There Goes the Neighborhood: Performance Degradation due to Nearby Jobs

By Abhinav Bhatele, Kathryn Mohror, Steven H. Langer, Katherine E. Isaacs

Presented by: Xiangyu Mao
4/1/2021

Performance Variability & Problems

- **-41% ~ +28% variation**
- Inaccurate job time estimation
- Inefficient scheduling
- Resource contention
- Lower throughput
- Higher energy consumption

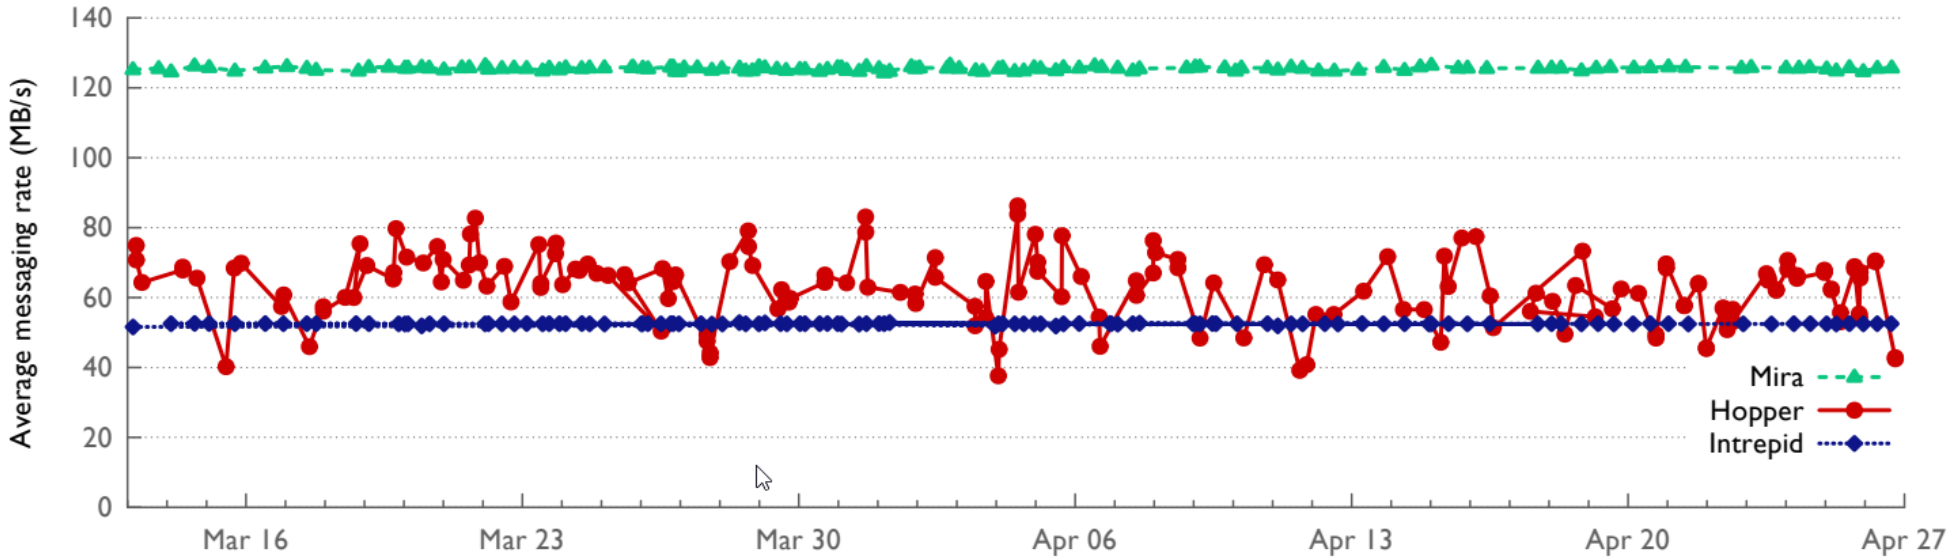

Problems for Developers

- Debugging performance issues
- Quantifying code changes' effectiveness & improvements

  (Usual performance gain: 5~10%)
- Measuring compiler or software changes
- **Multiple runs required for averaging**
- Job time estimation and submission

External Factors

- OS Daemons / Noise / Jitter
- Process Placement
- Interconnect Traffic
- File System Traffic
- I/O Traffic

