Petascale (and Beyond) Computing

Thom H. Dunning, Jr.

National Center for Supercomputing Applications

Institute for Advanced Computing Applications and Technologies

Department of Chemistry

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

Outline of Presentation

- **High-end Computing @ Illinois**
  - Blue Waters Petascale Computing Project
  - Center for Extreme-scale Computation in Science and Engineering
- **Fundamental Petascale Computing Technologies**
  - Directions in computing technologies
  - Petascale computing systems
  - Blue Waters Petascale Computing System
- **Moving to Exascale Computing**
  - Challenges in exascale computing

Science @ Petascale

Petascale computing will enable advances in a broad range of science and engineering disciplines:

Molecular Science

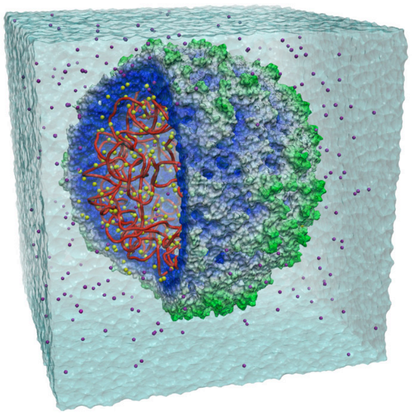
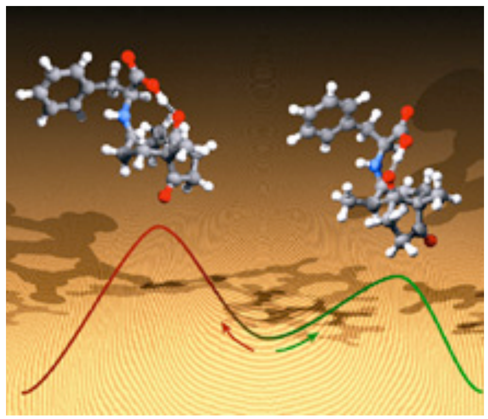

Weather & Climate Forecasting

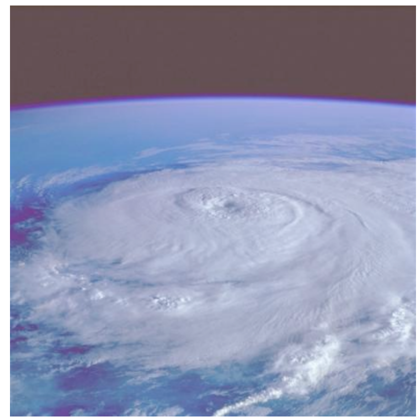
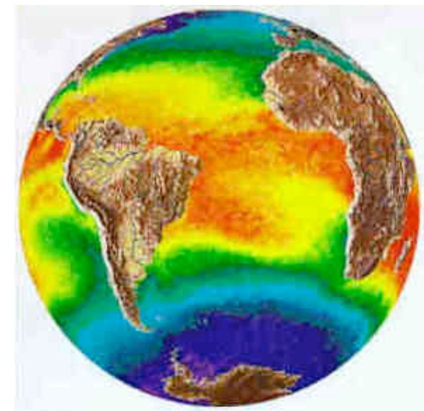

