

Data Intensive Distributed Computing

Ioan Raicu
Computer Science Department
Illinois Institute of Technology

CS 595
Hot Topics in Distributed Systems: Data-Intensive Computing
August 25th, 2010

Famous Quotes

The users should be able to focus their attention on the information content of the data, rather than how to discover, access, and use it.

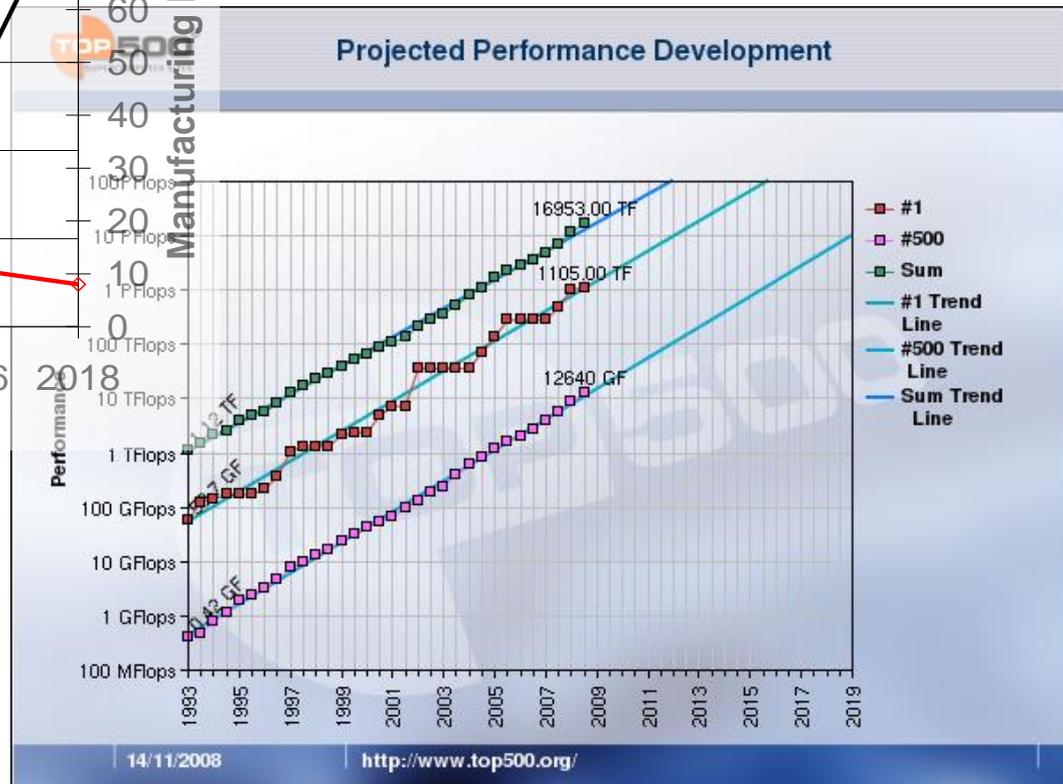
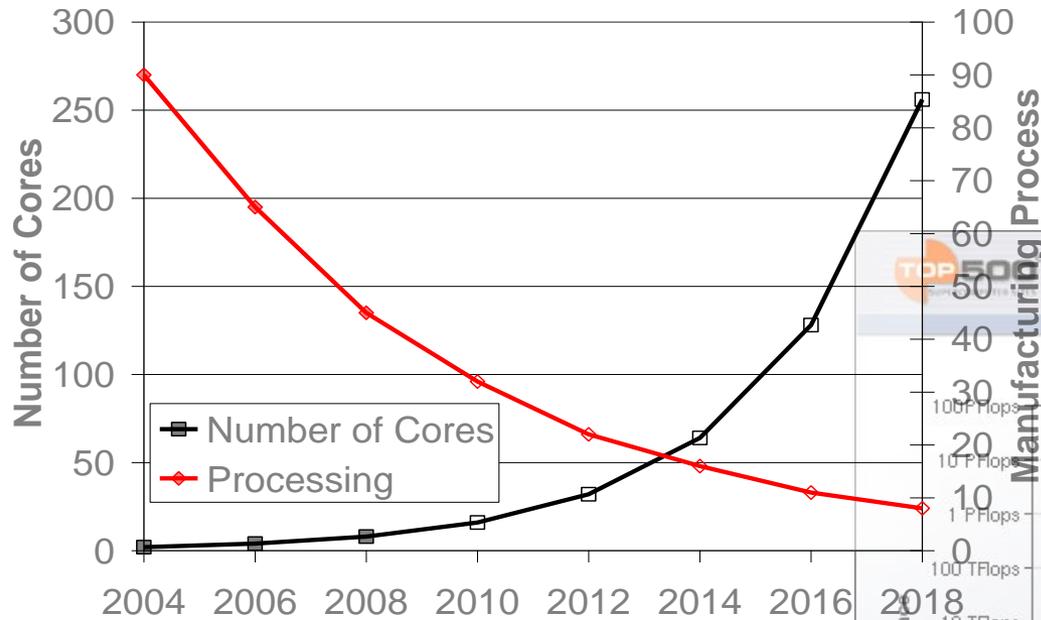
Climate Change Science Program report, 2003

