**

Government Sector	Industry
Candace Culhane, DoD	William Camp, Intel
Mike Foster, NSF	Thomas Theis, IBM
Jag Shah, DARPA	Stanley Williams, HP
Horst Simon, LBL	
Rick Stevens, Argonne	
David Womble, Sandia	

6. **General Conference Chairs:** Erik DeBenedictis, Sandia
   Thomas Sterling, Louisiana State University

Organization of the Workshop

Erik P. DeBenedictis

This note was written right before the workshop to give attendees insight into the overall organization and a context to interpret the talks. This workshop and the predecessor Petaflops Workshops dating back to 1994 include an unusually broad range of speakers – from the underlying device technology to the impact of computer applications on society. The workshops focus on whether computer technologies can be integrated at some point in the future in order to make powerful computers that solve important problems. With the foreknowledge available to the organizer, this workshop promises to reach important and perhaps unexpected conclusions. The workshop appears headed toward showing that some of the most ambitious, compute-intensive applications will require more than just the planned extrapolations of current computing technology. This document gives the reader some tips on what to look for in order that he or she can form their own conclusions.

Introduction

The diagram at the bottom of this page illustrates one part of the workshop's organization. The vertical axis is computer throughput for a supercomputer on the left and a laptop, robot controller, or spacecraft on the right. The workshop will include applications talks including climate modeling, personalized medicine, and mobile computing as shown on the left half of the diagram. Climate modelers are steadfast that they need 1 Zettaflops to meet their needs, whereas the needs of the other applications are unknown at present and may be determined at the workshop (symbol ⇳). Technologists will reveal various scenarios for the growth of computer performance (Moore's Law) and limits, as illustrated on the right half of the diagram. A key theme of the workshop will be to determine if the demands of specific applications are above or below the capabilities of the implementation technologies.

Historical Perspective

The 1994 Petaflops Workshop asked whether a petaflops computer was even possible. There was some question about this at the beginning of that workshop, but by the end it was clear that a petaflops computer would be possible. The 1994 workshop led to a "straw man" petaflops supercomputer called HTMT and based on superconducting technology. In the 13 years since that workshop, the community has essentially reached the petaflops goal with the 1/3 petaflops Blue Gene supercomputer. The 1994 CMOS

