Computational Science (aka Scientific Computing)

Shirley Moore
svmoore@utep.edu
http://www.cs.utep.edu/svmoore/
CPS 5401
August 25, 2014

Definition

- Computational science is an interdisciplinary field in which realistic mathematical models combined with scientific computing methods are used to study systems of real-world scientific or societal interest.
- The 2005 Report to the President, Computational Science: Ensuring America's Competitiveness, states that "the most scientifically important and economically promising research frontiers in the 21st century will be conquered by those most skilled with advanced computing technologies and computational science applications."

The Three Pillars of Modern Science

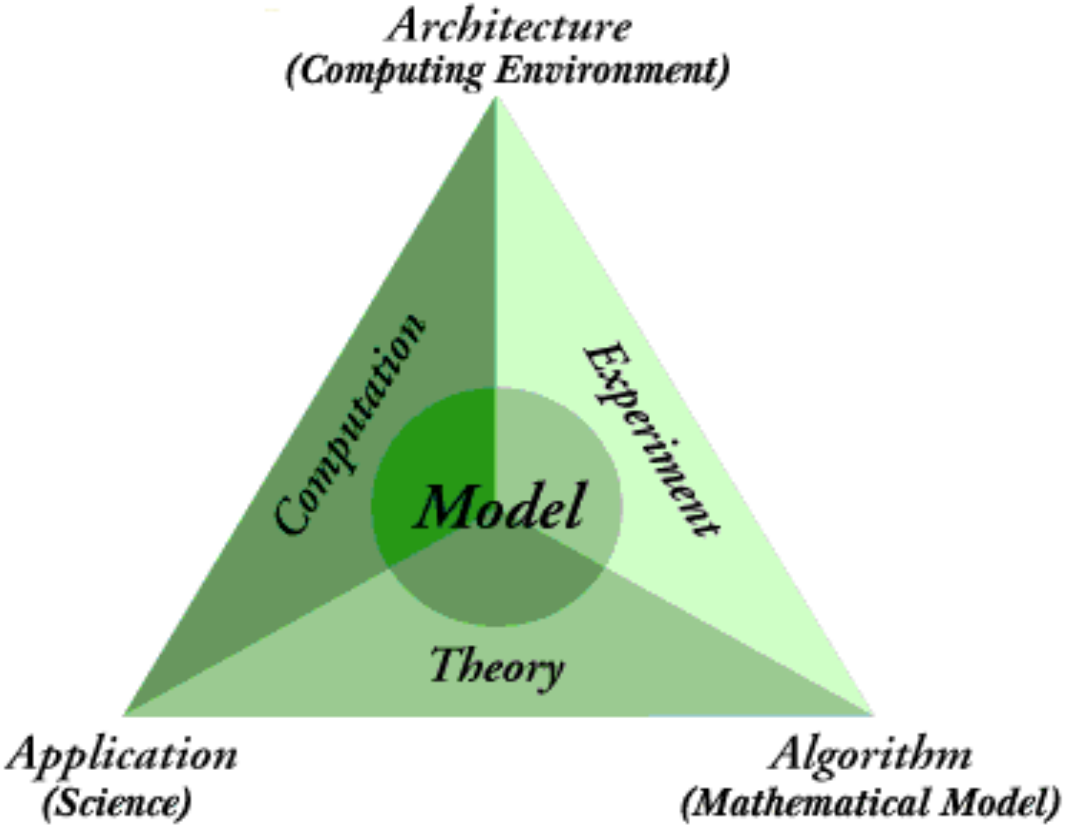

Computation has established itself as the third pillar of modern science. Computational science involves the appropriate use of a computational **architecture** to apply an **algorithm**, or method, to solve a scientific **application**, or problem. The combination of Application, Algorithm, and Architecture results in a Model.

Simulation: The Third Pillar of Science

- Traditional scientific and engineering paradigm:
  1. Do theory or paper design
  2. Perform experiments or build system
- Limitations
  - Too difficult – e.g., building large wind tunnels
  - Too expensive – e.g., build a throw-away passenger jet
  - Too slow – e.g., wait for climate change or galactic evolution
  - Too dangerous – e.g., weapons, drug design, climate experimentation
- Computational science paradigm
  3. User high performance computer systems to simulate the phenomenon
     - Base model on physical laws and efficient numerical methods

Computational Science Fuses Three Disciplines

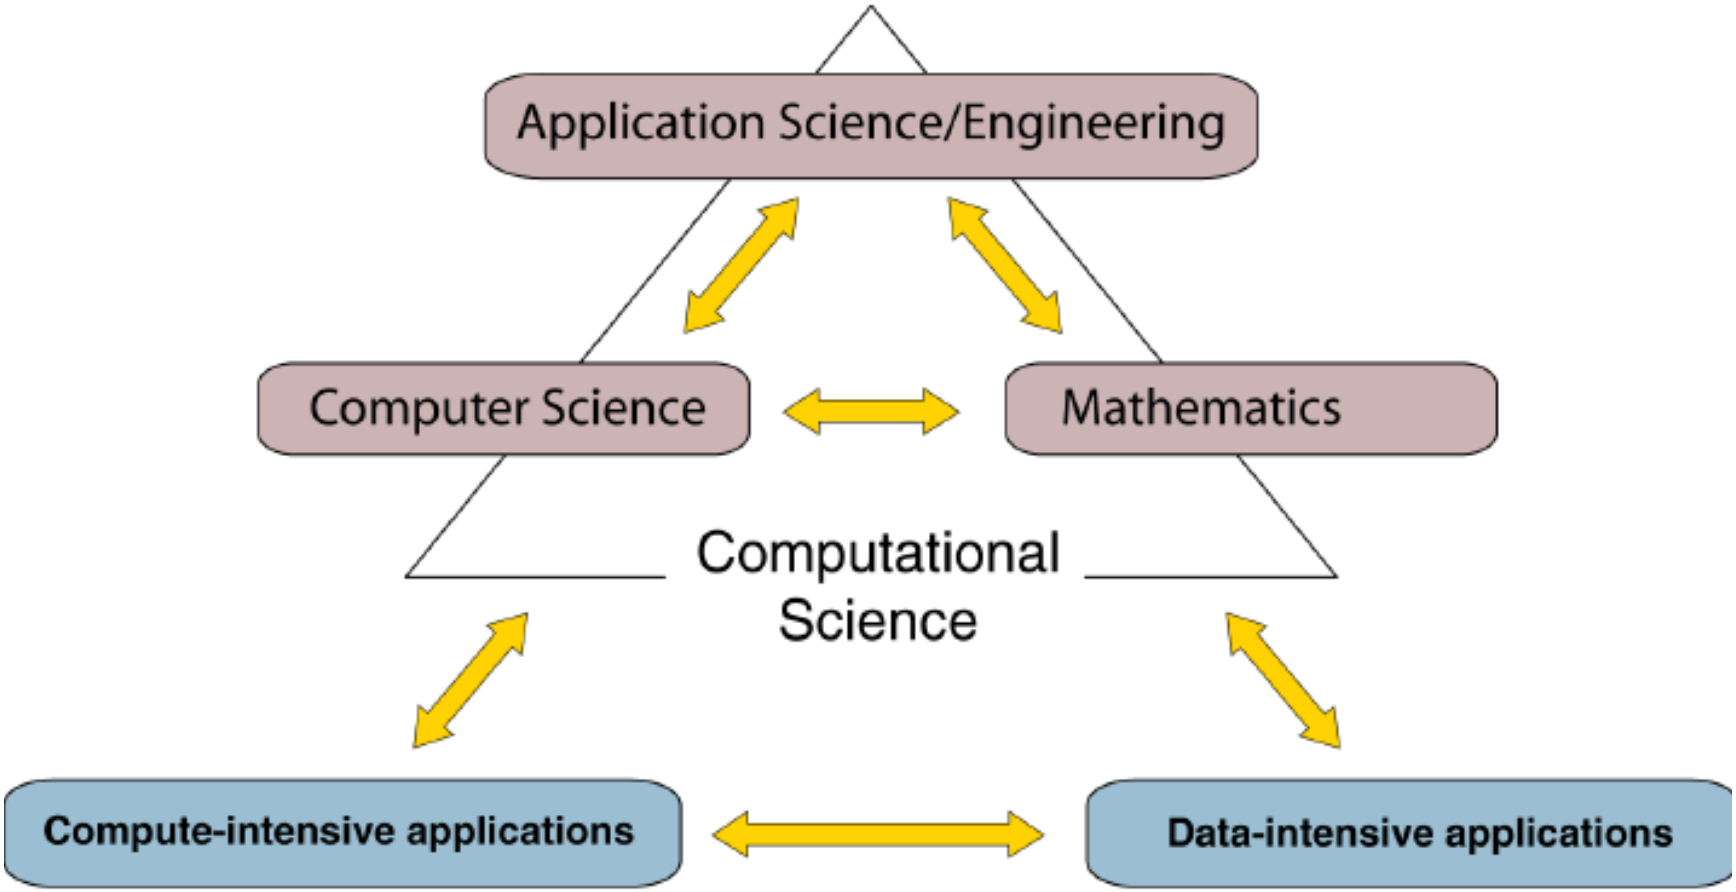

What do Computational Scientists Do?

