Enabling Data Intensive Science with PetaShare

Tevfik Kosar

Center for Computation & Technology
Louisiana State University

April 6, 2007

Center for Computation & Technology

AT LOUISIANA STATE UNIVERSITY

The Imminent Data “deluge”

Scientific data outpaced Moore’s Law!

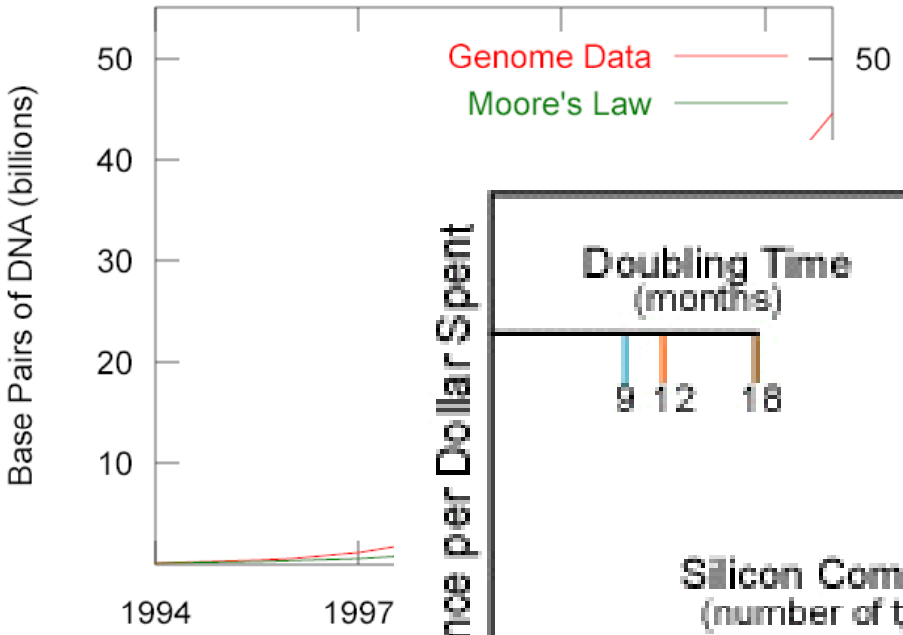

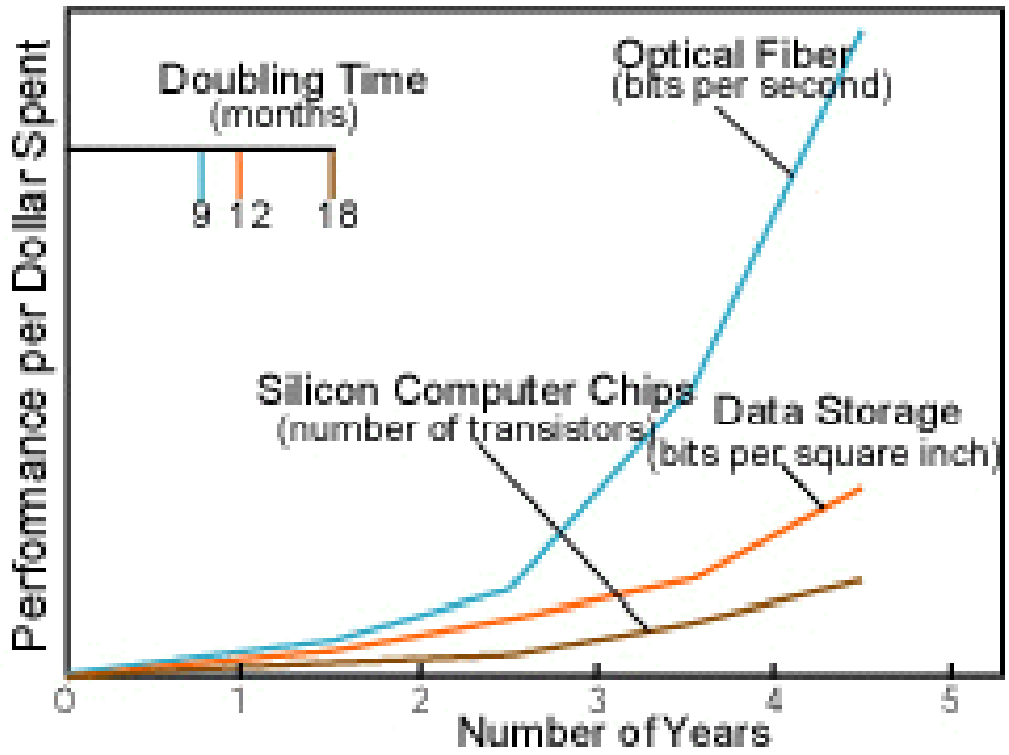

Demand for data in all areas of science!

Application	Area	Data Volume
VISTA	Astronomy	100 TB/year
		250 TB/year
		500 TB/year
		1000 TB/year
cs		1000 TB/year
		5000 TB/year

LONI

- Since September 2004, the State of Louisiana has committed $50M for a state-wide optical network.
- 40Gb/sec bandwidth
- Spanning 6 Universities and 2 Health Centers:
  - LSU
  - Latech
  - UL-Lafayette
  - Tulane
  - UNO
  - Southern University
  - LSU Health Centers in
    - New Orleans
    - Shreveport

- 112 processor IBM P5 servers being deployed at each site
- 540 processor Linux clusters will follow
- 100 TFlops in a couple of years

DONE?..

- We will have one of the
  - Fastest networks
  - Most powerful computational grids

  in the world..
- But this solves only part of the problem!
- Researchers at these institutions still not be able to share and even process their own data

GRID today

The Leading Source for Global News and Information from the evolving Grid ecosystem,
including Grid, SOA, Virtualization, Storage, Networking and Service-Oriented IT

Applications:

NSF Funds LSU $1 Million for PetaShare Development

The National Science Foundation (NSF) recently funded Louisiana State University (LSU) $1 million for the development of PetaShare, which is seen as "a system might become an important testbed for future grids, and a leading site in next-generation petascale research."

The unbounded increase in the size of data generated by scientific applications necessitates collaboration and sharing among the nation's education and research institutions. Simply purchasing high capacity, high performance storage systems and adding them to the existing infrastructure of the collaborating institutions does not solve the underlying and highly challenging data handling problem. Scientists are compelled to spend a great deal of time and energy on solving basic data-handling issues, such as the physical location of data, how to access it, and/or how to move it to visualization and/or compute resources for further analysis.

LSU assistant professor Tevfik Kosar and his team aim to develop an innovative distributed data archival, analysis and visualization cyberinfrastructure for data intensive collaborative research, which they call PetaShare. PetaShare will enable transparent handling of underlying data sharing, archival and retrieval mechanisms, and will make data available to scientists for analysis and visualization on demand. PetaShare will enable scientists to focus on their primary research problem, assured that the underlying infrastructure will manage the low-level data handling issues.

