U.S. Department of Energy

Office of Science

DOE Extreme Scale Computing Initiative

Michael Strayer

Associate Director, Advanced Scientific Computing Research

Where we are today

ASCR's Vision

- **Deliver Petascale Computing for Science Applications**
  - Continue to make the Leadership Computing Facilities available to the very best science through Innovative and Novel Computational Impact on Theory and Experiment (INCITE).
  - Continue to work with Pioneer Applications to deliver scientific results from day one.
- **Keep DOE Computational Science at the Forefront**
  - Continue to nurture applications critical to DOE missions through Scientific Discovery through Advanced Computing (SciDAC).
  - Provide direct support for "bleeding-edge" research groups willing to take on the risk of working with emerging languages and operating systems.
  - Foster innovative research at the ever blurring boundary between Applied Mathematics and Computer Science.
- **The Promise of Extreme Scale**
  - Work with key science applications to identify opportunities for new research areas only possible through extreme scale computing.
  - Support innovative research on advanced architectures and algorithms that accelerates the development of hardware and software that is well suited to extreme scale computational science.

ASCR
Facilities Strategy

- Providing the Tools – High-End Computing
  - High-Performance Production Computing - National Energy Research Scientific Computing Center **(NERSC)** at Lawrence Berkeley National Laboratory
    - Delivers high-end capacity computing to entire DOE SC research community
  - Leadership-Class Computing – **Leadership Computing Centers at Argonne National Laboratory and Oak Ridge National Laboratory**
    - Delivers highest computational capability to national and international researchers through peer-reviewed **Innovative and Novel Computational Impact on Theory and Computation (INCITE)** program (80% of resources)
- Investing in the Future - Research and Evaluation Prototypes
- Linking it all together – Energy Sciences Network (ESnet)

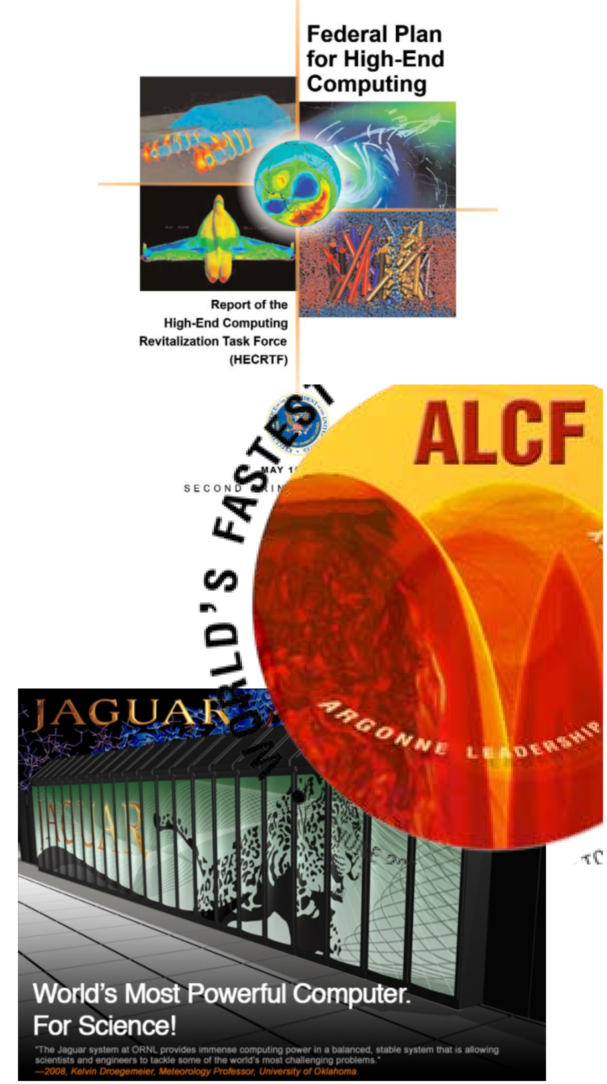

National Energy Research Scientific Computing Center (NERSC)

- Located at Lawrence Berkeley National Lab
  - Cray XT-4 Franklin: 102 Tflop/s, 9,660 nodes, 19,320 cores
  - IBM Power 5 Bassi: 6.7 Tflop/s, 888 cores
  - Linux Opteron Cluster Jacquard: 3.1 Tflop/s, 712 cores
- Franklin quad-core currently being upgraded 350 Tflop/s, 38,640 cores
- NERSC-6 Project
  - RFP issued in September 2008
  - Proposals are being reviewed

Franklin

Bassi and Jacquard

Argonne’s IBM Blue Gene/P – 556 TFs

U.S. Department of Energy

Office of Science

Oak Ridge's Cray XT5
Breaks the Petaflop Barrier

Jaguar	Total	XT5	XT4
Peak Performance	1,645	1,382	263
AMD Opteron Cores	**181,504**	**150,176**	**31,328**
System Memory (TB)	**362**	**300**	**62**
Disk Bandwidth (GB/s)	**284**	**240**	**44**
Disk Space (TB)	**10,750**	**10,000**	**750**
Interconnect Bandwidth (TB/s)	**532**	**374**	**157**

U.S. Department of Energy

Office of Science

ESnet
40 Gbps Core

- Leader in Networks for Science
  - OSCARS
  - PerfSONAR
  - DanteInternet2CanarieESnet

