GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

Sudheer Chunduri*

sudheer@anl.gov
Argonne National Lab

Brian Austin

baustin@lbl.gov
Lawrence Berkeley National Lab

Kalyan Kumaran

kumaran@anl.gov
Argonne National Lab

Steven Warren

swarren@cray.com
Cray Inc.

Taylor Groves*

tgroves@lbl.gov
Lawrence Berkeley National Lab

Jacob Balma

jbalma@cray.com
Cray Inc.

Glenn Lockwood

glock@lbl.gov
Lawrence Berkeley National Lab

Nathan Wichmann

wichmann@cray.com
Cray Inc.

Peter Mendygral*

pjm@cray.com
Cray Inc.

Krishna Kandalla

kkandalla@cray.com
Cray Inc.

Scott Parker

sparker@anl.gov
Argonne National Lab

Nicholas Wright

njwright@lbl.gov
Lawrence Berkeley National Lab

ABSTRACT

Network congestion is one of the biggest problems facing HPC systems today, affecting system throughput, performance, user experience, and reproducibility. Congestion manifests as run-to-run variability due to contention for shared resources (e.g., filesystems) or routes between compute endpoints. Despite its significance, current network benchmarks fail to proxy the real-world network utilization seen on congested systems. We propose a new open-source benchmark suite called the Global Performance and Congestion Network Tests (GPCNeT) to advance the state of the practice in this area. The guiding principles used in designing GPCNeT are described and the methodology employed to maximize its utility is presented. The capabilities of GPCNeT are evaluated by analyzing results from several world's largest HPC systems, including an evaluation of congestion management on a next-generation network. The results show that systems of all technologies and scales are susceptible to congestion and this work motivates the need for congestion control in next-generation networks.

*primary authors contributed equally to the paper

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

SC '19, November 17–22, 2019, Denver, CO, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6229-0/19/11...\$15.00

<https://doi.org/10.1145/3295500.3356215>

KEYWORDS

network benchmarking, congestion, performance variability

ACM Reference Format:

Sudheer Chunduri, Taylor Groves, Peter Mendygral, Brian Austin, Jacob Balma, Krishna Kandalla, Kalyan Kumaran, Glenn Lockwood, Scott Parker, Steven Warren, Nathan Wichmann, and Nicholas Wright. 2019. GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks. In *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '19)*, November 17–22, 2019, Denver, CO, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3295500.3356215>

1 INTRODUCTION

In High Performance Computing (HPC) systems, performant networks are essential to join resources efficiently together to enable scientific discovery at unprecedented scales. While simple metrics such as latency, injection bandwidth, and bisection bandwidth are useful indicators of a network's peak performance, typically they do not correspond to realized performance on production machines, particularly when diverse workloads contend for the same network resources. They often overemphasize best-case configurations on quiet systems and fail to expose behavioral differences between network topologies and architectures. In this work we examine the impact of contention for network resources and examine how such contention affects performance.

Network contention often manifests in application performance variation. Previous studies have shown that the impact of network jitter is substantial and increases super-linearly with process count for many systems [3–5, 12, 26, 28].

The purpose of this work is to introduce a benchmark suite that (a) induces network congestion and (b) measures its impact on latency and bandwidth. There has been no previous work on congestion inducing benchmarks because there was little that existing HPC networks did to address the issue. However, modern network architectures such as Cray Slingshot [6], Infiniband [16],

and Gen-Z [29] have made handling of congestion a priority with the acknowledgement that HPC centers want to provide a level of performance isolation between a breadth of heterogeneous workloads. As new networks incorporate features such as traffic classes and congestion control, HPC facilities require new benchmarks designed to evaluate their effectiveness. In this paper we seek to advance the state of the practice in this regard.

This paper makes several contributions towards these goals:

- We present GPCNeT – a generic, topology agnostic benchmark suite that captures the complex workloads anticipated on multitenant HPC networks. It uses multiple configurable congestion patterns concurrently to emulate multiple separate jobs and uses concurrent canary jobs that emulate boundary exchange algorithms to measure the impact of congestion.
- We show that GPCNeT works on production HPC machines and describe a systematic method to measure congestion intensity. Our results show that latency-bound communication can be very sensitive to congestion, bandwidth is generally less sensitive to congestion, and that congestion has a scale-dependent component. We also show that congestion manifests in different ways on the same machine depending on how an MPI library implements different interfaces.
- We demonstrate that by instrumenting the network counters of an HPC machine, it is possible to tune the benchmark to better represent a workload by emphasizing or de-emphasizing certain congestors or canaries.
- GPCNeT also introduces the *Congestion Impact* metric to study the impact of congestion across systems that are based on different topologies and architectures. Using this metric we perform an analysis of the impact of network congestion on several of the world’s largest HPC systems, including an evaluation of congestion management on a next-generation network showing that systems of all technologies and scales are susceptible to congestion effects. Networks with newer technologies (EDR IB, Slingshot), are less susceptible than older ones (Cray Aries).

In § 3 we present the GPCNeT benchmark which creates systematic congestion and measures the impact on latency and bandwidth for some common application-derived communication patterns. The benchmark was created to fulfill a set of design requirements (detailed in that section). We demonstrate the fulfillment of the design requirements in § 4, showing how GPCNeT provides meaningful results across a spectrum of architectures. We then use GPCNeT to compare the impact of congestion across architecture and topology for some of the world’s top HPC systems (§ 5).

2 BACKGROUND AND RELATED WORK

The performance efficiency of the interconnection network plays a significant role in determining the scaling and performance of communication-intensive applications. The contention for shared resources can have a significant impact on application performance. Hence, detailed performance evaluations of networks are essential to quantify the congestion impacts on production application performance.

Traditionally, MPI micro benchmarks such as Sandia [2], OSU [20], and Intel [17] microbenchmarks are used to measure the baseline latencies and bandwidths for data transfers across the network. Some benchmarks such as RMA-MT benchmarks [8, 31] are used to measure the specific aspects of MPI such as multi-threaded RMA. Hoefler et al. developed a microbenchmark for measuring both latency and overlap of computation and communication for non-blocking communication operations [13]. However, these benchmarks are intended to measure ideal baseline performance.

The impact of tail latency (extreme latency in a long-tailed distribution; e.g., greater than 99th percentile) on performance is not unique to HPC systems; this has been studied in the context of data center networks as well [22]. Yunqi et al. [32] developed a methodology using quantile regression to analyze, and attribute the sources of variance in tail latency to various hardware features of interest. The tail latency effects on production datacenter-scale file systems are also studied [24]. The congestion impact as measured using the LDMS telemetry data on two real production systems is analyzed [18].

There is a significant body of work that has measured the impact of congestion in HPC systems, but these do not include benchmarks that induce congestion. Hoefler et al. have developed benchmarks [14, 15] that consider the system noise aspects for better reflecting the application performance. Ates et al. developed a performance anomaly suite [1] that includes a simple network contention generator using pairs of communicating nodes. Performance modeling for the proposal, initial testing, and beginning deployment phase of the Blue Waters supercomputer system utilized LogGP model with parameters derived from Netgauge [14]. Grant et al. [11] developed a tool for measuring the impact of network contention due to interference from other jobs through a continually running benchmark application and the use of network performance counters. Prior works [30] showed the importance of considering the latency variation rather than just considering the latency mean for characterizing application performance.

These studies emphasize the importance of considering system noise for accurately quantifying the network efficiency for production application performance. However, there exists no standard benchmark suite for measuring the network performance in production scenarios. To be broadly applicable, these benchmarks are desired to be topology-agnostic, MPI implementation-agnostic, and network provider-agnostic. The benchmarks we propose in this study meet these criteria and can characterize network congestion impacts on various system architectures.

3 BENCHMARK DESIGN

We present the Global Performance and Congestion Network Tests. GPCNeT¹ differs from existing benchmarks by creating systematic congestion. Two types of communication patterns, sensitive (or “canary”) and potentially congesting (hereafter “congestors”), execute in a coordinated manner. GPCNeT compares the performance of sensitive communication in isolation and while congestors are running. Sensitive traffic is generated by latency and bandwidth-sensitive communication kernels motivated by HPC applications (e.g., multi-dimensional nearest-neighbor exchanges). Congestor

¹GPCNeT is publicly available at <https://github.com/netbench/GPCNET>

traffic is generated by kernels that stress the network (e.g., all-to-all and incasts). GPCNeT places a load on the network by running several congestor traffic patterns at the same time. These operations are all managed by a single executable in a single job launch. We identified additional requirements that affect the design of GPCNeT. The design requirements are that it

- (1) is portable across a variety of topologies and architectures,
- (2) is not easily manipulated to favor specific networks or implementations,
- (3) can generate complex communication patterns with simple tuning,
- (4) reports statistics representing the limiting communication performance of a latency or bandwidth sensitive HPC application,
- (5) is able to run on any number of nodes.

GPCNeT uses MPI to generate a variety of network traffic patterns and can run on any system that supports the MPI-3.0 specification (first design requirement). GPCNeT can be deployed at full system scale on a given system. This requirement affects how sensitive and congestor communication patterns are generated, which is described in more detail in later sections. Process placement is an important consideration because real-world systems that support a diverse workload tend to provide non-ideal placement of processes relative to the network topology. GPCNeT includes features to approximate non-ideal application placement so that sensitive traffic patterns cannot be isolated from congestors by relying on MPI rank reordering (second design requirement).

A key tuning parameter for GPCNeT is the number of MPI ranks executing on each node. By simply increasing the number of processes for a fixed number of nodes, GPCNeT will increase the load on the network and generate more complex communication patterns (additional detail in § 3.1 and § 3.2), which addresses the third design requirement. There are many configurations for integrating a node into an HPC network, however. A node can have multiple NIC ports for communication, for example. Throughout this paper we refer to the number of Processes per network Port (PPort) to describe the relationship between how GPCNeT is run and the network resources on a node.

Another important set of features in GPCNeT are the metrics it measures and reports (fourth design requirement). Many MPI benchmarks measure the minimum, maximum, and average performance of an operation measured across ranks. However, these values do not reflect the communication performance perceived by real applications in production environments. Tail latency (e.g., 99% or 99.9%) is commonly used to describe performance-limiting communication for datacenter distributed applications [7]. For the purposes of this work, we define tail latency as a measure of latency significantly greater than the center of the distribution. Any application with synchronization times approaching the tail latency on a network (under load) can experience performance impacts. GPCNeT measures and reports 99% tail latency and the 99% tail performance of all included communication kernels.

3.1 Sensitive (Canary) Traffic

It is essential for the contention measurement tests to span a large fraction of the system; congestion can be localized in space and

transient in time, and so might not be detected by monitoring a small number of links or node-pairs. We therefore selected network traffic patterns that can run on any number of nodes, and are sensitive to both latency and bandwidth effects. The communication pattern for these "canary" jobs alternates between all-reduce and a random ring. The all-reduce pattern is ubiquitous throughout HPC and measures the latency sensitivity of small-message collective operations, and the random ring targets bandwidth-sensitive point-to-point operations.

The random ring pattern was motivated by the prevalence of nearest-neighbor 3D communication in physics-based applications, but we have limited the dimensionality to a 1D ring so that the benchmark can run on any number of nodes. Multiple processes can be run on each node to better represent multi-dimensional communication patterns. Although randomizing the order of nodes in the ring eliminates locality (and related optimizations) from the communication pattern, it ensures that the benchmark is agnostic to network topology. A secondary benefit of randomization is the added flavor of irregular communication patterns generated by task-based programming models and graph processing.

3.1.1 Random Ring Infrastructure. A random ring infrastructure emulates an application or set of applications performing nearest-neighbor exchanges. The random ring design is based on a similar implementation in HPCC [23], where a one-dimensional periodic ring is generated by randomizing a list of MPI ranks and each rank communicates to its left and right neighbors in the ring. The random ring design satisfies the fifth design requirement that GPCNeT be able to run at any node count. Iterating over many different random rings also prevents optimizations based on MPI rank reordering.

Communication should only be off-node, preventing the results from being contaminated by on-node communication. This is achieved by creating PPort sub-communicators with each process on a node in a unique sub-communicator. Ranks in the same sub-communicator are randomized to form a random ring. With this design, a given node will communicate with up to $2 \times PPort$ other nodes, accounting for left and right communication. Communication patterns resembling a multidimensional decomposition or unstructured mesh are naturally produced by this behavior. The connectivity between nodes is less than all-to-all communication but more complex than a 2-dimensional decomposition. Running GPCNeT with increasing PPort is therefore a simple way to increase the number of network pathways traversed by sensitive traffic.

Listing 1 summarizes the random ring infrastructure and performance measurements in it. A timeout mechanism is included to prevent tests from running longer than 10 seconds, which is long enough to measure a statistically significant number of iterations of any of the communication kernels under most conditions. It is possible that a test runs longer than 10 seconds, however, when very heavy load is on the network. The generation of a random ring is based on the ring loop index. This means the same set of random rings is produced in each measurement iteration.

```

/* nM      = total number of measurements
   nR      = total number of random rings
             to generate in a sub-communicator
   nI      = number of iterations of a
             communication kernel
   myrank  = local MPI rank in a sub-communicator
   warmup  = number of warmup iterations

```

```

*/
timeout1 = timer();
for (m = 0; m < nM; m++) {

    /* test if running too long */
    timeout = timer() - timeout1;
    MPI_Allreduce(MPI_IN_PLACE, &timeout, 1, MPI_DOUBLE,
        MPI_MIN, global_comm);
    if (timeout > 10.0) continue;

    /* loop over random rings */
    for (n = 0; n < nR; n++) {
        randomize_list(&ranklist, n);
        neighbor_left = ranklist[myrank-1];
        neighbor_right = ranklist[myrank+1];
        for (i = -warmup; i < nI; i++) {
            if (i == 0) timer_bulk_start = timer();
            if (i >= 0) timer_fine_start = timer();

            communication_kernel(neighbor_left,
                neighbor_right);

            if (i >= 0)
                timer_fine[m*nR*nI + n*nI + i] = timer();
        }
        timer_bulk[m*nR + n] = timer() - timer_bulk_start;
    }
}

```

Listing 1: Pseudocode of the random ring infrastructure.