Famous Quotes

A supercomputer is a device for turning compute-bound problems into I/O-bound problems.

Seymour Cray

Projected Growth Trends



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

Top500 Projected Development,

http://www.top500.org/lists/2008/11/performance_development

Growing Storage/Compute Gap

- Local Disk:

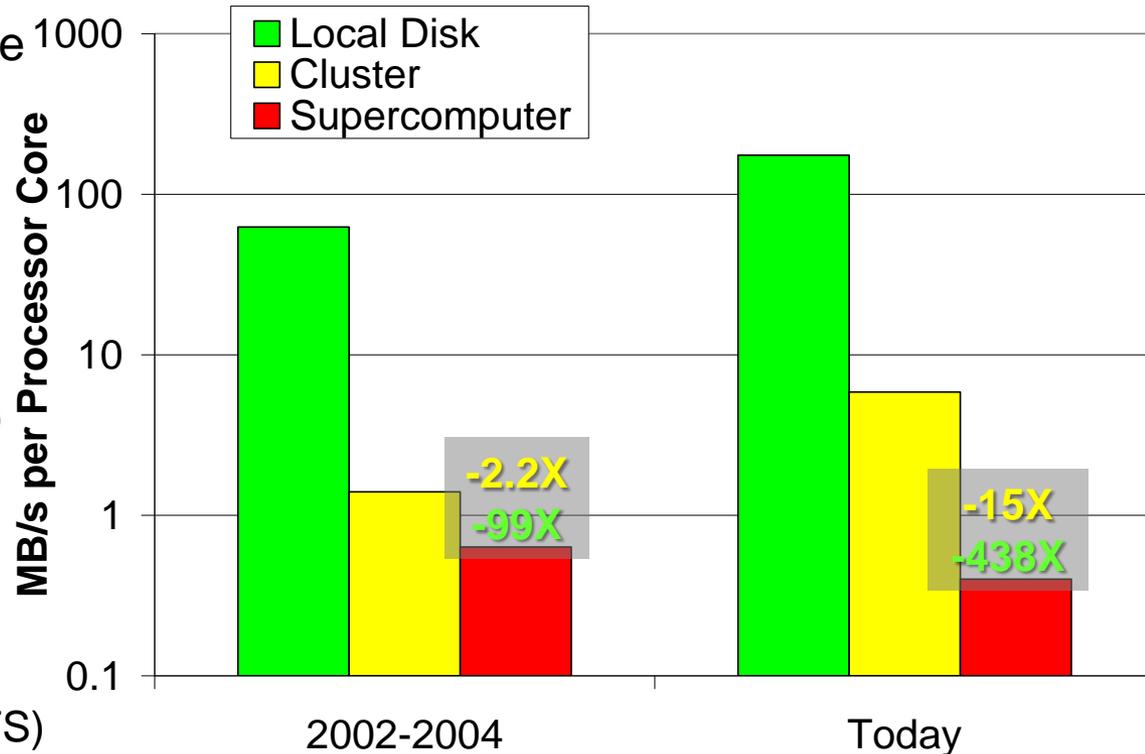
- 2002-2004: ANL/UC TG Site (70GB SCSI)
- Today: PADS (RAID-0, 6 drives 750GB SATA)

