U.S. DEPARTMENT OF ENERGY

A View from ASCR

DOE Computer Graphics Forum
April 12, 2010
Park City, Utah

Dr. Lucy Nowell

Computer Scientist & Program Manager

Advanced Scientific Computing Research

U.S. DEPARTMENT OF ENERGY

- Office of the Secretary
  Dr. Steven Chu
  Daniel B. Poneman
  Deputy Secretary
  - Office of the Under Secretary
    Kristina Johnson
    - Asst. Secretary for Energy Efficiency & Renewable Energy
      Cathy Zoi
    - Asst. Secretary for Fossil Energy
      James Markowski
    - Asst. Secretary for Nuclear Energy
      Pete Miller
    - Asst. Secretary for Energy Delivery
      Pat Hoffman (A)
  - Office of the Under Secretary for Science
    Steven E. Koonin
    - Office of Science
      William Brinkman
      Deputy for Programs
      Pat Dehmer
      - Basic Energy Sciences
        Harriet Kung
      - High Energy Physics
        Dennis Kovar
      - Advanced Scientific Computing Research
        Michael Strayer
      - Nuclear Physics
        Tim Hallman
      - Biological & Environmental Research
        Anna Palmisano
      - Fusion Energy Sciences
        Ed Synakowski
      - Workforce Development for Teachers & Scientists
        Bill Valdez
  - Advanced Research Projects Agency – Energy
    Arun Majumdar

U.S. DEPARTMENT OF ENERGY

ASCR Organization

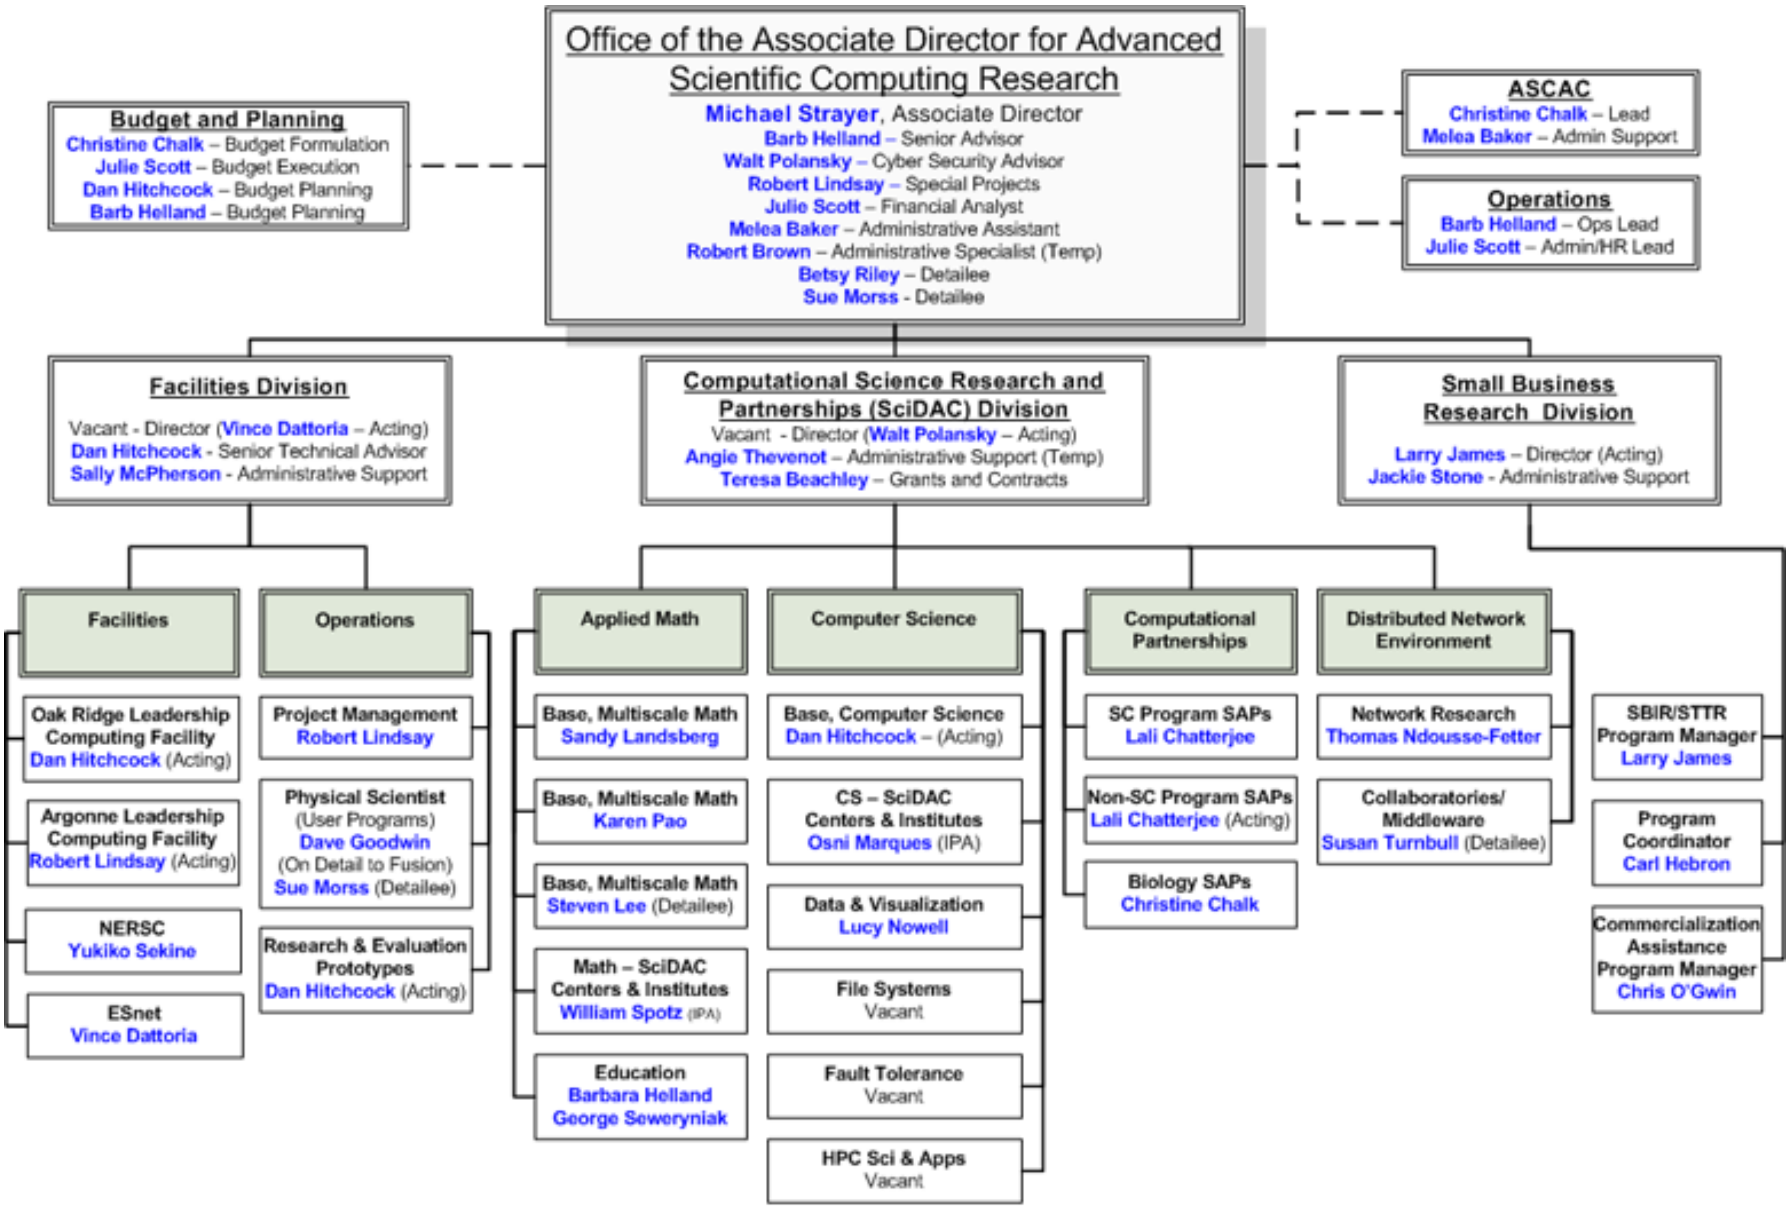

Office of Science (SC) FY 2011 Budget Request to Congress

U.S. DEPARTMENT OF ENERGY

(B/A in thousands)