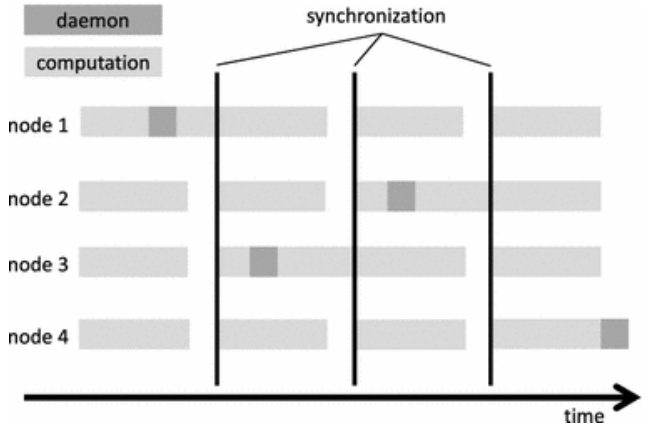

Yoshihiro Oyama, 2016

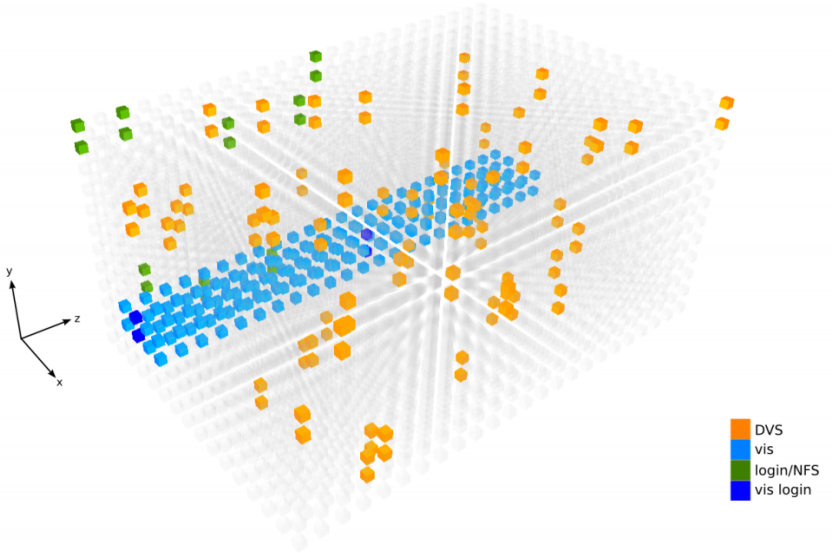

Experimental Setup

- **Machines**
  - IBM Blue Gene/Q
  - IBM Blue Gene/P
  - Cray XE6
- **Program: pF3D**
  - Laser & Plasma Simulation
  - Scalable
  - Good computation & communication load balance

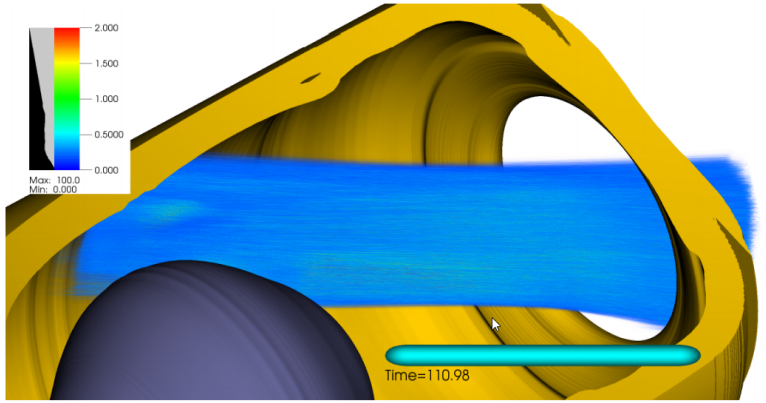

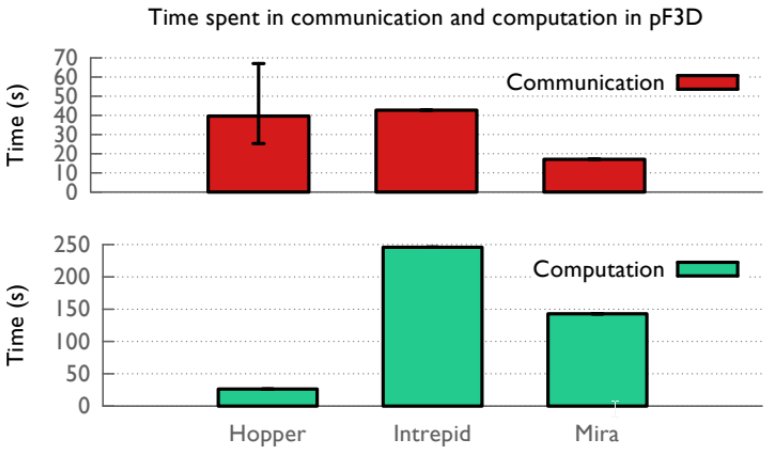

Cray XE6 (Cielo) Performance

- **CR1: (Cielo)**
  - 1x2x8 nodes
  - 80~100 MB/s
- **CR2:**
  - 2x2x4 nodes
  - 65~155 MB/s
- **CR3:**
  - 2x2x8 nodes
  - 35~100 MB/s
- **DR1&2: (Dawn)**
  - Less than 1%
  - Same partition size & shape