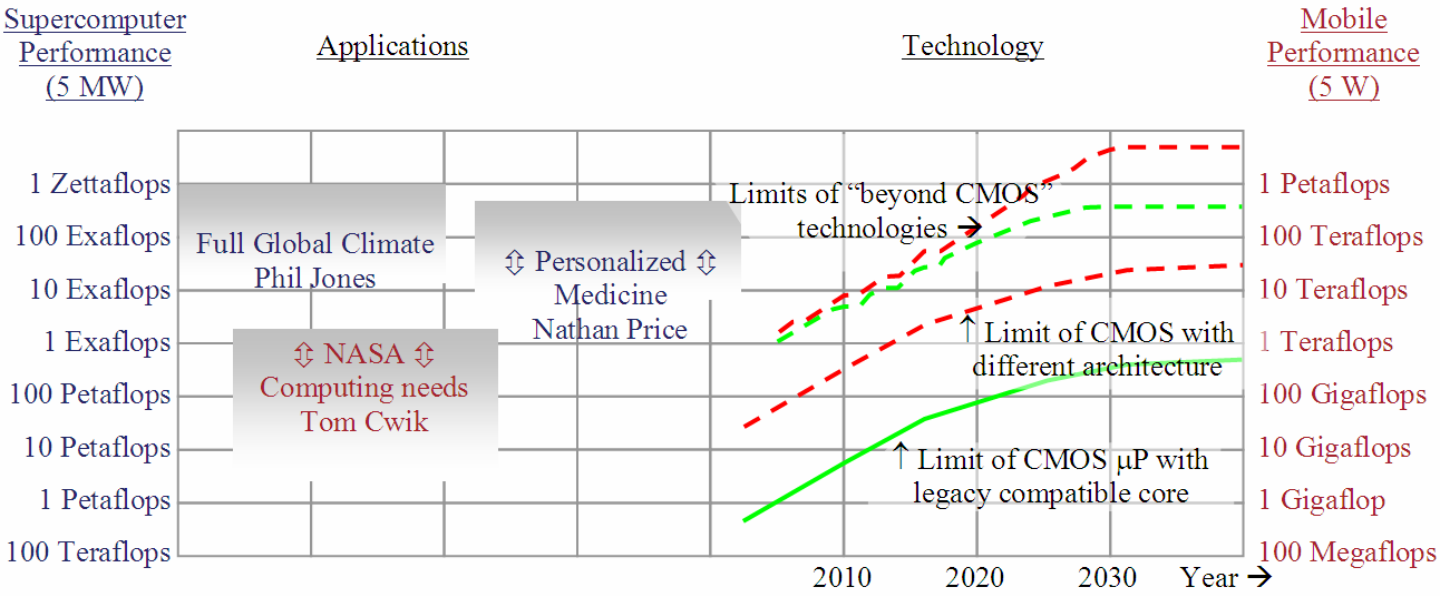

roadmap was destined for a long exponential rise, delivering performance gains in large part through trivially exploited clock rate increases. As a consequence, "regular" CMOS technology powers Blue Gene instead of any new specialized (superconducting) technology.

This is a relevant historical perspective because this workshop seems headed towards a different conclusion. The physical science speakers are poised to tell us CMOS is approaching maturity, and the computer systems speakers are coping with the fact that the remaining CMOS performance gains are coming in the form of hard to exploit parallelism.

Let me now give a preview and context for the presentations. I will establish the applications goal and then the technology available to implement it.

Applications

There are three talks on applications, one of which will call for a 1 zettaflops supercomputer. The other applications may be equally challenging.

- Phil Jones will give a talk on climate modeling on Monday at 11:30 AM. As the organizer, I have spoken to Phil and he says his talk will conclude that 1 zettaflops will be needed to understand the origins and mitigations of Global Warming. This is of course a problem of profound importance and interest.

- Nathan Price will speak 2 PM Tuesday on personalized medicine. This is a proposed form of medicine where computers calculate a cure for an individual's affliction. While it is not a zettascale calculation, its cost accrues to a single sick person and must be in the range of hundreds of dollars to be useful for improving human health on a broad scale. The cost limitation may make this application as challenging of the underlying technology as the construction of a zettaflops supercomputer

- Tom Cwik will speak Wednesday 11:30 AM on computing from the JPL perspective, which includes robotics and spacecraft computing. The mobile domain is power constrained to a few watts, or about one millionth the power constraint on supercomputers. A one-millionth scale supercomputer seems a viable model for future mobile computers for robotics and space applications. The power limitation may also make this application just as challenging as a zettaflops supercomputer.

Technology

There are three talks on the underlying technology, all of which discuss difficulties in meeting the zettaflops goal. These are talks by Tom Theis of IBM Monday 9:30 AM, George Bourianoff of Intel Monday 11:00 AM, and Stan Williams of HP Tuesday 9:15 AM. Two of the three abstracts appear later in this document, and they express optimism for continued progress but question whether a zettaflops is achievable. There also appears to be a consensus that a zettaflops supercomputer based on CMOS will not be possible. The third talk (I happen to know) has a similar message. I am urging each participant to listen to these technology talks and decide for them self how far CMOS or other logic options will go.

Other Technology, Architecture, and Software

The logic technology discussed might be powerful enough for the most challenging applications, but whether it really turns out that way will depend on the middle layers of the technology stack. In fact, the middle layers of the technology stack burn away several orders of magnitude through inefficiency, and a breakthrough on the positive side in the middle could have as much of an effect as a better physical technology. The key issues to look for are efficiency and compatibility, discussed below.

The talks are:

Interconnect	Jag Shah	Tue. 10 AM
Storage	Gian Luca Bona	Wed. 2 PM
Architecture	Karu Sankaralingam	Wed. noon
Software	Kathy Yelick	Wed. 11:30
Algorithms	David Keyes	Wed. noon
Systems software	Thomas Sterling	Mon. 2 PM

Today's CMOS technologies certainly work, and they work with moderate efficiency. As CMOS scales to it ultimate limits (per the talks by IBM, Intel, and HP), how will these intermediate technologies retain or lose efficiency? If any technology layer becomes very efficient with scaling, the ultimate potential of CMOS-based computers will decline. Ultimate performance can increase if efficiency in a layer increases. I suggest each participant listen to each of these talks and decide for them self how much efficiency there will be at say 1, 10, 100 petaflops and 1 exaflops (for a supercomputer, proportionately lower values for other classes of computers).

