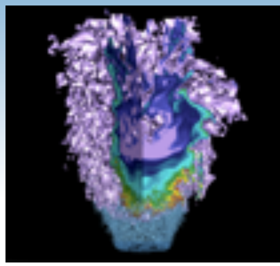 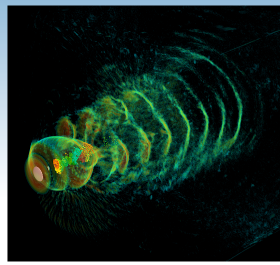 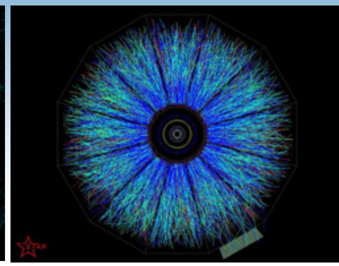 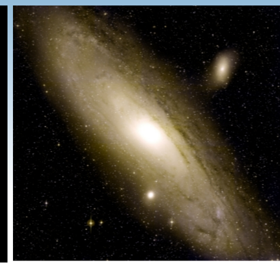 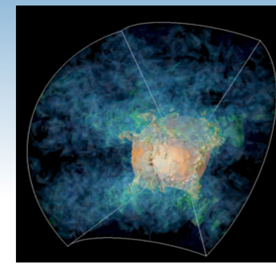 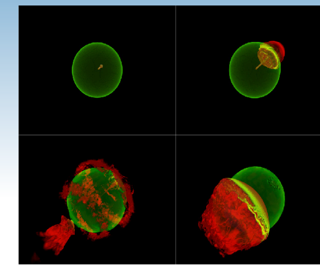

Nuclear Physics and Computing: Exascale Partnerships

Juan Meza
Senior Scientist
Lawrence Berkeley National Laboratory

Nuclear Science and Exascale

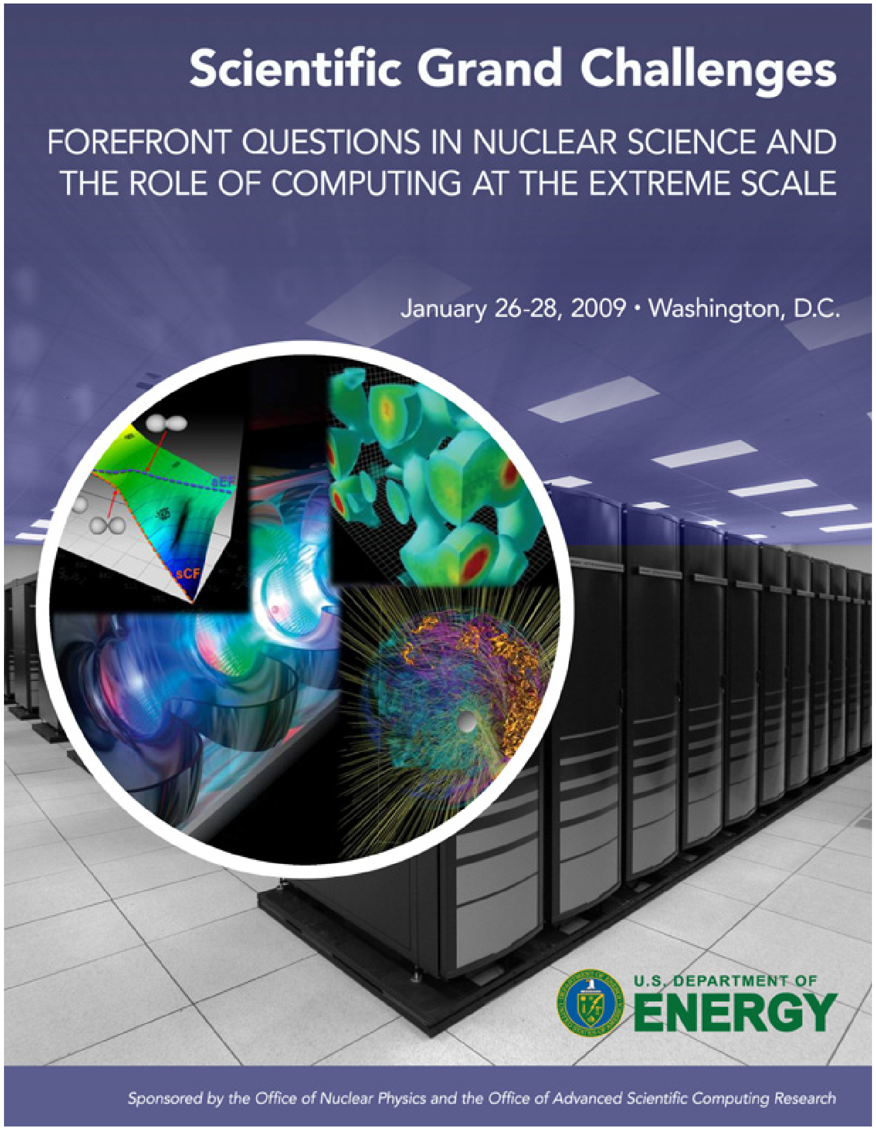

- Workshop held in DC to identify scientific challenges in nuclear physics
- Identify areas that could benefit from extreme scale computing
- Hosted by DOE Advanced Scientific Computing Research and Nuclear Physics