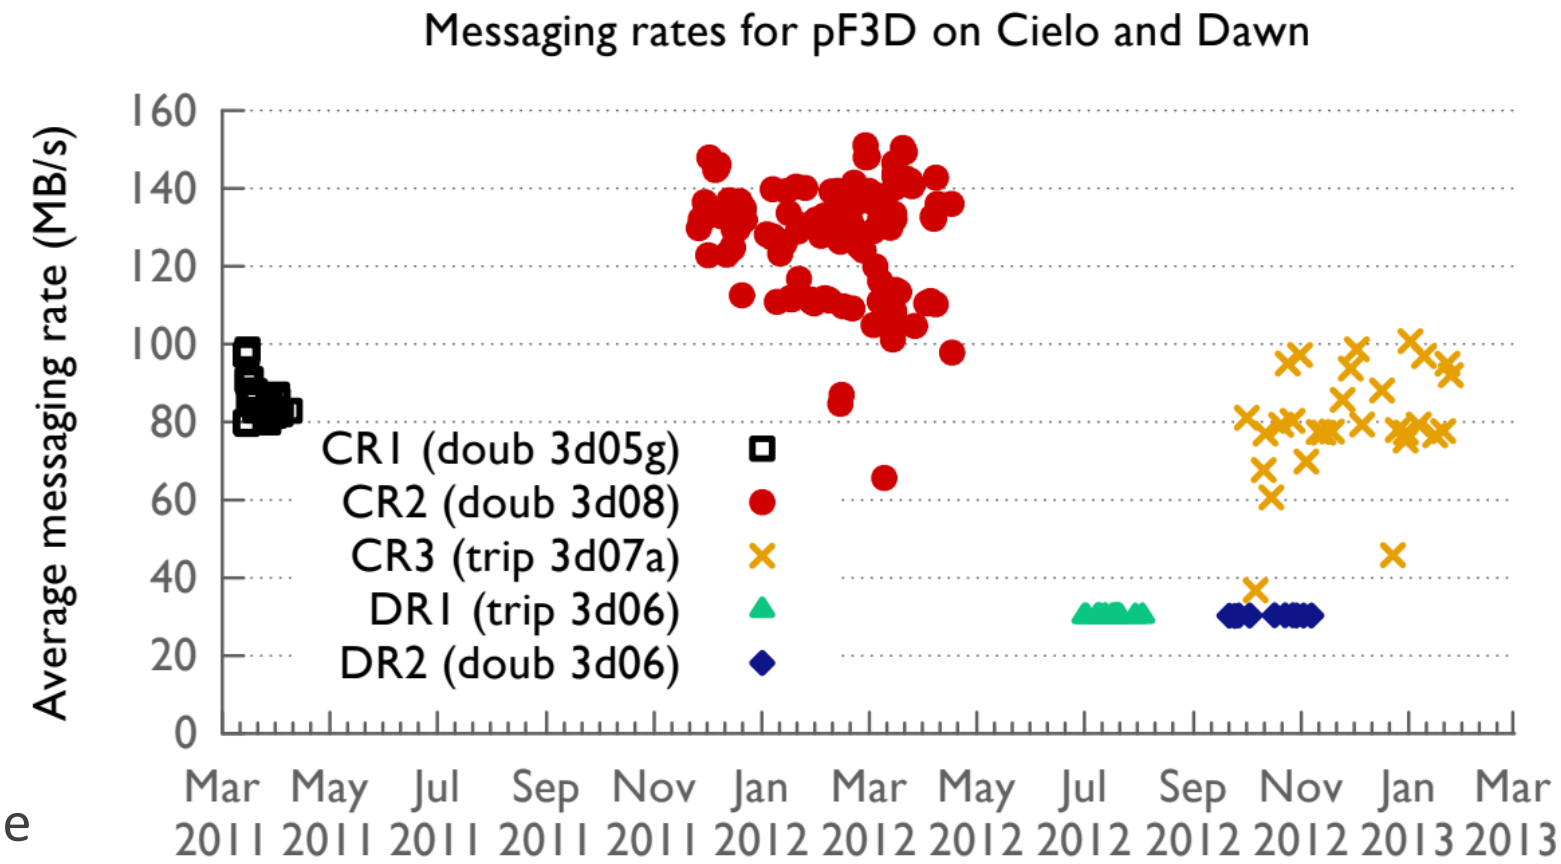

OS Jitter (Noise)

- 16 out of 24 nodes used (on Hopper)
- Core 0: Kernel activities
- Variation < 0.3%

  (Between using & NOT using core 0)
- **Insignificant variability**

  (compared to overall performance)