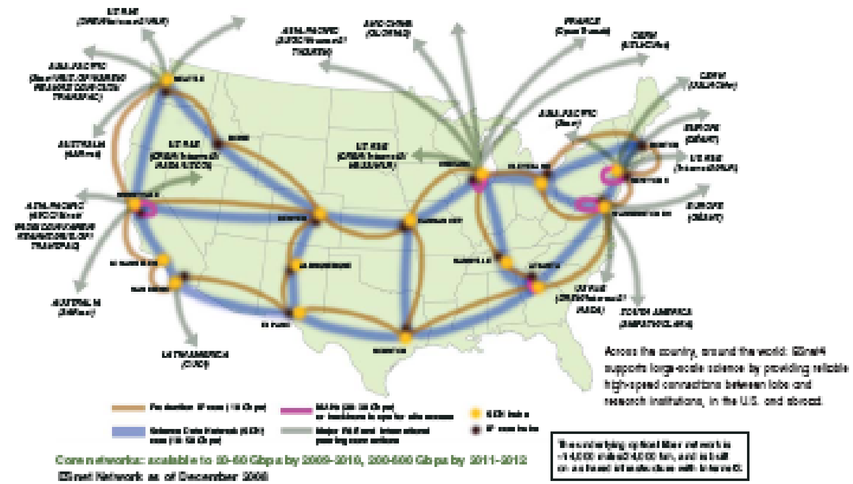

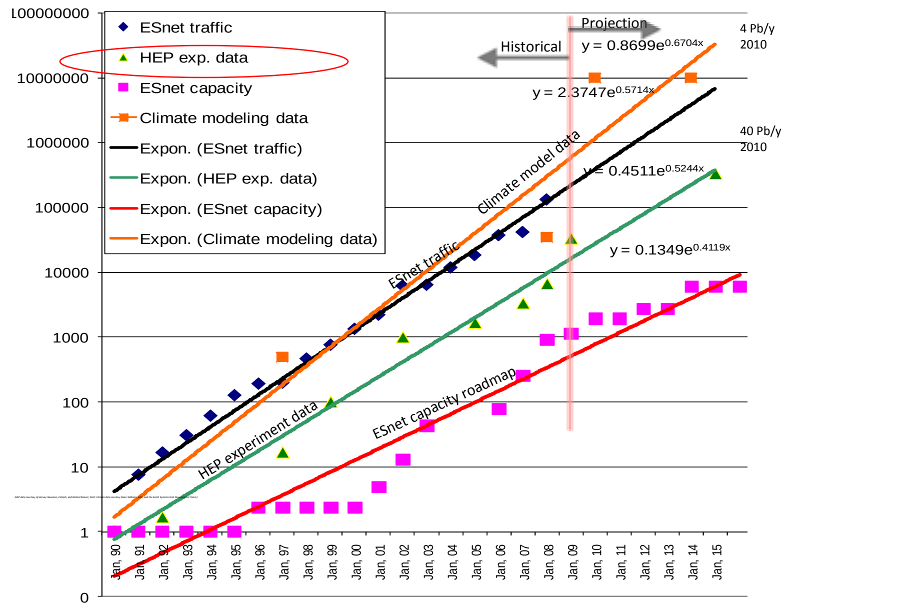

ASCR Research Strategy

- **Provide knowledge and foundational tools:**
  - **Applied Mathematics:** Develop algorithms for solving complex, mission-relevant science problems.
  - **Computer Science:** Facilitate the use of emerging Leadership-scale computing resources.
  - **Integrated Network Environment:** Develop tools to extract meaningful information from peta-byte data sets; Enable geographically distributed research teams to collaborate - share data, assess results, plan and conduct experiments.
- **Break new ground in science:**
  - **SciDAC:** Deliver computational tools and techniques to advance DOE-science through modeling and simulation
  - **Multiscale mathematics:** Discover methods and algorithms to fully-describe understanding of nature over vast scales of time and space.

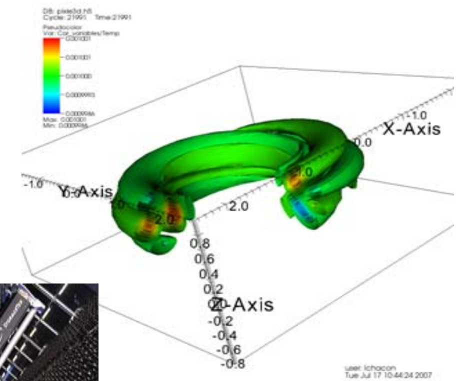

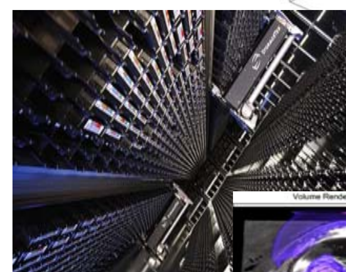

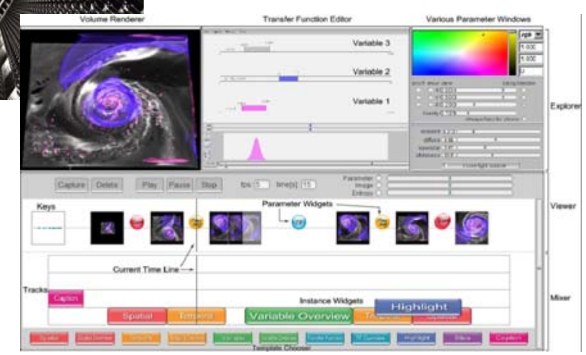

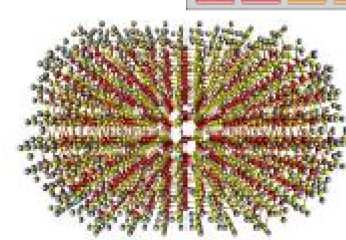

U.S. Department of Energy

Office of Science

Delivering the Science

Highlighting Scientific Discovery and the Role of High End Computing

Top 10 Computational Science Accomplishments

(Blue=SciDAC; Black=INCITE)

Rank	Title
1	**Modeling the Molecular Basis of Parkinson's Disease (Tsigelny)**
2	**Discovery of the Standing Accretion Shock Instability and Pulsar Birth Mechanism in a Core-Collapse Supernova Evolution and Explosion (Blondin)**
3	**Prediction and Design of Macromolecular Structures and Functions (Baker)**
4	**Understanding How Lifted Flame Stabilized in a Hot Coflow (Yoo)**
5	**New Insights from LCF-enable advanced kinetic simulations of global turbulence in fusion systems (Tang)**
6	**High Transition Temperature Superconductivity: A High-Temperature Superonductive State and a Pairing Mechanism in 2-D Hubbard Model (Scalapino)**
7	**PETsc: Providing the Solvers for DOE High-Performance Simulations (Smith)**
8	**Via Lactea II, A Billion Particle Simulation of the Dark Matter Halo of the Milky Way (Madau)**
9	**Probing the properties of water through advanced computing (Galli)**
10	**First Provably Scalable Maxwell Solver Enables Scalable Electromagnetic Simulations (Kovel)**

