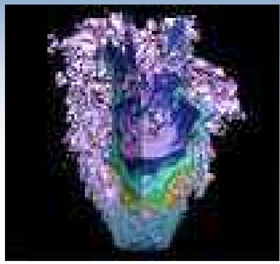
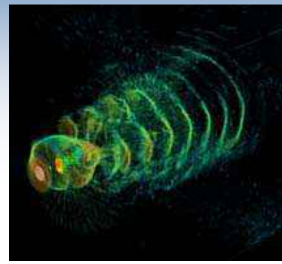
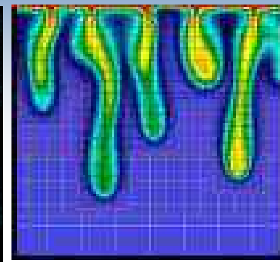

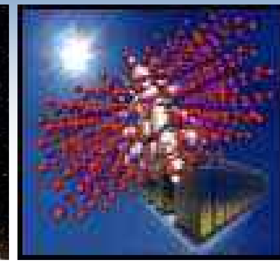
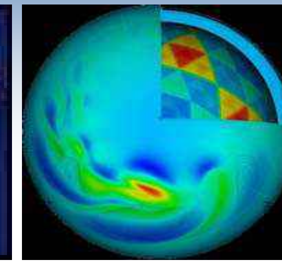

Exascale Computing: Technical Challenges

Kathy Yelick

Associate Laboratory Director and NERSC Director

Lawrence Berkeley National Laboratory

EECS Professor, UC Berkeley

Hopper System at NERSC

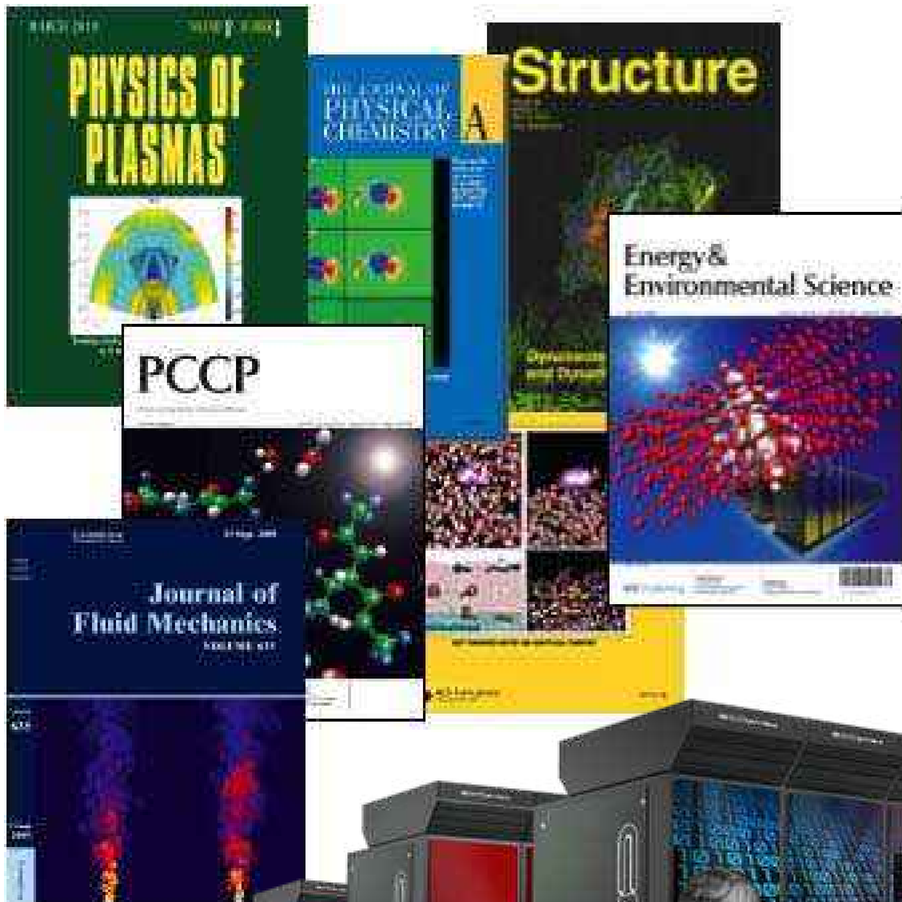

Represents broad science needs

- 4000 users, 500 science projects
- Over 1,500 publications each year
- More users than any other DOE Science facility: 65% from universities in 48 states