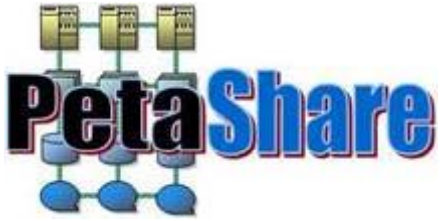

- Goal: enable domain scientists to focus on their primary research problem, assured that the underlying infrastructure will manage the low-level data handling issues.
- Novel approach: treat data storage resources and the tasks related to data access as first class entities just like computational resources and compute tasks.
- Key technologies being developed: data-aware storage systems, data-aware schedulers (i.e. Stork), and cross-domain meta-data scheme.
- Provides and additional 200TB disk, and 400TB tape storage

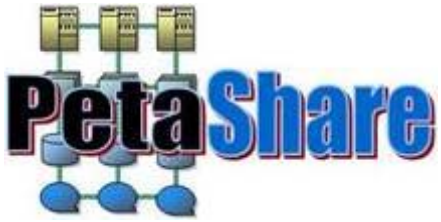

- PetaShare exploits 40 Gb/sec LONI connections between 5 LA institutions: LSU, LaTech, Tulane, ULL, and UNO.
- PetaShare links more than fifty senior researchers and two hundred graduate and undergraduate research students from ten different disciplines to perform multidisciplinary research.
- Application areas supported by PetaShare include coastal and environmental modeling, geospatial analysis, bioinformatics, medical imaging, fluid dynamics, petroleum engineering, numerical relativity, and high energy physics.

Participating institutions in the PetaShare project, connected through LONI. Sample research of the participating researchers pictured (i.e. biomechanics by Kodiyalam & Wischusen, tangible interaction by Ullmer, coastal studies by Walker, and molecular biology by Bishop).

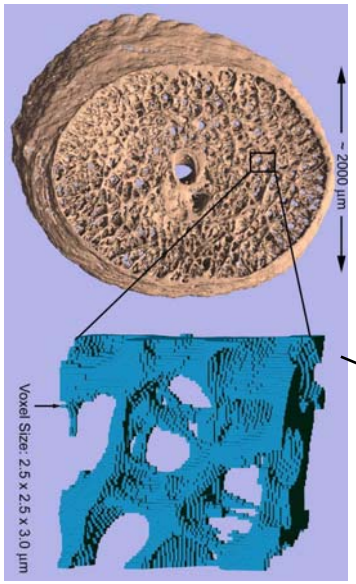

High Energy Physics
Biomedical Data Mining
LATECH

Coastal Modeling
Petroleum Engineering
Computational Fluid Dynamics
Synchrotron X-ray Microtomography
LSU

UNO
Biophysics

ULL
Geology
Petroleum Engineering

TULANE
Molecular Biology
Computational Cardiac Electrophysiology

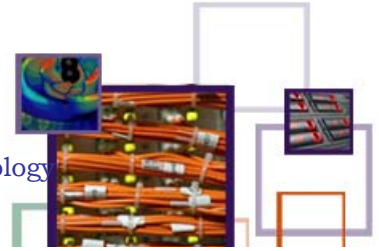

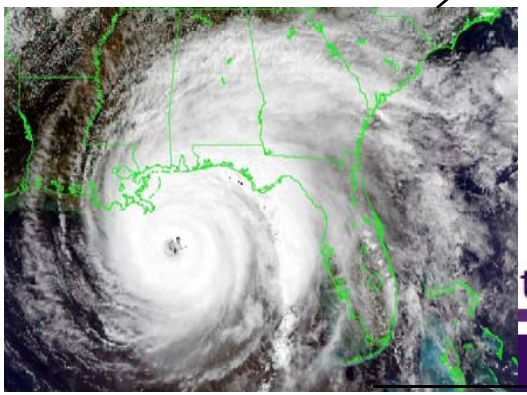

tation & Technology

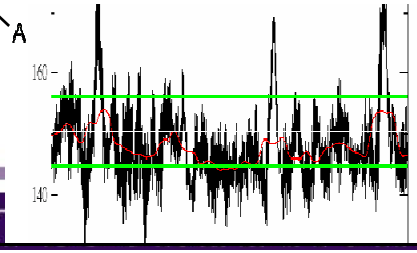

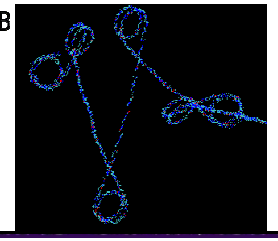

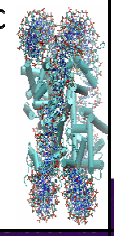

PetaShare Science Drivers

Coastal Studies

- Walker, Levitan, Mashriqui, Twilley (LSU)
- The Earth Scan Lab: with its three antennas, it captures 40GB of data from six satellites each day. (➔ 15 TB/year)
- Hurricane Center
  - Storm surge modeling, hurricane track prediction
- Wetland Biochemistry Institute
  - Coastal Ecosystem preservation
- SCOOP data archive