Five Major Areas

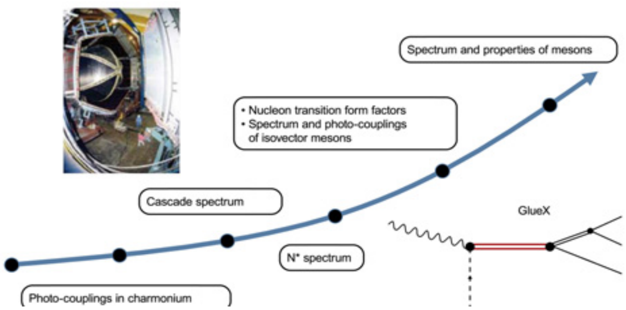

Nuclear Forces and Cold QCD

Nuclear Structure and Nuclear Reactions

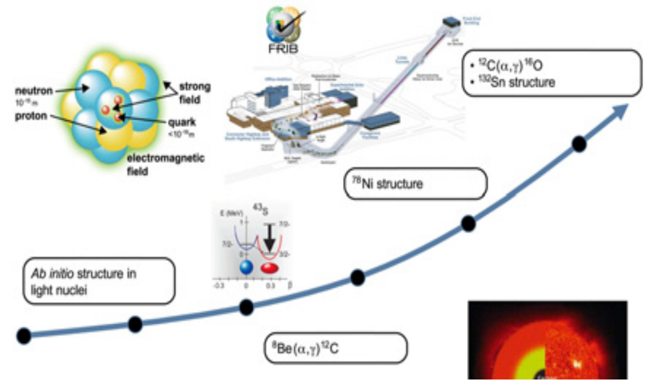

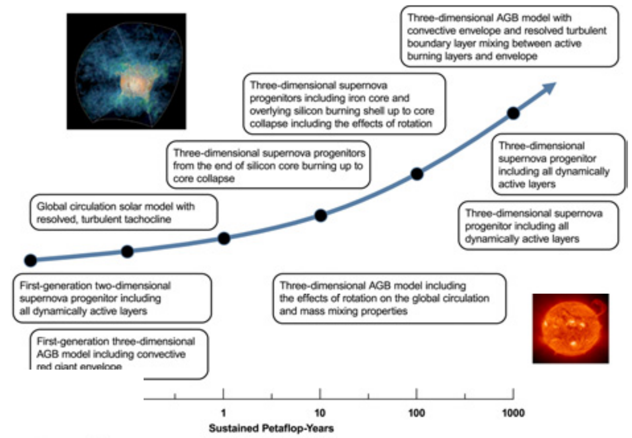

Nuclear Astrophysics

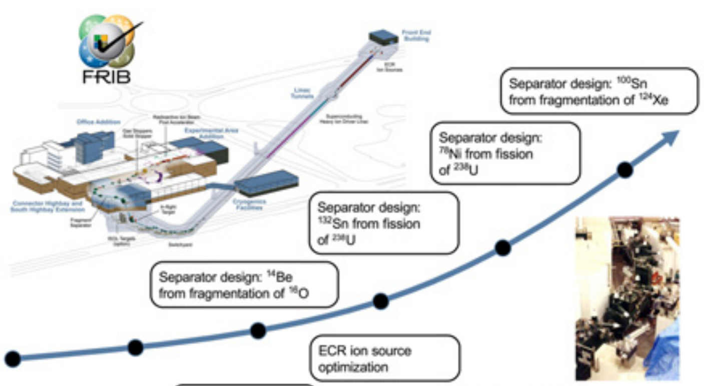

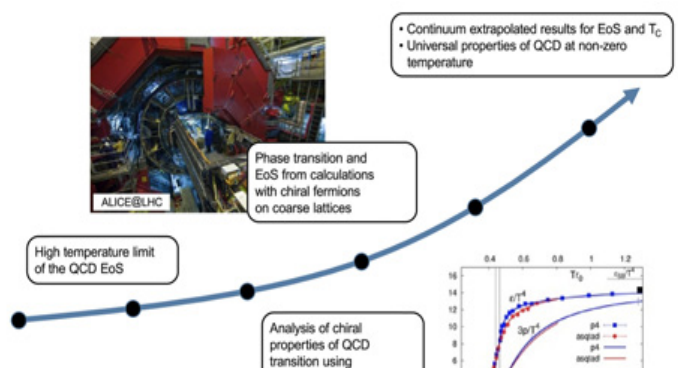

Accelerator Physics

Hot and Dense QCD

Some of the report's findings

- 19 Priority Research Directions identified
- Extreme scale computing will provide the computational resources required to perform calculations that will unify nuclear physics research
- Foster collaborations between nuclear physicists, computational scientists, applied mathematicians, and physicists outside of nuclear physics

Cold QCD and Nuclear Forces

- Spectrum of QCD
- How QCD makes a proton
- From QCD to nuclei
- Fundamental symmetries

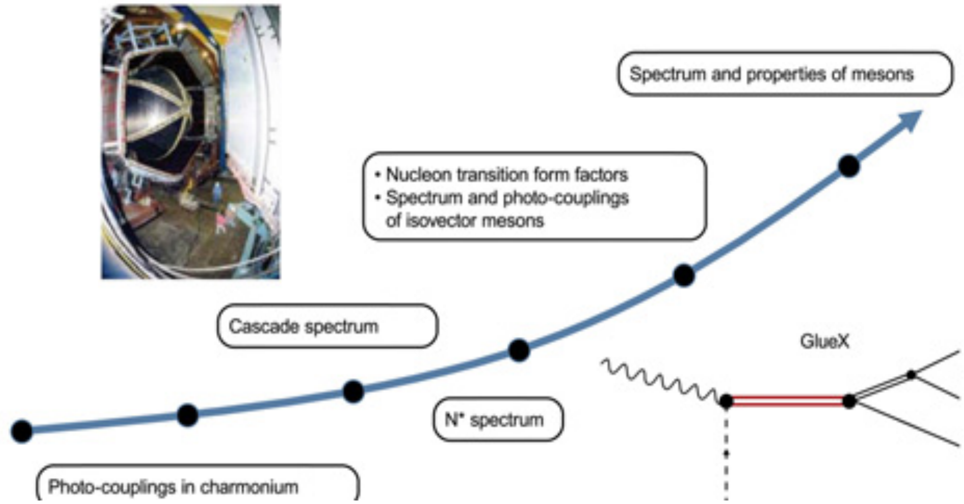

- Linear solvers for lattice -Dirac equation
- Deflation techniques and other preconditioners for the Dirac operator
- High-precision algorithms

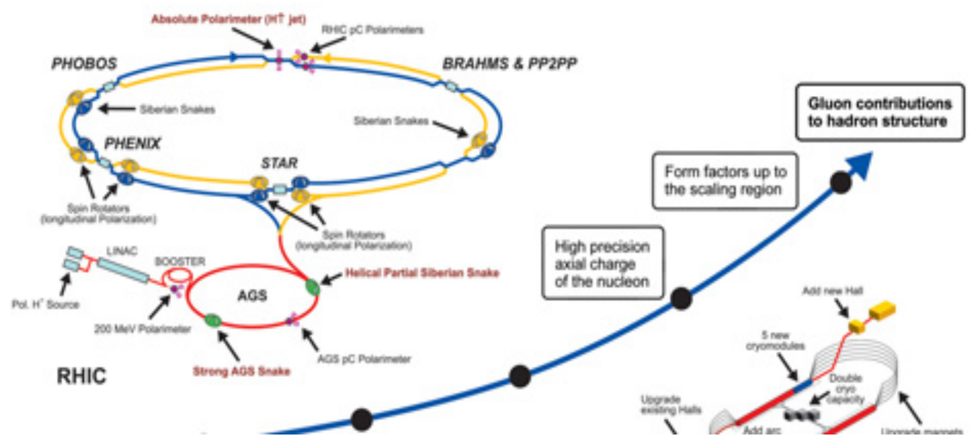

Nuclear Structure and Nuclear Reactions

- Predictive capability for the entire periodic table
- ab initio calculations of light nuclei and their reactions
- Green’s Function Monte Carlo
- ab initio no-core shell model, extended by the resonating group method (NCSM/RGM)
- Coupled cluster method

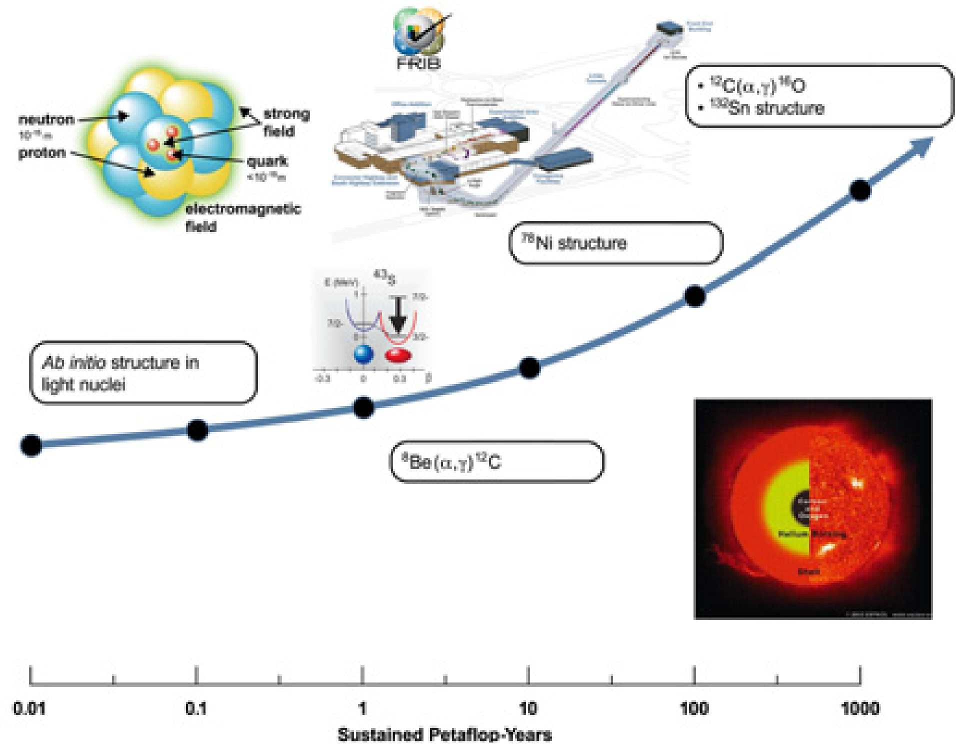

BERKELEY LAB

Nuclear Astrophysics

- Nonlinear algebraic equations
- Adaptive mesh refinement
- Sparse, structured linear systems of equations, with physics-based preconditioners
- Nonlinear, stiff, coupled, and large systems of ODEs
- Need multi-core versions of all of these