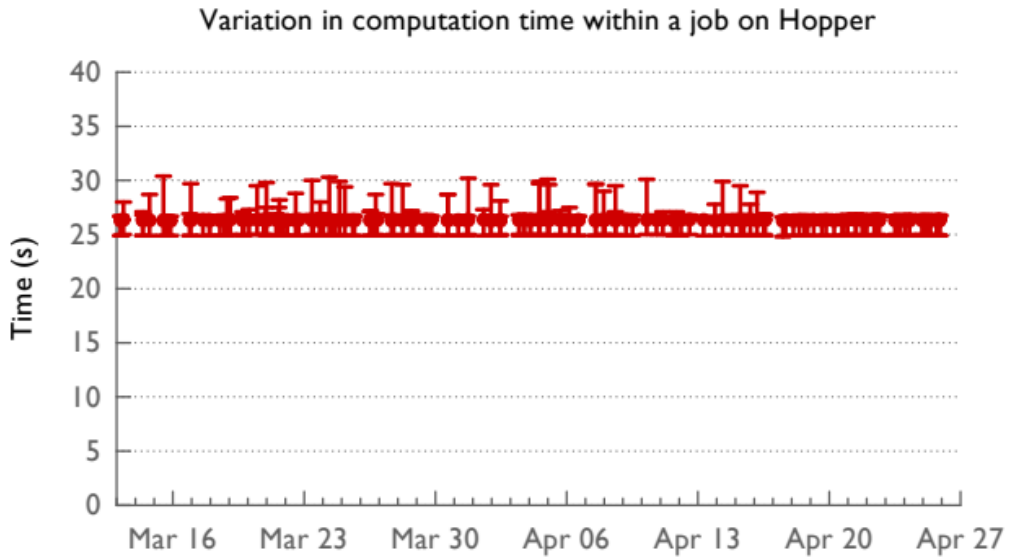

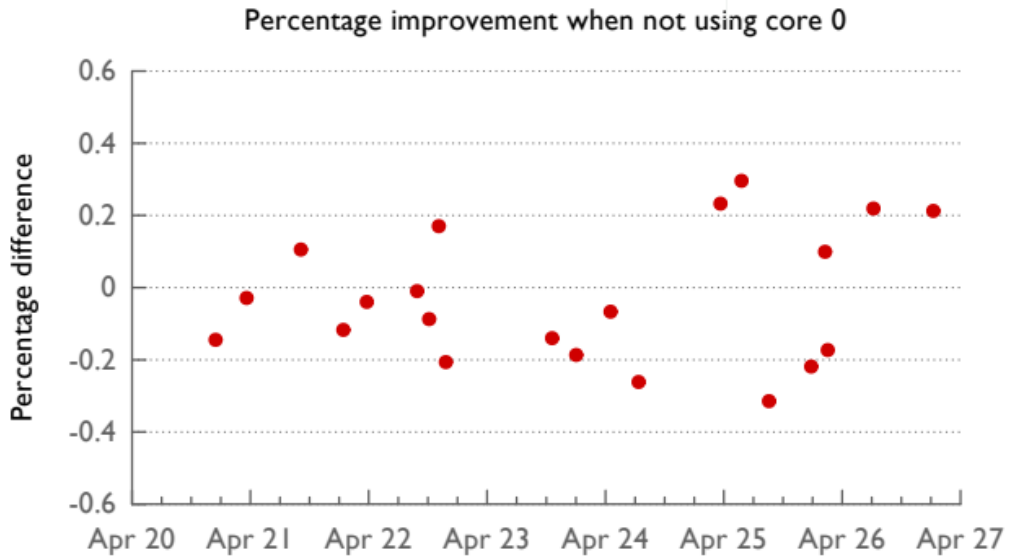

Nodes Contiguity (Compactness)

- Contiguity Metric =

  #nodes / #nodes in 3D mesh
  (farthest corners)

