U.S. DEPARTMENT OF ENERGY

Office of Science

Update on Exascale Research

William Harrod

William.Harrod@science.doe.gov

ASCR

August 15, 2011

U.S. DEPARTMENT OF ENERGY
Office of Science

Outline

- Current Programs
- Workshops
  - Programming Challenges
  - Architecture I
  - Architecture II
- Anticipated Future Programs
- Future Workshops

U.S. DEPARTMENT OF ENERGY
Office of Science

Current Exascale Programs

- Advanced Architectures and Critical Technologies for Exascale
  - 6 projects focused on power management, memory management, and reducing the cost of data movement
- R&E Prototpyes
- X-Stack Software Research
  - 10 projects focused on operating systems, fault tolerance, programming challenges, performance optimization, etc.
- Scientific Data Management and Analysis at Extreme Scale
  - 10 projects spanning file systems and I/O, data triage, feature detection and data analysis, and visualization

U.S. DEPARTMENT OF ENERGY
Office of Science

Exascale Co-Design Centers

Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx)

Director: Timothy Germann (LANL)

Center for Exascale Simulation of Advanced Reactors (CESAR)

Director: Robert Rosner (ANL)

Combustion Exascale Co-Design Center (CECDC)

Director: Jacqueline Chen (SNL)

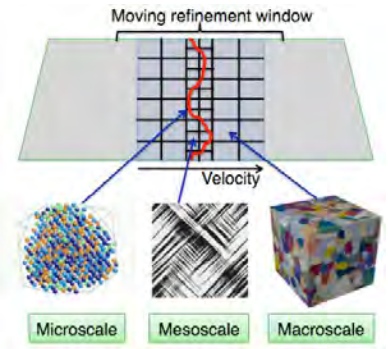

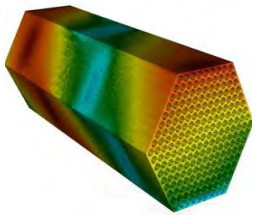

	ExMatEx (Germann)	CESAR (Rosner)	CECDC (Chen)
National Labs	LANL	ANL	SNL
	LLNL	PNNL	LBNL
	SNL	LANL	LANL
	ORNL	ORNL	ORNL
		LLNL	LLNL
			NREL
University & Industry Partners	Stanford	Studsvik	Stanford
	CalTech	TAMU	GA Tech
		Rice	Rutgers
		U Chicago	UT Austin
		IBM	Utah
		TerraPower	
		General Atomic	
		Areva	

U.S. DEPARTMENT OF ENERGY
Office of Science

Exascale Workshops

- Programming Challenges Workshop
  - Understand, prioritize, and shape the future research agenda that addresses challenges for Programming Exascale systems
  - Dates: July 27-29, 2011
  - Location: USC/ISI
  - http://science.energy.gov/ascr/research/computer-science/programming-challenges-workshop/
- Architectures I, Exascale and Beyond Gaps in Research, Gaps in our Thinking
  - Discuss and explore the lessons learned from Exascale Projects and develop reverse timeline for accepting Exascale system in 2020.
  - Dates: August 2 – 3, 2011
  - Location: Stanford University
  - http://www.orau.gov/archI2011/default.htm
- Architectures II, Exascale and Beyond Configuring, Reasoning, Scaling
  - Investigate how we design computers so future Exascale computers enable DOE critical applications
  - Dates: August 8 – 11, 2011
  - Location: Sandia National Laboratories
  - http://www.orau.gov/archII2011/default.htm

U.S. DEPARTMENT OF ENERGY

Office of Science

Programming Challenges Workshop

Workshop Agenda

July 27-29, 2011

USC/ISI

PM: Sonia Sachs

U.S. DEPARTMENT OF ENERGY
Office of Science

Programming Challenges Workshop: Why?

