Running 1 Million Jobs in 10 Minutes via the Falkon Fast and Light-weight tasK executiON framework

Ioan Raicu
Distributed Systems Laboratory
Computer Science Department
University of Chicago

In Collaboration with:
Ian Foster, University of Chicago and Argonne National Laboratory
Mike Wilde, University of Chicago and Argonne National Laboratory
Zhao Zhang, University of Chicago
Yong Zhao, Microsoft
Catalin Dumitrescu, Fermi National Laboratory

Megajob BOF at IEEE/ACM Supercomputing 2008
November 18th, 2008

Obstacles running MTC apps in Clusters/Grids

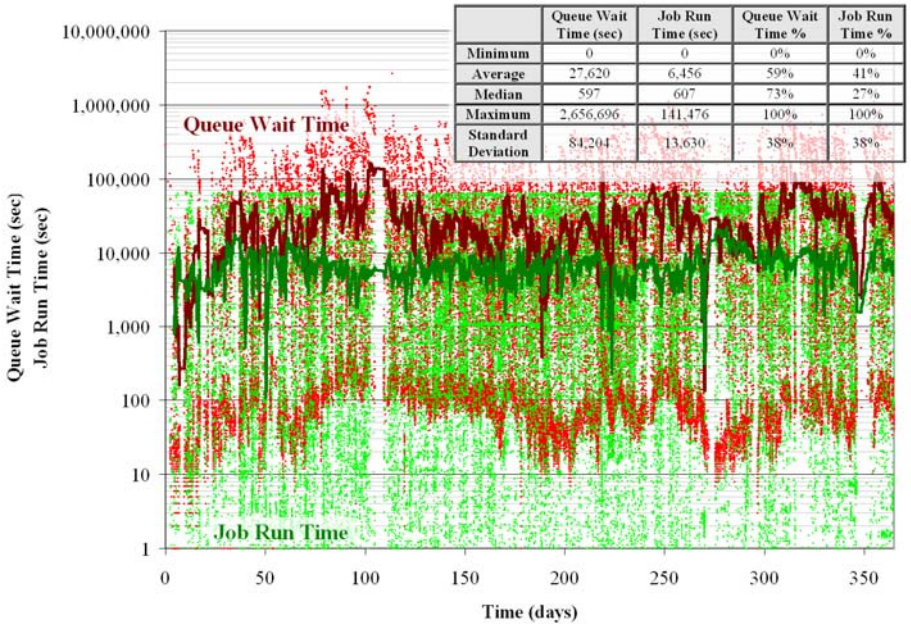

	Queue Wait Time (sec)	Job Run Time (sec)	Queue Wait Time %	Job Run Time %
Minimum	0	0	0%	0%
Average	27,620	6,456	59%	41%
Median	597	607	73%	27%
Maximum	2,656,696	141,476	100%	100%
Standard Deviation	84,204	13,630	38%	38%

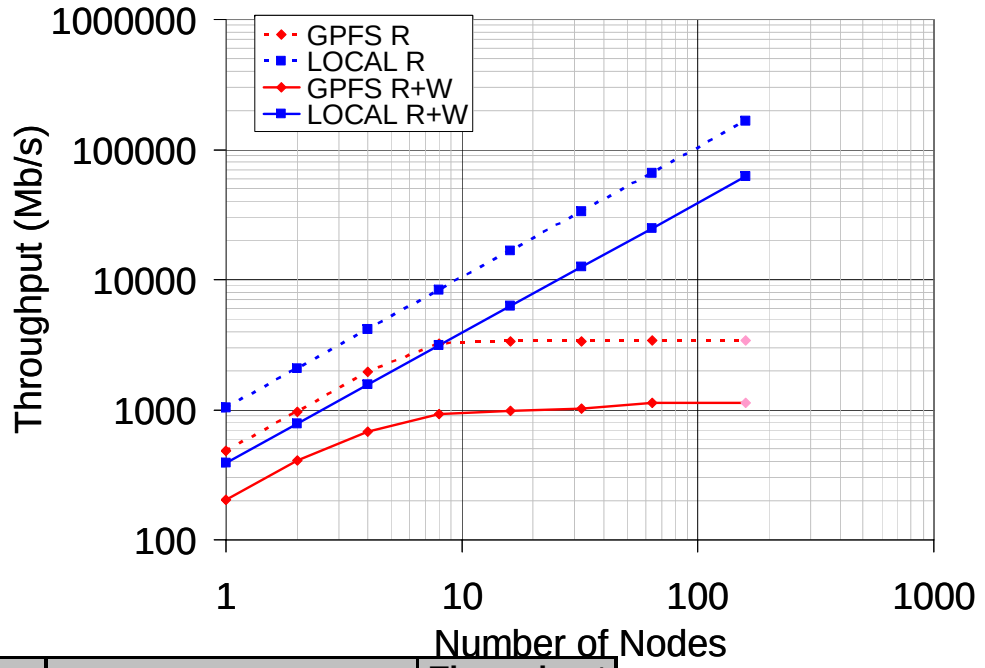

System	Comments	Throughput (tasks/sec)
Condor (v6.7.2) - Production	Dual Xeon 2.4GHz, 4GB	0.49
PBS (v2.1.8) - Production	Dual Xeon 2.4GHz, 4GB	0.45
Condor (v6.7.2) - Production	Quad Xeon 3 GHz, 4GB	2
Condor (v6.8.2) - Production		0.42
Condor (v6.9.3) - Development		11
Condor-J2 - Experimental	Quad Xeon 3 GHz, 4GB	22

Solution

- Falkon: A Fast and Light-weight tasK executiON framework
  - ***Goal:*** enable the ***rapid and efficient*** execution of many independent jobs on large compute clusters
  - Combines three components:
    - A ***streamlined task dispatcher***
    - ***Resource provisioning*** through multi-level scheduling techniques
    - ***Data diffusion*** and data-aware scheduling to leverage the co-located computational and storage resources

