Understanding Congestion on Omni-Path Fabrics

Extended Abstract*

Dan Cassidy
Indiana University South Bend
South Bend, IN
dcassidy@lanl.gov

Lauren Gillespie
Southwestern University
Georgetown, TX
gillespl@southwestern.edu

Christopher Leap
The University of New Mexico
Albuquerque, NM
cleap@unm.edu

ABSTRACT

High-performance computing systems require high-speed interconnects, such as InfiniBand (IB), to efficiently transmit data. Intel's Omni-Path Architecture (OPA) is a new interconnect similar to IB that is implemented on some of Los Alamos National Laboratory's recent clusters. Both interconnects suffer from degraded performance under heavy network traffic loads, resulting in packet discards. However, unlike IB, OPA specifically calls out these drops in the form of a new performance counter called congestion discards. Owing to the relative immaturity of the OPA fabric technology, the correlation between performance degradation and congestion discards has not been fully evaluated to date. This research aims to increase the level of understanding of the effects congestion has on cluster performance by presenting a sufficiently high data injection load to the OPA fabric such that performance degradation is induced and the cause of this performance degradation can be evaluated.

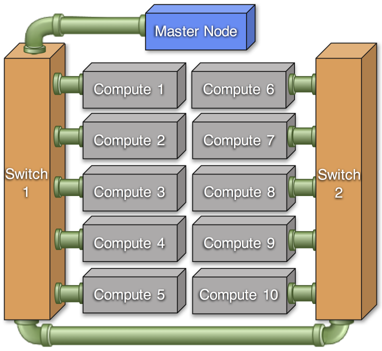

Figure 1: Fabric Topology. The fabric consists of 2 switches with a single interlink between and 5 compute nodes per switch. The compute nodes are paired with a counterpart on the other switch and communicate pairwise across the interlink to force congestion.

1 INTRODUCTION

For the past two decades, InfiniBand (IB), along with the Gemini and Aries interconnects, has served the high-speed networking community. Recently, Intel has entered the playing field with a new high-speed interconnect called Omni-Path Architecture (OPA). Similar to IB, OPA keeps track of how many packets have been dropped during transmission in a counter called "XmitDiscards" (Transmit Discards). This counter accounts for packets that have been dropped due to downed ports, malformed packet sizes, or congestion. Unlike IB, OPA has a counter that specifically calls out these discards due to congestion in a counter called "CongDiscards" (Congestion Discards). Los Alamos National Laboratories (LANL) has implemented OPA in some of its next generation systems, such as Fire, Ice, and Grizzly, which see numerous congestion discards on a daily basis. These discards brought about concern as to the effect congestion drops and consequently, congestion, have on cluster performance.

This research aims to shed light on the effect of congestion discards on OPA cluster performance by introducing heavy traffic into the network and modifying relevant congestion parameters in order to force congestion discards. Our results show that congestion discards themselves are not indicative of network performance. Over the course of this research, we also discovered other forms of traffic that affect network performance without inducing congestion discards. Specifically, traffic using both Intel's Performance Scaled Messaging 2 library (psm2) [3] and Remote Direct Memory Access (RDMA) [4] with queue pairs leads to significantly decreased performance yet little to no congestion discards.

2 CONGESTION DISCARDS

The performance counter congestion discards (CongDiscards) [2] is a counter unique to OPA that counts all dropped packets due solely to congestion. There are three performance parameters that affect when a congestion discard occurs. The first is head-of-queue lifetime (HoqLife) [2], which is the amount of time the head of queue packet has to enter the fabric. If the packet does not enter the fabric in time, it is discarded and resent, producing a congestion discard. Switch lifetime (SwitchLifetime) [2] is the amount of time that queued packets have to enter the fabric before being discarded and resent, also producing a congestion discard. Finally, virtual lanes stalled count (VLStallCount) [2] is the number of consecutive congestion discards allowed in a row before the whole queue is stalled and flushed, producing a congestion discard for each packet flushed.

3 EXPERIMENTAL SETUP

Our system was comprised of eleven nodes with two Intel® Xeon® CPU E5-2620 v4 processors using the ACPI CPU driver and two Intel® Omni-Path Edge Switch 100 Series 48 Port switches. As seen in Figure 1, each switch had five compute nodes which were paired across the interlink for testing purposes. This resulted in a large oversubscription to the interlink between switches, allowing us to force reliable congestion.

---

*LA-UR-17-27077

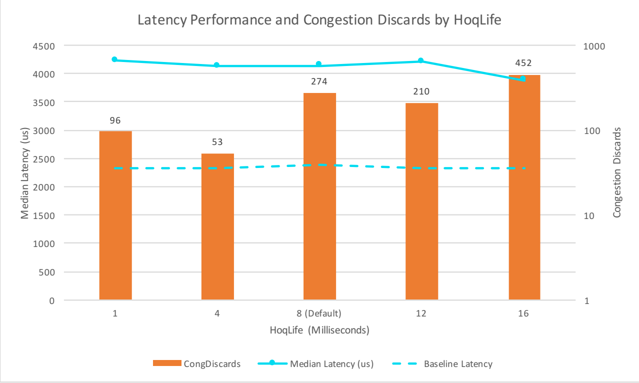

Figure 2: Latency vs. HoqLife. The orange bars represent the number of congestion discards at log 10 scale. The dotted line is the median latency in ms of the baseline latency benchmark. The solid line is the median latency of the latency benchmark and background dd congestion.