- **Challenges for Programming Exascale systems:**
  - The switch from bulk-synchronous computing to asynchronous computing.
  - The growth in explicit on-chip parallelism their ability to express and manage up to a billion separate threads—a factor of 10,000 greater than on current platforms.
  - Expressing and managing hierarchical locality and data movement
  - Dealing with heterogeneity across the system
- **Understand and prioritize methods for programming Exascale systems in the areas of**
  - Compilers
  - Programming Models
  - Programming Languages and environments
  - Runtime systems

U.S. DEPARTMENT OF ENERGY
Office of Science

Programming Challenges Workshop: Presentation Highlights

- **Keshav Pingali (U Texas) -**
  - We have failed at auto-parallelization of apps
    - Much of auto-parallelization must be done at runtime
  - Success requires multiple classes (levels) of programmers.
  - Solution to auto-parallelization must not require compiler to raise abstraction level to uncover high level structure
- **Richard Lethin (Reservoir Lab) -**
  - Programmers should not deal with code complexity!
    - Exascale programming will be too complicated –VLIW lessons
    - Opaque to any semantic or dependence analysis needed for optimization
    - Will over-specify the program and bake it to one architecture, defeating portability
  - What we need:
    - High-level semantics rich expression languages.
    - Automated transformations to address Exascale hardware issues.
    - Automatic scheduling and data layout optimization

U.S. DEPARTMENT OF ENERGY
Office of Science

Programming Challenges Workshop: Presentation Highlights

- **Vivek Sarkar (Rice) -**
  - Classes of intermediate-level programming constructs for bridging the gap:
    - Asynchronous tasks and data transfers
    - Collective and point-to-point synchronizations and reductions
    - Locality control for task and data distribution
  - Compiler and runtime implementation for these constructs
- **Kathy Yelick (LBNL) -**
  - What should be virtualized? Processors? Memory locality? Portable mechanisms for locality optimization? Issues for dynamic threads with locality?
  - Which level is responsible for virtualizing?
  - Does the programming model expose hardware resources or hide it
  - Solve the problems that must be solved: locality, heterogeneity, vertical communication mgmt., fault resilience, dynamic resource mgmt.

U.S. DEPARTMENT OF ENERGY
Office of Science

Programming Challenges Workshop: Conclusions

- Multiple levels of program and programmers
  - High level, semantics rich expression language constructs
  - Intermediate levels: auto-tunable representations
  - Transformations of high level representations to intermediate
- Performance portability is a requirement
- Domain Specific Languages (DSL) constructs will hide code complexity
- New mechanisms (at all levels!) for managing fine-grain dynamic parallelism, locality, heterogeneity, irregular data and computation, asynchronous execution environments, energy and resilience.
  - MPI in the node will not cut it!
- Runtime support and control of
  - Auto-parallelism and interactive, adaptive, dynamic mechanisms.
  - Energy, resilience, and performance

U.S. DEPARTMENT OF ENERGY

Office of Science

Architectures I Workshop:

Exascale and Beyond

Gaps in Research, Gaps in our Thinking

Workshop Agenda

HPCWire Article

August 2 – 3, 2011
Stanford University
Palo Alto, CA
PM Lenore M. Mullin

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures I Workshop: Why?

- It has been 2 years since DOE and DARPA completed their Exascale reports and with the commencement of various projects to develop Extremescale designs we want to query vendors (IBM, HP, SGI), DOE labs (LBL, Sandia), universities (MIT) and other funded projects to determine
  - What hardware and software challenges they face.
  - What lessons were learned from their efforts.
  - What gaps exist in research and/or our thinking about Exascale to achieve a production machine by 2020.
- We need to develop a reverse time line that will indicate possible development paths, including a variety of technical challenges that lead towards a 2020 delivery date of an Exascale computer.

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures I Workshop: Presentation Highlights

- **Bill Dally (Stanford & NVIDIA) –**
  - Historic scaling has ended
  - MUST focus on the challenges of power, i.e. locality and overhead, and programmability
  - Our economy depends on these efforts
  - He emphasize that it is NOT about FLOPS, it is about data movement and designing *new algorithms, and implementations,* that will exploit more work per data movement.
  - Legacy codes will continue to run faster, their lack of locality will cause them to bottle neck on global bandwidth.
  - Recommends having a research vehicle (experimental system) to collaboratively design architectures, metrics, analysis, programming systems, and applications.
  - Exascale system in 2020 is essential to the DOE effort and our overall economy.
- **Shekhar Borkar (Intel) –**
  - Presented a realistic technology roadmap towards Exascale
  - Numerous gaps in research: impact on system reliability (noise and soft errors), power delivery and management, energy efficient signaling, heterogeneous, hierarchical, systems interconnect and their integration with memory sub-systems, new DRAM architecture and micro-architecture, new memory subsystems with DRAM, NAND/PCM, and DISK.