Falkon Overview

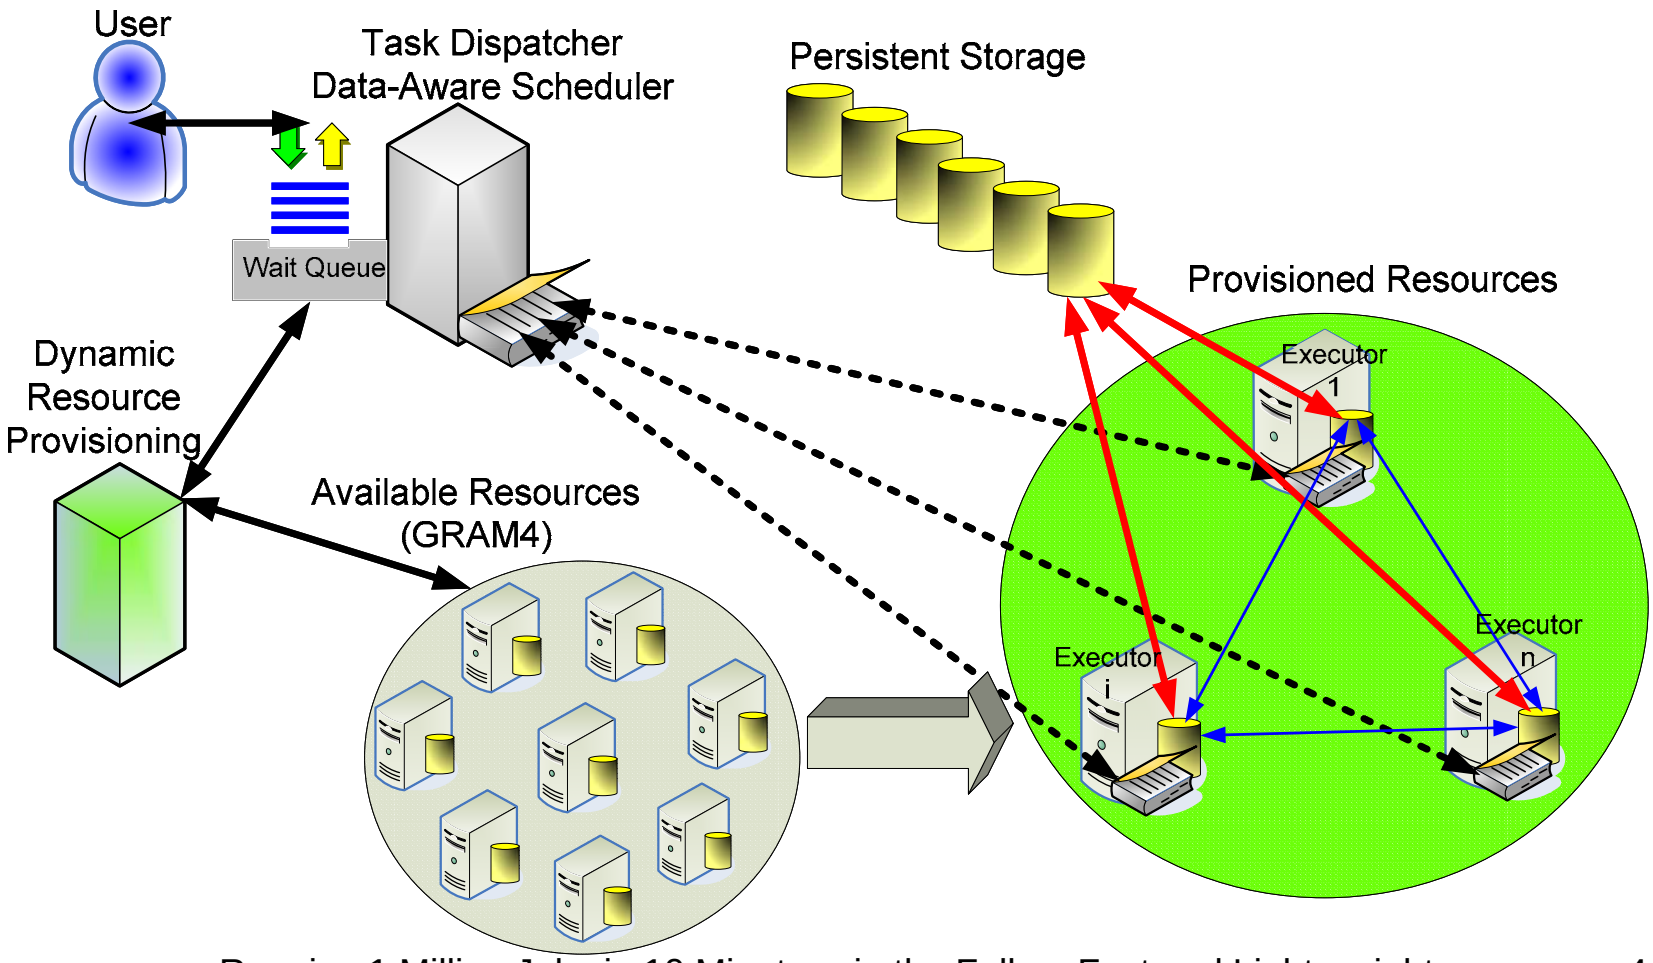

Distributed Falkon Architecture

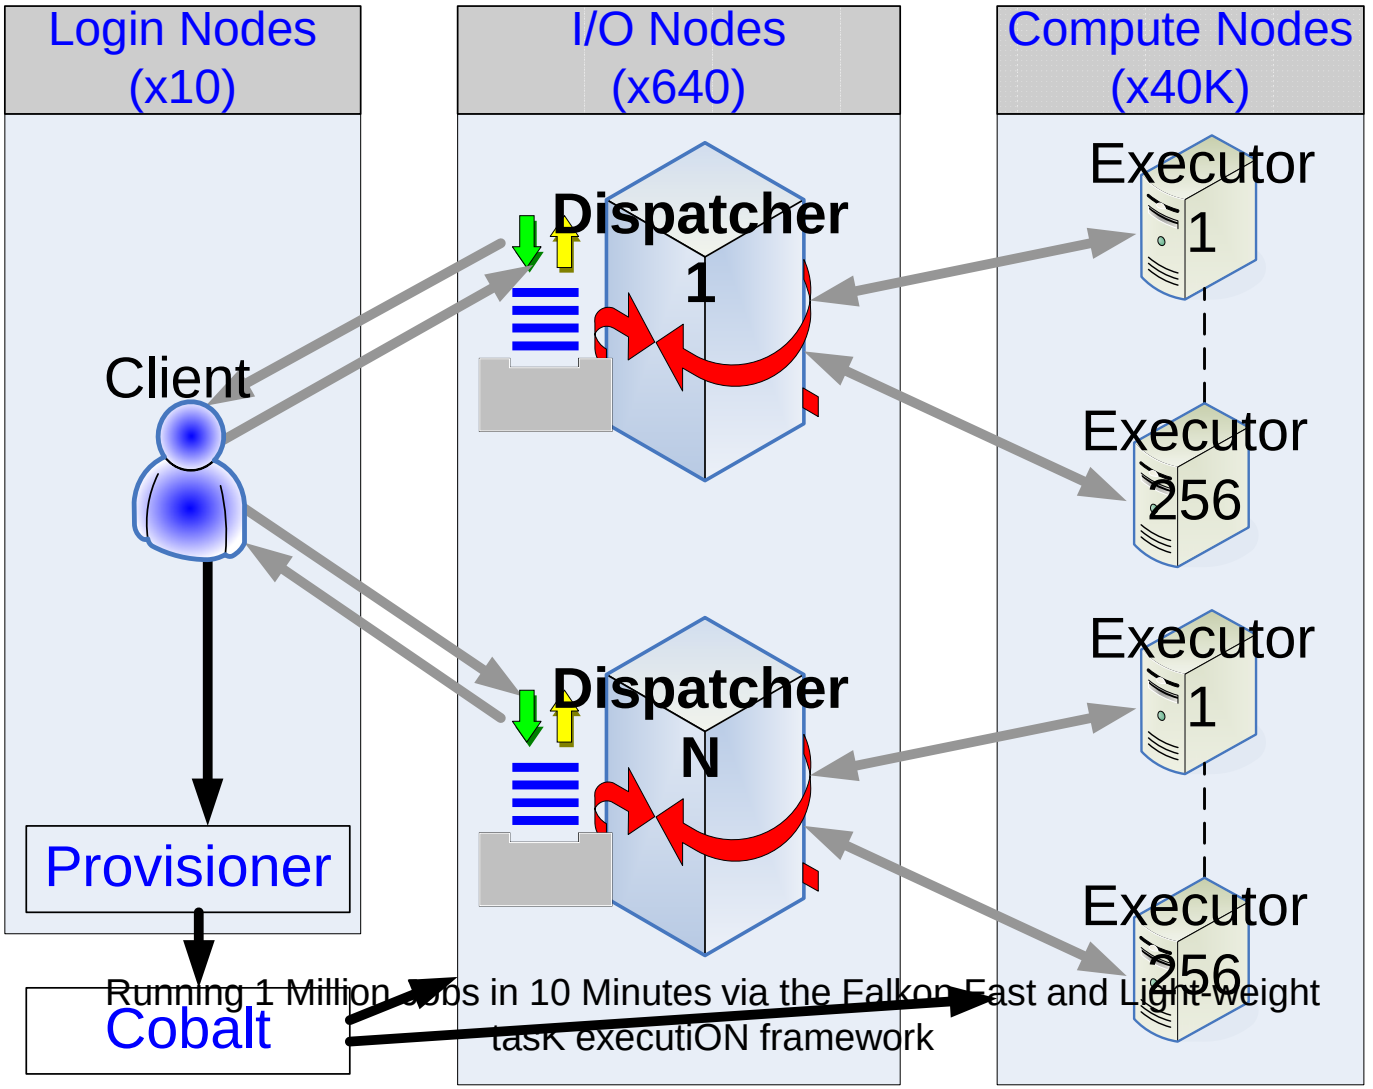

Dispatch Throughput

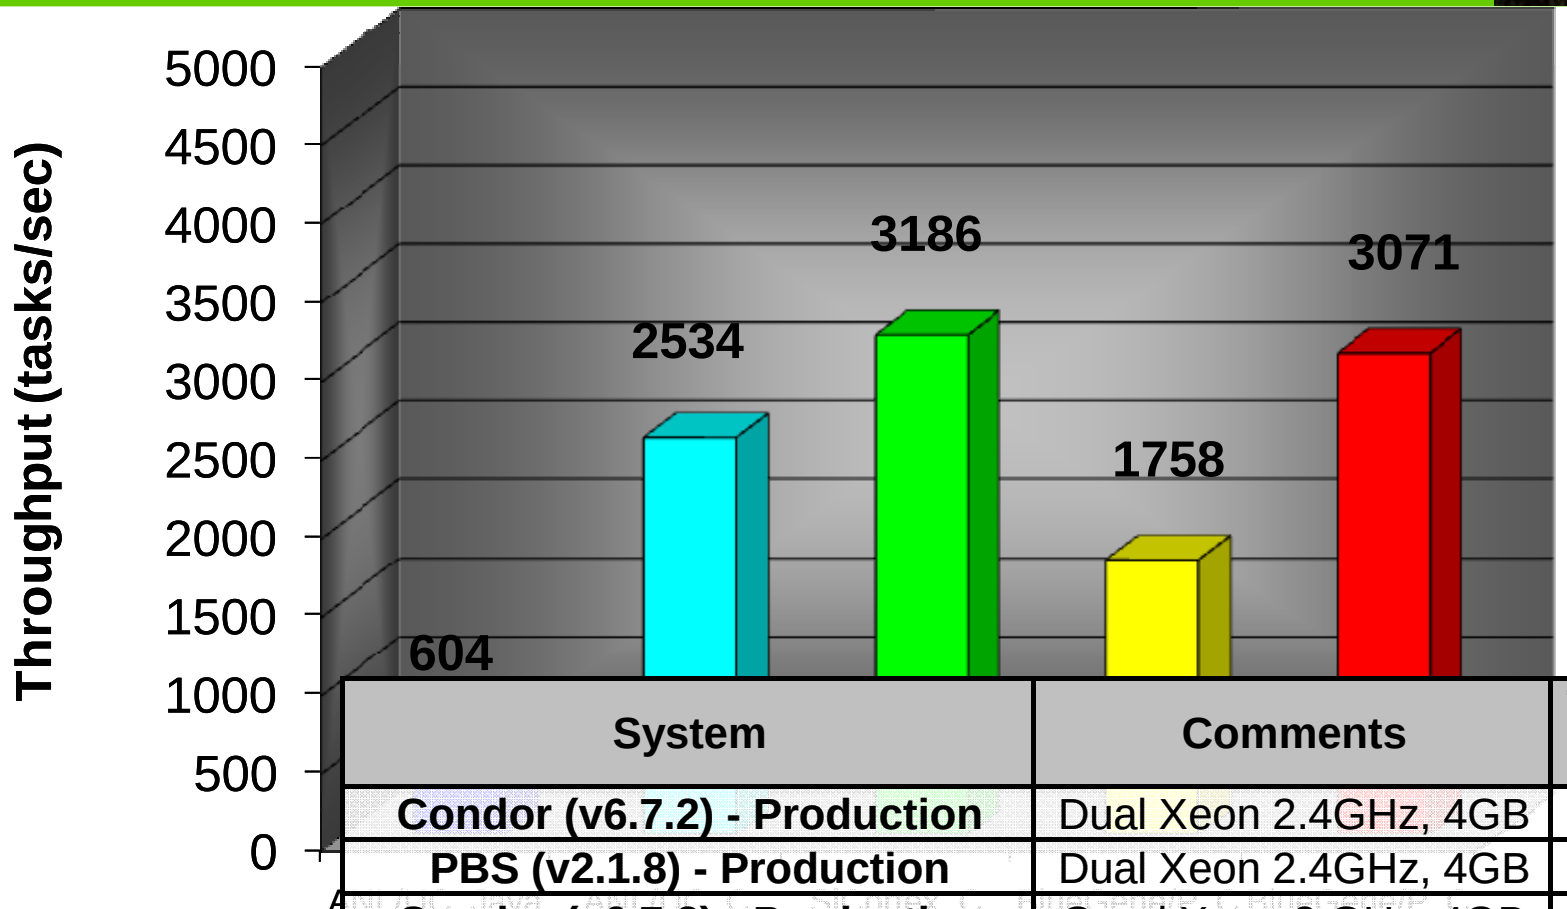

System	Comments	Throughput (tasks/sec)
Condor (v6.7.2) - Production	Dual Xeon 2.4GHz, 4GB	0.49
PBS (v2.1.8) - Production	Dual Xeon 2.4GHz, 4GB	0.45
Condor (v6.7.2) - Production	Quad Xeon 3 GHz, 4GB	2
Condor (v6.8.2) - Production		0.42
Condor (v6.9.3) - Development		11
Condor-J2 - Experimental	Quad Xeon 3 GHz, 4GB	22