- Combine scientific programming and mathematical skills with knowledge of application fields to computationally model systems of interest.
- Mathematical and programming skills
  - calculus, linear algebra, differential equations, statistics
  - simulation methods, parallel computing, scientific computing libraries and tools, visualization

Example 1: Climate Modeling

- Very, very simple climate model
- Community Earth System Model (CESM)
  - Huge DOE and NSF funded project
  - Project goals
    - to develop and to work continually to improve a comprehensive CESM that is at the forefront of international efforts in modeling the climate system, including the best possible component models coupled together in a balanced, harmonious modeling framework;
    - to make the model readily available to, and usable by, the climate research community, and to actively engage the community in the ongoing process of model development;
    - to use the CESM to address important scientific questions about the climate system, including global change and interdecadal and interannual variability; and
    - to use appropriate versions of the CESM for calculations in support of national and international policy decisions.

Example 2: Multiscale Quantum Simulations in Nano Science

- Goal: establish a suite of petascale simulation tools that enable transformational research at the fundamental quantum-mechanical level in nano science and technology.
- Development of such tools requires interdisciplinary, synergistic research in
  - methodology, implementation and prototype application of quantum methods,
  - profiling, performance modeling and automatic optimization of kernels.
- The tools are based on existing real-space multi-grid (RMG) method, which is now well-established and has been successfully applied to a large number of systems.
  - Biomolecular version of the code is now able to perform long-time quantum molecular dynamics simulations of solvated biomolecules.
  - Challenges are in reformulating and tuning the codes to run efficiently on petascale systems.
- Study of copper binding to the prion protein (PrP), which is responsible for a group of neurodegenerative diseases called the transmissible spongiform encephalopathies, such as Cruetzfeldt-Jakobs in humans or “mad cow” disease in cattle.
  - Implications for patients suffering from these diseases, as well as from other neurogical diseases, such as Alzheimer’s or Parkinson’s.
- Interfaced the quantum simulation results with NAMD; parts of the simulations for a biomolecular fragment are performed with the RMG code using quantum methods, and the results, including vibrational aspects, are fed to NAMD for molecular mechanics simulation of the entire protein
  - Expect to provide similar interfaces to other molecular dynamics codes
- Project Team
  - NCSU: J. Bernholc, W. Lu, C. T. Kelley, M. Hodak, E. Briggs
  - UTK: S. Moore, S. Tomov, H. Jagode

Significant Result

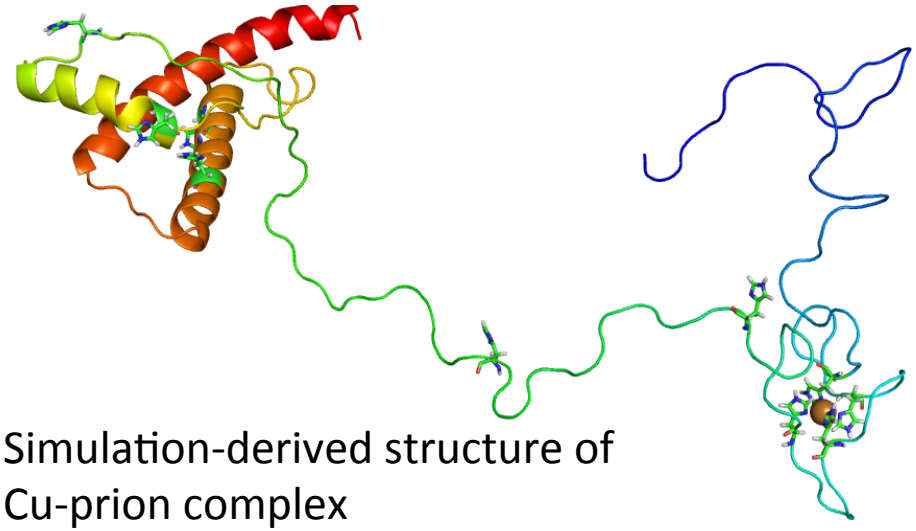

Simulation-derived structure of Cu-prion complex

Cu-prion complex is much stiffer, which lesser tendency for misfolding

Protective role of copper in prion diseases

- Study of copper binding to the prion protein (PrP), which is responsible for a group of neurodegenerative diseases called the transmissible spongiform encephalopathies, such as Cruetzfeldt-Jakobs in humans or “mad cow” disease in cattle.
- "Functional implications of multistage copper binding to the prion protein," M. Hodak, R. Chisnell, W. Lu, and J. Bernholc, Proc. Nat. Acad. Sci. 106, 11576 (2009).
- This work could have implications for patients suffering from these diseases, as well as from other neurological diseases, such as Alzheimer’s or Parkinson’s.

Supercomputers

- See a list of the world's fastest at www.top500.org
- Strategic importance of supercomputing
  - Essential for scientifc discovery
  - Critical for national security
  - Fundamental contributor to the economy and competitiveness through use in engineering and manufacturing
- Supercomputers are the tools for solving the most challenging scientific problems through large-scale simulations.

Units of Measure

- **High Performance Computing (HPC) units are:**
  - Flop: floating point operation, usually double precision unless noted
  - Flop/s: floating point operations per second
  - Bytes: size of data (a double precision floating point number is 8)
- **Typical sizes are millions, billions, trillions…**

Mega	Mflop/s = 10^6 flop/sec	Mbyte = 2^{20} = 1048576 ~ 10^6 bytes
Giga	Gflop/s = 10^9 flop/sec	Gbyte = 2^{30} ~ 10^9 bytes
Tera	Tflop/s = 10^{12} flop/sec	Tbyte = 2^{40} ~ 10^{12} bytes
Peta	Pflop/s = 10^{15} flop/sec	Pbyte = 2^{50} ~ 10^{15} bytes
Exa	Eflop/s = 10^{18} flop/sec	Ebyte = 2^{60} ~ 10^{18} bytes
Zetta	Zflop/s = 10^{21} flop/sec	Zbyte = 2^{70} ~ 10^{21} bytes
Yotta	Yflop/s = 10^{24} flop/sec	Ybyte = 2^{80} ~ 10^{24} bytes

- **Current fastest (public) machine ~ 33.9 Pflop/s**
  - Up-to-date list at www.top500.org

Moore's Law

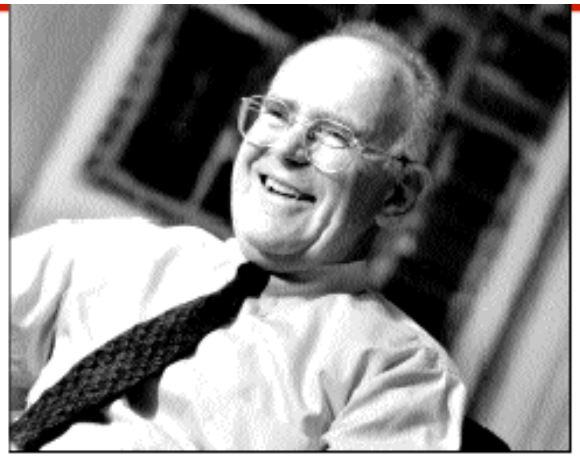

Gordon Moore (co-founder of Intel) Electronics Magazine, 1965

Number of devices/chip doubles every 18 months

2X transistors/Chip Every 1.5 years

Called "Moore's Law"