Astronomy

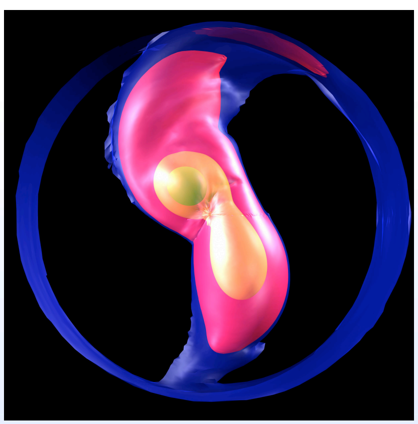

Earth Science

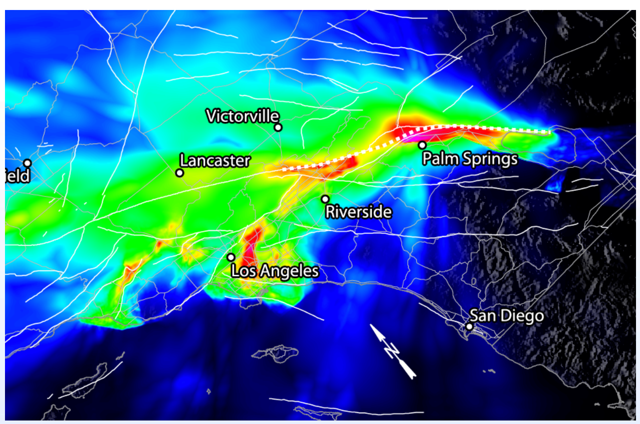

Health

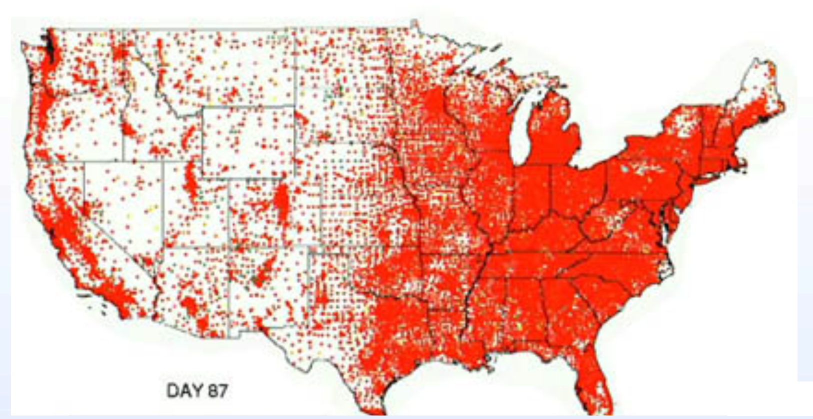

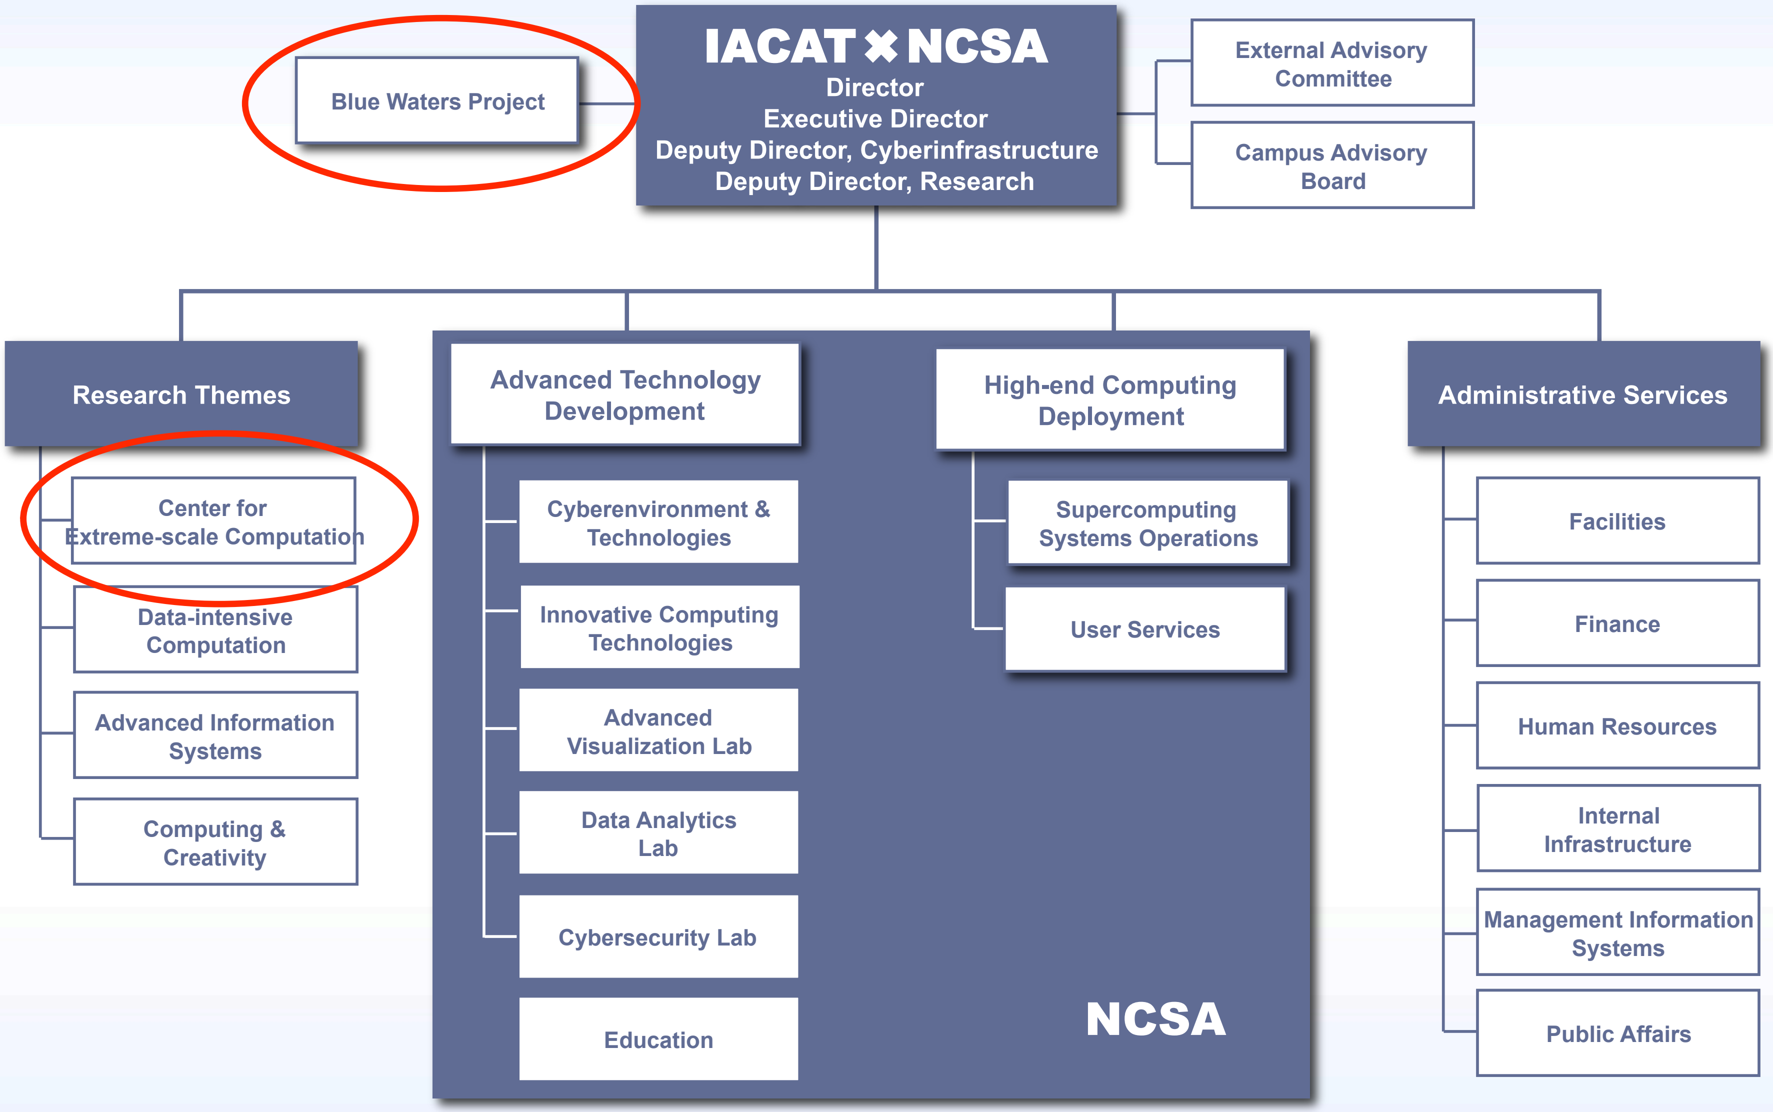

IACAT ✖ NCSA
Director
Executive Director
Deputy Director, Cyberinfrastructure
Deputy Director, Research
Blue Waters Project
External Advisory Committee
Campus Advisory Board
Research Themes