Falkon Endurance Test

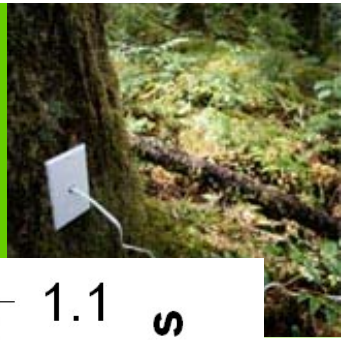

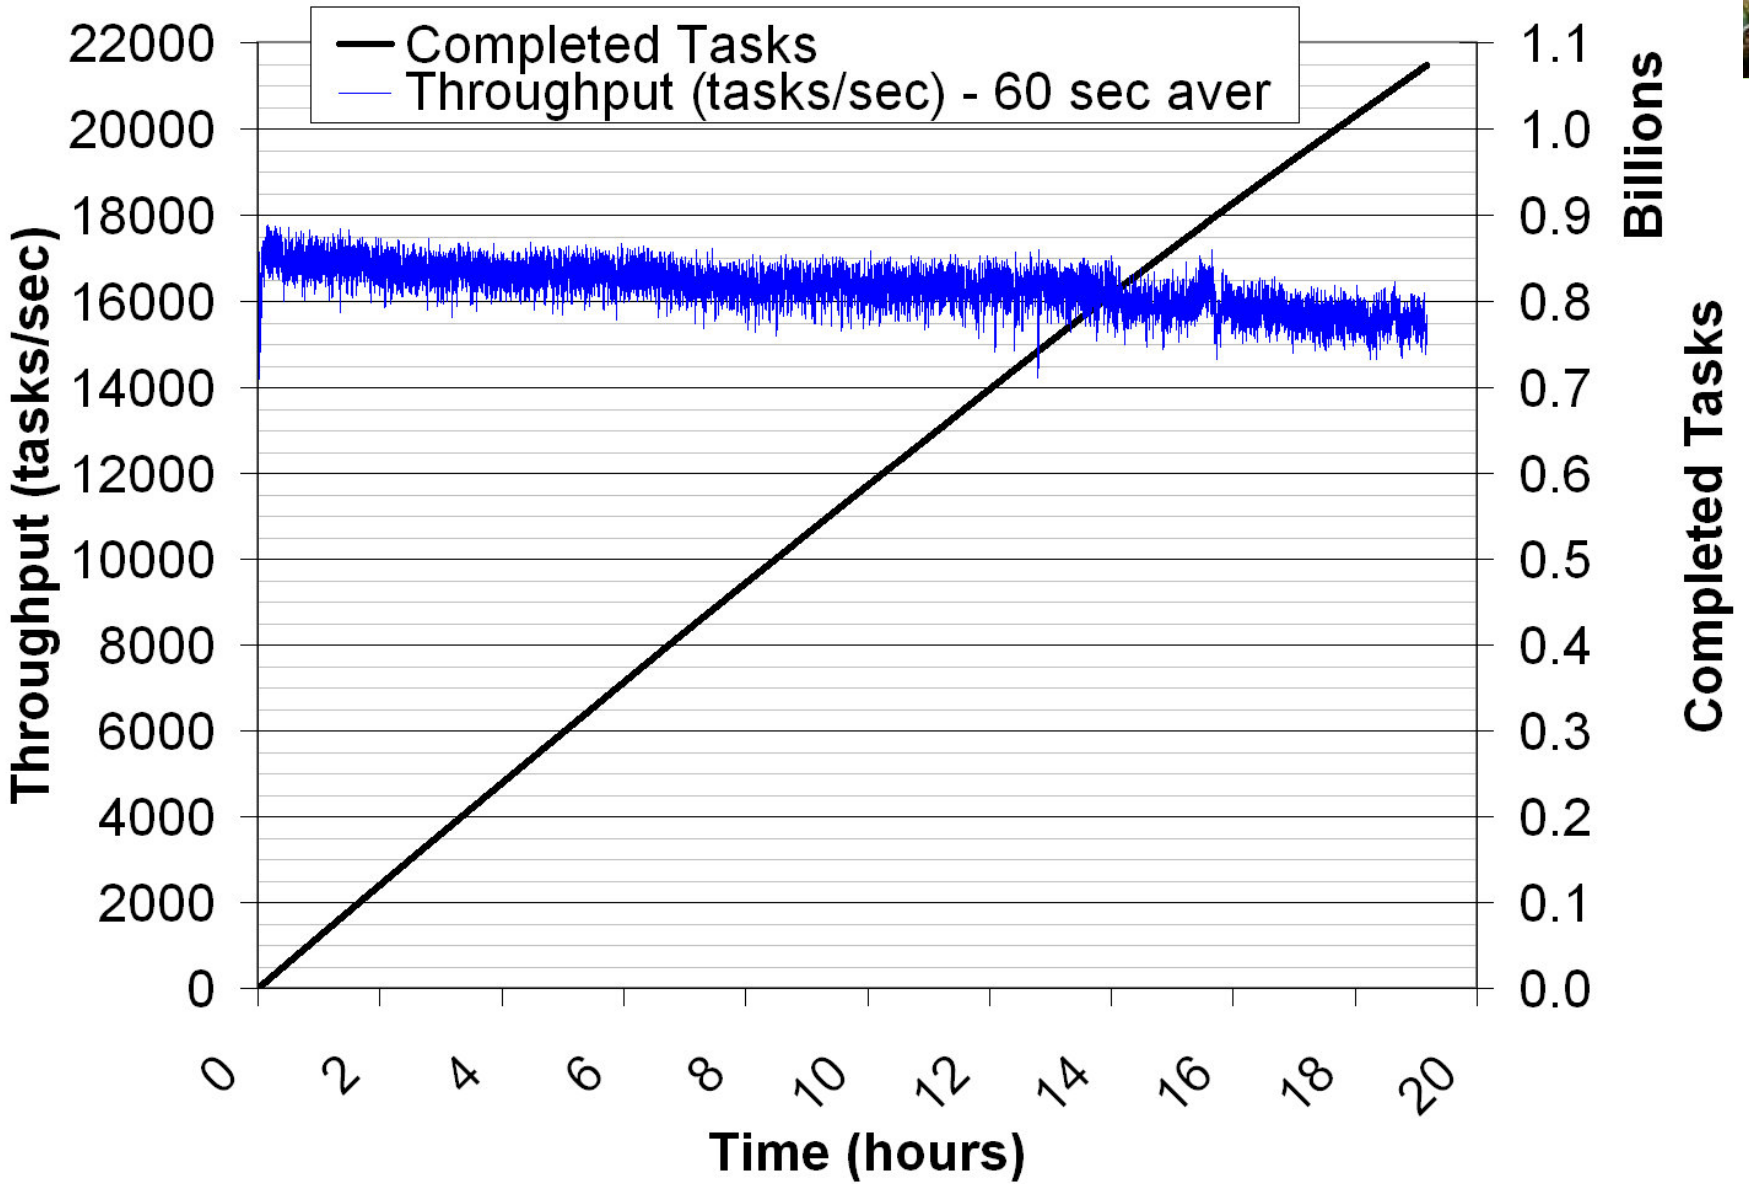

Efficiency

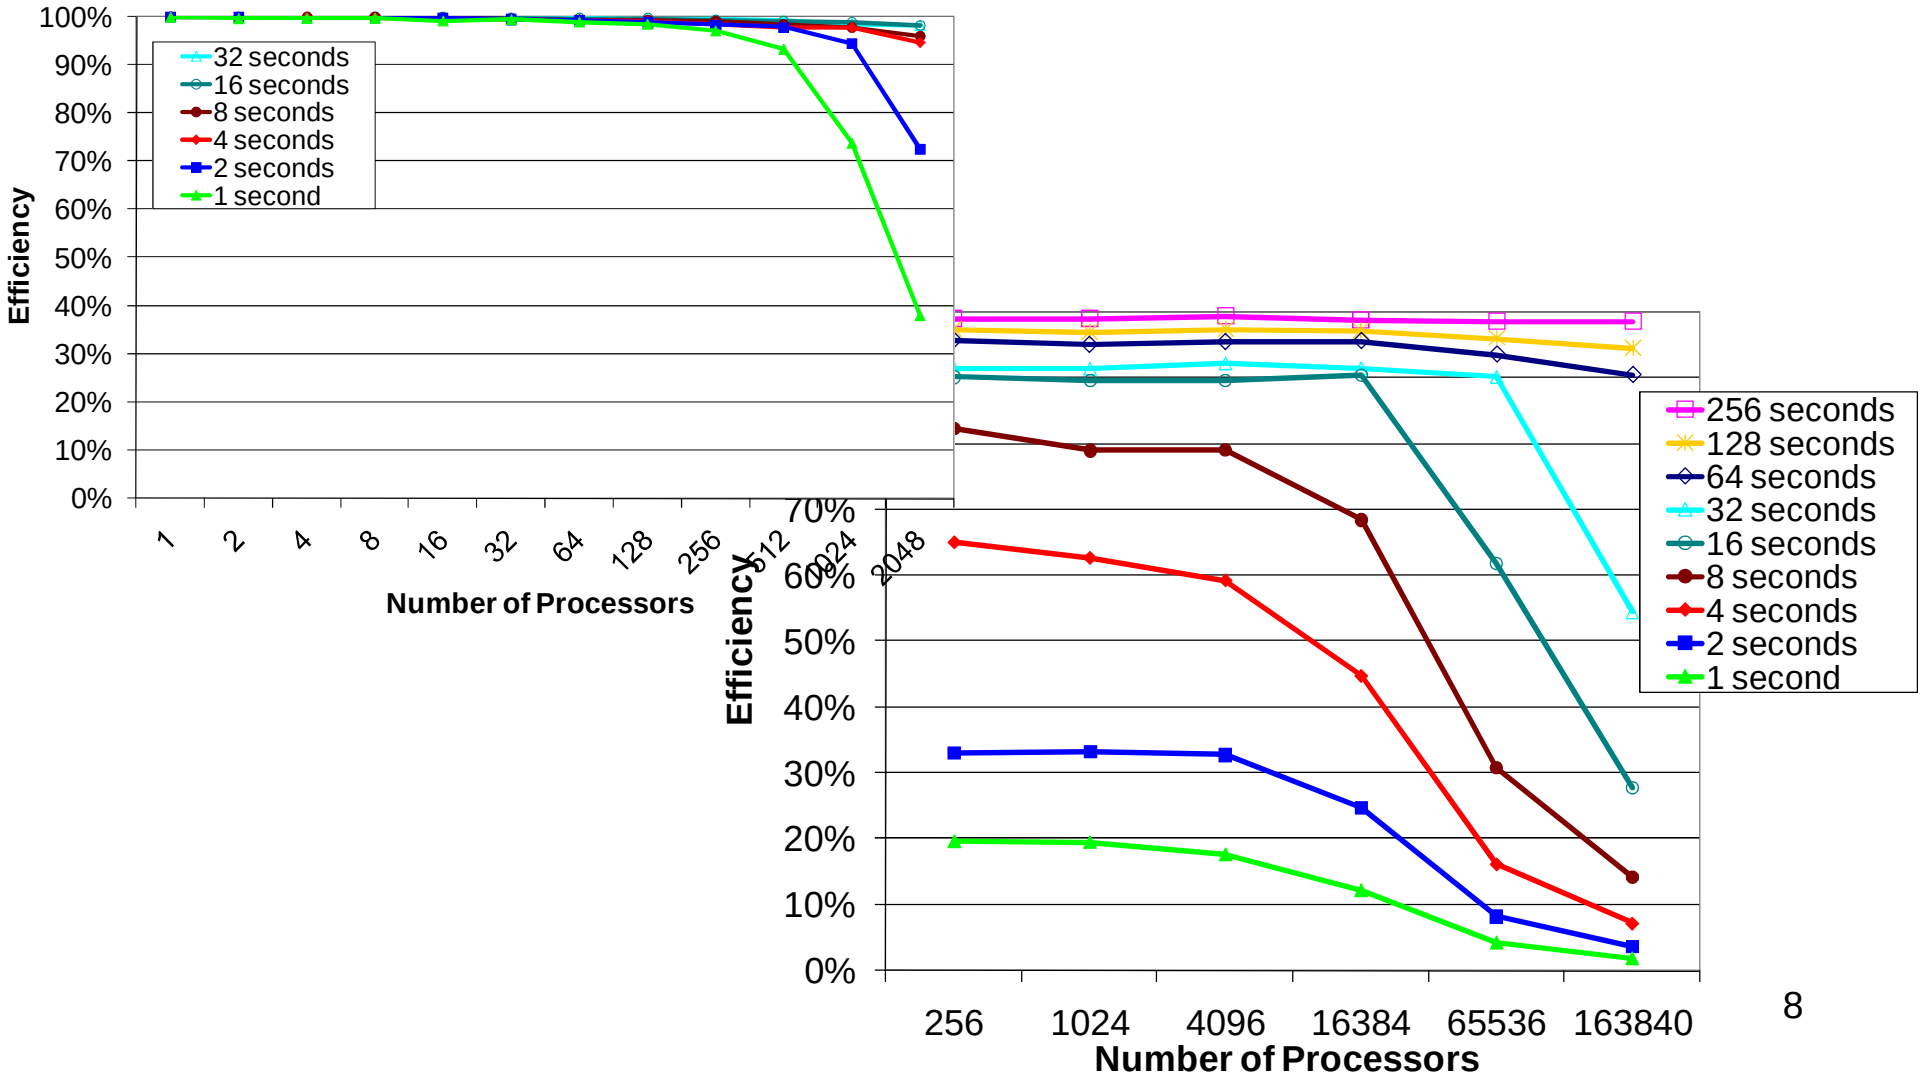
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Efficiency
32 seconds
16 seconds
8 seconds
4 seconds
2 seconds
1 second
1
2
4
8
16
32
64
128
256
512
1024
2048
Number of Processors
70%
60%
50%
40%
30%
20%
10%
0%
Efficiency
256 seconds
128 seconds
64 seconds
32 seconds
16 seconds
8 seconds
4 seconds
2 seconds
1 second
256
1024
4096
16384
65536
163840
Number of Processors

Falkon Activity History (10 months)

