Introduction to Distributed Systems

Ioan Raicu
Computer Science Department
Illinois Institute of Technology

CS 595: Data-Intensive Computing
August 29th, 2011

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

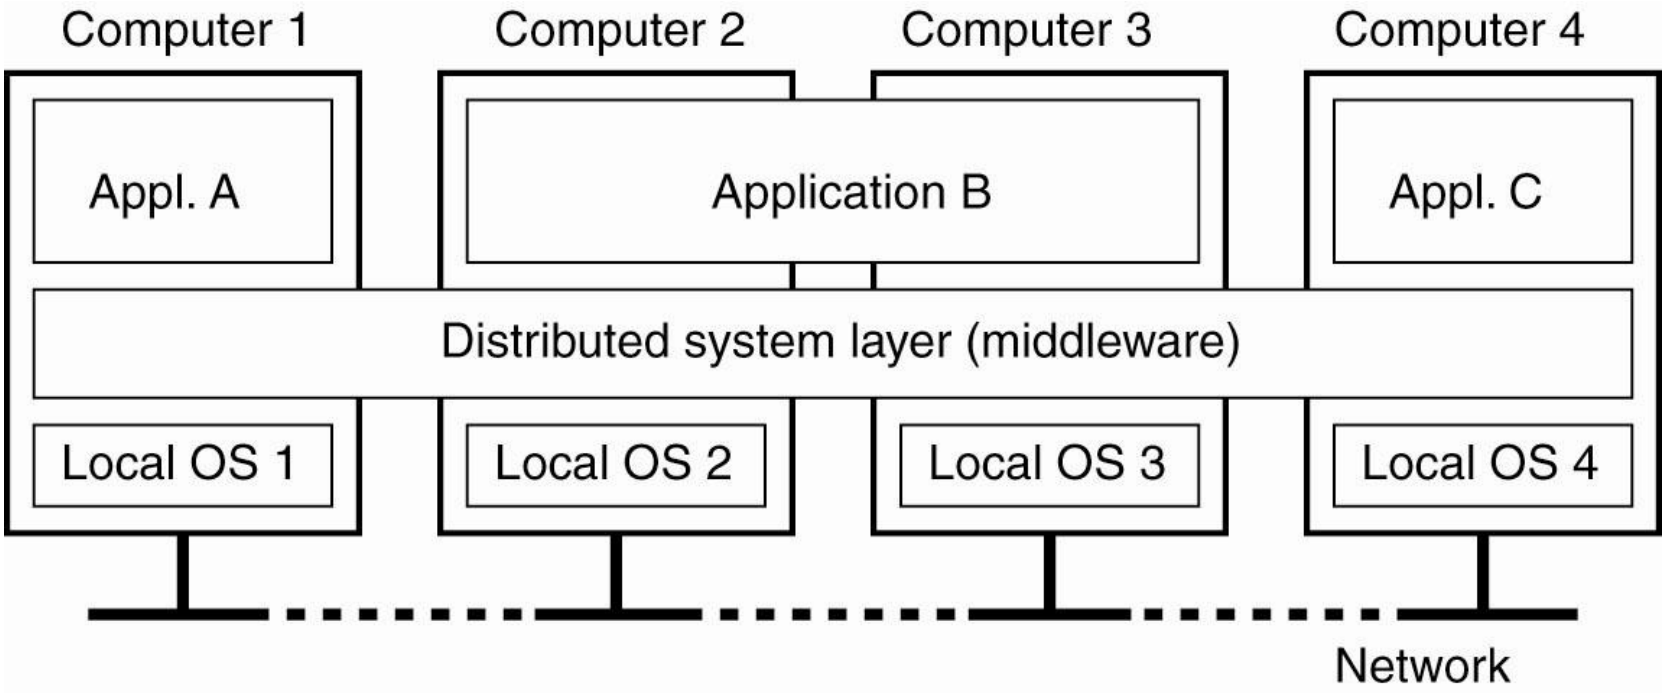

A distributed system organized as middleware. The middleware layer extends over multiple machines, and offers each application the same interface.