In order to compare congestion discards to bandwidth and latency performance, a reliable way to force heavy network traffic was necessary. In order to fill the fabric, 29 GB files were transferred using dd between node pairs. 29 GB size files were chosen since that was the largest file size usable by the system. Also, the Open Fabric Alliance's perftest tools [1], specifically ib_read_bw, were used to fill the fabric. To measure bandwidth and latency, Ohio State University's (OSU) microbenchmarks [5] were used. For bandwidth, the OSU bidirectional bandwidth benchmark (osu_bibw) was used, and for latency, the OSU all-to-all latency benchmark (osu_alltoall) was used. In addition, the congestion parameters HoqLife and SwitchLifetime were modified separately between tests. Each benchmark plus background congestion was run under the given parameters 50 times per test for statistical significance.

4 RESULTS AND CONCLUSIONS

Our results show that congestion discards alone are not indicative of network performance degradation when using dd to fill the fabric. As seen in Figure 2, when modifying HoqLife, congestion discards appreciably increased, yet latency still did not significantly increase, as was also true when modifying SwitchLifetime. The same trend was true with congestion discards and bandwidth where the presence of congestion discards led to no significant bandwidth decrease from the baseline when modifying both HoqLife and SwitchLifetime. However, when the SwitchLifetime and HoqLife parameters were disabled, there were appreciable congestion discards and network performance degraded so significantly that the benchmarks could not execute properly.

However, benchmarks run with perftest as the background congestion had much different results from dd. As seen in Figure 3, when running perftest and the OSU bandwidth benchmark, no congestion discards occurred, yet performance was significantly lower than the baseline and was also lower than the performance of the OSU bandwidth benchmark and dd tests. This clearly shows that there is a negative effect on network performance when running perftest, but this negative effect is not related to congestion discards.

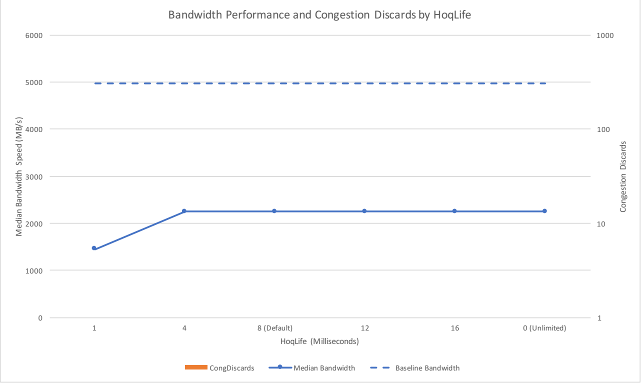

Figure 3: Bandwidth vs. HoqLife. There were no congestion discards. The dotted line represents the baseline bandwidth benchmark. The solid line is the median bandwidth of the bandwidth benchmark and background perftest congestion.

One possible reason for this significant performance decrease is that perftest uses the RDMA transport protocol, which offloads packet processing to the Network Interface Card (NIC). Therefore, the bottleneck is on the NIC, not the fabric, which leads to lower network performance but no congestion discards.

In conclusion, congestion discards alone are not a good measure of network performance and differing congestion parameter settings do not lead to appreciable performance changes. However, traffic that uses RDMA leads to no congestion discards but significant performance decreases due to unforeseen bottlenecks on the NIC.

5 FUTURE WORK

Areas of future exploration with OPA and congestion include experimenting with a larger cluster, which would allow for the congestive effects of different workload patterns to be explored. Understanding the effects of the VLStallCount parameter on network performance would be another area of interest. Finally, some results with psm2 showed a noticable dip at the 16KB message size, which is where OPA switches to the remote direct memory access (RDMA) protocol. Further investigating the interaction into this behavior could also be studied.

ACKNOWLEDGMENTS

We would like to thank our mentors Susan Coulter, Jesse Martinez, and Howard Pritchard. We would also like to thank Brian McGinnis from Intel for assisting us with tuning our systems. Finally, we would like to thank our instructor Joan Lucas and our teaching assistant Hunter Easterday along with our program coordinators Amanda Bonnie and Alfred Torrez.

REFERENCES

[1] OpenFabrics Alliance. 2017. perftest downloads. (2017). https://www.openfabrics.org/downloads/perftest/

[2] Intel Corporation. 2017. Intel® Omni-Path Fabric Suite Fabric Manager User Guide. (April 2017). https://www.intel.com/content/dam/support/us/en/documents/network-and-i-o/fabric-products/Intel_OP_FabricSuite_Fabric_Manager_UG_H76468_v6_0.pdf

[3] Intel Corporation. 2017. Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide. (April 2017). https://www.intel.com/content/dam/support/us/en/documents/network-and-i-o/fabric-products/Intel_PSM2_PG_H76473_v6_0.pdf
[4] Mellanox Technologies. 2015. RDMA Aware Networks Programming User Manu. (2015). http://www.mellanox.com/related-docs/prod_software/RDMA_Aware_Programming_user_manual.pdf
[5] Ohio State University. 2016. MVAPICH :: Benchmarks. (2016). http://mvapich.cse.ohio-state.edu/benchmarks/