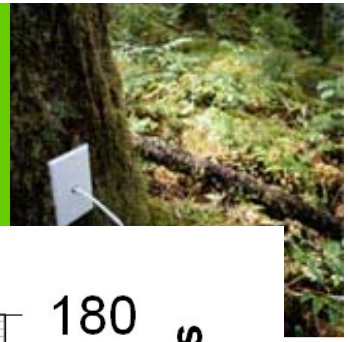

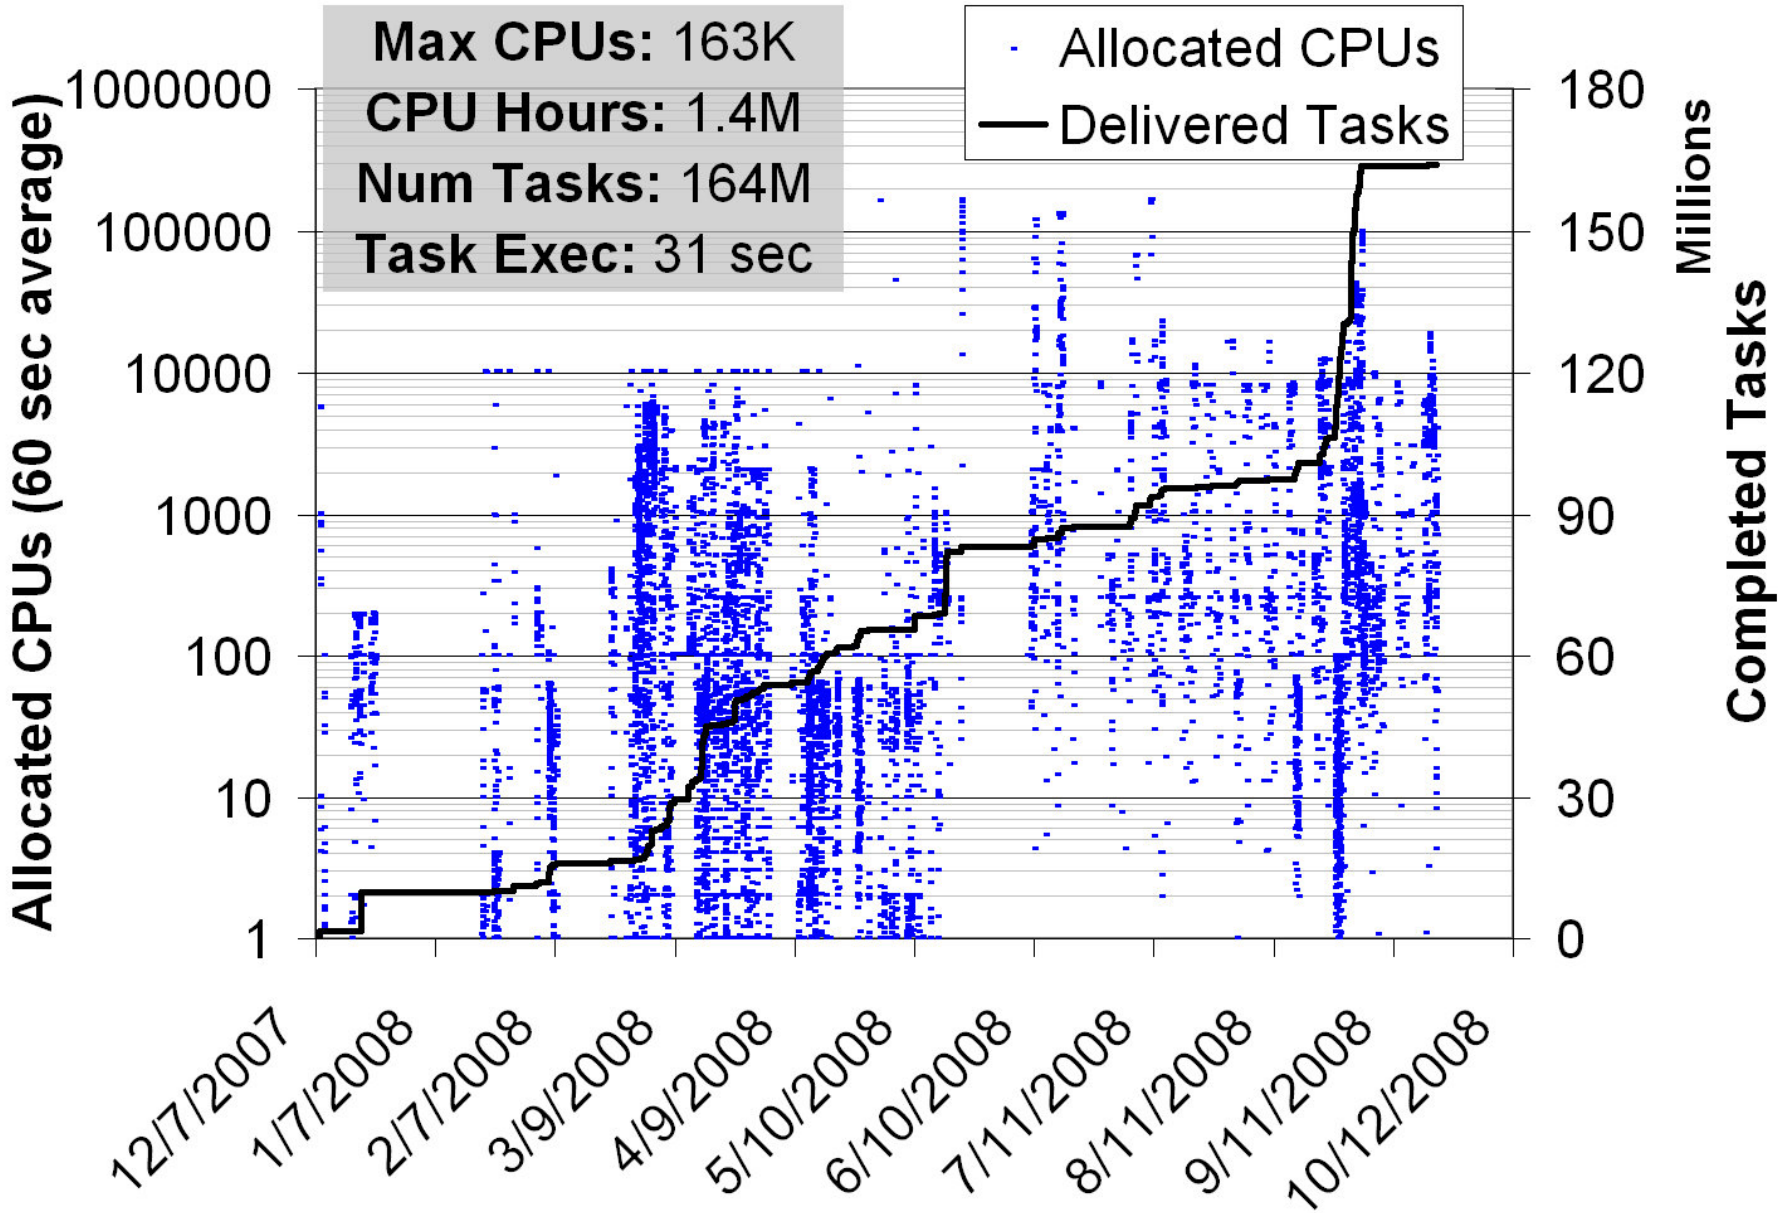

Falkon Demo

- Workload
  - 160K CPUs
  - 1M tasks
  - 60 sec per task
- 17.5K CPU hours in 7.5 min
- Throughput: 2312 tasks/sec
- 85% efficiency

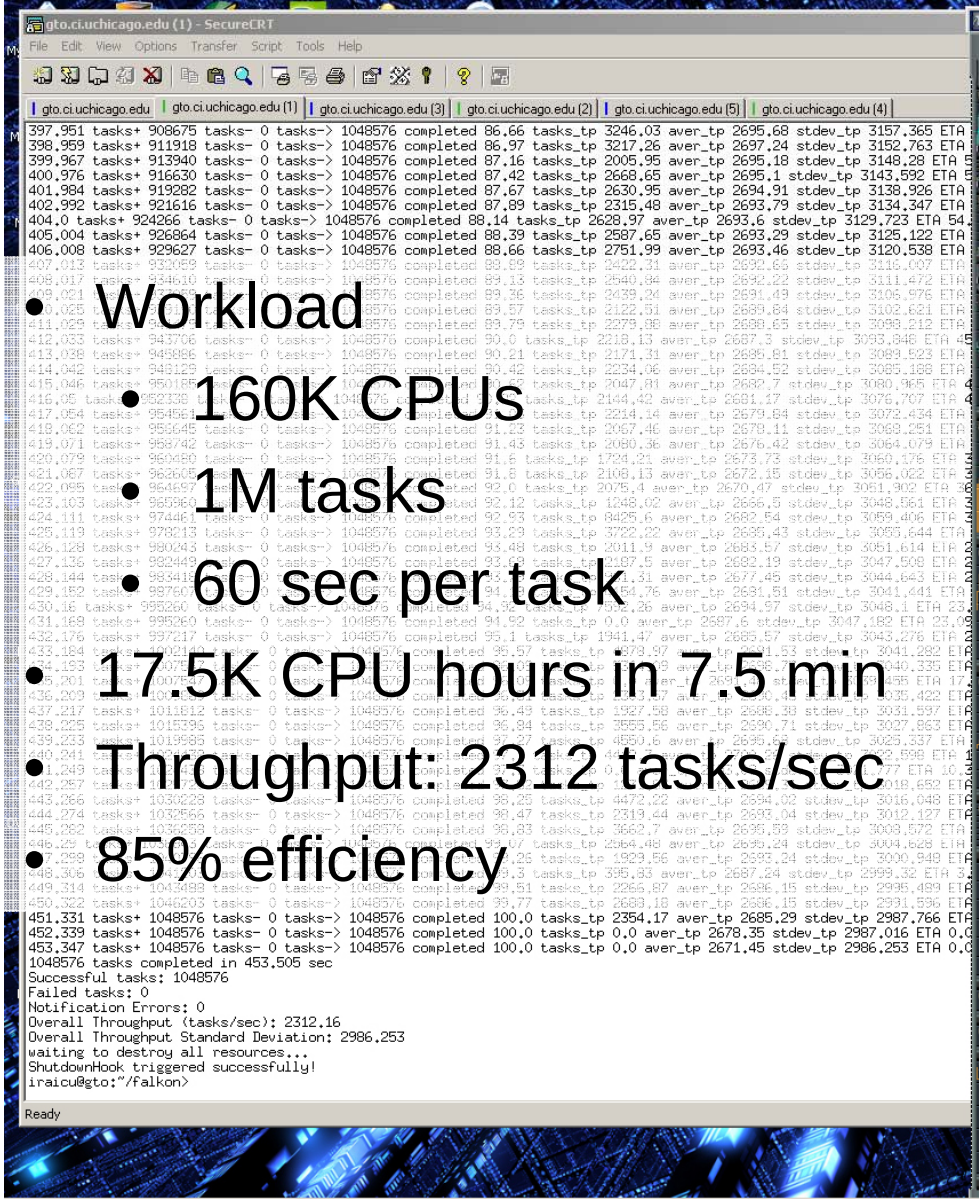

MARS Economic Modeling on IBM BG/P (128K CPUs)

