High Performance Computing Systems (CMSC714)

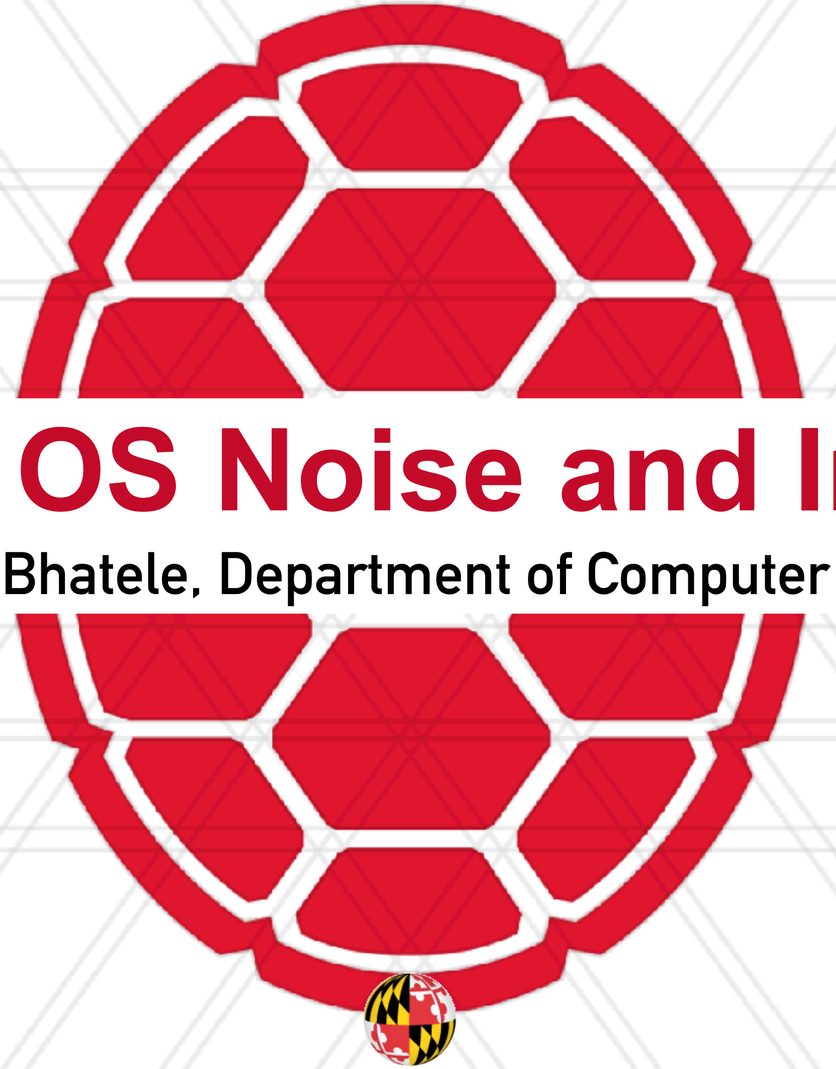

Lecture 17: OS Noise and Interference

Abhinav Bhatele, Department of Computer Science

UNIVERSITY OF MARYLAND

Summary of last lecture

- Dense matrix multiplication is a common operation in HPC codes
- Cannon's 2D algorithm and iso-efficiency for different algorithms
- Agarwal's 3D algorithm

Operating System

- Node on an HPC cluster may have:
  - A "full" linux kernel, or
  - A light-weight kernel
- Decides what services/daemons run
- Impacts performance predictability