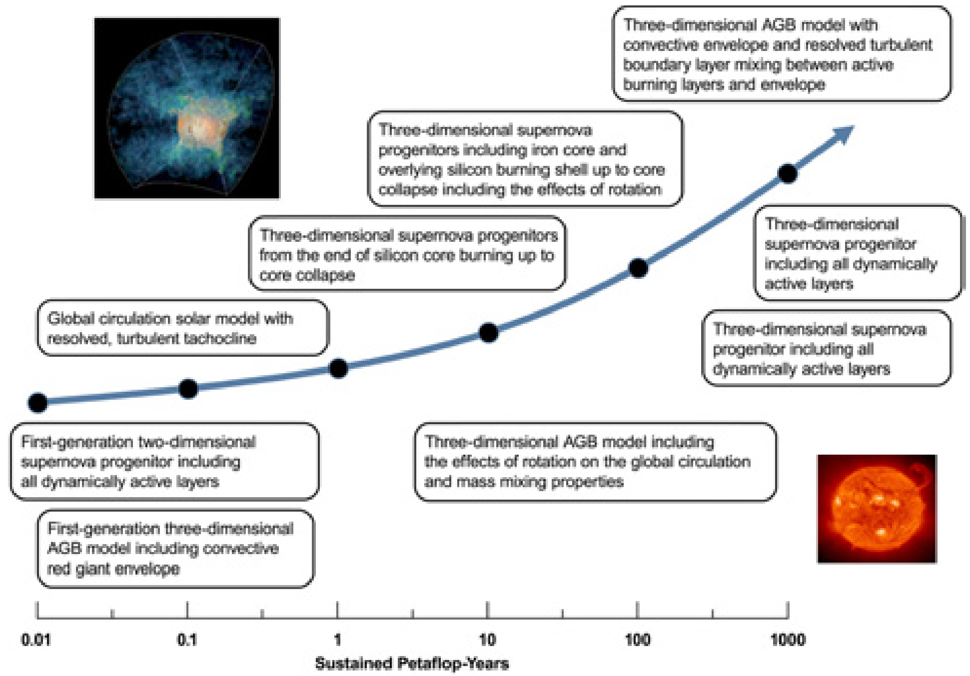

Hot and Dense QCD

- **Inversion of the fermion matrix, where non-zero entries fluctuate significantly during the calculation of new field configurations**
- **Need for deflation techniques, domain decomposition, and MG**
- **Notorious "sign problem"**

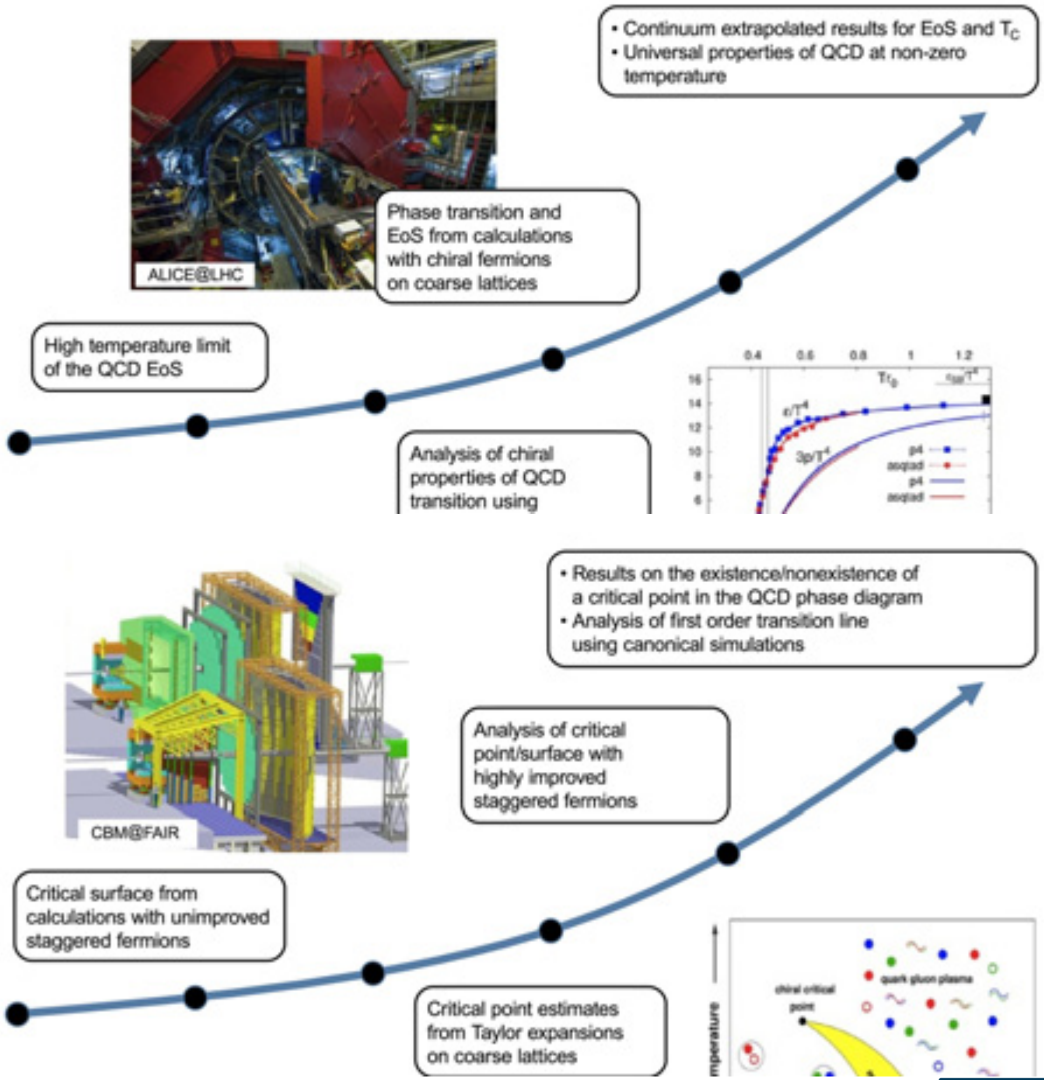

Accelerator Physics

- Maximize the production and efficiency and beam purity for rare isotope beams
- Optimal design for an electron-ion collider
- Design optimization of complex electromagnetic structures
- Advanced methods for accelerator simulation