The diagram at the bottom of the page is perhaps a helpful illustration of how the technologies' performance fits together. From the left, the application (ground-based supercomputer vs. mobile device) defines the total power available. The logic technology will turn this into binary logic throughput – including both the number of gates (which will define the amount of parallelism in the system) and the speed of each gate. The stages to the right in the diagram are described below.

- Gian-Luca Bona of IBM Wednesday 2 PM will discuss future storage options that have greater efficiency than hard drives. Storage, I/O, and some interconnect join logic as independent resources at the lowest level.

- Delays and power consumption of on-chip electrical wiring is causing loss of efficiency as chips scale. Jag Shah of DARPA will discuss Tuesday 10 AM a program to develop on-chip optical interconnect that will maintain efficiency as chips scale to larger numbers of components. This can stem the loss of throughput shown as "interconnect loss."

- Karu Sankaralingam of Wisconsin Wednesday noon will discuss the highly multicore TRIPS architecture and other nontraditional information processing architectures that may follow. These architectures strive to deliver as much of the raw information processing resource as possible to the end user, although features for easy programming (i. e. big caches) are desirable as well but have a cost.

- Kathy Yelick will speak on Tuesday

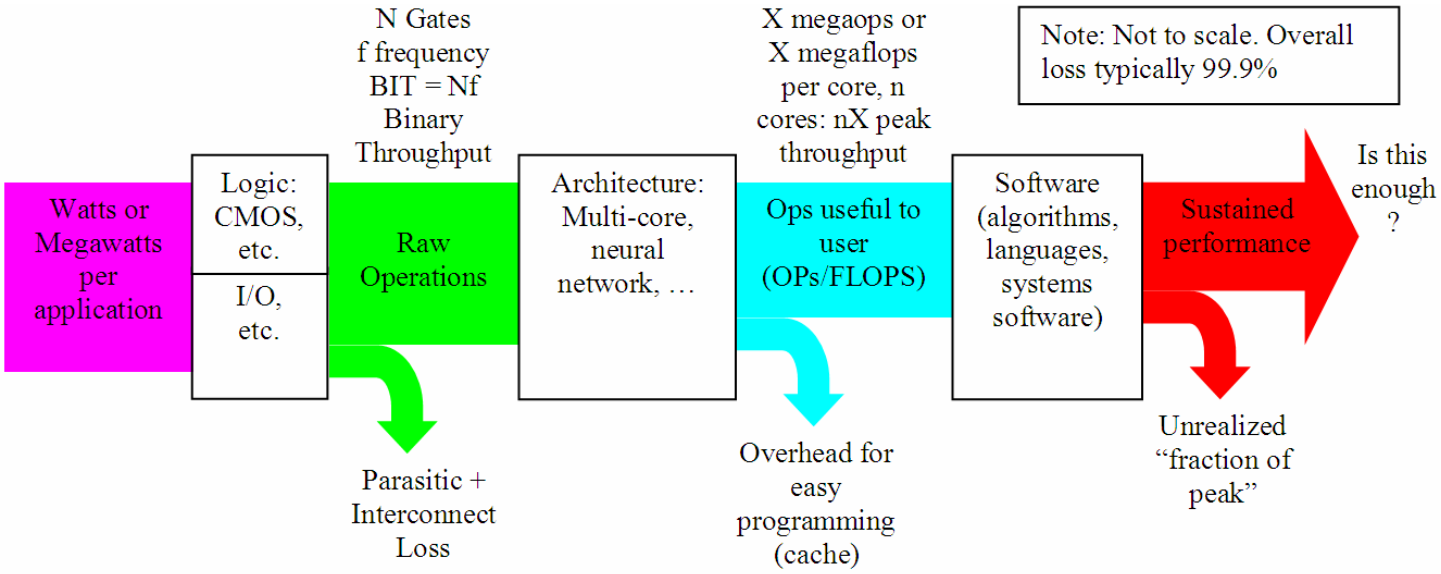

11:30 AM to the issue of programming highly parallel systems. As computers get more-and-more cores, programmers must exploit more-and-more parallelism to realize performance gains in the underlying gates. Aside from keeping programmers productive (in terms of getting applications coded in a reasonable time as parallelism increases), the problem to address here is loss of efficiency due to the underlying hardware being "stalled." Solutions will require better tools, languages and programmer training, which Kathy will discuss.