Distributed Systems

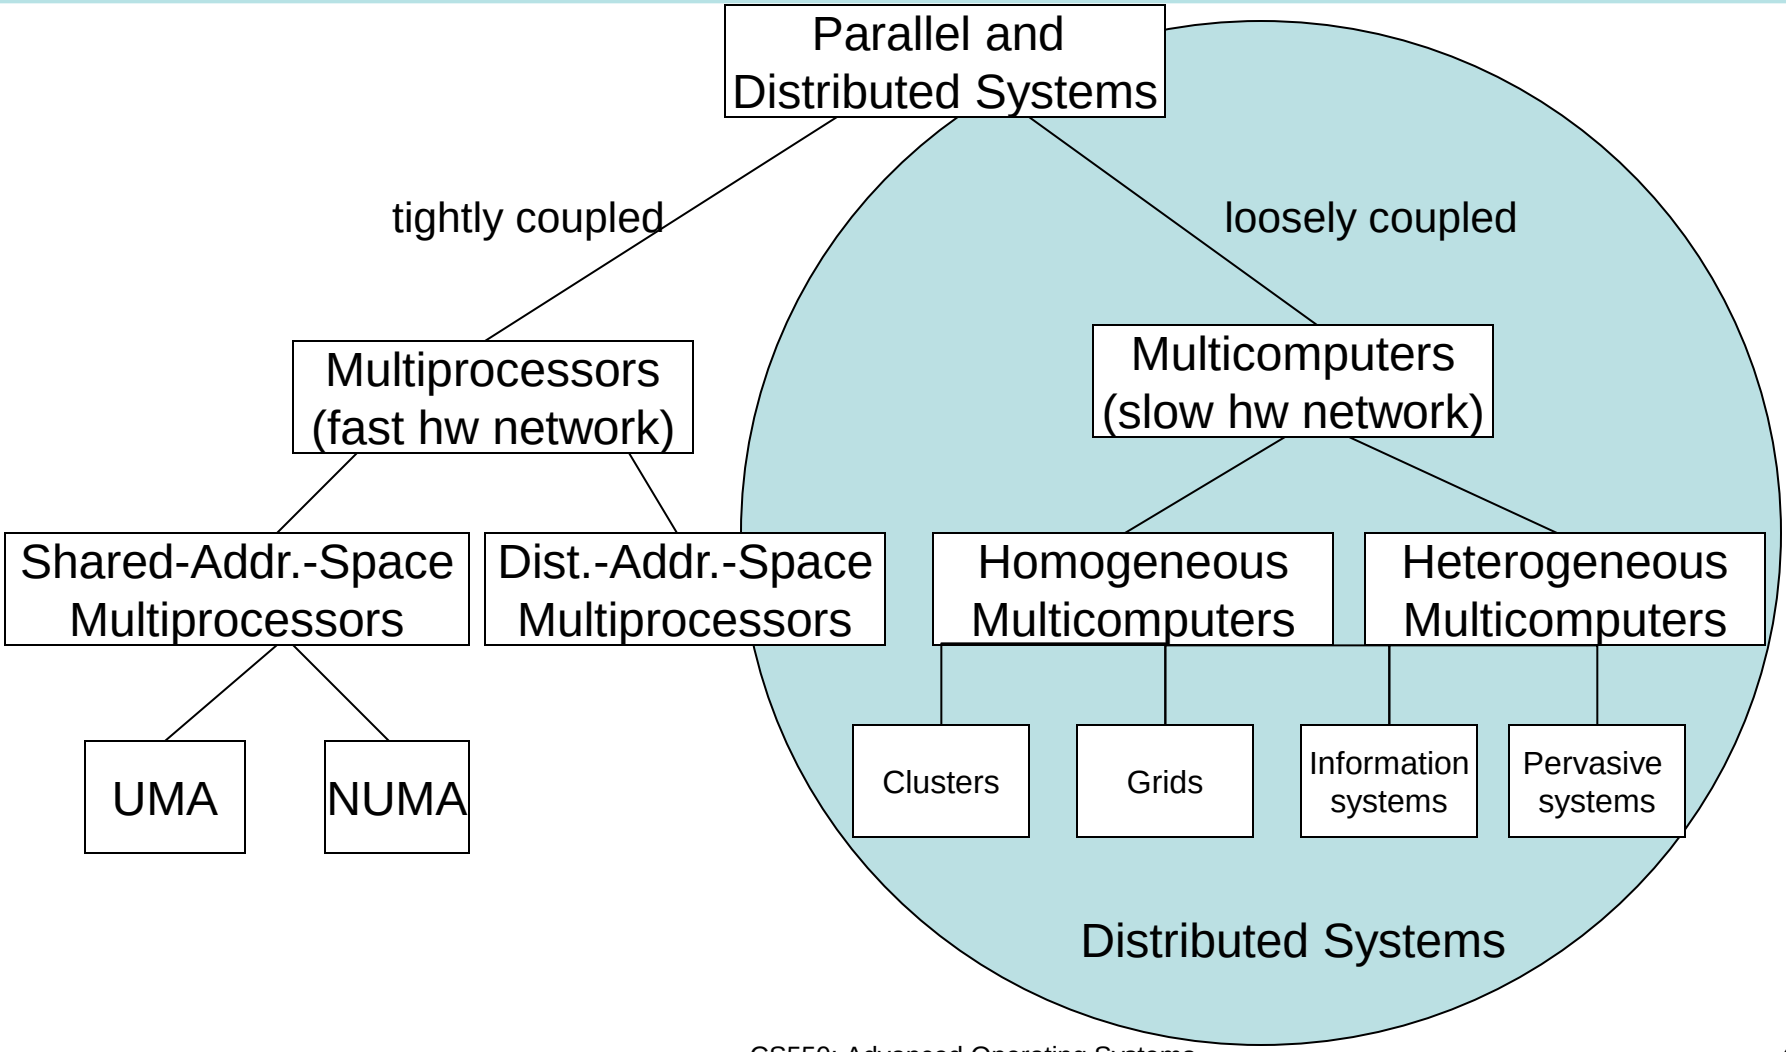

Distributed Systems:
Clusters, Grids, Clouds, and Supercomputers

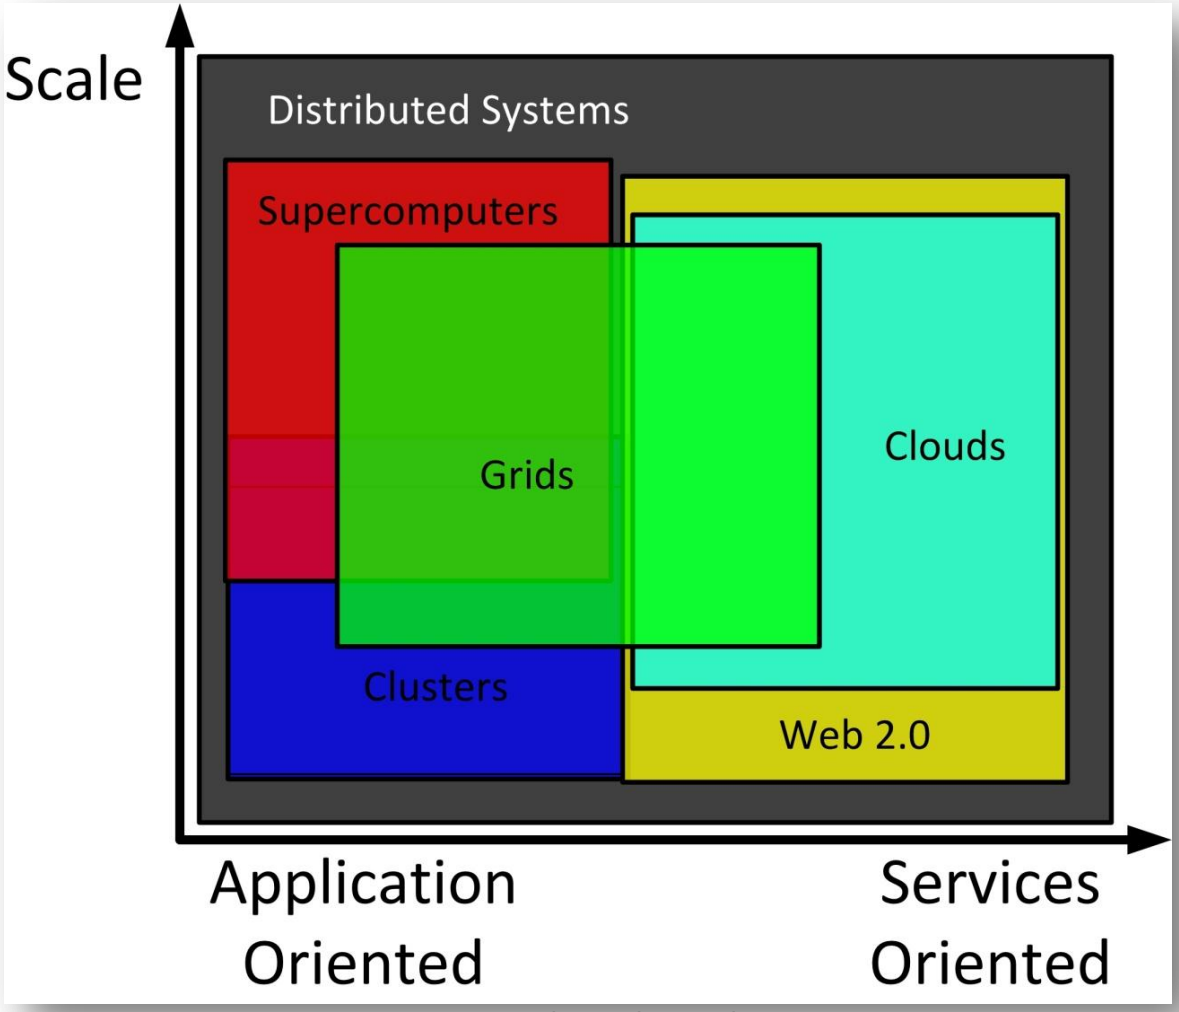

[GCE08] "Cloud Computing and Grid Computing 360-Degree Compared"

Cluster Computing

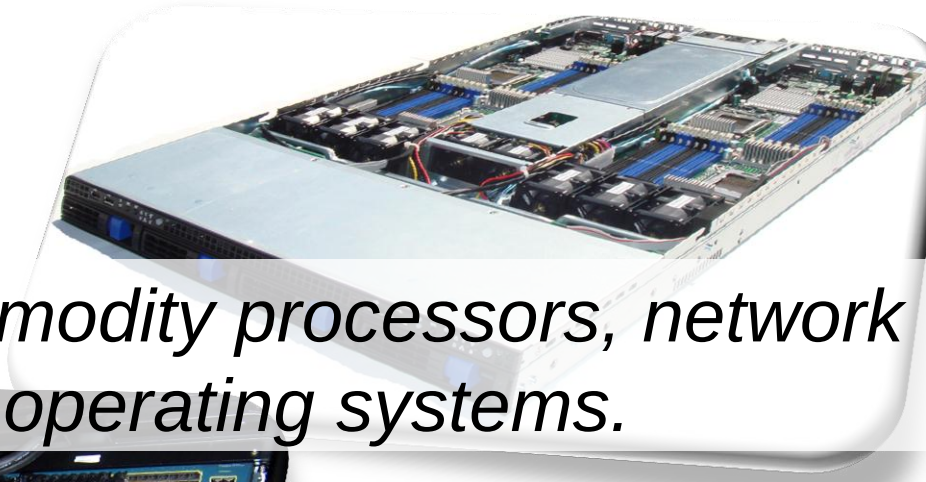

Computer clusters using commodity processors, network interconnects, and operating systems.