- Cluster:

- 2002-2004: ANL/UC TG Site (GPFS, 8 servers, 1Gb/s each)
- Today: PADS (GPFS, SAN)

- Supercomputer:

- 2002-2004: IBM Blue Gene/L (GPFS)
- Today: IBM Blue Gene/P (GPFS)



State of the Art: Storage Systems

- Segregated storage and compute
 - NFS, GPFS, PVFS, Lustre
 - Batch-scheduled systems: Clusters, Grids, and Supercomputers
 - Programming paradigm: HPC, MTC, and HTC
- Co-located storage and compute
 - HDFS, GFS
 - Data centers at Google, Yahoo, and others
 - Programming paradigm: MapReduce
 - Others from academia: Sector, MosaStore, Chirp

State of the Art: Storage Systems

- Segregated storage and compute

- NFS, GPFS, PVFS, Lustre

- Batch-scheduled
Supercomputers

- Programming pa

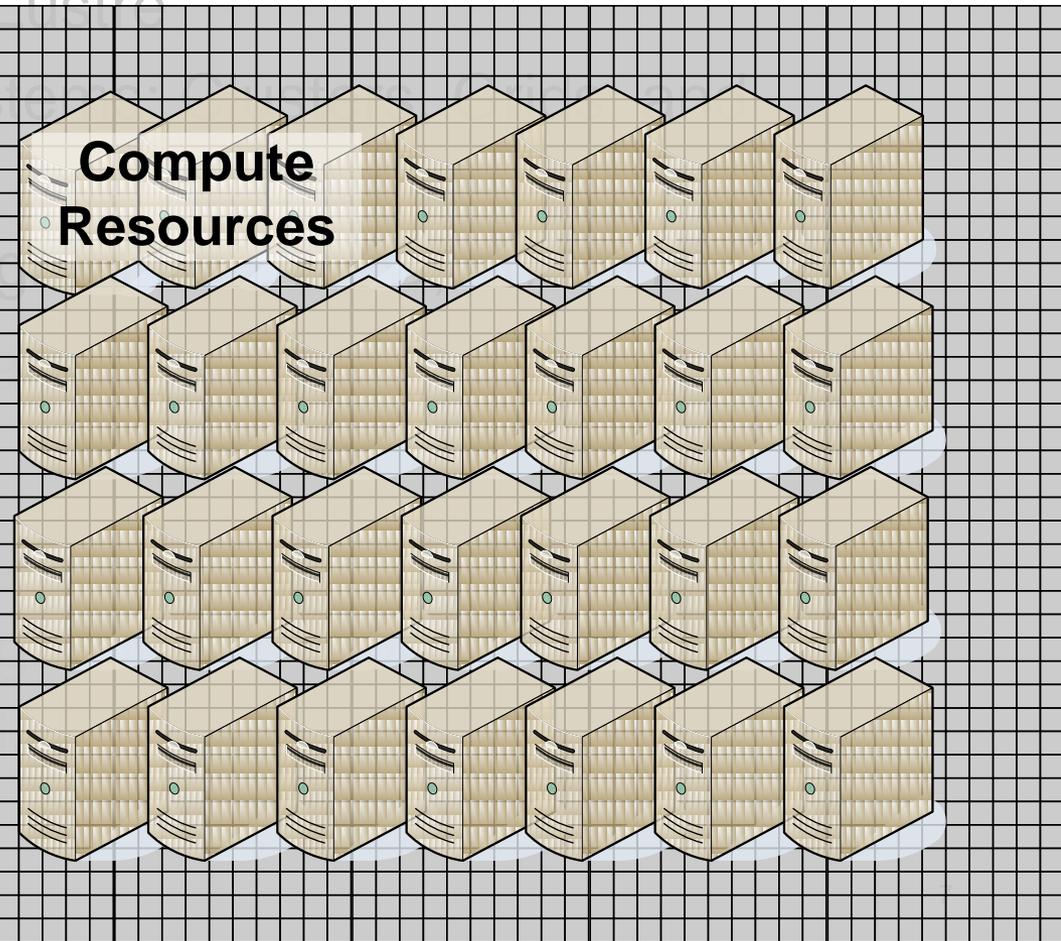
- Located stora

- Data centers at

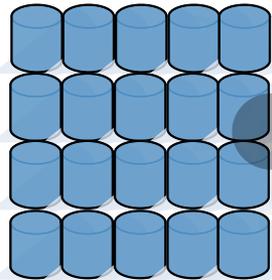
- Programming pa

- Others from aca

**Network
Fabric**



NAS



Network Link(s)