Operating System (OS) Noise

- Also called "jitter"
- Impacts computation due to interrupts by OS

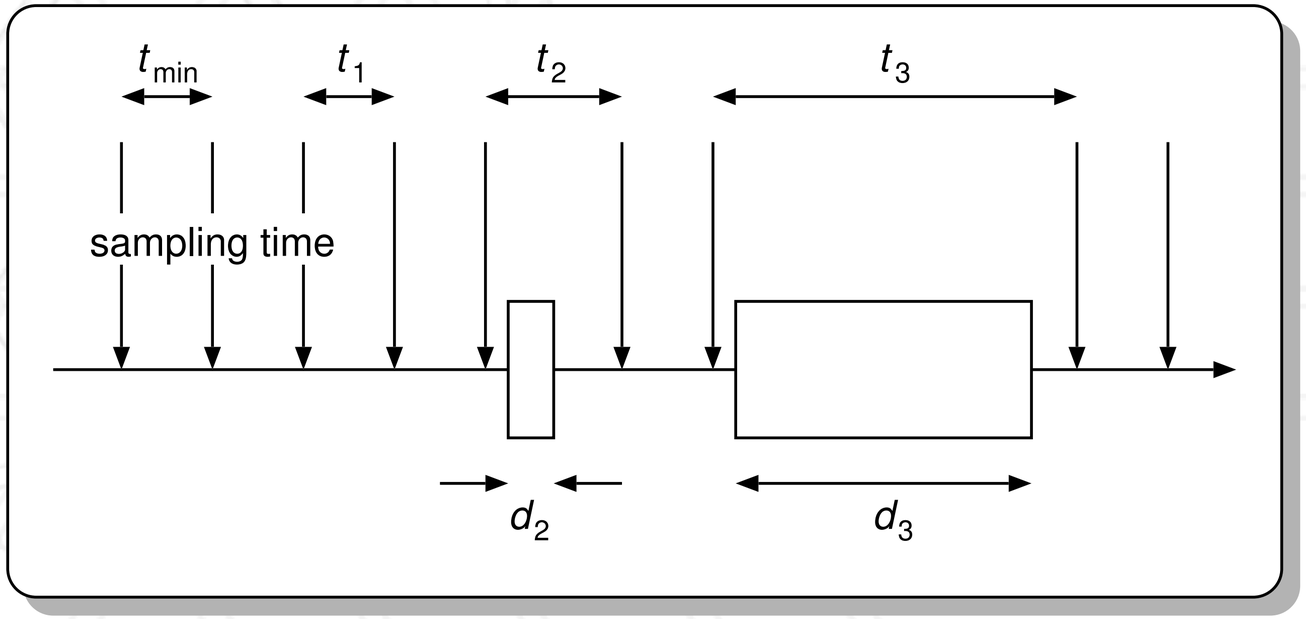

Measuring OS Noise

- Fixed Work Quanta (FTW) and Fixed Time Quanta (FTQ)

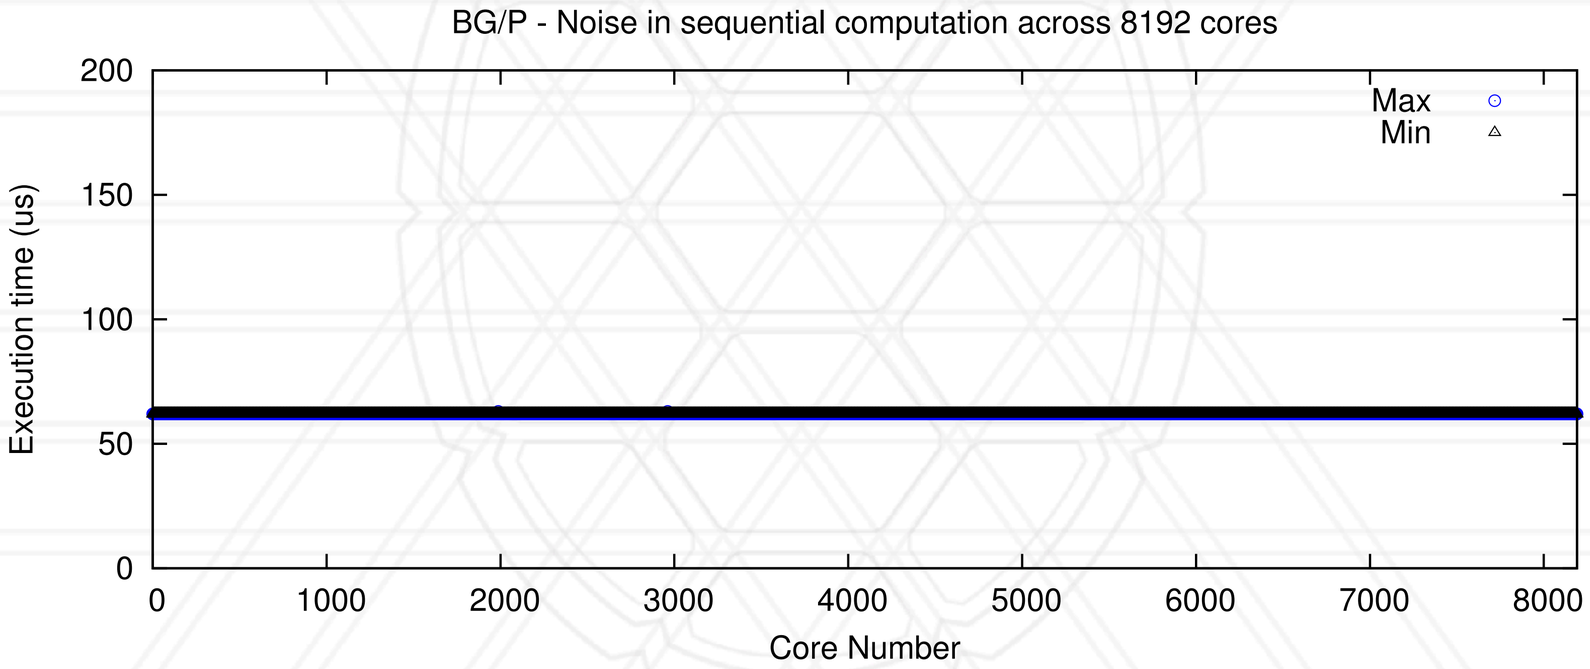

Benchmarks: https://asc.llnl.gov/sites/asc/files/2020-06/FTQFTW_Summary_v1.1.pdf

Measuring OS Noise

- Fixed Work Quanta (FTW) and Fixed Time Quanta (FTQ)

