Exploring Application Performance on Fat-Tree Networks in the Presence of Congestion

Philip Taffet, Sanil Rao, Ian Karlin

Abstract

Network congestion, which occurs when multiple applications simultaneously use shared links in cluster network, can cause poor communication performance, decreasing the performance and scalability of parallel applications. Many studies are performed while clusters also run other production workloads, which makes it harder for them to isolate causes and their effects. To look at congestion in a more controlled setting we used dedicated access time on an HPC cluster and measured the performance of three HPC applications with different communication patterns run with varying amounts and types of background traffic. This enables us to assess the relative sensitivity of the applications to congestion caused by different traffic patterns. Our tests show that the applications were not significantly impacted by even the most aggressive neighboring patterns, with all the performance degradation being 7% or less, pointing to the resiliency of the fat-tree topology.

Application Overview

Application	Physics	Grid Type
UMT	Radiation Transport	3D
AMG	Multi-grid used forElastic and Plastic Deformations	3D unstructured Grid
pF3D	Laser-plasma Interactions	3D Structured Grid

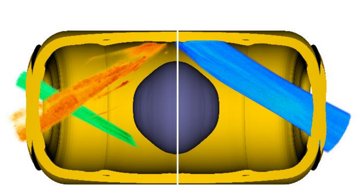

hypre

Communication Characteristics

Application	Communication Pattern	% Time spent in MPI⁺	Average Message Size for Primary Comm. Pattern	Message Count
UMT	7 pt Stencil	45%	53 kB	1.2M
AMG	27 pt Stencil + AllReduce	26%	4 kB	404M
pF3D	FFT	-	-	-

Applications were run on the Quartz system with 224 nodes and 36 tasks per node.
⁺Some of this is due to load imbalance and/or noise
Data for pF3D will be collected for the final poster, but typically the FFTs are 10's of KB

System Overview

Aspect	Description
Network Topology	3-level fat-tree w/ 2:1 taper
Interconnect	100 Gbps Omni-Path
Number of nodes	2688
Processor	Dual 18 core 2.1 GHz Broadwell

Experiments

Control

- Application runs on 8 leaf switches and 28 nodes per leaf switch
- Other nodes are idle
- No links shared with other applications

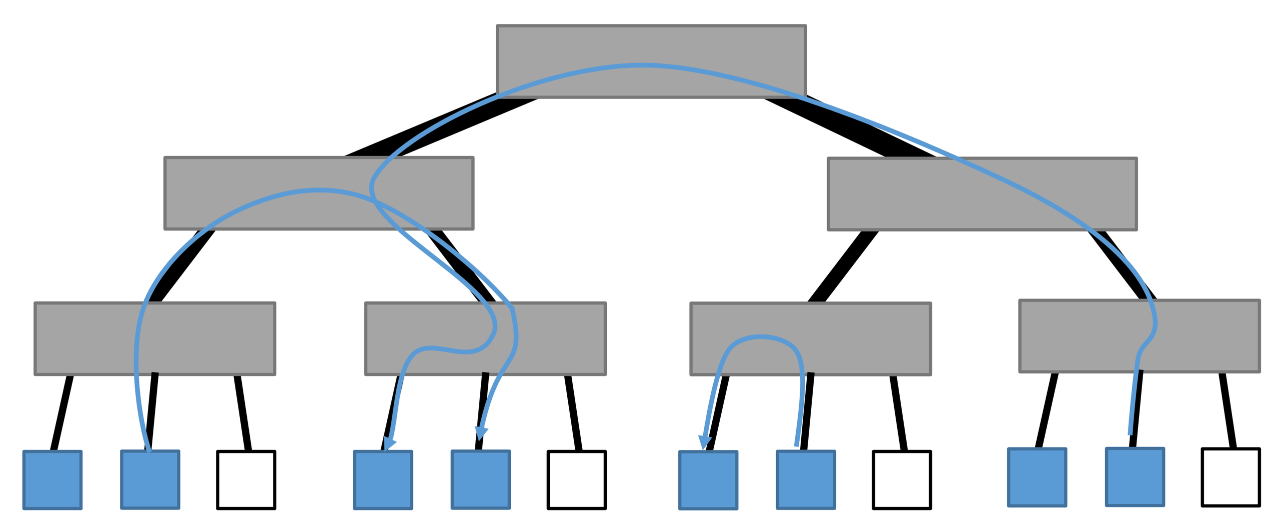

Bisection Bully

- Bully runs on 17 nodes per leaf switch
- Bully communication pairs are made up of two nodes connected to different 2nd level switch
- All bully messages traverse the diameter of the network