Cluster Computing Systems

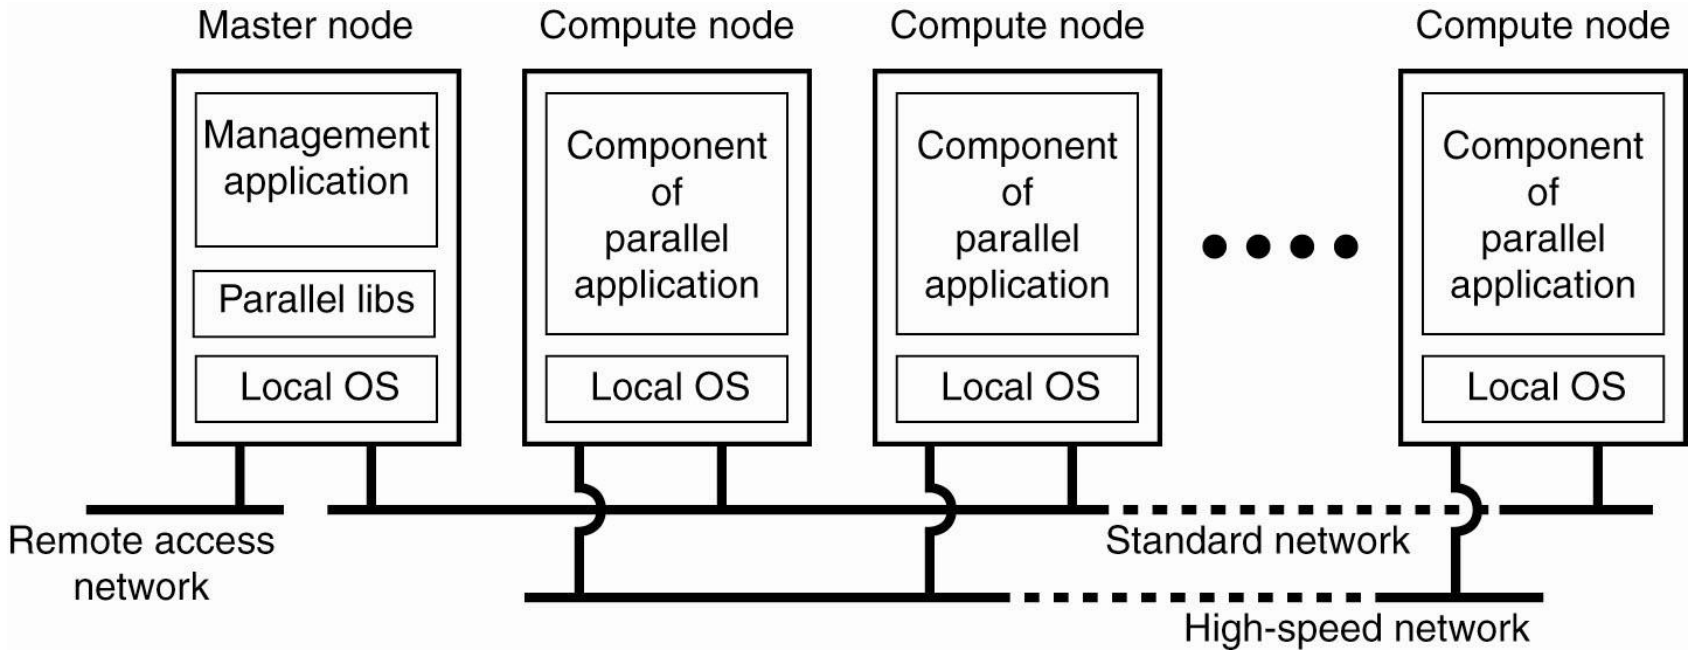

Supercomputing

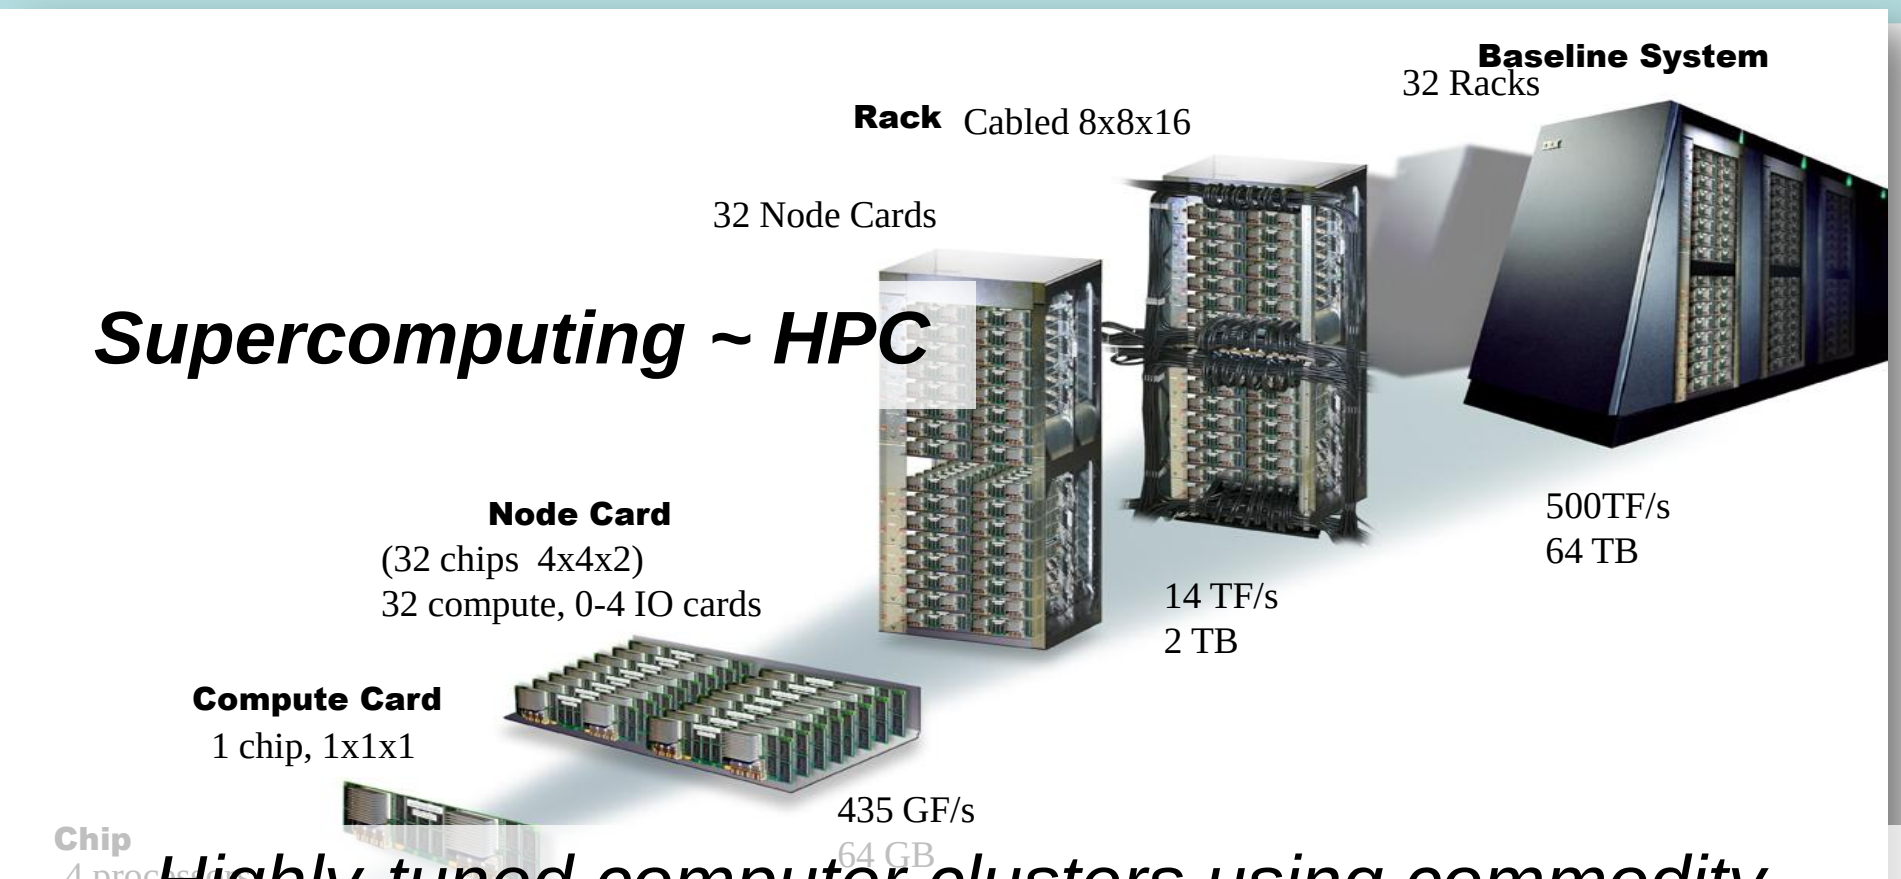

Supercomputing ~ HPC

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

Top 10 Supercomputers from Top500

- Cray XT4 & XT5
  - Jaguar #1
  - Kraken #3
- IBM BladeCenter Hybrid
  - Roadrunner #2
- IBM BlueGene/L & BlueGene/P
  - Jugene #4
  - Intrepid #8
  - BG/L #7
- NUDT (GPU based)
  - Tianhe-1 #5
- SGI Altix ICE
  - Plaiedas #6
- Sun Constellation
  - Ranger #9
  - Red Sky #10