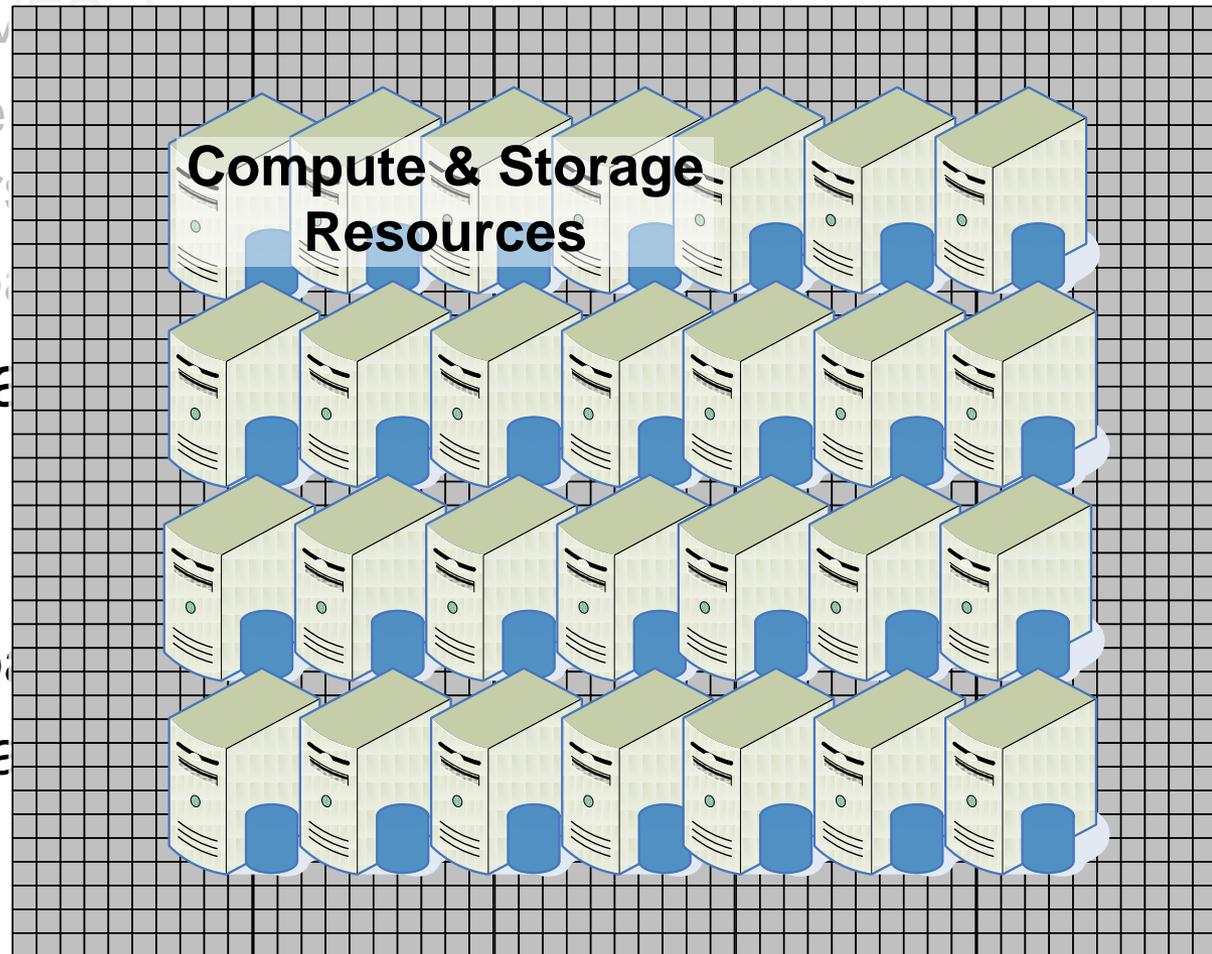
State of the Art: Storage Systems

- Segregated storage and compute
 - NFS, GPFS, PVFS, Lustre
 - Batch-scheduled systems: Clusters, Grids, and Supercomputers
 - Programming paradigm: HPC, MTC, and HTC
- Co-located storage and compute
 - HDFS, GFS
 - Data centers at Google, Yahoo, and others
 - Programming paradigm: MapReduce
 - Others from academia: Sector, MosaStore, Chirp

State of the Art: Storage Systems

- Segregated storage and compute
 - NFS, GPFS, PVFS
 - Batch-scheduling
 - Supercomputers
 - Programming paradigms
- Co-located storage
 - HDFS, GFS
 - Data centers at scale
 - Programming paradigms
 - Others from academia