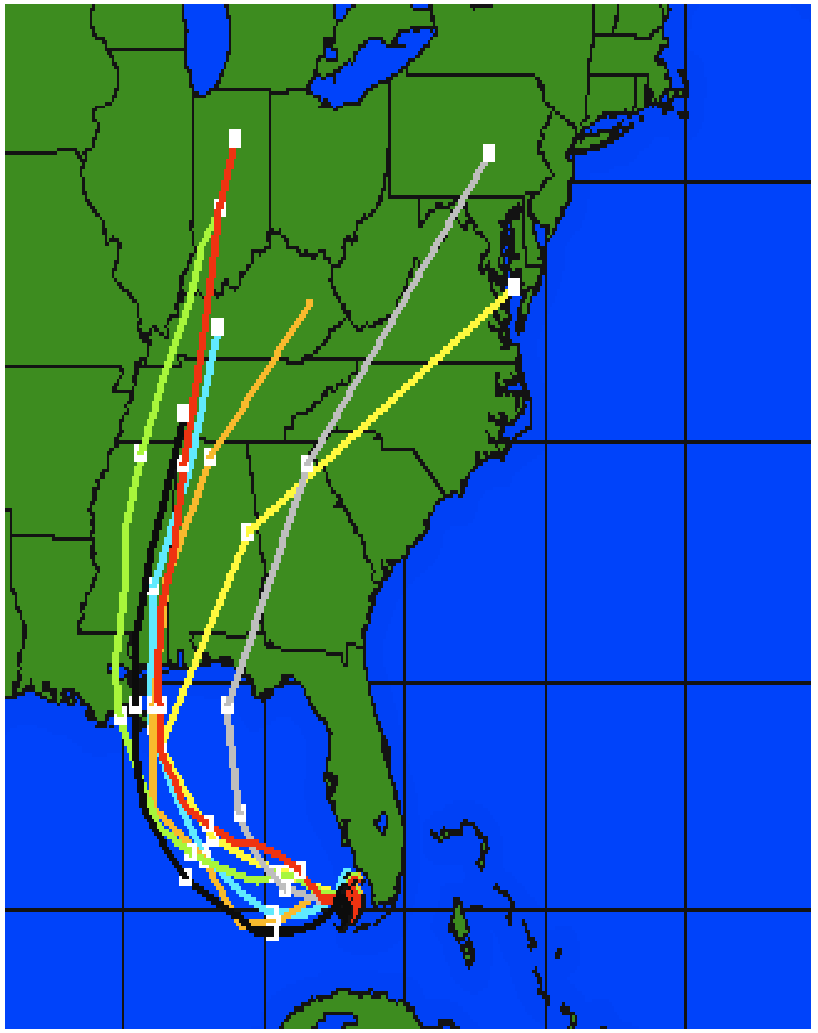

Petroleum Engineering

- White, Allen, Lei et al. (LSU, ULL, SUBR)
- UCoMS project – reservoir simulation and uncertainty analysis
- 26M simulations, each generating 50MB of data ➔ 1.3 PB of data total

- Drilling processing and real-time monitoring is data-intensive as well ➔ real-time visualization and analysis of TB's of streaming data

Computational Fluid Dynamics

- Acharya et al. (LSU)
- Focusing on simulation of turbulent flows including Direct Numerical Simulations (DNS), Large Eddy Simulations (LES), and Reynolds-Averaged Navier Stokes Simulations (RANS).
- In DNS, ~10,000 instances of flow field must be stored and analyzed, each instance may contain 150M discrete variables. Resulting data set ~ 10 TB.

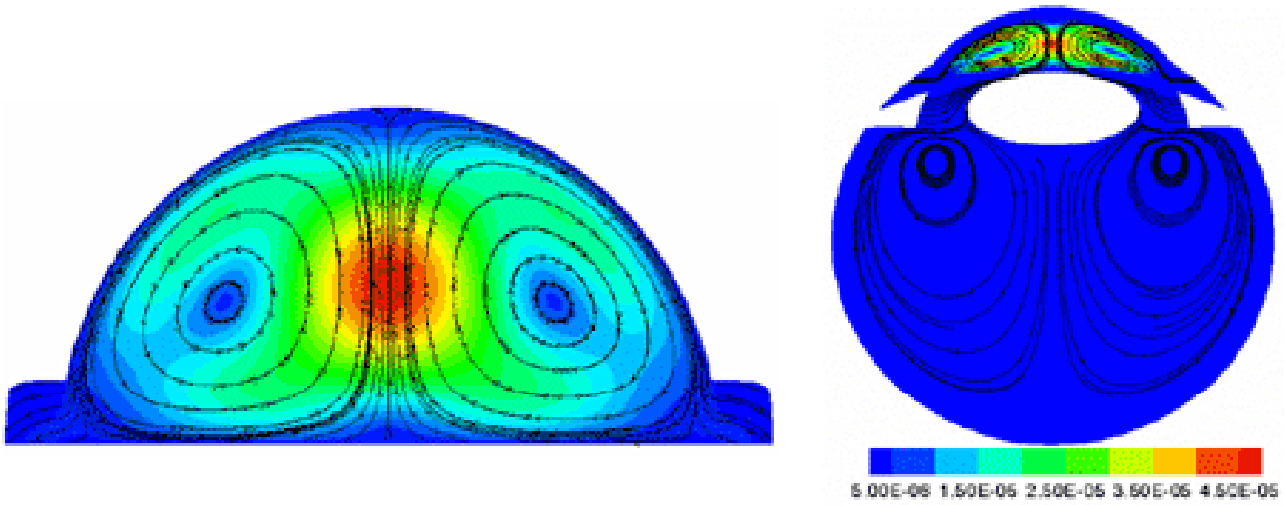

Molecular Biology

•Winters-Hilt (UNO)
•Biophysics and molecular biology – gene structure analysis
•Generates several terabytes of channel current measurements per month
•Generated data being sent to UC-Santa Cruz, Harvard and other groups

•Bishop (Tulane)
•Study the structure and dynamics of nucleosomes using all atom molecular dynamics simulations
•Each simulation requires 3 weeks of run time on a 24-node cluster, and 50-100 GB of storage ➔ 1-2 TB data per year

* Both access to the Genome database but separately!

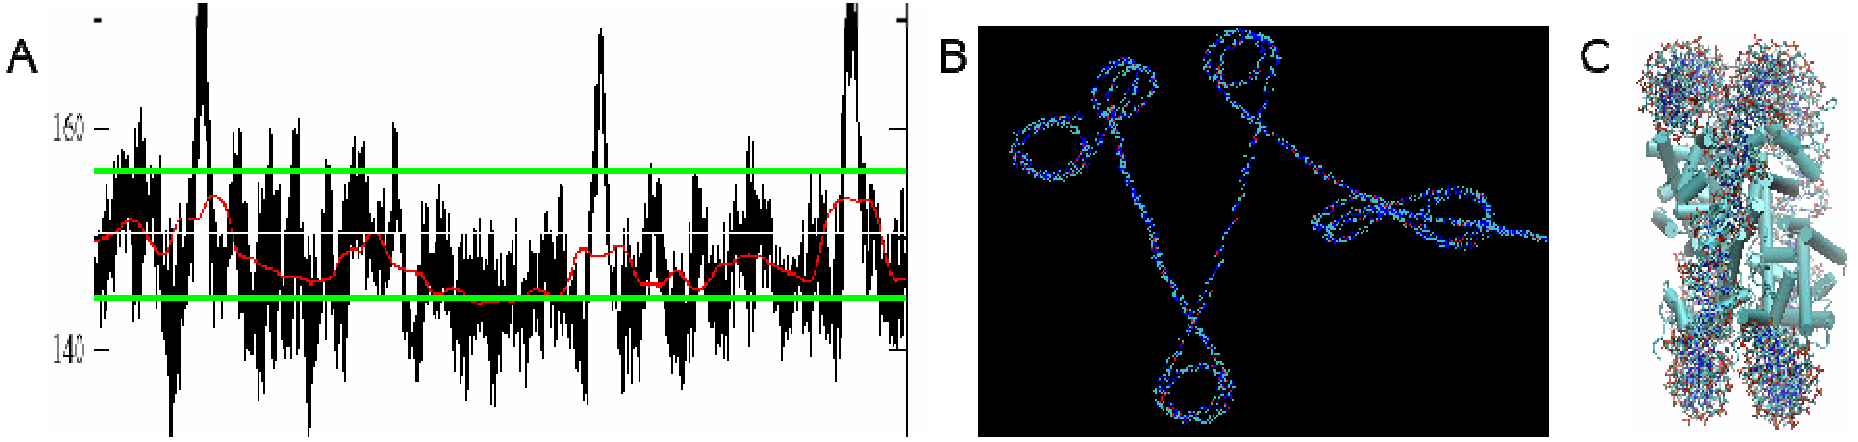

And Others…

- Numerical Relativity - Seidel et al (LSU)
- High Energy Physics – Greenwood, McNeil (LaTech, LSU)
- Computational Cardiac Electrophysiology – Trayanova (Tulane)
- Synchrotron X-ray Microtomography – Wilson, Butler (LSU)
- Bio Data Mining – Dua (LaTech)