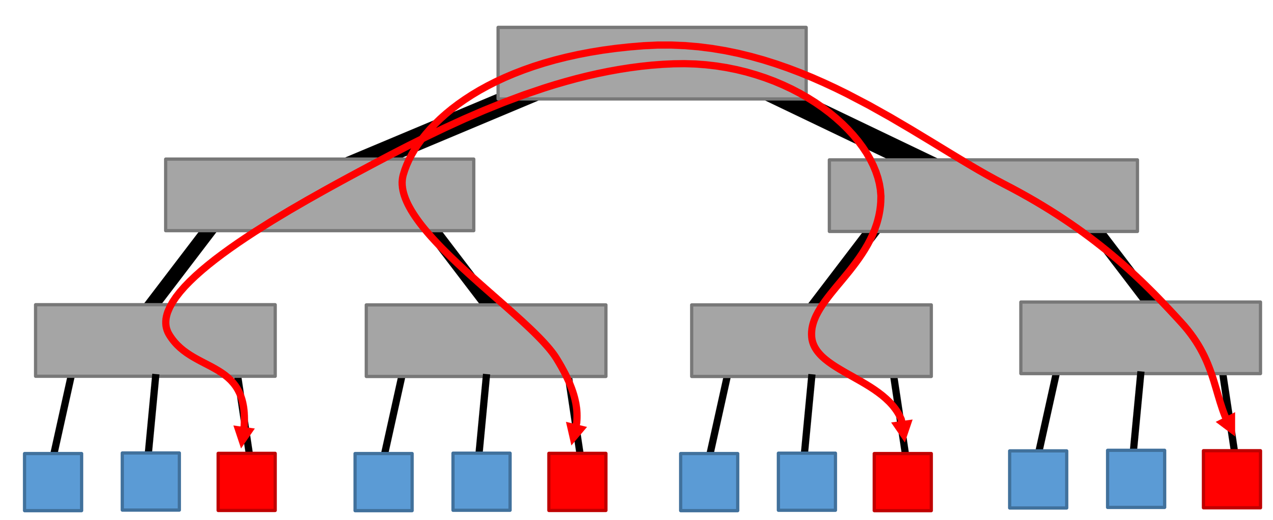

Legend

- Application
- Bully
- Network link
- Application traffic
- Bully traffic
- Switch

Real System Parameters

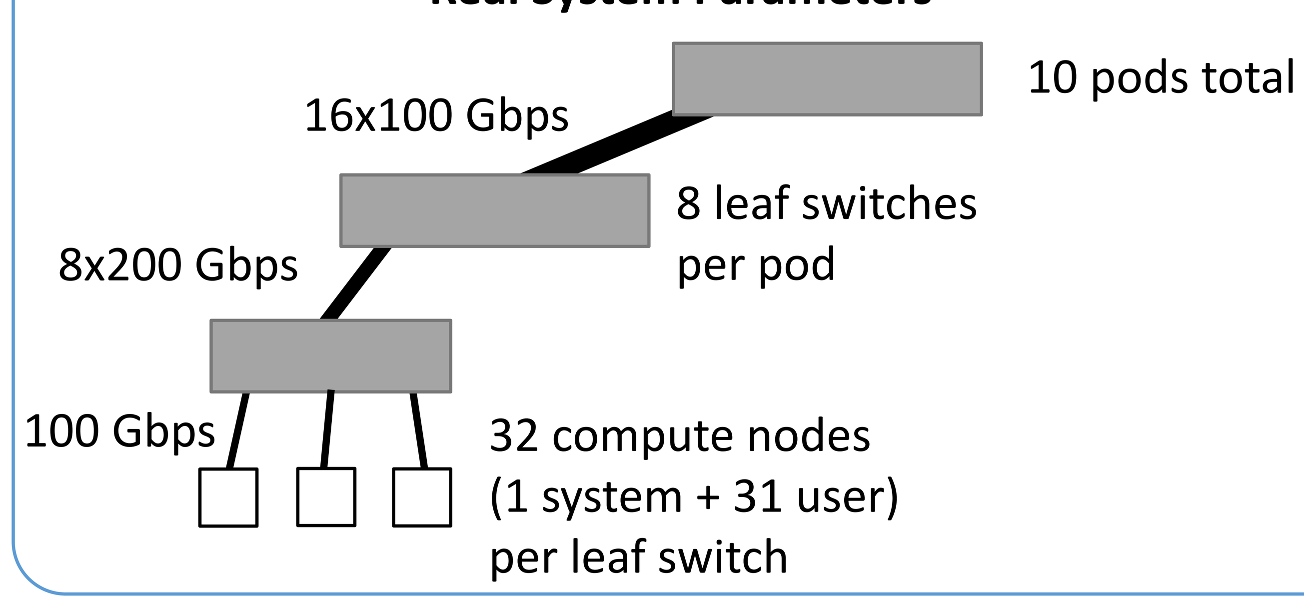

Adjacent Leaf Bully, Light Bully

- Bully communication pairs are made up of two nodes from different leaf switches but the same 2nd level switch
- Two levels of severity:
  - Normal runs on 17 nodes per leaf switch
  - Light runs on 10 nodes per leaf switch
    - Other nodes are idle