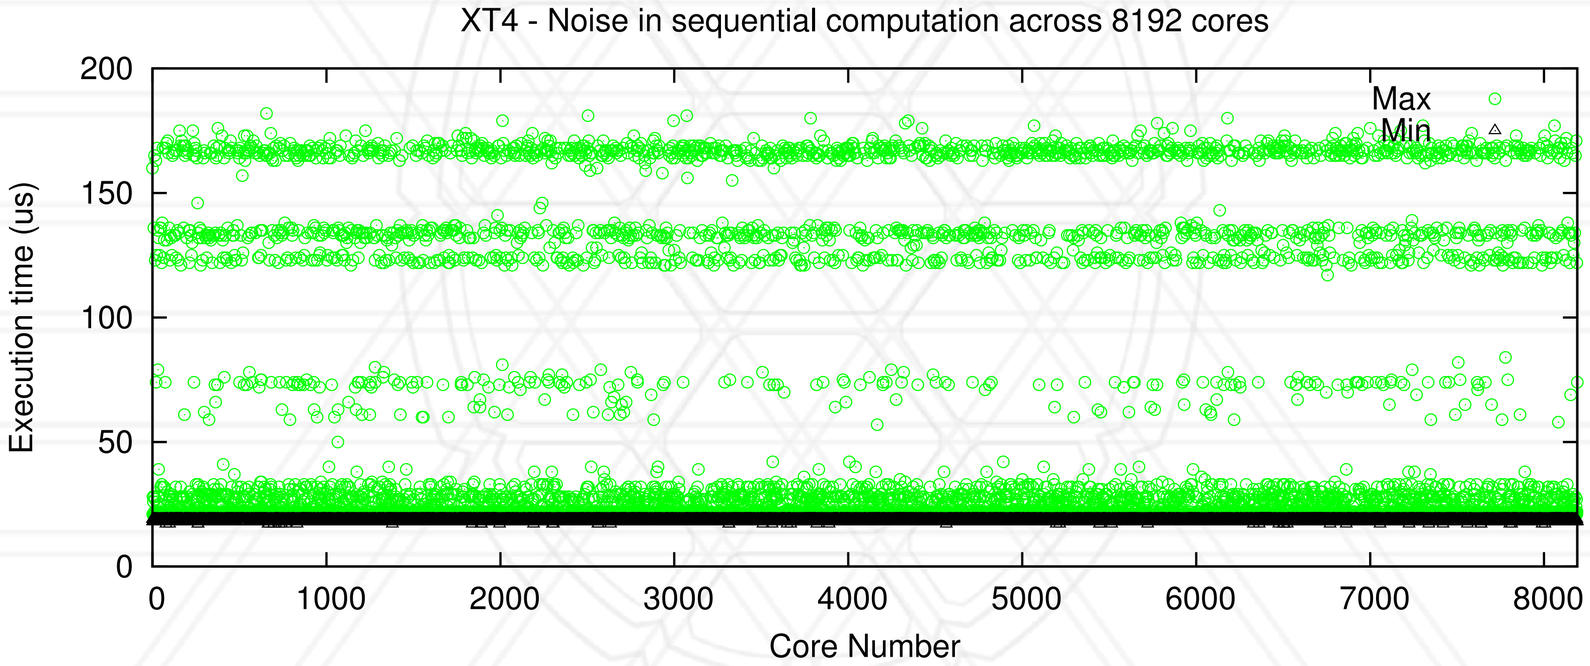

Benchmarks: https://asc.llnl.gov/sites/asc/files/2020-06/FTQFTW_Summary_v1.1.pdf

The Case of the Missing Supercomputer Performance

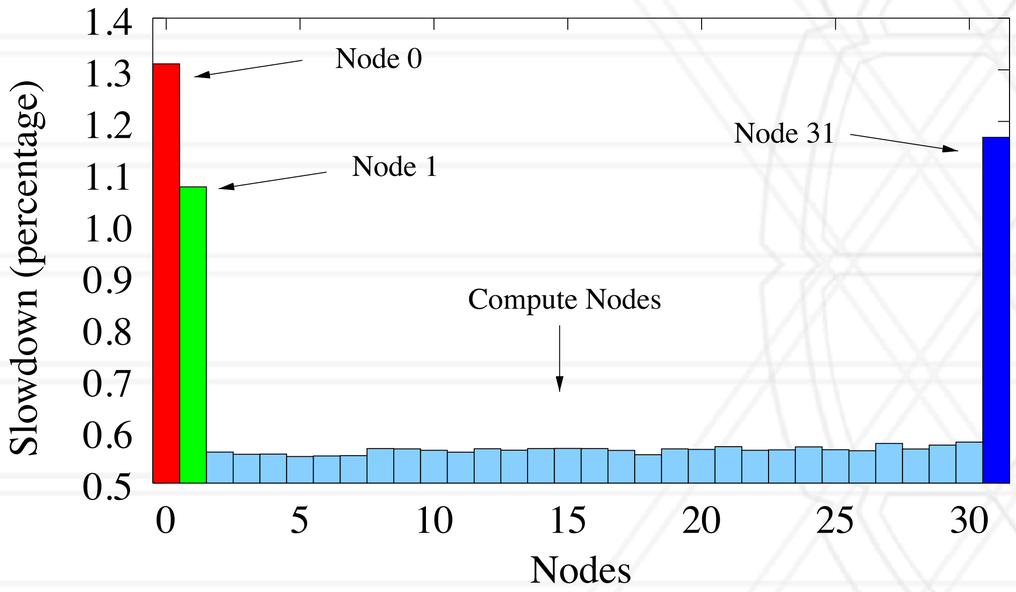

Fabrizio Petrini, Darren J. Kerbyson, and Scott Pakin. 2003. The Case of the Missing Supercomputer Performance: Achieving Optimal Performance on the 8,192 Processors of ASCI Q. In Proceedings of the 2003 ACM/IEEE conference on Supercomputing (SC '03). Association for Computing Machinery, New York, NY, USA, 55. DOI:https://doi.org/10.1145/1048935.1050204