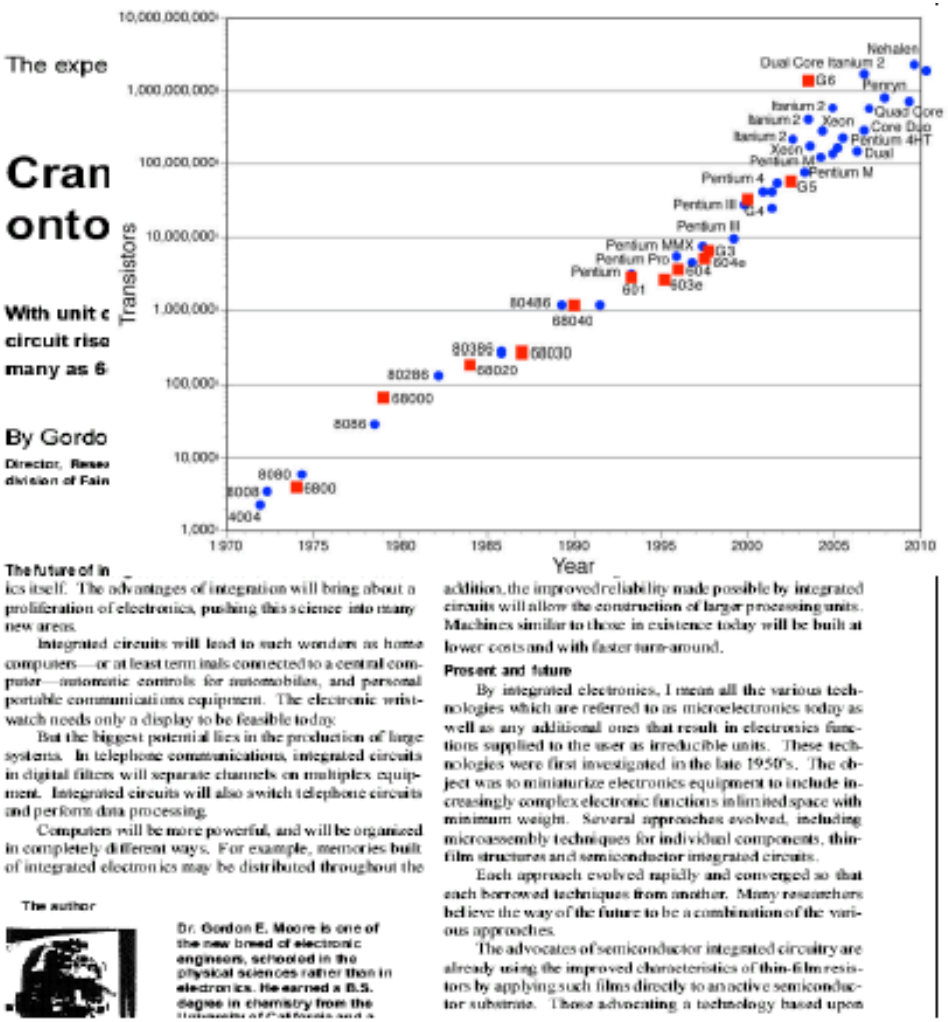

The expe

Cran onto

With unit c circuit rise many as 6

By Gordo

Director, Resea division of Fain

The future of in ics itself. The advantages of integration will bring about a proliferation of electronics, pushing this science into many new areas.

Integrated circuits will lead to such wonders as home computers—or at least terminals connected to a central computer—automatic controls for automobiles, and personal portable communications equipment. The electronic wristwatch needs only a display to be feasible today.

But the biggest potential lies in the production of large systems. In telephone communications, integrated circuits in digital filters will separate channels on multiplex equipment. Integrated circuits will also switch telephone circuits and perform data processing.

Computers will be more powerful, and will be organized in completely different ways. For example, memories built of integrated electronics may be distributed throughout the

addition, the improved reliability made possible by integrated circuits will allow the construction of larger processing units. Machines similar to those in existence today will be built at lower costs and with faster turn-around.

Present and future

By integrated electronics, I mean all the various technologies which are referred to as microelectronics today as well as any additional ones that result in electronics functions supplied to the user as irreducible units. These technologies were first investigated in the late 1950's. The object was to miniaturize electronics equipment to include increasingly complex electronic functions in limited space with minimum weight. Several approaches evolved, including microassembly techniques for individual components, thin-film structures and semiconductor integrated circuits.

Each approach evolved rapidly and converged so that each borrowed techniques from another. Many researchers believe the way of the future to be a combination of the various approaches.

The advocates of semiconductor integrated circuitry are already using the improved characteristics of thin-film resistors by applying such films directly to active semiconductor substrates. Those advocating a technology based upon

The author

Dr. Gordon E. Moore is one of the new breed of electronic engineers, schooled in the physical sciences rather than in electronics. He earned a B.S. degree in chemistry from the

Today's Multicores

99% of Top500 Systems Are Based on Multicore

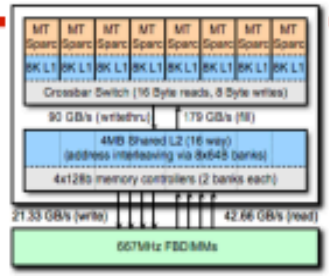

Sun Niagra2 (8 cores)

IBM Power 7 (8 cores)

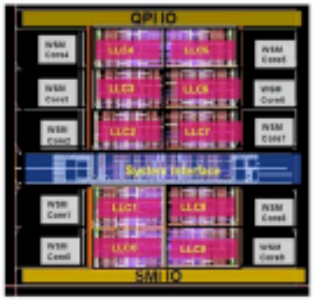

Intel Westmere (10 cores)

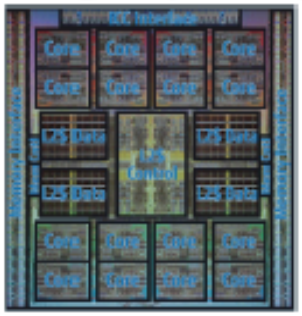

Fujitsu Venus (16 cores)

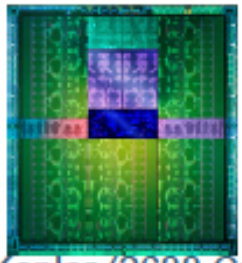

Nvidia Kepler (2688 Cuda cores)

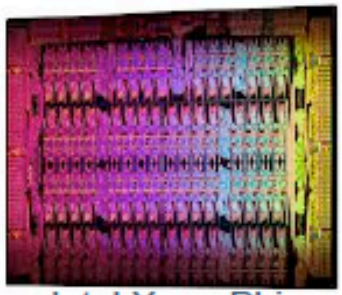

Intel Xeon Phi (60 cores)

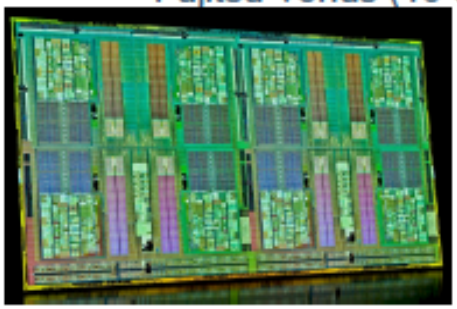

AMD Interlagos (16 cores)

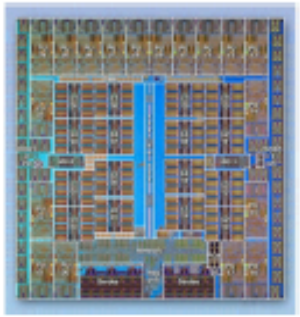

IBM BG/Q (18 cores)