Grid Computing

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

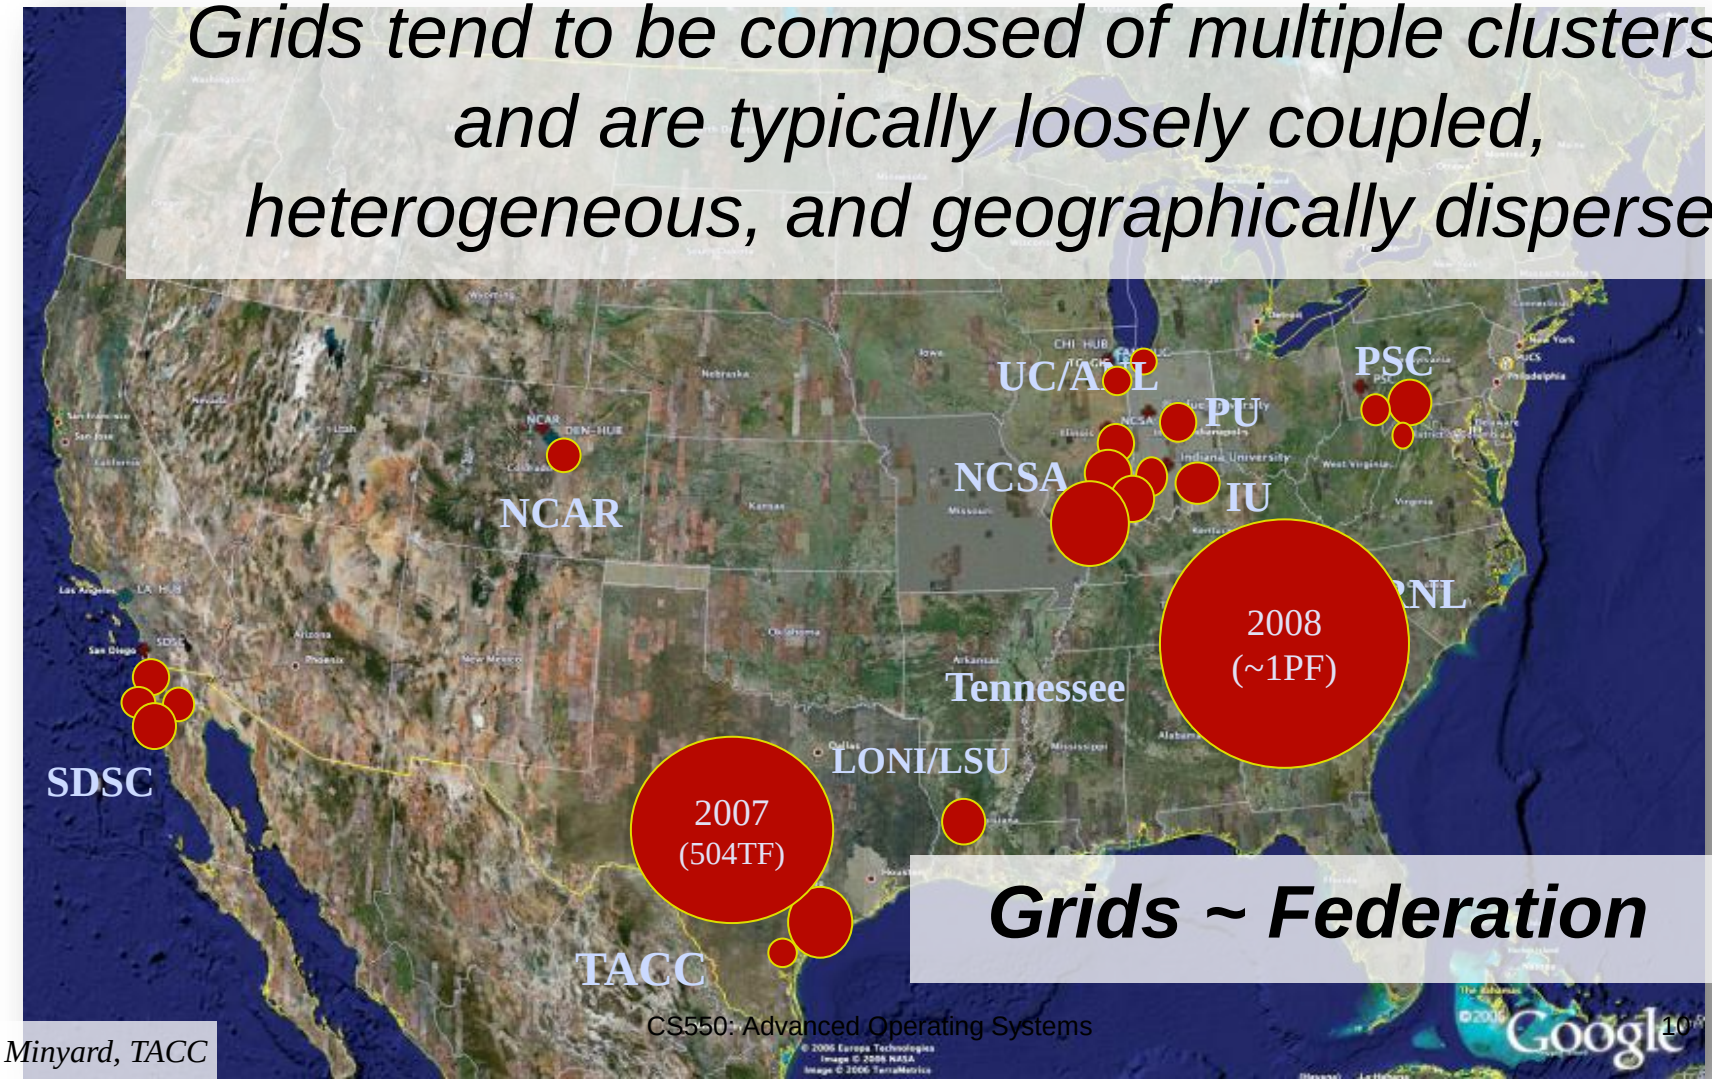

Grids ~ Federation

Tommy Minyard, TACC

Major Grids

- TeraGrid (TG)
  - 200K-cores across 11 institutions and 22 systems over the US
- Open Science Grid (OSG)
  - 43K-cores across 80 institutions over the US
- Enabling Grids for E-sciencE (EGEE)
- LHC Computing Grid from CERN
- Middleware
  - Globus Toolkit
  - Unicore