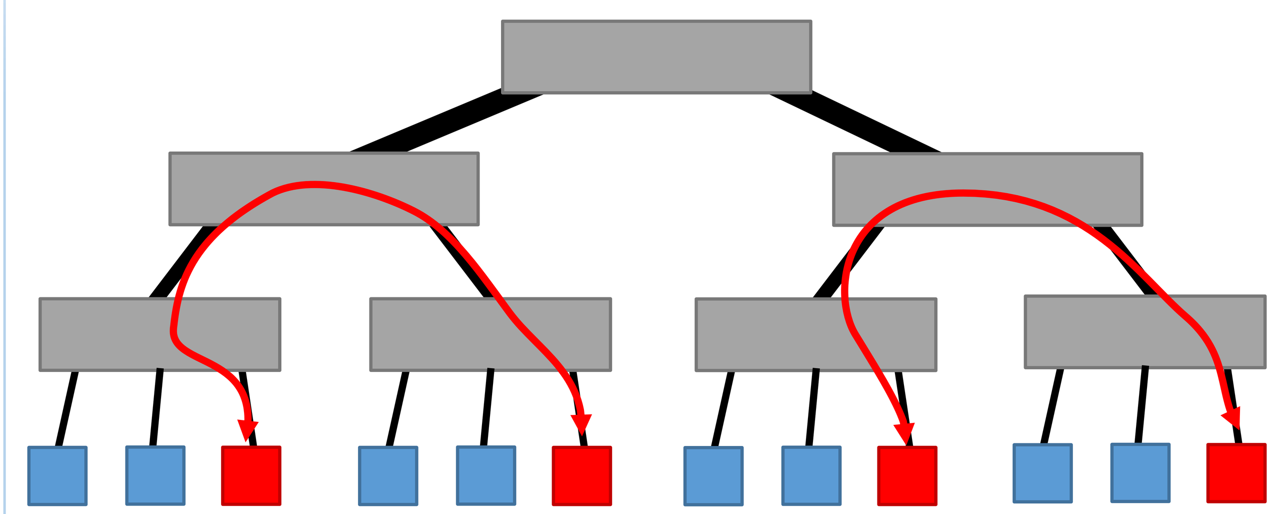

Incast Bully

- Bully runs on 3 nodes per leaf switch
- Bully nodes connected to leaf switch 2+ send messages to a bully node on leaf switch 1

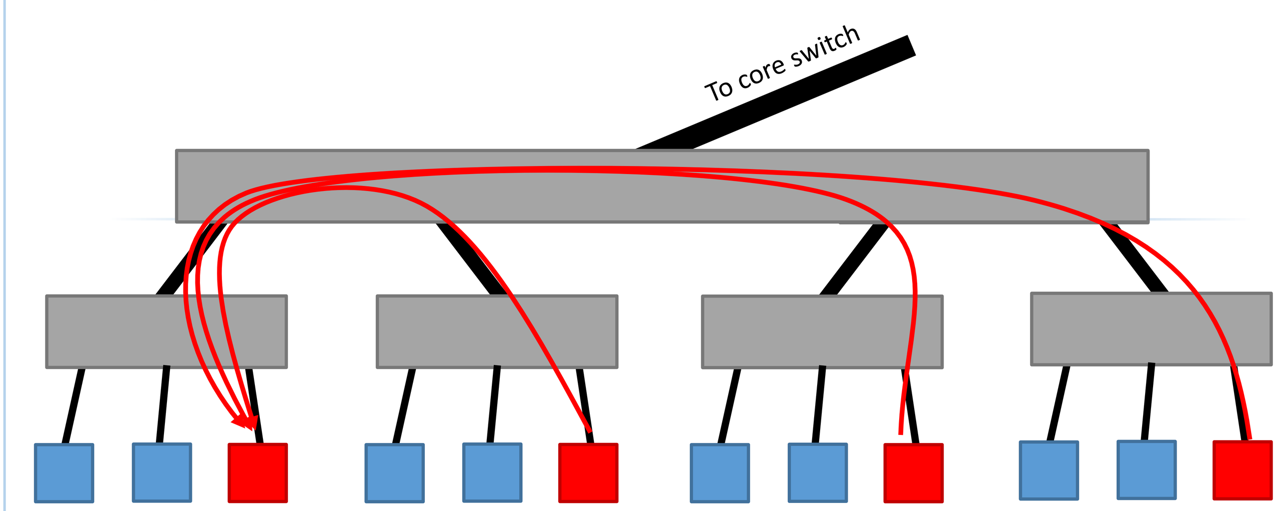

Interleaved

- Different applications interleaved in workload-like pattern
- Variable number of nodes per leaf switch run each application
  - More messages must take longer paths
- No bully application, but applications can cause congestion with each other

