Extended Research Abstract

Shafayat Rahman
Department of Computer Science, Florida State University
Tallahassee, Florida 32306
rahman@cs.fsu.edu

ACM Reference Format:
Shafayat Rahman. 2018. Extended Research Abstract. In *Proceedings of ACM Conference (ICPP'18)*. Eugene, OR, USA , 2 pages. https://doi.org/10.1145/n

1 INTRODUCTION

The performance of the interconnect network is massively important in the modern day supercomputer and data centers. As a PhD student in the FSU CS EXPLORER (EXtreme-scale comPuting, modeLing, netwORking & systEms Research) lab under the supervision of Dr. Xin Yuan, my research activity revolves around the analysis, improvement and performance evaluation of a number of topology and routing schemes widely used in the field of high performance computing. Over the years, I worked in a number of projects which I am very briefly describing over the next few paragraphs.

1.1 Load-Balanced slim-Fly Networks

The Slim Fly topology [1] has been proposed recently for future generation supercomputers. In this project [7], we investigated how the traffic is expected to disperse among the network, and discovered that in Slim Fly certain links are more likely to carry traffic than the rest of the links for both minimal and non-minimal routing. As a result, hot-spots are more likely to form in these links. To mitigate the issue, we came up with two different approaches. The first is to modify the topology and increase the bandwidth of the over-used links in such a rate so that the original load-imbalance goes away. The second approach is to modify the routing scheme to distribute the traffic in a more load-balanced fashion. We came up with the 'Weighted' VLB routing scheme, along with two strategies to tune the necessary weights to implement this approach. We validated our results with detailed analysis and simulation, and demonstrated that both the approaches result in a more effective Slim Fly network than its present form.

1.2 Dragonfly Design Space: Link Arrangement and Path Diversity

In Dragonfly [4] topology, the nodes are grouped together into clusters, and the clusters are connected to each other to form a diameter-three network. In this project, we investigate the design space of Dragonfly. Our objective is two-fold. First, we want to figure out how the inter-group connections should be implemented to get optimum performance from the network, specially when the network contains less than maximum number of groups. Second, we analyze the minimal path distribution within the entire design space, and investigate the performance of the existing routing algorithms. This project is currently on-going and we are using a number of modeling, simulation and learning techniques to attain our objectives.

1.3 Traffic-pattern based adaptive routing in Dragonfly Networks

For this project, we investigated the performance of adaptive routing in the Dragonfly network used in Cray Cascade [2]. Here we tried to infer the traffic pattern by observing the packet destinations within a certain window, and considered that information while making the adaptive routing decisions. The first phase of the project only considered the inter-group communications [3]. For the second phase, we are currently working to expand it to inter-group communications as well.

1.4 Performance modeling studies

I worked in a number of projects over the years which analyzed and devised scalable modeling methods to evaluate various topologies, routings and performance metrics. In one project [5], we modeled the UGAL routing over Dragonfly topology to get a better theoretical understanding on how the routing works. In another project [6], we evaluated a number of commonly-used throughput models and identified similar and contradictory trends in their performance. Finally, I worked in another project [6] that studied the performance characteristics of a number of topologies that provide either low diameter or high path diversity.
I started the program in FSU in 2011. I spent a number of years taking various core and elective courses, going through the PhD comprehensive exam, and exploring a number of research areas. I started working with Dr. Yuan from 2014. I acquired the background knowledge necessary to excel in the field of HPC, completed my area exam and worked in a summer internship at Oak Ridge National Lab. I am on course to wrap up the current projects and complete my degree by the end of 2018.

REFERENCES

[1] Maciej Besta and Torsten Hoefler. 2014. Slim fly: a cost effective low-diameter network topology. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press, 348–359.
[2] Greg Faanes, Abdulla Bataineh, Duncan Roweth, Edwin Froese, Bob Alverson, Tim Johnson, Joe Kopnick, Mike Higgins, James Reinhard, et al. 2012. Cray Cascade: A Scalable HPC System Based on a Dragonfly network. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society Press, 103.
[3] P. Faizian, M. S. Rahman, M. A. Mollah, X. Yuan, S. Pakin, and M. Lang. 2016. Traffic Pattern-Based Adaptive Routing for Intra-Group Communication in Dragonfly Networks. In *2016 IEEE 24th Annual Symposium on High-Performance Interconnects (HOTI)*. 19–26. https://doi.org/10.1109/HOTI.2016.017

ICPP'18, August 2018, Eugene, OR, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
https://doi.org/10.1145/n

[4] John Kim, Wiliam J Dally, Steve Scott, and Dennis Abts. 2008. Technology-Driven, Highly-Scalable Dragonfly Topology. In *ACM SIGARCH Computer Architecture News*, Vol. 36. IEEE Computer Society, 77–88.
[5] Md Atiqul Mollah, Peyman Faizian, Md Shafayat Rahman, Xin Yuan, Scott Pakin, and Michael Lang. 2017. Modeling UGAL on the Dragonfly Topology. In *the SC17 workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS17)*.
[6] Md Atiqul Mollah, Peyman Faizian, Md Shafayat Rahman, Xin Yuan, Scott Pakin, and Michael Lang. 2018. A Comparative Study of Topology Design Approaches for HPC Interconnects. In *the 18th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid'18)*.
[7] Md Shafayat Rahman, Md Atiqul Mollah, Peyman Faizian, and Xin Yuan. 2018. Load-balanced Slim Fly Networks. In *the 47th International Conference on Parallel Processing (ICPP)*.