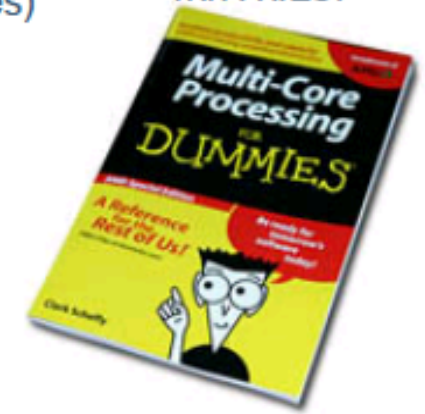

But Clock Frequency Scaling Replaced by Scaling Cores / Chip

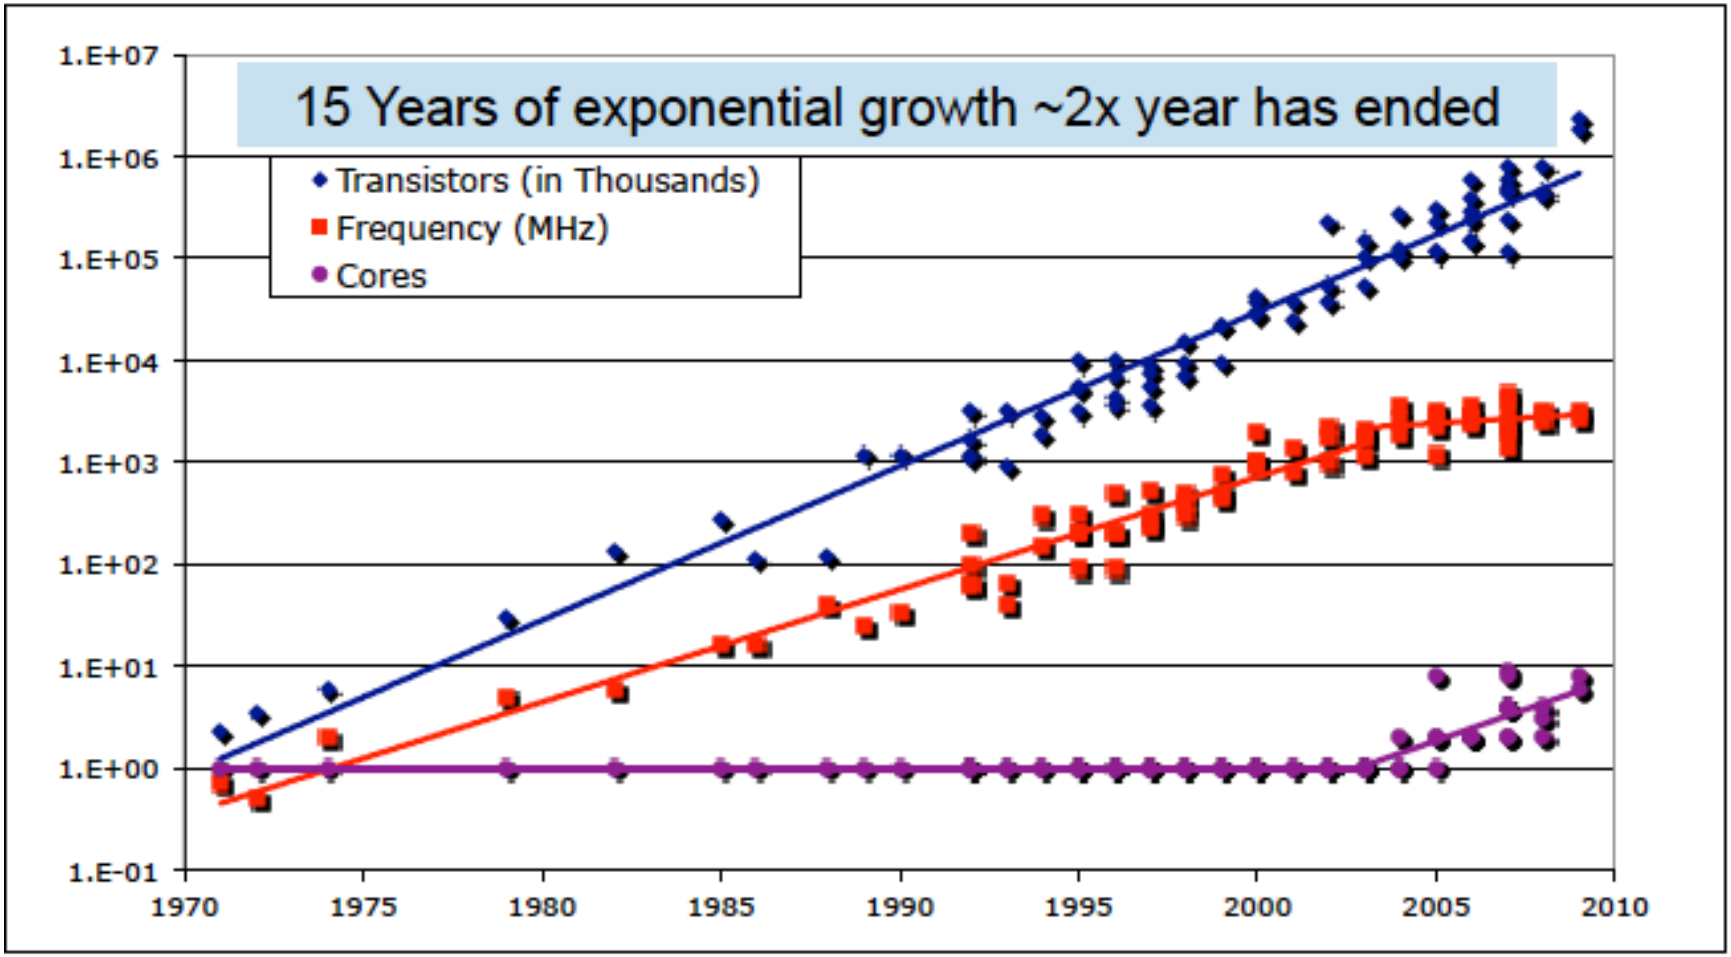

Data from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanoviç

Slide from Kathy Yelick

Performance Has Also Slowed, Along with Power

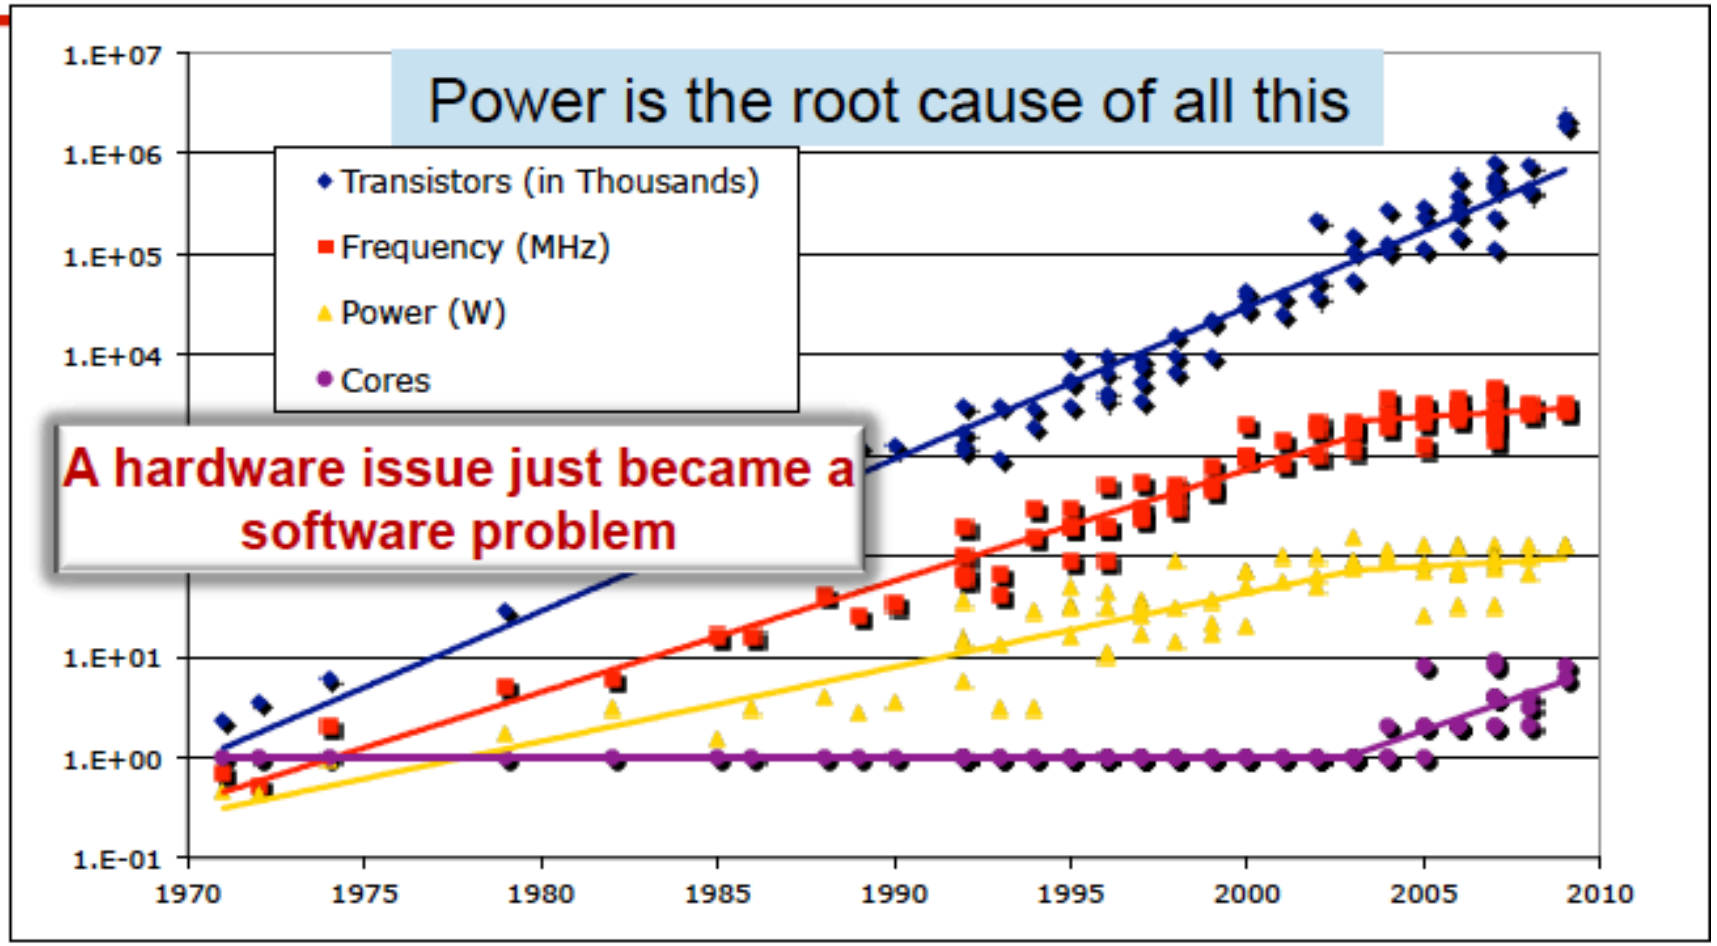

Data from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanoviç

Slide from Kathy Yelick

Performance Development of HPC Over the Last 20 Years

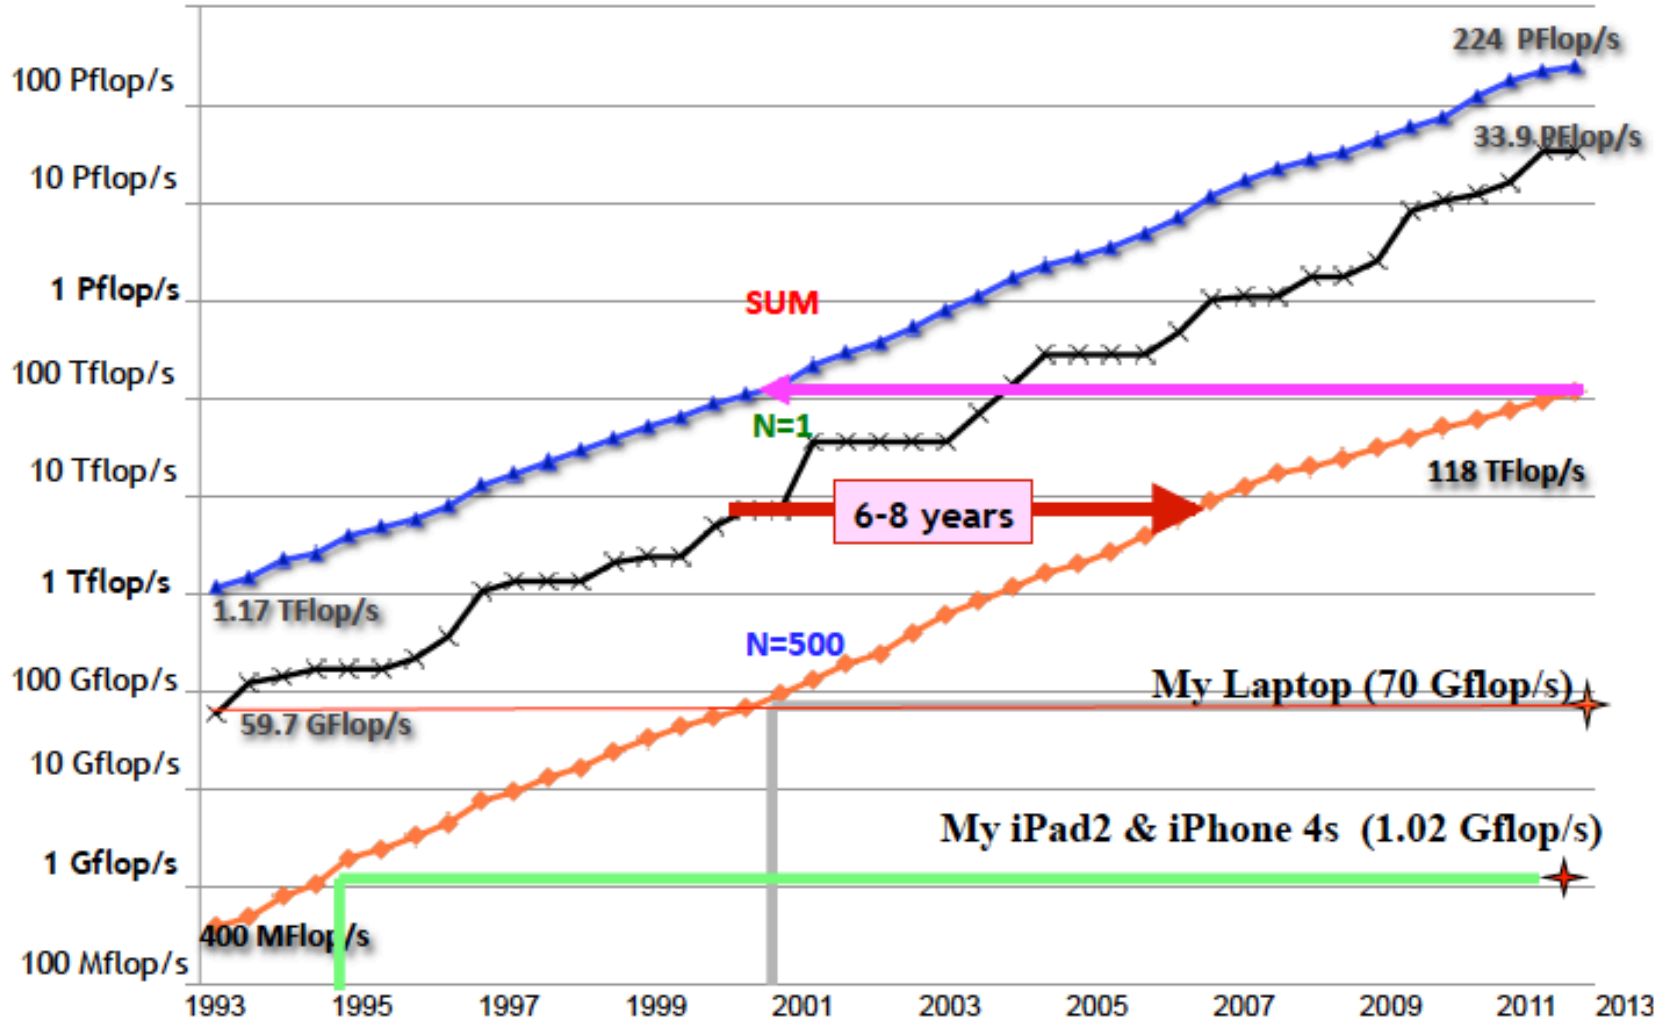

Slide courtesy of Jack Dongarra

Petaflops Club

31 Systems

13 4 3 3 3 3 1 1

Name	Rmax Linpack# Pflops	Country	
Tianhe-2 (MilkyWay-2)	33.9	China	NUDT: Hybrid Intel/Intel/Custom
Titan	17.6	US	Cray: Hybrid AMD/Nvidia/Custom
Sequoia	17.2	US	IBM: BG-Q/Custom
K Computer	10.5	Japan	Fujitsu: Sparc/Custom
Mira	8.59	US	IBM: BG-Q/Custom
Piz Daint	6.27	Switzerland	Cray: Hybrid AMD/Nvidia/Custom
Stampede	5.17	US	Dell: Hybrid/Intel/Intel/IB