Two communication kernels execute in the random ring infrastructure. One is a latency measurement and the other measures bandwidth with a synchronization. The infrastructure is able to support other kernels as well. GPCNeT can be easily modified to use MPI-RMA based latency and bandwidth kernels (included in the package), for example, as additional sensitive traffic.

Random Ring Point-to-point Latency (P2P Lat). P2P Lat measures the latency of sending and receiving 8 byte messages to and from the left and right neighbors. This kernel is intended to be more latency-sensitive than typical HPC applications to provide an upper limit on the performance impact from congestion. The loop limits in Listing 1 for this test are $nM = 10000$, $nR = 30$, $nI = 200$, and $warmup = 200$, which are sufficient for statistically rigorous results. Each canary process in the kernel performs two `MPI_Irecv` and `MPI_Isend` with respect to the two random neighbors per iteration followed by a `MPI_Waitall`. The walltime of each kernel iteration is measured and divided by two to account for round-trip time.

Random Ring Point-to-point Bandwidth with Synchronization (P2P BW+Sync). P2P BW+Sync is motivated by applications performing boundary value exchanges followed by a small reduction for a time step size or convergence test. Eight 128 Kbyte messages are sent to and received from the left and right neighbors (16 messages in total), which is typically enough messages to achieve peak bandwidth. Afterwards, a barrier on the sub-communicator synchronizes all ranks. This test is sensitive to bandwidth contention and latency from the synchronization. The loop limits in Listing 1 for this test are $nM = 10000$, $nR = 30$, $nI = 8$, and $warmup = 1$. The walltime of each iteration is measured and converted into a bandwidth value, accounting for 16 messages.

3.1.2 Allreduce. A separate test for small message `MPI_Allreduce` is included outside of the random ring infrastructure. This test is used to measure the effects from congestion on latency-sensitive collective communication. The same sub-communicators used in

the random ring infrastructure are used in this test. Therefore, there are `P2P MPI_Allreduce` operations occurring simultaneously, each with only off-node communication. The loop structure for this kernel is the same as Listing 1, excluding the loop over random rings, with $nM = 100000$, $nI = 200$, and $warmup = 1$.

3.1.3 Performance Metrics. GPCNeT uses the two timers, shown in Listing 1, to compute performance statistics on the communication kernel – a fine-grained timer that generates histograms for each configuration of random ring and a bulk timer that provides overview statistics.

3.2 Congestors

Congestion occurs on a network when there are insufficient resources available to accommodate the demanded load. Communication patterns that generate an imbalance between ingress and egress rate can cause congestion. Broadly, we can organize congestion into two categories: endpoint and intermediate. An example of endpoint is an incast operation where data from $N-1$ senders is directed to a single receiver, which is unable to accept data at the rate it arrives. It is also possible to induce side-effects of congestion, such as queuing delays, by placing heavy loads on a network. Many-to-many communication patterns (e.g., all-to-all with large messages) will stress a network and is an example of intermediate congestion. These patterns develop on networks supporting real workloads either explicitly by demanding applications or spontaneously through the cumulative effects of many applications and supporting subsystems, such as filesystems. To ensure portability (first design requirement), GPCNeT uses MPI operations to generate these communication patterns.

MPI operations define the logical flow of data but do not specify the underlying methods by which data is moved. For example, a logical incast can be written in MPI by initiating $N-1$ non-blocking receive operations at the root and a send operation from each sender. This pattern may or may not generate a true incast pattern depending on the underlying communication protocols (Get/Put) used in the MPI implementation. MPI-RMA exposes more explicit operations, such as `MPI_Get()` and `MPI_Put()`, but an MPI library can implement MPI-RMA using two-sided operations. To overcome these implementation complexities and ensure efficacy across architectures (first design requirement), GPCNeT runs four congestors simultaneously to create endpoint and intermediate congestion. A specific congestor may not produce heavy load or congestion on its own. But, in the absence of effective congestion management in the network, some combination of the congestors should produce heavy load. GPCNeT also uses a common message size for congestors, which makes it difficult to tune an MPI implementation to favor one type of a congestor without negatively affecting the rest of the kernels in the benchmark.

3.2.1 Congestor Infrastructure. Congestor communication kernels are executed within a common infrastructure, similar to the random ring infrastructure in § 3.1. Ranks on a congestor node are assigned one of the congestor kernels for the duration of execution. Listing 2 summarizes the congestor infrastructure. A difference from the random ring infrastructure is the use of `MPI_Iallreduce` to manage a timeout. This improves the chances that load from

a congestor is constant despite the presence of other collective communication. Similar to the sub-communicators used for sensitive traffic, a congestor kernel is executed with PPort different sub-communicators simultaneously to avoid any on-node traffic. Tests of GPCNeT with different PPort not only tune the behavior of sensitive traffic (see § 3.1.1), they also tune the intensity of the congestors. The performance of each congestor is measured and is included in GPCNeT output for verbose compilations.

```

/* nM      = total number of measurements
   nI      = number of iterations of the kernel
   myrank  = local MPI rank in a sub-communicator
*/
timeout1 = timer();
for (m = 0; m < nM; m++) {

    /* test if running too long */
    MPI_Wait(&req, MPI_STATUS_IGNORE);
    if (timeout > 10.0) continue;
    timeout = timer() - timeout1;
    MPI_Allreduce(MPI_IN_PLACE, &timeout, 1, MPI_DOUBLE,
        MPI_MIN, scom, &req);

    for (i = -1; i < nI; i++) {
        if (i == 0)
            timer_bulk_start = timer();
        if (i >= 0)
            timer_fine_start = timer();

        congestor_kernel(myrank);

        if (i >= 0)
            timer_fine[m*nR*nI + n*nI + i] = timer();
    }
    timer_bulk[m*nR + n] = timer() - timer_bulk_start;
}

```

Listing 2: Pseudocode of congestors infrastructure.

3.2.2 All-to-all (A2A). An all-to-all communication pattern is generated with a pair-wise exchange algorithm. This implementation provides more uniformity across MPI libraries than `MPI_Alltoall`, which might be highly tuned in some cases. This congestor uses 4 Kbyte messages, which are sufficient to generate significant traffic while limiting the size of the send and receive buffers. The loop limits in Listing 2 for this kernel are $nM = 256$ and $nI = 512$.

3.2.3 Point-to-point Incast (P2P Incast). P2P Incast is one congestor kernel that can generate an incast for some MPI implementations. An incast is generated for MPI libraries that use a RDMA Put-based protocol for send and receive operations to transfer a 4 Kbyte message, the message size used in this kernel. This message size is sufficient to generate an intense incast while limiting the time to complete the operation (roughly 4 Kbyte divided by the ejection rate at the root). The loop limits in Listing 2 for this kernel are $nM = 256$ and $nI = 512$.

3.2.4 One-sided RMA Incast (RMA Incast). Logical incasts using `MPI_Put` will generate a true incast for MPI libraries that implement MPI-RMA with one-sided RDMA operations for the tested network. If MPI-RMA is implemented with two-sided operations instead, this congestor kernel is approximately identical to the point-to-point incast kernel and also uses 4 Kbyte messages. The loop limits in Listing 2 for this kernel are $nM = 256$ and $nI = 512$.

3.2.5 One-sided RMA Broadcast (RMA Bcast). An incast can be implicitly generated by a naive one-sided broadcast. In this case,

the actual data is flowing from the root to the receivers, but request packets for that data from all of the receivers produce an incast. This pattern occurs for MPI libraries that implement MPI-RMA with one-sided RDMA operations for the tested network. MPI libraries that implement MPI-RMA with two-sided operations may also generate an incast for this kernel if an RDMA Get-based protocol is used to transfer a 4 Kbyte messages. The loop limits in Listing 2 for this kernel are $nM = 256$ and $nI = 512$.

3.3 Execution Sequence

The nodes on which GPCNeT is executing are divided into one group that executes sensitive communication kernels (S) and another that executes the congestors' communication kernels (C). The execution sequence is summarized as follows.

- (1) The full set of nodes is shuffled. 20% of the set are assigned to group S and the remaining 80%² to group C, which is further sub-divided evenly across the four congestor types (C0-C3).
- (2) The ranks on nodes from group S are arranged into sub-communicators such that the Nth rank on each node is in the same sub-communicator. The same is done for ranks in node groups C0 - C3. This approach prevents any on-node communication from contributing to measurements.
- (3) Ranks in group S loop through each of the tests, gathering baseline performance until reaching a barrier which initiates congestors C0-C3.
- (4) Ranks in group C begin executing their communication as a warm-up while ranks in group S wait. Once ranks in group C have executed their warm-up, they release ranks in group S.
- (5) Ranks in group S execute one of the sensitive communication tests and inform ranks in group C once the measurements are complete.
- (6) Steps 4-6 are repeated for each sensitive communication test.

The node list randomization in Step 1 uses the local time on rank 0 as the seed by default. The reason for this is to prevent application placement optimizations or node-level tuning (e.g., predictive QoS or specifying traffic classes) that would not typically be available to users or easily tuned for each job on production systems. This randomization does introduce potential run-to-run variability in the results from GPCNeT, however. Users can optionally modify GPCNeT to use a fixed seed if desired.

4 VALIDATION OF DESIGN REQUIREMENTS

In this section we perform a wide range of experiments on two test systems to evaluate whether our benchmarks fulfill the requirements set out in § 3. We explore the impact (1) of process and node count, (2) of MPI Library and (3) the isolated congestors have on GPCNeT results. We also examine the effect of network congestion on a real HPC application, MILC, and compare the effects of GPCNeT's synthetic congestion to production environments.

²Ratio of canary:congestors is adjustable.

Figure 1: Histogram of P2P Lat measurements at 64 (top) and 696 (bottom) nodes on a XC40 system. The upper and lower panel on each is without and with congestors, respectively. Runs use 1, 8, or 32 processes per node.

4.1 Test Systems

We use different HPC network architectures and MPI libraries (detailed in Table 1) to evaluate the effectiveness of GPCNeT at generating and measuring congestion impact.

Experiments on these systems cover a range of job sizes and choices of PPort. It is important to select the job size to fill a region of a network’s topology or to fill (or nearly fill) the entire network. For an idle system, adaptive routing features on some networks may provide sufficient extra bandwidth to a smaller job to avoid congestion impacts. For the Cray XC40 system, Crystal, the relevant scales are within a chassis (64 nodes), within a group (336 nodes), and the full system (696 nodes). Results for 64 and 696 nodes are included in this analysis. The only relevant scale on the CS500 system, Osprey, is all 128 nodes. These experiments used default settings for the MPI libraries and default site-specific settings for the network.

4.2 Effect of Varying PPort and Node Count

Figure 1 shows the histograms of P2P Lat, described in § 3.1.1, at different node counts with 1, 8, and 32 PPort. The peak of the isolated runs is between $1.7 \mu\text{s}$ and $2 \mu\text{s}$, and the tail increases in response to both increasing PPort and job size. Nearly all samples are well below $100 \mu\text{s}$. The distributions are multimodal, likely related to extra hops for minimal and non-minimal routes. When congestors load the network, the distributions spread and flatten out. This data shows that increasing node count and PPort result in larger impacts from congestion, with latency values exceeding 10 ms for 32 PPort on 696 nodes.

Histograms for P2P BW+Sync with synchronization are in Figure 2. Note that bandwidth is reported per MPI rank. Runs at 64 nodes show an expected drop in bandwidth per rank as PPort increases, but there is only a small effect from congestion. Runs at 696 nodes

Figure 2: Histogram of P2P BW+Sync measurements at 64 (top) and 696 (bottom) nodes on a XC40 system. The upper and lower panel on each is without and with congestors, respectively. Processes per node are 1, 8, and 32.

show more effect from congestion, especially at 32 PPort. Bandwidth drops from $\sim 100 \text{ MiB/s/rank}$ to only a few MiB/s/rank at 32 PPort and 696 nodes when congestors load the network.

Figure 3 shows the histograms for MPI_Allreduce. Cray MPICH uses a recursive doubling algorithm for MPI_Allreduce by default under these conditions. The node counts (64 and 696) result in a non-power of two number of ranks in all of the canary kernels, which adds two extra steps to the algorithm. These extra steps may account for additional peaks in the isolated distributions. Load from congestors results in significant impacts on performance, in particular at 696 nodes. From these results, it is clear that GPCNeT demonstrates the effects of congestion on tail latencies, and these effects increase with scale. The algorithmic complexity underlying MPI_Allreduce adds additional sensitivity, but the effects of congestion on bandwidth manifest only at larger concurrencies.

4.3 Effect of Each Congestor

GPCNeT deploys four congestors across 80% of the allocated processes by default. While 80% of the nodes acting as congestors represents worst-case behavior, GPCNeT is easily configurable to customize this ratio or use a subset of congestion patterns. Different congestors are needed to represent the variety of ways communication patterns can manifest as congestion in the network. For example, congestion may occur at network endpoints from incast or at internal bisection taperings from all-to-all patterns. Another reason to use multiple congestor patterns is that MPI libraries differ in underlying implementation. For example, on the Aries network, the GPCNeT point-to-point incast congestor does not induce congestion because the underlying library is performing Get operations. By deploying four unique congestion patterns, we provide a portable benchmark suite that is useful to a broad community of facilities and researchers (design requirement 1).

Table 1: HPC system configurations.

System Name	Site	Architecture	Processor	MPI Library	Interconnect	Network Topology	# Nodes	# Ports per Node
Crystal	Cray	XC40	Intel Broadwell (mix)	Cray MPICH v7.7.7	Cray Aries	3-level dragonfly with 2 groups and 100% global to injection bandwidth	696	1
Osprey	Cray	CS500	Intel Skylake and Broadwell (mix)	MVAPICH 2.3.1 and OpenMPI 4.0.1	Mellanox EDR	Tree with 100% global to injection bandwidth	128	1
Edison	NERSC	XC50	Intel Ivybridge	Cray MPICH v7.7.7(Pre)	Cray Aries	3-level dragonfly with 50% global to injection bandwidth	5586	1
Theta	ALCF	XC40	Intel KNL	Cray MPICH v7.7.7(Pre)	Cray Aries	3-level dragonfly with 55% global to injection bandwidth	4392	1
Sierra	LLNL	IBM	Power9 and NVIDIA V100	Spectrum MPI v10.2.0.11	Dual-port Mellanox EDR	Tree with 50% global to injection bandwidth	4320	2
Summit	ORNL	IBM	Power9 and NVIDIA V100	Spectrum MPI v10.2.0.11	Dual-port Mellanox EDR	Tree with 100% global to injection bandwidth	4608	2
Malbec	Cray	Slingshot	Intel Skylake (mix)	early Cray MPICH for Shasta	Cray Slingshot with Mellanox ConnectX-5 NICs	2-level dragonfly with 4 groups and 50% global to injection bandwidth	485	1

Figure 3: Histogram of MPI_Allreduce latency measurements at 64 (top) and 696 (bottom) nodes on a XC40 system. The upper and lower panel on each is without and with congestors, respectively. Processes per node are 1, 8, and 32.