	FY 2009		FY 2010	FY 2011		
	Current Base Approp.	Current Recovery Act	Current Approp.	Request to Congress	Request to Congress vs. FY 2010 Approp.	
Advanced Scientific Computing Research	358,772	161,795	394,000	426,000	+32,000	+8.1%
Basic Energy Sciences	1,535,765	555,406	1,636,500	1,835,000	+198,500	+12.1%
Biological & Environmental Research	585,176	165,653	604,182	626,900	+22,718	+3.8%
Fusion Energy Sciences	394,518	91,023	426,000	380,000	-46,000	-10.8%
High Energy Physics	775,868	232,390	810,483	829,000	+18,517	+2.3%
Nuclear Physics	500,307	154,800	535,000	562,000	+27,000	+5.0%
Workforce Development for Teachers & Scientists	13,583	12,500	20,678	35,600	+14,922	+72.2%
Science Laboratories Infrastructure	145,380	198,114	127,600	126,000	-1,600	-1.3%
Safeguards & Security	80,603	——	83,000	86,500	+3,500	+4.2%
Science Program Direction	186,695	5,600	189,377	214,437	+25,060	+13.2%
Small Business Innovation Research/Technology Transfer (SC)	104,905	18,719	——	——	——	——
Subtotal, Science	**4,681,572**	**1,596,000**	**4,826,820**	**5,121,437**	**+294,617**	**+6.1%**
Congressionally-directed projects	91,064	——	76,890	——	-76,890	-100.0%
Small Business Innovation Research/ Technology Transfer (DOE)	49,534	36,918	——	——	——	——
Use of prior year balances	-15,000	——	——	——	——	——
Total, Office of Science	**4,807,170**	**1,632,918**	**4,903,710**	**5,121,437**	**+217,727**	**+4.4%**

Office of Science FY 2011 Investment Highlights

The FY 2011 budget advances discovery science and invests in science for national needs in energy, climate, and the environment; national scientific user facilities; and education and workforce development.

Discovery science addressing national priorities

- **Energy Innovation Hub for Batteries and Energy Storage (+$34,020K, BES)**
- **Enhanced activities in climate science and modeling (Regional and Global Climate Modeling, +$6,495K; Earth System Modeling, +$9,015K; Atmospheric System Research, +$1,944K; ARM Climate Research Facility, +$3,961K; BER)**
- **Individual investigator, small group, and Energy Frontier Research Centers (EFRCs) in areas complementing the initial suite of 46 EFRCs awarded in FY 2009 (+$66,246K, BES)**
- **Leadership Computing Facilities operations and preparation for next generation of computer acquisitions for S&T modeling and simulation ($34,832K, ASCR)**
- **Multiscale modeling of combustion and advanced engine systems (+$20,000K, BES)**

Scientific user facilities—21st century tools of science, technology, and engineering

- **Facility construction is fully funded; projects are meeting baselines**
- **28 scientific user facilities will serve more than 26,000 users**
- **Several new projects and Major Items of Equipment are initiated (e.g., the Long Baseline Neutrino Experiment, +$12,000K, HEP)**

Education and workforce development

- **Expansions of the SC Graduate Fellowship Program (+$10,000K, 170 new awards, WDTS) and the SC Early Career Research Program (+$16,000K, 60 new awards, funded in all of the SC research programs)**