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures I Workshop: Presentation Highlights

- **Keren Bergman –**
  - Emphasized that photonics is an interconnect and is energy efficient; i.e. use a chip electrical router and switch.
  - Packaging must improve to provide a high bandwidth density off-chip interconnect switch fabric that is thermally robust.
  - To achieve scaled fabric of integrated CMOS-photonics we must develop this technology in small academic/industrial settings.
- **Dave Resnick –**
  - Use 3-d Memory components to get through the memory wall.
    - Projected memory is too small for Exascale.
  - Use TSV (Through Silicon Vias) to mount memory on top of logic layer. CMOS logic offers great opportunity for new memory functions: Gather/Scatter/Move, Coherency, Atomic operations, ALU functions

U.S. DEPARTMENT OF ENERGY
Office of Science

Exascale Reverse Timeline

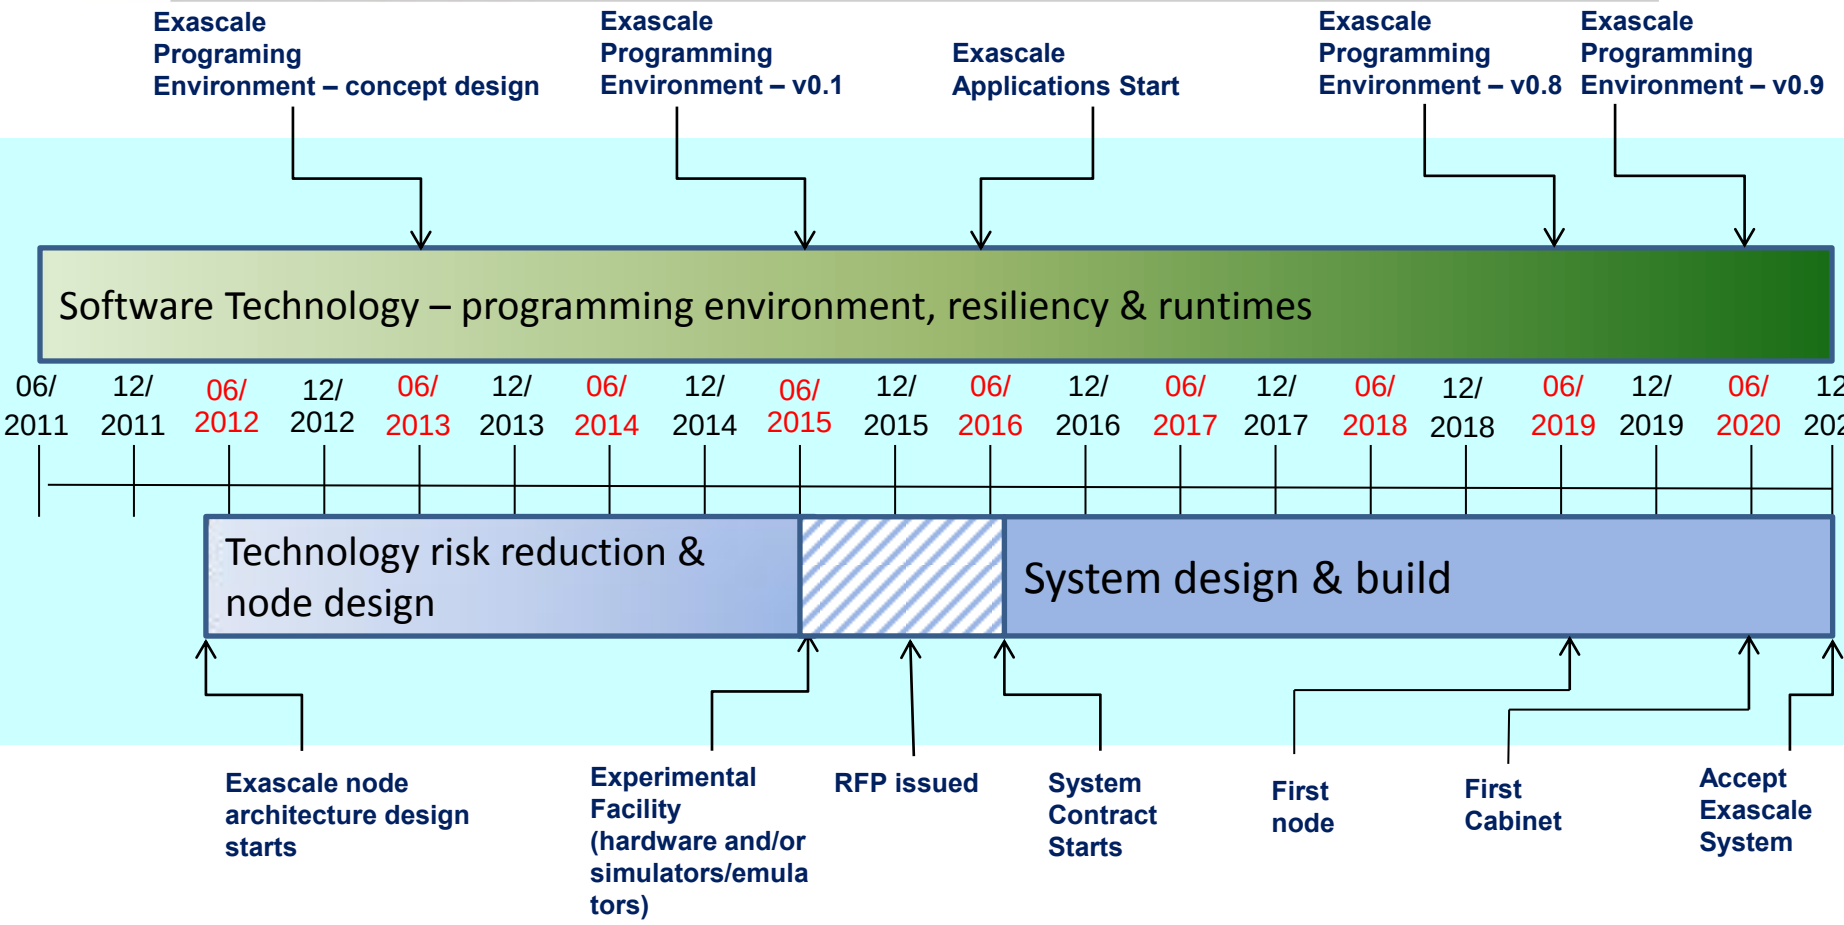

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures I Workshop: Conclusions

- **Energy efficiency is both a static and dynamic concern**
  - We want to optimize the cost of data movement
  - Gaps in:
    - Enabling photonics, demonstrating low-voltage electronics
    - New interconnect topologies and circuits
- **Programmability needs to be clarified and extended**
  - What and how should we express machine capabilities
  - Gaps in:
    - Observe and control performance attributes
    - Agile and hierarchical memories
    - Self aware OS
- **Resilience will ultimately limit the system utilization**
  - A system must be self-healing
  - Gaps in:
    - Adaptively isolate all malfunctioning system components and determine next steps given available resources.