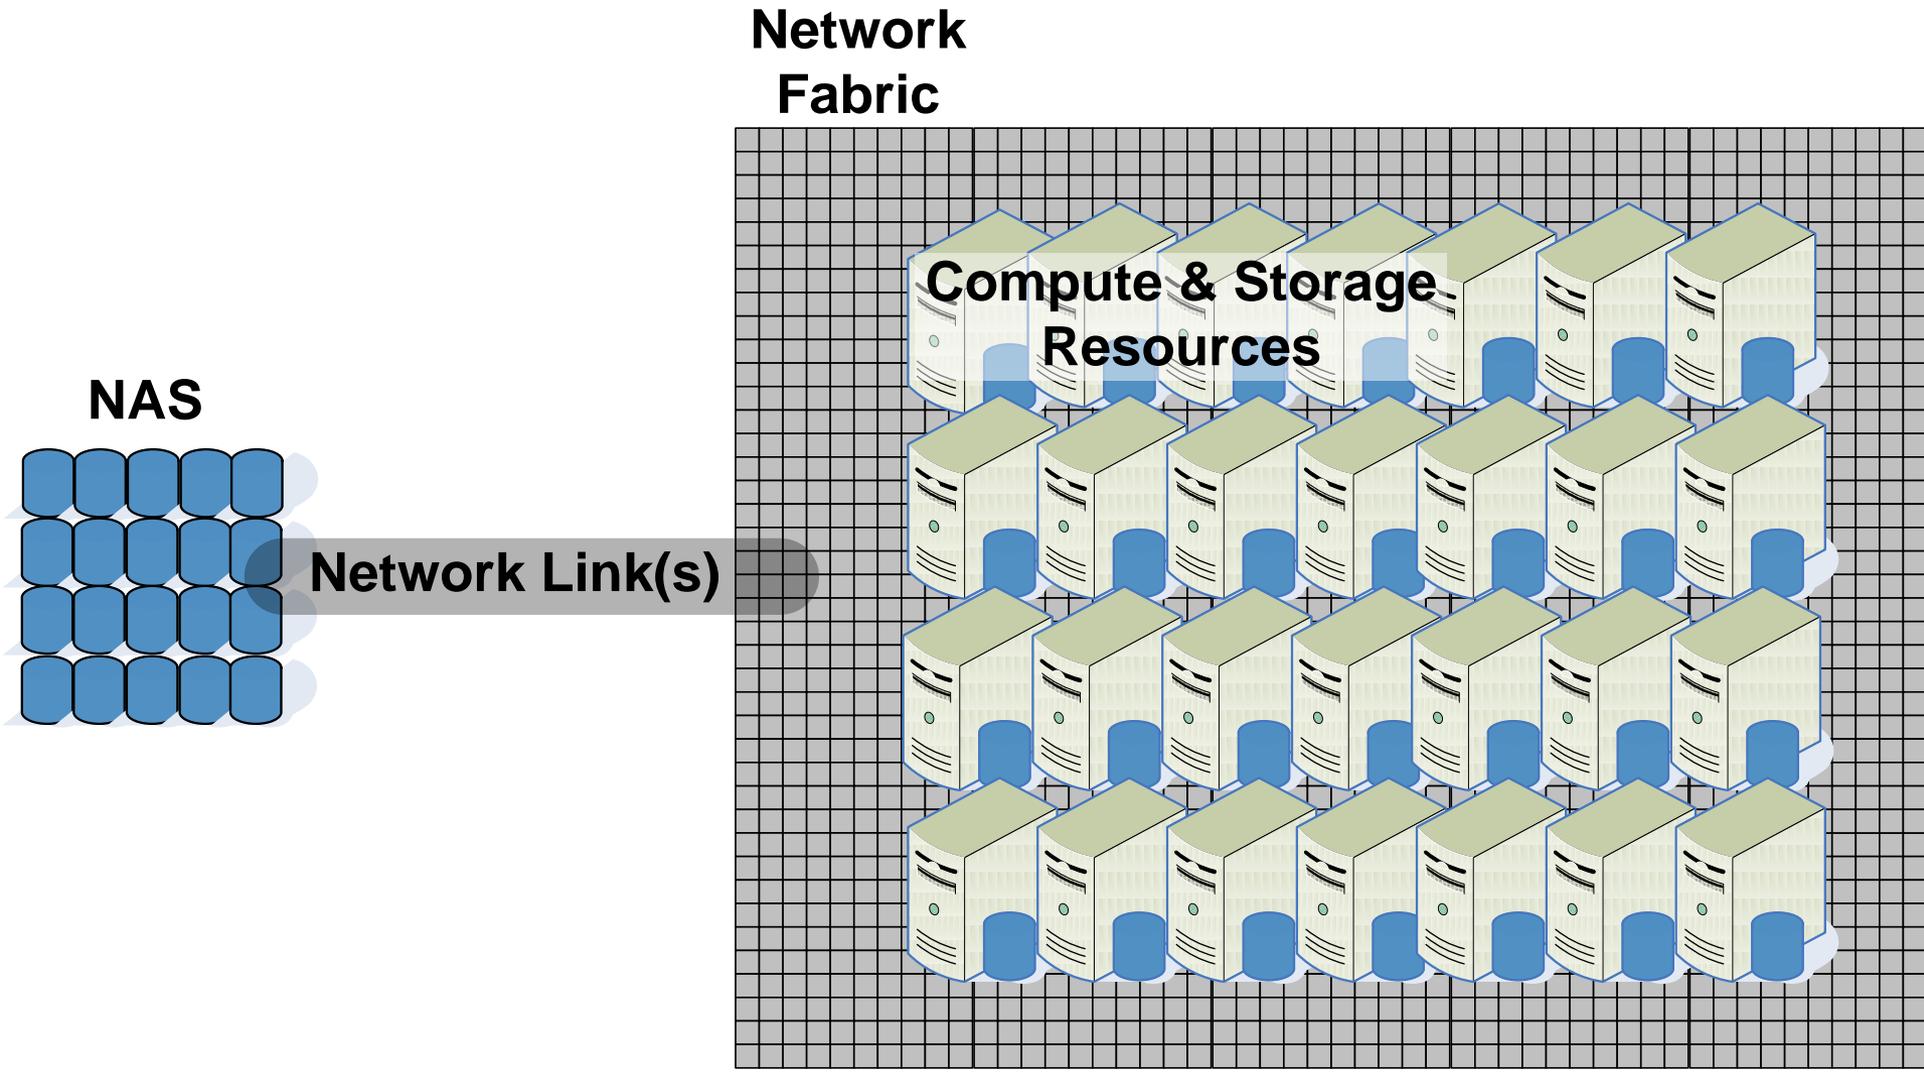
Network Fabric



Question

What if we could combine the scientific community's existing programming paradigms, but yet still exploit the data locality that naturally occurs in scientific workloads?

Combine State of the Art Systems



Questions