Leadership Computing Facilities

The Office of Science leads the World in supercomputing capabilities

U.S. DEPARTMENT OF ENERGY

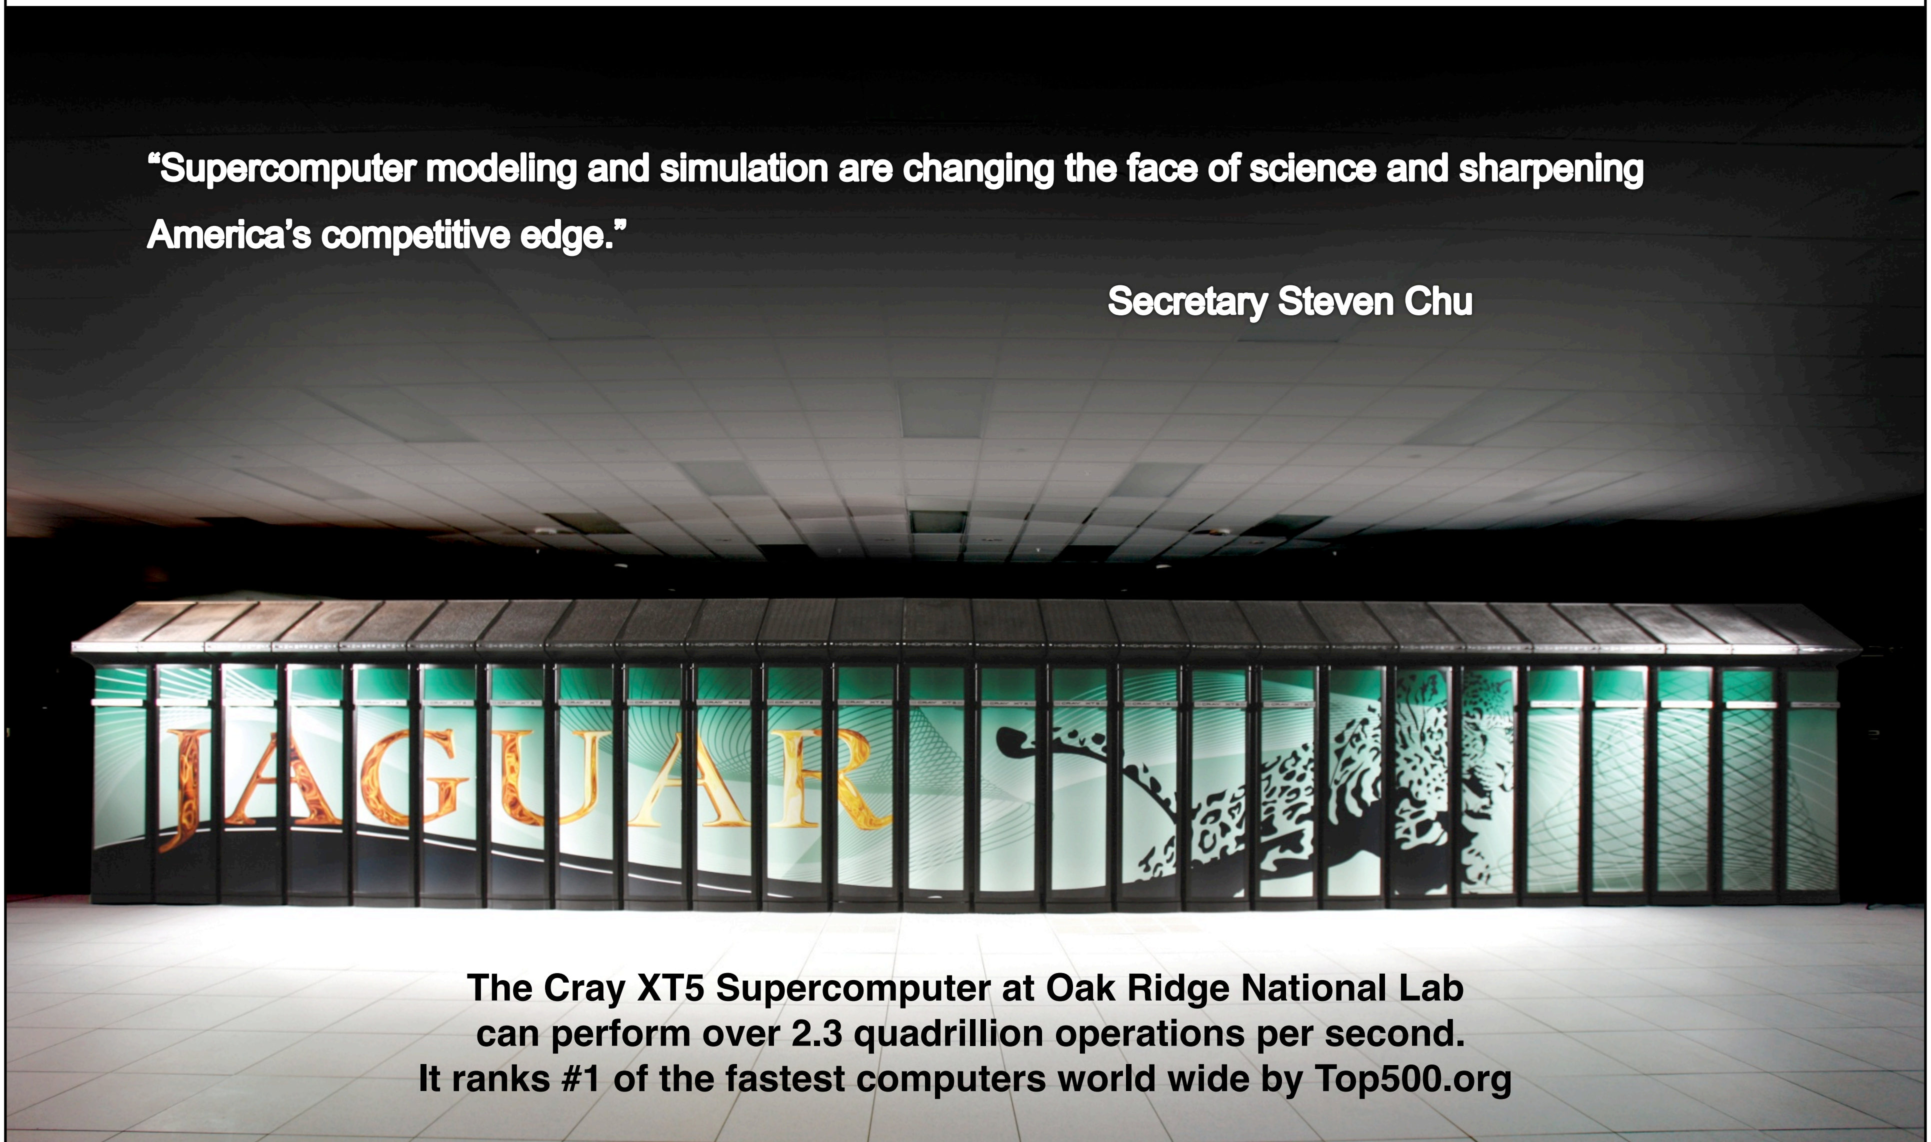

"Supercomputer modeling and simulation are changing the face of science and sharpening America's competitive edge."

Secretary Steven Chu

The Cray XT5 Supercomputer at Oak Ridge National Lab can perform over 2.3 quadrillion operations per second. It ranks #1 of the fastest computers world wide by Top500.org

Leadership Computing: Scientific Progress at the Petascale

U.S. DEPARTMENT OF ENERGY

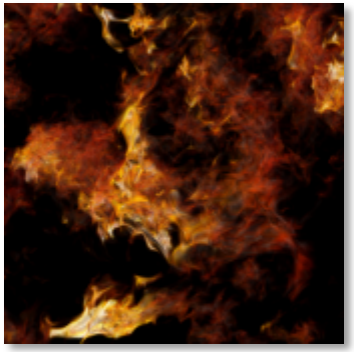

Turbulence
Understanding the statistical geometry of turbulent dispersion of pollutants in the environment.

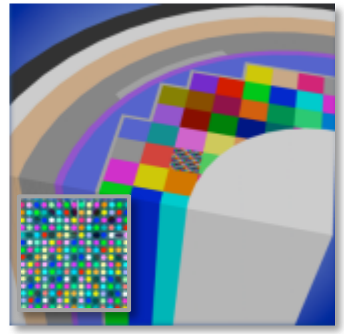

Nuclear Energy
High-fidelity predictive simulation tools for the design of next-generation nuclear reactors to safely increase operating margins.

Energy Storage
Understanding the storage and flow of energy in next-generation nanostructured carbon tube supercapacitors

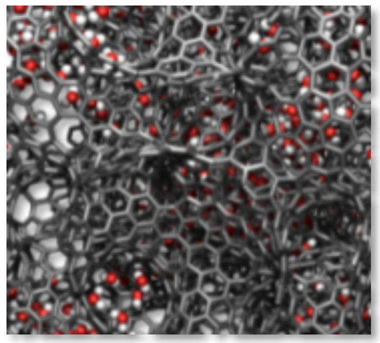

Fusion Energy
Substantial progress in the understanding of anomalous electron energy loss in the National Spherical Torus Experiment (NSTX).

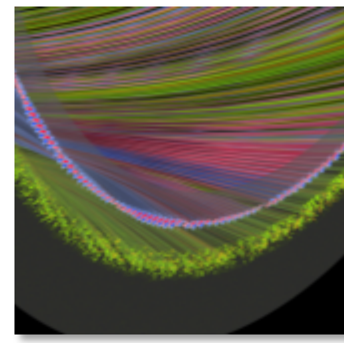

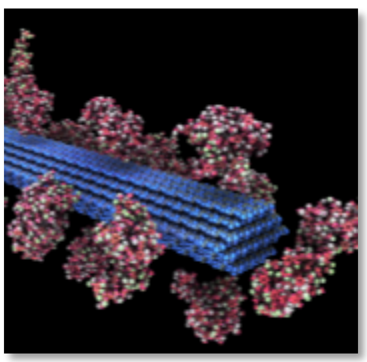

Biofuels
A comprehensive simulation model of lignocellulosic biomass to understand the bottleneck to sustainable and economical ethanol production.

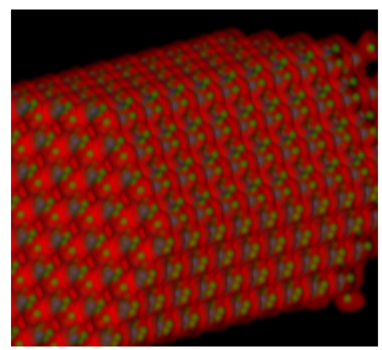

Nano Science
Understanding the atomic and electronic properties of nanostructures in next-generation photovoltaic solar cell materials.