Petaflop Cray XE6, Hopper system

- Selected for best *application performance* per dollar and per Watt
- About 3 megawatts for 1.25 Petaflops

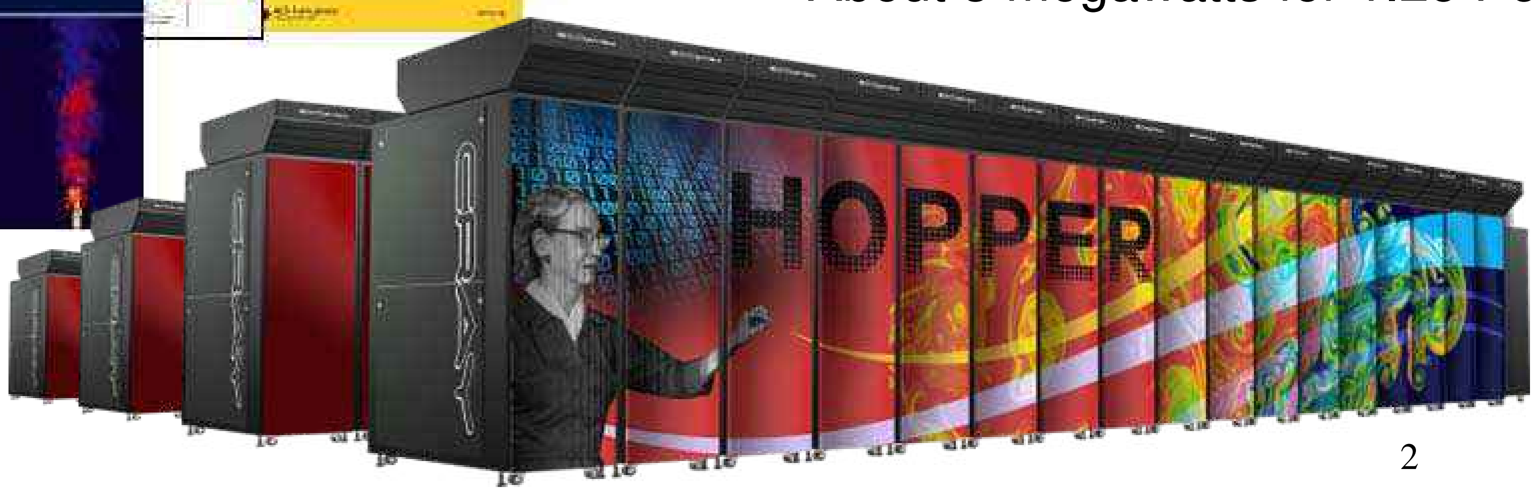

Exascale is Energy Efficient Computing

At $1M per MW, energy costs are substantial

- 1 petaflop in 2010 uses 3 MW
- 1 exaflop in 2018 at 200 MW with "usual" scaling
- 1 exaflop in 2018 at 20 MW is target