Cloud Computing: An Emerging Paradigm

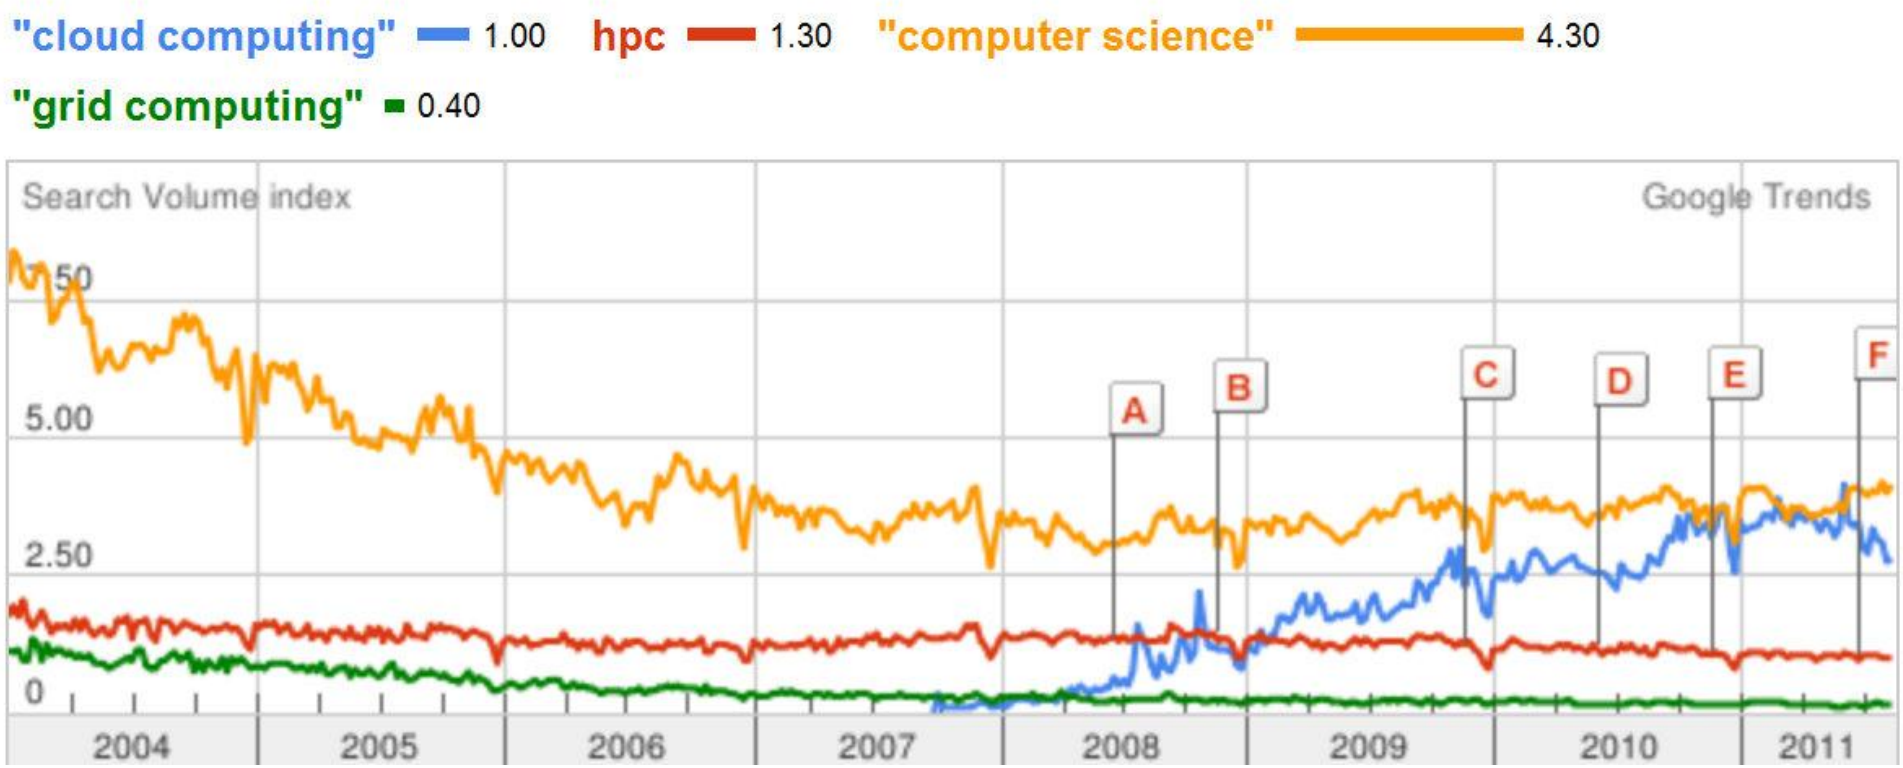

Cloud Computing

- *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

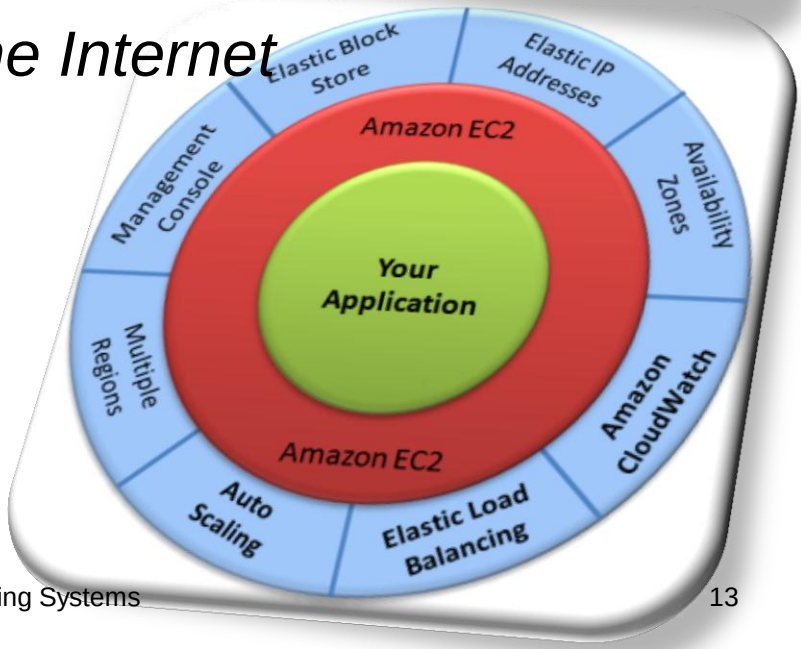

Magellan + DOE's Advanced Network Initiative