- David Keyes will discuss increasingly scalable algorithms Tuesday at noon. Most algorithms have a fixed amount of parallelism as a function of problem size, causing a drop in efficiency (reduced "fraction of peak") when the computer has more cores than there is parallelism in the program. With effort, it is often possible to find algorithms with more parallelism.

- Thomas Sterling will speak Monday at noon on operating systems that can manage the more highly parallel programs of the future. Thomas will also speak about operating systems methods that will retain expected percentages of uptime as the number of components that can fail increases.

What's left at the right side of the diagram is then available to the application for solving some important problem. Is it enough?

The talks by IBM, Intel, and HP also talk about "beyond CMOS" logic. The tip here is that none of the beyond CMOS options are particularly compatible with what we do today. With reference to the diagram on the previous page, beyond CMOS options (More than Moore, reversible logic, etc.) are not drop-in replacements for logic gates and cause incompatibilities as one moves across the diagram. It may be possible to address the incompatibilities with extra logic. However, this extra logic has a cost that should be weighed against the proposed advantages of the beyond CMOS logic.

I'm suggesting that each participant listen to the talks and decide for themselves the degree to which the technology is compatible with the various beyond CMOS options (spin-gain transistors, reversible logic, quantum computing, etc.) In many cases, I suspect you will find no compatibility and have no idea on how to proceed. These could become topics for future research programs.

Program

Monday 8:15 AM – Registration Desk Opens

Monday 9:00 AM – *Workshop Introduction, Overview, and Goals*
Erik P. DeBenedictis, Sandia

Monday 9:30 AM – *Prospects for Computing Beyond CMOS*
Thomas N. Theis, Director, Physical Sciences IBM Research, T.J. Watson Research Center

In the last few years, microprocessor clock speeds have ceased their historic rate of exponential advance. The root cause is the failure of traditional scaling rules for devices at the 90 nm generation and beyond, resulting in system designs that are severely constrained by economic limits on allowable power dissipation. While significant further improvements in CMOS FET technology are likely, I argue on physical grounds that dramatic advances in processor clock speeds will only be possible by moving to energy-conserving (also called "adiabatic" or "reversible") circuit designs. Careful studies in the mid-nineties showed that such circuits offer no advantage over conventional low-power circuits when built with CMOS transistors. Today, we see an explosion of interest in "post-CMOS" devices. I argue that a new device, much better suited to energy-conserving circuits than the FET, is physically possible, and I discuss the necessary attributes of such a device.

Monday 10:00 AM – *E3 Exascale Initiative*
Horst Simon, LBL

Monday 10:30 AM – Break

Monday 11:00 AM –*More Moore, More than Moore, beyond CMOS and the ITRS*
George Bourianoff, Intel

Monday 11:30 AM – *Climate Modeling on Future Architectures*
Philip Jones, Los Alamos National Laboratory

Future plans for climate models will be presented with a focus on future computing requirements. Higher resolution, addition of new physics and simulating biogeochemical cycles will all be important thrusts and each will stress new architectures in different ways. We will then speculate on how ideas for future architectures would impact how we design and implement models for future climate change studies.

Monday Noon – *Operating Systems for Exascale Computing*
Thomas Sterling, Louisiana State University

As computer systems technology and architecture have evolved, so have operating systems (OS). OS has historically scheduled jobs, allocated resources, virtualized low level devices, and managed file name spaces. Future high performance computing (HPC) systems will depart dramatically from conventional systems in that multicore, multithreaded, and heterogeneous accelerators will replace single microprocessors for total system parallelism that exceeds billion-way concurrency in hardware and software. In addition system reliability will be a major aspect of system control with system configuration adjusting potentially every few minutes. A new generation of system-wide OS may contend with these new structures replacing the ensemble computing strategy of one OS per node. This presentation will discuss one possible general strategy for future OS design and operation based on lightweight kernels on local

elements operating in synergy across systems to present a single albeit distributed set of functionalities to parallel applications. Such a strategy provides a virtual global system rather than a collection of virtualized local nodes, managing the changing underlying resource configuration to achieve graceful degradation in the presence of faults while responding to adaptive resource demands of new classes of dynamic applications.