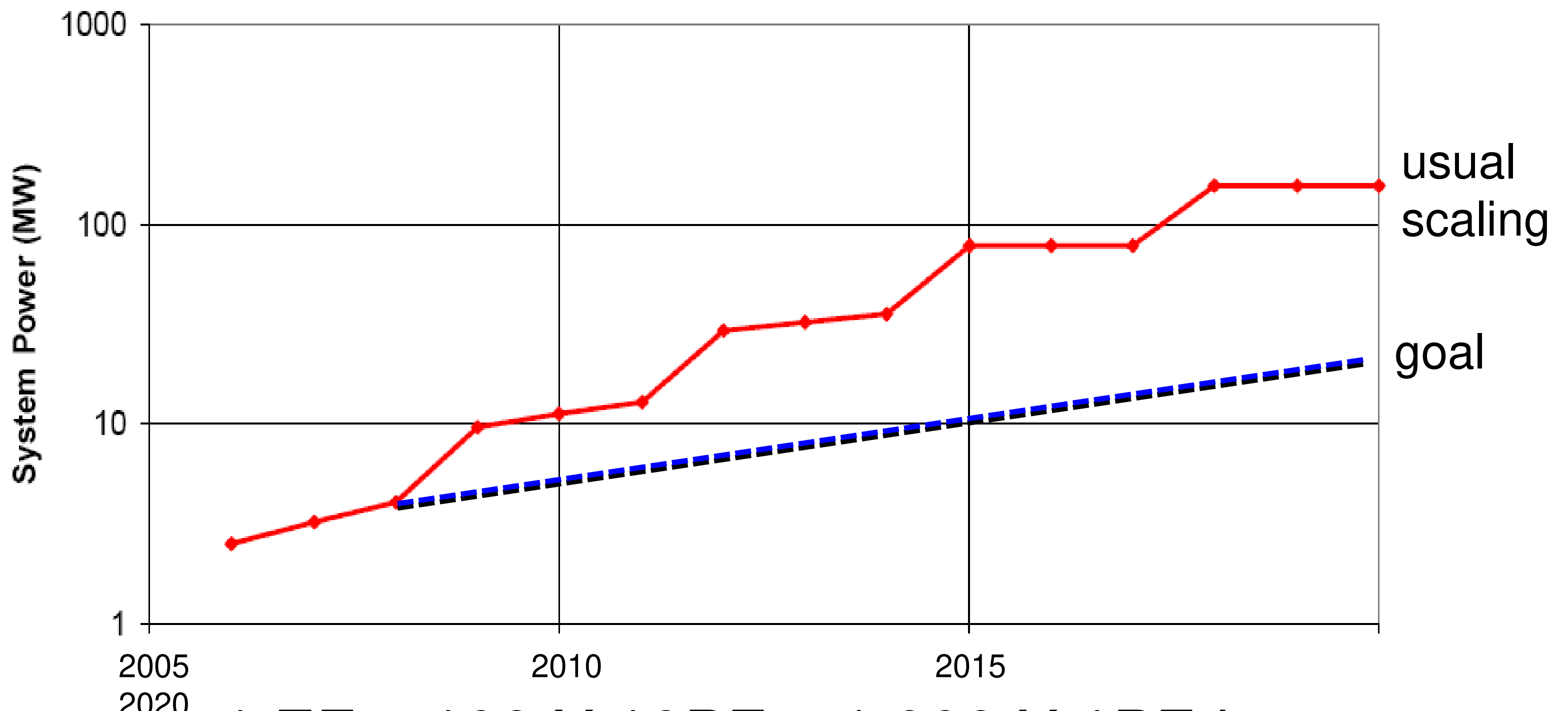

1 EF = 100 X 10PF = 1,000 X 1PF !

Computing Performance Improvements will be Harder than Ever

- Used to rely on processor speeds increases plus parallelism
- Single processors are not getting faster
- Next 1000x increase (peta to exa) will be harder than ever
- Key challenge is energy!