To demonstrate the importance of including multiple congestors, we present the impact of individual congestors on two systems (Crystal XC40 and Osprey CS500). Figures 4 and 5 show the results of each individual congestor running with the latency, bandwidth, and MPI_Allreduce benchmarks for these systems with different MPI implementations. For the results on the Osprey system we ran the experiments with both OpenMPI and MVAPICH; however, since the results are similar, we present data only for MVAPICH.

In Figure 4, the P2P Incast congestor has no significant impact on performance of either average or 99th percentile latency or bandwidth for the Crystal system. However, in Figure 5 the P2P Incast congestor has a significant impact on Osprey, increasing mean latency by two orders of magnitude. This shows the substantial differences observed across architectures and the importance of including multiple congestors. For both systems, Alltoall congestors produce a modest amount of congestion for the latency benchmarks, but bandwidth is not impacted. The greatest impact of congestion is from the RMA Incast congestor. On both systems this creates substantial bottlenecks at the network endpoints.

The default congestor ratio (20:20:20:20) is designed to elicit congestion from any HPC system, however facilities and researchers may wish to adjust the ratios to stress a particular type of congestion (e.g., endpoint or intermediate). These ratios, in conjunction with PPort, can be configured within GPCNeT to mimic specific production environments. Additional suggestions on how to analyze system wide congestion are provided in § 4.5.

4.4 Effect of MPI Implementation

Figures 6 - 7 show the same distributions as Figures 1 - 3 except for the CS500 system, Osprey. For Osprey results for both OpenMPI and MVAPICH are shown to demonstrate any effect differences in the MPI library have on the measurements. In Figure 6, isolated point-to-point latency is similar between the two libraries, except that OpenMPI is able to take advantage of potential overlap in the kernel and achieves slightly better latency. The impact on congestion at 1 PPort is negligible for both libraries. At 8 PPort some impact on the tail occurs with a sharp cutoff at 30 μ s and 50 μ s for MVAPICH and OpenMPI respectively. The tail goes out to $\sim 300 \mu$ s at 32 PPort.

Experiments at different PPort for P2P BW+Sync show similar effects on Osprey to the measurements on Crystal, described in §4.2. The bandwidth per rank drops with PPort, and the effect from congestion is only significant at 32 PPort. The differences between OpenMPI and MVAPICH are negligible.

Figure 7 shows the histograms for MPI_Allreduce. Similar to Cray MPICH, both OpenMPI and MVAPICH use a recursive doubling algorithm for the conditions of these tests, and the number of ranks running this kernel is not a power of two. The MVAPICH distribution shows a distinct second peak at 8 and 32 PPort in the isolated runs near 60 μ s. This may be due to the extra steps for non-power of two ranks, but OpenMPI does not have the same peak. Neither library shows any effect from congestors at 1 PPort. Both libraries show an effect at 8 PPort, but MVAPICH has a longer tail approaching 300 μ s and OpenMPI only 40 μ s. Similar behavior occurs at 32 PPort, with congestors inducing samples out to ~ 3 ms for MVAPICH and $\sim 500 \mu$ s for OpenMPI. Both libraries, on EDR Infiniband, are capable of producing and being susceptible to congestion. However, this data shows that subtle differences between MPI implementations can have a moderate impact on the degree of congestion experienced.

Figure 4: Histograms of P2P Lat (left), P2P BW+Sync (center), and MPI_Allreduce (right) at 696 nodes on a XC40 system. The isolated result is shown in the top panel. Result when only one congestor was running is on each panel below that. Each run was with 32 Processes per network Port (PPort). The mean and 99% values are indicated with red and green vertical lines respectively.

Figure 5: Histograms of P2P Lat (left), P2P BW+Sync (center), and MPI_Allreduce (right) at 128 nodes on a CS500 system using the MVAPICH. Isolated result is shown in the top panel. Result with individual congestors is in each panel below that. Each run was with 32 Processes per network Port (PPort). The mean and 99% values are indicated with red and green vertical lines respectively.

4.5 Understanding and Tuning the Scope of Congestion

In the development of GPCNeT, one of our goals was to develop a method to understand how the intensity of its congestors compares to "ambient" congestion when many jobs run simultaneously on a shared system. In this section, we present several techniques that HPC facilities can use to better understand where congestion manifests in the system and how to tune it to emulate specific scenarios. While we demonstrate this for the Cray Aries Network [9], the methodology is easily generalized to other network architectures. This level of evaluation requires system-level counters that all HPC networks provide, namely bytes transmitted and stalls/delays incurred at the router and node level.

These experiments utilized 5575 nodes of NERSC Edison. We present system level counters to examine the flow of traffic as well as the distribution of stalls³ (a proxy for congestion) in the network. This allows us to look at the detailed distributions of communication on a production system. Specifically, we examine the router tiles

and last-hop processor tile (PTile)⁴ counters for each router in the system. For the purposes of this work, we distinguish PTiles from other router tiles. Router tiles include intermediate traffic that may pass through the network, as it takes several hops to reach the final destination. PTile data is specific to the nodes that are locally attached to a given router. Router and PTile datasets complement each other by providing a view of both inter-router and endpoint traffic and congestion, respectively. We utilize these counters to compare performance for several configurations of the benchmark suite.

In each configuration, we reserve 20% of nodes for canaries. Across the runs, we vary (1) the number of MPI processes per node (either 1 or 24 PPort) and (2) the background traffic on the remaining 80% of the nodes. These remaining 80% of nodes operate in one of the following modes:

- Quiet:** remaining nodes are reserved and idle
- Wild:** remaining nodes contain production workloads
- Congestors:** 20% splits of each of the four congestor types

³A stall is a router cycle which no data is transferred even though data resides in the router buffers. This is normally due to a lack of credits or arbitration policies.

⁴Processor tiles are the last-hop within an Aries Network router, and are the entry point to the network for Aries NICs.

Figure 6: Histograms of random ring point-to-point latency at 128 nodes on a CS500 system. Results for both OpenMPI (top two panels) and MVAPICH (bottom two panels) are shown. Isolated measurements and measurements while congestors loaded the network are shown for each library. Processes per node are 1, 8, and 32.

Figure 7: Histograms of MPI_Allreduce latency at 128 nodes on a CS500 system. Results for both OpenMPI (top two panels) and MVAPICH (bottom two panels) are shown. Isolated measurements and measurements while congestors loaded the network are shown for each library. Processes per network Port (PPort) are 1, 8, and 32.

In Figures 8-10 we present the distributions of these counters collected over the entire Edison system for each run. Data is displayed in a violin plot format wherein each vertical bar shows the normalized magnitude of the of throughput or stalls, and the width of the bar signifies the number of routers achieving that measure. Each value is normalized to the mean recorded on the quiet canary run. The three hash marks in each bar show the min, median, and

Figure 8: Distribution of last-hop (endpoint) throughput with production and GPCNeT congestion.

Figure 9: Distribution of last-hop (endpoint) stalls under production and GPCNeT load.

Figure 10: Distribution of router stalls (intermediate congestion) under production and GPCNeT load.

maximum value. These figures show two sets of production runs (Wild-1 and Wild-2) separated over several days to demonstrate the shifts in network load as the variety of jobs changes over time. We obtained 18 wild runs on the Edison system, but space constraints

limits displaying all of them. The two runs shown are approximate of the inner quartile range for Edison bandwidth in the wild (1.7 to 2.0 Gbps for 1 PPort).

Figure 8 shows the endpoint throughput measured across the entire system. Comparing the first two bars we see that the first run utilizing one PPort is unable to fully utilize the bandwidth of the system with a peak bandwidth approximately 20% and a median less than 50% of the Quiet 24 PPort run. Adding production traffic causes peak and median throughput to decrease slightly. It is important to point out that each bar includes the throughput of the production traffic as well as the canary tasks. Therefore, unless congestion is occurring we would expect to see an increase in median traffic compared to the quiet runs. Network congestion also explains why in the wild we do not see any substantial difference in the distributions when comparing 1 PPort and 24 PPort for these runs. We can clearly see throughput diminish for the last two runs which include the congestor benchmarks (Cong 1 PPort and Cong 24 PPort).

Comparing Figure 8 to Figure 9 shows the impact endpoint stalls has on throughput. For a quiet system using 1 PPort, the distribution of endpoint congestion is fairly contained with the max stalls just 1.8X the mean. Increasing PPort to 24, the distribution of stalls begins to creep upward with a median 5.4X greater than 1 PPort. The Wild-1 and Wild-2 runs show very different distributions of endpoint congestion, particularly in the worst case. Wild-2 nears 60X increase to the worst-case stalls of a quiet system. Comparing the wild runs to runs with GPCNeT congestors, we see that worst-case endpoint congestion with GPCNeT is reasonably similar to what may be observed in a production system. The primary difference is that more nodes experience moderate endpoint congestion in the wild than with GPCNeT.

Lastly, we examine congestion for inter-router traffic in Figure 10. We observe that inter-router traffic does not necessarily result in bad performance of the system. The Quiet 24 PPort run shows that the median router stalls is 10X greater than Quiet 1 PPort; however, the 24 PPort run achieved 1.9X bandwidth in GPCNeT. One possible explanation for this is that the system is able to take advantage of adaptive routing, in which the stalls have limited impact on the application. Wild-1 and Wild-2 runs do not appear to be limited by system throughput as the distribution of router stalls is lower than the Quiet 24 PPort job. This is possible if, as Figure 9 suggests, the bottleneck is instead endpoint congestion. Both of the congested runs (Cong 1 PPort and Cong 24 PPort) incur a high number stalls in the worst-case. Additionally, in the 24 PPort run, back-pressure from congestion has spread throughout the entire system and creates a 14.2X increase in median stalls over Quiet 1 PPort.

Tuning congestion benchmarks to perfectly match all combinations of systems and workloads is not a feasible task. However, these charts illustrate how system architects might analyze the GPCNeT benchmarks to evaluate how the intensity of congestion relates to a production system and workload of interest. An example of this can be seen in Figures 9 and 10. We can see that for NERSC Edison workloads, the default configuration of the GPCNeT benchmarks provides a reasonable approximation of worst-case endpoint congestion. Furthermore we can see that the inter-router traffic generated by GPCNeT defaults appears to be 3X to 6X greater than

Figure 11: Mean and P99 Congestion Impact (Eq 1) across a range of systems and architectures for latency, bandwidth, and Allreduce. Raw values are in Table 2. Y-axis is log scale.

what we observe in the wild. However, each system is unique, and for systems running a single large capability job, the background traffic in the wild would place greater demands on inter-router traffic than observed in our runs. By using this information and adjusting the fraction of nodes devoted to each congestor pattern (§ 4.3), it is possible to tune GPCNeT's congestors to approximate a range of systems and workloads.

5 CONGESTION IMPACT ON HPC SYSTEMS

In this section we examine the impact of congestion on a range of systems, reporting the raw results in Table 2 and the congestion impact (see Equation 1). This includes the results on the world's most powerful HPC system, Summit [27], and a next-generation Slingshot [6] system, Malbec. We demonstrate the utility of GPCNeT for systems designers as we evaluate the impact of congestion across topologies and architectures.

5.1 Test Systems

These experiments include several production systems with different architectures and usage models. Table 1 shows details on the configuration of each. We use the default (production) network configurations on each system and a single traffic class for all traffic from GPCNeT following the discussion in § 3. We include CPU-only systems (Crystal, Osprey, Edison, Theta, and Malbec) as well as hybrid CPU-GPU systems (Summit and Sierra). Three network architectures are represented (Cray Aries, Mellanox EDR IB, and Cray Slingshot) with a range of topologies.

The Cray Slingshot network is a new Ethernet-compatible HPC interconnect for the Shasta architecture. Its Rosetta switch ASIC implements 64 ports operating at 200 Gbps/dir, allowing it to scale to over 250,000 endpoints using a dragonfly topology with a diameter of just three switch-to-switch hops. Slingshot provides extensive quality-of-service controls and is able to route packets or flows adaptively. It also has sophisticated congestion management mechanisms that detect contention and quickly provide targeted

Table 2: Sensitive communication kernel performance on test systems in isolation and with congestors loading the network. P2P Lat (Lat) values are in μs , P2P BW+Sync (BW) values are in MiB/s/rank, and Allreduce values are in μs . The data for Osprey is for MVAPICH only.