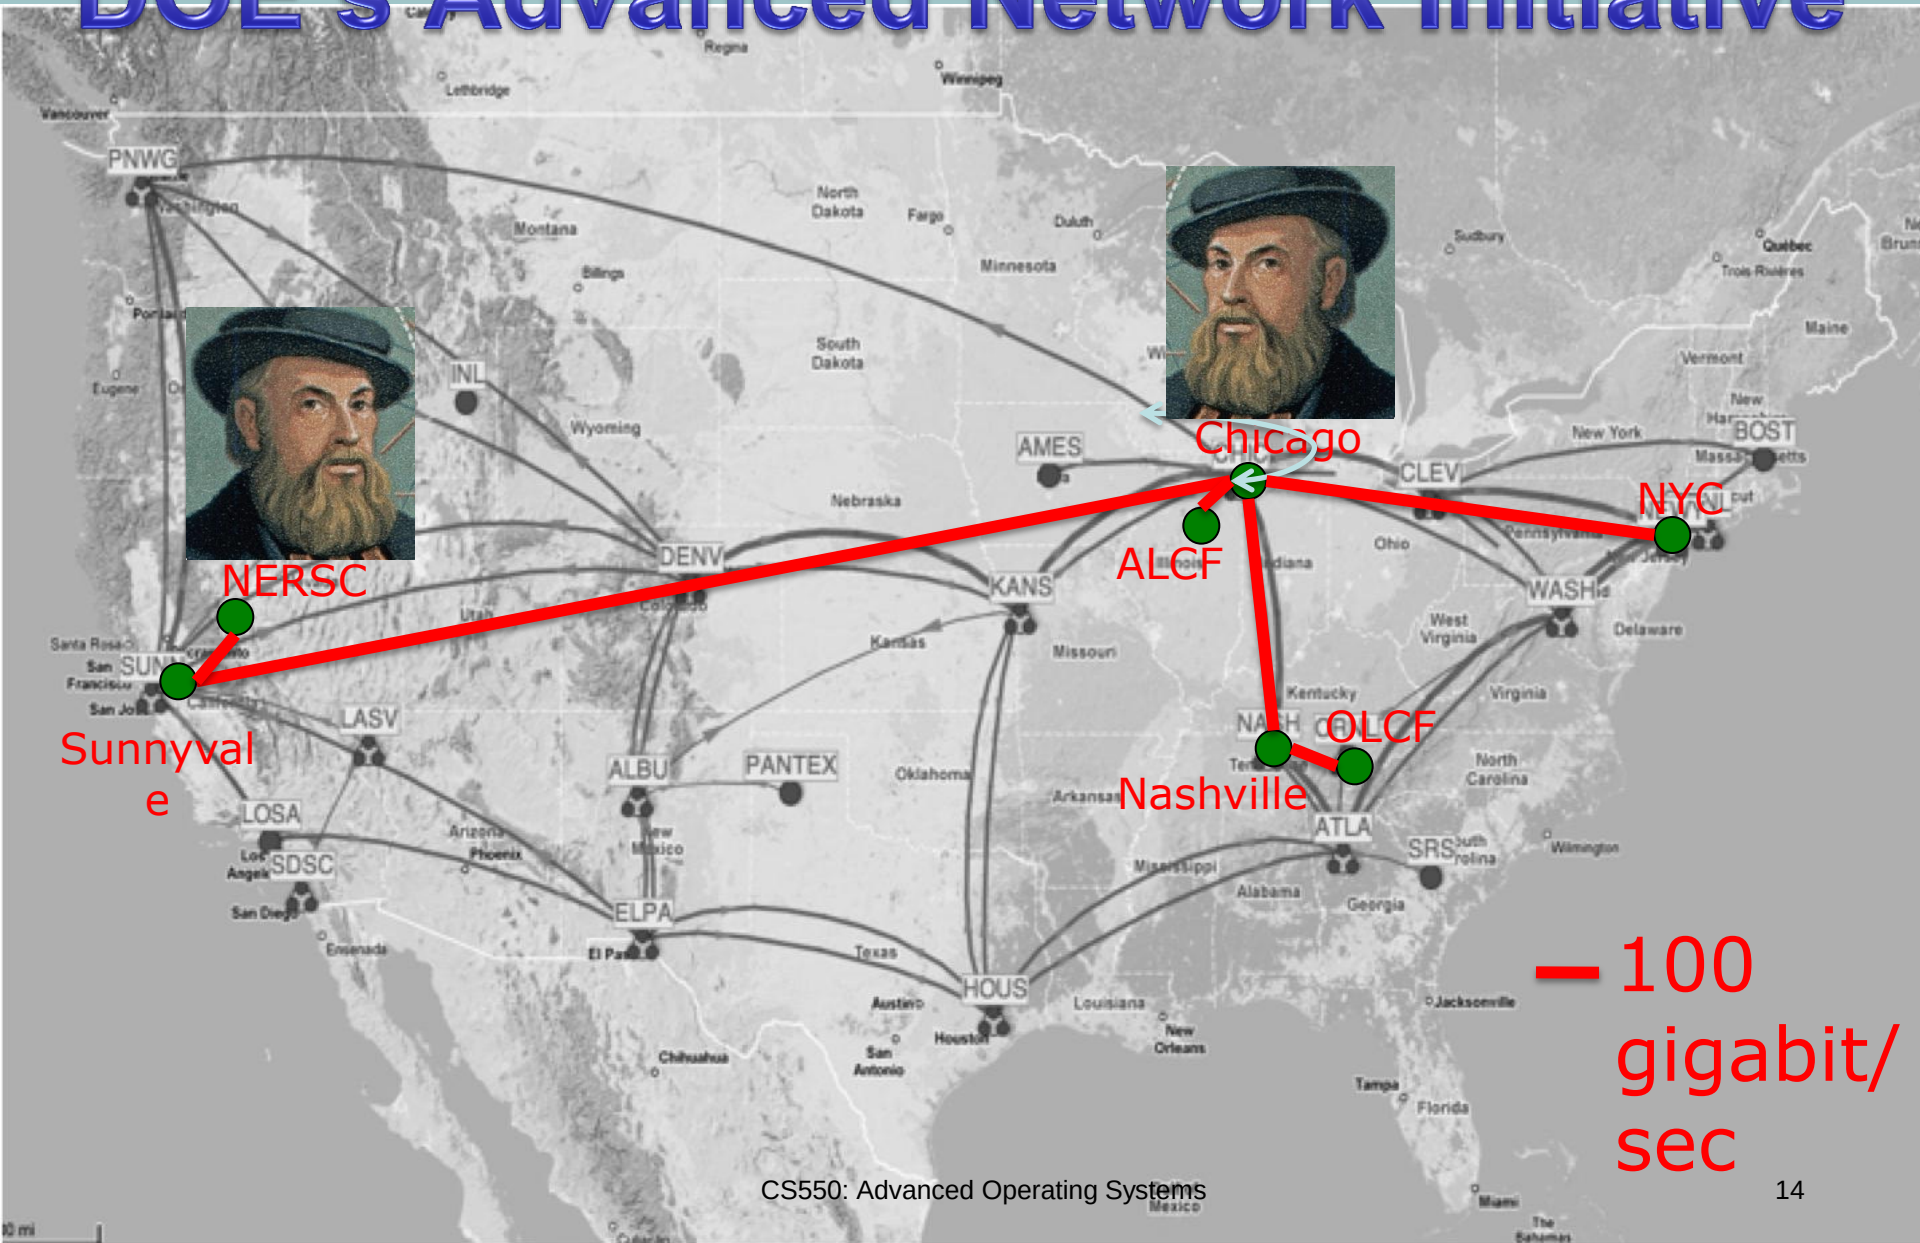

Major Clouds

- Industry
  - Google App Engine
  - Amazon
  - Windows Azure
  - Salesforce
- Academia/Government
  - Magellan
  - FutureGrid
- Opensource middleware
  - Nimbus
  - Eucalyptus
  - OpenNebula