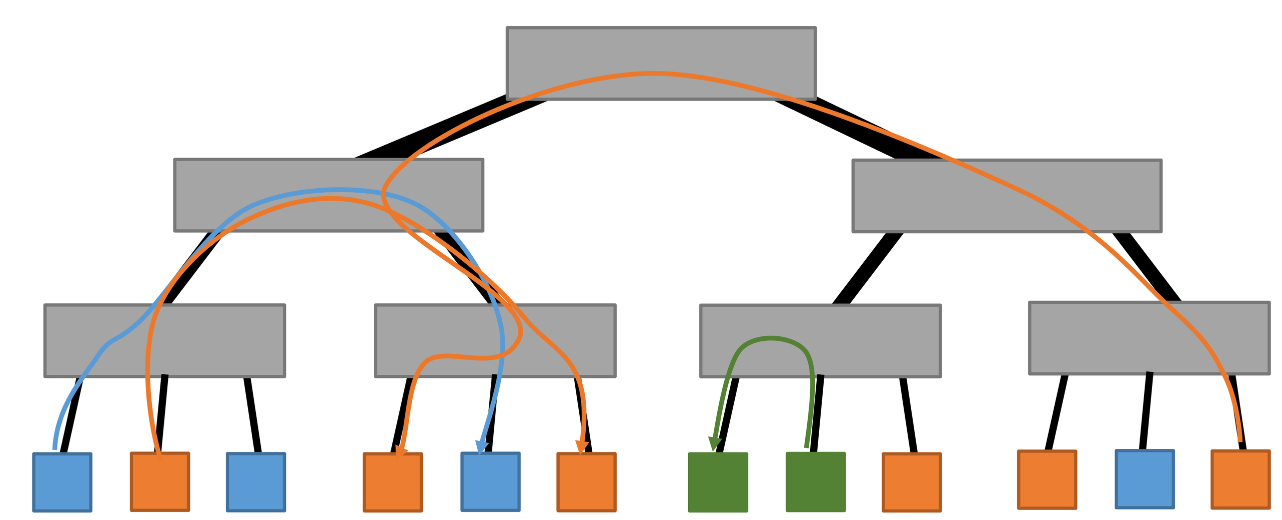

Results

Congestion Type	AMG	UMT	pF3D
Light Bully	-0.70%	**-0.68%**	0.00%
Adjacent Leaf Bully	**-1.74%**	-0.52%	-0.02%
Bisection Bully	**-6.47%**	**-1.94%**	0.00%
Incast Bully	*-0.74%*	0.01%	-0.14%
Interleaved	**-3.17%**	*-0.47%*	*-0.66%*

Mean performance difference compared to control for each application and type of congesting traffic. Negative values indicate that the application performed worse than the control. *Italic orange values* are statistically significant results at the α = 0.05 significance level, while the **bold green values** are also significant at α = 0.01.

Conclusions and Future Work

- Bisection bully caused the largest performance impact
  - Each message it sends traverses more hops than messages sent by adjacent leaf bully or light bully
- AMG exhibited the highest sensitivity to congestion
  - AMG's small messages are latency-bound
  - Other applications send bandwidth bound messages
- Bullies had negligible effect on pF3D
  - Default mapping localizes most communication
    - One major communication pattern node-local
    - Second major communication pattern local to a group of four nodes
- Incast bully caused the least performance impact
  - May be due to static routes or back pressure

- Test mappings for AMG and UMT that localize more communication
- Analyze data switch counter data from experiments to determine the root cause of the performance degradation
- Determine if adaptive routing decreases sensitivity to congestion by eliminating hot spots

References

- https://asc.llnl.gov/CORAL-benchmarks/Summaries/UMT2013_Summary_v1.2.pdf
- https://asc.llnl.gov/coral-2-benchmarks/downloads/AMG_Summary_v1_7.pdf
- http://www.idav.ucdavis.edu/~ki/publications/hipc2014.pdf

IM-POST-XXXXX
tid.llnl.gov/print