- 0: Most Sparse
- 1: Most Compact
- **Insignificant correlation**

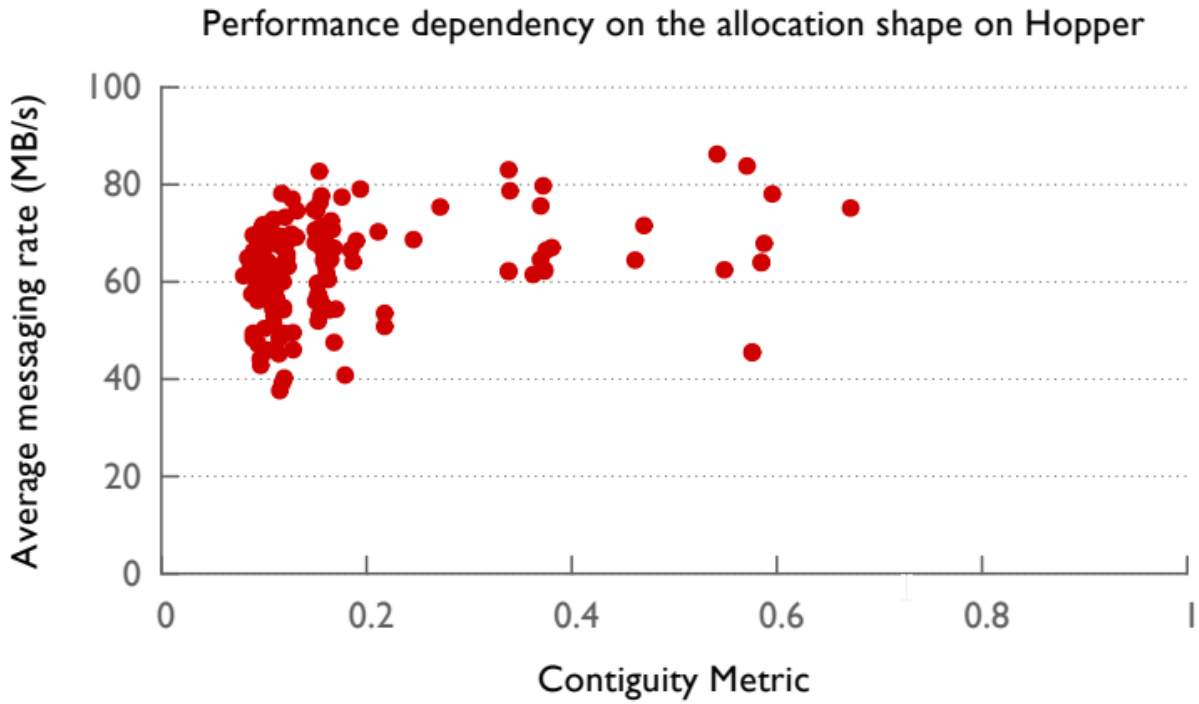

y-Communicator Hops

- x-communicator: shared memory (less variability)
- y-communicator: hops across links
- More hops -> Lower rate
- **Weak correlation**

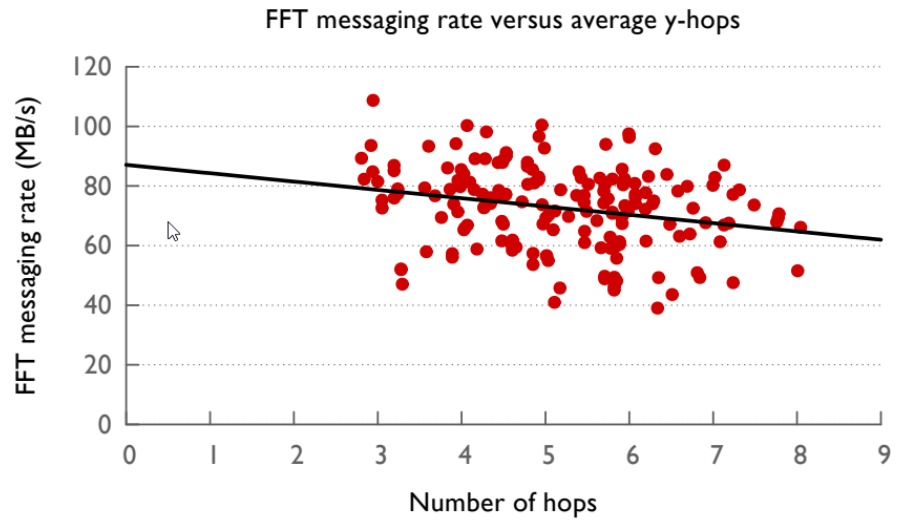

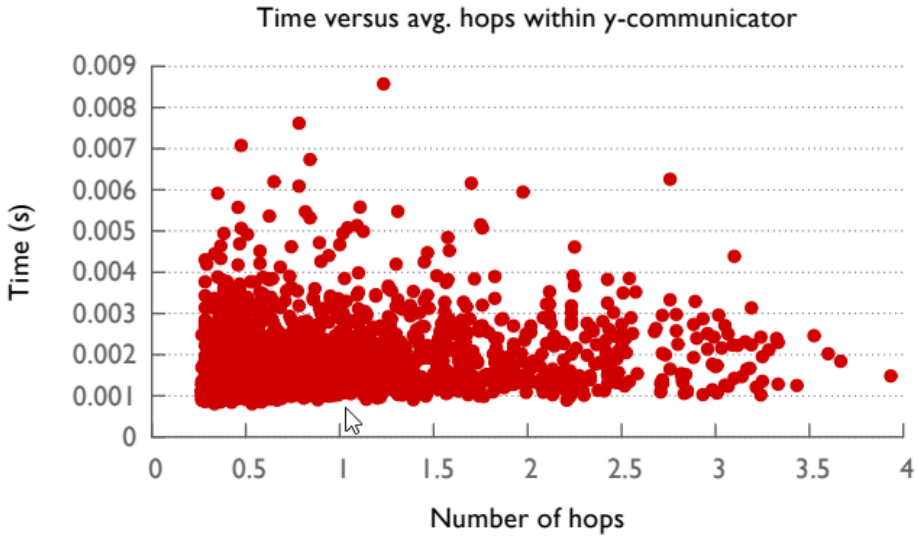

Contention from Other Jobs

- MILC: comm heavy application
- Contention for links
- w/ MILC: 58.0MB/s avg.
- w/o MILC: 66.0MB/s avg.