System Name	Nodes Used	PPort	Isolated Average			Loaded Network Average			Isolated 99%			Loaded Network 99%		
			Lat	BW	Allreduce	Lat	BW	Allreduce	Lat	BW	Allreduce	Lat	BW	Allreduce
Crystal	64	1	1.8	4764.5	9.4	1.9	4496.6	9.8	6.3	3813.8	10.1	5.4	3705.0	11.6
	64	8	1.9	606.0	10.0	9.5	593.7	33.7	7.5	437.7	11.6	25.6	438.4	78.7
	64	32	5.1	133.9	11.7	38.0	111.0	153.5	20.8	101.1	13.8	113.2	89.9	341.9
	696	1	2.0	3735.8	19.9	13.7	3268.4	69.5	6.6	3541.4	22.3	47.1	2756.8	130.4
	696	8	2.0	528.9	20.9	117.2	219.7	1060.6	2.7	380.2	23.3	484.7	113.5	2289.6
	696	20	5.3	187.3	24.2	391.6	20.6	3462.5	9.9	130.5	71.6	2474.9	9.2	9337.0
Osprey	696	32	8.0	111.0	26.0	3212.9	1.6	26765.9	25.1	81.0	102.0	9146.6	0.3	52551.7
	128	1	1.1	6310.1	9.5	1.1	6234.4	9.5	3.9	4677.2	21.5	3.3	4570.7	21.1
		8	1.3	1078.3	16.1	5.8	1022.5	64.0	3.8	708.7	56.9	27.1	628.6	290.7
		20	1.7	451.6	47.9	22.1	274.5	286.5	9.4	289.5	90.1	163.2	136.8	1129.2
Edison	5575	32	1.6	278.5	53.3	40.8	126.1	1088.7	9.7	175.8	88.8	228.9	38.9	2679.3
		1	3.1	1750.1	57.2	216.3	326.9	1427.7	10.1	1603.5	74.3	785.5	232	2492.7
Theta	4096	24	3.0	139.3	52	6014.4	2.1	87983.5	7.2	107.4	70	27117.9	1.3	158883.7
		1	11.2	2635.6	123.5	84.8	695.2	724.2	18.5	2533.0	153.8	370.7	463.3	1289.0
Summit	4500	16	11.1	230.0	119.8	8803.1	3.0	66491.9	16.7	170.7	145.3	36738.1	2.2	136865.4
		4	2.8	2314.6	26.7	32.8	166.1	1424.3	6.6	1971.9	39.2	613.5	33.5	3811.7
Sierra	4200	21	3.3	555.4	34.0	447.9	46.0	9274.4	10.8	378.7	79.4	7142.9	4.6	57770.2
		3	3.0	2747.3	28.8	98.3	672.3	1269.4	7.1	1884.6	37.7	1406.4	194.3	6446.0
Malbec	485	20	3.8	419.9	35.1	1686.6	11.5	24996.0	10.6	265.4	59.9	15276.2	4.6	92193.8
		1.8	565.0	26.6	2.1	523.2	27.9	9.1	475.9	43.7	9.7	437.2	45.2	

back pressure. We examine these capabilities, which are enabled by default on Slingshot, with the GPCNeT benchmark. Although Slingshot is a 200 Gbps/dir network, we limit it to 100 Gbps/dir for our experiments in order to match the 50% global to injection bandwidth of the Theta and Edison systems.

5.2 Congestion Impact

In order to standardize comparisons across architectures with different baseline performance, GPCNeT reports a normalized metric that is the impact on performance from the congestors. The Congestion Impact (CI) is defined for latency (l) and bandwidth (b) as follows:

$$CI_l = l_{\text{congested}} / l_{\text{isolated}} \quad (1)$$

$$CI_b = b_{\text{isolated}} / b_{\text{congested}} \quad (2)$$

where $l_{\text{congested}}$ and l_{isolated} are the latency measurements (e.g., 99%) for either P2P Lat or Allreduce with and without congestors, and $b_{\text{congested}}$ and b_{isolated} are the bandwidth measurements for P2P BW+Sync with and without congestors. In Figure 11 we explore the impact of congestion across multiple systems using high process count to test the limits of the network and how each system operates under pressure. We report the average and 99% Congestion Impact using the default ratio of congestors (20% canary and 80% congestors) and processes per network port ranging from 16 to 24. We use average isolated performance rather than minimum as the baseline for CI in our comparisons, because the mean is more representative of expected performance. In future work we will examine additional parameters, including the impact of reduced congestor ratios and the relative impact of individual congestors.

Figure 11 shows no system is completely immune to the impacts of congestion (results in log scale). In the worst case, latency degradation approaches four orders of magnitude on the Edison system. On the whole, the Aries systems (Crystal, Theta, and Edison) experience the worst congestion. These systems have a taper of approximately 50% for global bandwidth which translates into

a bisection to injection ratio of 25%. Aries does not deploy fine-grained congestion control. While modern InfiniBand systems do employ congestion control mechanisms, configuration of the settings can be difficult and misconfiguration can lead to poor baseline performance [10, 21, 33]. For these reasons, Summit and Sierra have disabled this feature in production. The Infiniband systems (Osprey, Sierra, Summit) perform better than the similarly sized Aries systems. Sierra has a reduced bisection to injection ratio compared to Summit and is more sensitive to congestion. Of all the systems evaluated, the Malbec system shows the lowest Congestion Impact, even when comparing to systems of similar or smaller size. The 99th percentile congestion impact remains under 10 for all tests. The primary reason for this is the advanced congestion control mechanisms of the Slingshot architecture, which detects the congestion and applies appropriate back pressure. In the future we would like to perform a larger set of experiments to show how congestion control algorithms may be optimized, but this is outside the scope of this work.

Last, we measured the application CI to relate the GPCNeT measurements to job performance. This measurement was conducted by running the Lulesh hydrodynamics proxy app [19] on 32 nodes simultaneously with a point-to-point ingress congestion pattern (not the full suite of GPCNeT congestors) on the remaining nodes of Crystal, Osprey, and Malbec. The time per Lulesh iteration is listed in Table 3, along with the application CI,

$$CI_{app} = t_{\text{isolated}} / t_{\text{congested}} \quad (3)$$

(the ratio of iteration timers). Comparing Tables 2 and 3 shows that the application CI follow the same trends as the GPCNeT CIs; networks with high GPCNeT CIs also have higher application CIs. These trends are merely qualitative because (a) Lulesh spends a smaller fraction of its runtime communicating than GPCNeT, making it less sensitive to network effects, and (b) processor differences between the systems also change their absolute runtimes.

Table 3: Lulesh congestion impact on different systems.

System Name	$t_{isolated}$ (ms)	$t_{congested}$ (ms)	CI
Crystal	45	85	1.9
Osprey	55	65	1.2
Malbec	30	30	1.0

6 CONCLUSIONS

Performance variation due to network congestion has been an issue on HPC platforms for many years. Our development of a benchmark suite to quantitatively induce and measure congestion on HPC machines yielded a number of insights into the nature of the causes of congestion and how different networks are affected by it. These insights, and the methods we used to derive them, have a number of implications for future studies of performance variation. By using the benchmark suite as it stands now, potentially with the calibration technique we describe, operators of HPC systems can determine what the sources of contention on their current machines are and actively work to minimize these effects. Additionally, they now have a benchmark and associated metric to evaluate the effectiveness of new network contention reduction mechanisms.

Our work also has implications for MPI (and other communication) library developers. By using the benchmark suite, they can determine how implementations can be designed to utilize protocols and algorithms that can tolerate the effects of network congestion more effectively. Also, by including feedback via the instrumentation mechanisms we describe into their libraries they could even provide runtime adaptivity to current network conditions.

In previous requests for proposals (RFPs) for leadership class machines, a desire for the minimization of network congestion has been expressed by specifying a simple co-efficient of variation between different runs of the same application [25]. In future work we plan to explore the inclusion of this benchmark into future RFPs.

ACKNOWLEDGMENTS

We are grateful to Ramesh Pankajakshan at LLNL for providing data from Sierra, Scott Atchley at ORNL for providing data from Summit and James Botts and Eric Roman at NERSC for providing access to Edison time and system data.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

This research used resources of the Argonne Leadership Computing Facility, which is a U.S. Department of Energy Office of Science User Facility operated under contract DE-AC02-06CH11357. Argonne National Laboratory's work was supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357.

REFERENCES

- [1] Emre Ates, Yijia Zhang, Burak Aksar, Jim Brandt, Vitus J. Leung, Manuel Egele, and Ayse K. Coskun. 2019. HPAS: An HPC Performance Anomaly Suite for Reproducing Performance Variations. In *Proceedings of the 48th International Conference on Parallel Processing (ICPP 2019)*. ACM, New York, NY, USA, Article 40, 10 pages. <https://doi.org/10.1145/3337821.3337907>
- [2] Brian W Barrett and K Scott Hemmert. 2009. An application based MPI message throughput benchmark. In *2009 IEEE International Conference on Cluster Computing and Workshops*. IEEE, 1–8.
- [3] Abhinav Bhatele, Kathryn Mohror, Steven H. Langer, and Katherine E. Isaacs. 2013. There Goes the Neighborhood: Performance Degradation Due to Nearby Jobs. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC '13)*. ACM, New York, NY, USA, Article 41, 12 pages. <https://doi.org/10.1145/2503210.2503247>
- [4] A. Bhatele, A.R. Titus, J.J. Thiagarajan, N. Jain, T. Gamblin, P.-T. Bremer, M. Schulz, and L.V. Kale. 2015. Identifying the Culprits Behind Network Congestion. In *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International*. 113–122. <https://doi.org/10.1109/IPDPS.2015.92>
- [5] Sudheer Chunduri, Kevin Harms, Scott Parker, Vitali Morozov, Samuel Oshin, Naveen Cherukuri, and Kalyan Kumaran. 2017. Run-to-run Variability on Xeon Phi Based Cray XC Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '17)*. ACM, New York, NY, USA, Article 52, 13 pages. <https://doi.org/10.1145/3126908.3126926>
- [6] Cray. 2019. Slingshot: The Interconnect for the Exascale Era. <https://www.cray.com/resources/slingshot-interconnect-for-exascale-era-download>
- [7] Jeffrey Dean and Luiz André Barroso. 2013. The Tail at Scale. *Commun. ACM* 56, 2 (Feb. 2013), 74–80. <https://doi.org/10.1145/2408776.2408794>
- [8] Matthew GF Dosanjh, Taylor Groves, Ryan E Grant, Ron Brightwell, and Patrick G Bridges. 2016. RMA-MT: a benchmark suite for assessing MPI multi-threaded RMA performance. In *2016 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*. IEEE, 550–559.
- [9] Greg Faanes, Abdulla Bataineh, Duncan Roweth, Tom Court, Edwin Froese, Bob Alverson, Tim Johnson, Joe Kopicnik, Mike Higgins, and James Reinhard. 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 (SC '12)*. IEEE Computer Society Press, Los Alamitos, CA, USA, Article 103, 9 pages. <http://dl.acm.org/citation.cfm?id=2388996.2389136>
- [10] Ernst Gunnar Gran, Magne Eimot, Sven-Arne Reinemo, Tor Skeie, Olav Lysne, Lars Paul Huse, and Gilad Shainer. 2010. First experiences with congestion control in InfiniBand hardware. In *2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS)*. IEEE, 1–12.
- [11] Ryan E. Grant, Kevin T. Pedretti, and Ann Gentile. 2015. Overtime: A Tool for Analyzing Performance Variation Due to Network Interference. In *Proceedings of the 3rd Workshop on Exascale MPI (ExaMPI '15)*. ACM, New York, NY, USA, Article 4, 10 pages. <https://doi.org/10.1145/2831129.2831133>
- [12] T. Groves, Y. Gu, and N. J. Wright. 2017. Understanding Performance Variability on the Aries Dragonfly Network. In *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. 809–813. <https://doi.org/10.1109/CLUSTER.2017.76>
- [13] Torsten Hoefer, Andrew Lumsdaine, and Wolfgang Rehm. 2007. Implementation and Performance Analysis of Non-blocking Collective Operations for MPI. In *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing (SC '07)*. ACM, New York, NY, USA, Article 52, 10 pages. <https://doi.org/10.1145/1362622.1362692>
- [14] Torsten Hoefer, Torsten Mehlman, Andrew Lumsdaine, and Wolfgang Rehm. 2007. Netgauge: A network performance measurement framework. In *International Conference on High Performance Computing and Communications*. Springer, 659–671.
- [15] T. Hoefer, T. Schneider, and A. Lumsdaine. 2008. Multistage switches are not crossbars: Effects of static routing in high-performance networks. In *2008 IEEE International Conference on Cluster Computing*. 116–125. <https://doi.org/10.1109/CLUSTER.2008.4663762>
- [16] InfiniBand Trade Association. 2019. <http://www.infinibandta.org/>.
- [17] Intel. 2015. Intel MPI benchmarks 4.0. <https://software.intel.com/en-us/articles/intel-mpi-benchmark>
- [18] S. Jha, A. Patke, J. Brandt, A. Gentile, M. Showerman, E. Roman, Z. Kalbarczyk, and R. Iyer. 2019. A Study of Network Congestion in Two Supercomputing High-Speed Interconnects. In *Proceedings of the 26th Symposium on High Performance Interconnects (HotI) (HotI '19)*.
- [19] Ian Karlin, Jeff Keasler, and Rob Neely. 2013. *LULESH 2.0 Updates and Changes*. Technical Report LLNL-TR-641973. 1–9 pages.
- [20] Jiuxing Liu, Balasubramanian Chandrasekaran, Weikuan Yu, Jiesheng Wu, Darius Buntinas, Sushmitha Kini, Dhableswar K Panda, and Pete Wyckoff. 2004. Microbenchmark performance comparison of high-speed cluster interconnects. *Ieee Micro* 24, 1 (2004), 42–51.
- [21] Qian Liu, Robert D Russell, and Ernst Gunnar Gran. 2016. Improvements to the InfiniBand congestion control mechanism. In *2016 IEEE 24th Annual Symposium on High-Performance Interconnects (HOTI)*. IEEE, 27–36.
- [22] S. Liu, H. Xu, L. Liu, W. Bai, K. Chen, and Z. Cai. 2018. RepNet: Cutting Latency with Flow Replication in Data Center Networks. *IEEE Transactions on Services Computing* (2018), 1–1. <https://doi.org/10.1109/TSC.2018.2793250>
- [23] Piotr R Luszczek, David H Bailey, Jack J Dongarra, Jeremy Kepner, Robert F Lucas, Rolf Rabenseifner, and Daisuke Takahashi. 2006. The HPC Challenge (HPCC) Benchmark Suite. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing (SC '06)*. ACM, New York, NY, USA, Article 213. <https://doi.org/>