JUQUEEN	5.01	Germany	IBM: BG-Q/Custom
Vulcan	4.29	US	IBM: BG-Q/Custom
SuperMUC	2.9	Germany	IBM: Intel/IB
TSUBAME 2.5	2.84	Japan	Cluster Pltf: Hybrid Intel/Nvidia/IB
Tianhe-1A	2.57	China	NUDT: Hybrid Intel/Nvidia/Custom
cascade	2.35	US	Atipa: Hybrid Intel/Intel/IB
Pangea	2.1	France	Bull: Intel/IB
Fermi	1.79	Italy	IBM: BG-Q/Custom
Pleiades	1.54	US	SGI Intel/IB
DARPA Trial Subset	1.52	US	IBM: Intel/IB
Spirit	1.42	US	SGI: Intel/IB
ARCHER	1.37	UK	Cray: Intel/Custom
Curie thin nodes	1.36	France	Bull: Intel/IB
Nebulae	1.27	China	Dawning: Hybrid Intel/Nvidia/IB
Yellowstone	1.26	US	IBM: BG-Q/Custom
Blue Joule	1.25	UK	IBM: BG-Q/Custom
Helios	1.24	Japan	Bull: Intel/IB
Garnet	1.17	US	Cray: AMD/Custom
Cielo	1.11	US	Cray: AMD/Custom
DiRAC	1.07	UK	IBM: BG-Q/Custom
Hopper	1.05	US	Cray: AMD/Custom
Tera-100	1.05	France	Bull: Intel/IB
Oakleaf-FX	1.04	Japan	Fujitsu: Sparc/Custom
MPI	1.03	Germany	iDataFlex: Intel/IB

8 Hybrid Architectu