CS Research

- Distributed Data Handling (Kosar)
- Grid Computing (Allen, Kosar)
- Visualization (Hutanu, Karki)
- Data Mining (Dua, Abdelguerfi)
- Database Systems (Triantaphyllou)

People involved with PetaShare Development and Usage

	Senior Personnel			Postdocs			Graduate Students			Undergraduates		
	Women	Minority	Total	Women	Minority	Total	Women	Minority	Total	Women	Minority	Total
LaTech			2						5			2
LSU	3		27			5	5	4	47	2	2	24
Tulane	1		3			2	3	5	11		1	3
ULL	7	4	18			8	13	5	31			4
UNO			7				5	4	32	3	8	17
Total	11	4	57			15	26	18	126	5	11	50

PetaShare Overview

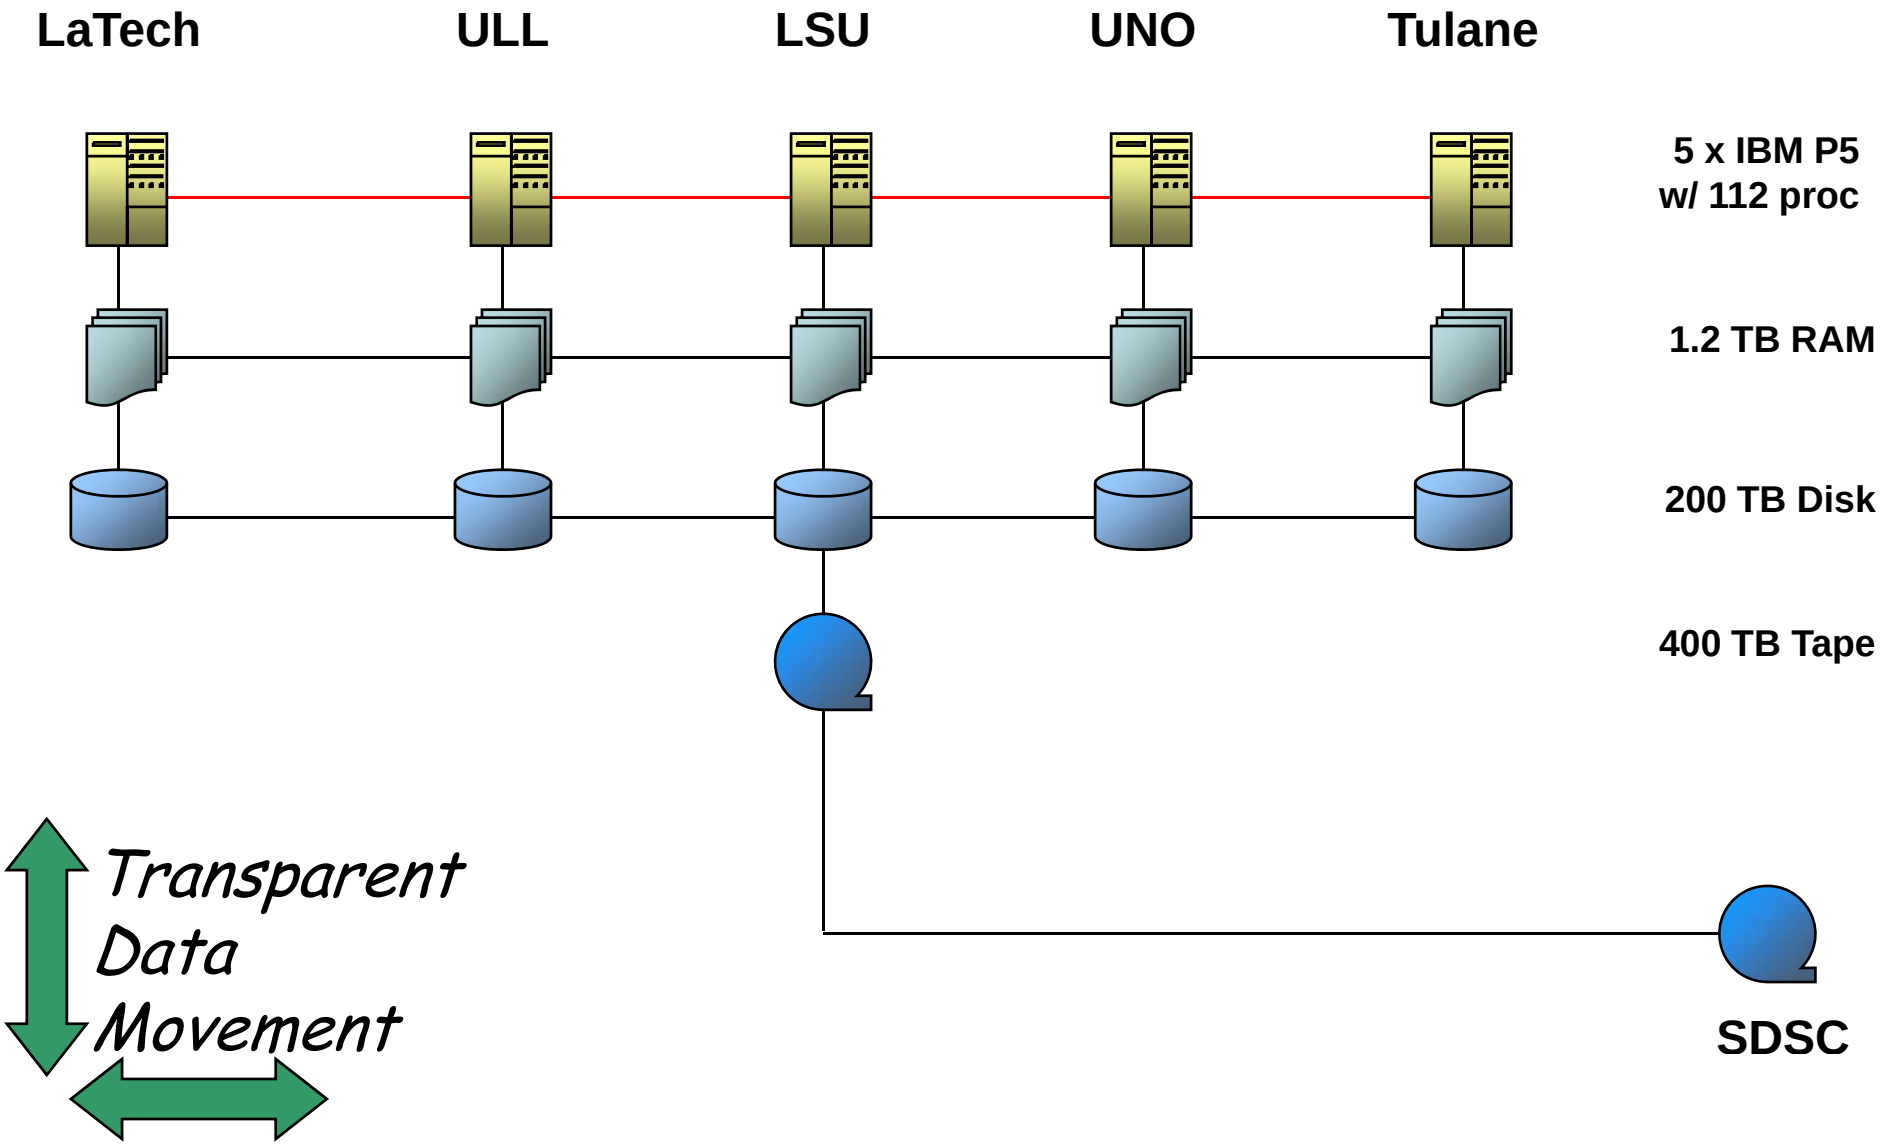

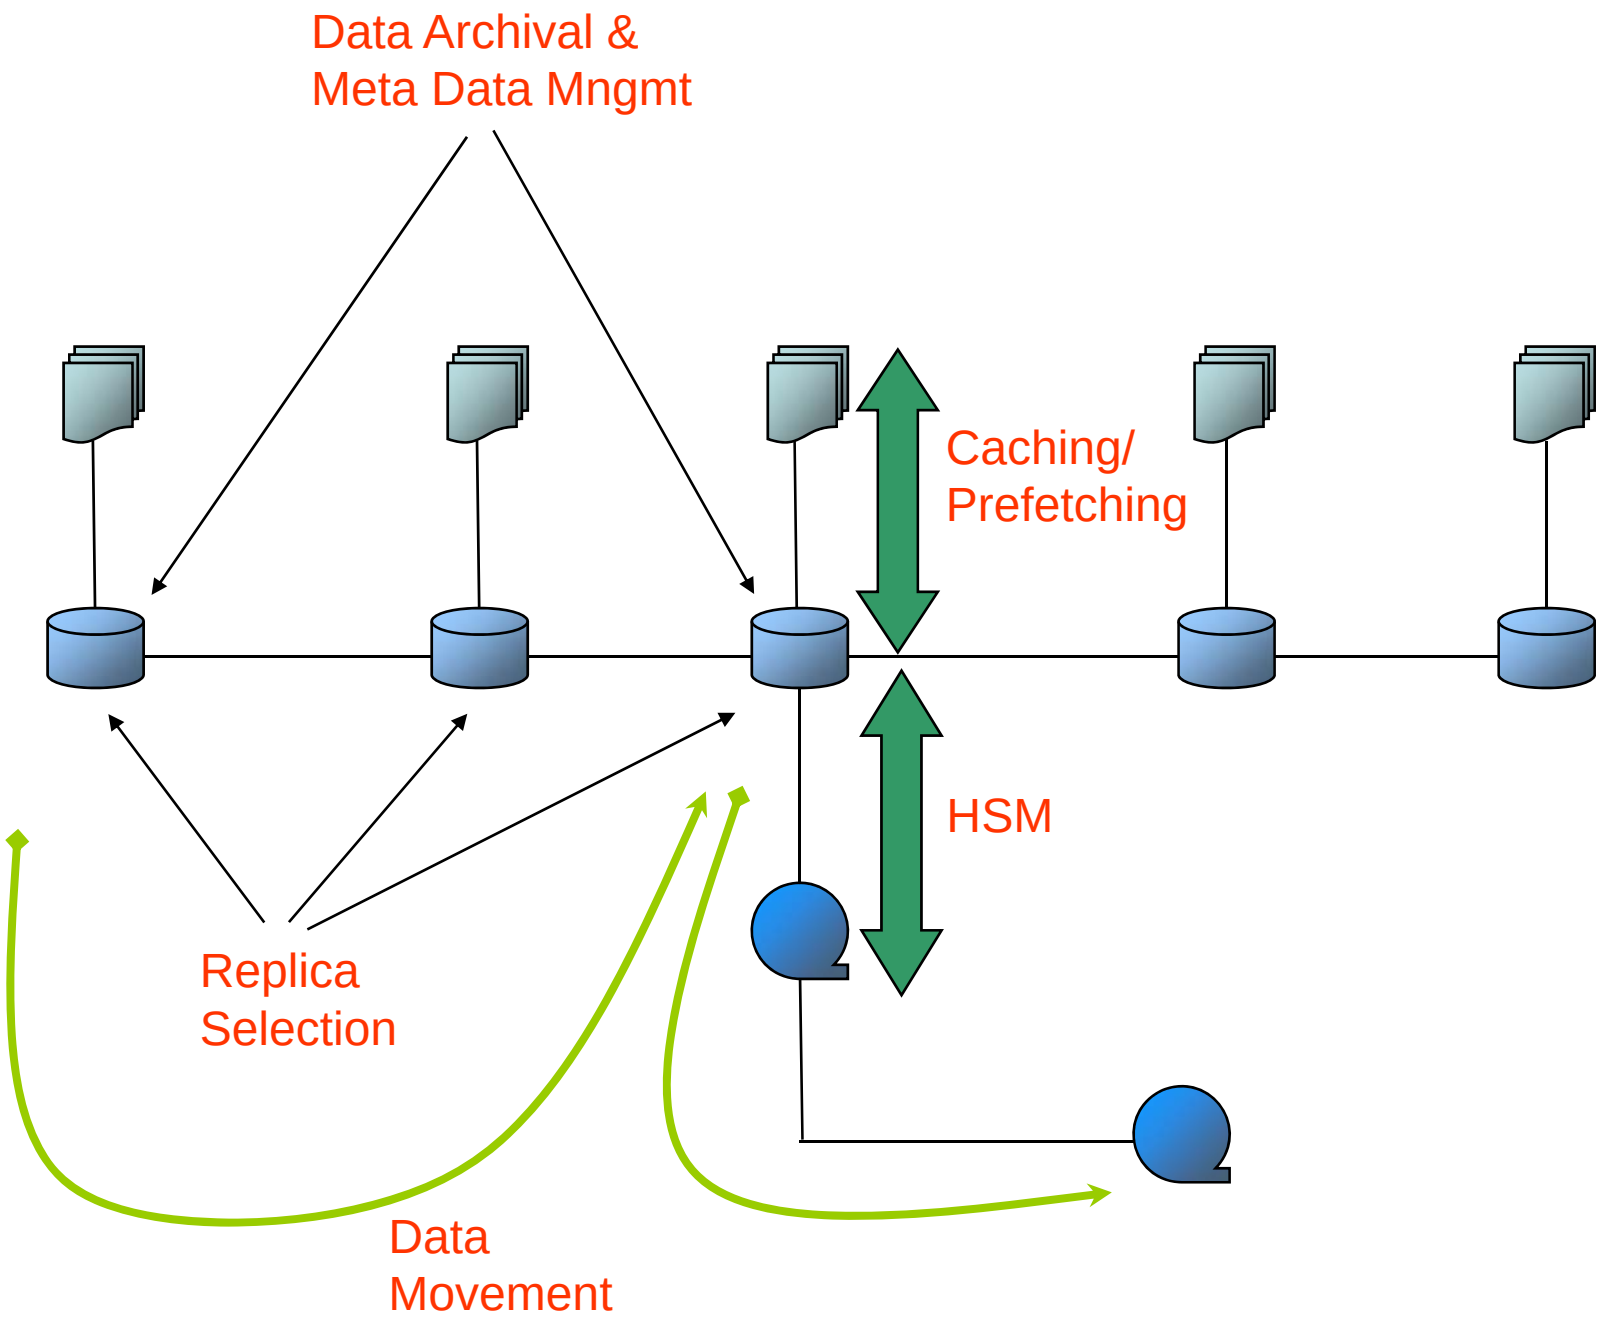
Data Archival &
Meta Data Mngmt
Caching/
Prefetching
HSM
Replica
Selection
Data
Movement