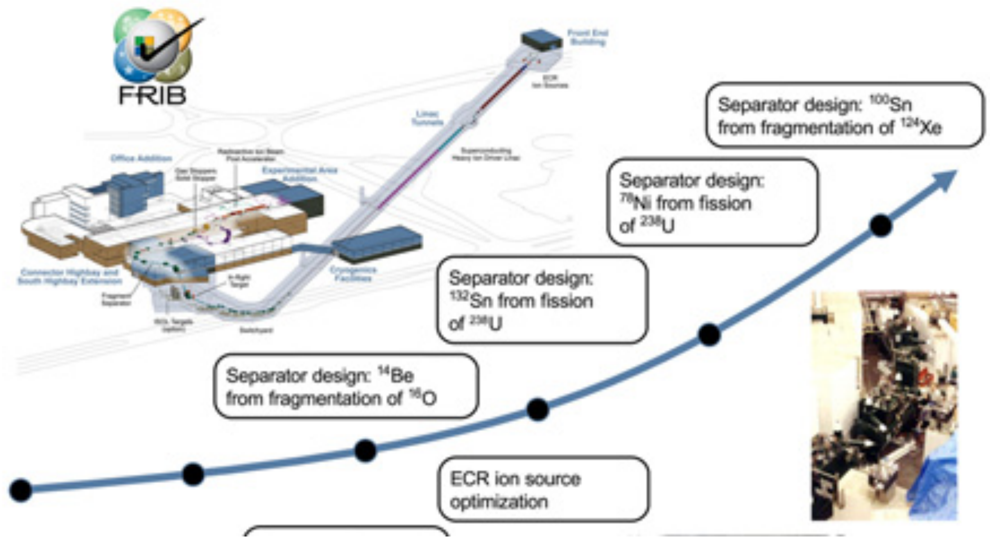

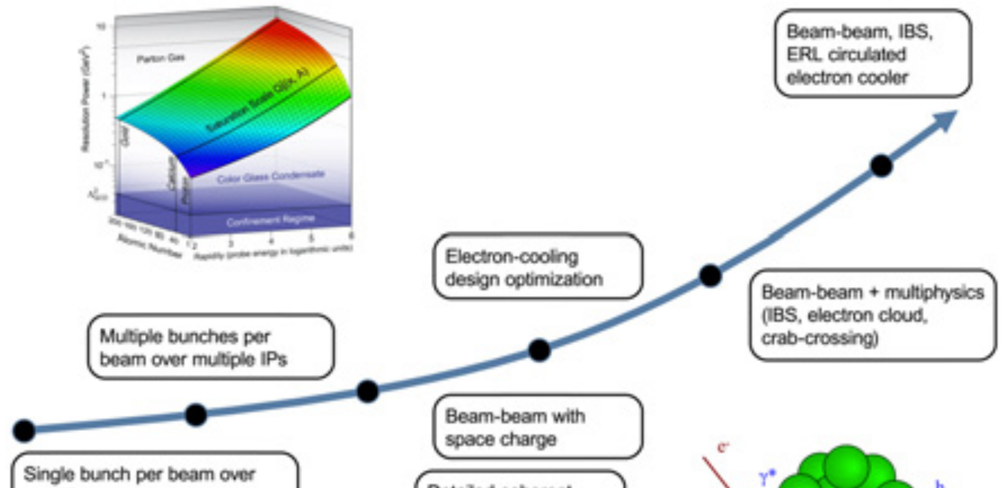

Sampling of Math/CS needs

- Solution of systems of linear equations
- Deflation techniques, domain decomposition, MG, effective preconditioners
- Monte Carlo for high-dimensional integrals
- Nonlinear algebraic equations
- Adaptive mesh refinement
- Nonlinear, stiff, coupled, and large systems of ODEs
- Global optimization, multi-objective optimization
- Multicore algorithms
- High precision algorithms
- Data management and Visualization

Computer Architecture Trends

37th List: The TOP10

Rank	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	548,352	8.162	9.90
2	National SuperComputer Center in Tianjin	NUDT	Tianhe-1A NUDT TH MPP, Xeon 6C, NVidia, FT-1000 8C	China	186,368	2.566	4.04
3	Oak Ridge National Laboratory	Cray	Jaguar Cray XT5, HC 2.6 GHz	USA	224,162	1.759	6.95
4	National Supercomputing Centre in Shenzhen	Dawning	Nebulae TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	China	120,640	1.271	2.58
5	GSIC, Tokyo Institute of Technology	NEC/HP	TSUBAME-2 HP ProLiant, Xeon 6C, NVidia, Linux/Windows	Japan	73,278	1.192	1.40
6	DOE/NNSA/LANL/SNL	Cray	Cielo Cray XE6, 8C 2.4 GHz	USA	142,272	1.110	3.98
7	NASA/Ames Research Center/NAS	SGI	Pleiades SGI Altix ICE 8200EX/8400EX	USA	111,104	1.088	4.10
8	DOE/SC/ LBNL/NERSC	Cray	Hopper Cray XE6, 6C 2.1 GHz	USA	153,408	1.054	2.91
9	Commissariat a l'Energie Atomique (CEA)	Bull	Tera 100 Bull bullx super-node S6010/ S6030	France	138.368	1.050	4.59
10	DOE/NNSA/LANL	IBM	Roadrunner BladeCenter QS22/LS21	USA	122,400	1.042	2.34