8 IBM BG/Q
18 Custom X
12 Infiniband X
9 Look like “cluste

Tianhe-2 (Milkyway-2)

China, 2013: *the 34 PetaFLOPS*

Developed in cooperation between NUDT and Inspur for National Supercomputer Center in Guangzhou

Peak performance of 54.9 PFLOPS

- 16,000 nodes contain 32,000 Xeon Ivy Bridge processors and 48,000 Xeon Phi accelerators totaling 3,120,000 cores
- 162 cabinets in 720m^2 footprint
- Total 1.404 PB memory (88GB per node)
- Each Xeon Phi board utilizes 57 cores for aggregate 1.003 TFLOPS at 1.1GHz clock
- Proprietary TH Express-2 interconnect (fat tree with thirteen 576-port switches)
- 12.4 PB parallel storage system
- 17.6MW power consumption under load; 24MW including (water) cooling
- 4096 SPARC V9 based Galaxy FT-1500 processors in front-end system

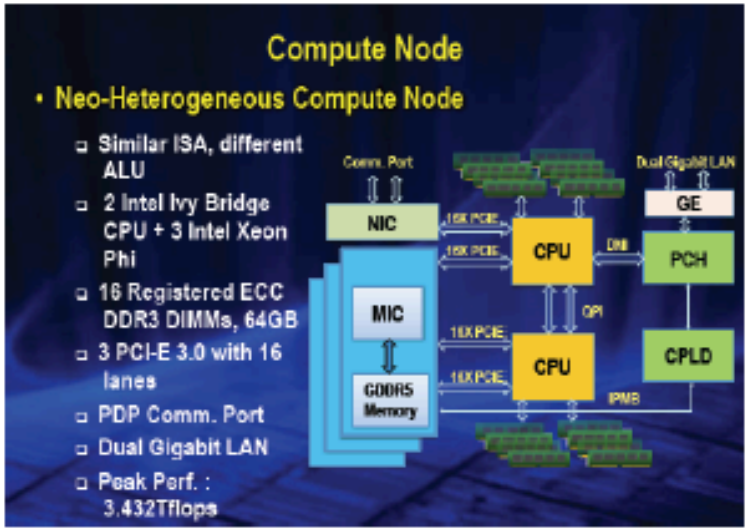

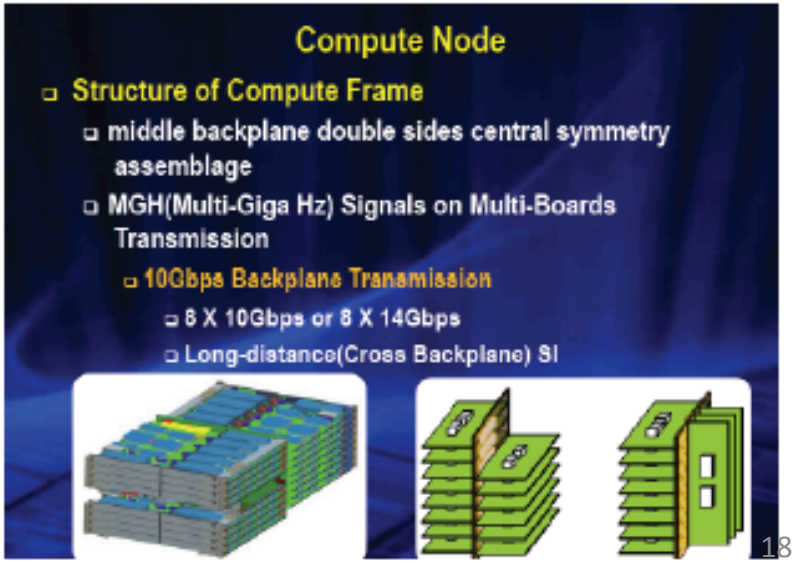

ORNL's "Titan" Hybrid System: Cray XK7 with AMD Opteron and NVIDIA Tesla processors