Storage Systems as First Class Entities

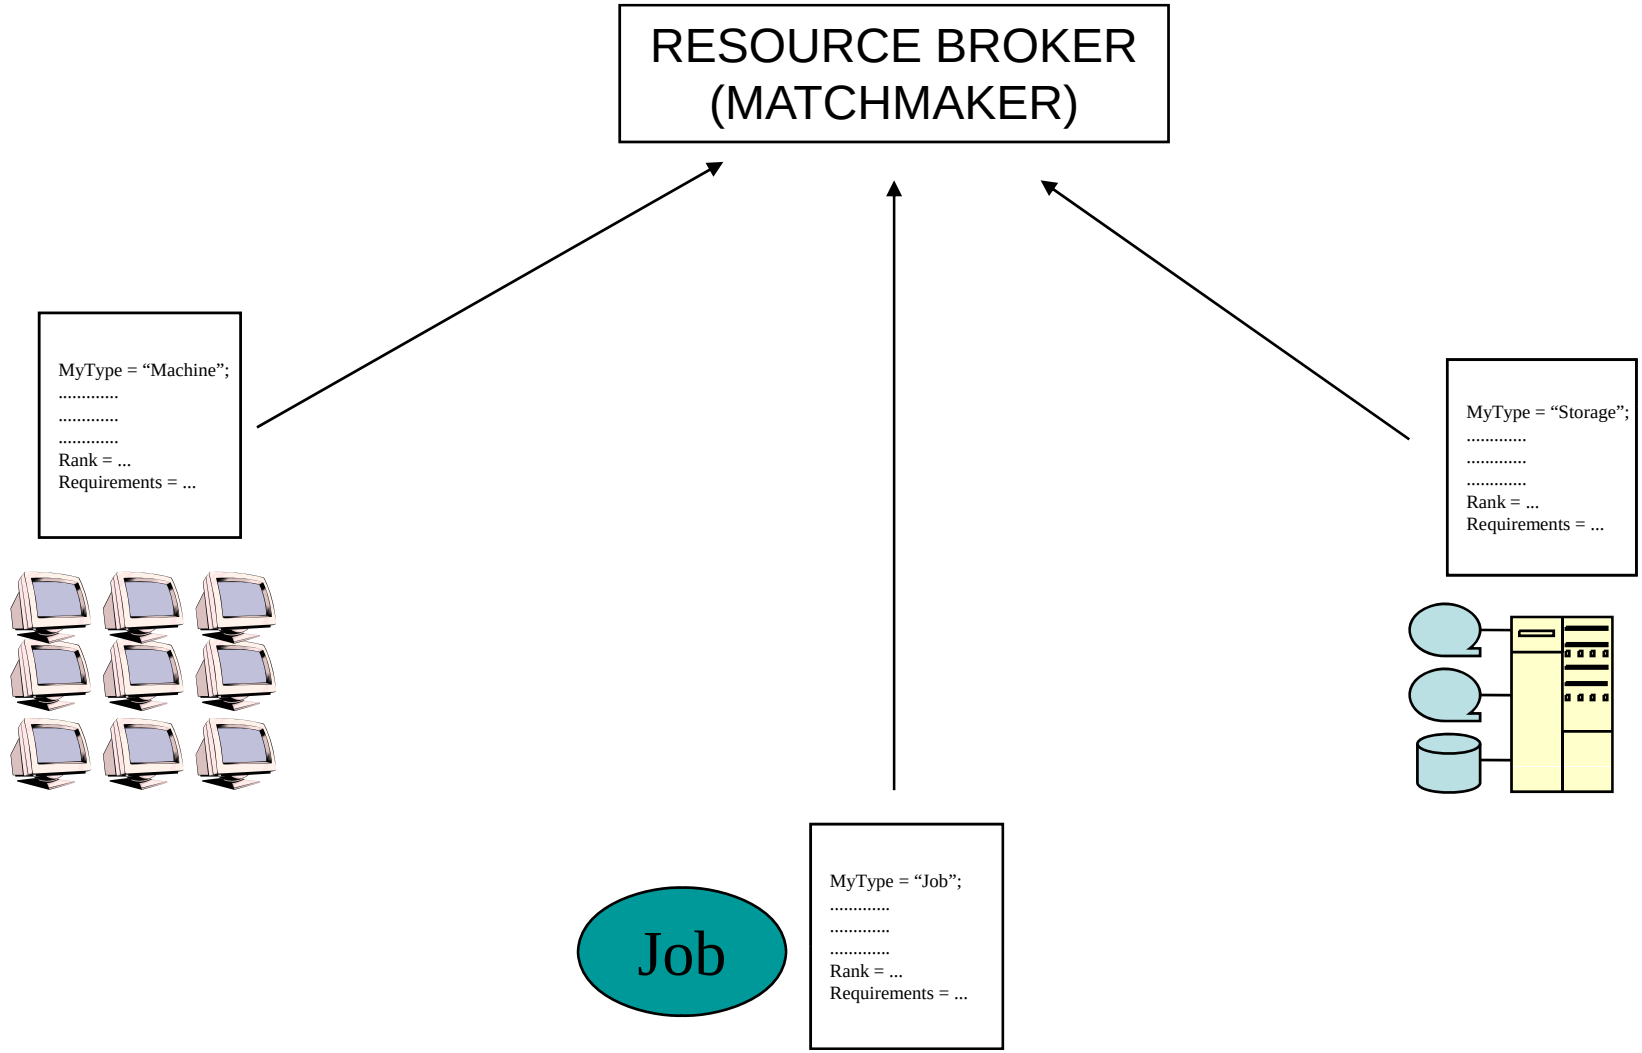

Data-Aware Storage

- Storage server advertises:
  - Metadata information
  - Location information
  - Available and used storage space
  - Maximum connections available (eg. Max FTP conn, Max GridFTP conn, Max HTTP conn)
- Scheduler takes these into account
  - Allocates a connection before data placement
  - Allocates storage

Data-Aware Schedulers

- Traditional schedulers not aware of characteristics and semantics of data placement jobs

```
Executable = genome.exe
Arguments  = a b c d
```

```
Executable = globus-url-copy
Arguments  = gsiftp://host1/f1
             gsiftp://host2/f2
             -p 4 -tcp-bs 1024
```

Any difference?

[ICDCS'04]

Data-Aware Schedulers

- What type of a job is it?
  - transfer, allocate, release, locate..
- What are the source and destination?
- Which protocols to use?
- What is available storage space?
- What is best concurrency level?
- What is the best route?
- What are the best network parameters?
  - tcp buffer size
  - I/O block size
  - # of parallel streams