U.S. Department of Energy
Office of Science

High-Transition Temperature Superconductivity

D. Scalapino (UCSB)
T. Maier, P. Kent, and T. Schulthess (ORNL)
M. Jarrell and A. Macridin (University of Cincinnati)
D. Poilblanc (Laboratoire de Physique Thˇorique, CNRS and Universitˇ de Toulouse)

Accomplishment

- Given new numerical techniques and leadership-class resources, 2D Hubbard Model solved computationally.
- Proved model does in fact describe high-temperature superconductivity.

  ⇨ *Settled a debate that raged for two decades.*

Broader Implications

- Provides a deeper understanding of high-T superconductivity.
- Key step toward the identification of high-T superconducting materials, *the Holy Grail of superconductivity research.*

Support: BES Base Program, SciDAC

Maier, T.A., A. Macridin, M. Jarrell, and D. J. Scalapino, ÑSystematic analysis of a spin-susceptibility representation of the pairing interaction in the two-dimensional Hubbard modelÓ *Phys. Rev. B.* **76**, 144516 (2007).
Maier, T.A., M. Jarrell, and D. J. Scalapino, ÑSpin susceptibility representation of the pairing interaction for the two-dimensional Hubbard modelÓ *Phys. Rev. B.* **75**, 134519 (2007).

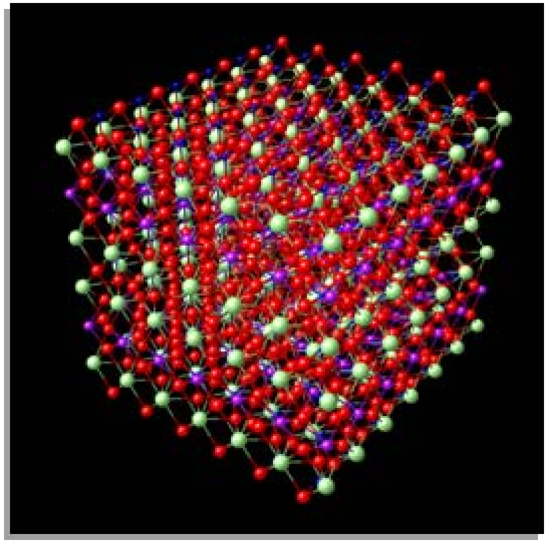

Model of a $YBa_2Cu_3O_7$ high-temperature superconductor crystal

Anthony Mezzacappa (ORNL), PACS Report, SciDAC 2008

Modeling the Molecular Basis of Parkinson's Disease

U.S. Department of Energy
Office of Science

⇨ Igor Tsigelny (SDSC)
⇨ Eliezer Masliah (UCSD)
⇨ Stanley Opella (UCSD)

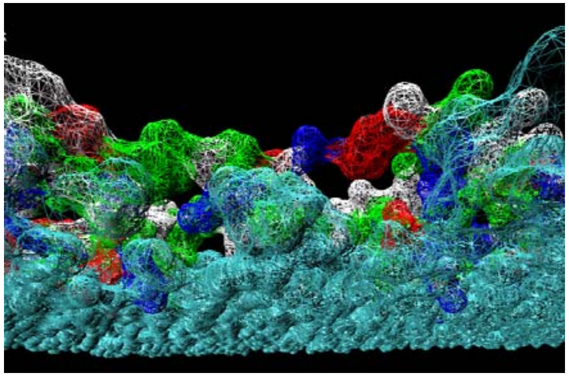

Accomplishment

• Elucidated the molecular mechanism of the progression of Parkinson's disease.
• Insights will help focus the search for treatment.

Broader Implications

• Provided a test bed for identifying possible therapeutic interventions through computational modeling.
• Overall approach has broad applicability to other diseases.

Support: INCITE

I.F. Tsigelny et al. 2007. Dynamics of α-synuclein aggregation and inhibition of pore-like oligomer development by β-synuclein, *FEBS J.*, **274**, 1862-1877.

Anthony Mezzacappa (ORNL), PACS Report, SciDAC 2008

Via Lactea II: A Billion Particle Simulation of the Dark Matter Halo of the Milky Way Galaxy

- ⇨ Piero Madau (UCSC)
- ⇨ Juerg Diemand (UCSC)
- ⇨ Michael Kuhlen (IAS)
- ⇨ Marcel Zemp (UCSC)

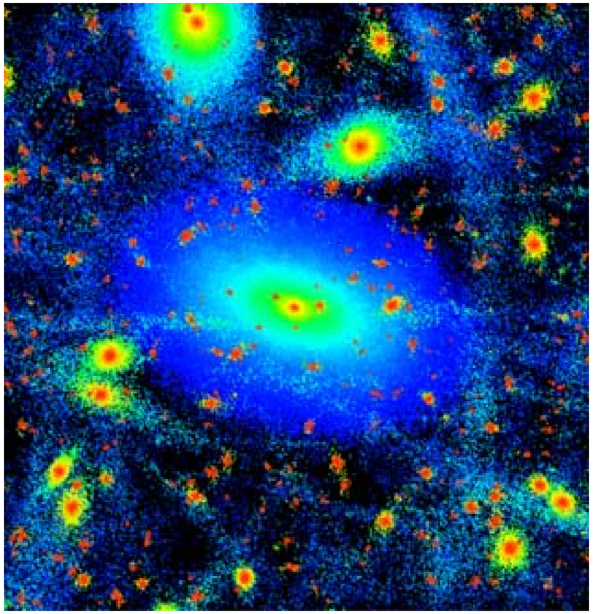

Accomplishment

Largest simulations ever performed of the formation of the dark matter halo of the Milky Way galaxy.