Monday 12:30 PM – *Lunch*

Monday 1:30 PM – *Challenges to reaching Exascale computing levels*
Horst Simon, LBL

Monday 2:00 PM – *Charter to the Working Groups*
Thomas Sterling, Louisiana State University

Monday 2:30 PM – Working Group Session 1
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Monday 3:30 PM – Break

Monday 4:00 PM – Working Group Session 1
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Monday 6:30 PM – Dinner

Tuesday 9:00 AM – *Workshop Introduction, Overview, and Goals*
Erik P. DeBenedictis, Sandia

Tuesday 9:15 AM – *Sprinting Toward the Practical Limits of Computation*
Stan Williams, HP Labs

There have been many incorrect predictions about the end of computational scaling, but this will not be one of them. It is within the realm of fundamental physics to improve computation dramatically beyond where it is now (petaflops), but if a zettaflop system ever exists, it is surely many decades in the future. We can only count on strict Moore's Law scaling to improve computational efficiency by a factor of 100 or so, and even then at the cost of having thousands of cores on a single chip. Replacing electronic interconnect by photonic interconnect at all length scales over a few microns might improve computational efficiency by another factor of 100 over the long term. However, getting that last factor of 100 will require a large number of tricks and technologies for which we have only an inkling now. This talk will provide a brief overview of some of the challenges and the tricks for building future ultra-scale computers.

Tuesday 10:00 AM – *Intrachip Photonic Communications*
Dr. Jag Shah, DARPA

This talk will discuss DARPA’s interest in exploring how ultradense functional units on a chip communicate with each other and the outside world

Tuesday 10:30 AM – Break

Tuesday 11:00 AM – *Future Supercomputer Architectures*
Steve Scott, CTO, Cray Inc.

This talk will explore the evolution of high-end computing architectures from the 100 TF scale we are at today through the petaflops and exaflops regimes. Key challenges and suggested paths to reach the exascale will discussed, and perhaps some musings about zettascale will creep in (though the speaker admits he really has no clue how we might get there).

Tuesday 11:30 AM – *Programming Exaflops*
Kathy Yelick, LBL/UC Berkeley

Tuesday Noon – *Scaling to Exaflop/s for Mesh-based Algorithms*
David Keyes, Columbia

Many scientific applications that push scientific supercomputing performance past its ever-rising high watermark possess as their central data structure a graph with local connections that comes from a mesh in some small number of dimensions (up to about six). This includes, for instance, systems of conservation laws in physical space, problems for distribution functions posed in phase space, lattice quantum chromodynamics, some circuit models, etc. The central kernels for the best solvers in such applications are essentially sparse matrix multiplications of dense vectors drawn from hierarchical subsets of the mesh, and vector-vector operations. Such applications are suitable for indefinite scaling, provided that the data structures are long-lived enough between adaptations to amortize the cost of creating good schedules for the required data movements. However, this does not mean that such applications are easy to program for performance. This talk is intended to lay out the relatively simple data motion consequences of the (sometimes relatively rich) underlying mathematics of the algorithms and help forge a dialog with architecture and language experts en route from petaflop/s, where programmer-directed message passing is a workable model, to exaflop/s.

Tuesday 12:30 PM – Lunch

Tuesday 1:30 PM – *Sandia's Programs in Supercomputing and Nanotechnology*
Sudip Dosanjh, Sandia National Laboratories

This presentation describes Sandia¹s work in capability systems and its plans for future systems. Application models and analyses are being used to analyze architectures and perform trade-off studies. Sanida¹s microprocessor fabrication facilities will also be discussed as well as opportunities to leverage these capabilities for High Performance Computing.