Projected Performance Development

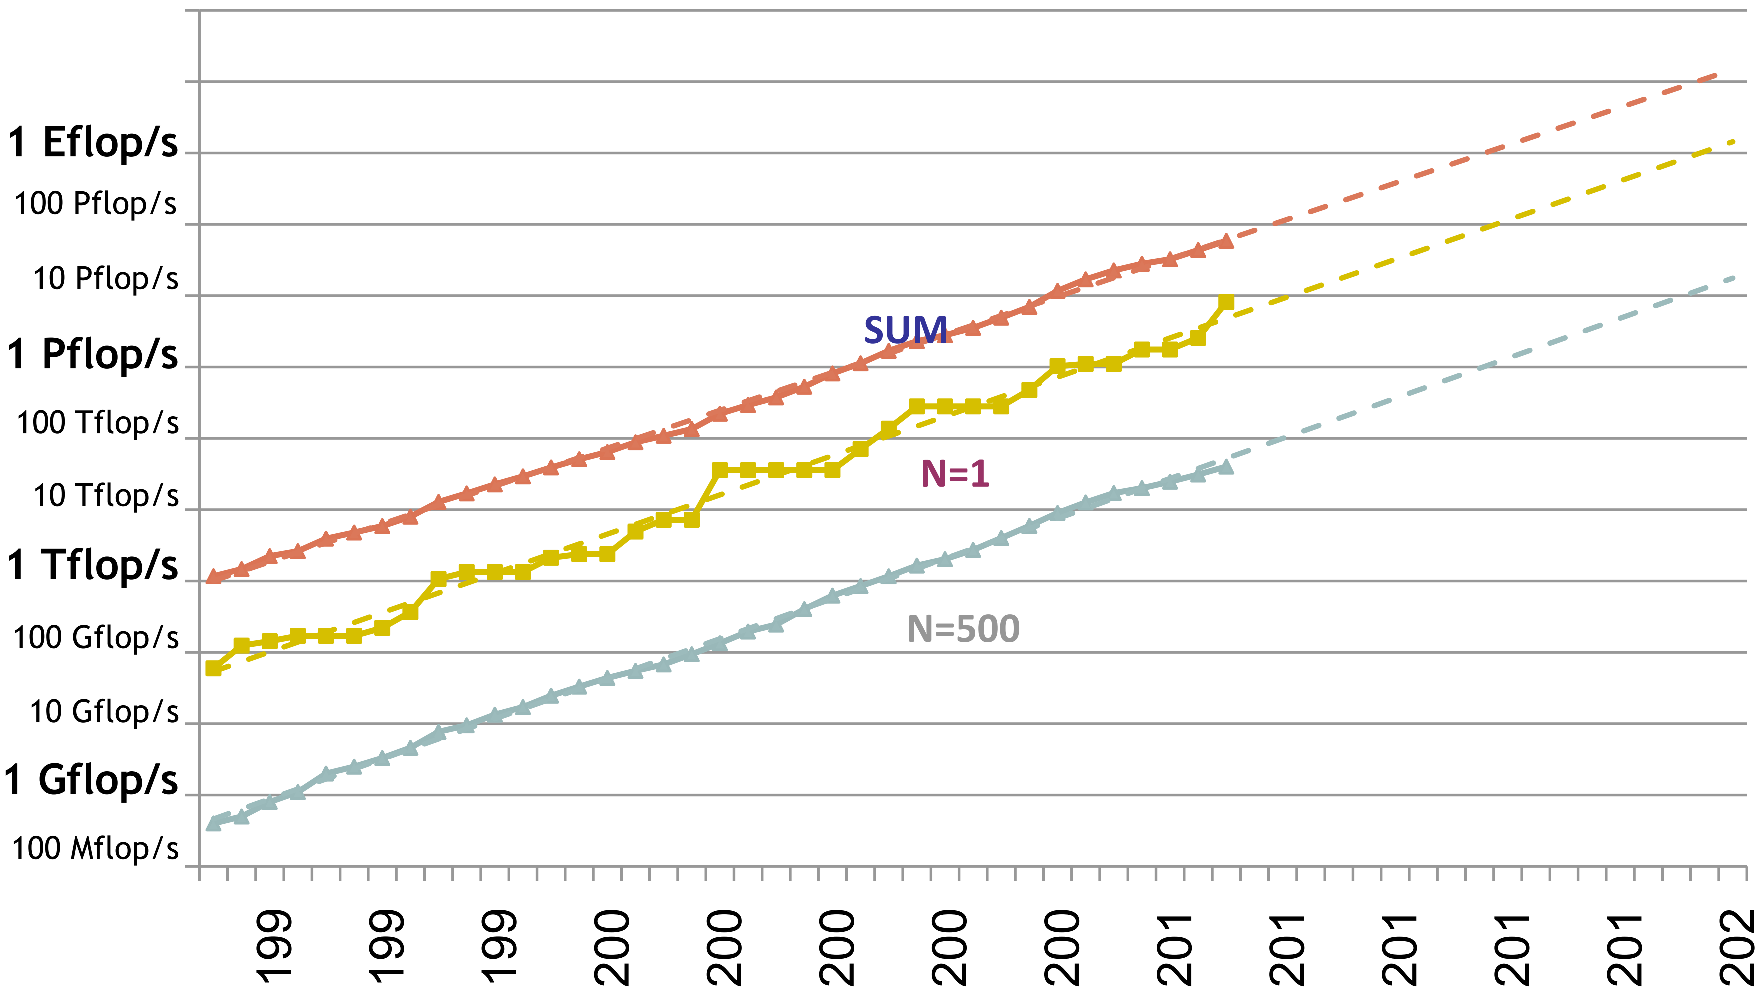

Top 500 slides courtesy of Erich Strohmaier, LBNL

Power Consumption

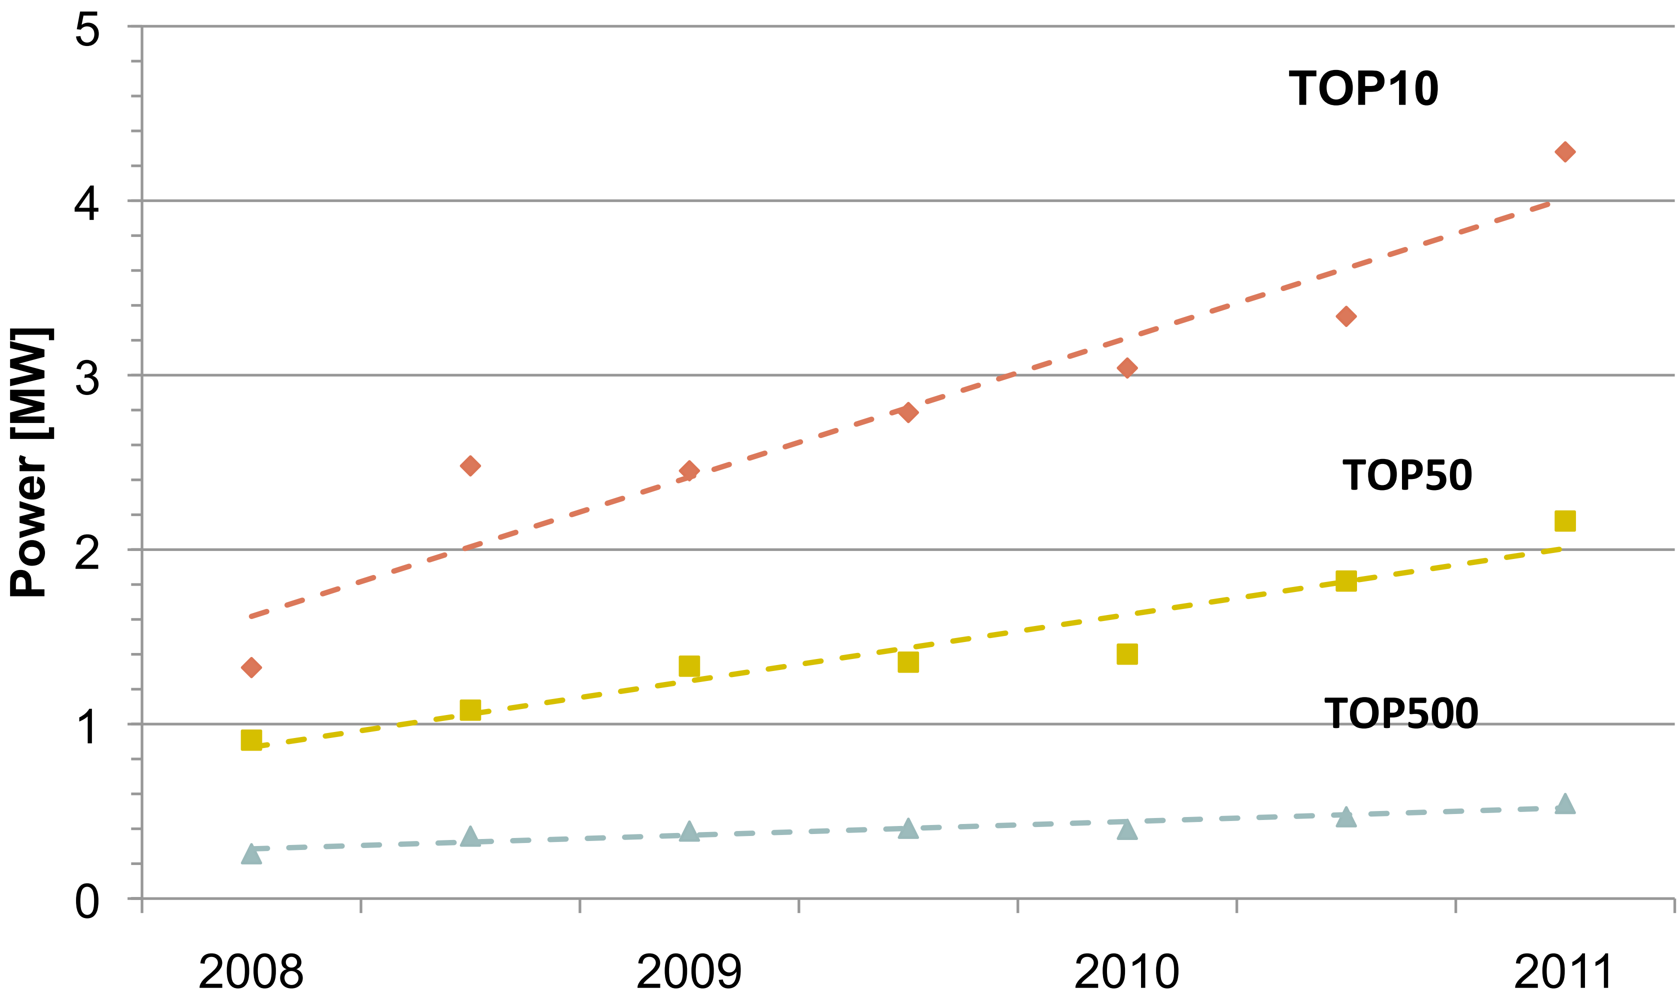

Top 500 slides courtesy of Erich Strohmaier, LBNL