⇨ *Resolved and predicted structure at small scales.*

Broader Implications

Structures predicted for small scales have observable gamma-ray signatures for certain classes of dark matter candidates.

✶ *Simulations may play an important role in identifying the dark matter in the Universe.*

Support: INCITE

Diemand, Kuhlen, Madau, Zemp, Moore, Potter, and Stadel 2008, Clumps and Streams in the Local Dark Matter Distribution, *Nature*, in press.

Anthony Mezzacappa (ORNL), PACS Report, SciDAC 2008

Where are we going?

Road to Extreme Scale

Computing is changing more rapidly than ever before, and scientists have the unprecedented opportunity to change computing directions

Traditional Sources of Performance Improvement are Flat-Lining (2004)

- New Constraints
  - 15 years of *exponential* clock rate growth has ended
- Moore's Law reinterpreted:
  - How do we use all of those transistors to keep performance increasing at historical rates?
  - Industry Response: #cores per chip doubles every 18 months *instead* of clock frequency!

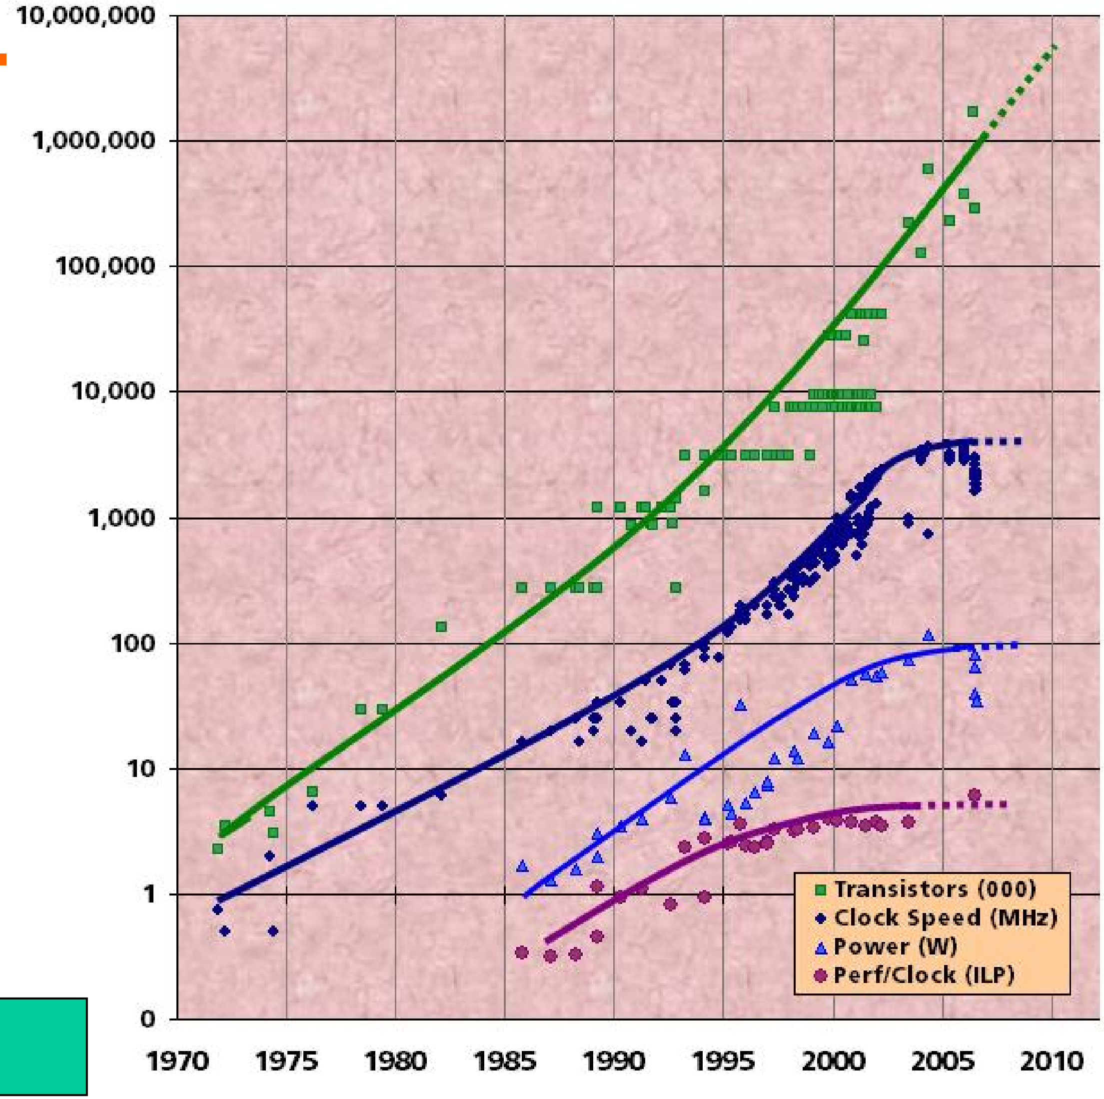

Figure courtesy of Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith

An Era of Challenge

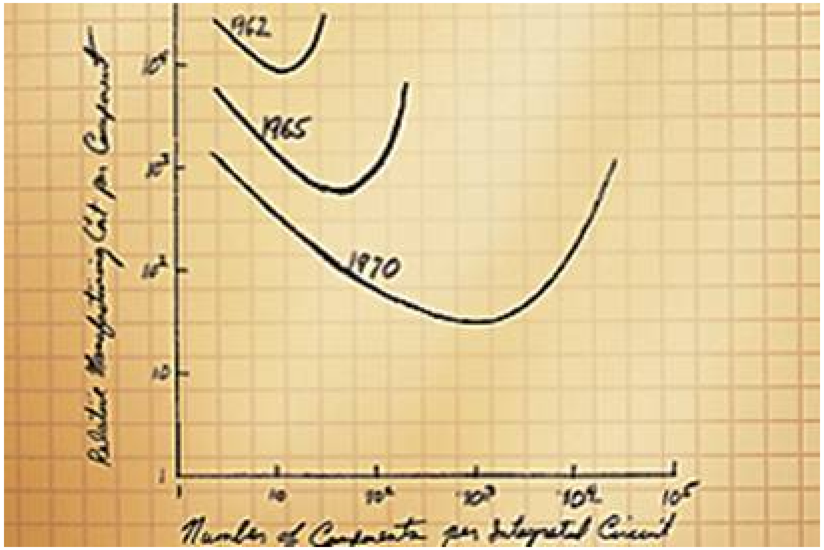

Moore's original graph predicting Moore's Law in 1965. Chip capacity will double every 2 yrs.