Tuesday 2:00 PM – *Computing Challenges for Systems Biology and Personalized Medicine*
Nathan Price, University of Illinois, Urbana-Champaign

The emergence of advanced technologies to generate large amounts of high-throughput biological data has catalyzed a phase transition in biology research where computation is becoming a key factor in biological discovery. These technologies, coupled with advanced computation, are also opening up opportunities for a future of personalized medicine, where patients receive care based on the computational analysis of information obtained from their personal genome sequence and the measurement of thousands of proteins in their blood at different time points. This talk will introduce some of the computational challenges for systems biology and personalized medicine that can be addressed through advances in high-end computing technologies.

Tuesday 2:30 PM – Working Group Session 2
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Tuesday 3:30 PM – Break

Tuesday 4:00 PM – Working Group Session 2
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Tuesday 6:30 PM – Dinner (Capitola Room)

Tuesday 7:00 PM – Banquet Talk *A Twenty Year Vision for the Convergence of Computation, Science and Computer Architecture*
Rick Stevens, ANL

I'll break out the crystal ball and chart a path for the next twenty years, I'll discuss what might be possible in computational science (with a focus on biological and environmental applications) and discuss the possibilities enabled by the melding of science and architecture. My goal is to challenge the current generation and perhaps inspire the next generation of computationally enabled scientists to think differently about what is possible.

Wednesday 9 AM – *eResearch in the Cloud: Data-Intensive High Performance Computing*
Tony Hey, Microsoft

This talk will be about the need for Extreme Computing to focus on Data-Intensive challenges as well as Petaflop systems. It will describe a future in which collaborative, multi-disciplinary frontier research is supported by data-intensive services 'in the cloud' as well as traditional software and systems.'

Wednesday 9:30 AM – *DARPA Exascale Initiative : goals, motivation, range of topics, players, and approach*
Bob Lucas, ISI

Wednesday 10:00 AM – *NSF High End Computing University Research Activity (HECURA) Program*
Almadena Chtchelkanova, NSF

In response to the High End Computing Revitalization Task Force (HECRTF) recommendations, http://www.nitrd.gov/subcommittee/hec/hecrtf-outreach/hec-ura/index.htm, an interagency group presently consisting of NSF, DOE Office of Science, NNSA, NASA, DARPA and NSA, started a pilot program called High End Computing University Research Activity (HECURA) to address basic research in the area of High-End Computing (HEC). In FY04 NSF and DARPA jointly funded Software Tools for High-End Computing program, focused on research and education projects in the area of software tools and compilers for HEC http://www.nsf.gov/pubs/2004/nsf04569/nsf04569.pdf.

In FY06 NSF, DARPA, EPSCoR and DOE Office of Science funded over 20 projects in the area of I/O, file and storage systems design for efficient, high throughput data storage, retrieval and management in the HEC environment submitted in response to the following solicitation http://www.nsf.gov/pubs/2006/nsf06503/nsf06503.pdf.

There are plans to continue HECURA program in FY08.

Wednesday 10:30 AM – Break

Wednesday 11:00 AM – *TBD*

Wednesday 11:30 AM – *NASA/JPL Future Computing Needs*
Tom Cwik, JPL

Wednesday Noon – *A System Perspective on End of Silicon*
Karu Sankaralingam, Univ. Wisconsin

The end for CMOS technology seems to be in sight. We are entering an era of non-ideal process scaling and unpredictable silicon which is causing disruptive changes at many levels of microelectronics system stack. Reliability, programmability, and efficiencies will dictate changes to these systems in addition to their massive computation needs. In this talk, I will discuss technology trends, implications for architects based on our experience in the TRIPS project, and promising directions for future systems.

Wednesday 12:30 PM – Lunch

Wednesday 1:30 PM – *Programming for the Future*
Rick Stevens

In this talk I'll discuss some ideas for progress on reducing the fragility of parallel programs, increasing the longevity of parallel implementations of applications, introduce some (hopefully) new ideas for scalability and introduce the seductive concept of scale invariant programming models.

Wednesday 2:00 PM – *Prospects for Solid State Data Storage: Beyond Flash Memory and the Hard Disk Drive*
Gian-Luca Bona, IBM

Wednesday 2:30 PM – Working Group Session 3
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Wednesday 3:30 PM – Break

Wednesday 4:00 PM – Working Group Session 3
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Wednesday 6:30 PM – Dinner

Thursday 9:00 AM – Working Group Session 4
Rooms: Software – Pebble, Technology – La Selva, Architecture – Santa Cruz, Software – Cowell

Thursday 10:30 AM – Break

Thursday 11:30 AM – Working Group reports
Santa Cruz room

Thursday 12:30 AM – Lunch

Thursday 1:30 PM – The workshop is officially over at the conclusion of lunch.

