An Overview of Distributed Systems

Ioan Raicu
Computer Science Department, Illinois Institute of Technology
Math and Computer Science Division, Argonne National Laboratory

September 13th, 2011

Famous Quotes

The advent of computation can be compared, in terms of the breadth and depth of its impact on research and scholarship, to the invention of writing and the development of modern mathematics.

Ian Foster, 2006

Science Paradigms

- Thousand years ago:
  science was **empirical**
  describing natural phenomena
- Last few hundred years:
  theoretical branch
  using models, generalizations

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{4\pi Gp}{3} - K\frac{c^2}{a^2}$$

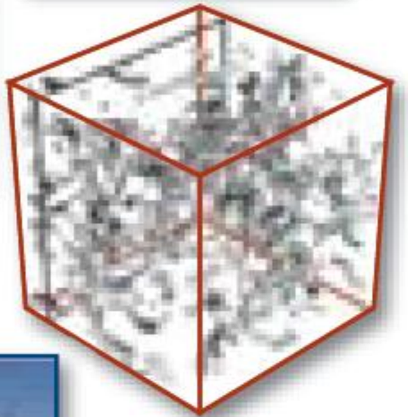

- Last few decades:
  a **computational** branch
  simulating complex phenomena
- Today: **data exploration** (eScience)
  unify theory, experiment, and simulation
  – Data captured by instruments
  or generated by simulator
  – Processed by software
  – Information/knowledge stored in computer
  – Scientist analyzes database/files
  using data management and statistics

Famous Quotes

Computational thinking will be a fundamental skill used by everyone in the world by the middle of the 21st Century.

Jeanette Wing, 2006

X-Info

- The evolution of X-Info and Comp-X for each discipline X
- How to codify and represent our knowledge

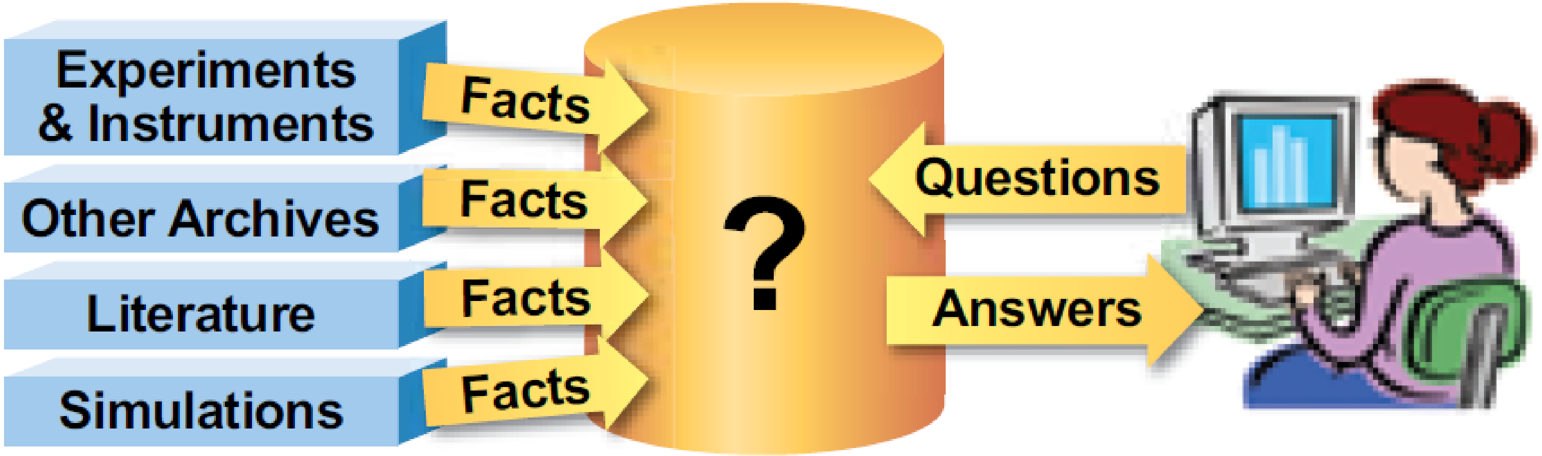

The Generic Problems

- Data ingest
- Managing a petabyte
- Common schema
- How to organize it
- How to reorganize it
- How to share it with others
- Query and Vis tools
- Building and executing models
- Integrating data and literature
- Documenting experiments
- Curation and long-term preservation

All Scientific Data Online

- Many disciplines overlap and use data from other sciences
- Internet can unify all literature and data
- Go from literature to computation to data back to literature
- Information at your fingertips for everyone-everywhere

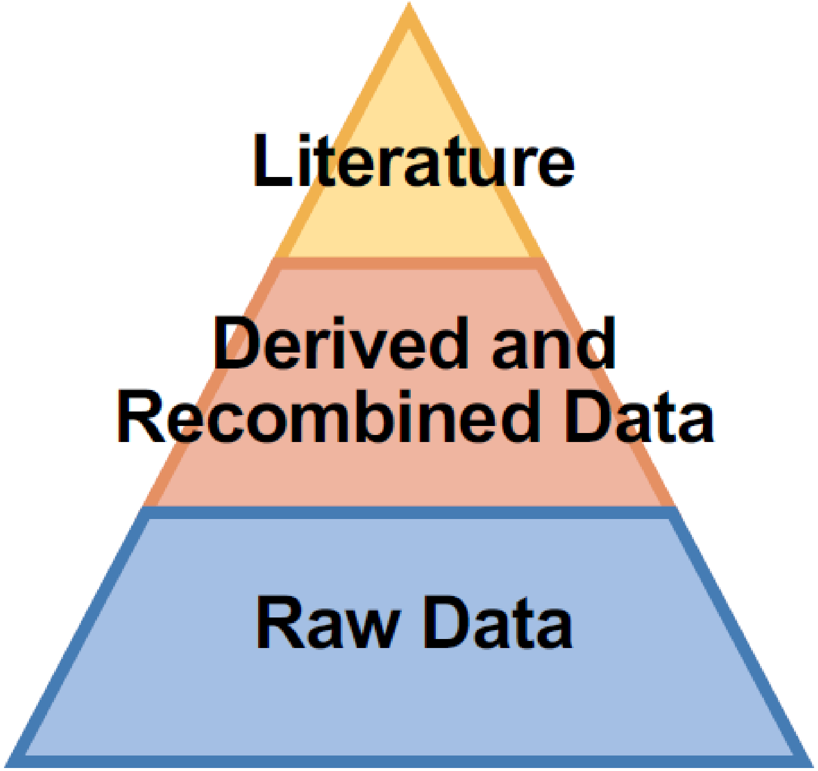

- Increase Scientific Information Velocity
- Huge increase in Science Productivity

Distributed Systems

- What is a distributed system?

"A collection of independent computers that appears to its users as a single coherent system"

-A. Tanenbaum

Distributed Systems

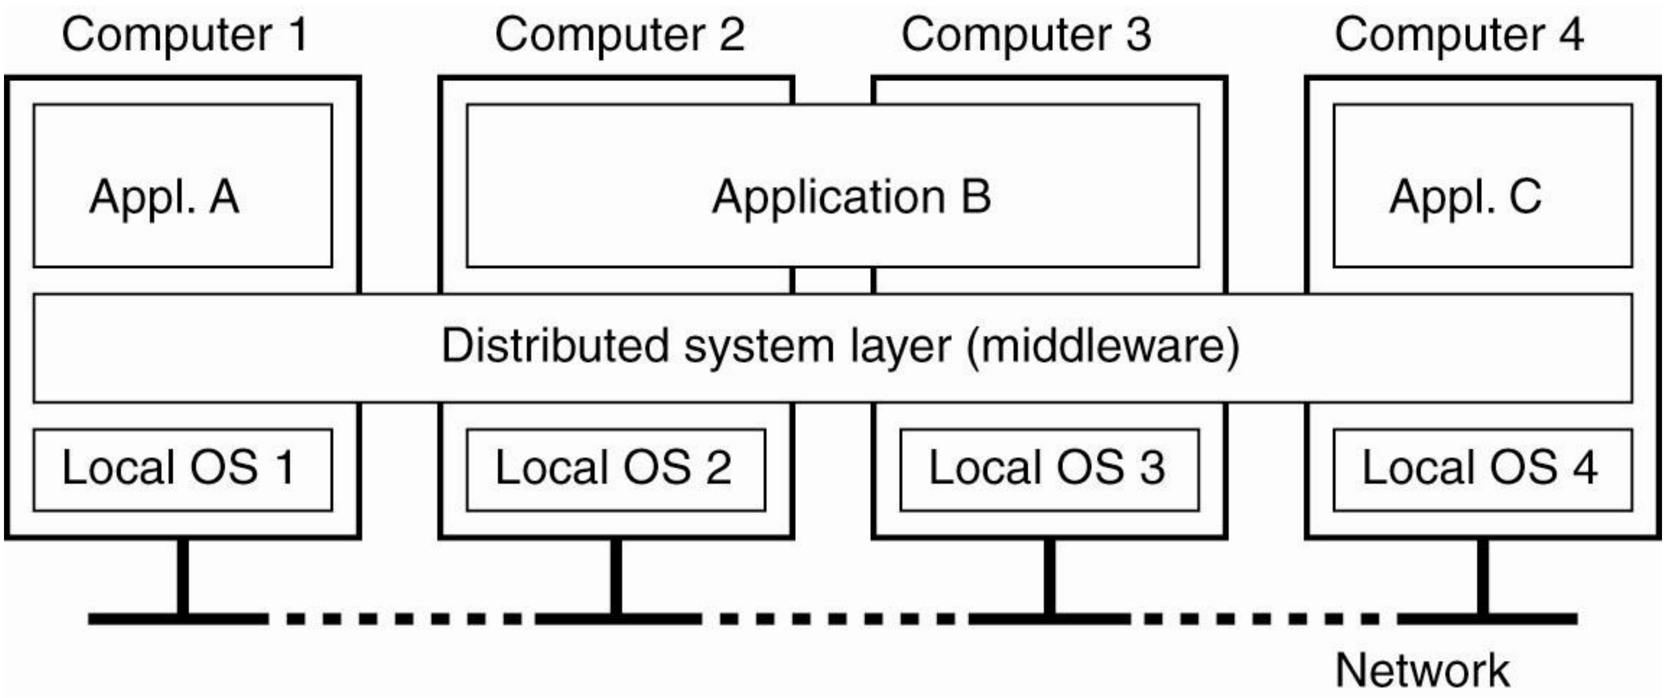

A distributed system organized as middleware. The middleware layer extends over multiple machines, and offers each application the same interface.