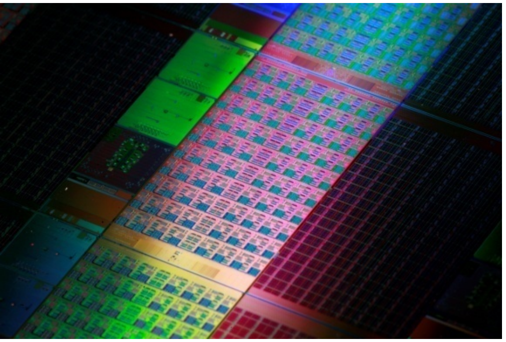

Intel Teraflops Research Chip

IBM Stacked Chip

- **Unpredictable evolution of hardware**
- **Multilevel and heterogeneous parallelism; memory hierarchies**
- **Programming models must work at scale (numbers of core, lines of code, numbers of components)**
- **Managing data, simulation, experimental and observed**
- **Communications: synchronous → asynchronous**
- **Reliability**

It's not just extreme scale, it's also extreme complexity

Multicore comes in a wide variety

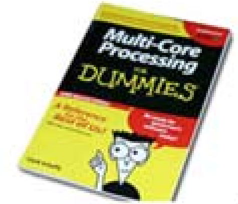

- Multiple parallel general-purpose processors (GPPs)
- Multiple application-specific processors (ASPs)

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

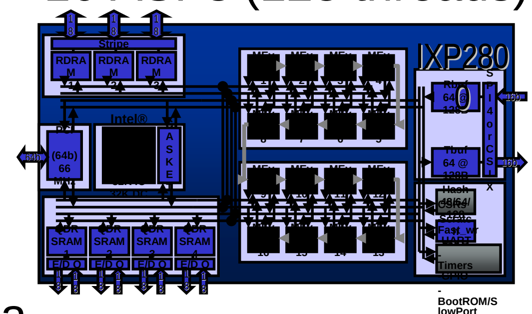

IBM Cell
1 GPP (2 threads)
8 ASPs

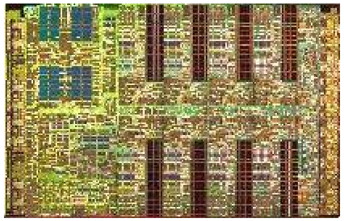

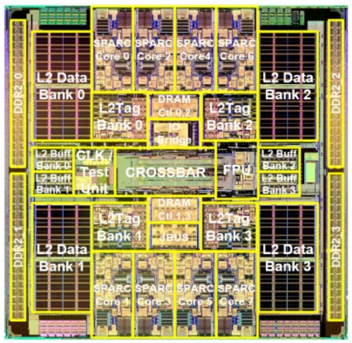

Sun Niagara
8 GPP cores (32 threads)

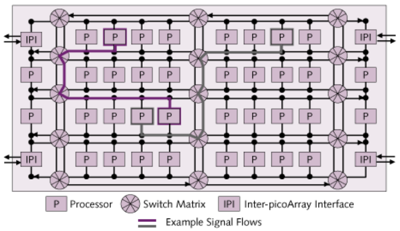

Picochip DSP
1 GPP core
248 ASPs

Cisco CRS-1
188 Tensilica GPPs

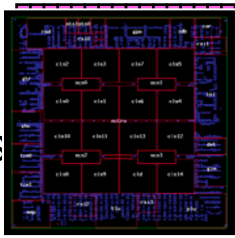

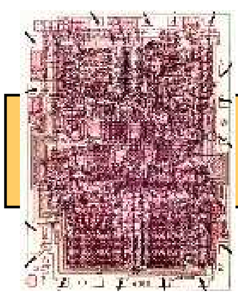

Intel 4004 (1971):
4-bit processor,
2312 transistors,
~100 KIPS,
10 micron PMOS,
11 mm^2 chip

1000s of processor cores per die

"The Processor is the new Transistor" [Rowen]

What’s Next?