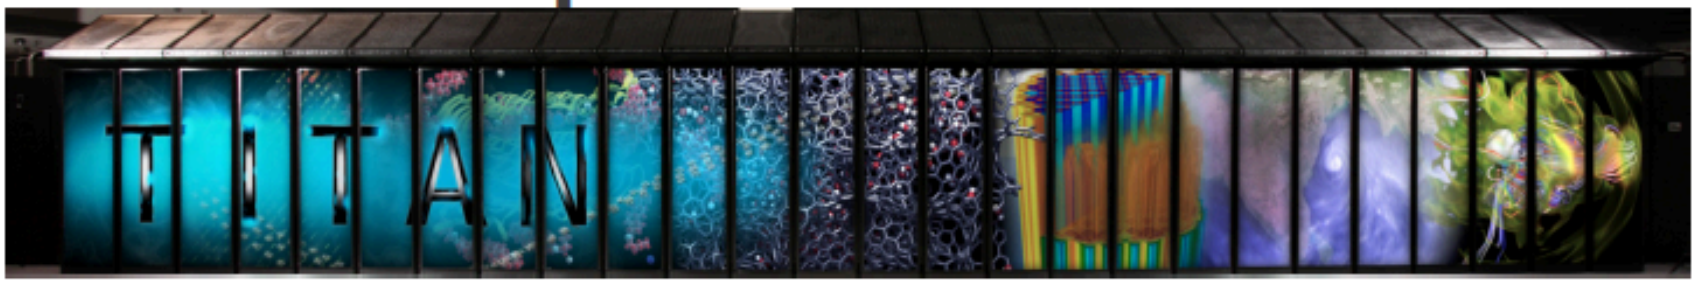

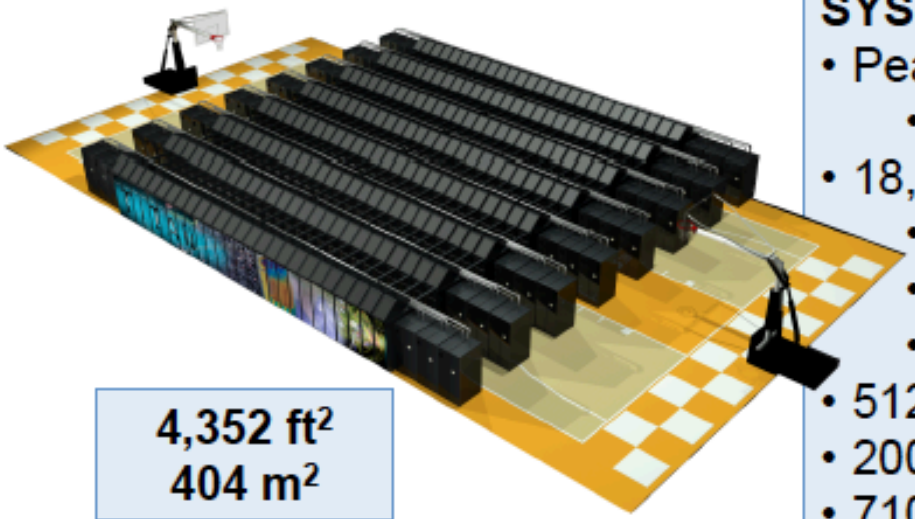

4,352 ft²
404 m²

SYSTEM SPECIFICATIONS:

- Peak performance of 27 PF
  - 24.5 Pflop/s GPU + 2.6 Pflop/s AMD
- 18,688 Compute Nodes each with:
  - 16-Core AMD Opteron CPU
  - NVIDIA Tesla "K20x" GPU
  - 32 + 6 GB memory
- 512 Service and I/O nodes
- 200 Cabinets
- 710 TB total system memory
- Cray Gemini 3D Torus Interconnect
- 9 MW peak power

OAK RIDGE

Industrial Use of Supercomputers

- **Of the 500 Fastest Supercomputer**
  - **Worldwide, Industrial Use is ~ 50%**

- Aerospace
- Automotive
- Biology
- CFD
- Database
- Defense
- Digital Content Creation
- Digital Media
- Electronics
- Energy
- Environment
- Finance
- Gaming
- Geophysics
- Image Proc./Rendering
- Information Processing Service
- Information Service
- Life Science
- Media
- Medicine
- Pharmaceutics
- Research
- Retail
- Semiconductor
- Telecomm
- Weather and Climate Research
- Weather Forecasting

















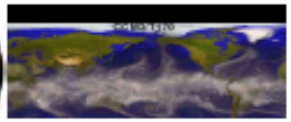

Hierarchical Heterogeneous Architectures

Shared memory programming between processes on a board and a combination of shared memory and distributed memory programming between nodes and cabinets

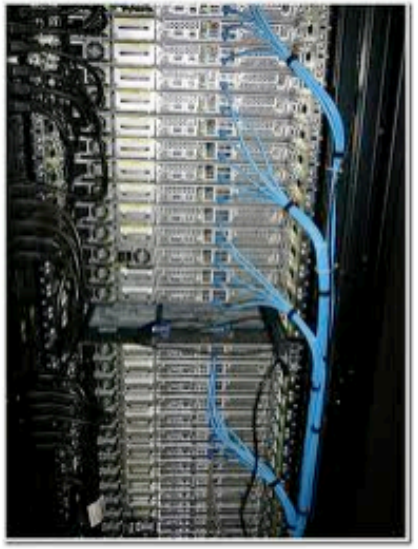

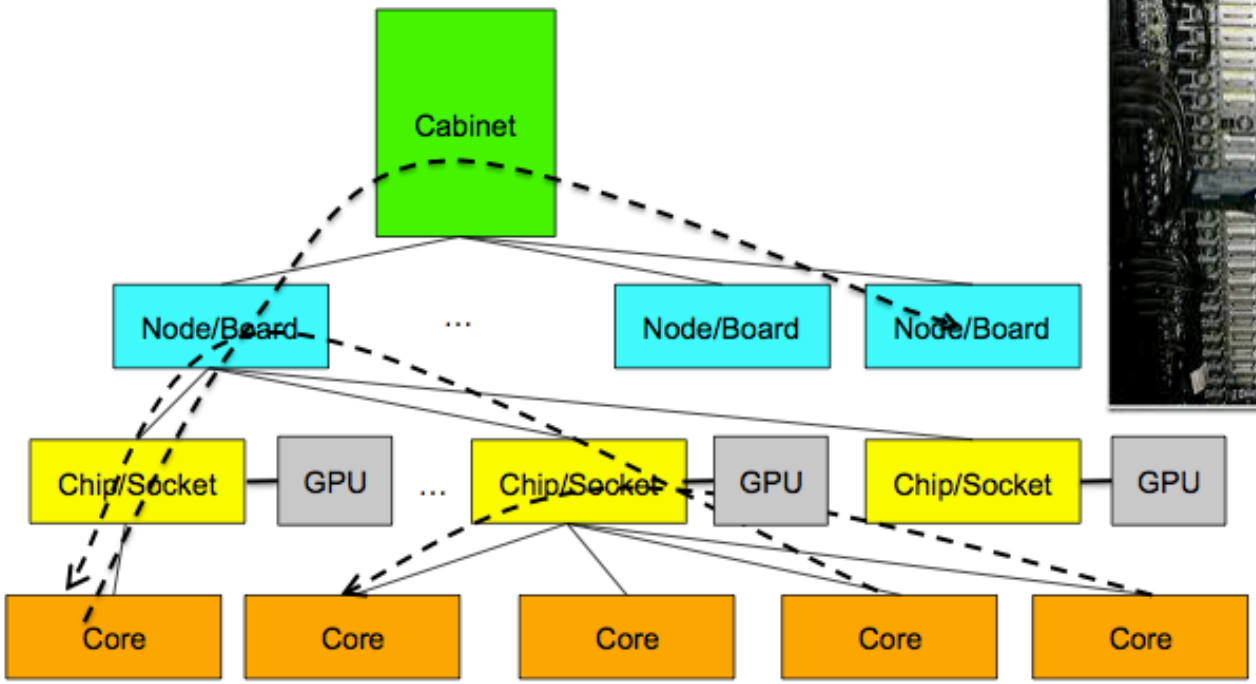

Moore's Law Reinterpreted

- Number of cores per chip doubles every 2 year, while clock speed decreases (not increases).
  - Need to deal with systems with millions of concurrent threads
    - Future generation will have billions of threads!
  - Need to be able to easily replace inter-chip parallelism with intro-chip parallelism
- Number of threads of execution doubles every 2 year