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

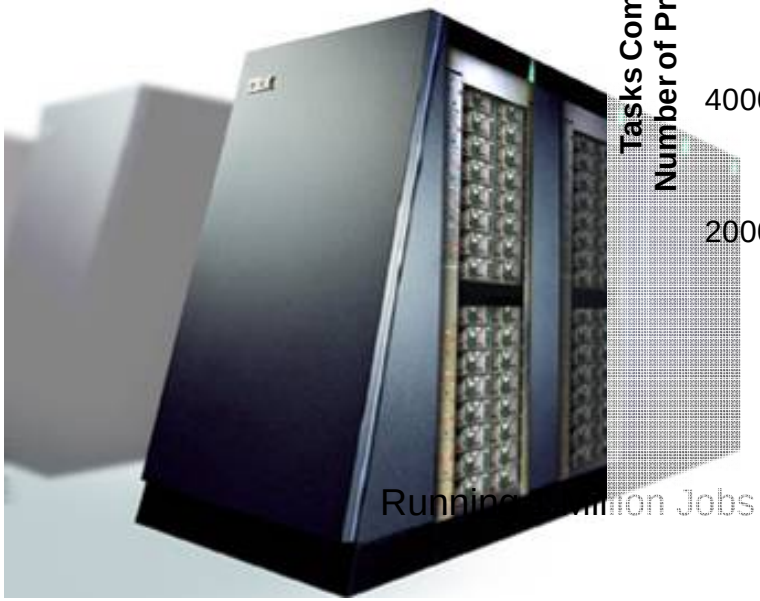

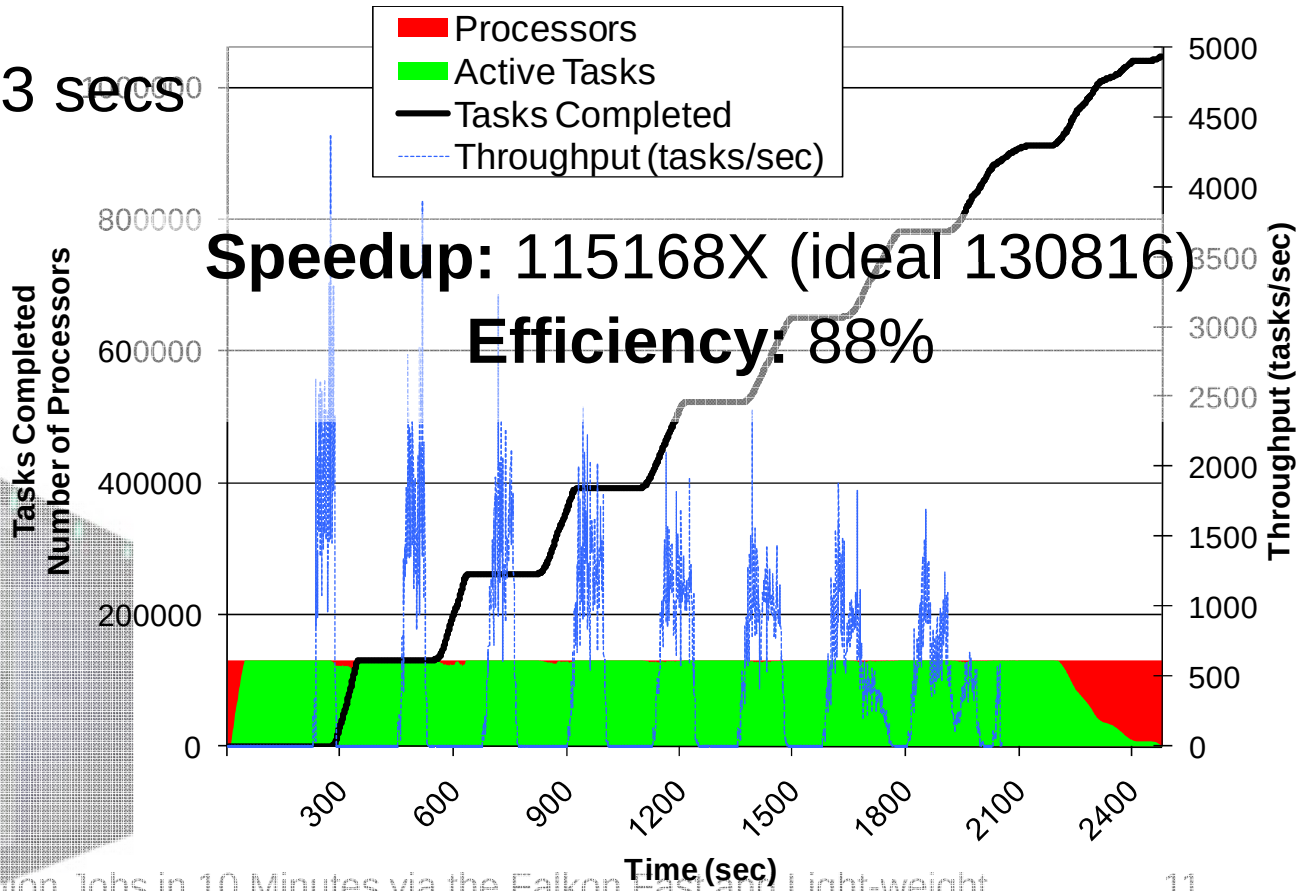

DOCK on the BG/P

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%

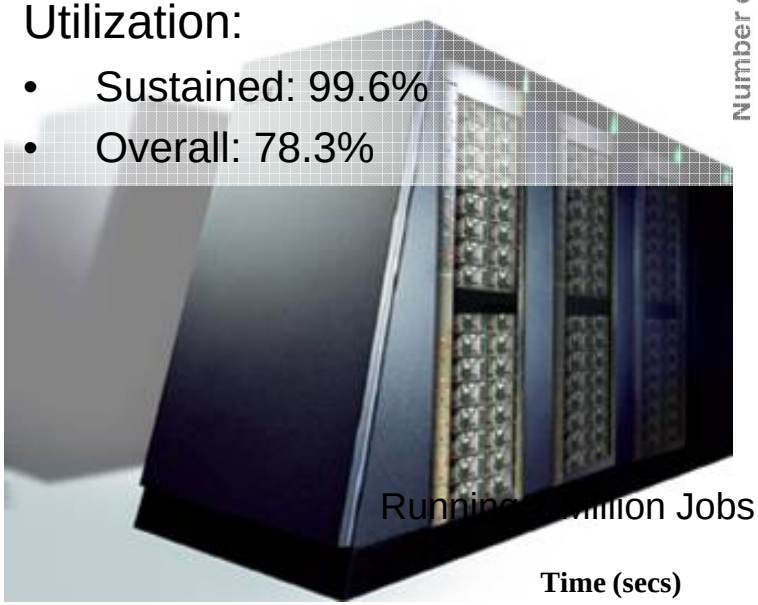

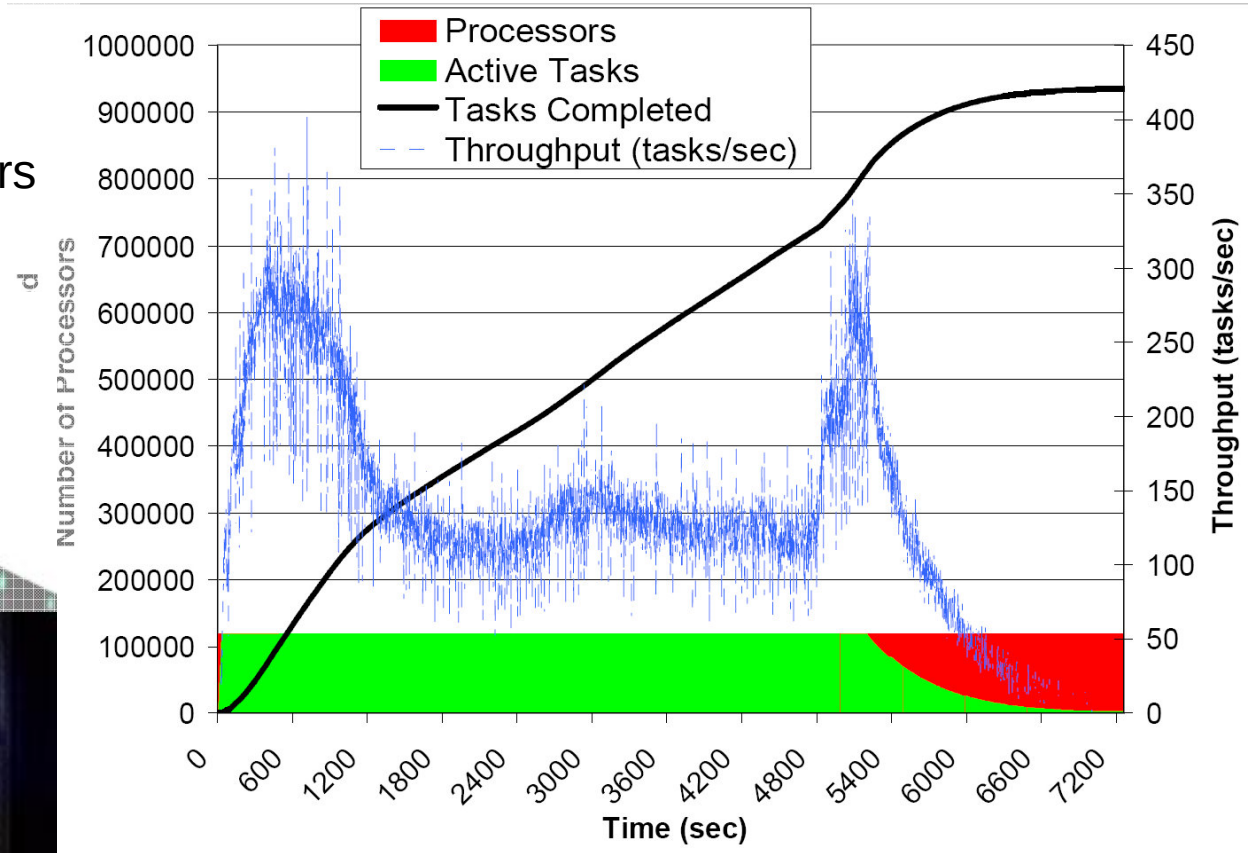

Time (secs)