U.S. DEPARTMENT OF ENERGY

Multi-scale Simulation of Internal Combustion Engines

A new initiative to develop the science base for computational design of advanced engines

Predictive simulation of combustion in an evolving fuel environment is essential for developing more efficient and cleaner engines.

The scientific community has provided a roadmap via:

- BES workshop: *Basic Research Needs for Clean and Efficient Combustion*, October 2006
- ASCR/BES workshop: *Discovery in Basic Energy Sciences: The Role of Computing at the Extreme Scale*, August 2009
- SC ongoing collaboration with EERE's Vehicle Technology Program

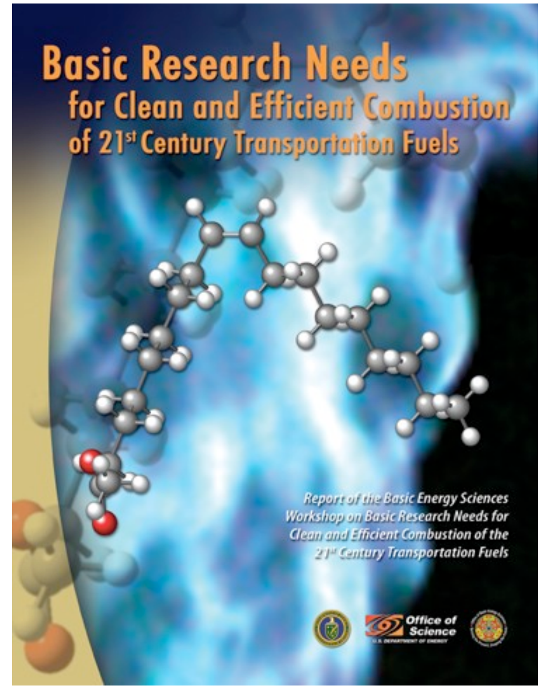

The new BES activity (+$20,000K) will provide:

- **Models that span vast scale ranges:** coupling of combustion chemistry with turbulent flow requiring simulation over 9 orders of magnitude in space and time.
- **Improved understanding of fundamental physical and chemical properties:** multi-phase fluid dynamics, thermodynamic properties, heat transfer, and chemical reactivity.
- **Engine simulation:** science-based predictive simulation and modeling design

U.S. DEPARTMENT OF ENERGY

Office of Science Early Career Research Program

Investment in FY 2011 will bring 62 new scientists into the program

$16 million will be available in FY 2011 to fund about 60 additional Early Career Research Program awards at universities and DOE national laboratories.

Purpose: To support individual research programs of outstanding scientists early in their careers and to stimulate research careers in the disciplines supported by the Office of Science

Eligibility: Within 10 years of receiving a Ph.D., either untenured academic assistant professors on the tenure track or full-time DOE national lab employees

Award Size:

- University grants $150,000 per year for 5 years to cover summer salary and expenses
- National lab awards $500,000 per year for five years to cover full salary and expenses

FY 2010 Results:

- 69 awards funded via the American Recovery and Reinvestment Act
- 1,750 proposals peer reviewed to select the awardees
- 47 university grants and 22 DOE national laboratory awards
- Awardees are from 44 separate institutions in 20 states

FY 2011 Application Process:

- Funding Opportunity Announcement issued in Spring 2010
- Awards made in the Second Quarter of 2011

http://www.science.doe.gov/SC-2/early_career.htm

U.S. DEPARTMENT OF ENERGY

DOE Office of Science Graduate Fellowships

The FY 2011 request doubles the number of graduate fellowships in basic science

$10 million will be available in FY 2011 to fund about 170 additional fellowships

Purpose: To educate and train a skilled scientific and technical workforce in order to stay at the forefront of science and innovation and to meet our energy and environmental challenges

Eligibility:

- Candidates must be U.S. citizens and a senior undergraduate or first or second year graduate student to apply
- Candidates must be pursuing advanced degrees in areas of physics, chemistry, mathematics, biology, computational sciences, areas of climate and environmental sciences important to the Office of Science and DOE mission

Award Size:

- The three-year fellowship award, totaling $50,500 annually, provides support towards tuition, a stipend for living expenses, and support for expenses such as travel to conferences and to DOE user facilities.

FY 2010 Results:

- 160 awards will be made this Spring with FY 2010 and American Recovery and Reinvestment Act funds

FY 2011 Application Process:

- Funding Opportunity Announcement issued in Fall 2010
- Awards made in March 2011

ASCR's Context Within DOE and SC

- **ASCR Contributes to 2 of DOE's Strategic Goals**
  - **Goal 3.1 – Scientific Breakthroughs: Achieve the major scientific discoveries that will drive U.S. competitiveness, inspire America, and revolutionize approaches to the Nation's energy, national security, and environmental quality challenges.**
    - **Advance the computational sciences and the leadership-class computational capabilities required for today's frontiers of scientific discovery.**
  - **Goal 3.2 – Foundations of Science: Deliver the scientific facilities, train the next generation of scientists and engineers, and provide the laboratory capabilities and infrastructure required for U.S. scientific primacy.**

U.S. DEPARTMENT OF ENERGY

ASCR Mission

The mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop, and deploy the computational and networking capabilities that enable researchers to analyze, model, simulate, and predict complex phenomena important to the Department of Energy. A particular challenge of this program is fulfilling the science potential of emerging multi-core computing systems and other novel "extreme-scale" computing architectures, which will require significant modifications to today's tools and techniques.

U.S. DEPARTMENT OF ENERGY

ASCR Computer Science Base Research

- **ASCR Base CS Program tries to address two fundamental questions:**
  - How can we make today's and tomorrow's leading edge computers tools for science?
  - How do we extract scientific information from petascale data from experiments and simulation?

- **There are several factors that provide important context for the ASCR Base CS program:**
  - SciDAC Centers and Institutes
  - Research and Evaluation Partnerships
  - ASCR Facilities

Scientific Discovery through Advanced Computing (SciDAC)

- Create comprehensive, scientific computing software infrastructure to enable scientific discovery in the physical, biological, and environmental sciences at the petascale

- Develop new generation of data management and knowledge discovery tools for large data sets (obtained from scientific users and simulations)

http://www.scidac.gov

Challenges for the Future
Path to Extreme Scale

- **Science at Extreme Scales**
- **Complexity**
- **Engineering Large Computer Systems**
- **Cybersecurity**