U.S. DEPARTMENT OF ENERGY

Office of Science

Architectures II Workshop:

Exascale and Beyond

Configuring, Reasoning, Scaling

Workshop Agenda

August 8 – 11, 2011
Sandia National Laboratories
Albuquerque, NM
PM: Lenore M. Mullin

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures II Workshop: Why?

- The time has come to rethink how we design computers
- Determine if investments in abstractions will eventually lead to automated, energy efficient, easily scalable and portable scientific software.
- We must Identify ways
  - to abstract architectures, in general, to meet the ever changing complexities inherent in Exascale, and beyond.
  - Develop a process that will enable hardware and software architectures design space exploration, including simulation and emulation modeling
  - Build upon previous attempts of machine abstraction by learning from our successes and failures.

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures II Workshop: Presentation Highlights

- **Rich Lethin (Reservoir Lab) –**
  - Studies of heterogeneous reconfigurable computing platforms and lesson learned.
  - Using an abstract model, with simulation, is a valuable and valid approach to architectural research and design.
  - Core and network scheduling should be integrated.
- **Marc Snir (UIUC) –**
  - Current program stack preserves operation count and this CAN NOT predict neither time nor energy. Most energy consumption is due to communication.
  - Reduction of time and energy will be achieved (mostly/entirely) by new algorithms that control communications.
  - There needs to be a unified programming model that manages locality and a unified complexity theory for communication.