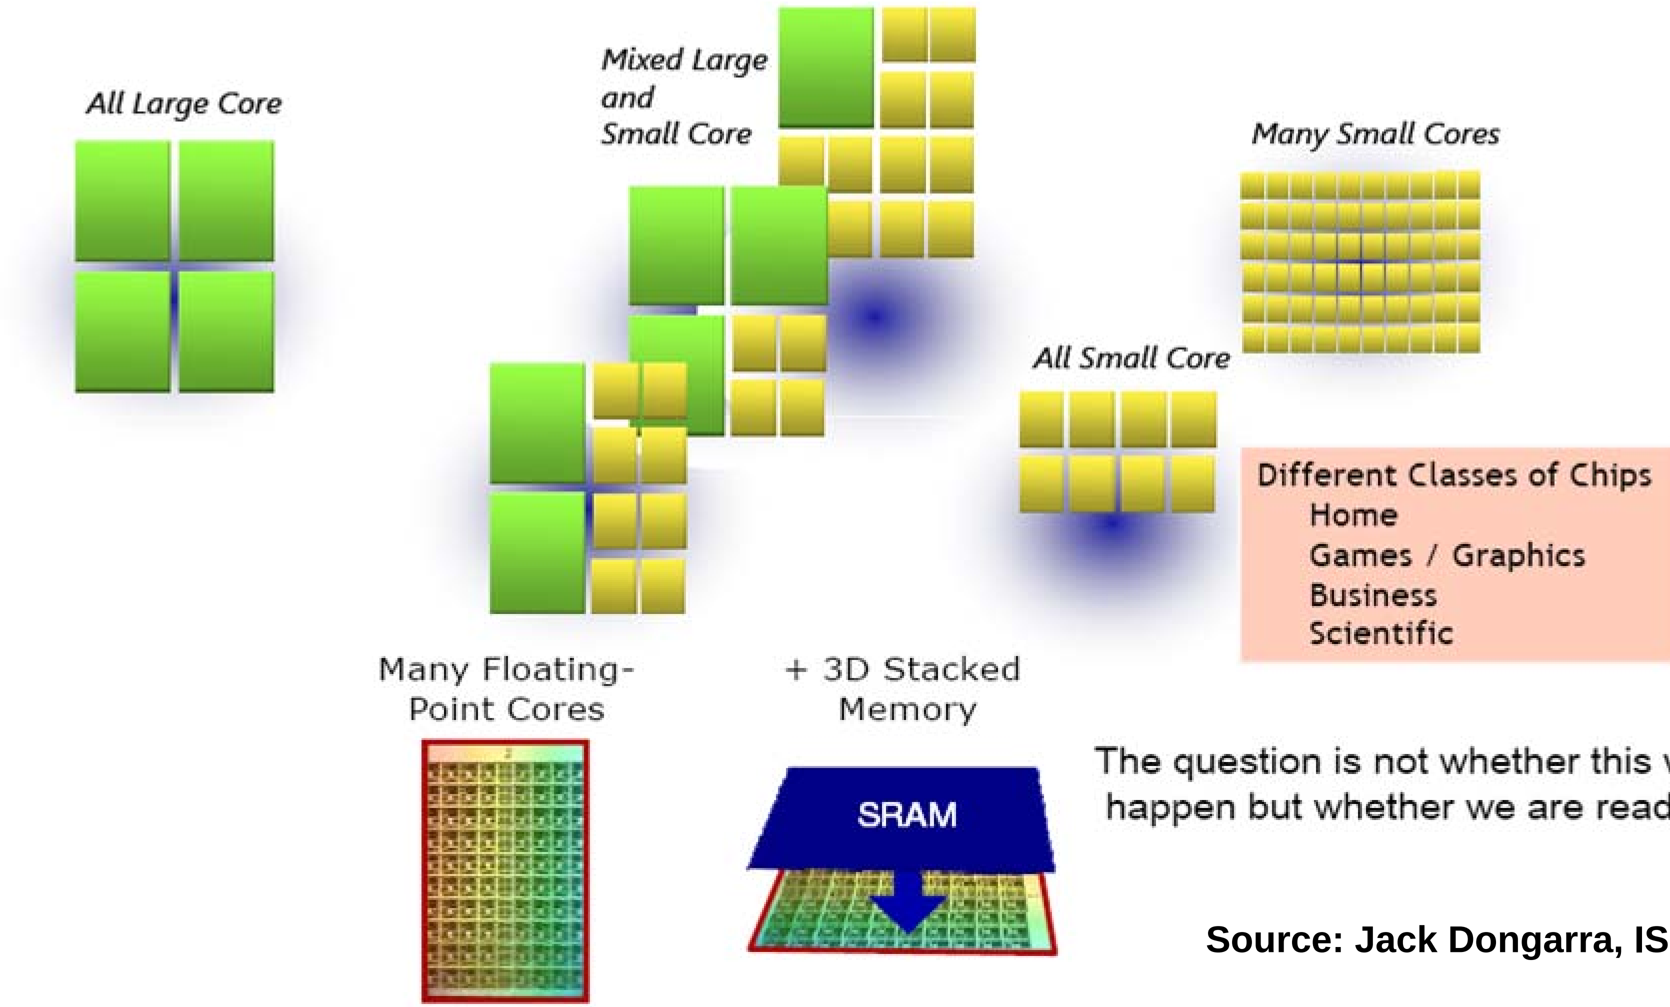

The question is not whether this will happen but whether we are ready

Source: Jack Dongarra, ISC 2008

HPC Resources

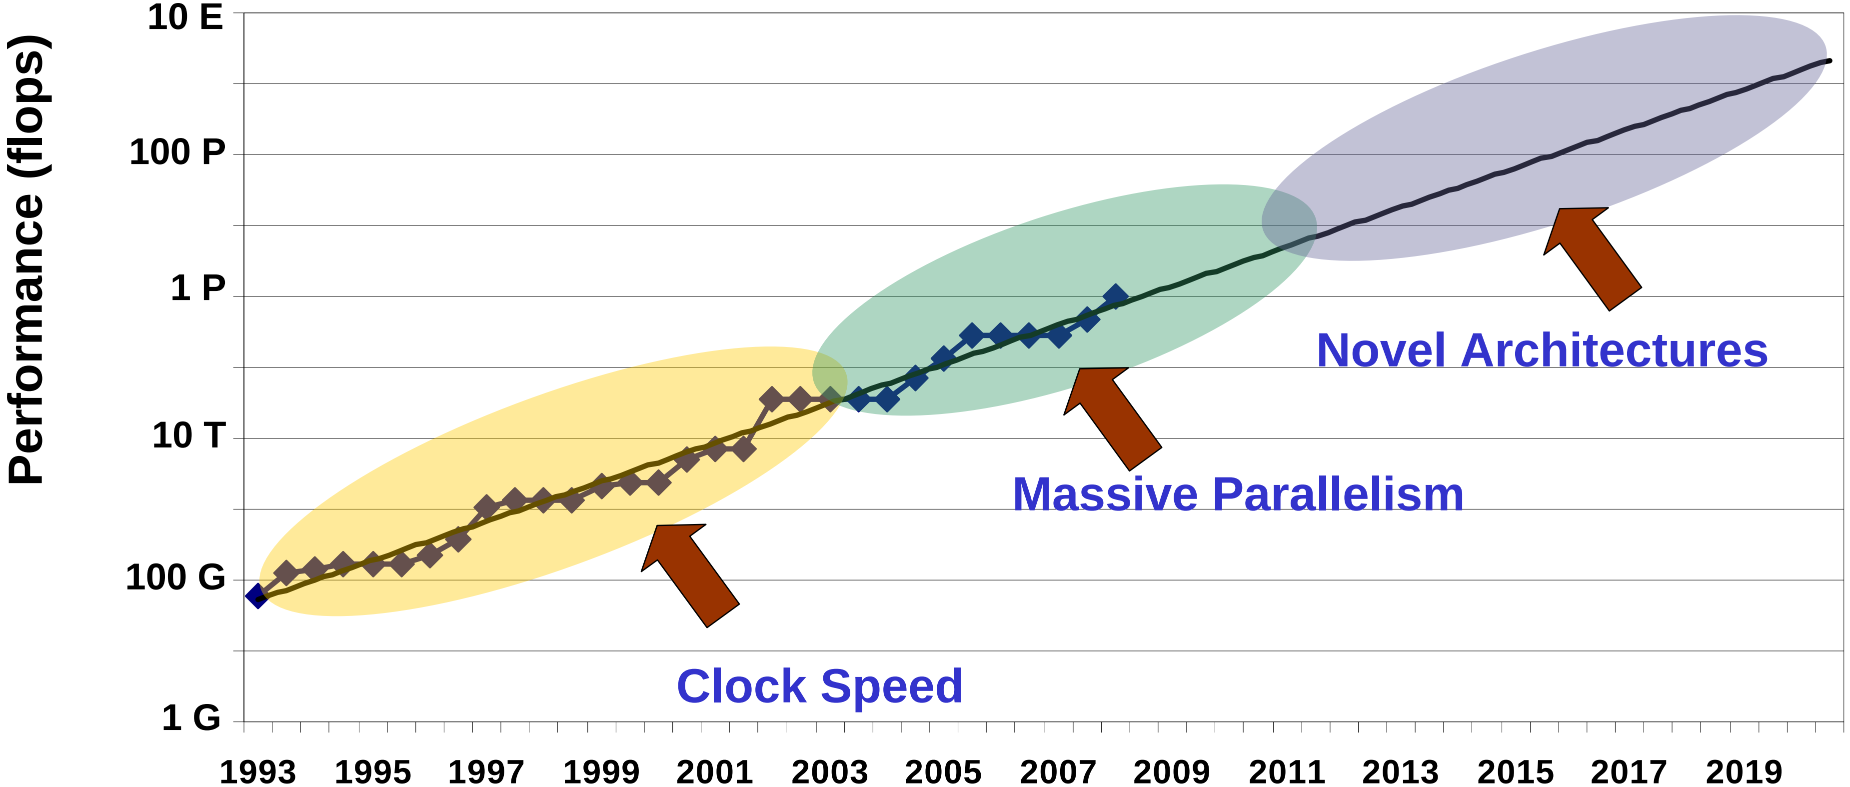

What next?

First Steps
Listening to the Community

Three Town Hall Meetings held April-June, 2007

Climate, Combustion, Fusion, Fission Solar, Biology, Socioeconomic Modeling and Astrophysics

Mathematics, Computer Science Algorithms, Software infrastructure and Cyberinfrastructure

Integrated program- investments in hardware and software research and development