- 10.1145/1188455.1188677
- [24] Pulkit A. Misra, María F. Borge, Íñigo Goiri, Alvin R. Lebeck, Willy Zwaenepoel, and Ricardo Bianchini. 2019. Managing Tail Latency in Datacenter-Scale File Systems Under Production Constraints. In *Proceedings of the Fourteenth EuroSys Conference 2019 (EuroSys '19)*. ACM, New York, NY, USA, Article 17, 15 pages. <https://doi.org/10.1145/3302424.3303973>
 - [25] Department of Energy. 2013. Trinity / NERSC-8 RFP. <https://www.nersc.gov/users/computational-systems/cori/nersc-8-procurement/trinity-nersc-8-rfp/>
 - [26] David Skinner and William Kramer. 2005. Understanding the causes of performance variability in HPC workloads. In *Workload Characterization Symposium, 2005. Proceedings of the IEEE International*. IEEE, 137–149.
 - [27] Summit: Oak Ridge National Laboratory's 200 petaflop supercomputer . 2019. <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>.
 - [28] Nathan R Tallent, Abhinav Vishnu, Hubertus Van Dam, Jeff Daily, Darren J Kerbyson, and Adolfo Hoisie. 2015. Diagnosing the causes and severity of one-sided message contention. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 130–139.
 - [29] The Gen-Z Consortium. 2019. <https://genzconsortium.org/>.
 - [30] Robert Underwood, Jason Anderson, and Amy Apon. 2018. Measuring Network Latency Variation Impacts to High Performance Computing Application Performance. In *Proceedings of the 2018 ACM/SPEC International Conference on Performance Engineering (ICPE '18)*. ACM, New York, NY, USA, 68–79. <https://doi.org/10.1145/3184407.3184427>
 - [31] Hans Weeks, Matthew GF Dosanjh, Patrick G Bridges, and Ryan E Grant. 2016. SHMEM-MT: a benchmark suite for assessing multi-threaded SHMEM performance. In *Workshop on OpenSHMEM and Related Technologies*. Springer, 227–231.
 - [32] Yunqi Zhang, David Meisner, Jason Mars, and Lingjia Tang. 2016. Treadmill: Attributing the Source of Tail Latency Through Precise Load Testing and Statistical Inference. In *Proceedings of the 43rd International Symposium on Computer Architecture (ISCA '16)*. IEEE Press, Piscataway, NJ, USA, 456–468. <https://doi.org/10.1109/ISCA.2016.47>
 - [33] Yibo Zhu, Haggai Eran, Daniel Firestone, Chuanxiong Guo, Marina Lipshteyn, Yehonatan Liron, Jitendra Padhye, Shachar Raindel, Mohamad Haj Yahia, and Ming Zhang. 2015. Congestion control for large-scale RDMA deployments. *ACM SIGCOMM Computer Communication Review* 45, 4 (2015), 523–536.

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

Throughout the development paper we ran several versions of GPCNeT as we incorporated new features and output additional statistics. The underlying structure of the benchmark remains the same as described in Section 3.1.1 and 3.2.1 for all of the runs. The source for each version is saved with the data that it output for a given system. And a history of the code is retained in git. GPCNeT requires an MPI-3 compatible MPI library. The specific versions used as well as details of each system are listed in Table 1 of the paper and here:

| System Name, Site, Architecture, Processor, MPI Library, Interconnect & Network Topology, # Nodes, # Ports per Node |

| Crystal, Cray, XC40, Intel Broadwell (mix), Cray Aries, 3-level dragonfly with 2 groups and 50% global links, 696 nodes | Osprey, Cray, CS500, Intel Skylake and Broadwell (mix), Mellanox EDR, Non-blocking tree with adaptive routing enabled, 128 nodes | Edison, NERSC, XC50, Intel Ivybridge, Cray Aries, 3-level dragonfly with 50% global links, 5586 nodes | Theta, ALCF, XC40, Intel KNL, Cray Aries, 3-level dragonfly with 55% global links, 4392 nodes | Sierra, LLNL, IBM, Power9 and NVIDIA V100, Dual-rail Mellanox EDR, Tree with 2:1 taper, 4320 nodes | Summit, ORNL, IBM, Power9 and NVIDIA V100, Dual-rail Mellanox EDR, Non-blocking tree, 4608 nodes | Malbec, Cray, Intel Skylake (mix), Cray Slingshot with Mellanox ConnectX-5 NICs, 2-level dragonfly with 4 groups and 100% global links reduced to 25 Gbps, 485 nodes |

The degree of congestion created by GPCNeT is dependent on the number of processes and nodes that it runs on in the system. These details are included in Table 2 and where applicable the Figures in the paper.

The Edison System was utilized from January to mid May of 2019 for the results in this paper.

ARTIFACT AVAILABILITY

Software Artifact Availability: All author-created software artifacts are maintained in a public repository under an OSI-approved license.

Hardware Artifact Availability: There are no author-created hardware artifacts.

Data Artifact Availability: All author-created data artifacts are maintained in a public repository under an OSI-approved license.

Proprietary Artifacts: No author-created artifacts are proprietary.

List of URLs and/or DOIs where artifacts are available:

<https://github.com/netbench/gpcnet>

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: Crystal, XC40, Intel Broadwell (mix), Cray Aries, 3-level dragonfly with 2 groups and 50% global links, 696 nodes | Osprey, CS500, Intel Skylake and Broadwell (mix), Mellanox

EDR, Non-blocking tree with adaptive routing enabled, 128 nodes | Edison, XC50, Intel Ivybridge, Cray Aries, 3-level dragonfly with 50% global links, 5586 nodes | Theta, XC40, Intel KNL, Cray Aries, 3-level dragonfly with 55% global links, 4392 nodes | Sierra, IBM, Power9 and NVIDIA V100, Dual-rail Mellanox EDR, Tree with 2:1 taper, 4320 nodes | Summit, IBM, Power9 and NVIDIA V100, Dual-rail Mellanox EDR, Non-blocking tree, 4608 nodes | Malbec, Cray, Intel Skylake (mix), Cray Slingshot with Mellanox ConnectX-5 NICs, 2-level dragonfly with 4 groups and 100% global links reduced to 25 Gbps, 485 nodes

Libraries and versions: Crystal, Cray MPICH v7.7.7 | Osprey, MVAPICH 2.3.1 and OpenMPI 4.0.1 | Edison, Cray MPICH v7.7.7(Pre) | Theta, Cray MPICH v7.7.7(Pre) | Sierra, Spectrum MPI v10.2.0.11 | Summit, Spectrum MPI v10.2.0.11 | Malbec, early Cray MPICH for Shasta

Output from scripts that gathers execution environment information.

Edison:

```
SLURM_NODELIST=nid00188
LESSKEY=/etc/lesskey.bin
PE_LIBSCI_VOLATILE_PRGENV=CRAY GNU INTEL
MKLR00T=/opt/intel/compilers_and_libraries_2018.1.16
  ↪ 3/linux/mkl
PE_SMA_DEFAULT_PKGCONFIG_VARIABLES=PE_SMA_COMPFLAG_@
  ↪ prgenv@
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_mic_knl=160
MODULE_VERSION_STACK=3.2.10.6
GREP_COLOR=1;32
KSH_AUTOLOAD=1
ZAP_LIBPATH=/opt/ovis/lib64/ovis-lib
SLURM_CHECKPOINT_IMAGE_DIR=/var/slurm/checkpoint
SLURM_JOB_NAME=sh
MANPATH=/usr/common/software/man:/usr/common/mss/man
  ↪ :/usr/common/nsg/man:/opt/cray/pe/mpt/7.7.3/gni/
  ↪ man/mpich:/opt/cray/pe/atp/2.1.3/man:/opt/cray/a
  ↪ lps/6.6.43-6.0.7.0_26.4__ga796da3.ari/man:/opt/c
  ↪ ray/job/2.2.3-6.0.7.0_44.1__g6c4e934.ari/man:/op
  ↪ t/cray/pe/pmi/5.0.14/man:/opt/cray/pe/libsci/18.
  ↪ 07.1/man:/opt/cray/pe/man/csmlversion:/opt/cray/
  ↪ pe/craype/2.5.15/man:/opt/intel/compilers_and_li
  ↪ braries_2018.1.163/linux/man/common:/usr/syscom/
  ↪ nsg/man:/opt/cray/pe/modules/3.2.10.6/share/man:
  ↪ /usr/local/man:/usr/share/man:/opt/cray/share/ma
  ↪ n:/opt/cray/pe/man:/opt/cray/share/man:/opt/cray
  ↪ /share/man
PE_CXX_PKGCONFIG_LIBS=mpichcxx
PE_PETSC_DEFAULT_GENCOMPS_CRAY_skylake=86
PE_PAPI_DEFAULT_ACCEL_FAMILY_LIBS_nvidia=-lcupti,-l
  ↪ cudart,-lcuda
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
NNTPSERVER=news
```



```

DVS_MAXNODES=1__
PE_TPSL_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
XDG_SESSION_ID=74437
PE_TRILINOS_DEFAULT_GENCOMPS_CRAY_x86_64=87
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_FFTW_DEFAULT_TARGET_mic_knl=mic_knl
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
HOSTNAME=edison04
CRAY_UDREG_INCLUDE_OPTS=-I/opt/cray/udreg/2.3.2-6.0.
↪ 7.0.33.18__g5196236.ari/include
SLURM_TOPOLOGY_ADDR=s15.s13.nid00188
SLURMD_NODENAME=nid00188
CRAY_SITE_LIST_DIR=/etc/opt/cray/pe/modules
PE_SMA_DEFAULT_COMPFLAG_GNU=-fcray-pointer
LIBRARYMODULES=acml:alps:cray-dwarf:cray-fftw:cray-g
↪ a:cray-hdf5:cray-hdf5-parallel:cray-libsci:cray-
↪ libsci_acc:cray-mpich:cray-mpich-abi:cray-mpich2
↪ :cray-netcdf:cray-netcdf-hdf5parallel:cray-paral
↪ lel-netcdf:cray-petsc:cray-petsc-complex:cray-sh
↪ mem:cray-tpsl:cray-trilinos:cuda-toolkit:fftw:ga:
↪ hdf5:hdf5-parallel:iobuf:libfast:netcdf:netcdf-h
↪ df5parallel:ntk:onesided:papi:petsc:petsc-comple
↪ x:pmi:tpsl:trilinos:xt-libsci:xt-mpich2:xt-mpt:x
↪ t-papi
XKEYSYMDB=/usr/X11R6/lib/X11/XKeysymDB
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/
↪ opt/cray/pe/parallel-netcdf/1.8.1.3/@PRGENV@/PE
↪ _PARALLEL_NETCDF_DEFAULT_GENCOMPS@/lib/pkgconfig
PE_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/
↪ pe/netcdf/4.6.1.3/@PRGENV@/PE_NETCDF_DEFAULT_GE
↪ NCOMPS@/lib/pkgconfig
PE_TRILINOS_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cra
↪ y/pe/trilinos/12.12.1.1/@PRGENV@/PE_TRILINOS_DE
↪ FAULT_GENCOMPS@/PE_TRILINOS_DEFAULT_TARGET@/lib
↪ /pkgconfig
SLURM_Prio_PROCESS=0
RCLOCAL_BASEOPTS=true
PE_MPICH_ALTERNATE_LIBS_dpm=_dpm
PE_ENV=INTEL
INTEL_LICENSE_FILE=28518@crayintel.licenses.nersc.go
↪ v:28518@intel.licenses.nersc.gov
PE_SMA_DEFAULT_COMPFLAG=
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_HDF5_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
SLURM_SRUN_COMM_PORT=62703
SHELL=/bin/bash
PE_TPSL_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
PKGCONFIG_ENABLED=1
TERM=xterm-color
HOST=edison04
SLURM_JOB_QOS=debug
PROFILEREAD=true
PE_PETSC_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2 7.1
↪ 6.1
CLICOLOR=1

```

```

PROJECT=/project/projectdirs/m888/USER
INTEL_MINOR_VERSION=18
HISTSIZE=1000
SLURM_PTY_WIN_ROW=40
PE_PETSC_DEFAULT_GENCOMPS_GNU_haswell=71 53 49
PE_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
PE_TPSL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSCI
PE_TPSL_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53 49
TMPDIR=/tmp
CRAYPE_DIR=/opt/cray/pe/craype/2.5.15
SSH_CLIENT=128.55.209.26 49906 22
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
CRAY_UGNI_POST_LINK_OPTS=-L/opt/cray/ugni/6.0.14.0-6
↪ .0.7.0.23.1__gea11d3d.ari/lib64
CRAY_XPMEM_POST_LINK_OPTS=-L/opt/cray/xpmem/2.2.15-6
↪ .0.7.1.5.10__g7549d06.ari/lib64
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_TRILINOS_DEFAULT_VOLATILE_PRGENV=CRAY GNU INTEL
PE_PETSC_DEFAULT_GENCOMPS_INTEL_haswell=160
SLURM_TOPOLOGY_ADDR_PATTERN=switch.switch.node
LIBRARY_PATH=/opt/intel/compilers_and_libraries_2018
↪ .1.163/linux/compiler/lib/intel64:/opt/intel/com
↪ pilers_and_libraries_2018.1.163/linux/mkl/lib/in
↪ tel64
PE_HDF5_DEFAULT_VOLATILE_PRGENV=GNU
CRAY_MPICH2_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-int
↪ el/16.0
ALTD_SELECT_OFF_USERS=
PE_PETSC_DEFAULT_GENCOMPS_CRAY_interlagos=86
ALTD_LINKER=/usr/common/software/altd/2.0/bin/ld
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/op
↪ t/cray/pe/hdf5-parallel/1.10.2.0/@PRGENV@/PE_HD
↪ F5_PARALLEL_DEFAULT_GENCOMPS@/lib/pkgconfig
PE_FFTW_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /fftw/3.3.8.1/@PE_FFTW_DEFAULT_TARGET@/lib/pkgco
↪ nfig
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PA
↪ TH=/opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/@PRG
↪ ENV@/PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS@/l
↪ ib/pkgconfig
PE_TPSL_DEFAULT_GENCOMPS_CRAY_mic_knl=86
INTEL_PATH=/opt/intel/compilers_and_libraries_2018.1
↪ .163
ALTD_SELECT_ON=0
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_GA_DEFAULT_VOLATILE_PRGENV=GNU
PE_LIBSCI_DEFAULT_GENCOMPS_GNU_x86_64=71 61 51 49
PMI_CONTROL_PORT=63823
PE_MPICH_GENCOMPS_GNU=71 51 49
MORE--s1
PE_TPSL_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_PKGCONFIG_PRODUCTS=PE_MPICH:PE_LIBSCI
PE_MPICH_DEFAULT_GENCOMPILERS_GNU=7.1 5.1 4.9