Evolution of HPC Resources

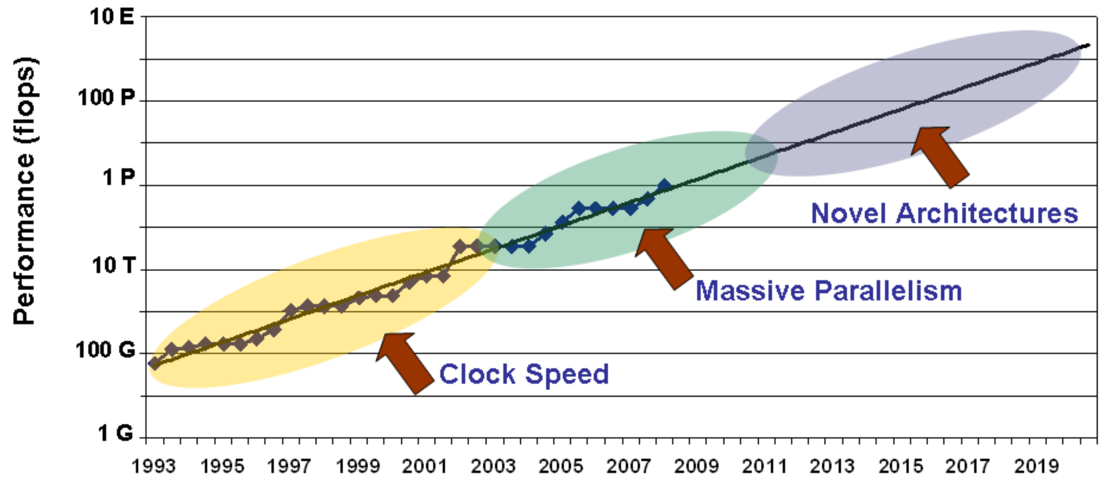

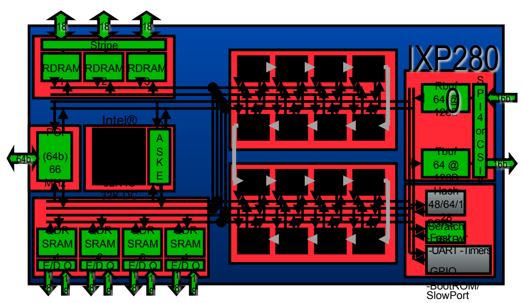

Intel Network Processor
1 GPP Core
16 ASPs (128 threads)

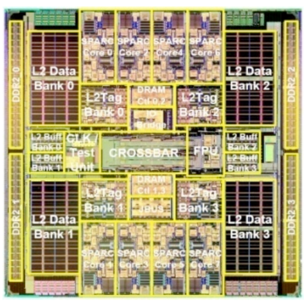

Sun Niagara
8 GPP cores (32 threads)

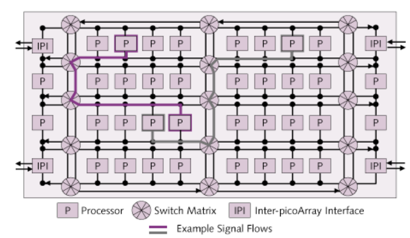

IBM Cell
1 GPP (2 threads)
8 ASPs

U.S. DEPARTMENT OF ENERGY

Challenges for the Future
Mountains of data

- **Storing:**
  - **Long term: where do we put 500TB?**
  - **Short term: scratch ~ 1TB, but need ~ 10TB!**
- **Moving:**
  - **Archive to scratch (~ 2 weeks to move 10TB)**
  - **HPC facility to local analysis cluster (longer)**
- **Processing:**
  - **Everything must be parallel, scalable.**
  - **IO speed, memory are the bottlenecks.**
- **Transforming Data into Insight**
  - **Physics are more complex**
  - **Wider range of scales, manual sifting is impossible.**
  - **Multi-scale analysis methods**
  - **Feature detection, growing, and tracking**

HPSS storage facility at NERSC

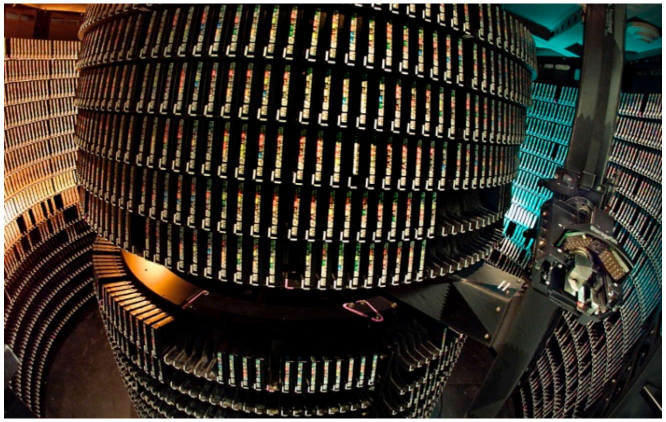

"Where is the wisdom that is lost in knowledge? Where is the knowledge we have lost in information?"

-T.S. Eliot

Exascale Initiative High Level Targets

- **The Exascale Initiative targets platform deliveries in 2018 and a robust Exascale simulation environment for the science exemplars by 2020**
- **Co-development of hardware, system software, programming model and applications require intermediate (100-200 PF/s) platforms in 2015**

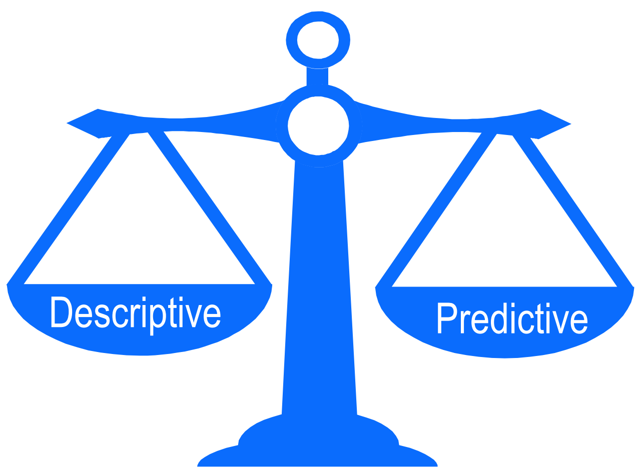

We are at the tipping point
for predictive capability

U.S. DEPARTMENT OF ENERGY

Identifying exascale applications and technology for DOE missions

- **Town Hall Meetings April-June 2007**
- **Scientific Grand Challenges Workshops November 2008 – October 2009**
  - **Climate Science (11/08),**
  - **High Energy Physics (12/08),**
  - **Nuclear Physics (1/09),**
  - **Fusion Energy (3/09),**
  - **Nuclear Energy (5/09),**
  - **Biology (8/09),**
  - **Material Science and Chemistry (8/09),**
  - **National Security (10/09)**
- **Cross-cutting workshops**
  - **Architecture and Technology (12/09)**
  - **Architecture, Applied Mathematics and Computer Science (2/10)**
- **Meetings with industry (8/09, 11/09)**