Tight coupling to a selected set of science communities and the associated applied mathematics R&D.

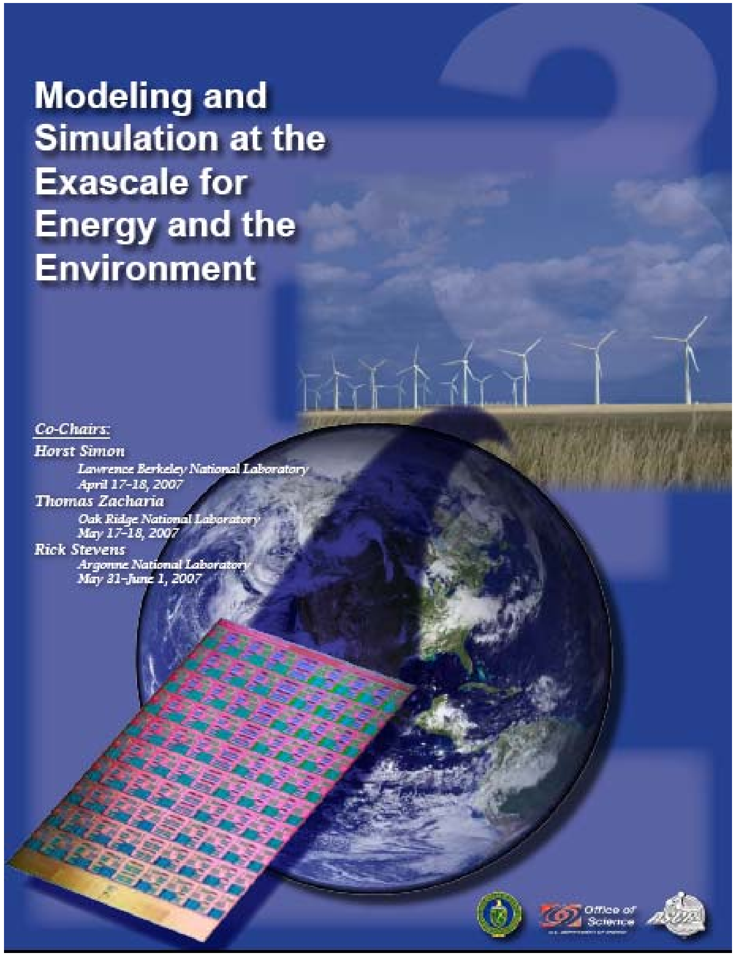

Extreme Scale Computing Workshops
Guidelines

- **I will not (and cannot) tell you what an Extreme system will look like**
- **The old approach (until about 2004) was: here is a new computer system, what science can you do with it?**
- **We want to turn the process around**
  - **Ask “What machine capability do we need to answer science questions?”**
  - **Not “What can we answer with that machine?”**
- **The Extreme Scale workshops follows the new approach**