Distributed Systems:
Clusters, Grids, Clouds, and Supercomputers

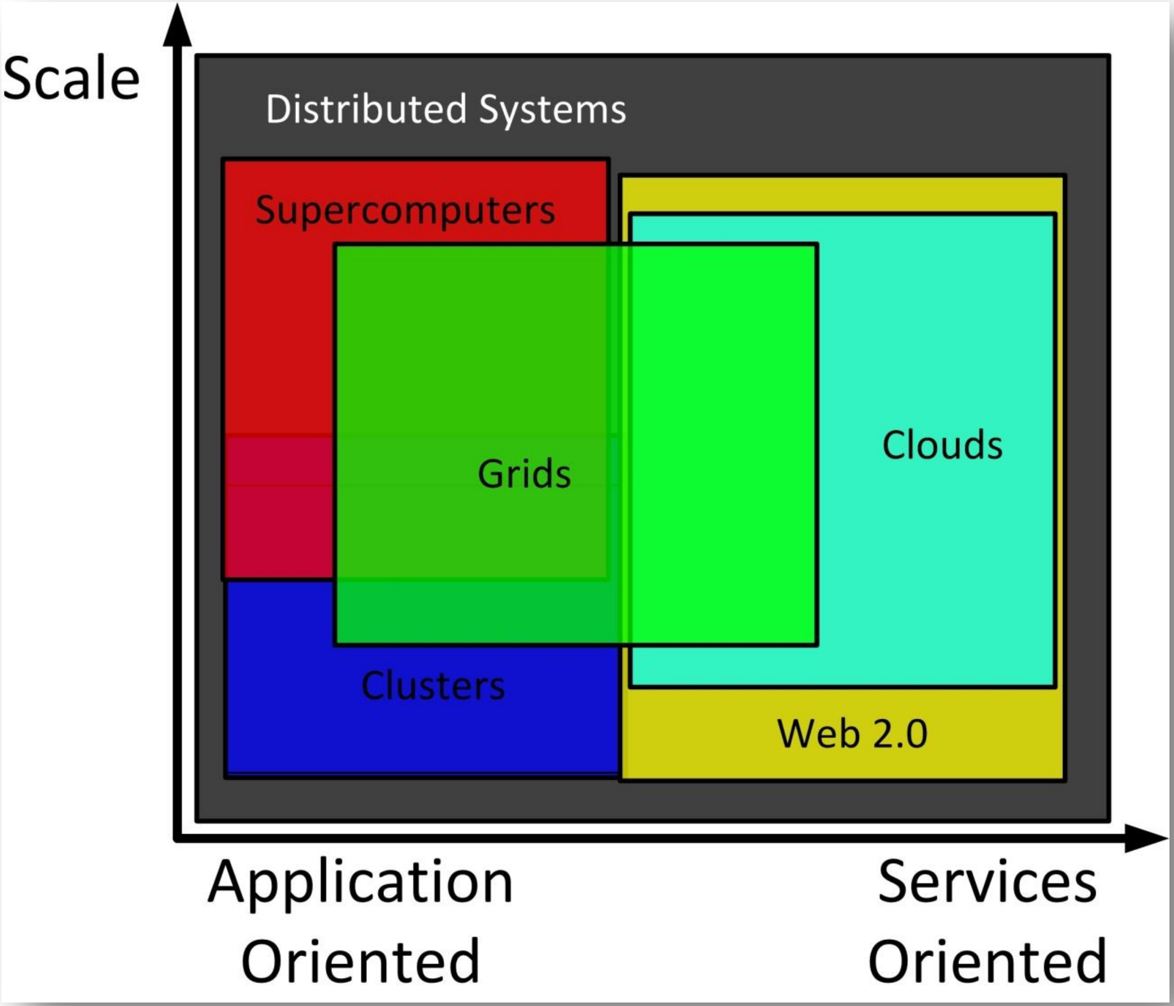

[GCE08] "Cloud Computing and Grid Computing 360-Degree Compared"

Cluster Computing

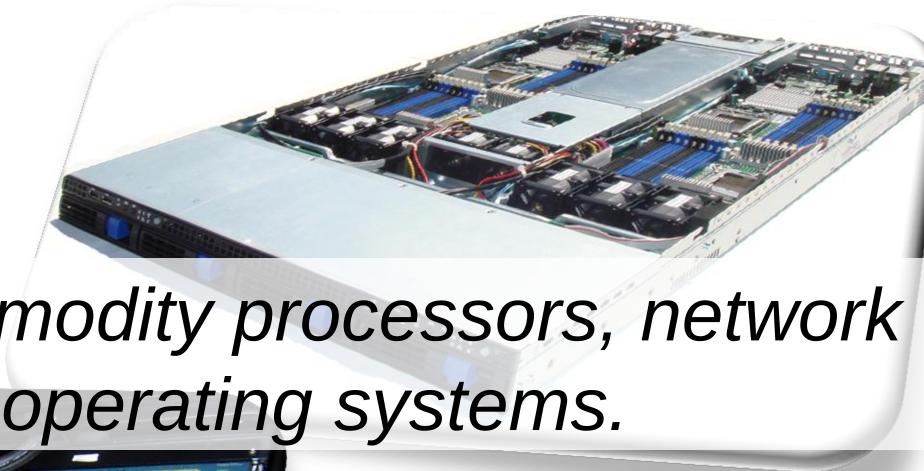

Computer clusters using commodity processors, network interconnects, and operating systems.