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures II Workshop: Presentation Highlights

- **Dan Campbell (GaTech) –**
  - Concepts of heterogeneity, metrics, and ability to adapt (morphing) inherent in DARPA PCA program are similar/identical to Exascale
  - Believes there should be both a high level compiler (close to the abstract machine) and low level compiler (close to particular machine specifics) and these needed to be integrated with performance metrics for adapting.
- **Arun Rodrigues (SNL) –**
  - Better simulation and modeling tools will ease our fears of programming and executing on Exascale.
  - Simulation should be a co-design tool.
    - We need an "ideal model" for modeling the model and it should include: universal machine abstraction that provides performance, energy, and cost functions.
    - We should be able to model multiple machines and pick the best for our desired requirements.

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures II Workshop: Conclusions

- Abstract machines are critical to the future
  - We need to anchor architecture R&D in sound co-design practices, based on a realistic ASCR application workload.
  - Modeling (of all flavors) needs to replace in practice naïve metrics that skew the design space and lead to sub-optimal solutions.
- The timeline for Exascale research is extremely short.
  - We should have started over a year ago but if we start a parallel effort now we may not face this situation in the future.
  - The complexity of Exascale systems and applications will require a concerted approach for optimality
    - from the HW to the runtime system to the programming models to the application, i.e. an abstract machine.
- Abstract machines need to be backed by detailed modeling, simulation, emulation, and verification on real platforms to reduce potential error in the final result.
- There was a mixing of ideas from scientists across national labs, funding agencies, international scientists, industry leaders, and academic institutions.

U.S. DEPARTMENT OF ENERGY
Office of Science

Outcome of All Workshops

- These workshops were an excellent start, and it is critically important that these discussions continue.
- Building a sense of community, and identifying a good core group of people that can further the R&D, and play a prominent role in ensuring dissemination and practical application of the results.
- All workshops emphasized the need for abstractions and the ability to assess performance and measure progress. Performance metrics must be available to enable comparisons of strategies and to provide validation of predictions.
- An experimental facility/center with professional development must exist for collaborative design, experimentation, fast prototyping of new concepts (from architecture, to languages, to compilers, to runtime systems) and validation of Exascale systems.
- **Each workshop will generate a report which will include the prioritized challenges and strategies for Exascale.**

Anticipated Future Programs

- Programming Models, Languages, Compilers, and Tools
  - Minimize exposure of system complexity
  - Extreme concurrency
  - Heterogeneous system
  - Minimize data movement
- X-Stack
  - Strong focus on runtimes for efficiency and resiliency
  - Self-aware OS
- Exascale Architectures
  - Abstract machine models for design space exploration, utilizing simulation
  - Driven by DOE selected applications
- Extreme Scale Solver Algorithms
  - Fine grain parallelism
  - Data movement & locality

U.S. DEPARTMENT OF ENERGY
Office of Science

Upcoming Workshops and Principal Investigator Meetings

- Joint DoD-NNSA-ASCR Workshop on Resilience, Oct 17-21, Catonsville, MD
- Exascale Research PI Meeting, Oct 11-13, Annapolis, MD
- Applied Math PI Meeting, Oct 17-19, Reston, VA
- Next Gen: Workshop on Scientific Collaborations for Extreme-Scale Science, Dec. 6-7, Gaithersburg, MD
- Electrical vs. Optical Interconnects, Early January, D.C. area
- Operating Systems for Exascale, study group, dates TBD.
- Exascale Research PI Meeting, April 17-19, 2012, TBD, CA

BACKUP

U.S. DEPARTMENT OF ENERGY
Office of Science

Programming Challenges Workshop: Participants

- **Workshop Committee**
  - Saman Amarasinghe (MIT)
  - Mary Hall (U. Utah)
  - Pat McCormick (LANL)
  - Richard Murphy (Sandia)
  - Keshav Pingali (U. Texas)
  - Dan Quinlan (LLNL)
  - Vivek Sarkar (Rice)
  - John Shalf(LBNL)
- **Advisory Committee:**
  - Bob Lucas (USC/ISI)
  - Kathy Yelick (LBNL/UCB)