Center for Extreme-scale Computation
Data-intensive Computation
Advanced Information Systems
Computing & Creativity
Advanced Technology Development
Cyberenvironment & Technologies
Innovative Computing Technologies
Advanced Visualization Lab
Data Analytics Lab
Cybersecurity Lab
Education
High-end Computing Deployment
Supercomputing Systems Operations
User Services
NCSA
Administrative Services
Facilities
Finance
Human Resources
Internal Infrastructure
Management Information Systems
Public Affairs

Fundament Petascale Computing Technologies

Switch to Multicore Chips

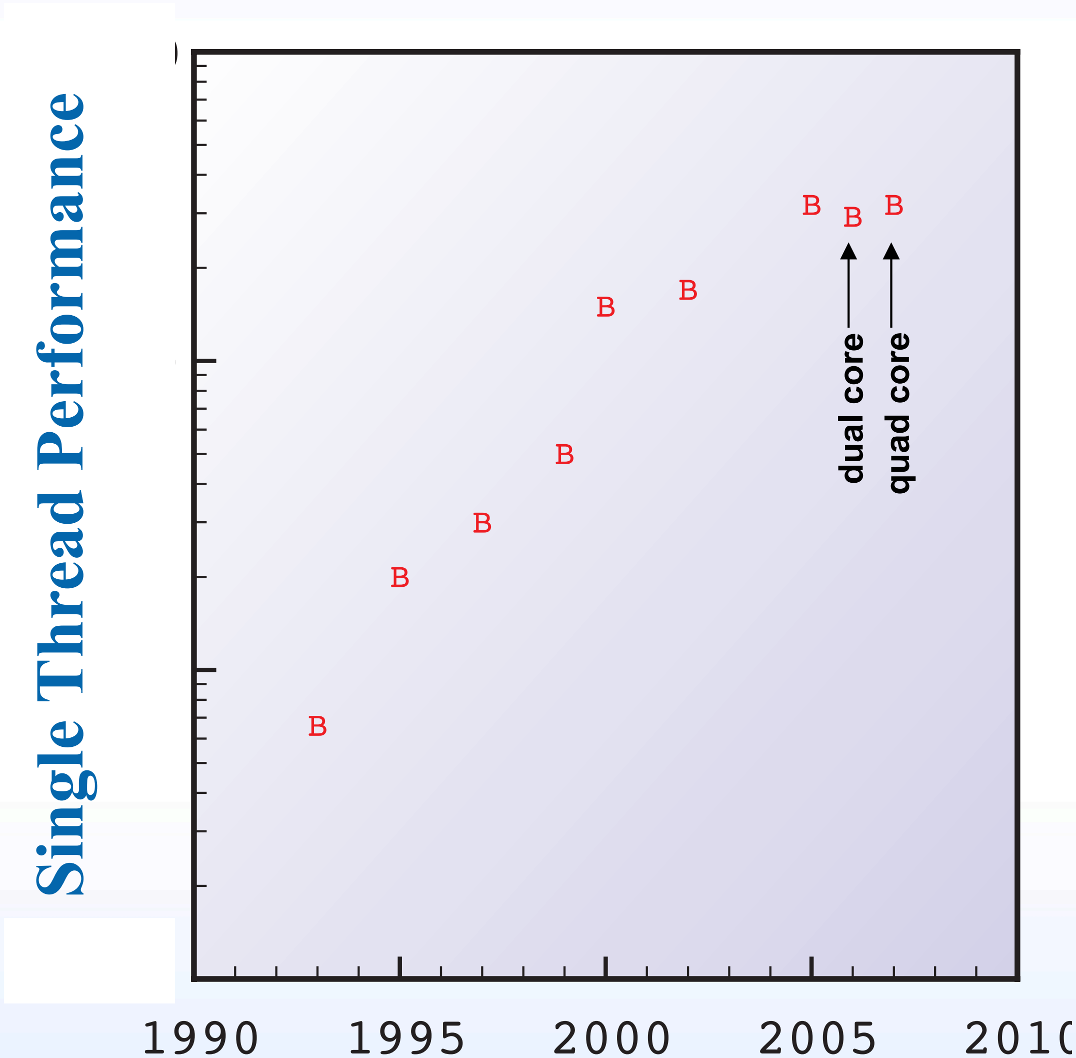

Increases in single thread performance has stalled. For the next several years the only way to obtain significant increases in performance will be through increasing use of parallelism:

- 8× in 2009
- 16× in 2011
- 32× in 2013
- *etc.*

On to Many-core Chips

NVIDIA Telsa
(240 cores)

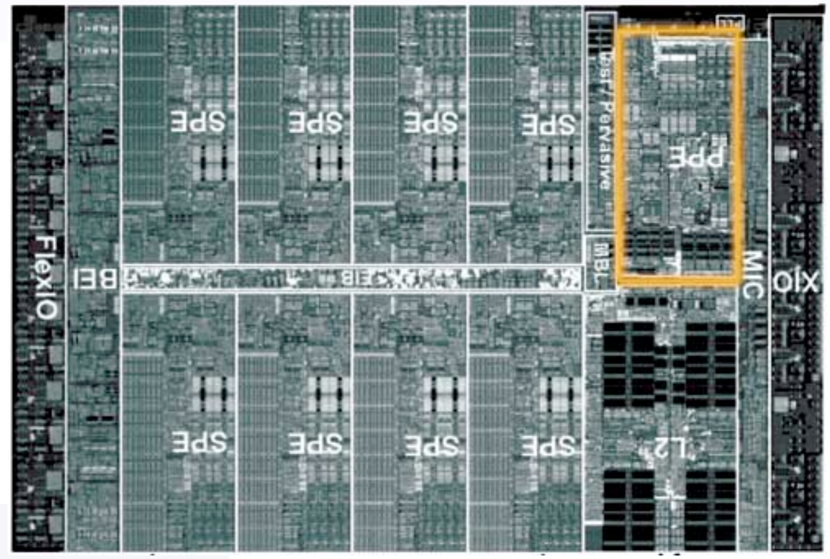

IBM Cell
(1+8 cores)

AMD Firestream
(800 cores)

Benchmarks: Direct SCF Calculations*

	Time/Iter (s)		Energy		
Molecule	GPU	CPU**	GPU	CPU**	Speedup
Caffeine	0.16	4.1	-1605.91827	-1605.91825	25
Cholesterol	1.36	67.4	-3898.82189	-3898.82189	50
Buckyball	7.32	279.4	-10709.0757	-10709.0839	40
Taxol	4.91	269.2	-12560.6830	-12560.6828	55
Valinomycin	8.44	691.2	-20351.9813	-20351.9904	80

Differences due to use of 32-bit precision, will be eliminated in 64-bit version of INVIDIA chip

* *GeForce 8800 @ 346 GF (SP), I. Ufimtsev and T. Martinez, CiSE* **10**, 26-34 (2008).

** *Using GAMESS on AMD Opteron 175 CPU.*

NVIDIA: Tesla S1070

- 4 Tesla T10s
- Frequency: 1.44 GHz
- 960 cores (240/T10)
- Performance
  - SP: 4.14 TF
  - DP: 0.34 TF
- 16 GB memory (4/T10)
- 408 GB/s memory bandwidth (104/T10)
- CUDA programming environment