The Case of the Missing Supercomputer Performance

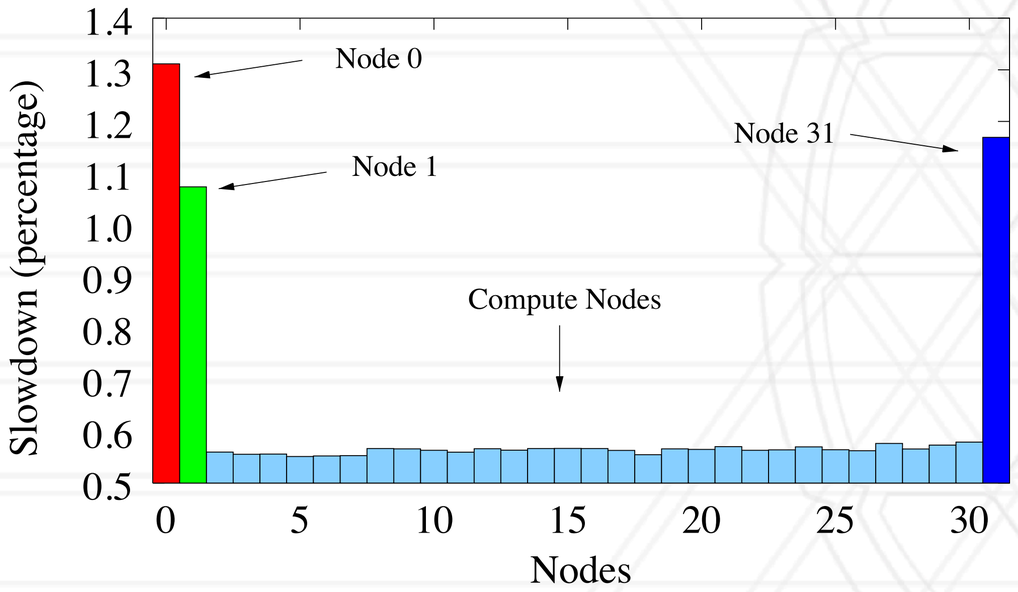

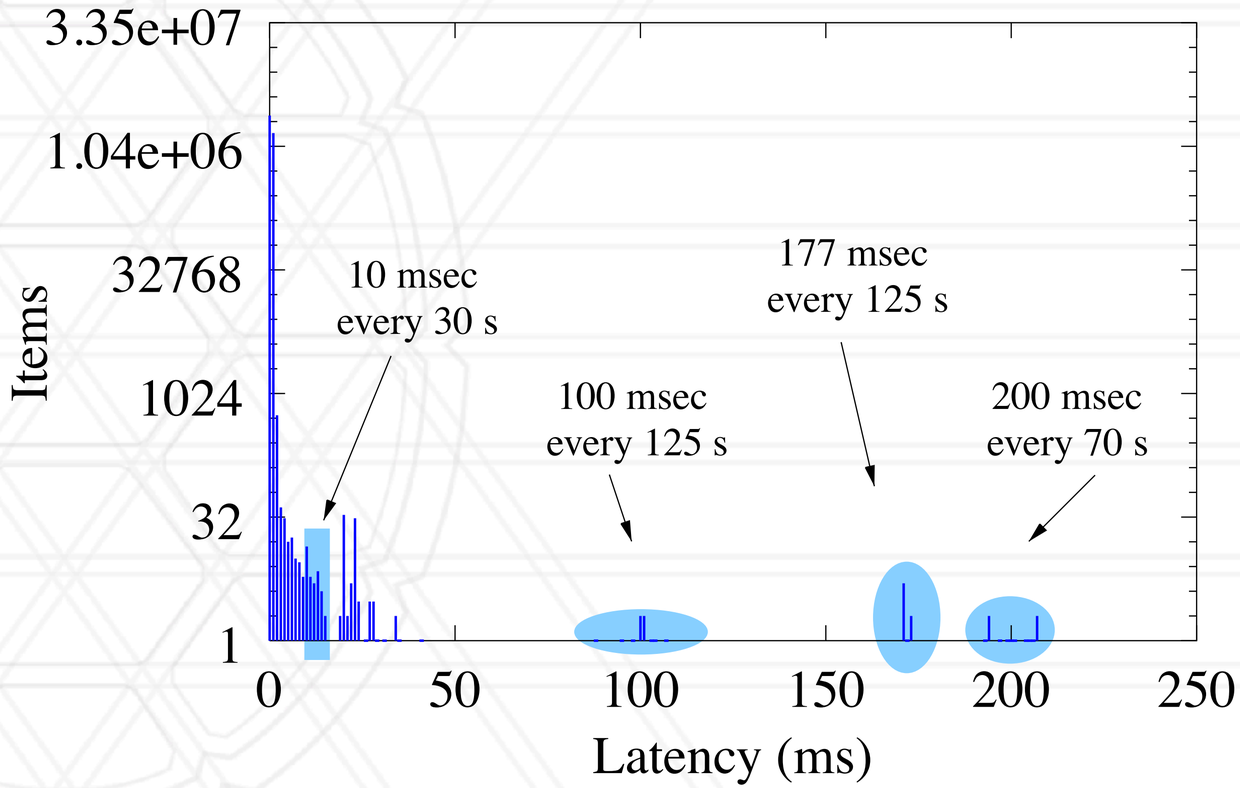

Fabrizio Petrini, Darren J. Kerbyson, and Scott Pakin. 2003. The Case of the Missing Supercomputer Performance: Achieving Optimal Performance on the 8,192 Processors of ASCI Q. In Proceedings of the 2003 ACM/IEEE conference on Supercomputing (SC '03). Association for Computing Machinery, New York, NY, USA, 55. DOI:https://doi.org/10.1145/1048935.1050204

