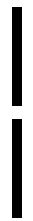


MIT on the Top500: Benchmarking a large-scale heterogeneous cluster

Daniel Loreto, Erik Nordlander, and
Adam Oliner



Benchmarking a large-scale cluster: Why MIT is not on the Top500

Daniel Loreto, Erik Nordlander, and
Adam Oliner

Acknowledgements

- Chris Hill
- Constantinos Evangelinos

Outline



- Benchmark
- Cluster
- Results
- What We Learned

Benchmark: HPL

- High Performance Linpack
- NxN Linear System
- LU Factorization

Benchmark: Settings

- N - problem size
- NB - block size
- PxQ - processor array
- BCAST - message pattern
- ...and many more!



Cluster: MIT ACES Grid

- 4 Clusters
 - 566 CPUs on 344 nodes
 - 829 GB memory
 - Gigabit Ethernet
- $R_{\text{peak}} = 2.706 \text{ TFlops}$
- $N_{\text{peak}} = 750,000$


Cluster: MIT ACES Grid (really)

- 4 Clusters
 - ~~566 CPUs on 344 nodes~~ 305 CPUs
 - ~~829 GB memory~~ 339 GB
 - ~~Gigabit Ethernet~~ Gigabit/Megabit
- ~~$R_{\text{peak}} = 2.706$ TFlops~~ 1.464 TFlops
- ~~$N_{\text{peak}} = 750,000$~~ 480,000


Cluster: MIT ACES Grid (really)

- Frequent outages
- Heavily used
- Moving machines
- Occasionally on fire


Results: Xeons



Results: Pentium 4s



Results: Mixed



Results: Performance Projection

- 305 CPUs: 547 GFlops
 - 1464 GFlops peak
 - 37.4% efficiency
- 566 CPUs: 966 GFlops
 - 2706 GFlops peak
 - 35.7% efficiency

Ultimately...

- Ran HPL heterogeneously
- Achieved 0.442 TFlops
 - 184 CPUs
 - $N = 100,000$
- No Top500

Anticipated Problems

- Heterogeneity
- Network Topology
- Parameter Selection
- Algorithmic Limitations

Actual Problems

- Fire
- Life and Death
- Politics
- Failures, Errors, and Broken Pipes

Benchmarking Rule #1

- Know thy cluster
 - Find out for yourself

Benchmarking Rule #2

- Get Root
 - Installing vicariously is intractable

Benchmarking Rule #3

- Get Keys
 - Physical access may be needed

Benchmarking Rule #4

- Be self-important
 - Benchmarking is top priority
 - Schedule time, kick people off

Benchmarking Rule #5

- Understand Politics
 - Benchmarking is not popular

Contributions

- Compiled and ran HPL across multiple hardware architectures
- Projected performance to the ACES Grid
- Proposed The Five Rules of Benchmarking