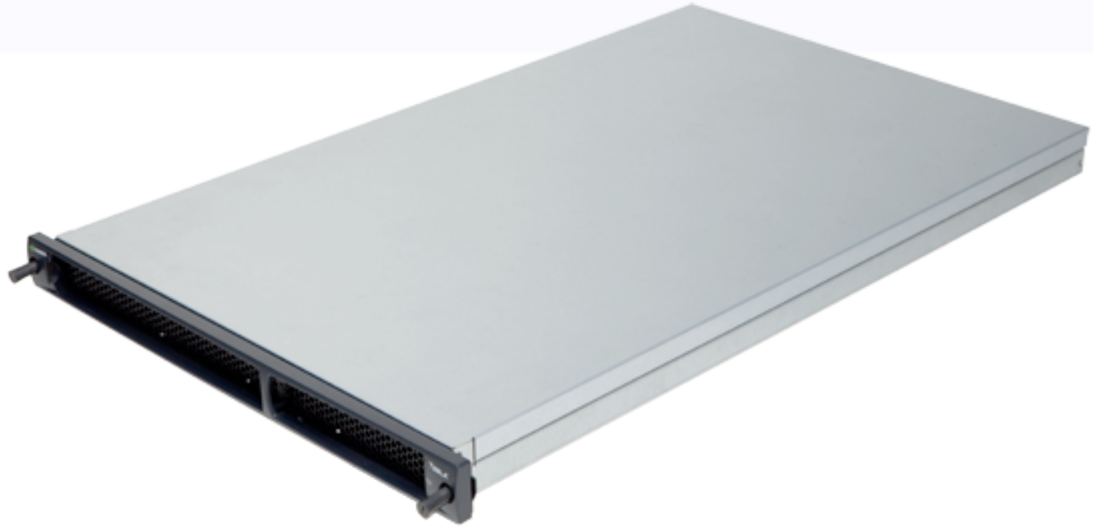

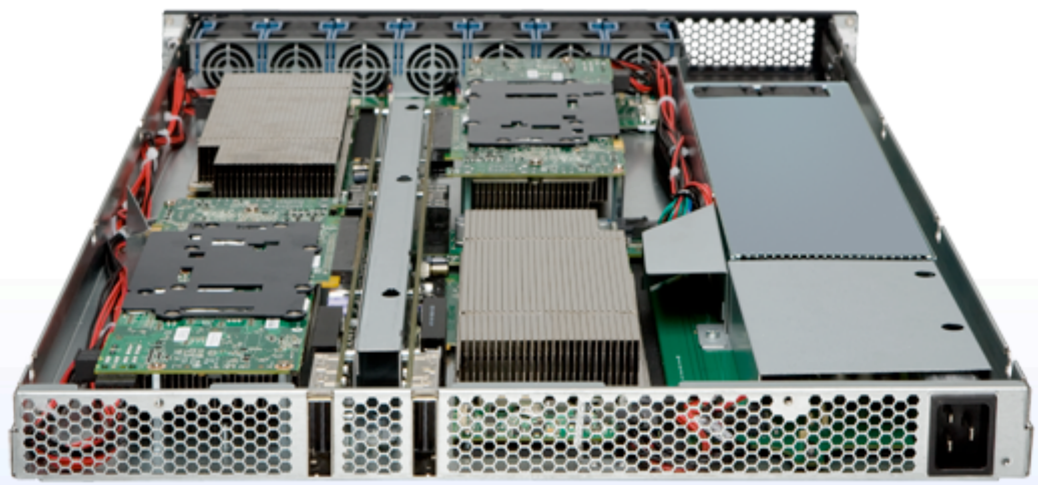

Era of Petascale Computing

LANL Roadrunner Computer System

- **Computing resources**
  - 12,960 IBM PowerXCell 8i accelerators (116,640 cores)
  - 6,480 AMD dual-core Opterons (12,960 cores)
  - 1.46 PF peak
  - 1.1 Petaflop/s Linpack
- **Memory**
  - 52 TB (accelerators)
  - 104 TB total
- **Electrical power**
  - 3.9 MW (maximum)
  - ≥ 250 Megaflops/Watt
- **Floor space**
  - 296 racks, 6800 ft²

IBM Roadrunner Petascale System

ORNL Jaguar Computer System

- **Computing resources**
  - 37,544 AMD quad-core Opterons
  - 150,176 cores
  - 1.38 PF peak
  - 1.06 Petaflop/s Linpack
- **Memory**
  - 300 TB
- **I/O Storage and Bandwidth**
  - 10 PB
  - 240 GB/s
- **Interconnect Bandwidth**
  - 374 TB/s
- **Floor space**
  - 4400 ft^2