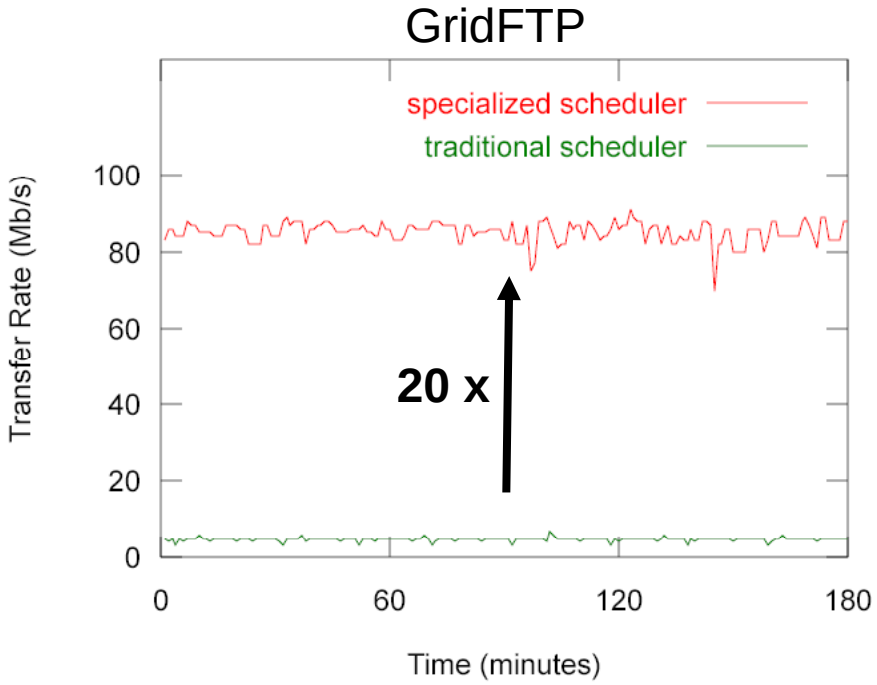

Optimizing Throughput and CPU Utilization at the same Time

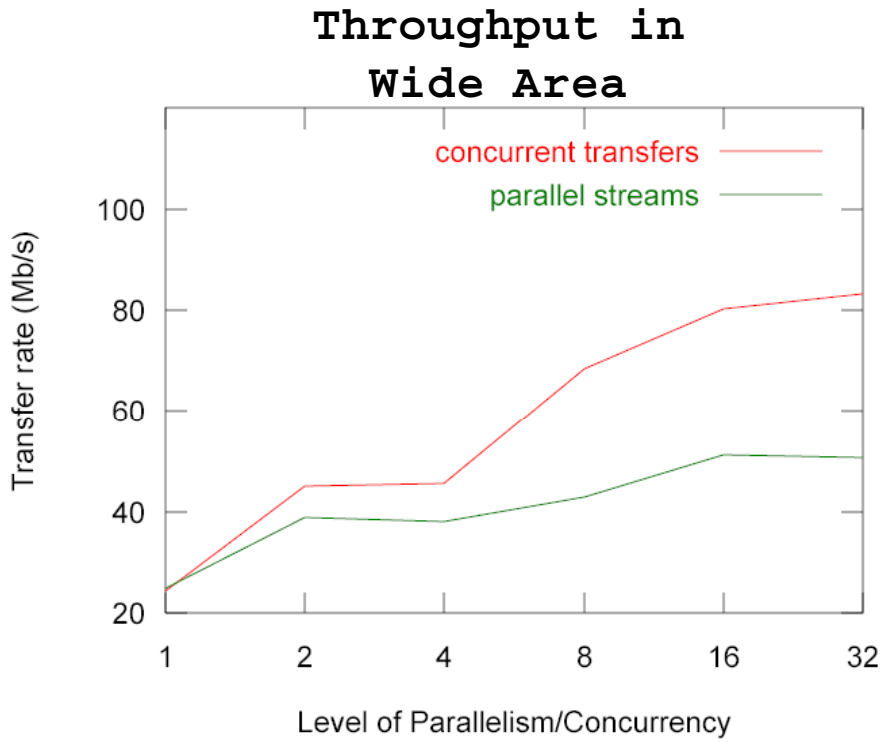

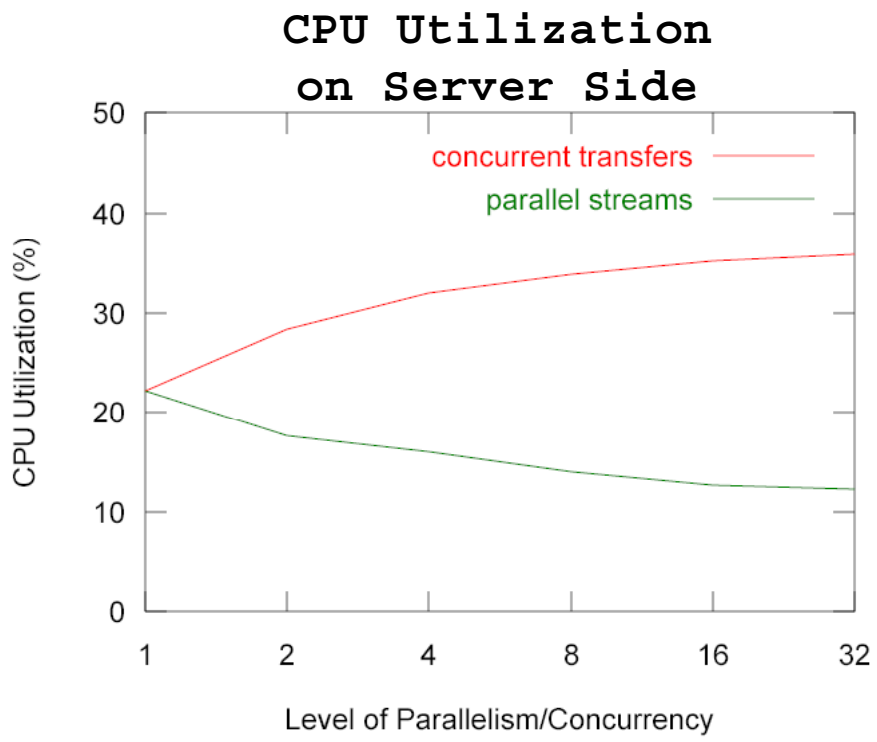

- Definitions:
  - Concurrency: transfer n files at the same time