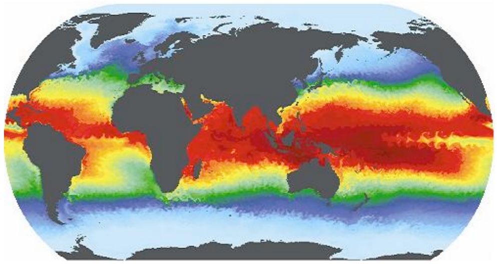

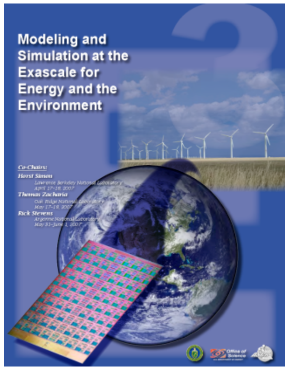

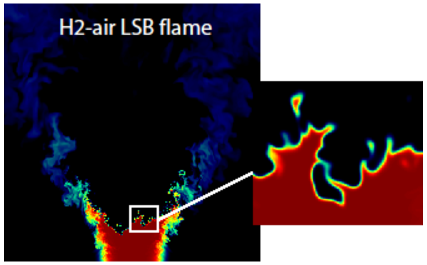

MISSION IMPERATIVES

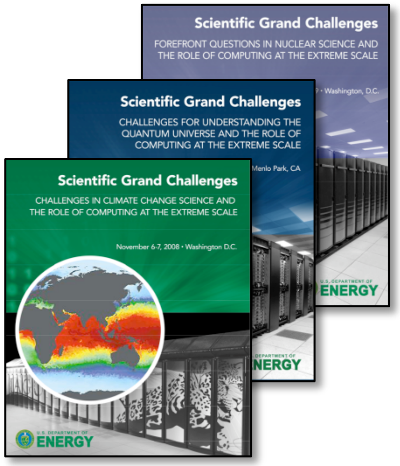

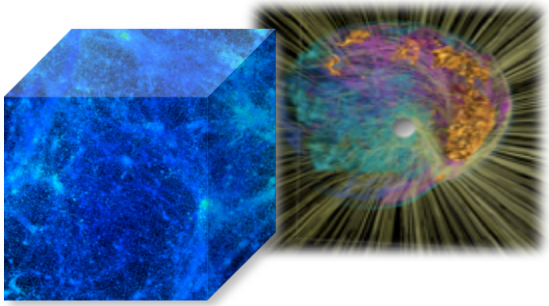

FUNDAMENTAL SCIENCE

DOE mission imperatives require simulation and analysis for policy and decision making

- ***Climate Change***: **Understanding and mitigating the effects of global warming**
  - Sea level rise
  - Severe weather
  - Regional climate change
  - Geologic carbon sequestration
- ***National Nuclear Security***: **Maintaining a safe, secure and reliable nuclear stockpile**
  - Stockpile certification
  - Predictive scientific challenges
  - Real-time evaluation of urban nuclear detonation
- ***Energy***: **Reducing U.S. reliance on foreign energy sources and reducing the carbon footprint of energy production**
  - Reducing time and cost of reactor design and deployment
  - Improving the efficiency of combustion energy sources

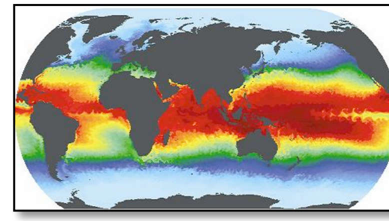

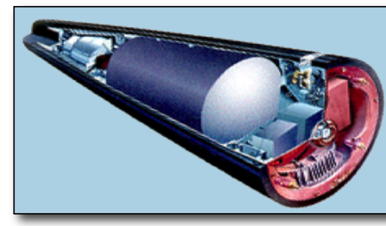

Accomplishing these missions requires exascale resources.

Climate change science is focused on providing effective tools for decision makers

Time scale:	Present	Decades	Centuries
Issue	Decisions		
Options for mitigation of greenhouse warming	**Promote needed technologies; economic incentives; costs. Conservation; patterns of Consumption. Social institutions.**	**Mitigation consequences for food production, water resources, etc. Mid-course correction for mitigation. Technological diffusion. Geo-engineering needed?**	**Optimum mitigation path to limit damage from global warming. Adaptation options.**
Coastal/river infrastructures and flooding hazards	**Insurance rates and policies**	**Safety provided over design lifetime.**	**Risk judgment**
Water resources	**Limit irrigation and lawn watering. Deploy conservation measures. Relate to political stability.**	**Deploy new dams and Reservoirs. National security-prevent regional conflict**	**Valuation of consequences.**
Agricultural production	**Impact of bio-fuel production on world food supplies.**	**Optimum strategies for land and new crop development. Increase in occurrence of droughts**	
Ecosystem management	**Impacts realized-habitat loss. Thresholds.**	**Selection of biological reserves.**	
Protection of Lives	**Sensitivities to severe Weather.**	**Threats of famine. Conflict avoidance. Abrupt change.**	

Identifying Outstanding Grand Challenges in Climate Change Research: Guiding DOE's Strategic Planning Sept, 2008

Power Consumption

- **Barriers**
  - **Power is leading design constraint for computing technology**
  - **Target ~20MW, estimated > 100MW required for Exascale systems (DARPA, DOE)**
  - **Efficiency is industry-wide problem (IT technology >2% of US energy consumption and growing)**
- **Technical Focus Areas**
  - **Energy efficient hardware building blocks (CPU, memory, interconnect)**
  - **Novel cooling and packaging**
  - **Si-Photonic Communication**
  - **Power Aware Runtime Software and Algorithms**
- **Technical Gap**
  - **Need 5X improvement in power efficiency over projections that include technological advancements**

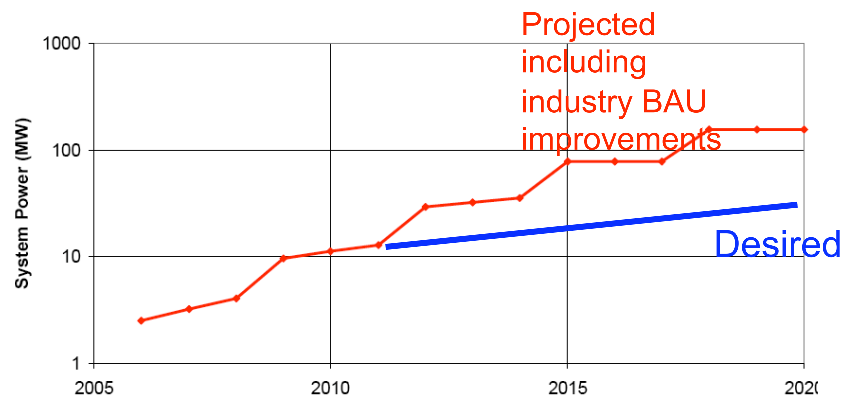

Possible Leadership class power requirements
From Peter Kogge (on behalf of Exascale Working Group), "Architectural *Challenges* at the Exascale Frontier", June 20, 2008

Projected Power Usage

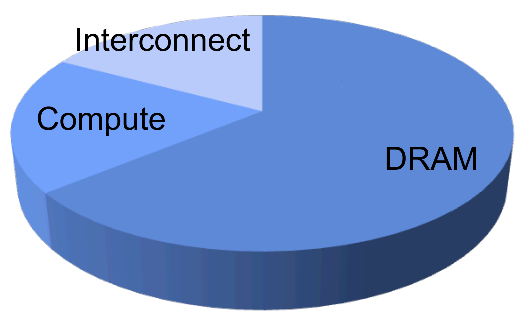

System memory dominates energy budget

Memory and Storage Bandwidth

- Barriers
  - Per-disk performance, failure rates, and energy efficiency no longer improving
  - Linear extrapolation of DRAM vs. Multi-core performance means the height of the memory wall is accelerating
  - Off-chip bandwidth, latency throttling delivered performance
- Technical Focus Areas
  - *Efficient Data Movement*
    - Photonic DRAM interfaces
    - Optical interconnects / routers
    - Communications optimal algorithms
  - *New Storage Approaches*
    - Non-volatile memory gap fillers
    - Advanced packaging (chip stacking)
    - Storage efficient programming models (Global Address Space)
- Technical Gap
  - Need 5X improvement in memory access speeds to keep current balance with computation.

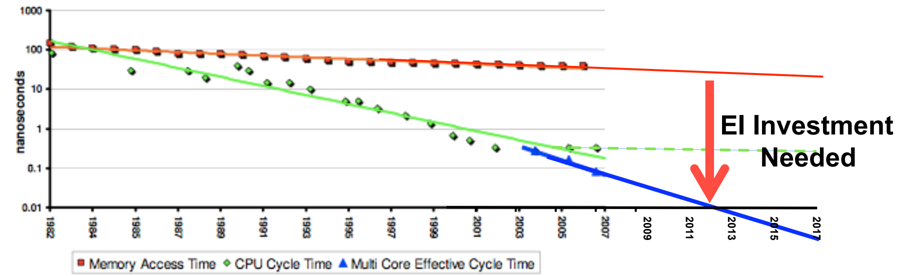

Figure 6.12: CPU and memory cycle time trends.