Mitigating OS noise

- Running a light-weight OS
- Turn off unnecessary daemons
- Reduce the frequency of daemons
- Dedicated cores for OS daemons
- User programs can avoid using certain cores

Impact on communication

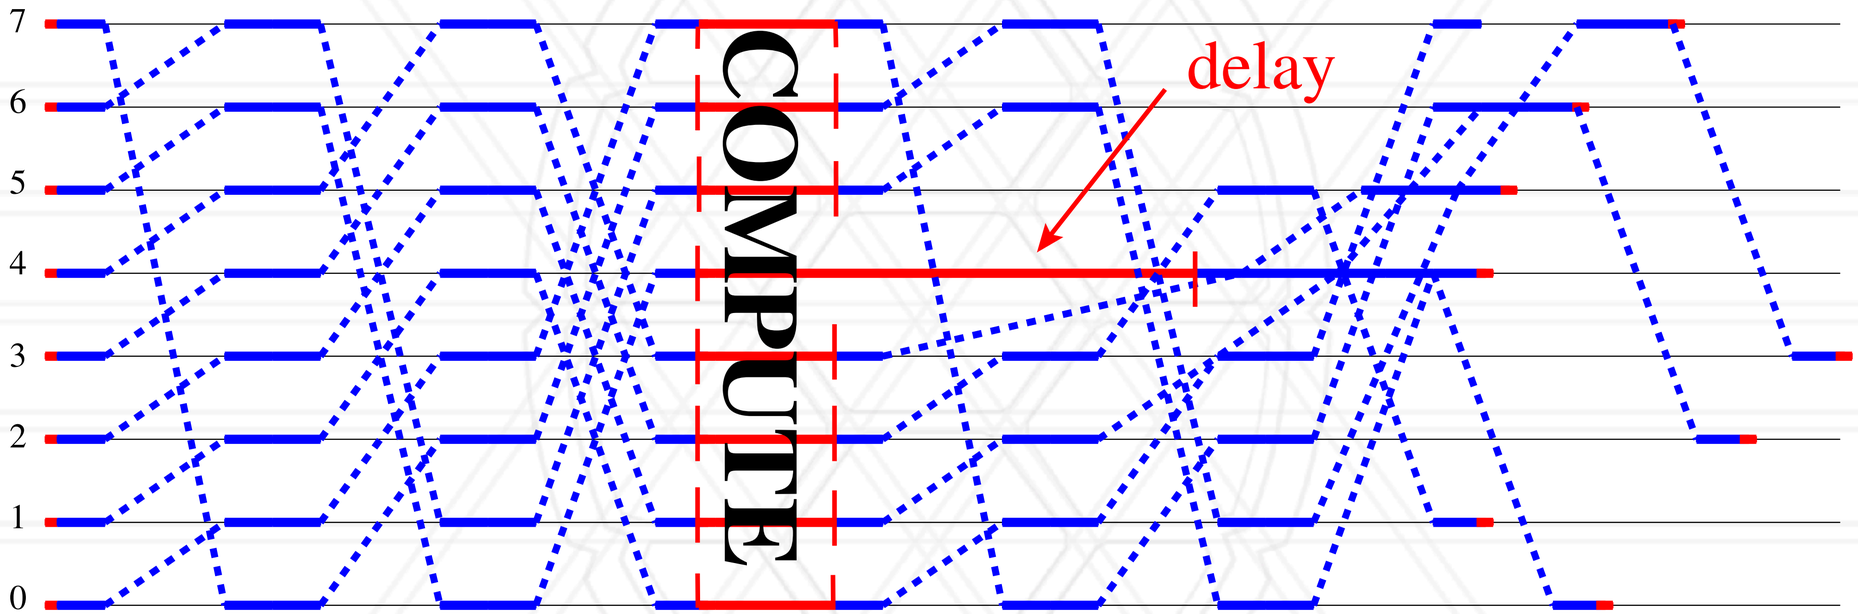

Hoefler et al.: https://htor.inf.ethz.ch/publications/img/hoefler-noise-sim.pdf