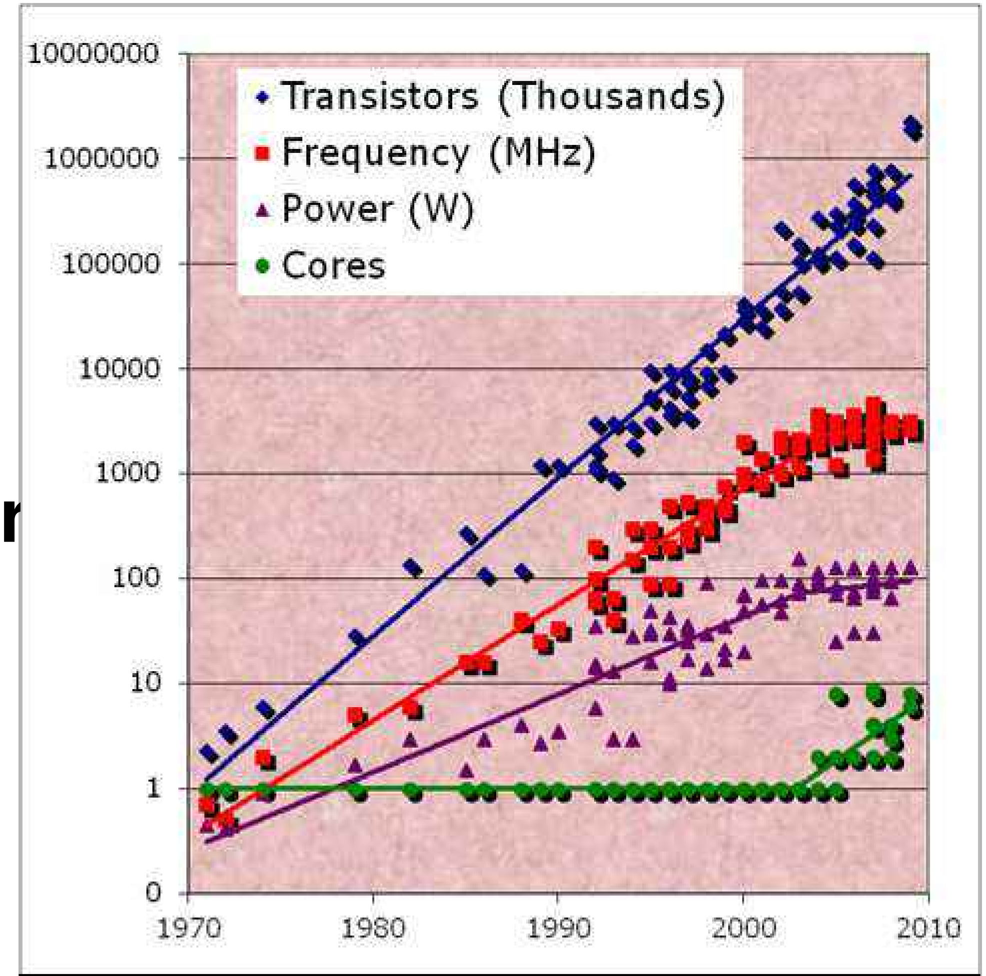

All future performance in added concurrency, including relatively new on-chip concurrency

National Academies Report on Computing Performance

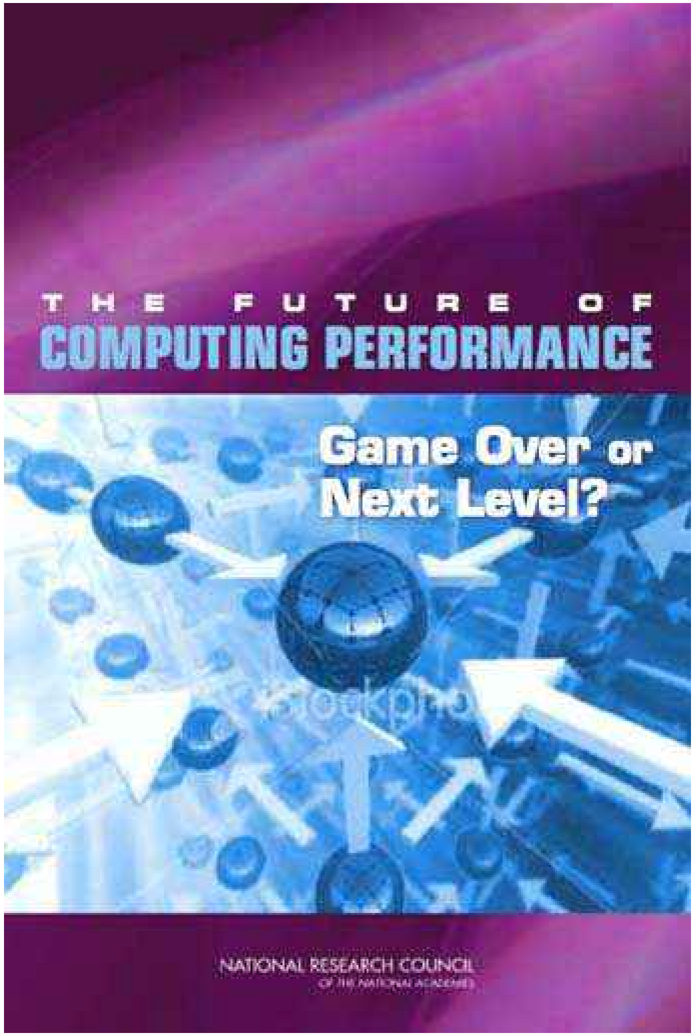

- Past performance increases have driven innovation
  - Commercial innovations, science, engineering, defense
  - All computing: data analysis, simulation and control
- Challenges in report
  - Processor speeds stalled
  - Energy is limitation
- Report symposium in DC
  - 3/22, http://www.cstb.org

New Processor Designs are Needed to Save Energy

Cell phone processor
(0.1 Watt, 4 Gflop/s)