Supercomputing

Supercomputing ~ HPC

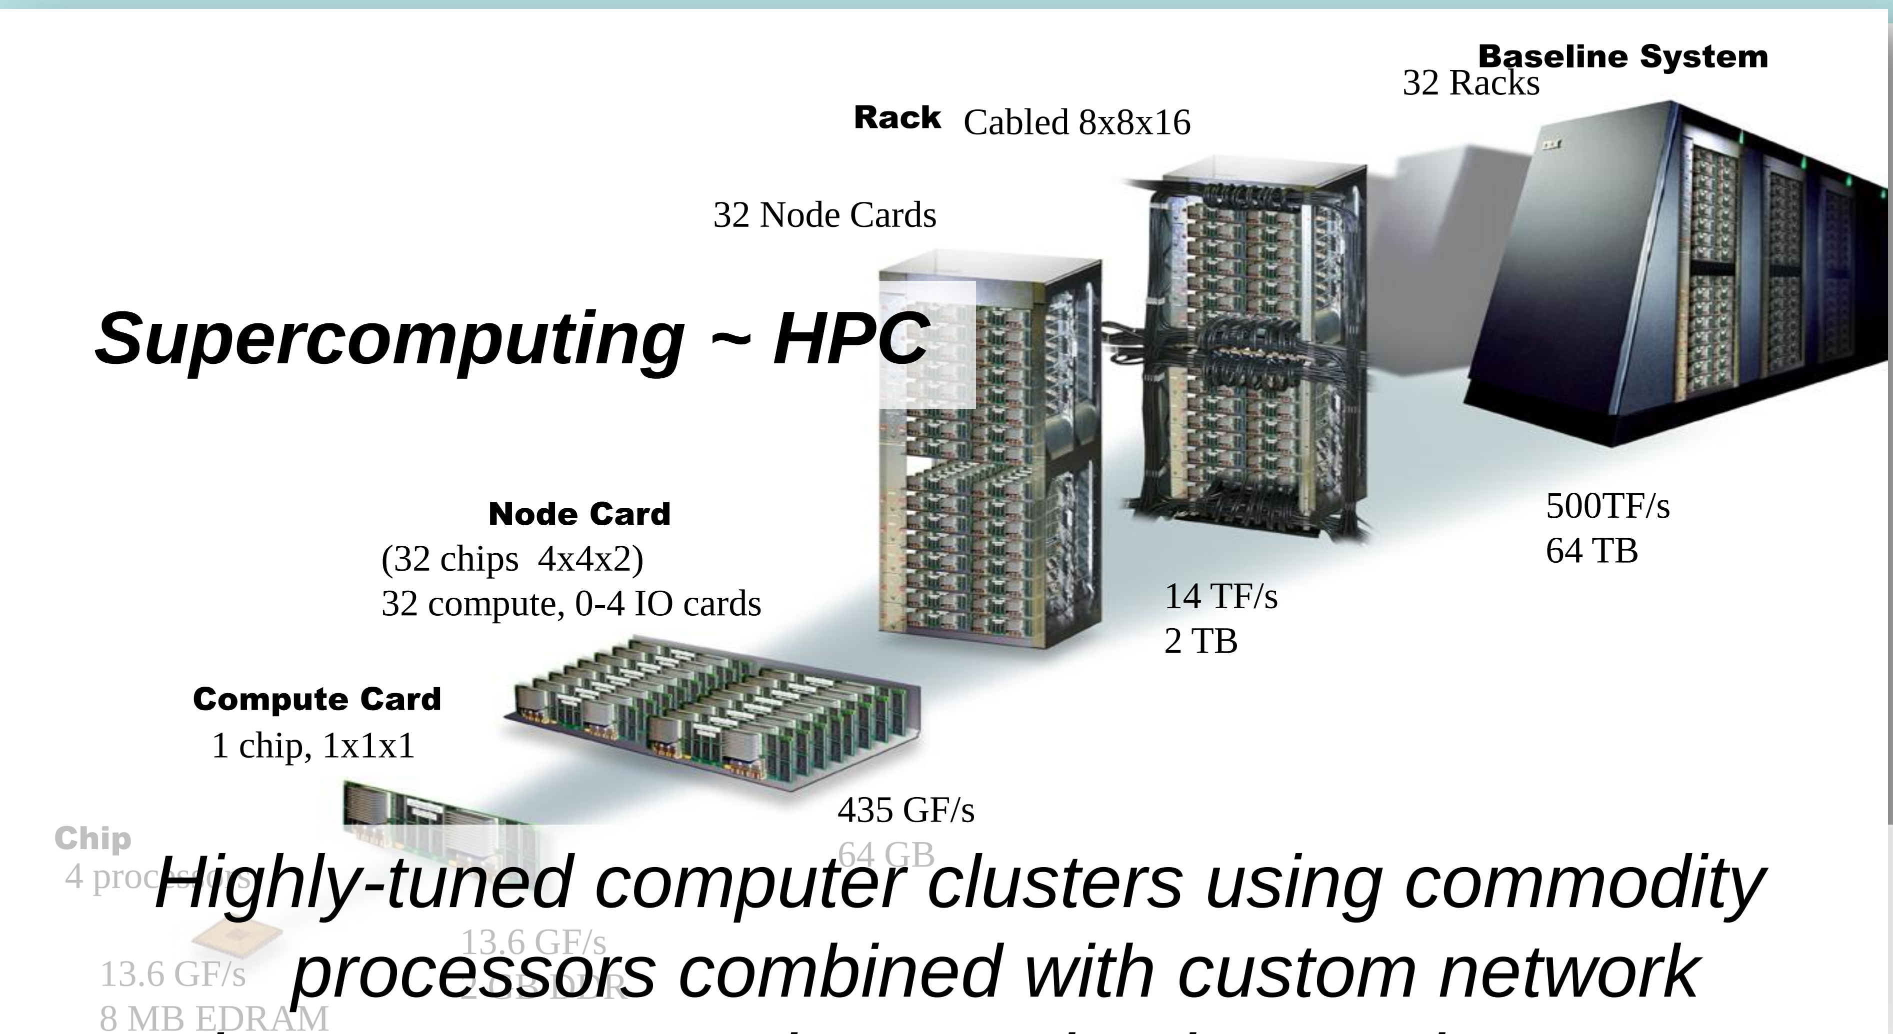

Highly-tuned computer clusters using commodity processors combined with custom network interconnects and customized operating system

Grid Computing

Grids tend to be composed of multiple clusters, and are typically loosely coupled, heterogeneous, and geographically dispersed

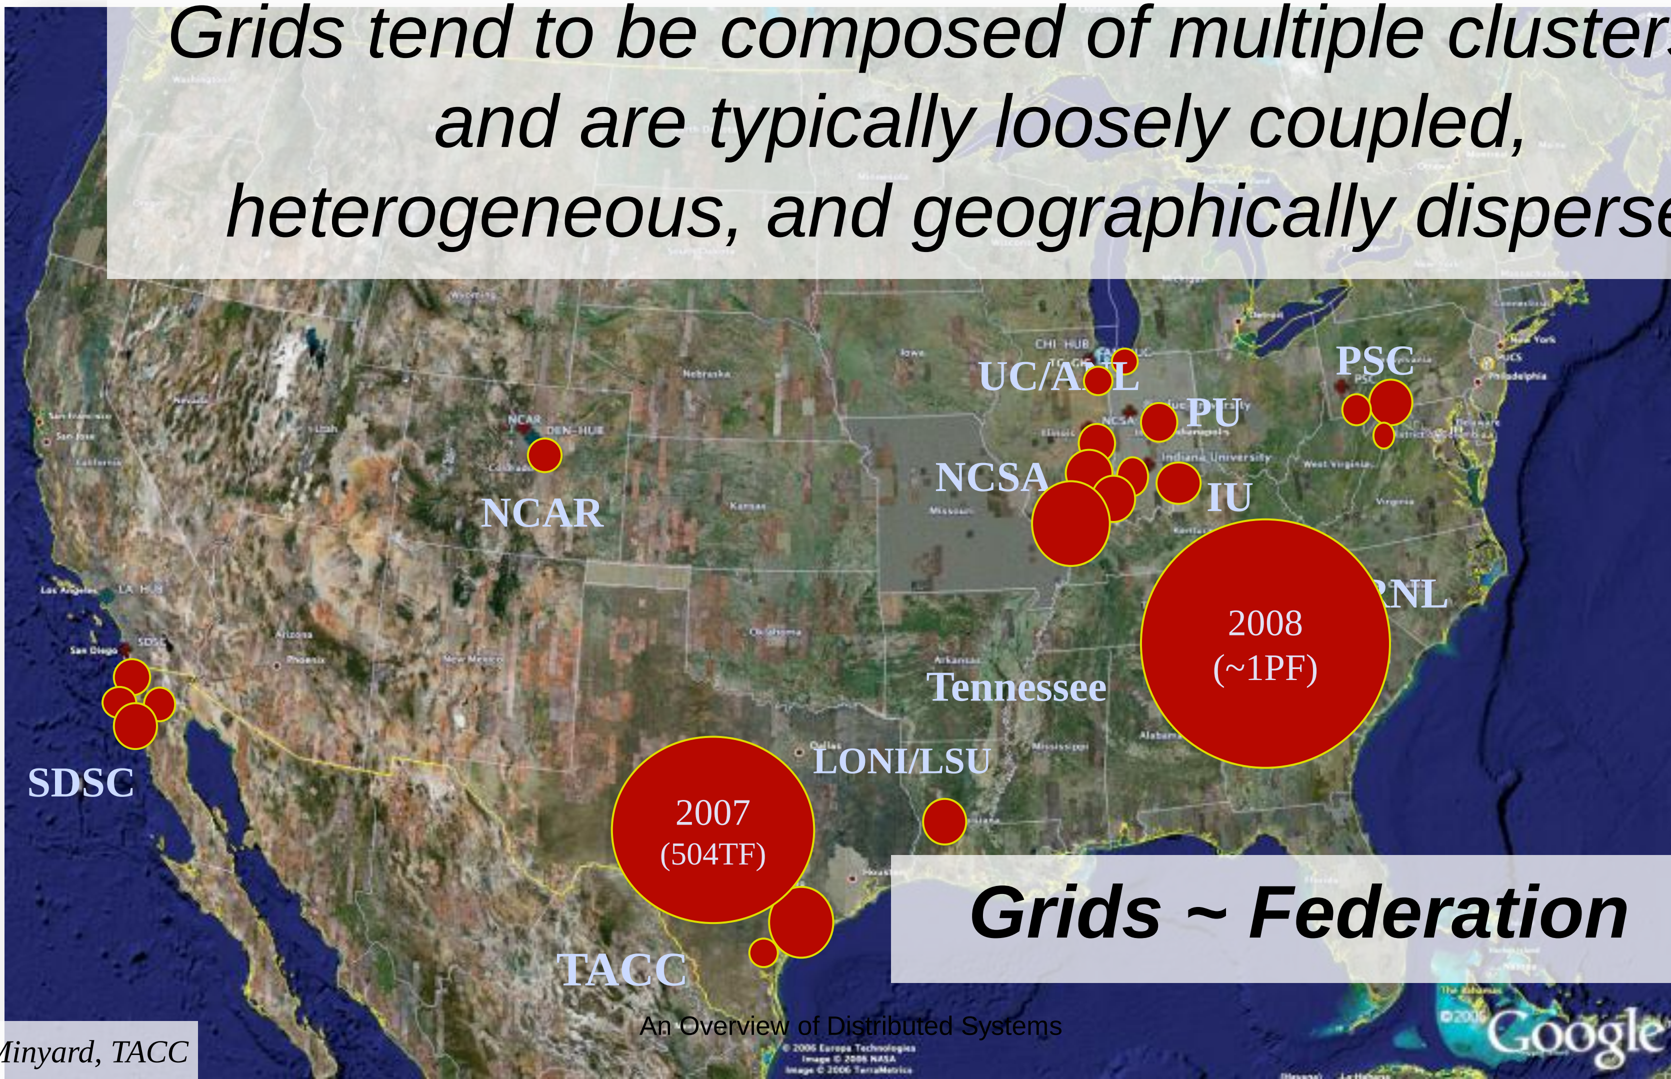

Grids ~ Federation

Tommy Minyard, TACC

Cloud Computing

- *A large-scale distributed computing paradigm driven by:*
  1. *economies of scale*
  2. *virtualization*
  3. *dynamically-scalable resources*
  4. *delivered on demand over the Internet*