Power Efficiency

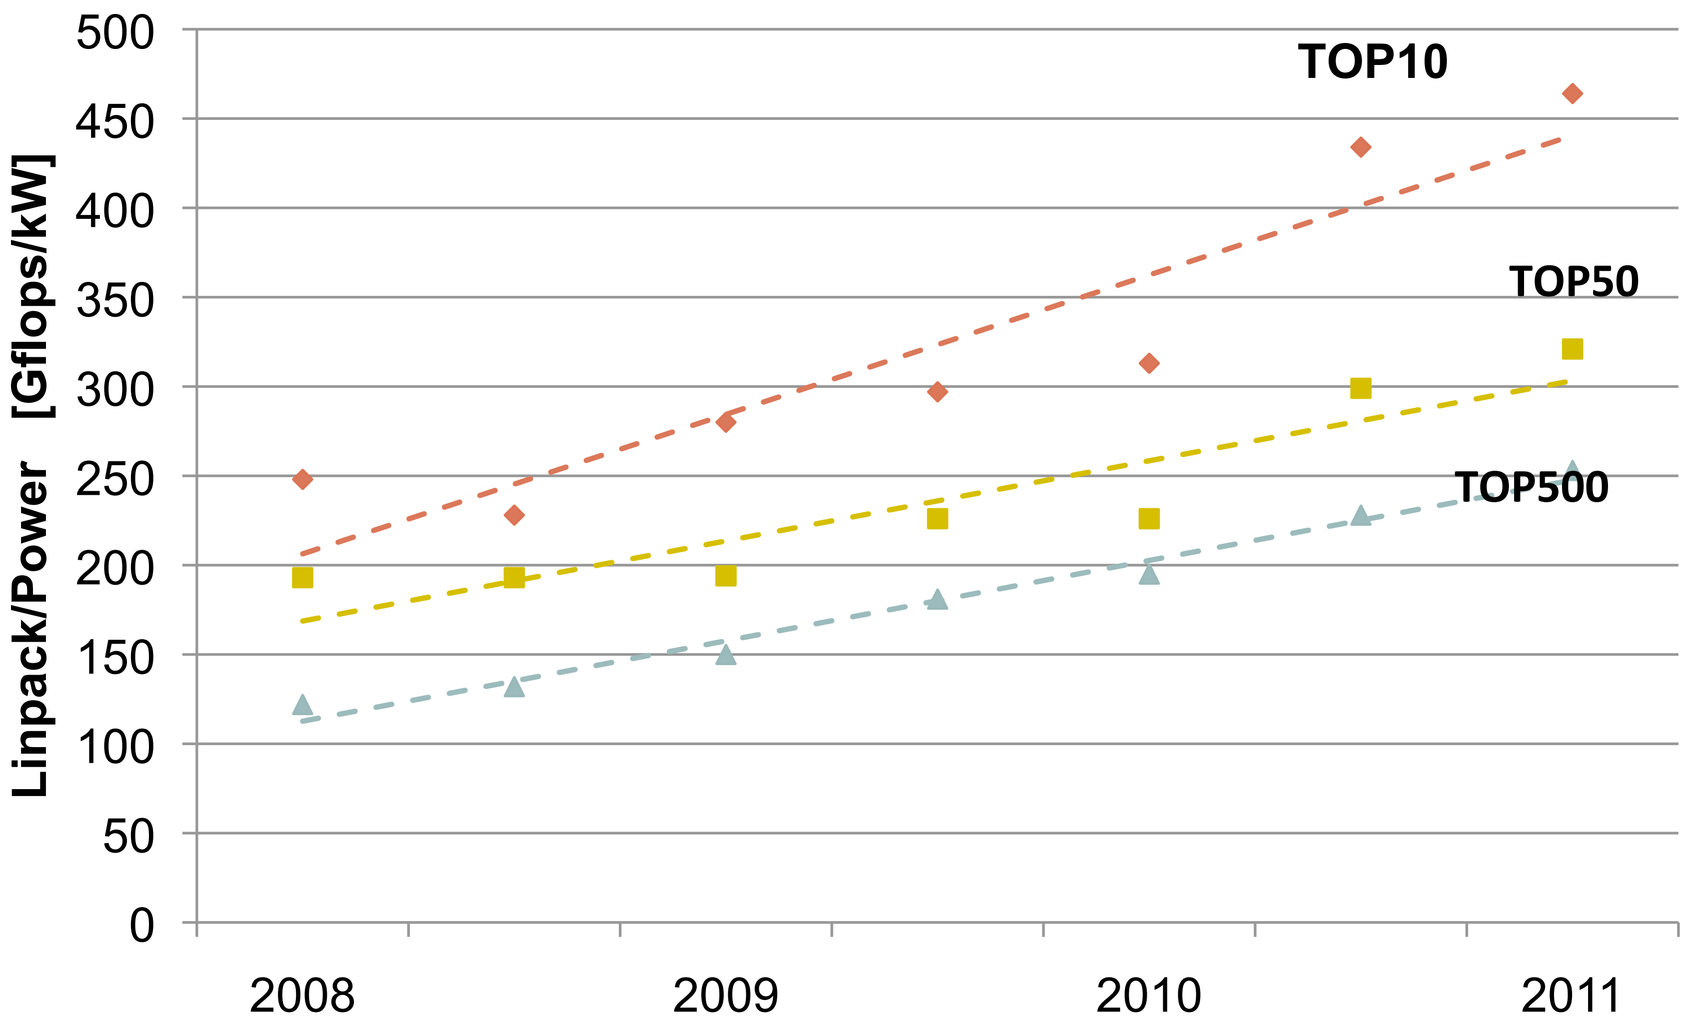

Top 500 slides courtesy of Erich Strohmaier, LBNL

Cores per Socket

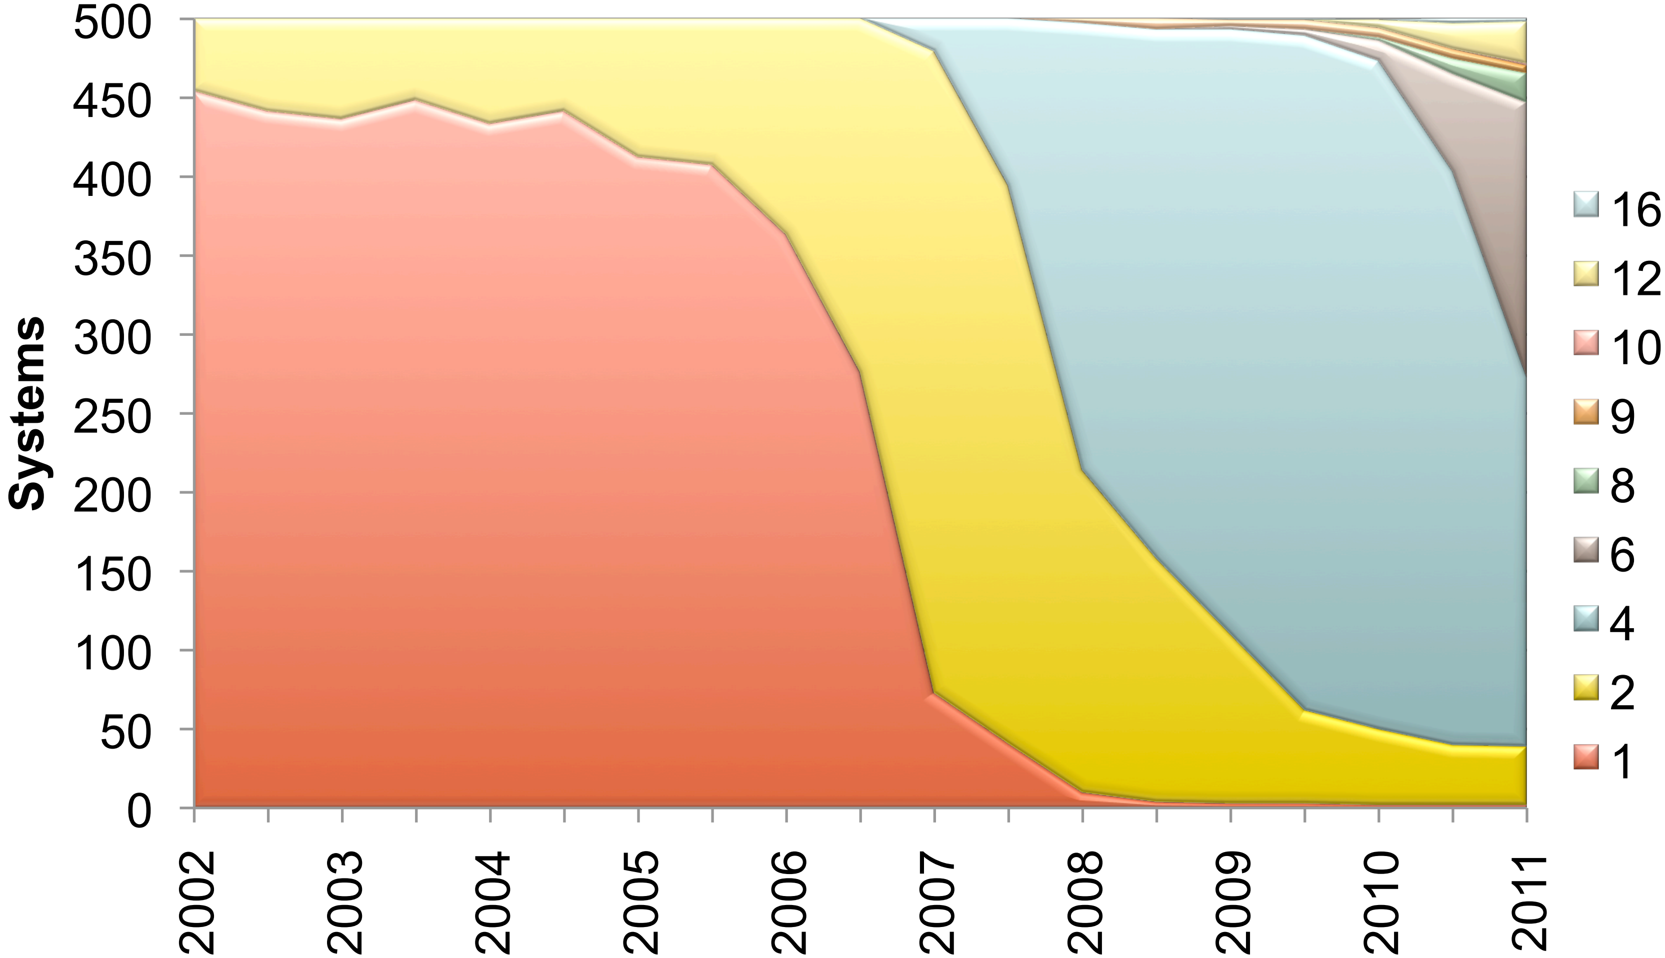

Top 500 slides courtesy of Erich Strohmaier, LBNL

Computing Performance Improvements will be Harder than Ever

- Used to rely on processor speeds increases plus parallelism
- Single processors are not getting faster
- Key challenge is energy!