Server processor
(100 Watts, 50 Gflop/s)

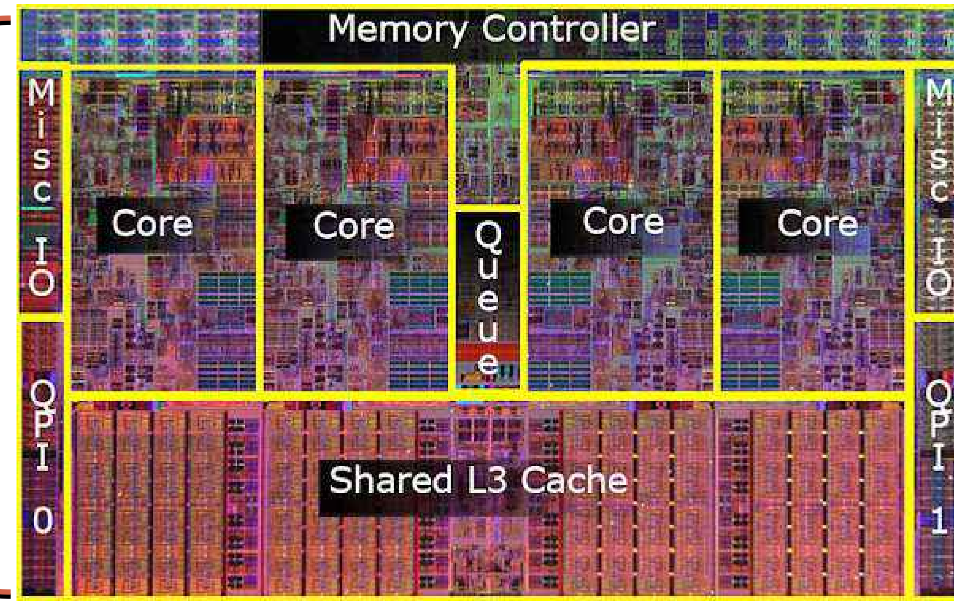

- **Server processors have been designed for performance, not energy**
  - **Graphics processors are 10-100x more efficient**
  - **Embedded processors are 100-1000x**
  - **Need manycore chips with thousands of cores**

New Processors Means New Software

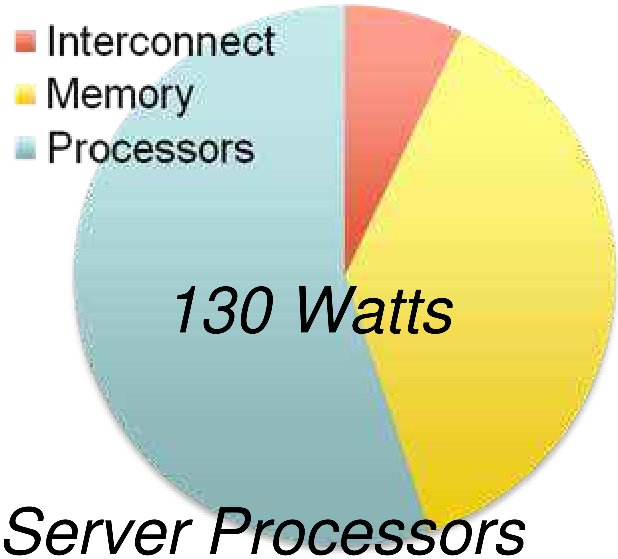

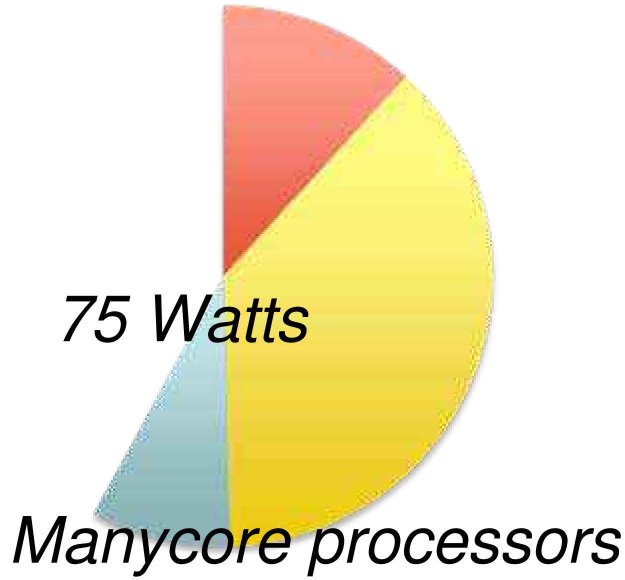

- **Exascale systems will be built from chips with thousands of tiny processor cores**
  - The architecture (how they will be organized) is still an R&D problem, but likely a mixture of core types
  - They will require a different kind of programming and new software