Clouds ~ hosting

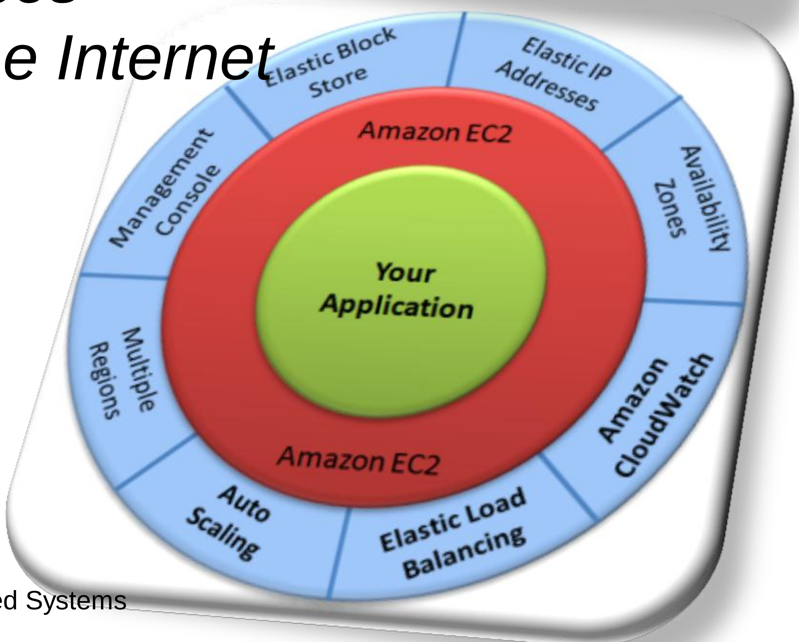

Distributed Pervasive Systems

- Electronic health care systems

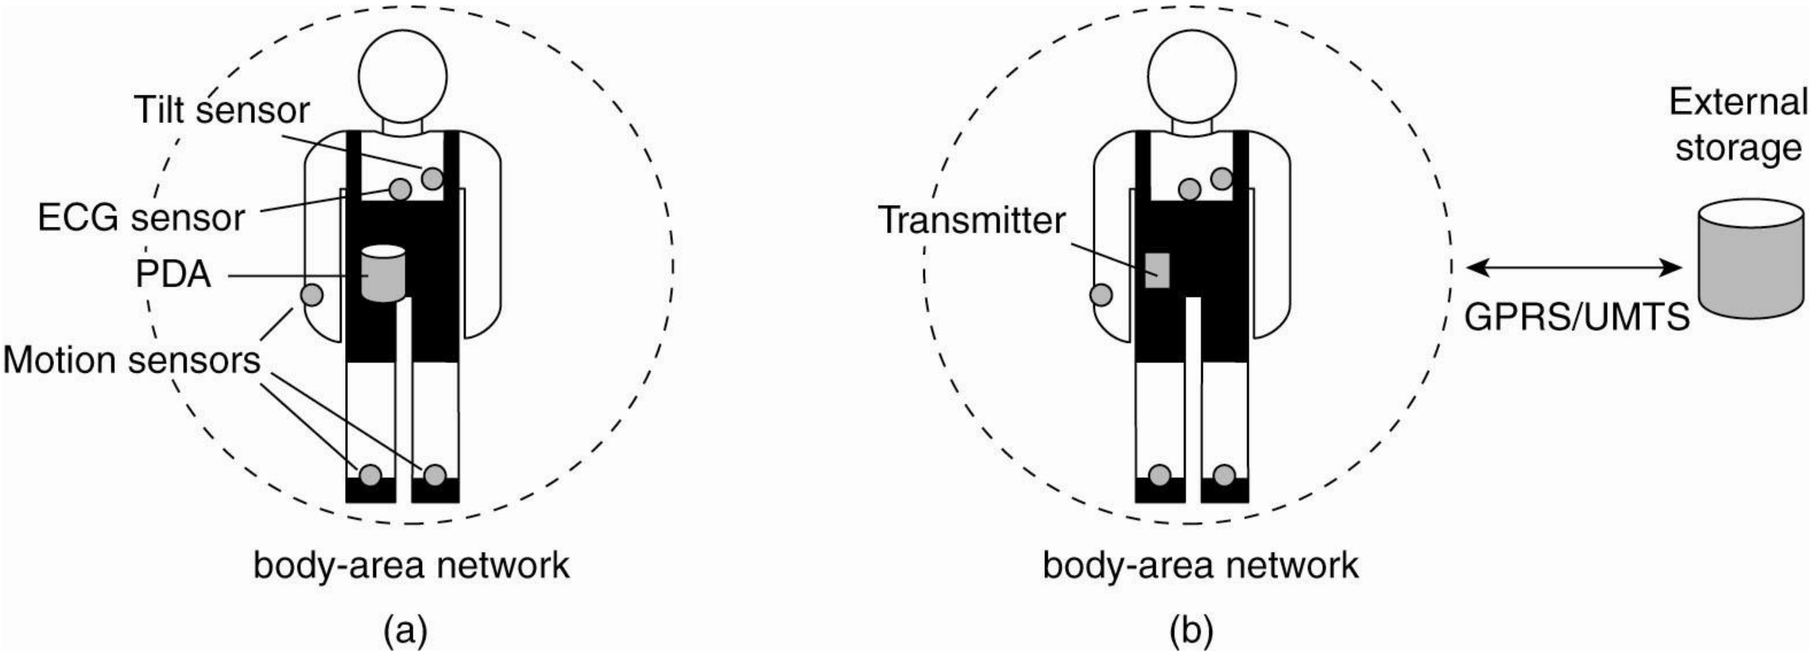

Monitoring a person in a pervasive electronic health care system, using (a) a local hub or (b) a continuous wireless connection.

Distributed Pervasive Systems

- Sensor systems

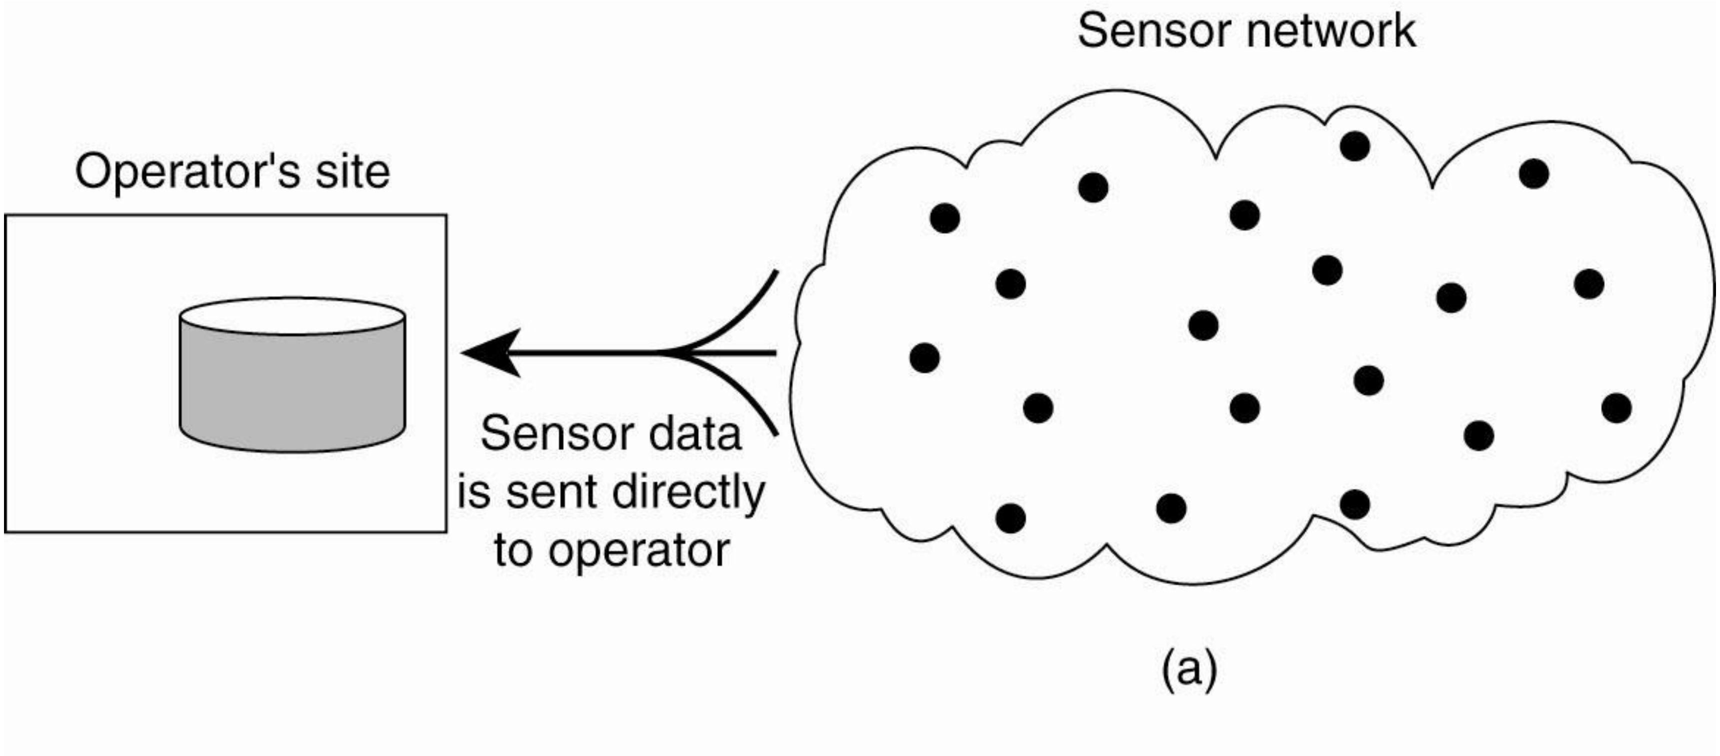

Organizing a sensor network database, while storing and processing data (a) only at the operator's site or ….