Collective IO Model

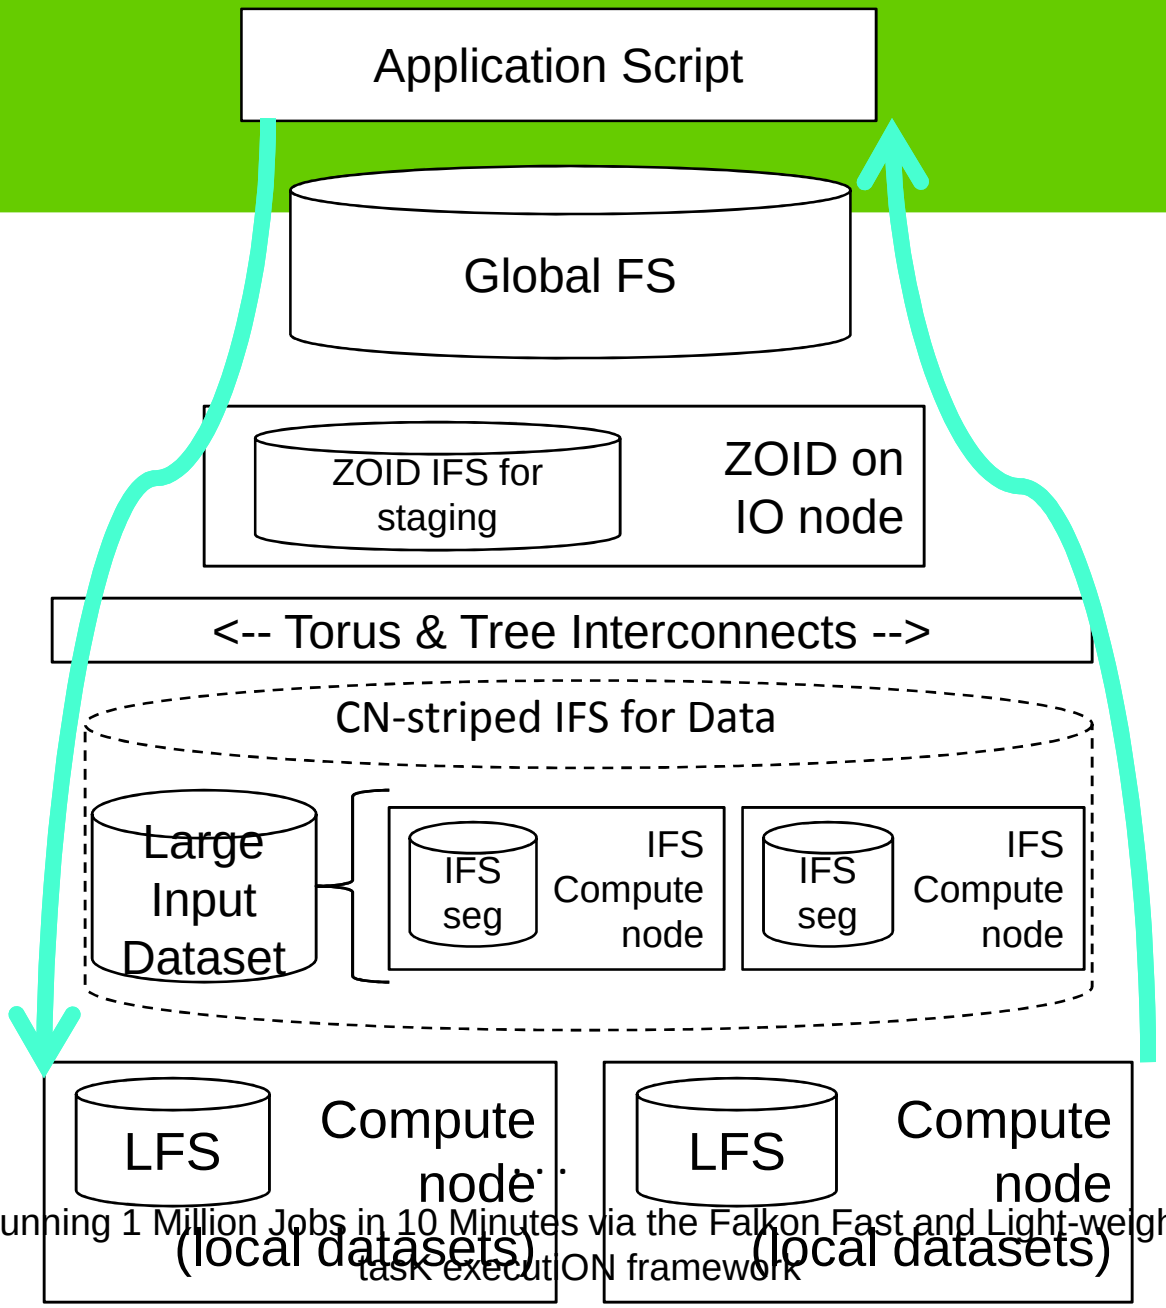

Write Performance
CIO vs. GFS efficiency

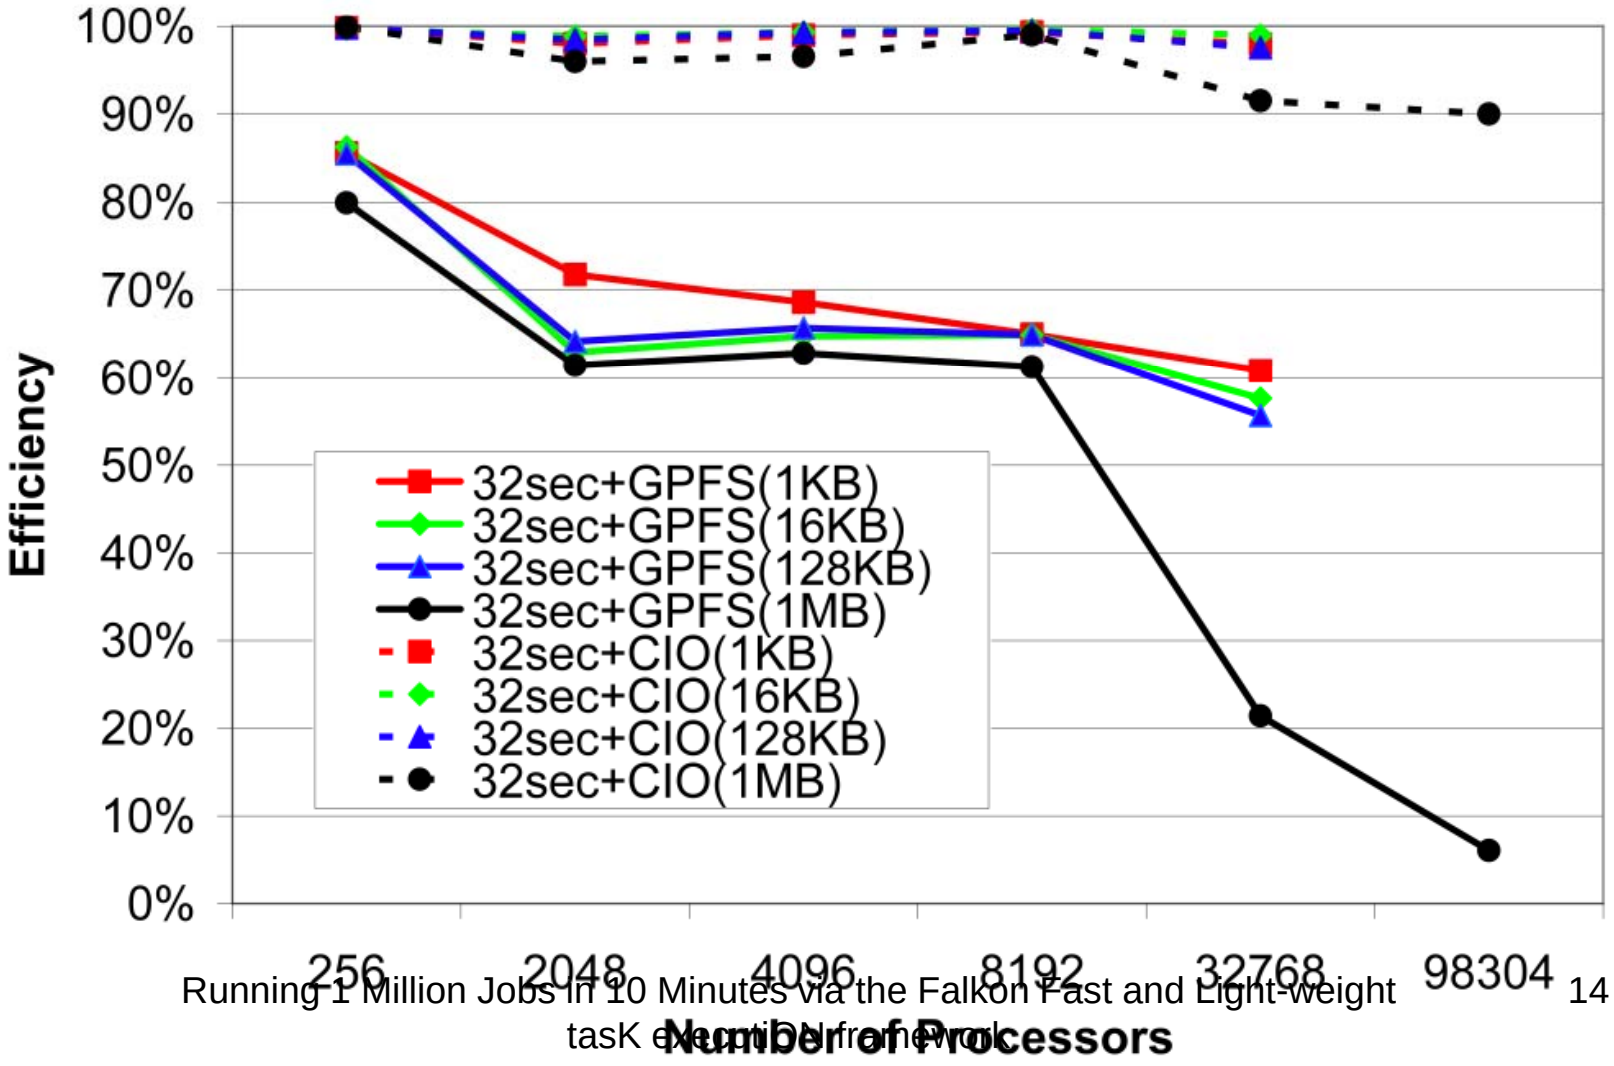

Data Diffusion

- Resource acquired in response to demand
- Data and applications diffuse from archival storage to newly acquired resources
- Resource “caching” allows faster responses to subsequent requests
  - Cache Eviction Strategies: RANDOM, FIFO, LRU, LFU
- Resources are released when demand drops