```

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

```

FPATH=/opt/cray/pe/modules/3.2.10.6/init/sh_funcs/n
↳ o_redirect:/opt/cray/pe/modules/3.2.10.6/init/sh
↳ _funcs/no_redirect:/opt/cray/pe/modules/3.2.10.6
↳ /init/sh_funcs/no_redirect
SLURM_CPU_BIND_VERBOSE=quiet
SLURM_SPANK_NERSC_ZONESORT_INTERVAL=0
PE_TPSL_64_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray
↳ /pe/tpsl/18.06.1/@PRGENV@64/@PE_TPSL_64_DEFAULT_
↳ GENCOMPS@/@PE_TPSL_64_DEFAULT_TARGET@/lib/pkgcon
↳ fig
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_haswell=86
ALTD_VERBOSE=0
PE_PETSC_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSC
↳ I:PE_HDF5_PARALLEL:PE_TPSL
PE_PAPI_DEFAULT_ACCEL_LIBS_nvidia35=-lcupti,-lcudar
↳ t,-lcuda
PE_CRAY_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/pa
↳ rallel-netcdf/1.8.1.3/CRAY/8.6/lib/pkgconfig:/op
↳ t/cray/pe/netcdf-hdf5parallel/4.6.1.3/CRAY/8.6/l
↳ ib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/CRAY/8.
↳ 6/lib/pkgconfig:/opt/cray/pe/hdf5-parallel/1.10.
↳ 2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/hdf5/1.1
↳ 0.2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/ga/5.3
↳ .0.8/CRAY/8.6/lib/pkgconfig
PE_TRILINOS_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.7
PE_PETSC_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
SSH_TTY=/dev/pts/40
PE_FORTRAN_PKGCONFIG_LIBS=mpichf90
PE_LIBSCI_DEFAULT_OMP_REQUIRES_openmp=_mp
PE_SMA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/
↳ mpt/7.7.3/gni/sma@PE_SMA_DEFAULT_DIR_DEFAULT64@/
↳ lib64/pkgconfig
ALLINEA_QUEUE_DLL=/opt/cray/pe/mpt/7.7.3/gni/mpich-i
↳ ntel/16.0/lib/libtvmmpich.so.3.0.1
SLURM_CPU_BIND_LIST=0xFFFFFFFFFFFF
PE_TRILINOS_DEFAULT_GENCOMPS_INTEL_x86_64=160
MPICH_SHARED_MEM_COLL_OPT=0
CRAY_MPICH_BASEDIR=/opt/cray/pe/mpt/7.7.3/gni
ALPS_APP_ID=12937886
SLURM_NNODES=1
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2
↳ 7.1 6.1 5.3 4.9
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_haswell=160
JRE_HOME=/usr/lib64/jvm/java/jre
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
PE_HDF5_PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1
↳ 5.3 4.9
USER=USER
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
LS_COLORS=
LD_LIBRARY_PATH=/opt/cray/job/2.2.3-6.0.7.0_44.1__g6
↳ c4e934.ari/lib64:/opt/intel/compilers_and_librar
↳ ies_2018.1.163/linux/compiler/lib/intel64:/opt/i
↳ ntel/compilers_and_libraries_2018.1.163/linux/mk
↳ l/lib/intel64

```

```

PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_TRILINOS_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.3
↳ 5.1 4.9
PE_FFTW_DEFAULT_TARGET_interlagos=interlagos
PE_LIBSCI_DEFAULT_VOLATILE_PRGENV=CRAY GNU INTEL
PE_TPSL_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TRILINOS_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
CRAY_RCA_POST_LINK_OPTS=-L/opt/cray/rca/2.2.18-6.0.7
↳ .0_33.3__g2aa4f39.ari/lib64
↳ -lrca
PE_PETSC_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↳ GNU64 INTEL INTEL64
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2
↳ 7.1 5.3 4.9
PE_PKGCONFIG_LIBS=mpich:AtpSigHandler:cray-rca:libsc
↳ i_mpi:libsci
PE_MPICH_FIXED_PRGENV=INTEL
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_LIBSCI_PKGCONFIG_VARIABLES=PE_LIBSCI_OMP_REQUIRES
↳ _@openmp@:PE_SCI_EXT_LIBPATH:PE_SCI_EXT_LIBNAME
CSCRATCH=/global/cscratch1/sd/USER
CRAY_IAA_INFO_FILE=/tmp/cray_iaa_info.12937886
PE_TPSL_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
XNLSPATH=/usr/share/X11/nls
PE_PETSC_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_PETSC_DEFAULT_GENCOMPS_GNU_interlagos=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_GNU_sandybridge=71 53 49
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_sandybridge=160
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_skylake=16.0
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_mic_knl=160
PE_LIBSCI_DEFAULT_GENCOMPS_CRAY_x86_64=86
MPICH_ABORT_ON_ERROR=1
INTEL_VERSION=18.0.1.163
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_PETSC_DEFAULT_GENCOMPS_GNU_mic_knl=53
MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-intel/16.0
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2
↳ 7.1 5.3 4.9
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_sandybridge=160
LDAPTLS_REQCERT=never
PE_PAPI_DEFAULT_PKGCONFIG_VARIABLES=PE_PAPI_ACCEL_LI
↳ BS_@accelerator@
ALTD_ON=1
SLURM_STEP_NUM_NODES=1
SLURM_JOBID=12937886
PE_MPICH_FORTRAN_PKGCONFIG_LIBS=mpichf90
MPICH_MPIIO_DVS_MAXNODES=14
PE_FFTW_DEFAULT_TARGET_sandybridge=sandybridge
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
HOSTTYPE=x86_64
PE_NETCDF_HDF5PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_
↳ HDF5_PARALLEL
PE_HDF5_PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH

```

```

ATP_POST_LINK_OPTS=-Wl,-L/opt/cray/pe/atp/2.1.3/libA
↪ pp/
CPATH=/opt/intel/compilers_and_libraries_2018.1.163/
↪ linux/mkl/include
PE_FFTW_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
SRUN_DEBUG=3
RCLOCAL_PRGENV=true
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_mic_knl=5.3
PE_LIBSCI_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_TPSSL_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53 49
PE_PRODUCT_LIST=CRAYPE_IVYBRIDGE:CRAY_RCA:CRAY_ALPS:
↪ DVS:CRAY_XPMEM:CRAY_DMAPP:CRAY_PMI:CRAY_UGNI:CRAY
↪ Y_UDREG:CRAY_LIBSCI:CRAYPE:INTEL
FROM_HEADER=
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_LIBSCI_GENCOMPS_INTEL_x86_64=160
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 5.3 4.9
PE_TPSSL_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
CRAY_MPICH_ROOTDIR=/opt/cray/pe/mpt/7.7.3
PAGER=less
SLURM_LAUNCH_NODE_IPADDR=10.128.6.142
ALPS_LLI_STATUS_OFFSET=1
HDF5_USE_FILE_LOCKING=FALSE
PE_MPICH_GENCOMPILERS_CRAY=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_skylake=160
PE_LIBSCI_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1 4.9
PE_MPICH_MODULE_NAME=cray-mpich
PE_PETSC_DEFAULT_GENCOMPS_GNU_skylake=61
PE_TPSSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
CSHEDIT=emacs
SLURM_STEP_ID=0
ALPS_APP_PE=0
PE_MPICH_DEFAULT_VOLATILE_PRGENV=CRAY GNU
PE_MPICH_TARGET_VAR_nvidia20=-lcudart
XDG_CONFIG_DIRS=/etc/xdg
PE_LIBSCI_GENCOMPS_CRAY_x86_64=86
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_LIBSCI_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_TPSSL_64_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIB
↪ SCI
LSCOLORS=ExFxCxDxBxegeabagacad
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_sandybridge=86
INTEL_MAJOR_VERSION=18
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49
NLSPATH=/opt/intel/compilers_and_libraries_2018.1.16
↪ 3/linux/compiler/lib/intel64/locale/%l_%t/%N:/op
↪ t/intel/compilers_and_libraries_2018.1.163/linux
↪ /mkl/lib/intel64/locale/%l_%t/%N
CRAY_DMAPP_INCLUDE_OPTS=-I/opt/cray/dmapp/7.1.1-6.0.
↪ 7.0_34.3__g5a674e0.ari/include
↪ -I/opt/cray/gni-headers/5.0.12.0-6.0.7.0_24.1__g
↪ 3b1768f.ari/include
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPS_GNU=51 49
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
LIBGL_DEBUG=quiet
PE_NETCDF_DEFAULT_GENCOMPS_GNU=

```

```

DVS_VERSION=0.9.0
MINICOM=-c on
PE_LIBSCI_PKGCONFIG_LIBS=libsci_mpi:libsci
USERMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:PrgE
↪ nv-pathscale:PrgEnv-pgi:acml:alps:apprentice:app
↪ rentice2:atp:blcr:cce:chapel:cray-cddb:cray-fftw
↪ :cray-ga:cray-hdf5:cray-hdf5-parallel:cray-lgdb:
↪ cray-libsci:cray-libsci_acc:cray-mpich:cray-mpic
↪ h-compat:cray-mpich2:cray-netcdf:cray-netcdf-hdf
↪ 5parallel:cray-parallel-netcdf:cray-petsc:cray-p
↪ etsc-complex:cray-shmem:cray-snp-launcher:cray-tp
↪ sl:cray-trilinos:craypat:craype:craypkg-gen:cuda
↪ toolkit:ddt:fftw:ga:gcc:hdf5:hdf5-parallel:intel
↪ :iobuf:java:lgdb:libfast:libsci_acc:mpich1:netcd
↪ f:netcdf-hdf5parallel:netcdf-nofsync:netcdf-nofs
↪ ync-hdf5parallel:ntk:onesided:papi:parallel-netc
↪ df:pathscale:perftools:perftools-lite:petsc:pets
↪ c-complex:pgi:pmi:stat:totalview:tpsl:trilinos:x
↪ t-asyncpe:xt-craypat:xt-lgdb:xt-libsci:xt-mpich2
↪ :xt-mpt:xt-papi:xt-shmem:xt-totalview
CRAY_LIBSCI_DIR=/opt/cray/pe/libsci/18.07.1
CRAY_LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
SLURM_TASKS_PER_NODE=48
MAIL=/var/mail/USER
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1 5.3
↪ 4.9
PATH=/usr/common/software/altd/2.0/bin:/usr/common/s
↪ oftware/bin:/usr/common/mss/bin:/usr/common/nsg/
↪ bin:/opt/cray/pe/mpt/7.7.3/gni/bin:/opt/cray/rca
↪ /2.2.18-6.0.7.0_33.3__g2aa4f39.ari/bin:/opt/cray
↪ /alps/6.6.43-6.0.7.0_26.4__ga796da3.ari/sbin:/op
↪ t/cray/job/2.2.3-6.0.7.0_44.1__g6c4e934.ari/bin:
↪ /opt/cray/pe/craype/2.5.15/bin:/opt/intel/compil
↪ ers_and_libraries_2018.1.163/linux/bin/intel64:/
↪ opt/ovis/bin:/opt/ovis/sbin:/usr/syscom/nsg/sbin
↪ :/usr/syscom/nsg/bin:/opt/cray/pe/modules/3.2.10
↪ .6/bin:/usr/local/bin:/usr/bin:/bin:/usr/bin/X11
↪ :/usr/games:/usr/lib/mit/bin:/usr/lib/mit/sbin:/
↪ opt/cray/pe/bin
PE_MPICH_GENCOMPILERS_GNU=7.1 5.1 4.9
LIBGL_ALWAYS_INDIRECT=1
MODULE_VERSION=3.2.10.6
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_interlagos=86
PE_HDF5_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /hdf5/1.10.2.0/@PRGENV@/@PE_HDF5_DEFAULT_GENCOMP
↪ S@/lib/pkgconfig
PE_PKGCONFIG_DEFAULT_PRODUCTS=PE_TRILINOS:PE_TPSSL_64
↪ :PE_TPSSL:PE_PETSC:PE_PARALLEL_NETCDF:PE_NETCDF_H
↪ DF5PARALLEL:PE_NETCDF:PE_MPICH:PE_LIBSCI:PE_HDF5
↪ _PARALLEL:PE_HDF5:PE_GA:PE_FFTW2:PE_FFTW
SLURM_STEP_LAUNCHER_PORT=62703
PE_NETCDF_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPILERS_GNU=5.1 4.9
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1 5.3
↪ 4.9
PE_FFTW_DEFAULT_TARGET_abudhabi=abudhabi