New Memory and Network Technology Needed to Lower Energy

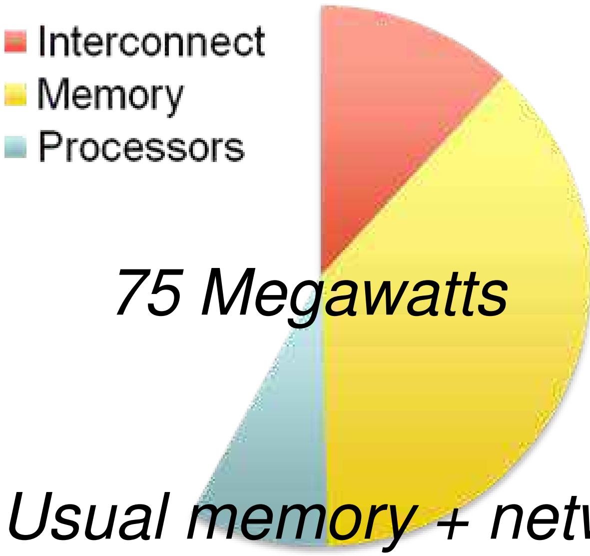

75 Megawatts

Usual memory + network

25 Megawatts

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**

Exascale Machines will Require New Algorithms and Applications

- **Even with innovations:**
  - Arithmetic (floating point) is essentially "free"
  - Data movement is expensive (time and energy)
- **This is a new model for algorithms and applied mathematics on these machines**
  - Algorithms avoid data movement
  - This can be done, but…
  - Significant change from tradition of counting arithmetic operations

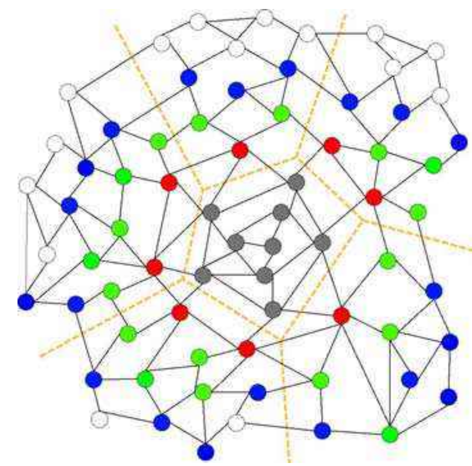

- **Exascale will enable new science problems, which will also require new algorithms**

Hardware and Software Scaling Require New Resilience Models

- **Resiliency challenges**
  - Chance of component failure grows with system
  - Failure and power management → irregular performance behavior
- **Hardware / software**
  - Component values should not cause system-wide or application-wide outages
  - Affects all software levels

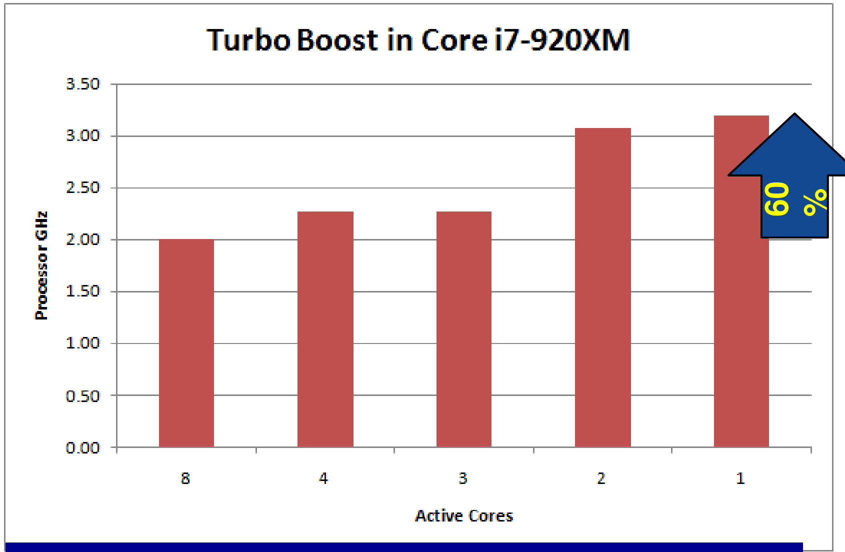

Software assumption that all processors run at the same speed:

- Clock speed may change due to temperature and power
- Failures in memory system may also affect performance

Coordinated Program in Exascale

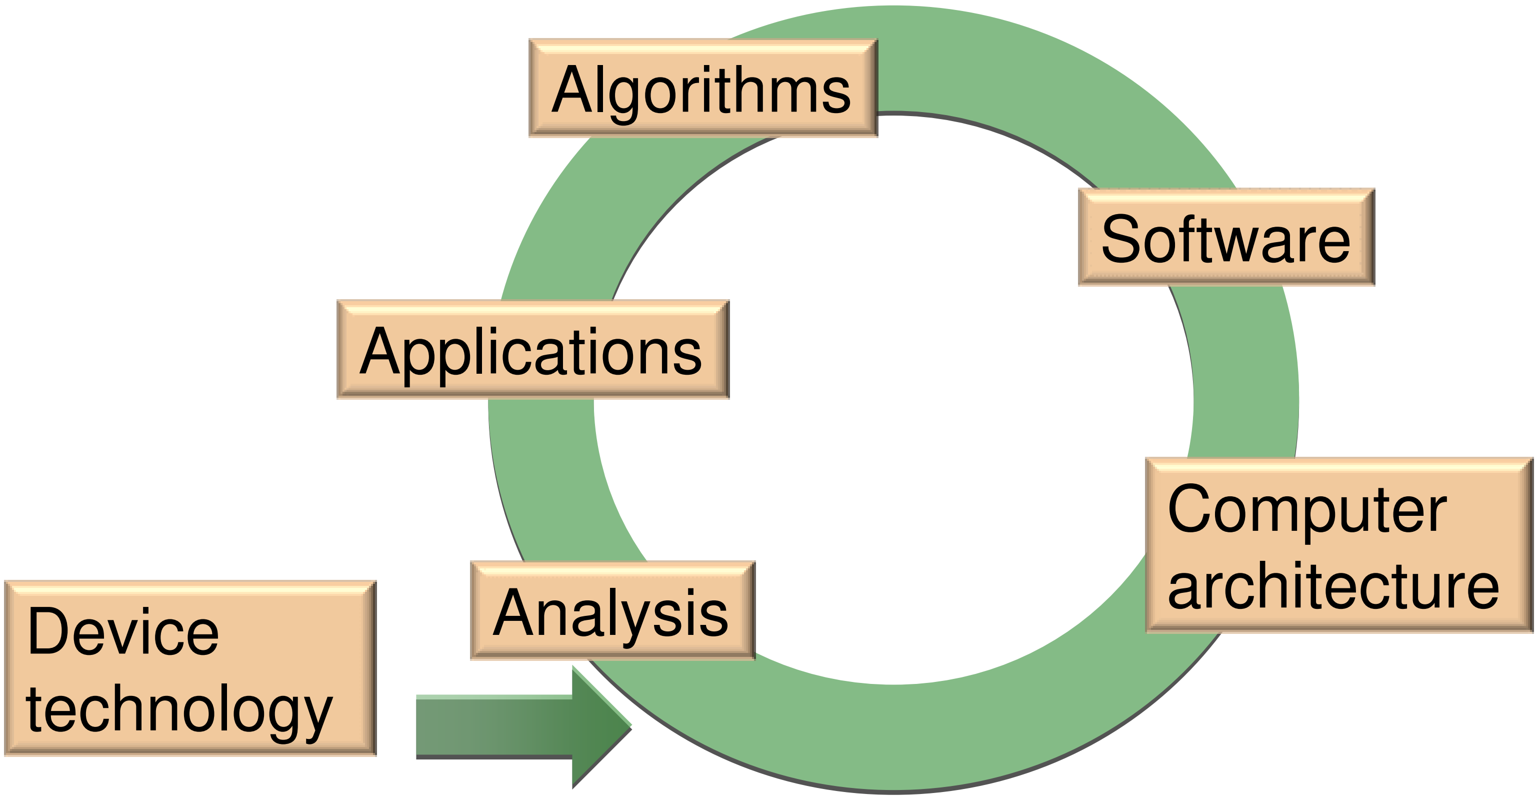

- **Need to Co-Design hardware, software, algorithms, and applications with the goal of:**
- **Effective machines; exascale-ready science**