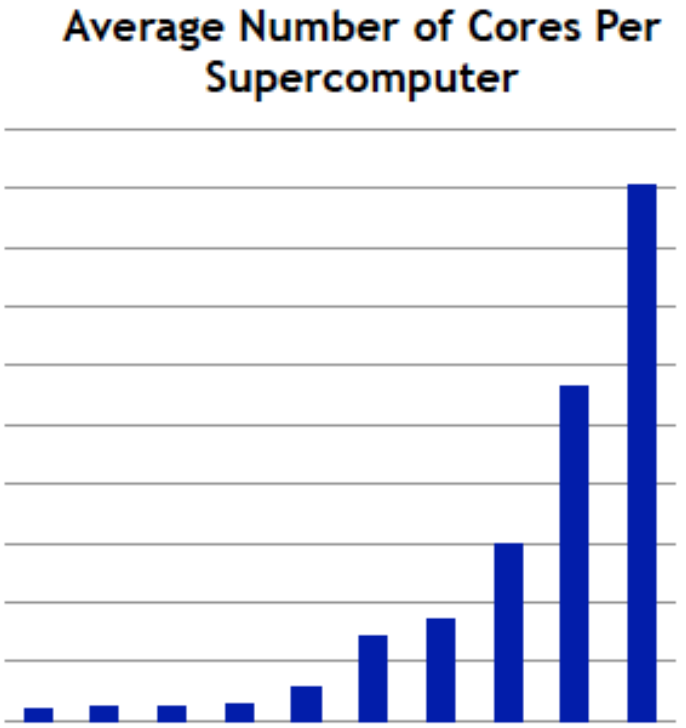

Power Efficiency

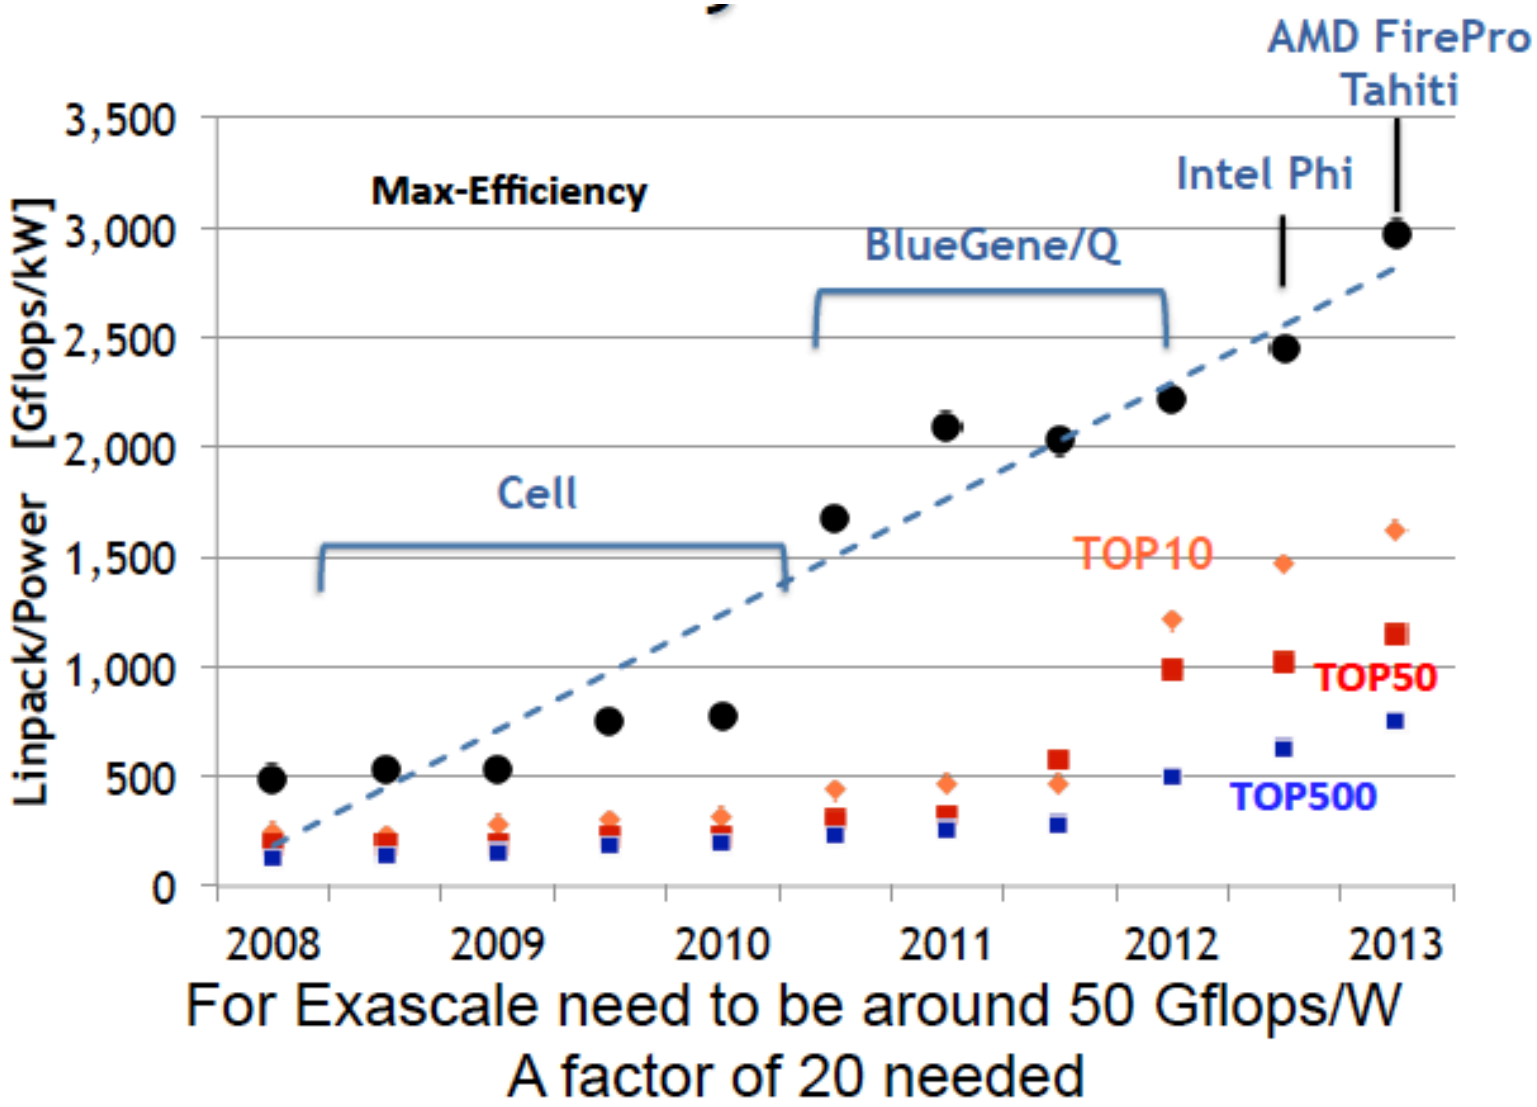

High Cost of Data Movement

- Flop/s or percentage of peak flop/s become much less relevant

Approximate power costs (in picoJoules)

	2011
DP FMADD flop	100 pJ
DP DRAM read	4800 pJ
Local Interconnect	7500 pJ
Cross System	9000 pJ

Source: John Shalf, LBNL

- Algorithms & Software: minimize data movement; perform more work per unit data movement.

Conventional Wisdom is Changing

Old Conventional Wisdom

- Peak clock frequency as primary limiter for performance improvement
- Cost: FLOPs are biggest cost for system: optimize for compute
- Concurrency: Modest growth of parallelism by adding nodes
- Memory scaling: maintain byte per flop capacity and bandwidth
- Uniformity: Assume uniform system performance
- Reliability: It's the hardware's problem

New Conventional Wisdom

- Power is primary design constraint for future HPC system design
- Cost: Data movement dominates optimize to minimize data movement
- Concurrency: Exponential growth of parallelism within chips
- Memory Scaling: Compute growing 2x faster than capacity or bandwidth
- Heterogeneity: Architectural and performance non-uniformity increase
- Reliability: Cannot count on hardware protection alone

Goals of This Course

- Acquire a broad range of skills and knowledge needed to use computationally intensive methods
- Develop an awareness of the resources available to keep abreast of the rapidly changing field of computational science

Major Course Topics

- Basic skills in Unix/Linux
- Scientific programming languages
- Parallel computing architectures
- Parallel programming paradigms
- Performance optimization
- Scalability
- Dense and sparse linear algebra libraries
- Libraries for other numerical methods
  - Eigenvalue problems
  - Fast Fourier transform
  - Numerical optimization
  - Partial differential equations

Course Info

- Website: http://svmoore.pbworks.com/
  - Syllabus
  - Lecture material
  - Labs
- Instructor: Shirley Moore
  - Office: CCSB 3.0422
  - Office hours 3-4pm MW, others by appointment
- TA: Henry Moncada
  - Office: CCSB 3.1202H
  - Office hours: TBD
- Course format
  - One hour lecture/theory
  - One hour lab/hands-on practice (starts Wed., Aug 27)
- Class preparation
  - Read material prior to class

The Future of High Performance Computing

- Lecture by Kathy Yelick, UCBerkeley/NERSC
  - https://www.youtube.com/watch?v=-O-C_d5VDko
  - minutes 5:30-12:30
- Will what you are learning in this course become obsolete – if so, how soon?
- Is the Linpack benchmark used to determine the TOP500 obsolete?