Distributed Information Systems

- Transaction processing systems

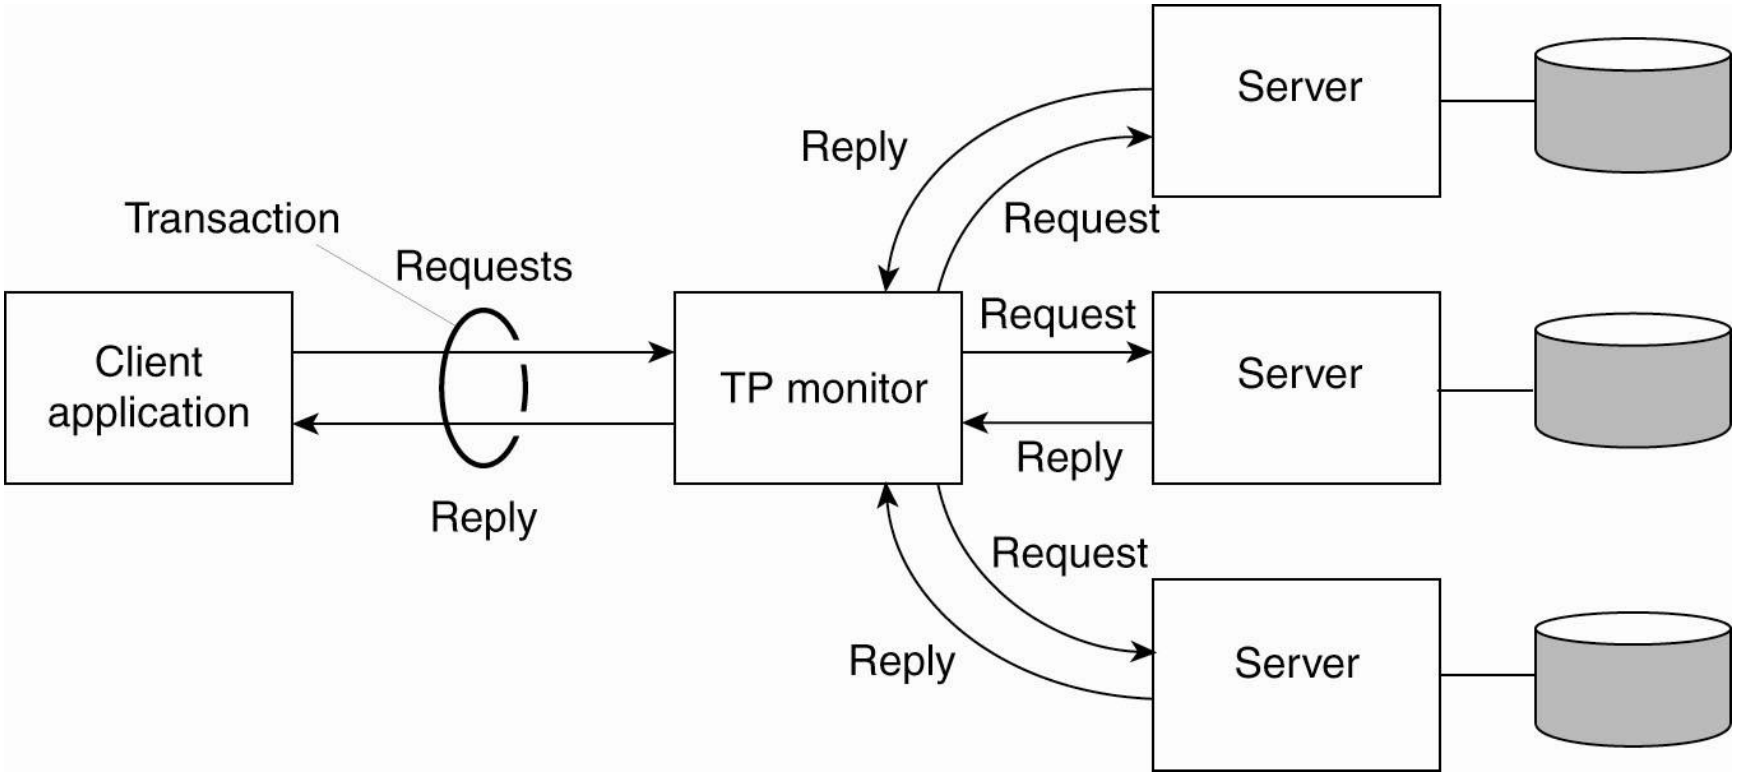

Distributed Pervasive Systems

- Electronic health care systems

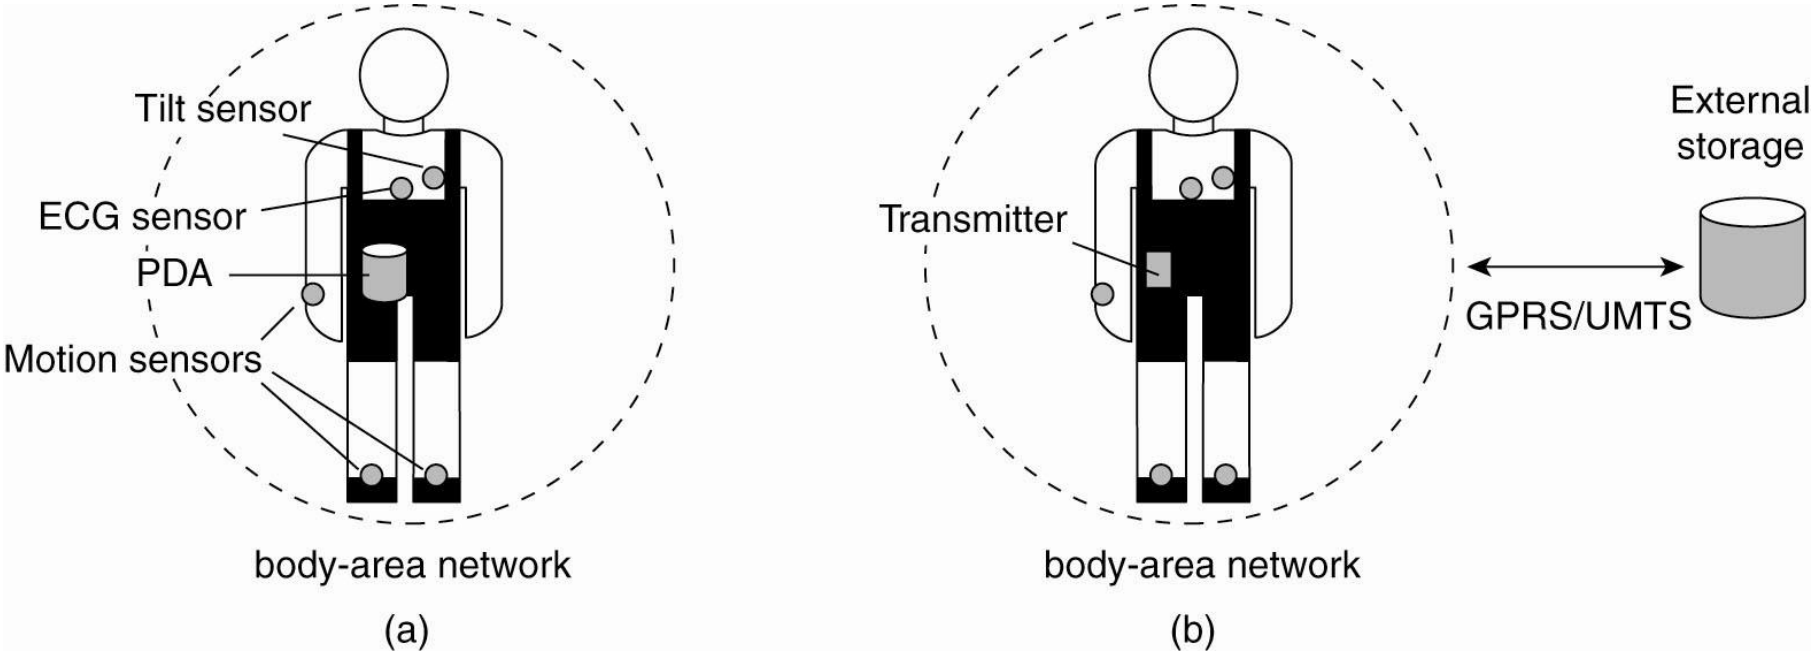

Monitoring a person in a pervasive electronic health care system, using (a) a local hub or (b) a continuous wireless connection.