U.S. Department of Energy

Office of Science

Extreme Scale Workshops

- Led by other Program Offices in DOE
- Previous workshops
  - BER/Climate Workshop: Challenges in Climate Change Science and the Role of Computing at the Extreme Scale, November 6-7, 2008
  - HEP/High Energy Physics Workshop: Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at the Extreme Scale, December 9-11, 2008
  - NP/Nuclear Physics Workshop: Forefront Questions in Nuclear Science and the Role of High Performance Computing

HEP Extreme Scale Workshop
An Example

- Chair: Roger Blandford, SLAC
- Co-Chairs:
  - Norman Christ, Columbia University
  - Young-Kee Kim, University of Chicago
- Breakout Sessions and Leads
  - Astrosphysics Data, Alex Szalay, Johns Hopkins University
  - Cosmology and Astrophysics Simulations, Mike Norman, USCD
  - Experimental Particle Physics, Jim Shank, Boston University and Frank Wuerthewein, UCSD
  - High Energy Theoretical Physics, Steve Sharpe, University of Washington
  - Accelerator Simulation, Panagiotis Spentzouris, FNAL

Breakout Sessions
Identify Priorities

Accelerator Simulation

Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale
December 9-11, 2008 · Menlo Park, CA

Priority Research Directions

- Design & optimize a high-energy lepton collider linac module for cost and risk reduction
  - Unscaled beam-structure simulations for the first time
- Predict beam loss and resulting activation in Intensity Frontier accelerators
  - Multi-scale, multi-physics beam dynamics simulations
    - From 10^{-3} m beams, to 10 m wakefields, to many 10^{3} m propagation
- Shorten design and build cycle of accelerator structure
  - Multi-scale, multi-physics accelerator structure simulations
    - Integrated thermal, mechanical, and electromagnetic models
- Develop and design an ultra compact plasma based collider
  - Non-linear wakefields, model experiments, self-consistent beam dynamics at full scale collider parameters

Slide 5

List of Priority Research Directions

- The Energy Frontier, using high-energy colliders to discover new particles and directly probe the architecture of the fundamental forces.
- The Intensity Frontier, using intense particle beams to uncover properties of neutrinos and observe rare processes that will tell us about new physics beyond the Standard Model.
- The Cosmic Frontier, using underground experiments and telescopes, both ground and space based, to reveal the natures of dark matter and dark energy and using high-energy particles from space to probe new phenomena.
- What is Dark Matter ?
- Does the Higgs field exist in nature ?
- Are there additional dimensions of space ?
- What are neutrinos telling us ?

List of Priority Research Directions

- Cosmic Microwave Background
- Baryon Acoustic Oscillations
- Large Scale Structure
- Weak Lensing
- The High Redshift Universe
- Extreme Environments
- Data Intensive Scalable Computing
- Next Generation of Computational Scientists

List of Priority Research Directions

- Searching for physics beyond the SM
- Testing QCD at the sub-percent level
- Simulating possible theories of physics beyond the SM
- Enhancing algorithmic performance
- Non-lattice directions

Next Steps
Making a Change

- In the President's FY2009 Budget request to Congress, funds were identified in the ASCR program for "direct support for science application leading edge developers willing to take on the risks of working with new and emerging languages and tools"
  - In partnership with other Program Offices
  - Call expected in late summer

SciDAC - X

Applications

Computing/ Networking

SciDAC X

- Materials, chemistry
- Nuclear physics
- Fusion science
- Combustion
- Alternative fuels & transportation
- Carbon management
- Groundwater
- Astrophysics
- Computational Biology
- High-energy physics
- Advanced Fuel Cycles
- Climate

- Institutes for applied mathematics and computer science
- Software & algorithm development centers
- Cyber environment for open, but secure science

INCITE X
NERSC Allocation

Multi-peta flop and Exaflop Compute Resources	300-1,000 PF, 20-80MW
ESnet	Best-in-class, reliable, robust, on demand bandwidth and network services

What computing will be needed to enable the grand challenges in High Energy Physics?

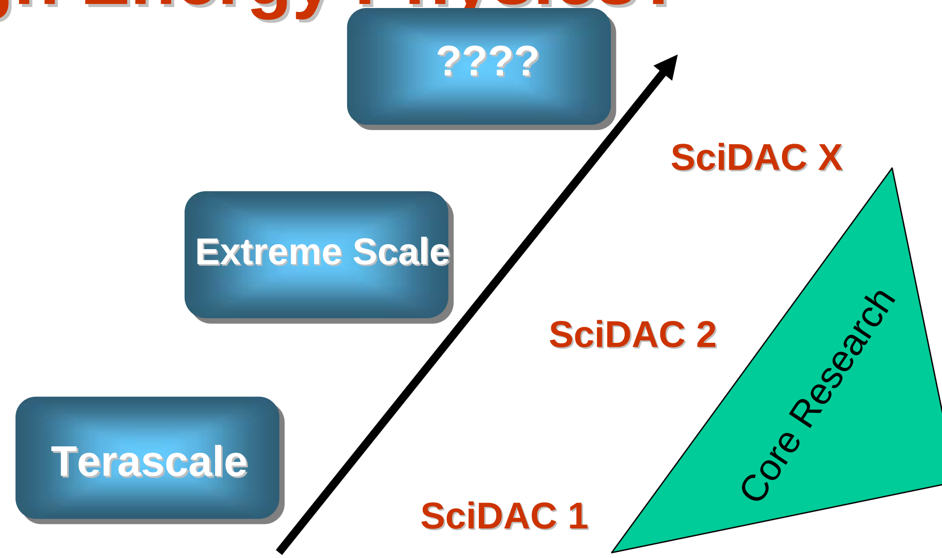

I climb the "Hill of Science,"
I "view the landscape o'er;"
Such transcendental prospect,
I ne'er beheld before!

Emily Dickinson