Working Groups

Guiding Questions

Cross Cutting Themes:

- Power
- Parallelism
- Performance
- Execution Models
- Efficiency
- Cost
- Reliability

Applications and Algorithms – Pebble Room
Chair: Horst Simon
Co-Chair: Mark Stalzer

Strategic Issues:

- Application drivers for Exascale and potential Science impact
- Application system resource requirements
- Algorithmic approaches to exploit Exascale system capability and structures

Detailed Questions:

- What will be important end-user applications requiring Exascale computing?
- What will be the needed memory capacity to support the range of Exascale applications?
- How much parallelism can be exposed for these applications and is it enough to exploit billion-way parallelism architectures?
- What are the expected I/O demands for mass storage including checkpoint/restart, initial conditions, intermediate results for runtime steering, and final results?
- What are the requirements for and possibilities of algorithms suitable for Exascale applications?
- For Exascale applications and their enabling algorithms what semantic constructs, basic mechanisms, and execution model attributes are needed?
- What discoveries or advances might be possible with Exascale computing that are not possible with Petascale or below. Also, what advances will be possible by applying Exascale-capable technology to smaller systems (i.e., robots, spacecraft)?
- What algorithm classes will become more or less important at Exascale as opposed to Petascale or lower. For example, will finite elements become more or less prevalent? How about Monte Carlo? Will larger computers be applied to larger simulations or large parameter sweeps over current-size simulations? While we're on the topic, will simulations be the prevalent application class for Exascale supercomputers?

Enabling Technologies – La Selva Room
Chair: Sheila Vaydia
Co-Chair: Erik Debenedictis

Strategic Issues:

- Limitations of conventional semiconductor and wired technologies and roadmaps in achieving Exascale computing – hence the need for alternatives
- Promising device technology innovations in logic, primary and secondary storage, and interconnect
- Nano-scale technology approaches and end of Moore's Law
- Over the horizon concepts

Detailed Questions:

- What are achievable clock rates or switching times at reasonable power consumption of competing logic technologies?
- What are the contending technologies that will lead to nano-scale and the end of Moore's law?
- What are the likely paths to intra-system communications including chip-to-chip and system-wide interconnections?
- If memory capacity will be one of the major constraining factors to application performance, what are the likely technologies and their operational characteristics towards Exascale and beyond?
- Is there life after disks? With the monotonically and potentially exponential growth in storable information, what is the likely or potential media technologies that will get us through the 21st century?
- What information processing structures may come to offer advantage over Boolean logic gates. Example: alternatives circuits, non charge-based information, quantum effects, ...
- Is there a future in systems that blend information processing and non-volatile memory? Examples may include artificial neural networks and molecular-scale FPGAs where the same physical structures learn at a low speed and process at a higher speed.
- How should quantum computing be integrated with the general information processing? What about quantum effects short of a traditional quantum computer?

Parallel Architectures – Santa Cruz room (main meeting room)
Chair: Larry Bergman
Co-Chair: Loring Craymer

Strategic Issues:

- Exploitation of emerging technologies
- Scalability to billion-way parallelism
- Addressing sources of performance degradation
- Usability

Detailed Questions:

- What architecture strategies will enable Exascale system implementation?
- What architecture approaches will address challenges of latency, overhead, contention, and starvation expected at Exascale system structures?
- It is estimated that multi billion-way parallelism (> 100 million cores) will be required for Exascale systems. What architectures can be expected to deliver, exploit, and coordinate such unprecedented parallelism?

- What architecture techniques will minimize power consumptions?
- Are new execution models required to guide architecture structure and operations, and if so what?
- In the new Exascale architectures, how will memory and address spaces be handled?
- Will the future hold more highly parallel traditional architectures or something new. To be specific, the traditional architecture is a homogeneous parallel system comprising x86-compatible processors with MMU and a large cache. Something new might be a heterogeneous system of heavy and light cores, where the light cores are GPUs, cell, attached neural networks, etc.