Distributed Pervasive Systems

- Sensor systems

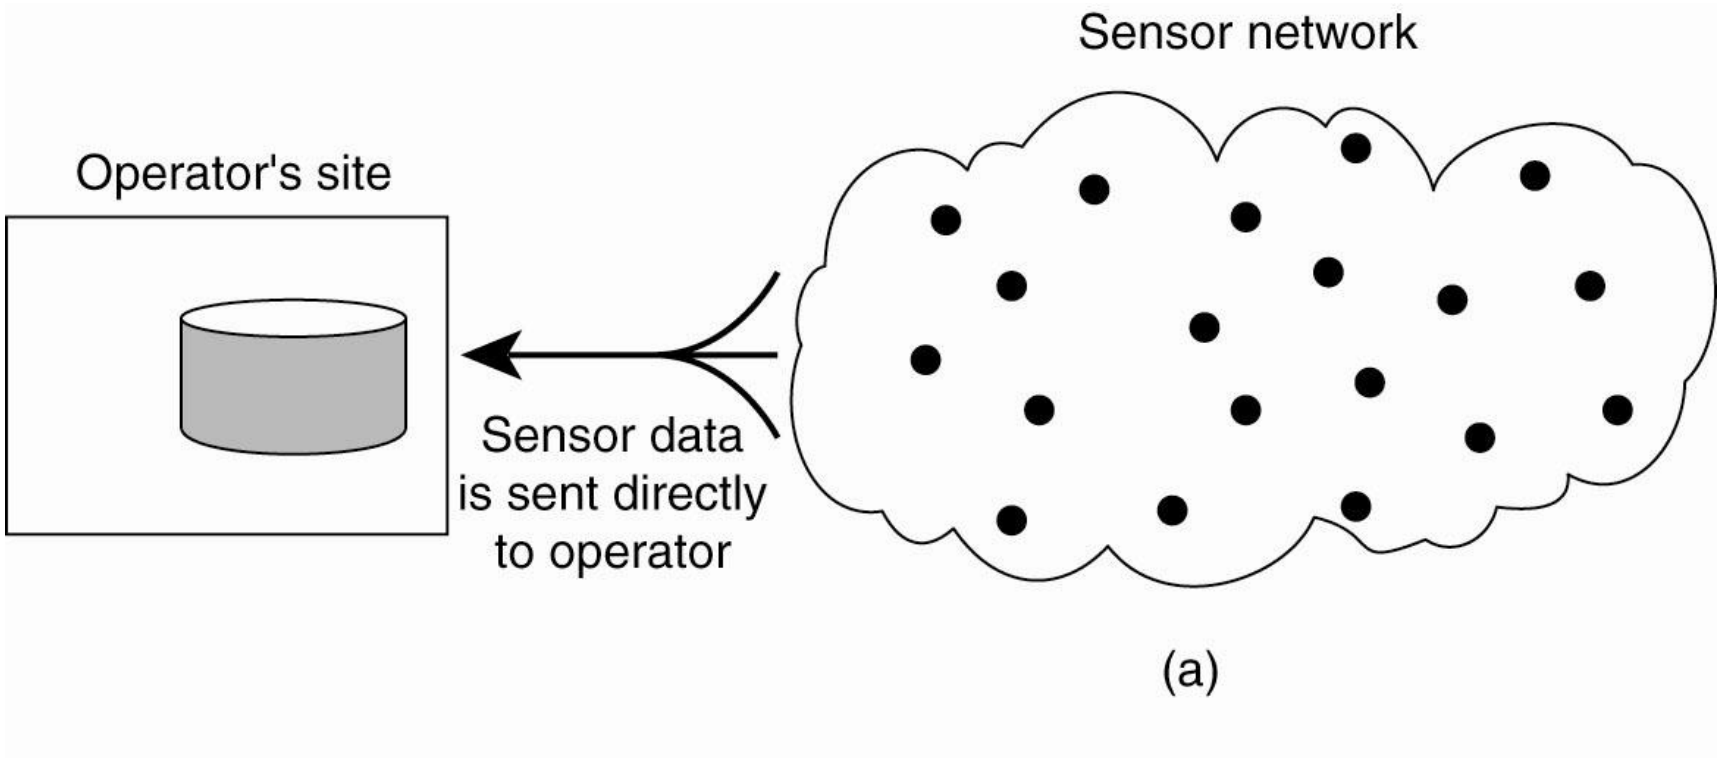

Organizing a sensor network database, while storing and processing data (a) only at the operator's site or ….

Distributed Pervasive Systems

- Sensor systems

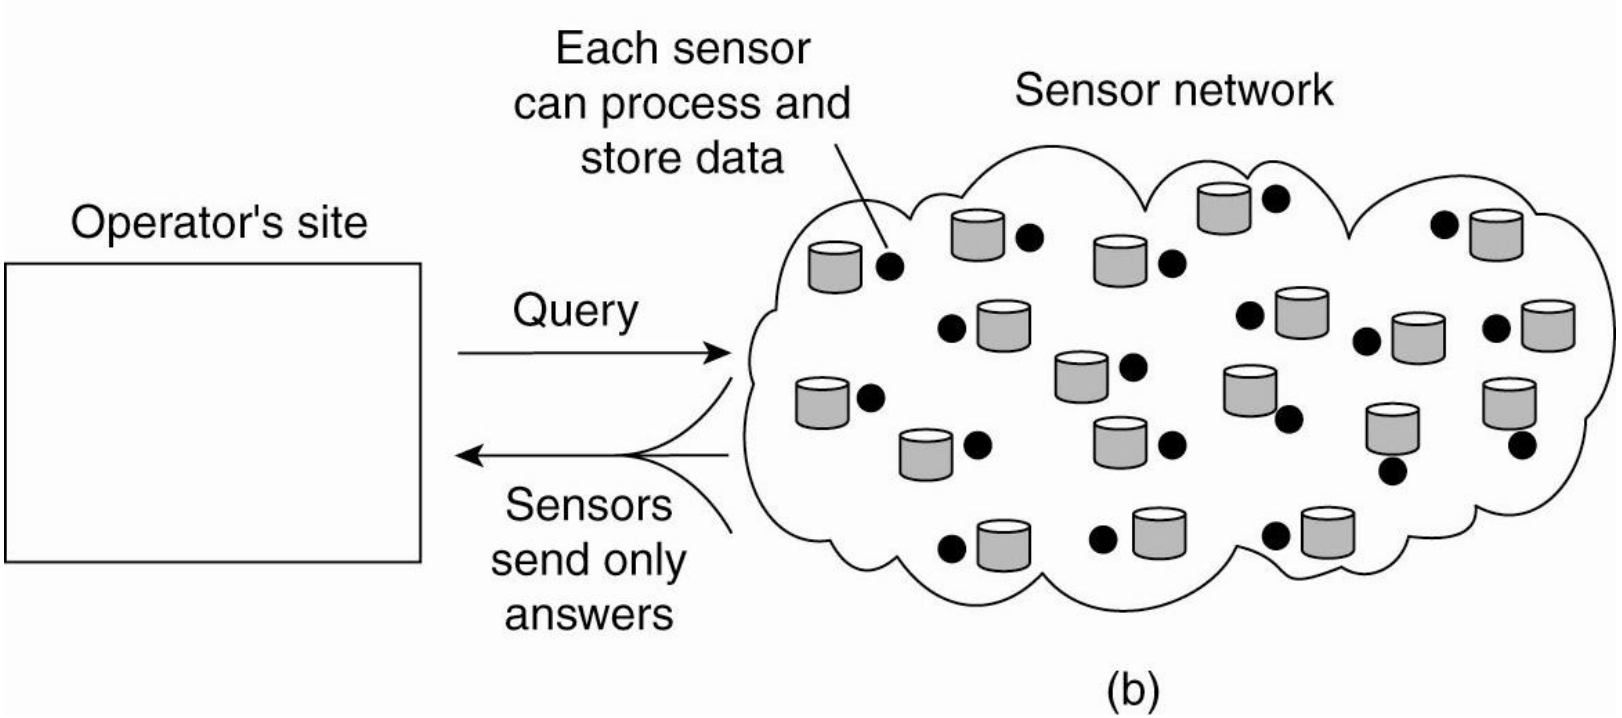

Organizing a sensor network database, while storing and processing data … or (b) only at the sensors.

Distributed vs. Single Systems

- Data sharing
  - Multiple users can access common database, data files,…
- Device/resource sharing
  - Printers, servers, CPUs,….
- Communication
  - Communication with other machines…
- Flexibility
  - Spread workload to different & most appropriate machines
- Extensibility
  - Add resources and software as needed

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

Key Characteristics of Distributed Systems

- Support for resource sharing
- Openness
- Concurrency
- Scalability
- Fault tolerance (reliability)
- Transparence

Resource Sharing

- Share hardware,software,data and information
- Hardware devices
  - Printers,disks,memory,…
- Software sharing
  - Compilers,libraries,toolkits,…
- Data
  - Databases,files,…

Openness

- Definition?
- Hardware extensions
  - Adding peripherals,memory,communication interfaces…
- Software extensions
  - Operating systems features
  - Communication protocols

Concurrency

- In a single system several processes are interleaved
- In distributed systems: there are many systems with one or more processors
  - Many users simultaneously invoke commands or applications
  - Many servers processes run concurrently, each responding to different client request

Scalability

- Scale of system
  - Few PCs servers ->dept level systems->local area networks->internetworked systems->wide are network…
  - Ideally, system and application software should not change as systems scales
- Scalability depends on all aspects
  - Hardware
  - Software
  - networks

Fault Tolerance

- Definition?
- Two approaches:
  - Hardware redundancy
  - Software recovery
- In distributed systems:
  - Servers can be replicated
  - Databases may be replicated
  - Software recovery involves the design so that state of permanent data can be recovered

Transparency in a Distributed System

Transparency	Description
Access	Hide differences in data representation and how a resource is accessed
Location	Hide where a resource is located
Migration	Hide that a resource may move to another location
Relocation	Hide that a resource may be moved to another location while in use
Replication	Means that users do not know whether a replica or a master provides a service.
Concurrency	Hide that a resource may be shared by several competitive users
Failure	Hide the failure and recovery of a resource
Persistence	Hide whether a (software) resource is in memory or on disk

Pitfalls When Developing Distributed Systems

- False assumptions made by first time developer:
  - The network is reliable.
  - The network is secure.
  - The network is homogeneous.
  - The topology does not change.
  - Latency is zero.
  - Bandwidth is infinite.
  - Transport cost is zero.
  - There is one administrator.

Questions