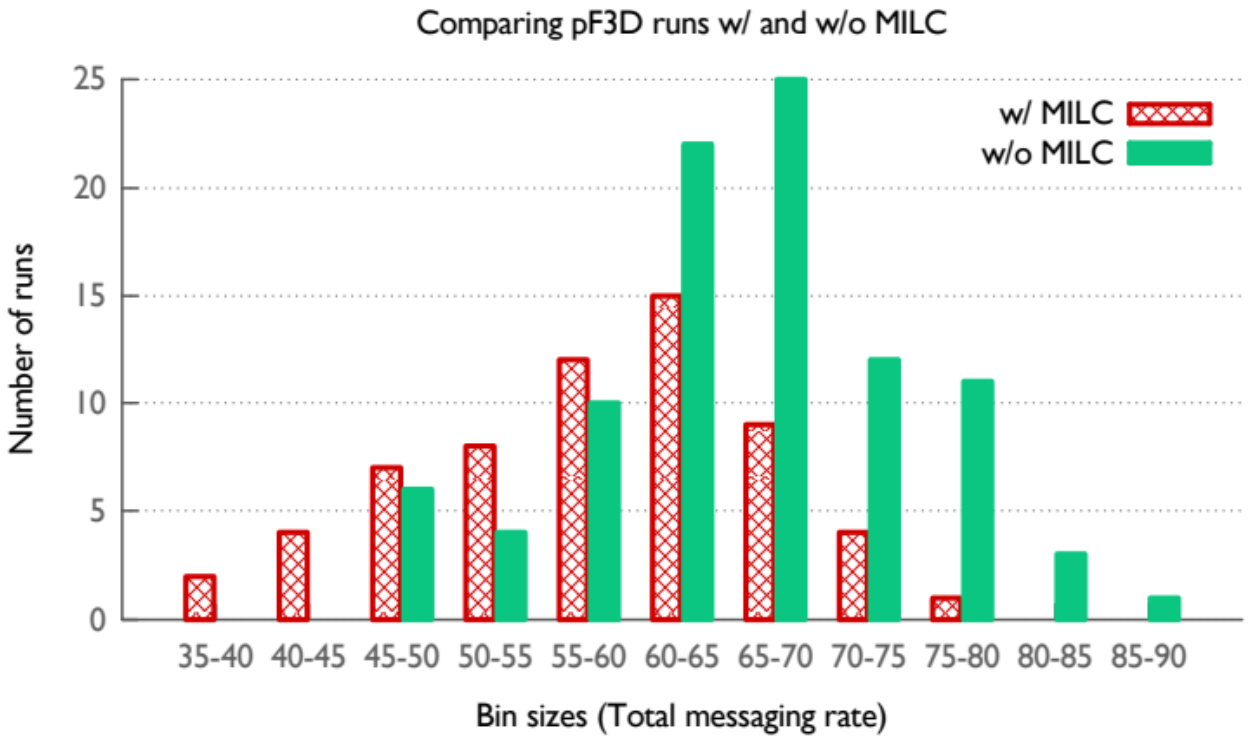

Visualization of Job Placement

pF3D (Blue) surrounded

25% Lower messaging rate

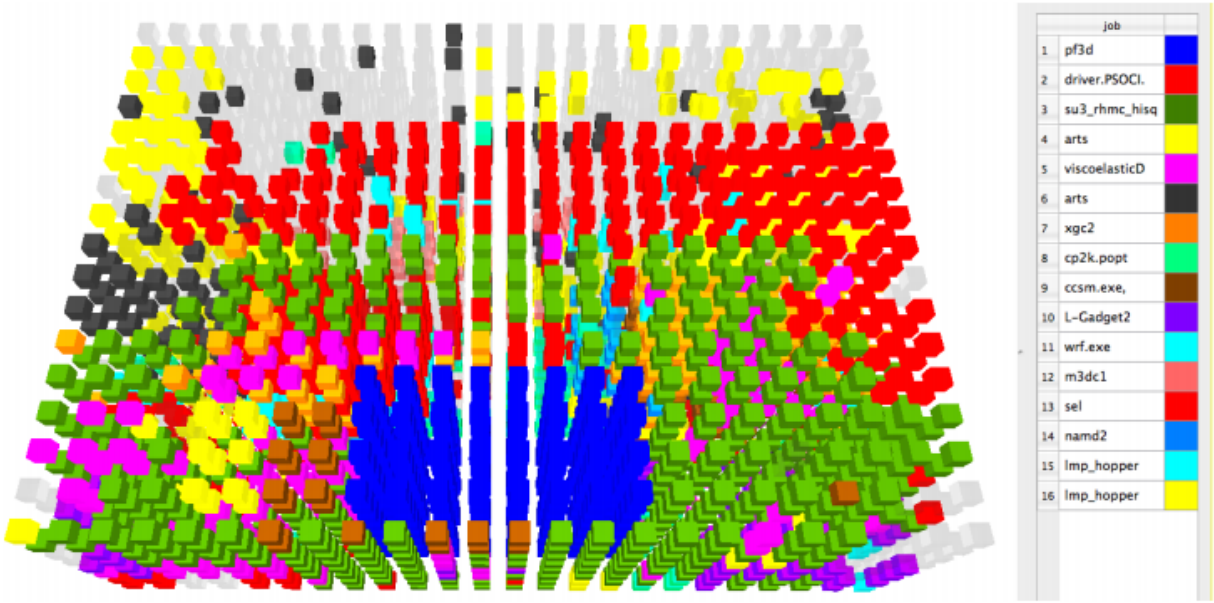

pF3D (Blue) not surrounded

High messaging rate

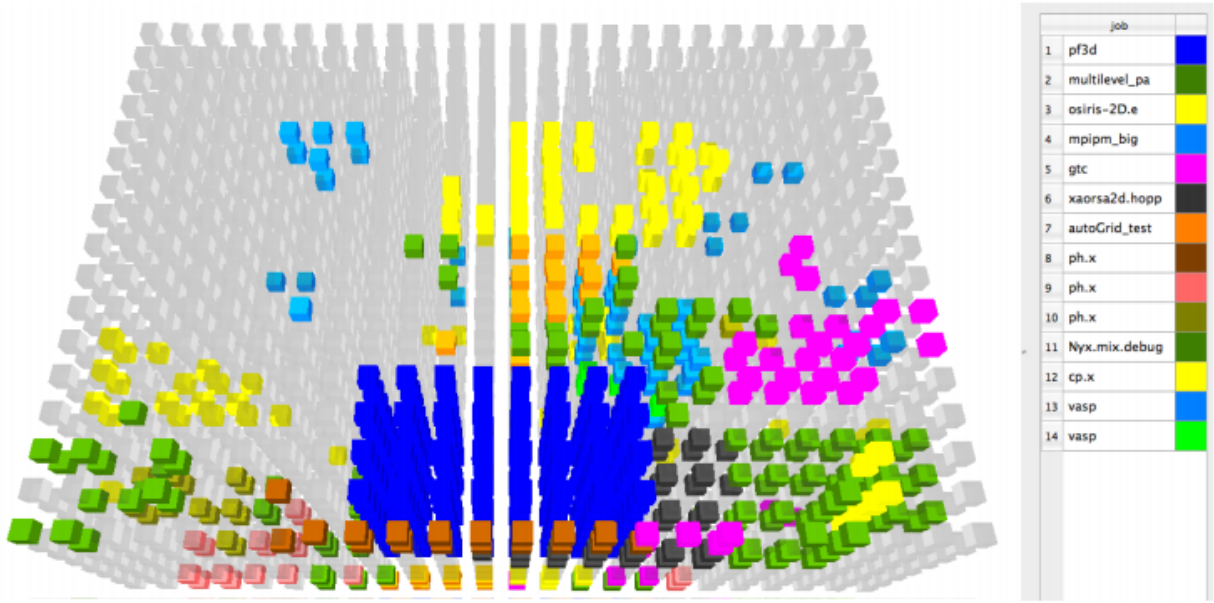