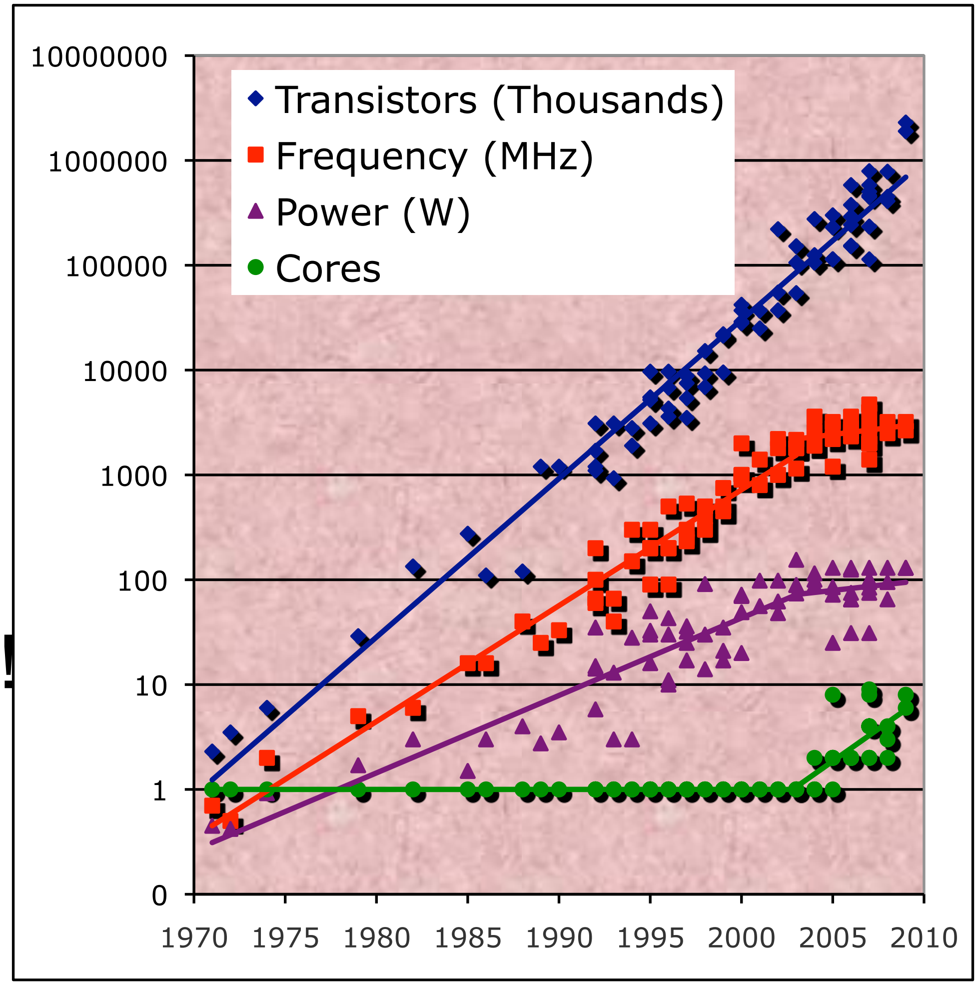

Future: All in added concurrency, include new on-chip concurrency

New Processors Means New Software

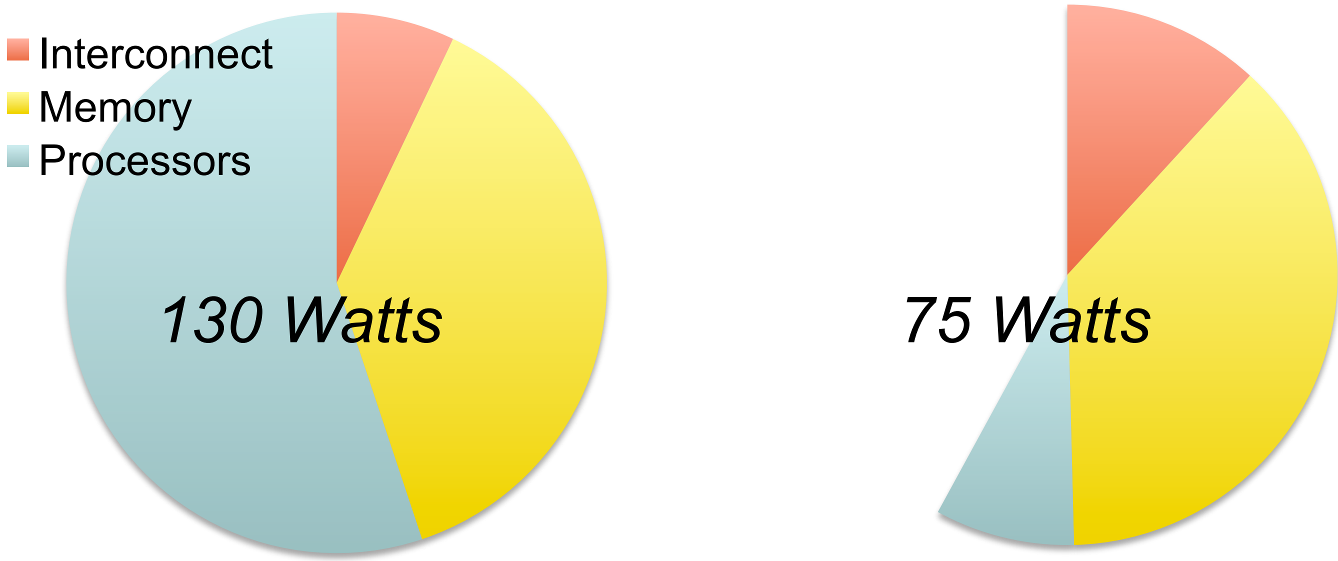

Usual Processors *Manycore processors*

- **Exascale systems will be built from chips with thousands of tiny processor cores**
  - **The architecture (how they will be organized) is still an R&D problem, but likely a mixture of core types**
  - **They will require a different kind of programming and new software**

Memory and Network Also Important

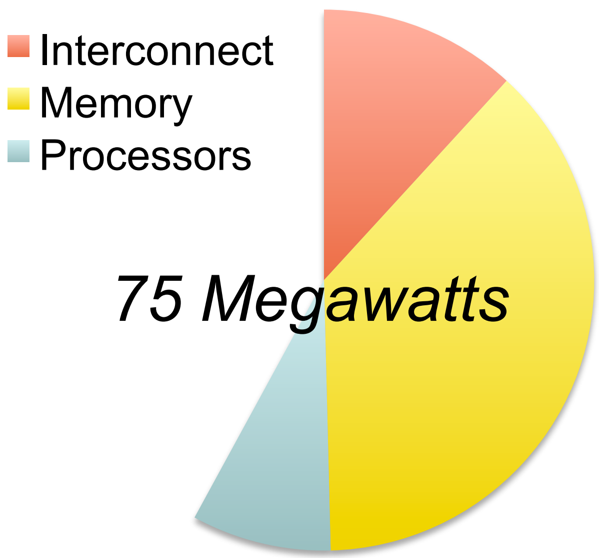

Usual memory + network

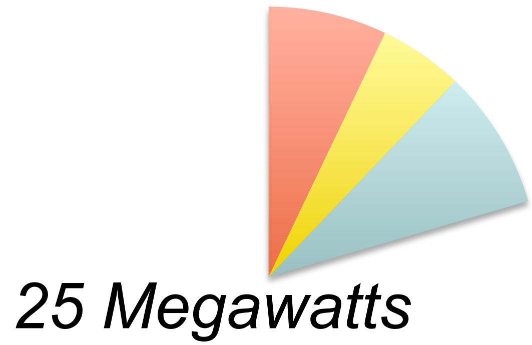

New memory + network

- **Memory as important as processors in energy**
  - **Requires basic R&D to lower energy use by memory stacking and other innovations**
- **True for all computational problems, but especially data intensive ones**

Math and Computing Science Trends

Strong-Scaling Drives Change in Algorithm Requirements

- Parallel computing has thrived on weak-scaling for past 15 years
- Flat CPU performance increases emphasis on strong-scaling
- Focus on Strong Scaling will change requirements:
  - Concurrency: Will double every 18 months
  - Implicit Methods: Improve time-to-solution (pay for allreduce)
  - Multiscale/AMR methods: Only apply computation where it is required – (need better approaches for load balancing)

Exascale Machines Require New Algorithms

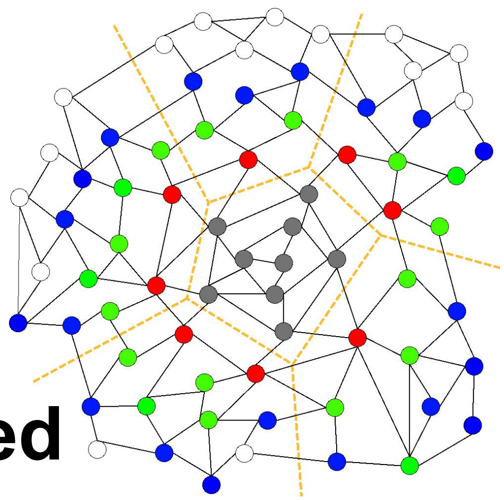

- Even with innovations:
  - arithmetic (flop) is "free"
  - data movement is expensive (time and energy)
- New model for algorithms and applied mathematics on these machines
  - Algorithms avoid data movement
  - Significant change from tradition of counting arithmetic operations
- Exascale will enable new science problems, which will also require new algorithms

Programming Models

- **Programming Languages are a reflection of the underlying hardware**
  - **Model is rapidly diverging from our current machine model**
  - **If the programming model doesn't reflect the costs on underlying machine, then cannot fix this problem**
- **Important Language Features for Future**
  - **Hierarchical data layout statements**
  - **Loops that execute based on data locality (not thread #)**
  - **Support for non-SPMD programming model**
  - **Annotations to make functional guarantees**
  - **Abstractions for disjoint memory spaces**

Slide courtesy of John Shalf, LBNL

Predictive Computational Science

- **Accuracy + Uncertainty Quantification = Predictive**
  - High-fidelity algorithms, high-resolution models
  - Quantification of uncertainty due to mathematical models, input parameters, algorithms, codes
  - First necessary step in validating simulations
- **Need to develop capability in UQ and Verification & Validation and tools to facilitate workflows**

BERKELEY LAB
Lawrence Berkeley National Laboratory

Data Enabled Science

- Many scientific problems are now rate-limited by access to and analysis of data
  - Climate models must integrate petascale data sets
  - Genome sequencing is outpacing computing and algorithms
  - Observational data from instruments will also outpace current methods
- Data mining and analysis methods will not scale with growth of current data sets
- Need to develop capabilities to rapidly access, search, and mine datasets

Summary of Challenges to Exascale

- **System power is the primary constraint**
- **Concurrency (1000x today)**
- **Memory bandwidth and capacity are not keeping pace**
- **Processor architecture is an open question**
- **Software needs to change to match architecture**
- **Algorithms need to minimize data movement, not flops**
- **I/O and Data analytics also need to keep pace**
- **Reliability and resiliency will be critical at this scale**

Unlike the last 20 years most of these (1-7) are equally important across scales, e.g., 100 10-PF machines

Final thoughts

"Nuclear physicists are benefitting greatly from interactions and collaborations with computer scientists, and such collaborations must be further embraced and strengthened as researchers move toward the era of extreme computing"

Scientific Grand Challenges: Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale, 2009

Questions to jump-start discussion

- How far can we go by just fine-tuning our existing algorithms for new architectures?
- Will we have to change the way we program large simulations?
- Can we reproduce results of simulations that use 1,10, 100M cores?
- And if we can reproduce them, how do we quantify the uncertainty in the results?

Backup Slides