Impact on application codes

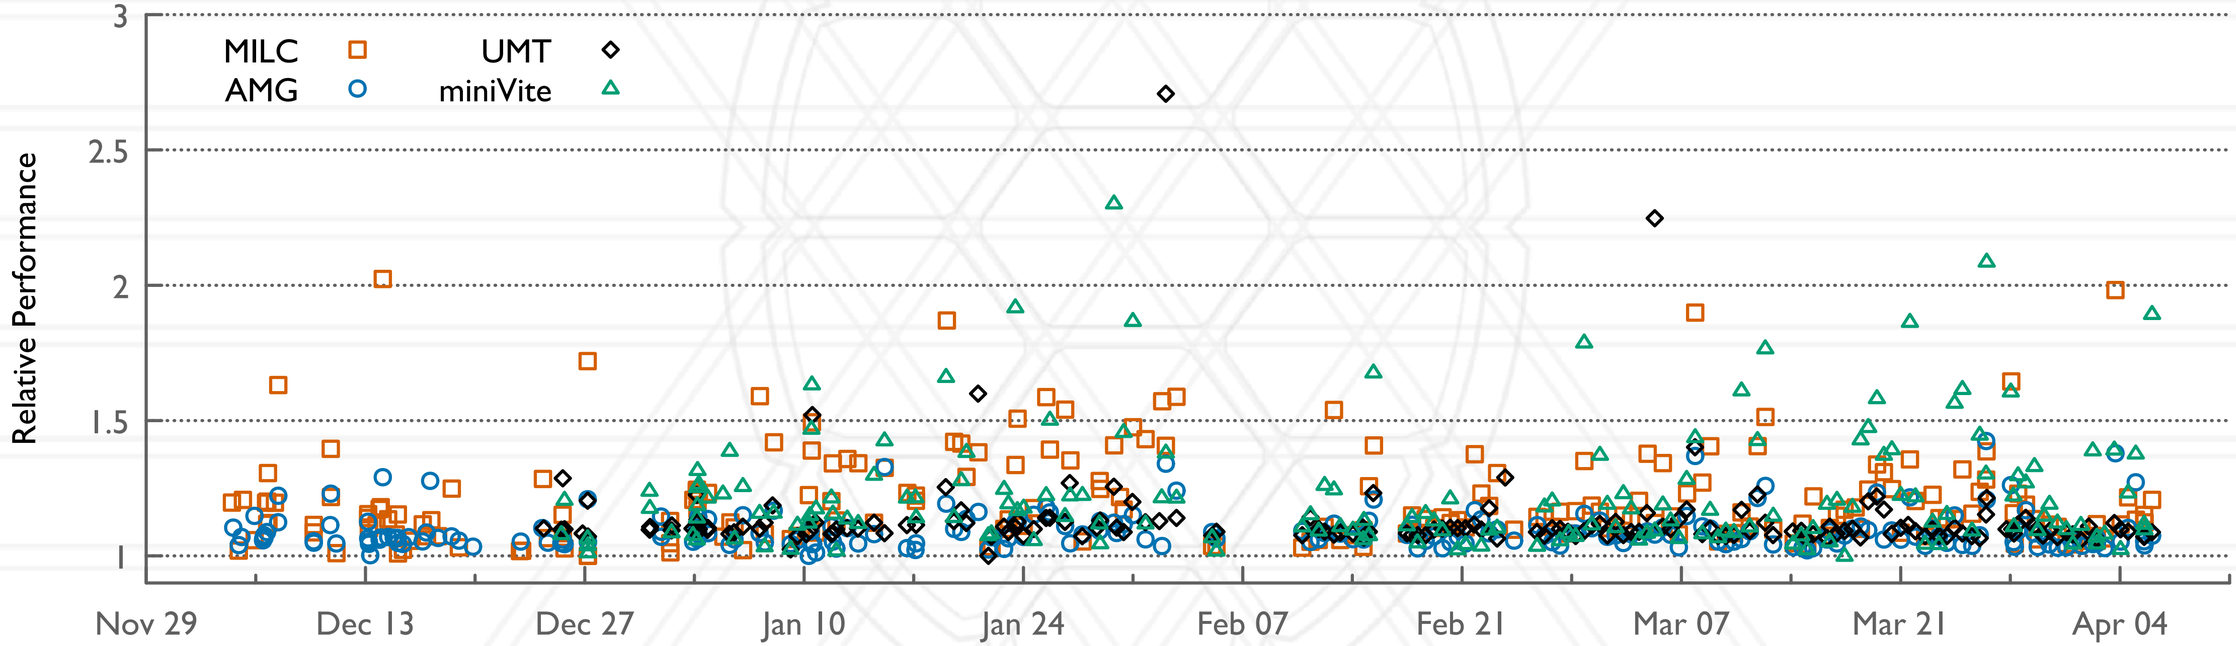

Performance variability due to congestion

- No variability in computation time
- All of the variability can be attributed to communication performance
- Factors:
  - Placement of jobs
  - Contention for network resources