Distributed vs. Centralized Systems

- Economics
  - Microprocessors have better price/performance than mainframes
- Speed
  - Collective power of large number of systems
- Geographic and responsibility distribution
- Reliability
  - One machine's failure need not bring down the system
- Extensibility
  - Computers and software can be added incrementally

Disadvantages of Distributed Systems

- Software
  - Little software exists compared to PCs
- Networking
  - Still slow and can cause other problems (e.g. when disconnected)
- Security
  - Data may be accessed by unauthorized users

Applications
Economic Modeling: MARS

- CPU Cores: 130816
- Tasks: 1048576
- Elapsed time: 2483 secs
- CPU Years: 9.3

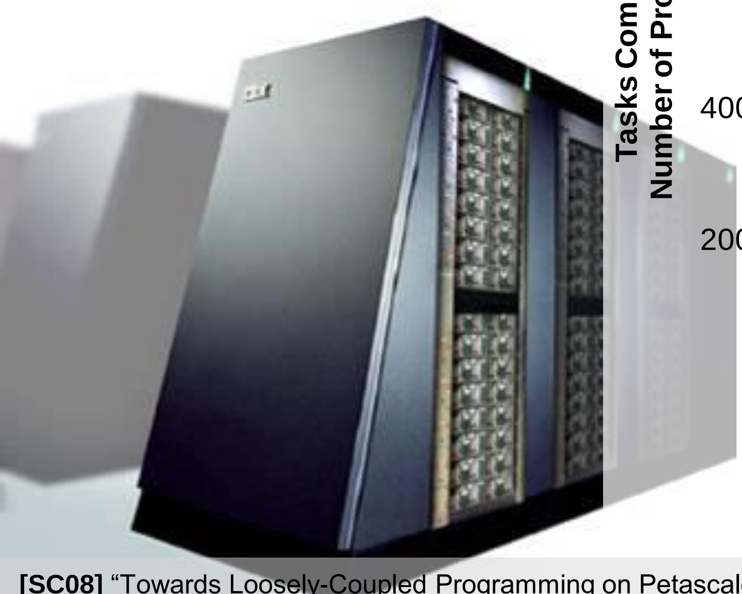

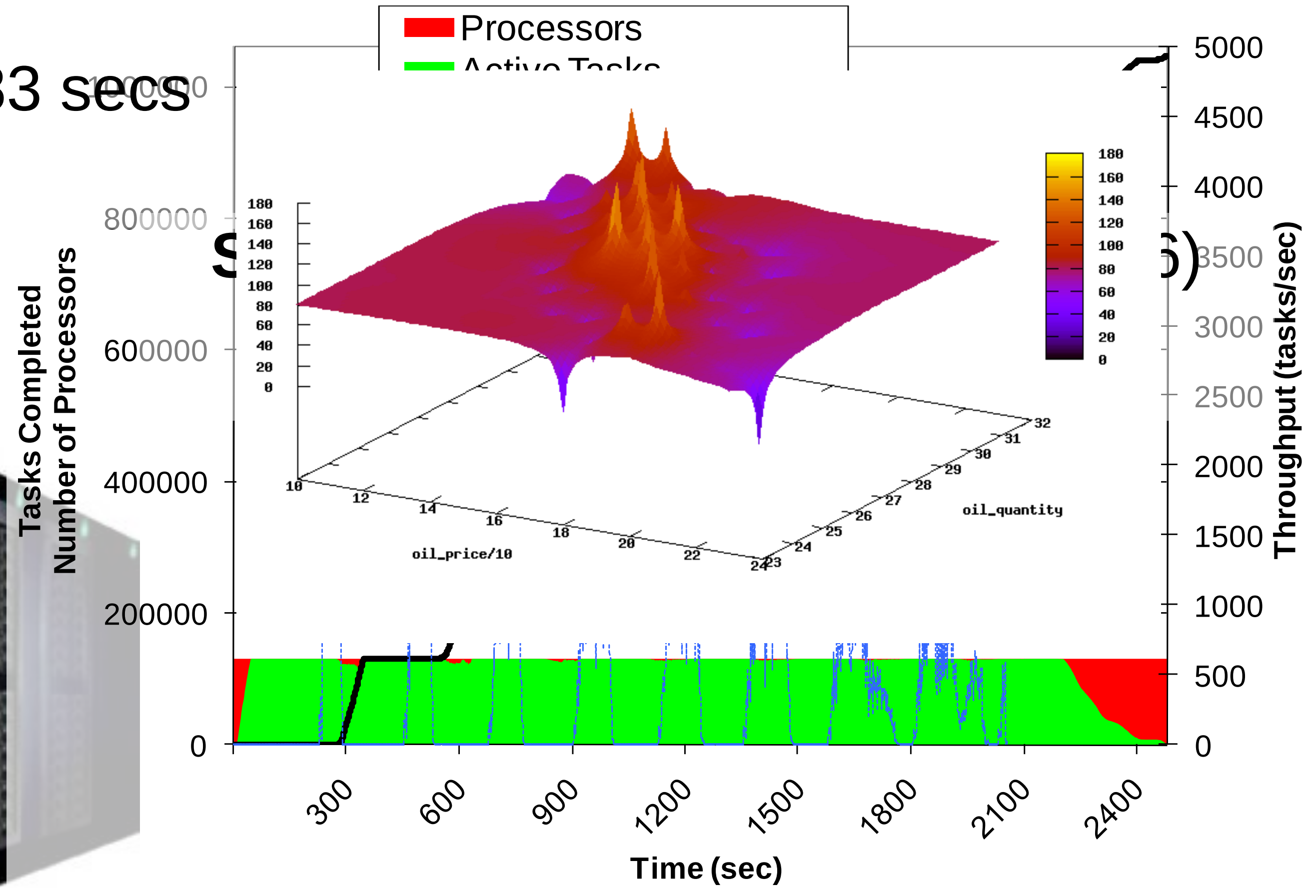

Applications: Pharmaceuticals

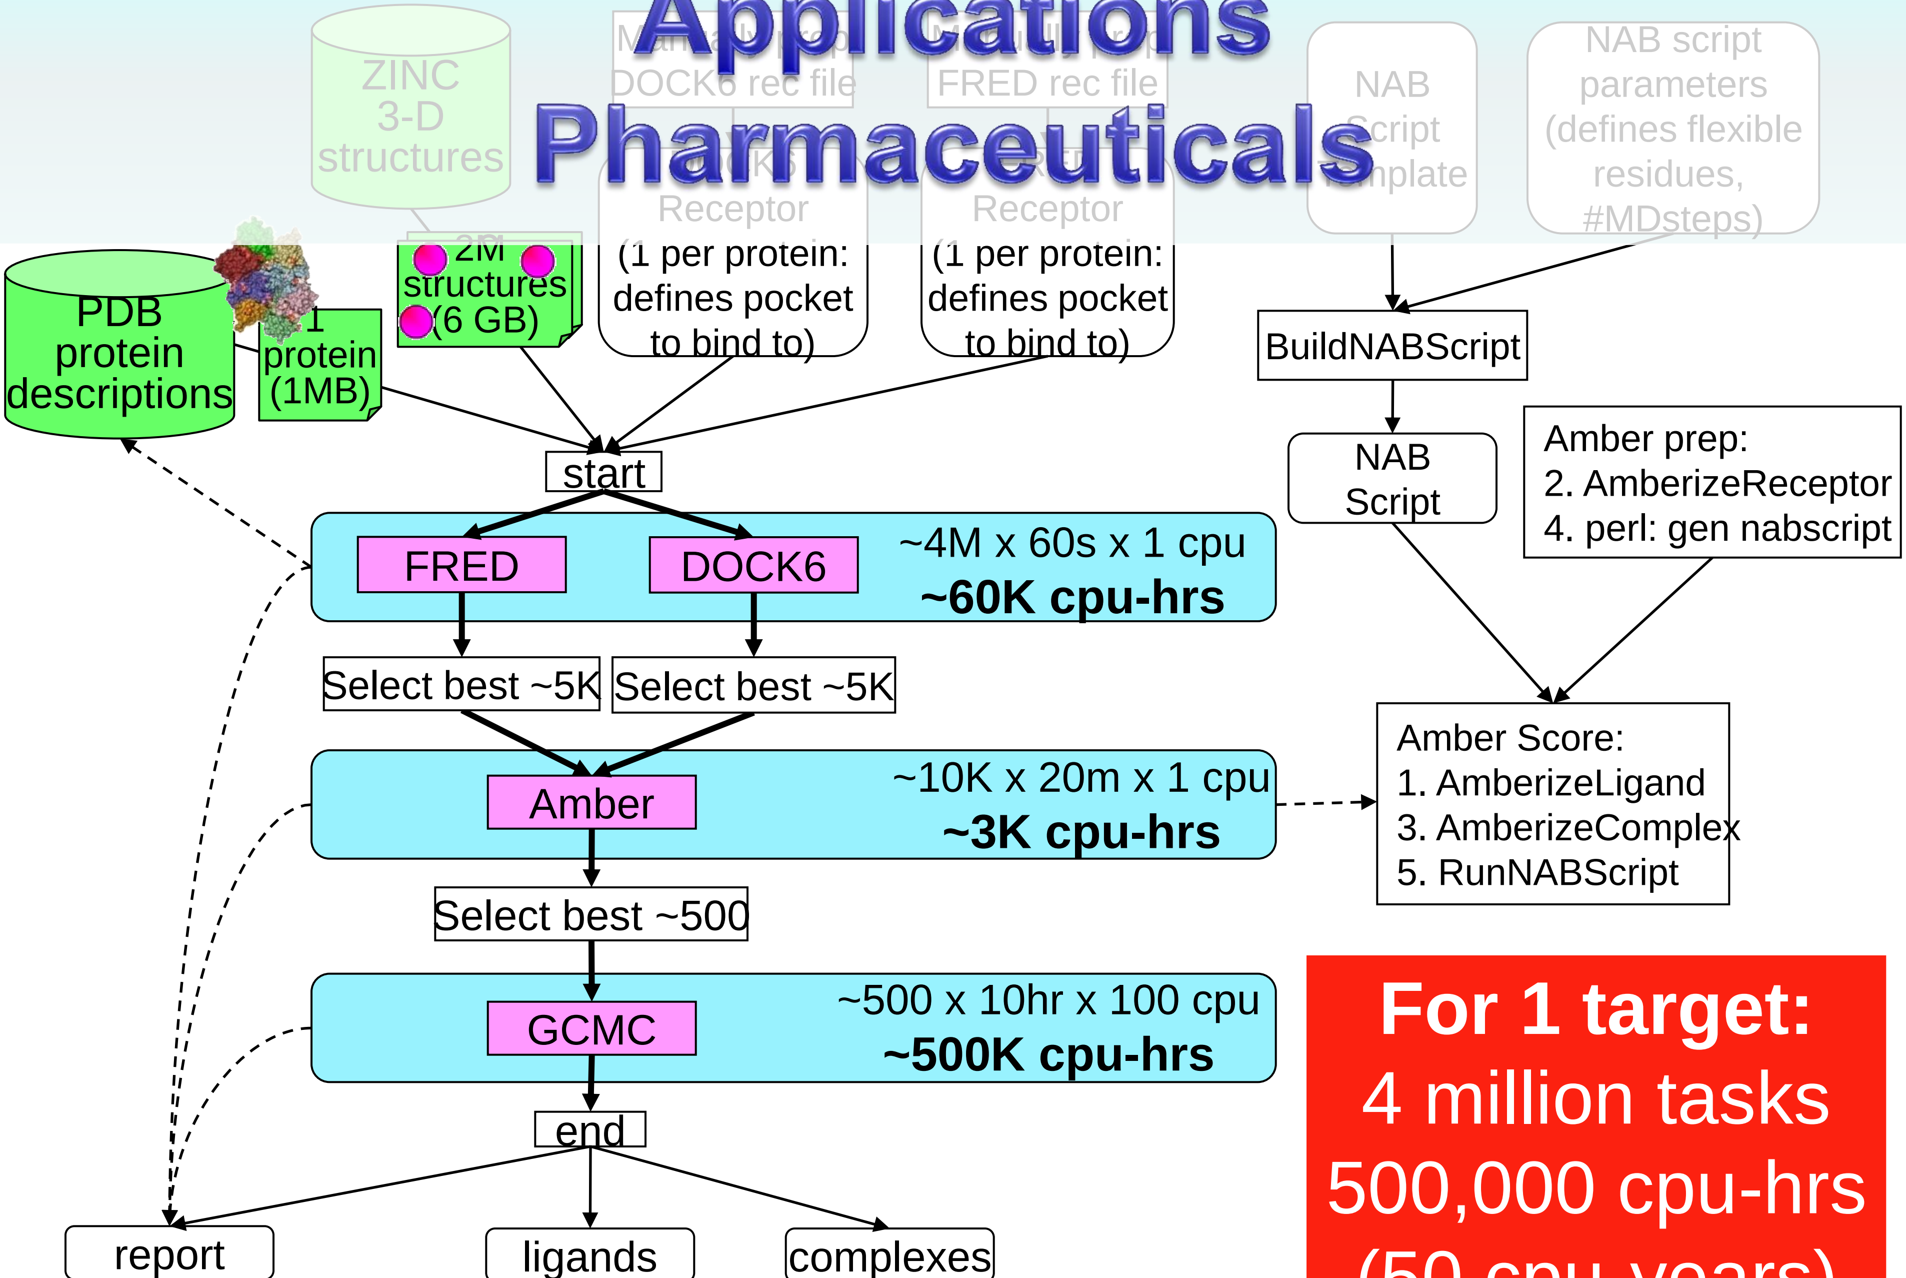

For 1 target:
4 million tasks
500,000 cpu-hrs
(50 cpu-years)

Applications Pharmaceuticals: DOCK

CPU cores: 118784

Tasks: 934803

Elapsed time: 2.01 hours

Compute time: 21.43 CPU years

Average task time: 667 sec

Relative Efficiency: 99.7%

(from 16 to 32 racks)

Utilization:

- Sustained: 99.6%
- Overall: 78.3%

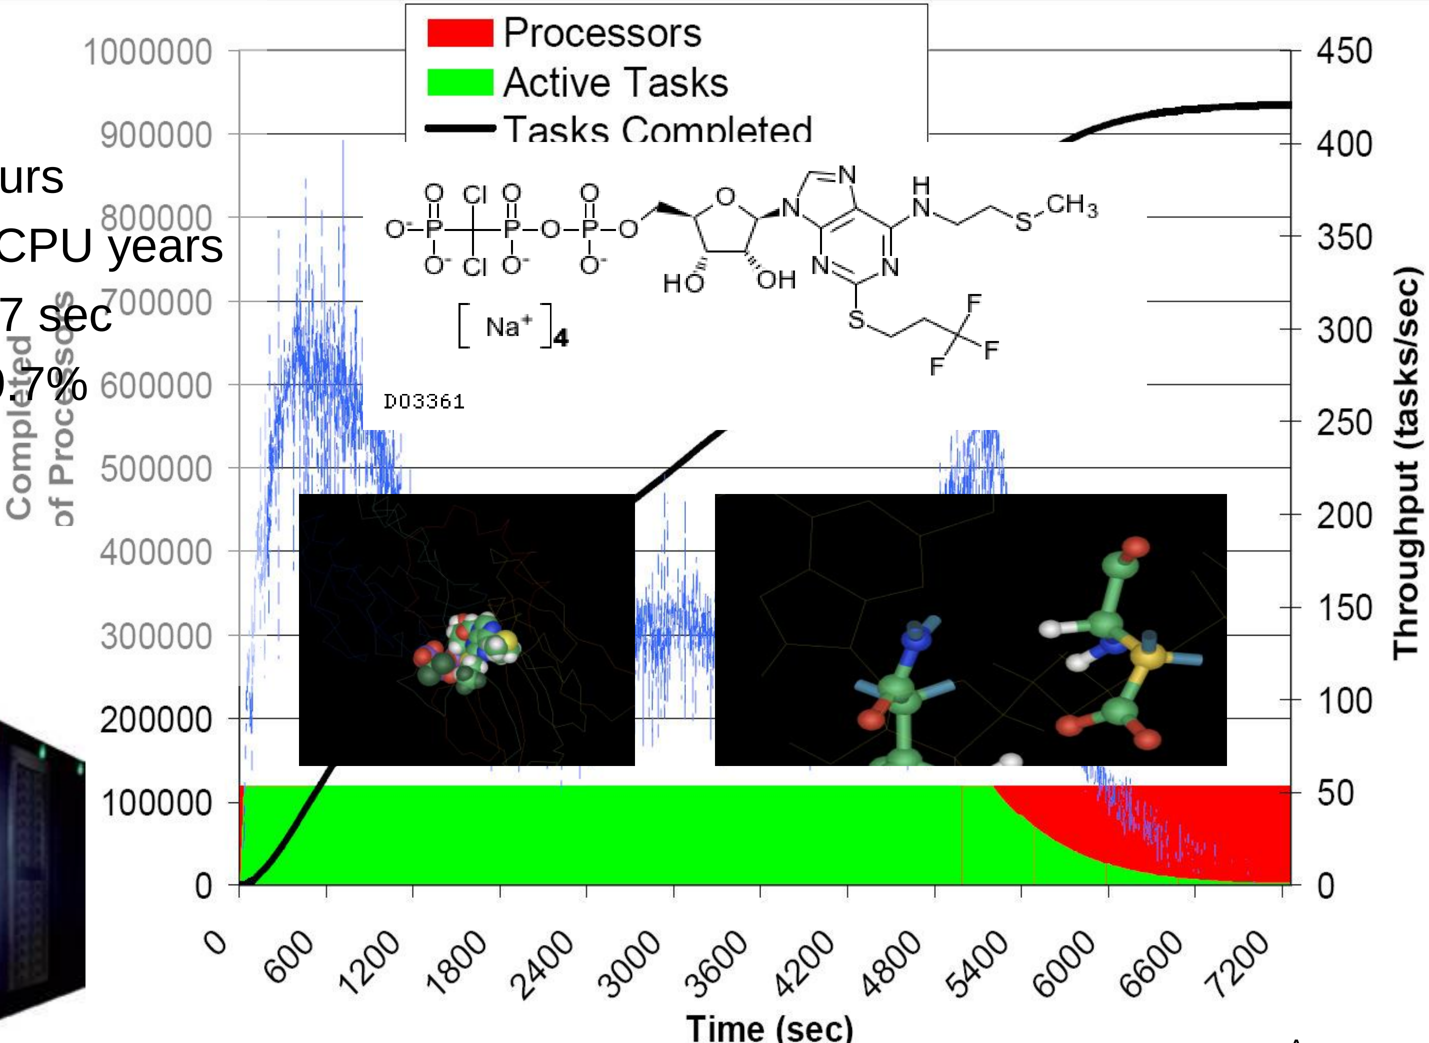

Applications
Astronomy: AstroPortal

- Purpose
  - On-demand “stacks” of random locations within ~10TB dataset
- Challenge
  - Processing Costs:
    - O(100ms) per object
  - Data Intensive:
    - 40MB:1sec
  - Rapid access to 10-10K “random” files
  - Time-varying load

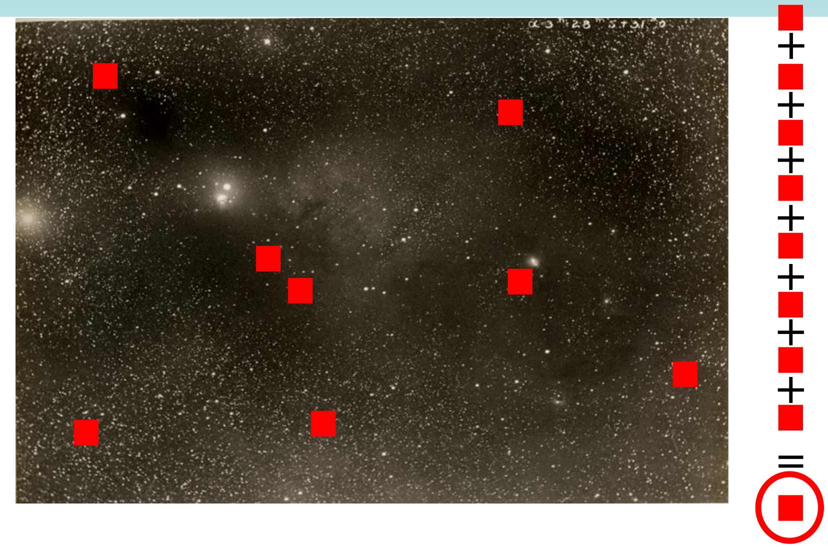

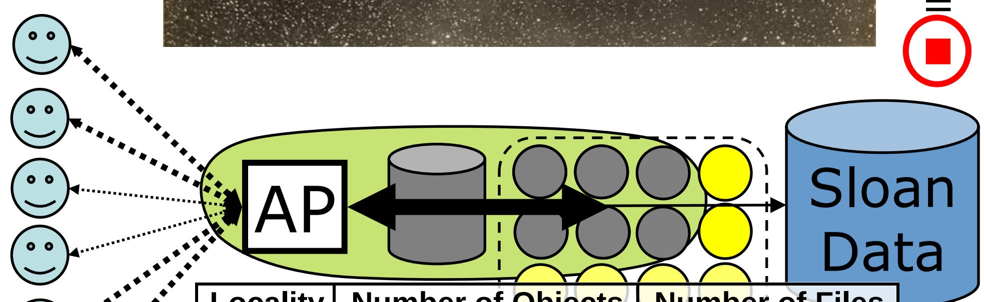

Locality	Number of Objects	Number of Files
1	111700	111700
1.38	154345	111699
2	97999	49000
3	88857	29620
4	76575	19145
5	60590	12120
10	46480	4650
20	40460	2025
30	23695	790

[DADC08] “Accelerating Large-scale Data Exploration through Data Diffusion”
[TG06] “AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis”

DataSys
Data-Intensive Distributed
Systems Laboratory

DataSys: Data-Intensive Distributed Systems Laboratory

- **Research Focus**
  - Emphasize designing, implementing, and evaluating systems, protocols, and middleware with the goal of supporting **data-intensive applications on extreme scale distributed systems**, from many-core systems, clusters, grids, clouds, and supercomputers
- **People**
  - 1 Faculty member
  - 5 PhD students
  - 3 MS students
  - 1 UG student
- **More information**
  - http://datasys.cs.iit.edu/