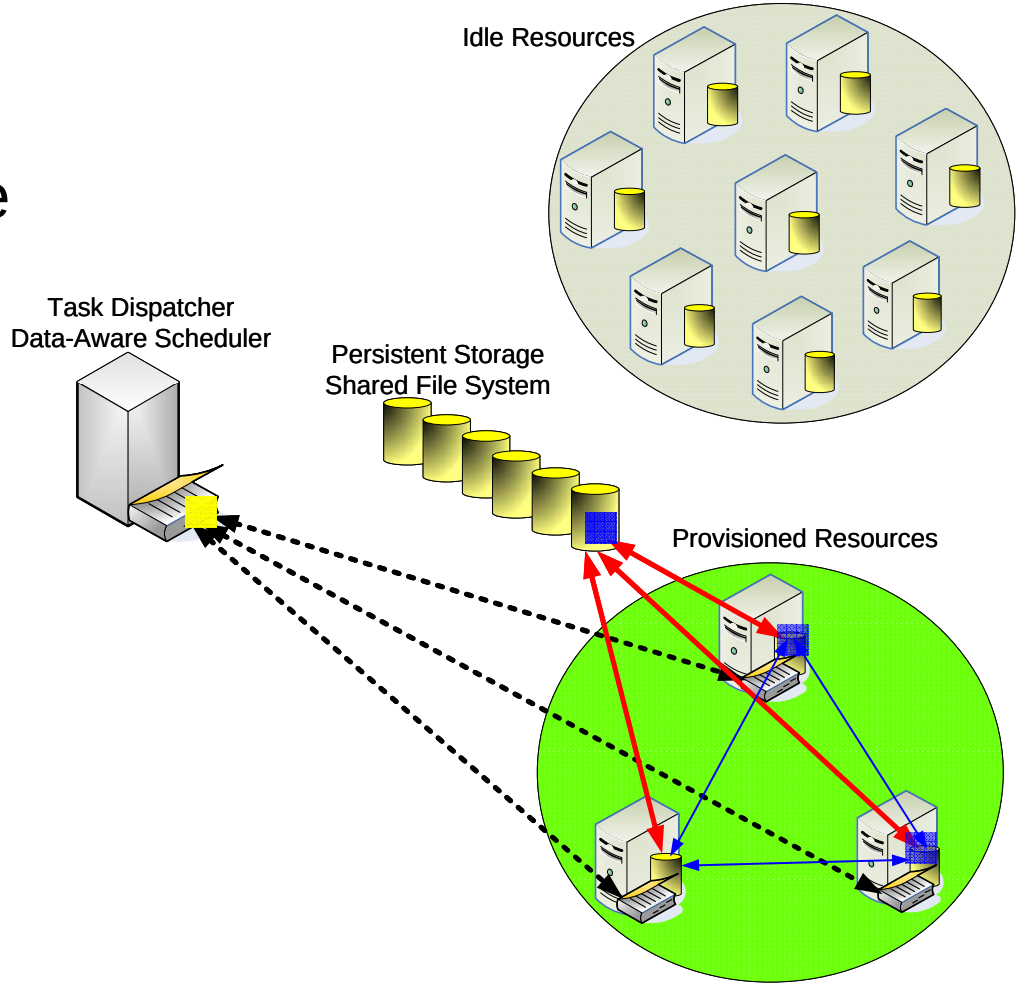

All-Pairs Workload 500x500 on 200 CPUs

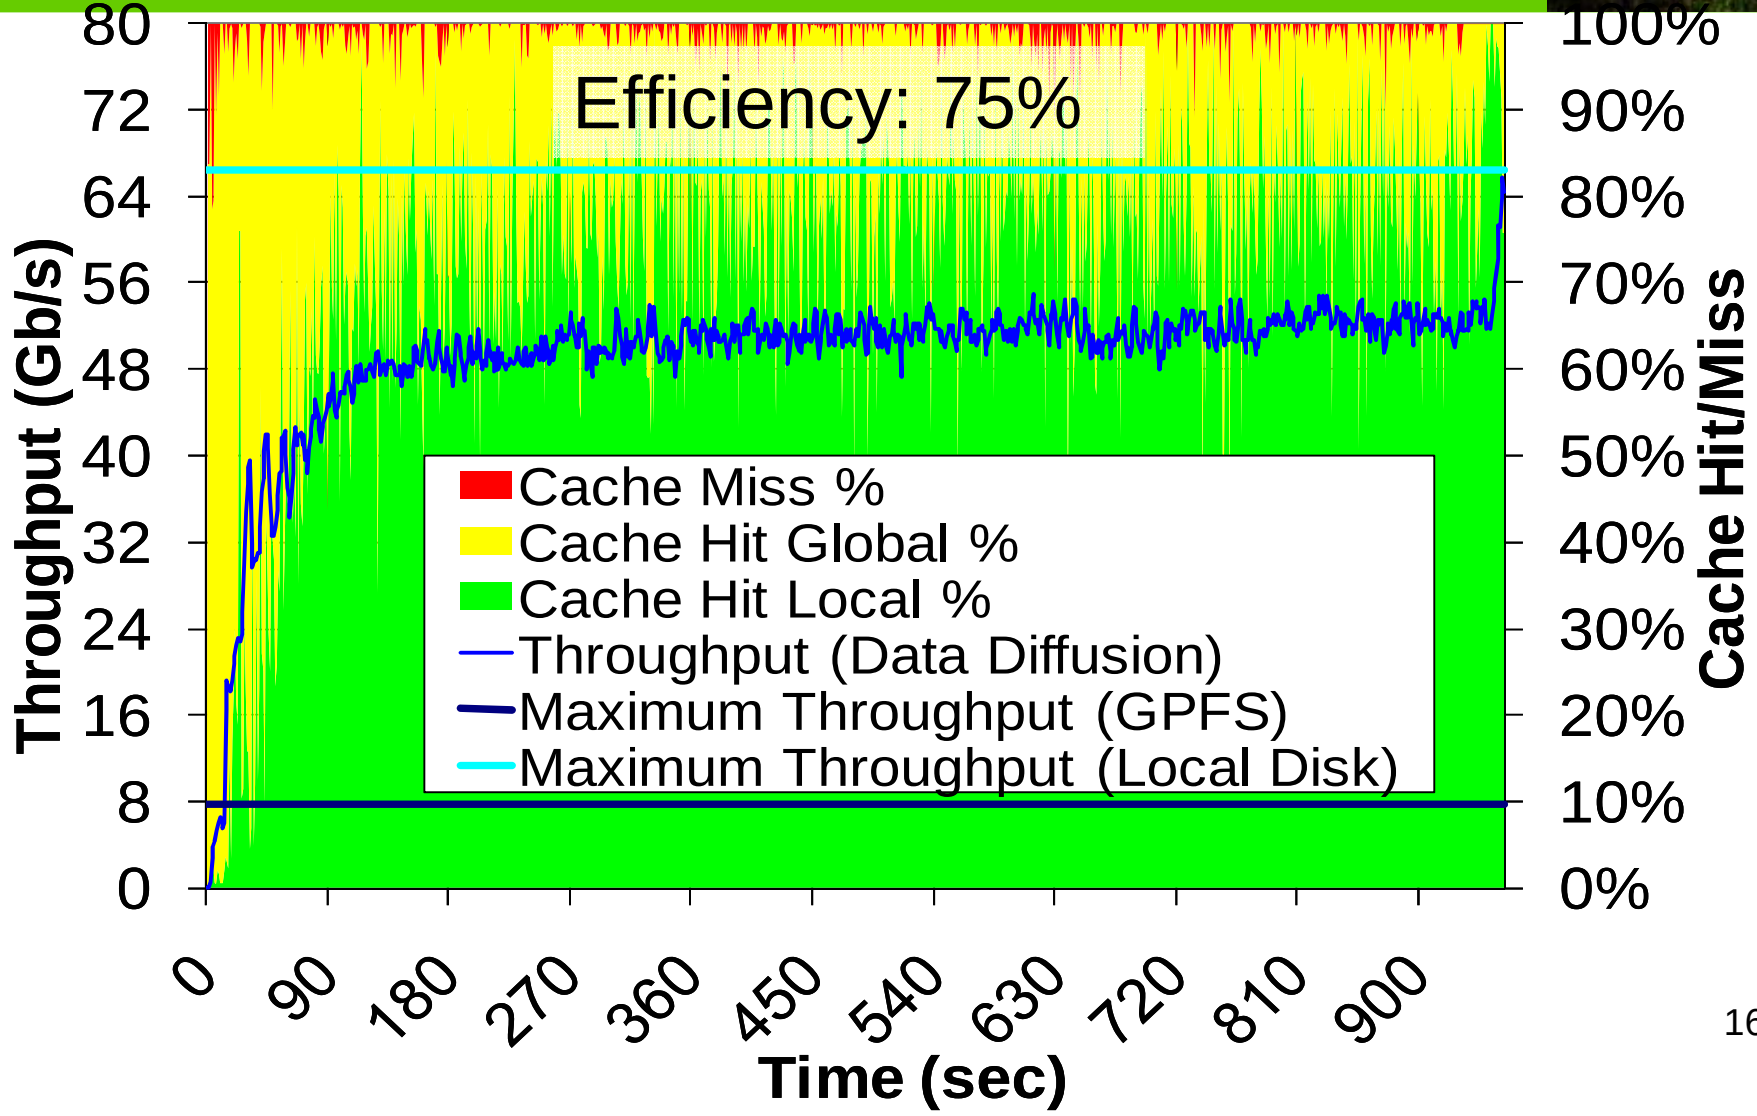

Mythbusting

- ~~Embarrassingly~~ Happily parallel apps are trivial to run
  - Logistical problems can be tremendous
- Loosely coupled apps do not require “supercomputers”
  - Total computational requirements can be enormous
  - Individual tasks may be tightly coupled
  - Workloads frequently involve large amounts of I/O
  - Make use of idle resources from “supercomputers” via backfilling
  - Costs to run “supercomputers” per FLOP is among the best
    - BG/P: 0.35 gigaflops/watt **(higher is better)**
    - SiCortex: 0.32 gigaflops/watt
    - BG/L: 0.23 gigaflops/watt
    - x86-based HPC systems: an order of magnitude lower
- Loosely coupled apps do not require specialized system software
- Shared file systems are good for all applications
  - They don’t scale proportionally with the compute resources
  - Data intensive applications don’t perform and scale well

More Information

- More information: http://people.cs.uchicago.edu/~iraicu/
- Falkon: http://dev.globus.org/wiki/Incubator/Falkon
- Funding:
  - **NASA**: Ames Research Center, Graduate Student Research Program
    - Jerry C. Yan, NASA GSRP Research Advisor
  - **DOE**: Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy
  - **NSF**: TeraGrid
- Check out Falkon:
  - "svn co https://svn.globus.org/repos/falkon"