```

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

```

PE_PETSC_DEFAULT_GENCOMPS_CRAY_mic_knl=86
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2
↪ 7.1 6.1
XTPE_NETWORK_TARGET=aries
ATP_IGNORE_SIGTERM=1
CPU=x86_64
SLURM_WORKING_CLUSTER=edison:ctl1:6817:8448
PMI_CRAY_NO_SMP_ORDER=0
_=/usr/bin/env
SLURM_JOB_ID=12937886
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS_GNU=
LDMSD_PLUGIN_LIBPATH=/opt/ovis/lib64/ovis-ldms
PE_SMA_DEFAULT_DIR_CRAY_DEFAULT64=64
PE_NETCDF_HDF5PARALLEL_DEFAULT_FIXED_PRGENV=CRAY
↪ INTEL
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
PE_HDF5_PARALLEL_DEFAULT_GENCOMPS_GNU=
PE_HDF5_PARALLEL_DEFAULT_FIXED_PRGENV=CRAY INTEL
JAVA_BINDIR=/usr/lib64/jvm/java/bin
PMI_NO_FORK=1
CRAY_UDREG_POST_LINK_OPTS=-L/opt/cray/udreg/2.3.2-6.
↪ 0.7.0_33.18__g5196236.ari/lib64
PE_TPSL_64_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↪ GNU64 INTEL INTEL64
PE_TPSL_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_sandybridge=86
INPUTRC=/etc/inputrc
PE_MPICH_VOLATILE_PRGENV=CRAY GNU
PWD=/global/homes/t/USER
CRAYPE_VERSION=2.5.15
PE_TPSL_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
CRAY_ALPS_POST_LINK_OPTS=-L/opt/cray/alps/6.6.43-6.0
↪ .7.0_26.4__ga796da3.ari/lib64
ARLOG=/global/cscratch1/cori/aries
SLURM_STEPID=0
SLURM_JOB_USER=USER
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_sandybridge=7.1
↪ 5.3 4.9
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_MPICH_DEFAULT_GENCOMPS_CRAY=86

TARGETMODULES=craype-abudhabi:craype-abudhabi-cu:craype-accel-host:craype-accel-nvidia20:craype-accel-nvidia30:craype-accel-nvidia35:craype-barcelona:craype-broadwell:craype-haswell:craype-hugepages128K:craype-hugepages128M:craype-hugepages16M:craype-hugepages256M:craype-hugepages2M:craype-hugepages32M:craype-hugepages4M:craype-hugepages512K:craype-hugepages512M:craype-hugepages64M:craype-hugepages8M:craype-intel-knc:craype-interlagos:craype-interlagos-cu:craype-istanbul:craype-ivybridge:craype-mc12:craype-mc8:craype-mic-knl:craype-network-aries:craype-network-gemini:craype-network-infiniband:craype-network-none:craype-network-seastar:craype-sandybridge:craype-shanghai:craype-target-compute_node:craype-target-local_host:craype-target-native:craype-xeon:xtpe-barcelona:xtpe-interlagos:xtpe-interlagos-cu:xtpe-istanbul:xtpe-mc12:xtpe-mc8:xtpe-network-gemini:xtpe-network-seastar:xtpe-shanghai:xtpe-target-native:xtpe-xeon
_LMFFILES=/opt/cray/pe/modulefiles/modules/3.2.10.6:/usr/syscom/nsg/modulefiles/nsg/1.2.0:/opt/modulefiles/intel/18.0.1.163:/opt/cray/pe/craype/2.5.15/modulefiles/craype-network-aries:/opt/cray/pe/modulefiles/craype/2.5.15:/opt/cray/pe/modulefiles/cray-libsci/18.07.1:/opt/cray/ari/modulefiles/udreg/2.3.2-6.0.7.0_33.18__g5196236.ari:/opt/cray/ari/modulefiles/ugni/6.0.14.0-6.0.7.0_23.1__gea11d3d.ari:/opt/cray/pe/modulefiles/pmi/5.0.14:/opt/cray/ari/modulefiles/dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari:/opt/cray/ari/modulefiles/gni-headers/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari:/opt/cray/ari/modulefiles/xpmem/2.2.15-6.0.7.1_5.10__g7549d06.ari:/opt/cray/ari/modulefiles/job/2.2.3-6.0.7.0_44.1__g6c4e934.ari:/opt/cray/ari/modulefiles/dvs/2.7_2.2.116-6.0.7.1_8.11__g0dac9eb:/opt/cray/ari/modulefiles/alps/6.6.43-6.0.7.0_26.4__ga796da3.ari:/opt/cray/ari/modulefiles/rca/2.2.18-6.0.7.0_33.3__g2aa4f39.ari:/opt/cray/pe/modulefiles/atp/2.1.3:/opt/cray/pe/modulefiles/PrGenv-intel/6.0.4:/opt/cray/pe/craype/2.5.15/modulefiles/craype-ivybridge:/opt/cray/pe/modulefiles/cray-mpich/7.7.3:/usr/common/software/modulefiles/altd/2.0:/opt/modulefiles/Base-opts/2.4.135-6.0.7.0_38.1__g718f891.ari
PE_LIBSCI_DEFAULT_OMP_REQUIRES=
JAVA_HOME=/usr/lib64/jvm/java
SLURM_SRUN_COMM_HOST=10.128.6.142
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_skylake=8.6
PE_LIBSCI_MODULE_NAME=cray-libsci/18.07.1
SLURM_CPU_BIND_TYPE=mask-cpu:
PE_INTEL_FIXED_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.7.3/gni/mpich-intel/16.0/lib/pkgconfig
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61

```

```

PE_MPICH_NV_LIBS_nvidia20=-lcudart
PYTHONSTARTUP=/etc/pythonstart
PE_MPICH_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.
↳ 7.3/gni/mpich-@PRGENV@PE_MPICH_DIR_DEFAULT64@/@
↳ PE_MPICH_GENCOMPS@/lib/pkgconfig
MODULEPATH=/global/homes/t/USER/modules:/opt/cray/pe
↳ /craype/2.5.15/modulefiles:/opt/cray/pe/modulefi
↳ les:/opt/cray/modulefiles:/opt/modulefiles:/usr/
↳ common/software/modulefiles:/usr/syscom/nsg/modu
↳ lefiles:/usr/common/das/modulefiles:/usr/common/
↳ ftg/modulefiles:/usr/common/graphics/modulefiles
↳ :/usr/common/jgi/modulefiles:/usr/common/tig/mod
↳ ulefiles:/opt/cray/craype/default/modulefiles:/o
↳ pt/cray/ari/modulefiles
PE_LIBSCI_GENCOMPILERS_CRAY_x86_64=8.6
SLURM_UMASK=0022
SLURM_PTY_WIN_COL=158
TZ=US/Pacific
SDK_HOME=/usr/lib64/jvm/java
CRAYPE_USE_BUILDTOOLS=0
NSG_HOME=/usr/syscom/nsg
SHMEM_ABORT_ON_ERROR=1
MAN_POSIXLY_CORRECT=1
LOADEDMODULES=modules/3.2.10.6:nsg/1.2.0:intel/18.0.
↳ 1.163:craype-network-aries:craype/2.5.15:cray-li
↳ bsci/18.07.1:udreg/2.3.2-6.0.7.0_33.18__g5196236
↳ .ari:ugni/6.0.14.0-6.0.7.0_23.1__gea11d3d.ari:pm
↳ i/5.0.14:dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari:
↳ gni-headers/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari:
↳ xpmem/2.2.15-6.0.7.1_5.10__g7549d06.ari:job/2.2.
↳ 3-6.0.7.0_44.1__g6c4e934.ari:dvs/2.7.2.2.116-6.0
↳ .7.1_8.11__g0dac9eb:alps/6.6.43-6.0.7.0_26.4__ga
↳ 796da3.ari:rca/2.2.18-6.0.7.0_33.3__g2aa4f39.ari
↳ :atp/2.1.3:PrgEnv-intel/6.0.4:craype-ivybridge:c
↳ ray-mpich/7.7.3:altd/2.0:Base-opts/2.4.135-6.0.7
↳ .0_38.1__g718f891.ari
SLURM_JOB_UID=73766
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_PKG_CONFIG_PATH=/opt/cray/pe/cti/1.0.7/lib/pkgcon
↳ fig:/opt/cray/pe/cti/1.0.6/lib/pkgconfig
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53
↳ 49
PE_FFTW_DEFAULT_TARGET_ivybridge=ivybridge
CRAY_DMAPP_POST_LINK_OPTS=-L/opt/cray/dmapp/7.1.1-6.
↳ 0.7.0_34.3__g5a674e0.ari/lib64
PE_FFTW_DEFAULT_TARGET_x86_skylake=x86_skylake
PE_FFTW_DEFAULT_TARGET_share=share
SLURM_NODEID=0
SLURM_SUBMIT_DIR=/global/u1/t/USER
PE_LIBSCI_OMP_REQUIRES_openmp=_mp
CRAY_RCA_INCLUDE_OPTS=-I/opt/cray/rca/2.2.18-6.0.7.0
↳ _33.3__g2aa4f39.ari/include
↳ -I/opt/cray/krca/2.2.4-6.0.7.1_5.40__g8505b97.ar
↳ i/include
↳ -I/opt/cray-hss-devel/8.0.0/include
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6

```

```

PE_PETSC_DEFAULT_GENCOMPILERS_GNU_skylake=6.1
SLURM_STEP_RESV_PORTS=63823
PE_MPICH_CXX_PKGCONFIG_LIBS=mpichcxx
CRAY_MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-inte
↳ 1/16.0
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_mic_knl=86
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
SLURM_TASK_PID=29377
PE_MPICH_PKGCONFIG_VARIABLES=PE_MPICH_NV_LIBS_@accel
↳ erator@:PE_MPICH_ALTERNATE_LIBS_@multithreaded@:
↳ PE_MPICH_ALTERNATE_LIBS_@dpm@
PE_LIBSCI_DEFAULT_GENCOMPS_INTEL_x86_64=160
SLURM_CPUS_ON_NODE=48
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
CRAY_MPICH2_VER=7.7.3
CRAY_PMI_POST_LINK_OPTS=-L/opt/cray/pe/pmi/5.0.14/li
↳ b64
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16.0
PE_MPICH_PKGCONFIG_LIBS=mpich
PE_HDF5_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_TPSSL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ /tpsl/18.06.1/@PRGENV@/@PE_TPSSL_DEFAULT_GENCOMPS
↳ @/@PE_TPSSL_DEFAULT_TARGET@/lib/pkgconfig
SLURM_PROCID=0
SLURM_JOB_NODELIST=nid00188
PE_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
GPG_TTY=/dev/pts/0
PE_PARALLEL_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_MPICH_ALTERNATE_LIBS_multithreaded=_mt
PE_GA_DEFAULT_GENCOMPILERS_GNU=5.3 4.9
PE_LIBSCI_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/libsc
↳ i/18.07.1/@PRGENV@/@PE_LIBSCI_GENCOMPS@/@PE_LIBS
↳ CI_TARGET@/lib/pkgconfig
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
QT_SYSTEM_DIR=/usr/share/desktop-data
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
PE_MPICH_TARGET_VAR_nvidia35=-lcudart
SHLVL=4
HOME=/global/homes/t/USER
PE_PKGCONFIG_PRODUCTS_DEFAULT=PE_PAPI
JDK_HOME=/usr/lib64/jvm/java
CRAY_LIBSCI_VERSION=18.07.1
SLURM_PTY_PORT=62702
OSTYPE=linux
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
LESS_ADVANCED_PREPROCESSOR=no
SLURM_LOCALID=0
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_MPICH_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/p
↳ e/mpt/7.7.3/gni/mpich-@PRGENV@PE_MPICH_DEFAULT_
↳ DIR_DEFAULT64@/@PE_MPICH_DEFAULT_GENCOMPS@/lib/p
↳ kgconfig
ALTD_PATH=/usr/common/software/altd/2.0
PE_TPSSL_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU GNU64
↳ INTEL_INTEL64

```

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

```

GREP_OPTIONS=--color=auto
SLURM_CLUSTER_NAME=edison
SLURM_JOB_CPUS_PER_NODE=48
PE_TPSL_DEFAULT_GENCOMPS_INTEL_sandybridge=160
CRAY_PMI_INCLUDE_OPTS=-I/opt/cray/pe/pmi/5.0.14/include
↳ ude
XCURSOR_THEME=DMZ
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_interlagos=86
LS_OPTIONS=-N --color=none -T 0
SLURM_JOB_GID=73766
SLURM_SUBMIT_HOST=mom2
WINDOWMANAGER=
PE_MPICH_DIR_CRAY_DEFAULT64=64
PRGENVMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:PrgEnv-ppc
↳ gEnv-pathscales:PrgEnv-pgi
ATP_MRNET_COMM_PATH=/opt/cray/pe/atp/2.1.3/libexec/atp
↳ tp_mrnet_commnodewrapper
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PKG_CONFIG_PATH_DEFAULT=/opt/cray/pe/papi/5.6.0.3/lib64
↳ b64/pkgconfig
MPICH_ALLOC_MEM_PG_SZ=2M
CRAYPE_NETWORK_TARGET=aries
SLURM_GTIDS=0
SLURM_JOB_PARTITION=debug
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
BASH_ENV=/global/homes/t/USER/.bashrc
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2 7.1
↳ 5.3 4.9
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_haswell=7.1 5.3 4.9
PE_NETCDF_DEFAULT_REQUIRED_PRODUCTS=PE_HDF5
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16
↳ .0
CRAY_GNI_HEADERS_INCLUDE_OPTS=-I/opt/cray/gni-headers
↳ s/5.0.12.0-6.0.7.0.24.1__g3b1768f.ari/include
CRAY_LIBSCI_PREFIX_DIR=/opt/cray/pe/libsci/18.07.1/INTEL
↳ 16.0/x86_64
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1
↳ 5.3 4.9
PE_MPICH_NV_LIBS=
MACHTYPE=x86_64-suse-linux
ALTD_SELECT_USERS=
LOGNAME=USER
PE_TPSL_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49
LESS=-M -I -R
PYTHONPATH=/opt/ovis/lib/python2.7/site-packages:/project
↳ /project/projectdirs/m888/USER/aries/processing_scripts
↳ /project/projectdirs/m888/USER/aries/notebooks
↳ /project/projectdirs/m888/USER/aries/process
↳ ing_scripts:/project/projectdirs/m888/USER/aries
↳ /notebooks:/project/projectdirs/m888/USER/aries/
↳ processing_scripts:/project/projectdirs/m888/USER
↳ R/aries/notebooks
PE_TRILINOS_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_HDF5
↳ F5_PARALLEL:PE_NETCDF:HDF5PARALLEL:PE_LIBSCI:PE_
↳ TPSL