Future of Distributed Systems Manycore Computing

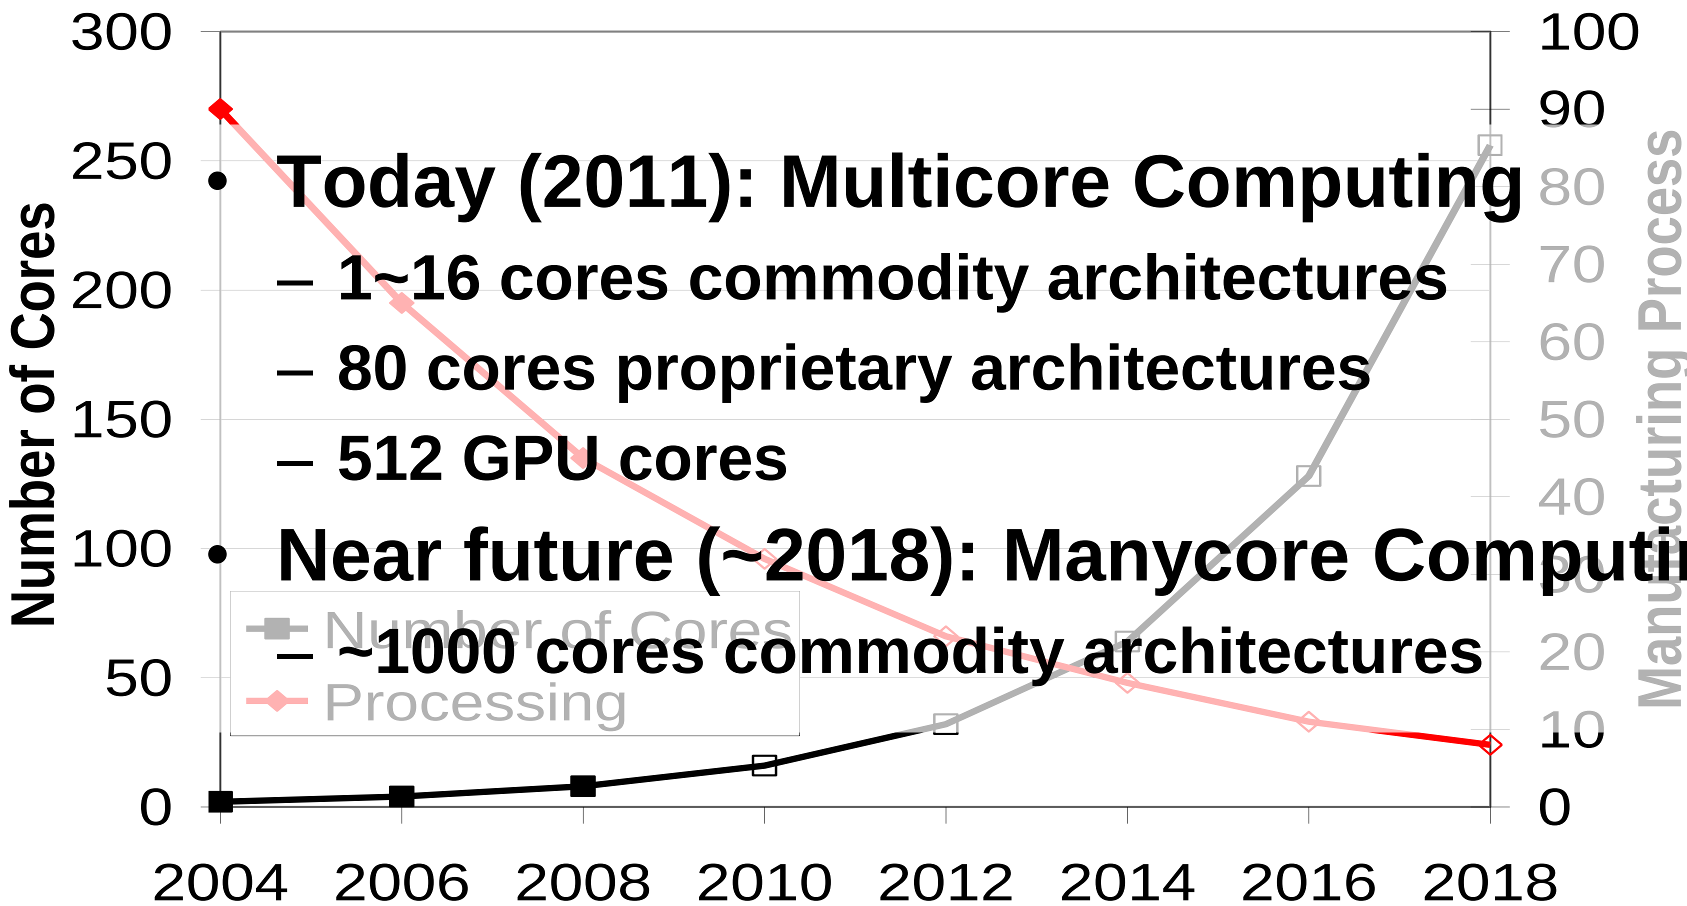

- **Today (2011): Multicore Computing**
  - **1~16 cores commodity architectures**
  - **80 cores proprietary architectures**
  - **512 GPU cores**
- **Near future (~2018): Manycore Computing**
  - **~1000 cores commodity architectures**

Pat Helland, Microsoft, The Irresistible Forces Meet the Movable Objects, November 9th, 2007

Future of Distributed Systems
Exascale Computing

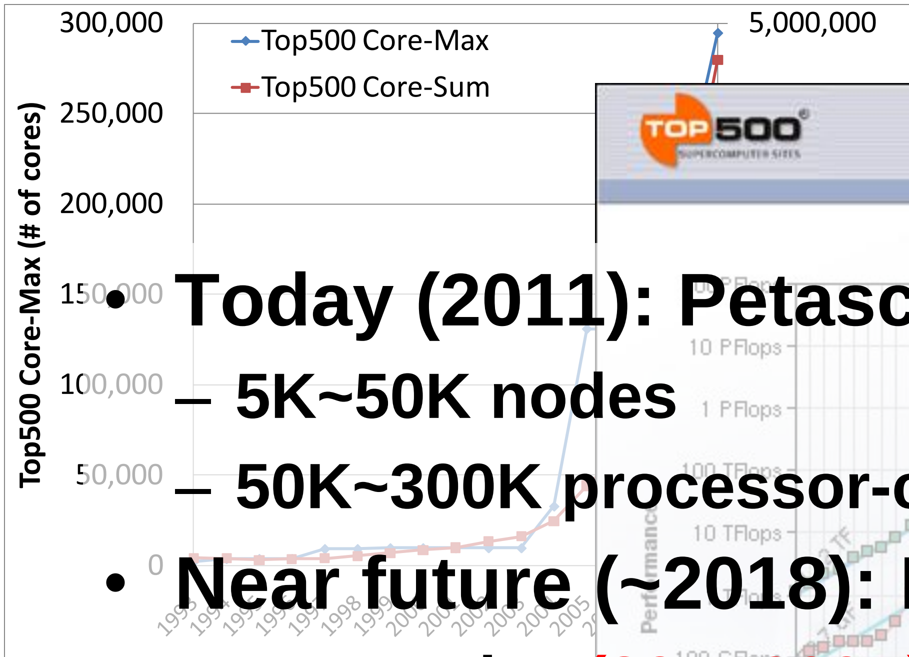

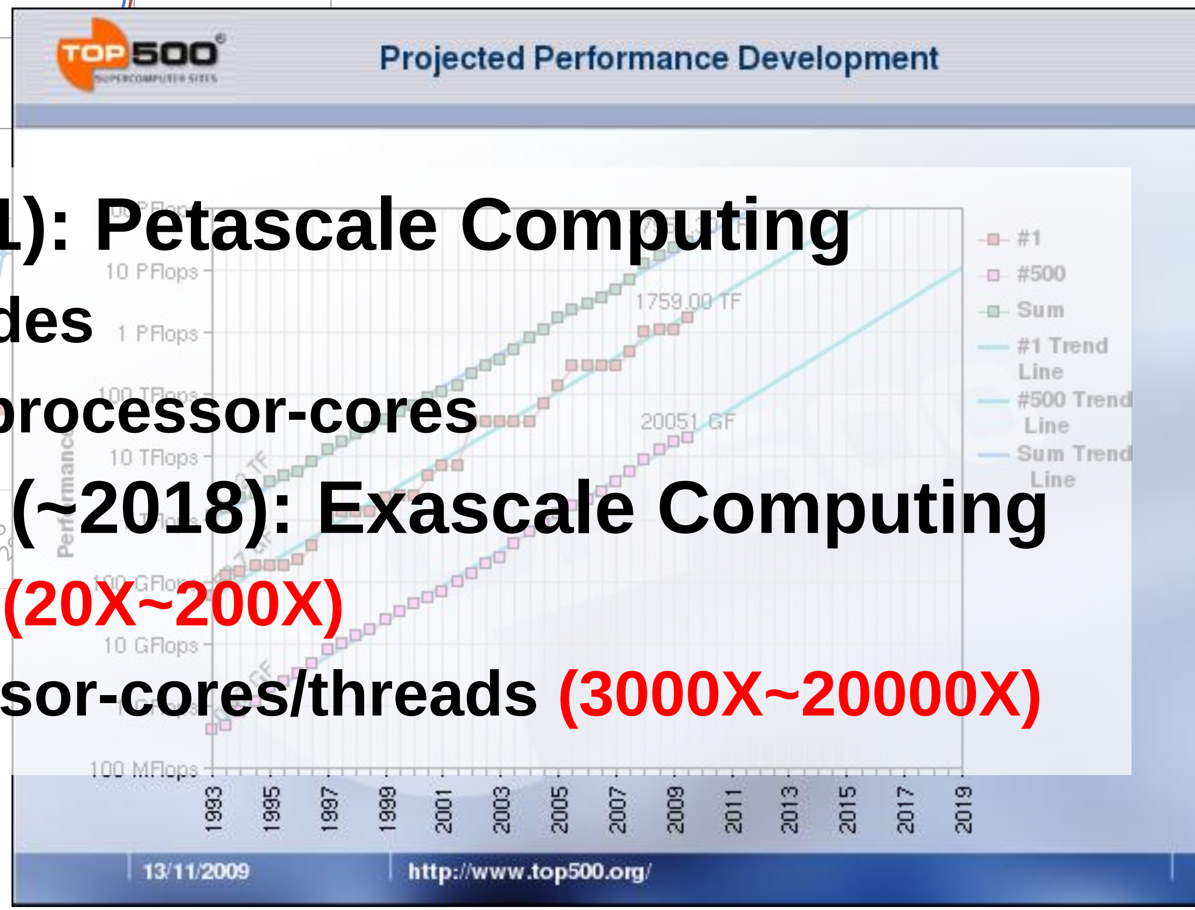

- **Today (2011): Petascale Computing**
  - **5K~50K nodes**
  - **50K~300K processor-cores**
- **Near future (~2018): Exascale Computing**
  - **~1M nodes (20X~200X)**
  - **~1B processor-cores/threads (3000X~20000X)**

Top500 Projected Development, http://www.top500.org/lists/2009/11/performance_development

Future of Distributed Systems Cloud Computing

- Relatively new paradigm… 3 years old
- Amazon in 2009
  - 40K servers split over 6 zones
    - 320K-cores, 320K disks
    - $100M costs + $12M/year in energy costs
    - Revenues about $250M/year
- Amazon in 2018
  - Will likely look similar to exascale computing
    - 100K~1M nodes, ~1B-cores, ~1M disks
    - $100M~$200M costs + $10M~$20M/year in energy
    - Revenues 100X~1000X of what they are today

Future of Distributed Systems Common Challenges

- Power efficiency
  - Will limit the number of cores on a chip (Manycore)
  - Will limit the number of nodes in cluster (Exascale and Cloud)
  - Will dictate a significant part of the cost of ownership
- Programming models/languages
  - Automatic parallelization
  - Threads, MPI, workflow systems, etc
  - Functional, imperative
  - Languages vs. Middlewares

Future of Distributed Systems Common Challenges

- Bottlenecks in scarce resources
  - Storage (Exascale and Clouds)
  - Memory (Manycore)
- Reliability
  - How to keep systems operational in face of failures
  - Checkpointing (Exascale)
  - Node-level replication enabled by virtualization (Exascale and Clouds)
  - Hardware redundancy and hardware error correction (Manycore)

Summary

- Everything about science is changing because of the impact of information technology
- Experimental, theoretical, and computational science are all being affected by the data deluge, and a fourth, “data-intensive” science paradigm is emerging.
- Goal
  - A world in which all of the science literature is online
  - all of the science data is online
  - They interoperate with each other
- Computing and storage will increase in scale exponentially over the next decade
- Data will play central role in future computing systems
- Lots of new tools are needed to make this happen!

More Information

- More information:
  - http://www.cs.iit.edu/~iraicu/
  - http://datasys.cs.iit.edu/
- Contact:
  - iraicu@cs.iit.edu
- Questions?