Challenges to ~~Exascale~~ Performance Growth

1) **System power** is the primary constraint
2) **Concurrency** (1000x today)
3) **Memory** bandwidth and capacity are not keeping pace
4) **Processor** architecture is an open question
5) **Interconnect** for high bandwidth with low cost and power
6) **Software** needs to change to match architecture
7) **Algorithms** need to minimize data movement, not flops
8) **I/O and Data analytics to** keep pace with machine speed
9) **Reliability and resiliency** will be critical at this scale

Unlike the last 20 years most of these (1-8) are equally important across scales, e.g., 100 10-PF machines

Cloud Computing Doesn't Solve Science Need

Traditional clouds are not suited to science

- Up to 50x slower even for small parallel jobs
- 5-50x more expensive depending on problem
- Lack HPC networks, scheduling, high utilization; add profit

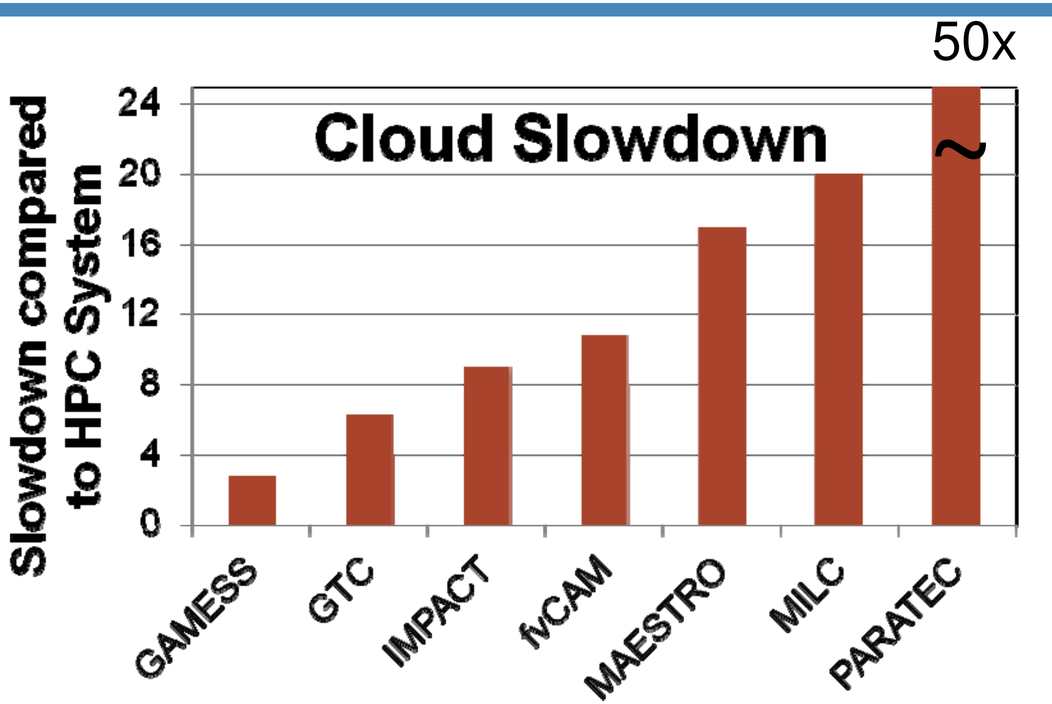

Workload differs, but energy challenges are common:

- *April 2008 – Microsoft's new Chicago data center will have a total capacity of 198 megawatts of power.*
- *January 2011 -- SuperNAP envisions a 500-megawatt Las Vegas campus with 31,000 cabinets of servers*

National Need for Exascale Program

- Computing performance is in crisis
- R&D investments are needed to improve energy efficiency along with performance
  - Processor and memory technology
  - Computer architecture, systems and software
  - Applied mathematic and algorithms
- A coordinated program is necessary to produced complete and effective systems
  - Natural hardware evolution will not work
- Challenges must be met for the computing performance needed for US competitiveness