  - Parallelism: transfer 1 file using n parallel streams

Storage Space Management

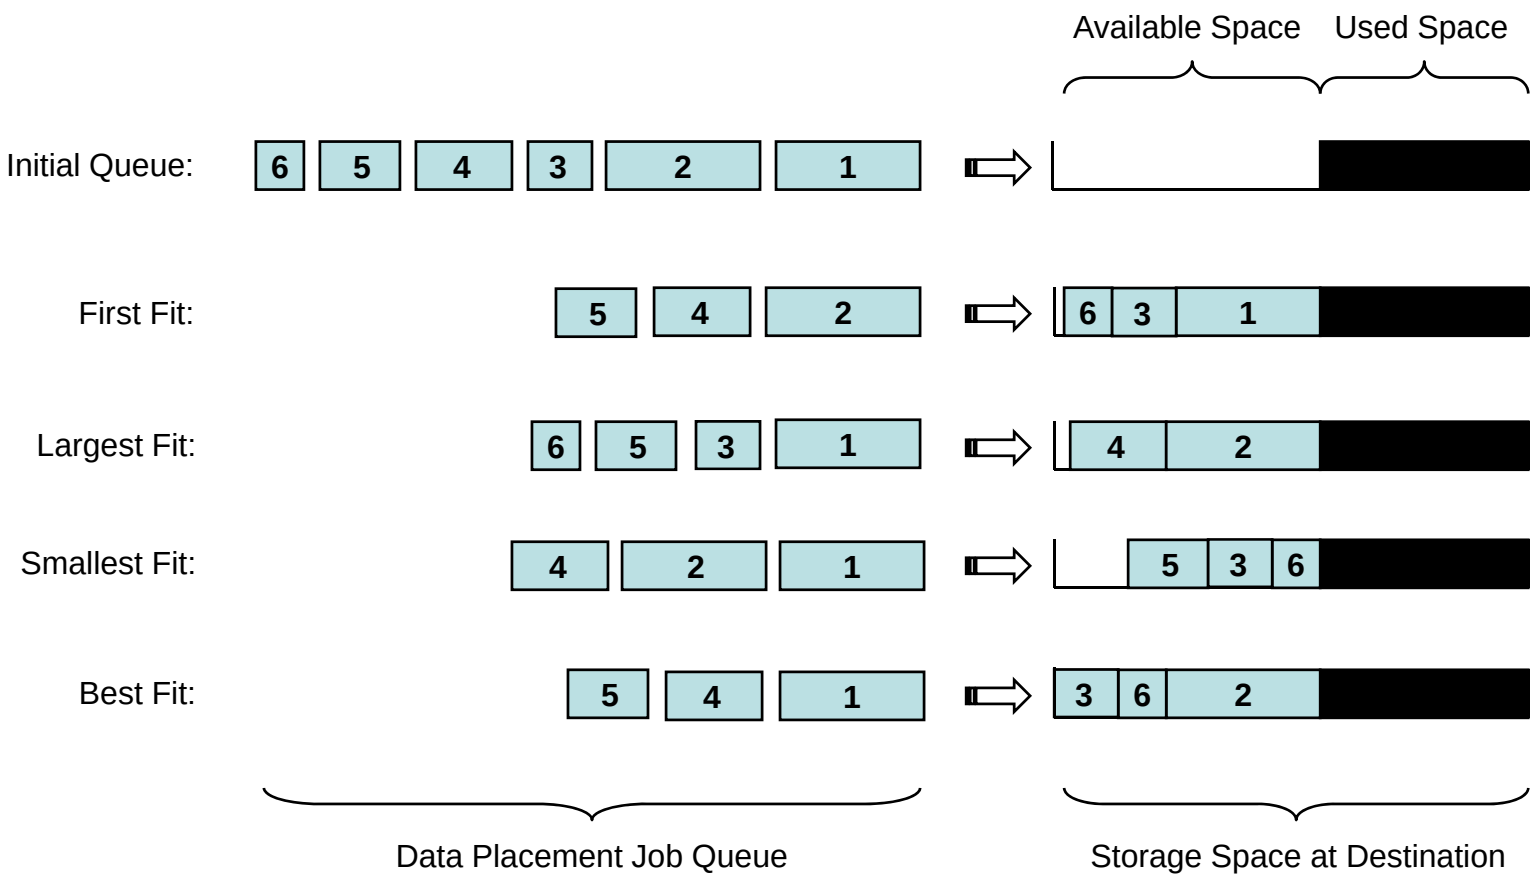

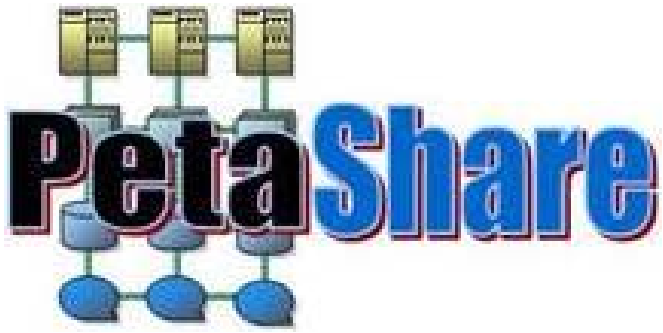

A system driven by the local needs (in LA), but has potential to be a generic solution for the broader community!