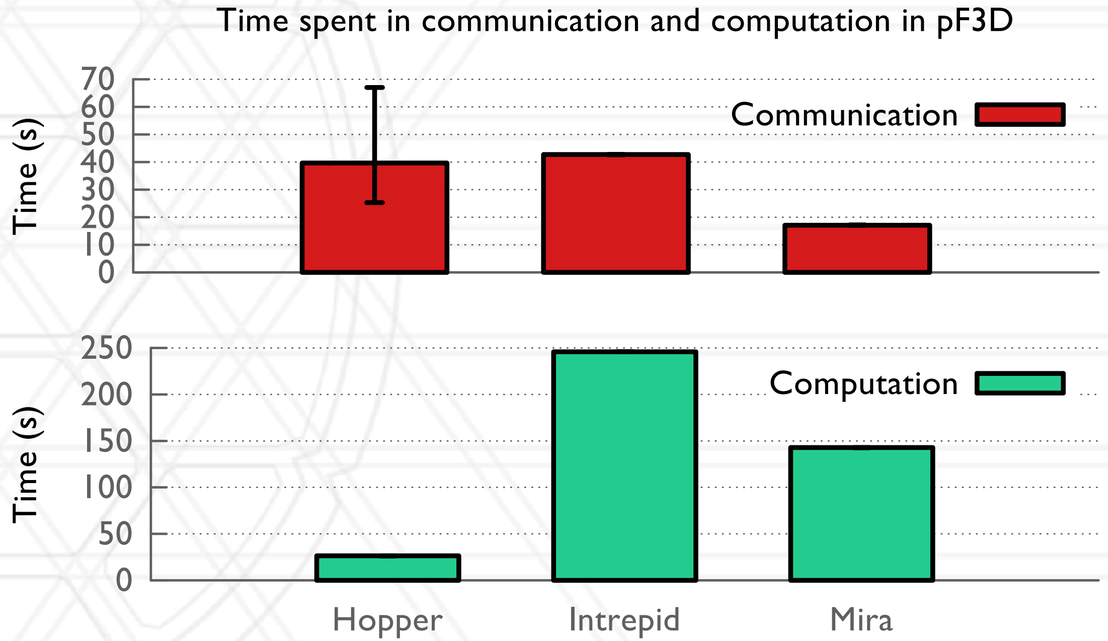

Bhatele et al. http://www.cs.umd.edu/~bhatele/pubs/pdf/2013/sc2013a.pdf

Impact of other jobs

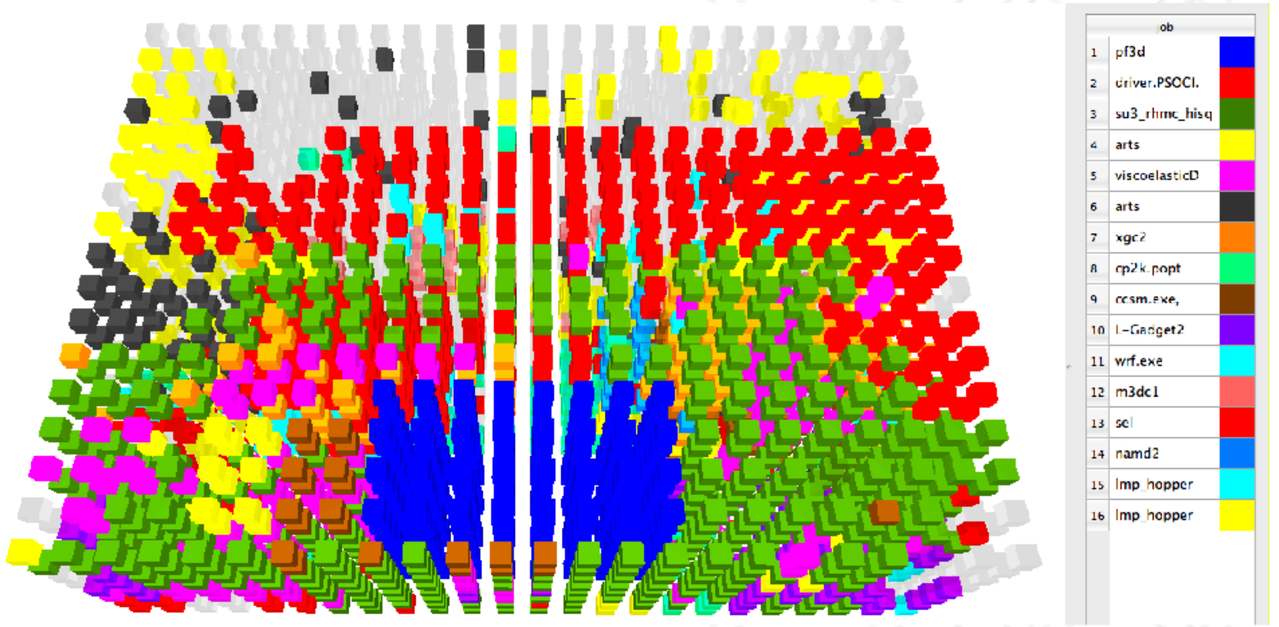

April 11

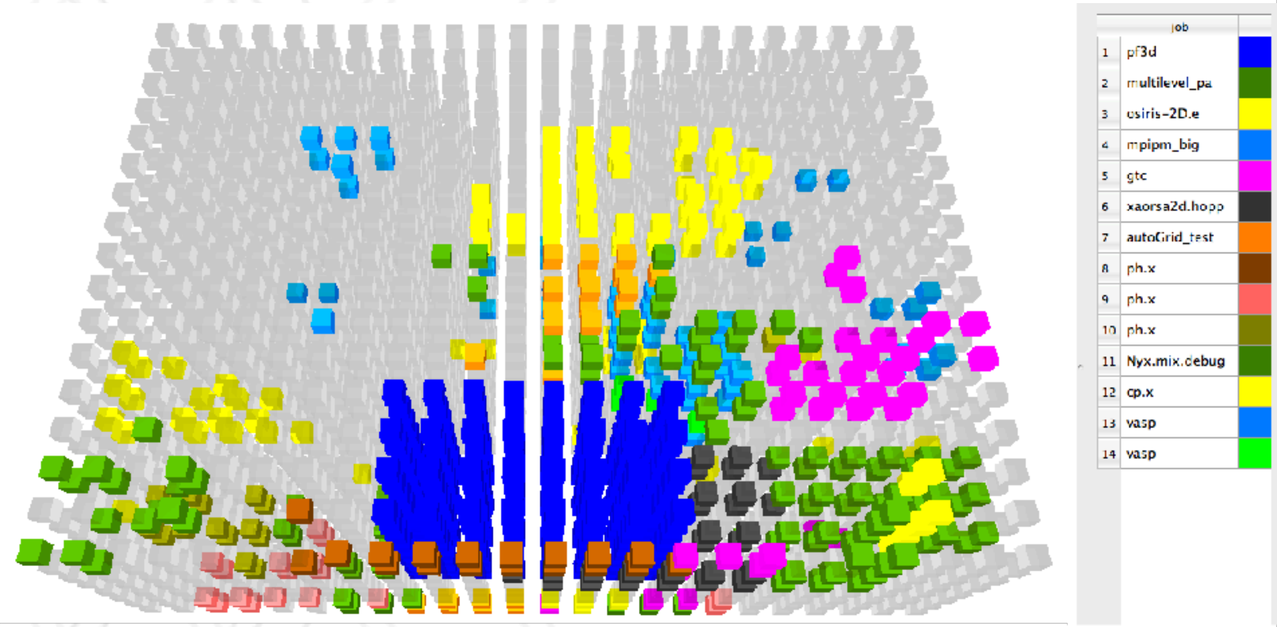

April 16

Impact of other jobs

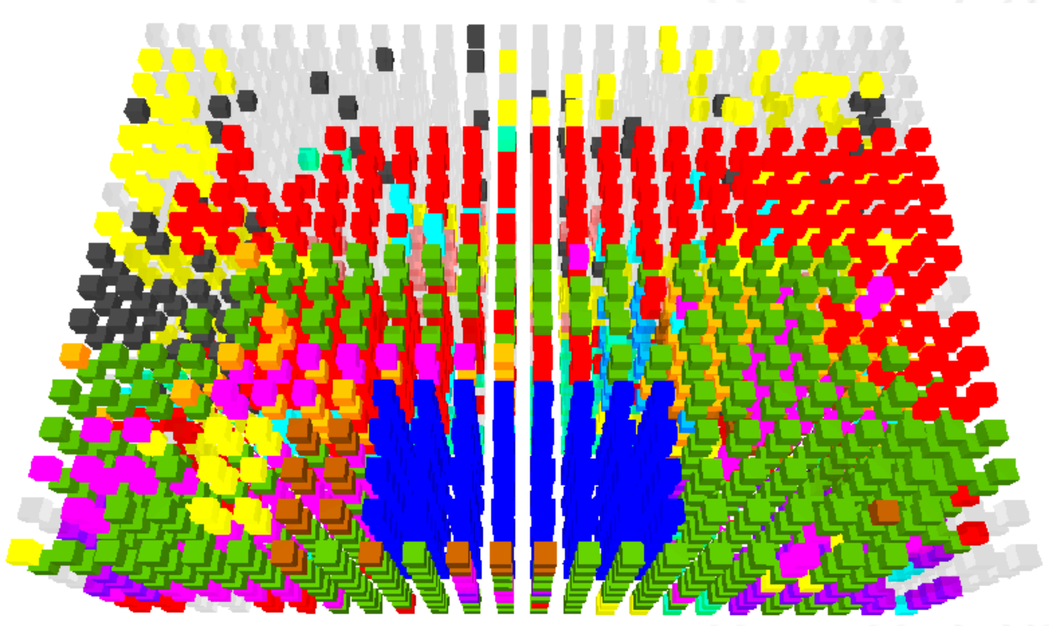

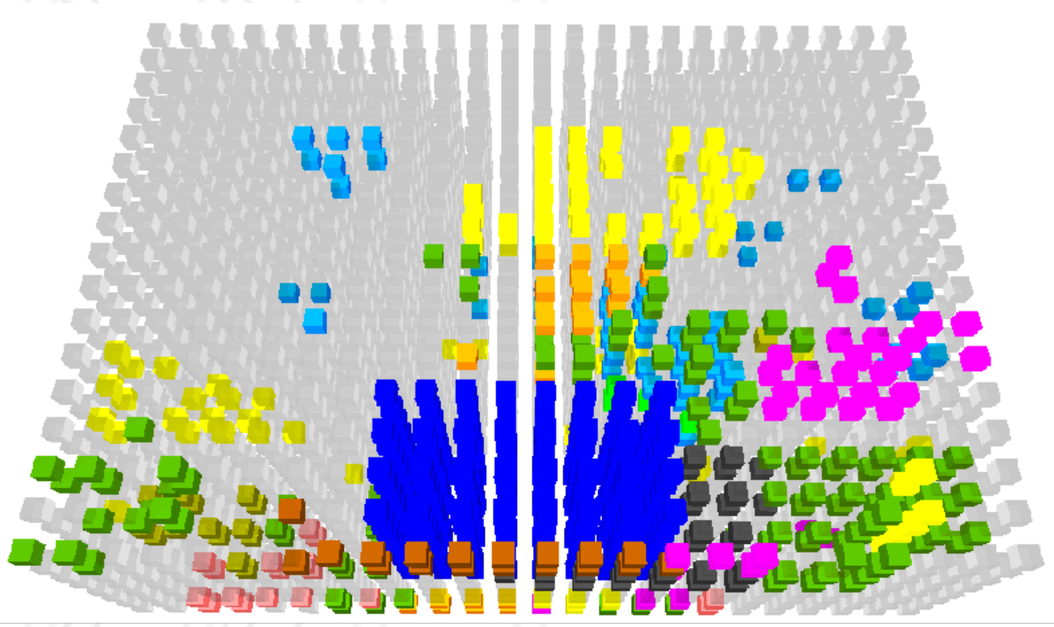

April 11

MILC job in green

April 16

25% higher messaging rate

Leads to several problems ...

- Individual jobs run slower:
  - More time to complete science simulations
  - Increased wait time in job queues
  - Inefficient use of machine time allocation/core-hours
- Overall lower throughput
- Increase energy usage/costs

Affects software development cycle

- Debugging performance issues
- Quantifying the effect of various software changes on performance
  - Code changes
  - System software changes
- Estimating time for a batch job or simulation

Questions?

UNIVERSITY OF
MARYLAND

Abhinav Bhatele

5218 Brendan Iribe Center (IRB) / College Park, MD 20742

phone: 301.405.4507 / e-mail: bhatele@cs.umd.edu