Visualization of Job Placement

pF3D (Blue) interspersed

2.29x Lower messaging rate

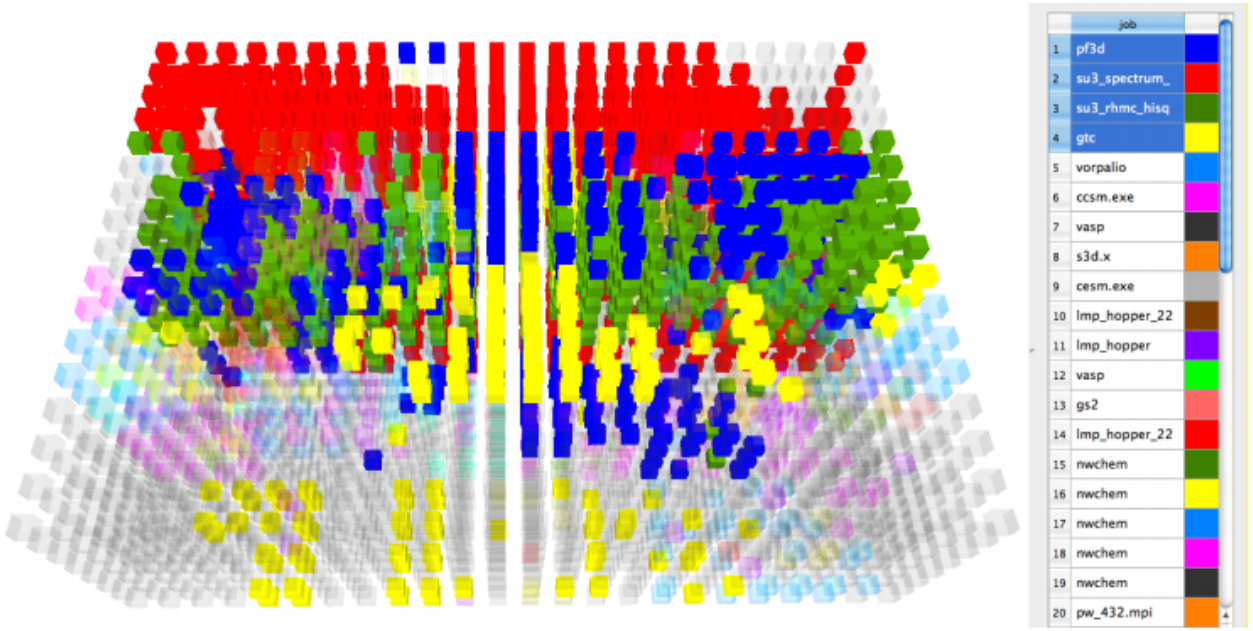

pF3D (Blue) not separated

High messaging rate

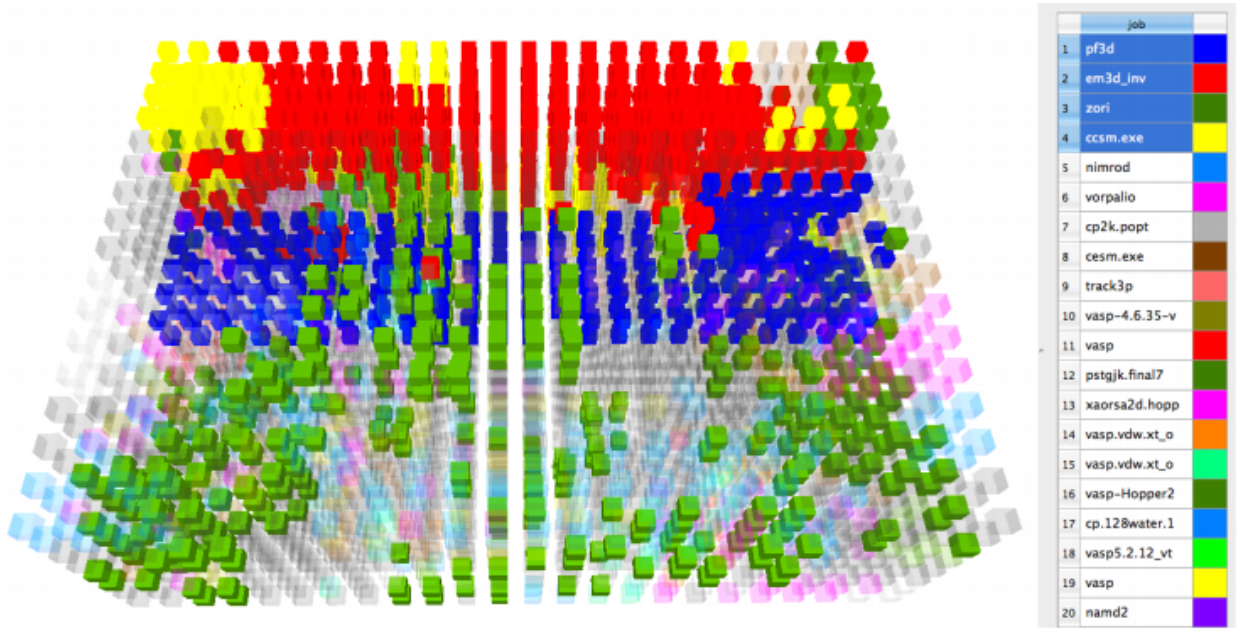

Conclusion

- OS Jitter: **Minor variability**
- Contiguity (Compactness): **Weak correlation**
- y-Communicator Hops: **Weak correlation**
- Contention with other jobs
  - **Surrounded**: **25% Slower**
  - **Interspersed**: **2.29x Slower**