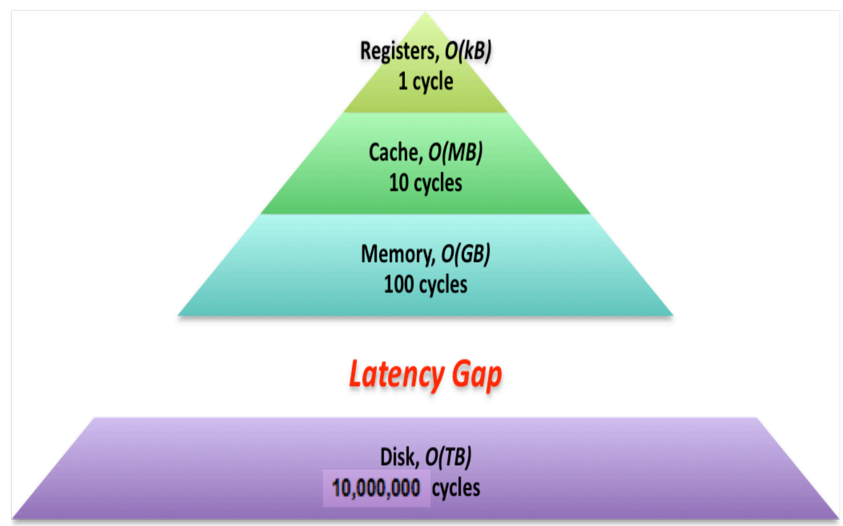

Reliability and Resilience

- **Barriers**
  - Number of system components increasing faster than component reliability
  - Mean time between failures of minutes or seconds for exascale
  - Silent error rates increasing
  - No job progress due to fault recovery if we use existing checkpoint/restart
- **Technical Focus Areas**
  - Improved hardware and software reliability
    - Better RAS collection and analysis (root cause)
    - Greater integration
  - Fault resilient algorithms and applications
  - Local recovery and migration
- **Technical Gap**
  - Need 1000X improvement in MTTI so that applications can run for many hours. Goal is 10X improvement in hardware reliability. Local recovery may and migration may yield another 10X. However, for exascale, applications will need to be fault resilient.

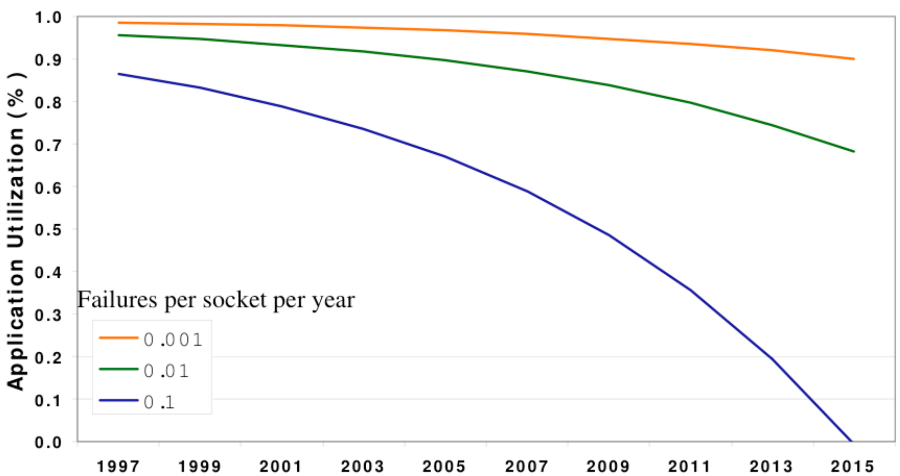

Effective application utilization (including checkpoint overhead) at 3 rates of hardware failure

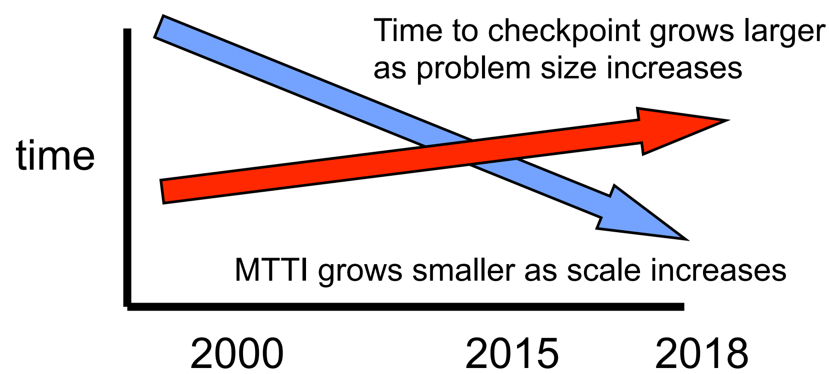

By exascale checkpoint/restart no longer viable

System Software Scalability

- **Barriers**
  - Fundamental assumptions of system software architecture did not anticipate exponential growth in parallelism
  - Requirements for resilience at scale
  - IO wall reducing effectiveness of simulation environment
- **Technical Focus Areas**
  - System Hardware Manageability
  - System Software Scalability
  - Applications Scalability
  - Supporting investments in infrastructure to support systems
  - Initial deliveries to validate software and operations path
- **Technical Gap**
  - 1000x improvement in system software scaling
  - 100x improvement in system software reliability
  - Need application hooks into RAS system

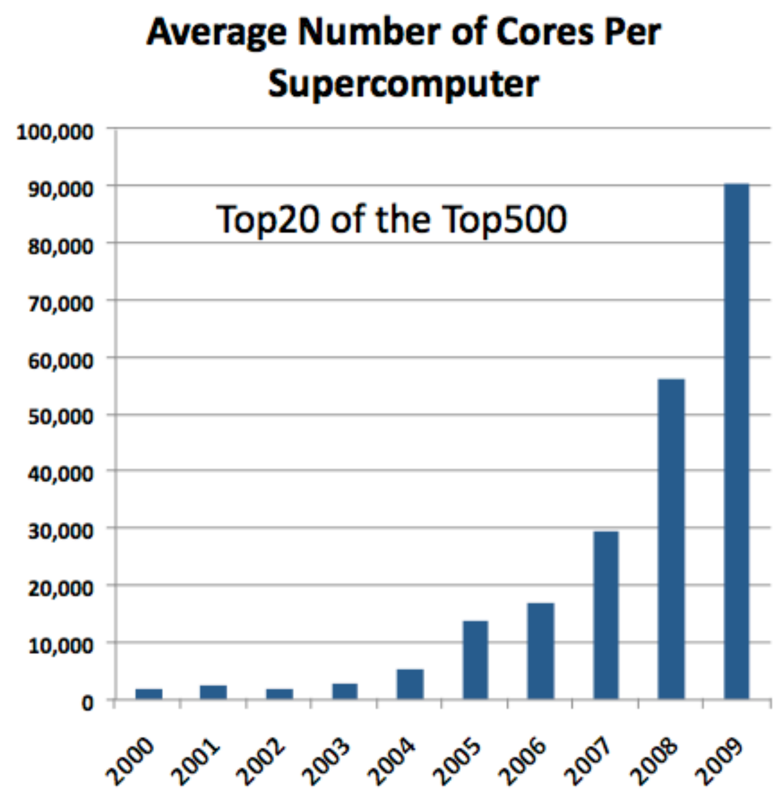

Programming Models and Environments

- **Barriers: Delivering a complex large-scale scientific instrument that is productive and fast.**
  - O(1B) way parallelism in Exascale system
    - Massive lightweight cores for low power
    - Some "full-feature" cores lead to heterogeneity
  - O(1K) way parallelism in a processor
    - Data and independent thread parallelism
  - Effective management of locality
    - Software-managed memory (local store)
    - Effective abstractions for explicitly managed memory hierarchies
    - Communication avoiding algorithms
    - Communication optimized for architecture
  - Complexity of scientific applications
  - Programming for resilience
  - Science goals require complex codes