System Software and Programming Environments – Cowell Room
Chair: Ron Brightwell
Co-Chair: Jeff Vetter

Strategic Issues:

- Advanced execution models
- Parallel programming models, methods, and tools
- Resource management, allocation, and scheduling
- Mass storage and I/O management

Detailed Questions:

- What new semantic constructs and execution models will be needed by Exascale applications and algorithms? What will the programming languages of the future look like?
- How will programmers contend with billion-way parallelism? What will the new languages look like?
- What operating system organizing strategy will effectively manage 100 million or more processing cores efficiently and reliably?
- How will compilers and runtime systems support the new classes of applications dominated by dynamic meta data structures?
- What will be the balance in the future between user direct control of resources and system automation for ease of use?
- Will programmers continue to program with arithmetic statements, or may a different paradigm become prevalent? Examples, neural networks, fuzzy logic, graph algorithms, image processing, real time?

Program at a Glance

	Monday October 22	Tuesday October 23	Wednesday October 24	Thursday October 25
8:15 am	Registration			
9:00 am	*Workshop Introduction, Overview, and Goals*, E. DeBenedictis, Sandia	*Workshop Introduction, Overview, and Goals* E. DeBenedictis, Sandia	Industry: *eResearch in the Cloud: Data-Intensive High Performance Computing*, Tony Hey, Microsoft	WG Session 4: Develop report Software: Pebble Technology: La Selva Architecture: Santa Cruz Software: Cowell
9:30 am	Industry: *Prospects for Computing Beyond CMOS Logic*, Thomas Theis, IBM	Industry: *Sprinting Toward the Practical Limits of Computation*, Stan Williams, HP	Government: *DARPA Exascale Initiative: goals, motivation, topics, players, and approach*, Bob Lucas, ISI	
10:00 am	Government: *E3 Exascale Initiative*, Horst Simon, LBL	Government: *UNÍC: Intrachip Photonic Communications*, Jag Shah, DARPA MTO	Government: *NSF High End Computing Univ. Research Activity Program*, Almadena Chtchelkanova	
10:30 am	Break	Break	Break	Break
11:00 am	Technology: *More Moore, More than Moore, beyond CMOS and the ITRS*, George Bourianoff, Intel	Industry: *Future Supercomputer Architectures*, Steve Scott, Cray	Technology: Vacant	Working Group Reports Santa Cruz Room
11:30 am	Applications: *Climate Modeling on Future Architectures*, Phil Jones, LANL	Technology: *Exaflops Programming (title TBD)*, Kathy Yelick	Applications: *NASA/JPL Future Computing Needs*, Tom Cwik, JPL	
Noon	Technology: *Operating Systems for Exascale Computing*, Thomas Sterling, LSU CCT	Technology: *Scaling to Exaflop/s for Mesh-based Algorithms*, David Keyes, Columbia	Technology: *A System Perspective on End of Silicon*, Karu Sankaralingam, U. Wisconsin	
12:30 pm	Lunch	Lunch	Lunch	Lunch
1:30 pm	*Challenges to reaching Exascale computing levels*, Horst Simon, LBL	*Sandia's Programs in Supercomputing and Nanotechnology*, Sudip Dosanjh, Sandia	*Programming for the Future*, Rick Stevens, Argonne National Laboratory	
2:00 pm	*Charter to the Working Groups*, Thomas. Sterling, LSU	Applications: *Computing Challenges for Systems Biology and Personalized Medicine*, Nathan Price, UIUC	Technology: *Prospects for Solid State Data Storage: Beyond Flash Memory and the Hard Disk Drive*, Gian-Luca Bona, IBM	
2:30 pm	WG Session 1 Software: Pebble Technology: La Selva Architecture: Santa Cruz Software: Cowell	WG Session 2 Software: Pebble Technology: La Selva Architecture: Santa Cruz Software: Cowell	WG Session 3 Software: Pebble Technology: La Selva Architecture: Santa Cruz Software: Cowell	
3:30 pm	Break	Break	Break	
4:00 pm	WG Session 1	WG Session 2	WG Session 3	
6:30 pm	Dinner	Dinner, Capitola Room	Dinner	
7:00 pm		Banquet: A Twenty-year vision… Speaker: Rick Stevens, Argonne Nat. Lab.		
7:45 pm	Social	Social	Social	