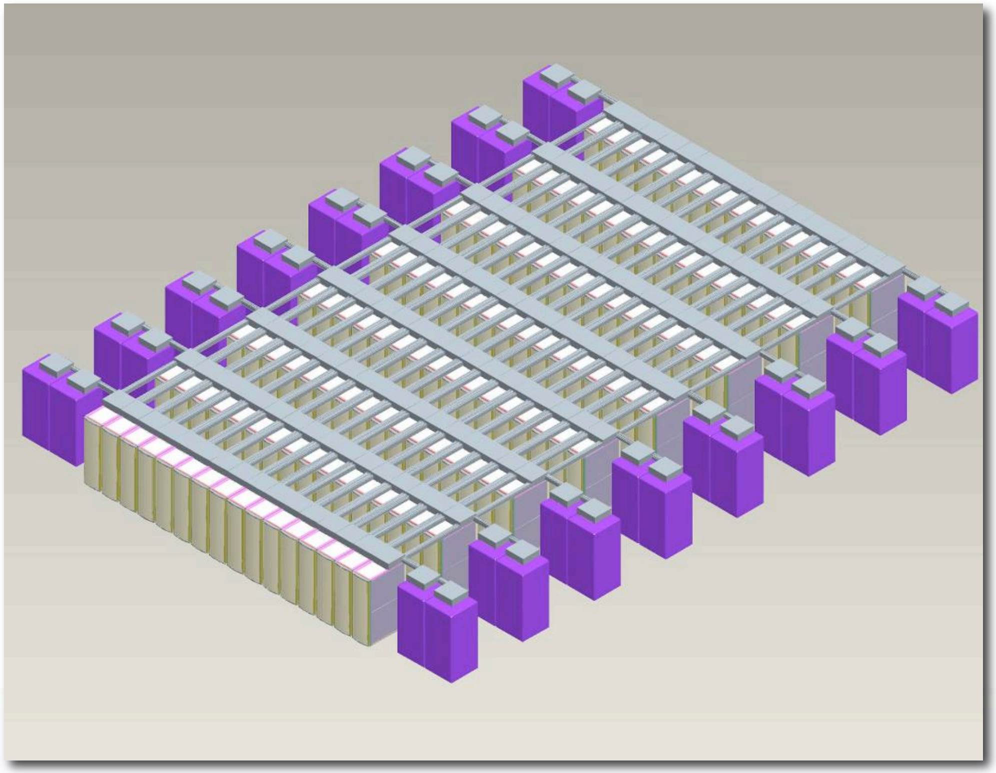

Cray Jaguar (XT5) Petascale System

Blue Waters Computing System

System Attribute	Abe	Blue Waters*
Vendor	Dell	IBM
Processor	Intel Xeon 5300	IBM
Power7		
Peak Performance (PF)	0.090	
Sustained Performance (PF)	0.005	~1
Number of Cores/Chip	4	–
Number of Processor Cores	9,600	>200,000
Amount of Memory (TB)	14.4	>800
Amount of Disk Storage (PB)	0.1	>10
Amount of Archival Storage (PB)	5	>500
External Bandwidth (Gbps)	40	100-400