- **Participants**

Almasi, George	Harrod, Bill	Macaluso, Antoinette	Sachs, Sonia R.
Amarasinghe, Saman	Henning, Paul	Mahlke, Scott	Sadayappan, P. (Saday)
Balaji, Pavan	Heroux, Mike	Mattson, Tim	Saraswat, Vijay
Barrett, Richard	Janssen, Curtis	McCormick, Pat	Sarkar, Vivek
Bernholdt, David	Johnson, Fred	Mellor-Crummey, John	Schuster, Vince
Blelloch, Guy	Kale, Sanjay	Misra, Jayadev	Shalf, John
Bronevetsky, Greg	Kandemir, Mahmut	Mukhopadhyay, Supratik	Shen, Xipeng
Chame, Jacqueline	Knobe, Kath	Mullin, Lenore	Sottile, Matt
Choi, Sung-Eun	Koniges, Alice	Murphy, Richard	Snir, Marc
De Supinski, Bronis	Krishnamoorthy, Sriram	Nowell, Lucy	Swaminarayan, Sriram
Diniz, Pedro	Krishnan, Manojkumar	Padua, David	Thakur, Rajeev
Garland, Michael	Kulkarni, Milind	Pakin, Scott	Vuduc, Richard
Gao, Professor	Lethin, Rich	Pingali, Keshav	Yan, Yonghong
Godinez, Larry	Lucas, Bob	Quinlan, Dan	Yelick, Kathy
Hall, Mary	Lusk, Rusty	Rajopadhye, Sanjay	Yi, Qing

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures I Workshop: Participants

Organizing Committee

- William Dally
- Peter Kogge
- Shekhar Borkar
- Richard Murphy
- Lenore Mullin
- William Harrod
- Sonia Sachs

- **Participants**

Anant Agarwal	MIT
Jim Ang	SNL
Keren Bergman	Columbia
Shekhar Borkar	Intel
Robert Colwell	DARPA
William Dally	Stanford
Mootaz Elnozahy	IBM
Guang Gao	Delaware
Al Geist	ORNL
John Gustafson	Intel
William Harrod	DOE/ASCR
Scott Hemmert	SNL
Ron Ho	Oracle Labs
Thuc Hoang	DOE/NNSA
Adolfy Hoisie	PNNL
Charlie Holland	DARPA
Fred Johnson	SAIC
Norm Jouppi	HP
Steve Keckler	Nvidia
Peter Kogge	Notre Dame
Bob Lucas	USC/ISI
David Mountain	NSA
Richard Murphy	SNL
Nowell, Lucy	ASCR
Sanjay Patel	UIUC
Wilf Pinfold	Intel
Irene Qualters	NSF
Dave Resnick	SNL
Sonia Sachs	DOE/ASCR
Steve Scott	Cray, Inc.
John Shalf	LBNL
Alan Snavely	UCSD
Thomas Sterling	Indiana

U.S. DEPARTMENT OF ENERGY
Office of Science

Architectures II Workshop: Participants

Organizing Committee:

Richard Lethin
Marc Snir
Thomas Sterling
Arun Rodrigues
John Shalf
William Harrod
Lenore Mullin
Sonia Sachs

Participants

Sarita Adve	UIUC
Stephen Booth	Edinburgh
Ron Brightwell	SNL
Dan Campbell	GTRI
Andrew Chien	Chicago
Srini Devadas	MIT
Pedro Diniz	USC/ISI
Sudip Dosanjh	SNL
Jean-Luc Gaudiot	UC, Irvine
William Harrod	ASCR DOE
Ron Ho	Oracle
Adolfy Hoisie	PNNL
Fred Johnson	SAIC
Dean Klein	Micron
Jeffrey Kuskin	DE Shaw
Christopher Lamb	Nvidia
Anita LaSalle	NSF
Richard Lethin	Reservoir
Andrew Lumsdaine	Indiana
Ron Minnich	SNL
Lenore Mullin	ASCR DOE
Richard Murphy	SNL
Keshav Pingali	Texas
Arun Rodrigues	SNL
Sonia Sachs	ASCR DOE
Ron Sass	UNC, Charlotte
Steve Scott	Cray
John Shalf	LBNL
Marc Snir	UIUC
Thomas Sterling	Indiana
Michael Wolfe	PGI