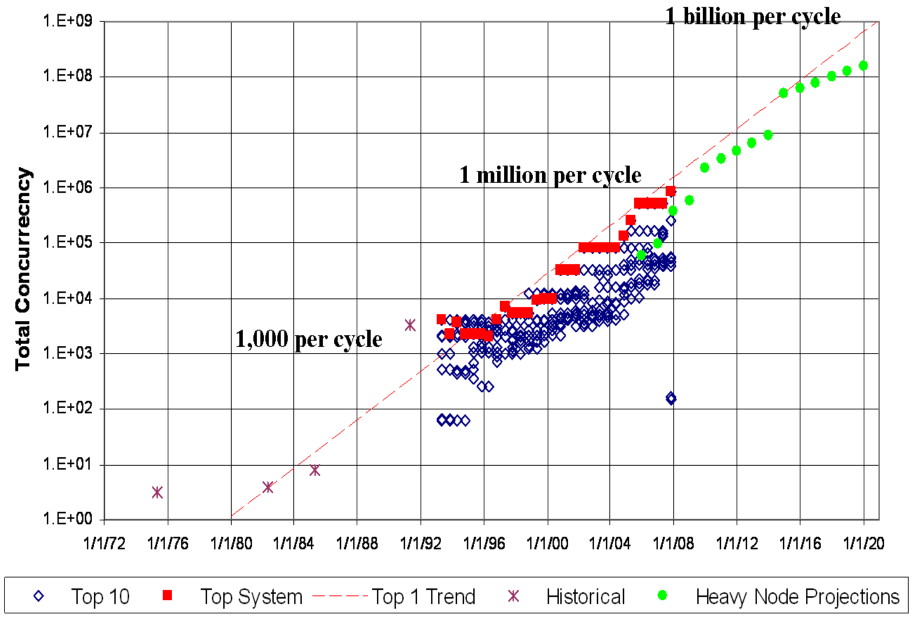

How much parallelism must be handled by the program?
From Peter Kogge (on behalf of Exascale Working Group), "Architectural *Challenges* at the Exascale Frontier", June 20, 2008

- **Technology Investments**
  - Evolutionary: extend existing models used in science for scalability and to hide system complexity, e.g., heterogeneity and failures
  - Moderate: leverage emerging models in scientific computing
  - Revolutionary: develop a new paradigm for high usability at extreme scales
- **Technical Gap:** Productivity, Performance and Correctness for 1000x more parallelism while increasing programming productivity of scientists by 10x

Potential System Architectures
In case you were sleeping

Systems	2009	2015 +1/-0	2018 +1/-0
System peak	**2 Peta**	**100-300 Peta**	**1 Exa**
Power	**6 MW**	**~15 MW**	**~20 MW**
System memory	0.3 PB	5 PB	64 PB (+)
Node performance	125 GF	0.5 TF or 7 TF	1-2 or 10TF
Node memory BW	25 GB/s	1-2TB/s	2-4TB/s
Node concurrency	12	O(100)	O(1k) or 10k
Total Node Interconnect BW	3.5 GB/s	100-200 GB/s 10:1 vs memory bandwidth 2:1 alternative	200-400GB/s (1:4 or 1:8 from memory BW)
System size (nodes)	18,700	50,000 or 500,000	O(100,000) or O(1M)
Total concurrency	225,000	O(100,000,000) *O(10)-O(50) to hide latency	O(billion) * O(10) to O(100) for latency hiding
Storage	15 PB	150 PB	500-1000 PB (>10x system memory is min)
IO	0.2 TB	10 TB/s	60 TB/s (how long to drain the machine)
MTTI	days	O(1day)	O(1 day)

U.S. DEPARTMENT OF ENERGY

On-Chip Architecture: *different approaches to on-chip clustering*

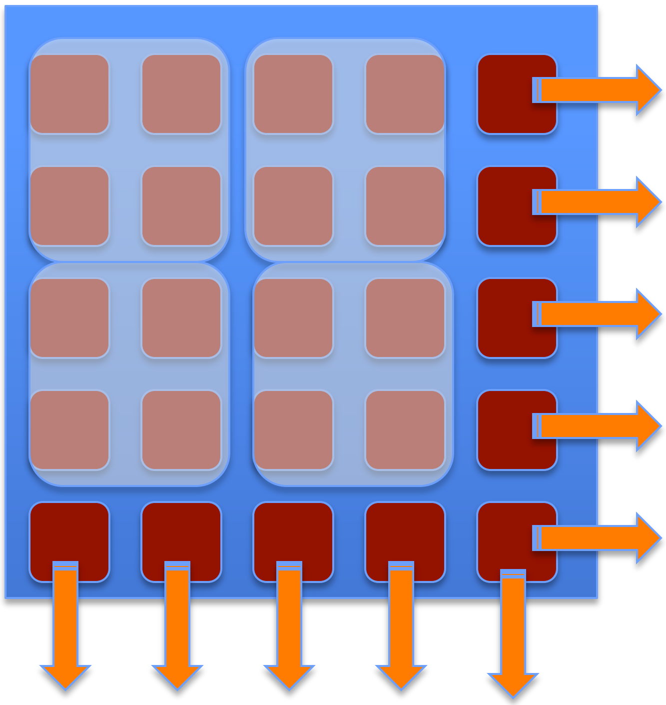

- **Cost of moving long-distances on chip motivates clustering on-chip**
  - 1mm costs ~6pj (today & 2018)
  - 20mm costs ~120 pj (today & 2018)
  - FLOP costs ~100pj today
  - FLOP costs ~25pj in 2018
- **Different Architectural Directions**
  - GPU: WARPs of hardware threads clustered around shared register file
  - CMP: limited area cache-coherence
  - CMT: hardware multithreading clusters

U.S. DEPARTMENT OF ENERGY

More Bad News

- I/O to disk will be relatively slower than it is today;
- Part of the file system may be on the node;
- There will be silent errors;
- Weak scaling approaches (constant memory/flop) probably will not work.

U.S. DEPARTMENT OF ENERGY

Approaches

- Locality, Locality, Locality!
- Billion Way Concurrency;
- Uncertainty Quantification including hardware variability;
- Flops free, but data movement expensive so:
  - Remap multiphysics to put as much work per location on same die;
  - Include embedded UQ to increase concurrency;
  - Include data analysis if you can for more concurrency
  - Trigger output to move only important data off machine;
  - Reformulate to trade flops for memory use.

New Science for a Secure and Sustainable Energy Future

- Significant discoveries will come at the intersection of control science with advanced materials and chemical phenomena, and there is a clear first-mover advantage to those who focus their research efforts there
- It will take “dream teams” of highly educated talent, equipped with forefront tools, and focused on the most pressing challenges to increase the rate of discovery