* *Reference petascale computing system (no accelerators).*

Petascale Computing Facility

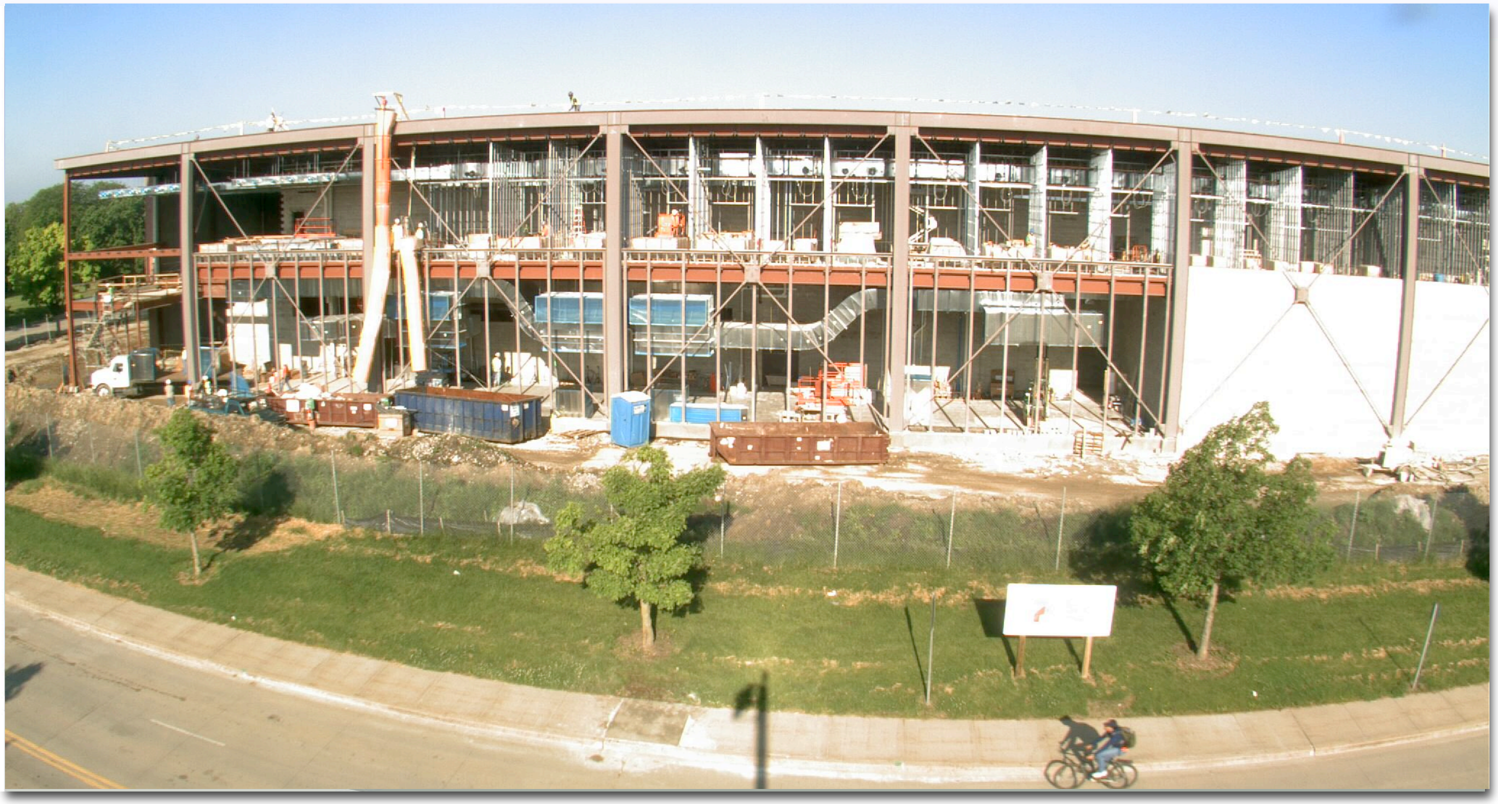

Partners

EYP MCF/
Gensler
IBM
Yahoo!

- **Modern Data Center**
  - 90,000+ ft^2 total
  - 20,000 ft^2 machine room
- **Energy Efficiency**
  - LEED certified (silver)
  - Efficient cooling system

www.ncsa.uiuc.edu/BlueWaters

Selected Computer Science Challenges

- **Programming Models and Languages**

  … will MPI be adequate

  - PGAS (partitioned global address space) programming model
  - Universal parallel C (UPC), Co-array Fortran (CAF)

- **Application Development**

  … developing the next generation of science and engineering applications

  - Integrated Development Environment
  - Scalable algorithms and mathematical libraries

- **Enhanced Reliability**

  … need to minimize impact of/ride through failure

  - Systems level (e.g., virtualization)
  - Applications level

Moving to Exascale Computing

Challenges in Exascale Computing

- **Computational Metrics**
  - Computation
    - Chip architecture
  - Storage capacity
    - Main memory
    - Scratch storage
    - Archival storage
  - Communication
    - Interprocessor network
    - I/O bandwidth
- **Physical Attributes**
  - Physical size
  - Power consumption
- **Balance of Design**
  - Range of applicability

Thank You!