```

```

PE_HDF5_DEFAULT_GENCOMPS_GNU=
PE_MPICH_GENCOMPS_CRAY=86
DMAPP_ABORT_ON_ERROR=1
PE_MPICH_DEFAULT_GENCOMPILERS_CRAY=8.6
CVS_RSH=ssh
PE_LIBSCI_OMP_REQUIRES=
PE_TRILINOS_DEFAULT_GENCOMPS_GNU_x86_64=82 73 51 49
SLURM_STEP_NUM_TASKS=1
SLURM_JOB_ACCOUNT=nstaff
NETWORKS=/project/projectdirs/m3229/USER
PE_LIBSCI_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
TOOLMODULES=apprentice:apprentice2:atp:chapel:cray-l
↳ gdb:cray-snp-launcher:craypat:craypkg-gen:ddt:gdb
↳ :iobuf:papi:perftools:perftools-lite:stat:totalview
↳ iew:xt:craypat:xt-lgdb:xt-papi:xt-totalview
SSH_CONNECTION=128.55.209.26 49906 128.55.209.26 22
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
XDG_DATA_DIRS=/usr/share
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
DVS_INCLUDE_OPTS=-I/opt/cray/dvs/2.7.2.2.116-6.0.7.1
↳ .8.11__g0dac9eb/include
PE_MPICH_DEFAULT_GENCOMPS_GNU=71 51 49
PE_MPICH_DEFAULT_FIXED_PRGENV=INTEL
SLURM_JOB_NUM_NODES=1
CRAY_PRGENVINTEL=loaded
PE_LIBSCI_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/
↳ pe/libsci/18.07.1/@PRGENV@/@PE_LIBSCI_DEFAULT_GEN
↳ COMPS@/@PE_LIBSCI_DEFAULT_TARGET@/lib/pkgconfig
MODULESHOME=/opt/cray/pe/modules/3.2.10.6
PE_FFTW2_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2 7.1
↳ 5.3 4.9
PE_GA_DEFAULT_FIXED_PRGENV=CRAY INTEL
PKG_CONFIG_PATH=/opt/cray/rca/2.2.18-6.0.7.0.33.3__g
↳ 2aa4f39.ari/lib64/pkgconfig:/opt/cray/alps/6.6.4
↳ 3-6.0.7.0.26.4__ga796da3.ari/lib64/pkgconfig:/opt
↳ t/cray/xpmmem/2.2.15-6.0.7.1.5.10__g7549d06.ari/lib
↳ ib64/pkgconfig:/opt/cray/gni-headers/5.0.12.0-6.0
↳ 7.0.24.1__g3b1768f.ari/lib64/pkgconfig:/opt/cr
↳ ay/dmapp/7.1.1-6.0.7.0.34.3__g5a674e0.ari/lib64/
↳ pkgconfig:/opt/cray/pe/pmi/5.0.14/lib64/pkgconfi
↳ g:/opt/cray/ugni/6.0.14.0-6.0.7.0.23.1__gea11d3d
↳ .ari/lib64/pkgconfig:/opt/cray/udreg/2.3.2-6.0.7
↳ .0.33.18__g5196236.ari/lib64/pkgconfig:/opt/cray
↳ /pe/craype/2.5.15/pkg-config:/opt/cray/pe/iobuf/
↳ 2.0.8/lib/pkgconfig:/opt/cray/pe/fftw/2.1.5.9/li
↳ b/pkgconfig:/opt/cray/pe/atp/2.1.3/lib/pkgconfig
PE_MPICH_NV_LIBS_nvidia35=-lcudart
LESSOPEN=lessopen.sh %s
PELOCAL_PRGENV=true
PE_PETSC_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/p
↳ e/petsc/3.8.4.0/complex/@PRGENV@/@PE_PETSC_DEFAU
↳ LT_GENCOMPS@/@PE_PETSC_DEFAULT_TARGET@/lib/pkgco
↳ nfig
LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_x86_64=160

```



```

SLURM_STEP_TASKS_PER_NODE=1
CRAY_NUM_COOKIES=2
PE_LIBSCI_DEFAULT_PKGCONFIG_VARIABLES=PE_LIBSCI_DEFA
↳ ULT_OMP_REQUIRES=@openmp@:PE_SCI_EXT_LIBPATH:PE_
↳ SCI_EXT_LIBNAME
PE_MPICH_NV_LIBS_nvidia60=-lcudart
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53
↳ 49
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_mic_knl=160
LIBSCI_VERSION=18.07.1
SLURM_STEP_NODELIST=nid00188
CRAY_COOKIES=4014866432,4014931968
XDG_RUNTIME_DIR=/run/user/73766
PE_LIBSCI_GENCOMPILERS_INTEL_x86_64=16.0
CRAY_PRE_COMPILE_OPTS=-hnetwork=aries
CRAY_ALPS_INCLUDE_OPTS=-I/opt/cray/alps/6.6.43-6.0.7
↳ .0.26.4__ga796da3.ari/include
CRAY_CPU_TARGET=sandybridge
PE_FFTW_DEFAULT_TARGET_broadwell=broadwell
NERSC_HOST=edison
PE_TPSSL_64_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1
↳ 5.3 4.9
MPICH_GNI_ROUTING_MODE=ADAPTIVE_1
CRAY_UGNI_INCLUDE_OPTS=-I/opt/cray/ugni/6.0.14.0-6.0
↳ .7.0.23.1__gea11d3d.ari/include
PE_PAPI_DEFAULT_ACCELL_FAMILY_LIBS=
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_LIBSCI_REQUIRED_PRODUCTS=PE_MPICH
CRAY_XPMEM_INCLUDE_OPTS=-I/opt/cray/xpmem/2.2.15-6.0
↳ .7.1.5.10__g7549d06.ari/include
craype_already_loaded=0
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_haswell=160
PE_LIBSCI_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1
↳ 4.9
PE_LIBSCI_GENCOMPS_GNU_x86_64=71 61 51 49
SLURM_CPU_BIND=quiet,mask-cpu:0xFFFFFFFFFFFF
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_FFTW_DEFAULT_TARGET_x86_64=x86_64
LESSCLOSE=lessclose.sh %s %s
LDMSLOGS=/global/cscratch1/ovis/csv
ATP_HOME=/opt/cray/pe/atp/2.1.3
SLURM_MEM_PER_NODE=57344
SCRATCH=/scratch2/scratchdirs/USER
PE_TPSSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16
↳ .0
ALTD_WORKDIR=/scratch1/scratchdirs/altdlogs
PE_FFTW_DEFAULT_TARGET_haswell=haswell
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_GA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/g
↳ a/5.3.0.8@PRGENV@/PE_GA_DEFAULT_GENCOMPS@lib/
↳ pkgconfig
G_BROKEN_FILENAMES=1
PE_GA_DEFAULT_GENCOMPS_GNU=53 49
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_interlagos=7.1 5.3
↳ 4.9
PE_PAPI_DEFAULT_ACCEL_LIBS=

```

```

CRAY_LD_LIBRARY_PATH=/opt/cray/pe/mpt/7.7.3/gni/mpic
↳ h-intel/16.0/lib:/opt/cray/rca/2.2.18-6.0.7.0_33
↳ .3__g2aa4f39.ari/lib64:/opt/cray/alps/6.6.43-6.0
↳ .7.0.26.4__ga796da3.ari/lib64:/opt/cray/xpmem/2.
↳ 2.15-6.0.7.1.5.10__g7549d06.ari/lib64:/opt/cray/
↳ dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari/lib64:/op
↳ t/cray/pe/pmi/5.0.14/lib64:/opt/cray/ugni/6.0.14
↳ .0-6.0.7.0.23.1__gea11d3d.ari/lib64:/opt/cray/ud
↳ reg/2.3.2-6.0.7.0_33.18__g5196236.ari/lib64:/opt
↳ /cray/pe/libsci/18.07.1/INTEL/16.0/x86_64/lib
PE_SMA_DEFAULT_DIR_PGI_DEFAULT64=64
PE_INTEL_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/p
↳ arallel-netcdf/1.8.1.3/INTEL/16.0/lib/pkgconfig:
↳ /opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/INTEL/1
↳ 6.0/lib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/IN
↳ TEL/16.0/lib/pkgconfig:/opt/cray/pe/mpt/7.7.3/gn
↳ i/mpich-intel/16.0/lib/pkgconfig:/opt/cray/pe/hd
↳ f5-parallel/1.10.2.0/INTEL/16.0/lib/pkgconfig:/o
↳ pt/cray/pe/hdf5/1.10.2.0/INTEL/16.0/lib/pkgconfi
↳ g:/opt/cray/pe/ga/5.3.0.8/INTEL/18.0/lib/pkgconf
↳ ig
COLORTERM=1
PE_PETSC_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_PETSC_DEFAULT_GENCOMPS_GNU_x86_64=71 53 49
intel_already_loaded=0
PE_PETSC_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_MPICH_DEFAULT_DIR_CRAY_DEFAULT64=64
JAVA_ROOT=/usr/lib64/jvm/java
BASH_FUNC_module%=() { eval
↳ ` /opt/cray/pe/modules/3.2.10.6/bin/modulecmd
↳ bash $* `
}
LSB Version:          n/a
Distributor ID:       SUSE
Description:          SUSE Linux Enterprise Server 12
↳ SP3
Release:              12.3
Codename:              n/a
Linux nid00188 4.4.103-6.38.4.0.153-cray_ari_c #1 SMP
↳ Thu Nov 1 16:05:05 UTC 2018 (6ef8fef) x86_64
↳ x86_64 x86_64 GNU/Linux
Architecture:         x86_64
CPU op-mode(s):       32-bit, 64-bit
Byte Order:           Little Endian
CPU(s):               48
On-line CPU(s) list:  0-47
Thread(s) per core:   2
Core(s) per socket:   12
Socket(s):            2
NUMA node(s):         2
Vendor ID:            GenuineIntel
CPU family:           6
Model:                62
Model name:           Intel(R) Xeon(R) CPU E5-2695 v2
↳ @ 2.40GHz
Stepping:             4

```

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

CPU MHz:	2401.000	HugePages_Total:	0
CPU max MHz:	2401.0000	HugePages_Free:	0
CPU min MHz:	1200.0000	HugePages_Rsvd:	0
BogoMIPS:	4799.63	HugePages_Surp:	0
Virtualization:	VT-x	Hugepagesize:	2048 kB
L1d cache:	32K	DirectMap4k:	12028 kB
L1i cache:	32K	DirectMap2M:	2048000 kB
L2 cache:	256K	DirectMap1G:	67108864 kB
L3 cache:	30720K	NAME	MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
NUMA node0 CPU(s):	0-11,24-35	loop0	7:0 0 128K 0 loop /var/opt/cray/imps-d
NUMA node1 CPU(s):	12-23,36-47	↪	istribution/squash/mounts/global
Flags:	fpu vme de pse tsc msr pae mce	loop1	7:1 0 1.1M 0 loop
↪ cx8 apic sep mtrr pge mca cmov pat pse36 clflush		↪	/var/opt/cray/imps-distribution/squash/mounts/cl
↪ dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx		↪	e.edison.cle6.0.up07.20190226110832
↪ pdpe1gb rdtscp lm constant_tsc arch_perfmon pebs		loop2	7:2 0 2.4G 1 loop /.rootfs_lower_ro
↪ bts rep_good nopl xtopology nonstop_tsc aperfmperf		loop3	7:3 0 23.8G 1 loop
↪ eagerfpu pni pclmulqdq dtes64 monitor ds_cpl vmx		↪	/var/opt/cray/imps-image-binding/PE_x86_64/squas
↪ smx est tm2 sse3 cx16 xtpr pdcm pcid dca sse4_1		↪	h_mounts/squashfs_1vPMsu_mount_point
↪ sse4_2 x2apic popcnt tsc_deadline_timer aes xsave		loop4	7:4 0 1 loop
↪ avx f16c rdrand lahf_lm ida arat epb pln pts		loop5	7:5 0 1 loop
↪ dtherm spec_ctrl kaiser tpr_shadow vnmi		loop6	7:6 0 0 loop
↪ flexpriority ept vpid fsgsbase smep erms xsaveopt		loop7	7:7 0 0 loop
MemTotal:	65938820 kB	loop8	7:8 0 0 loop
MemFree:	62480052 kB	loop9	7:9 0 0 loop
MemAvailable:	62204336 kB	loop10	7:10 0 0 loop
Buffers:	3368 kB	loop11	7:11 0 0 loop
Cached:	251848 kB	loop12	7:12 0 0 loop
SwapCached:	0 kB	loop13	7:13 0 0 loop
Active:	190728 kB	loop14	7:14 0 0 loop
Inactive:	214056 kB	loop15	7:15 0 0 loop
Active(anon):	178628 kB	loop16	7:16 0 0 loop
Inactive(anon):	198888 kB	loop17	7:17 0 0 loop
Active(file):	12100 kB	loop18	7:18 0 0 loop
Inactive(file):	15168 kB	loop19	7:19 0 0 loop
Unevictable:	2576 kB	loop20	7:20 0 0 loop
Mlocked:	2576 kB	loop21	7:21 0 0 loop
SwapTotal:	0 kB	loop22	7:22 0 0 loop
SwapFree:	0 kB	loop23	7:23 0 0 loop
Dirty:	0 kB	loop24	7:24 0 0 loop
Writeback:	0 kB	loop25	7:25 0 0 loop
AnonPages:	152384 kB	loop26	7:26 0 0 loop
Mapped:	57320 kB	loop27	7:27 0 0 loop
Shmem:	227720 kB	loop28	7:28 0 0 loop
Slab:	381332 kB	loop29	7:29 0 0 loop
SReclaimable:	22648 kB	loop30	7:30 0 0 loop
SUnreclaim:	358684 kB	loop31	7:31 0 0 loop
KernelStack:	14080 kB	loop32	7:32 0 0 loop
PageTables:	4824 kB	loop33	7:33 0 0 loop
NFS_Unstable:	0 kB	loop34	7:34 0 0 loop
Bounce:	0 kB	loop35	7:35 0 0 loop
WritebackTmp:	0 kB	loop36	7:36 0 0 loop
CommitLimit:	32969408 kB	loop37	7:37 0 0 loop
Committed_AS:	571856 kB	loop38	7:38 0 0 loop
VmallocTotal:	34359738367 kB	loop39	7:39 0 0 loop
VmallocUsed:	0 kB	loop40	7:40 0 0 loop
VmallocChunk:	0 kB	loop41	7:41 0 0 loop
HardwareCorrupted:	0 kB		

```

loop42  7:42  0      0 loop
loop43  7:43  0      0 loop
loop44  7:44  0      0 loop
loop45  7:45  0      0 loop
loop46  7:46  0      0 loop
loop47  7:47  0      0 loop
loop48  7:48  0      0 loop
loop49  7:49  0      0 loop
loop50  7:50  0      0 loop
loop51  7:51  0      0 loop
loop52  7:52  0      0 loop
loop53  7:53  0      0 loop
loop54  7:54  0      0 loop
loop55  7:55  0      0 loop
loop56  7:56  0      0 loop
loop57  7:57  0      0 loop
loop58  7:58  0      0 loop
loop59  7:59  0      0 loop
loop60  7:60  0      0 loop
loop61  7:61  0      0 loop
loop62  7:62  0      0 loop
loop63  7:63  0      0 loop
loop64  7:64  0      0 loop
loop65  7:65  0      0 loop
loop66  7:66  0      0 loop
loop67  7:67  0      0 loop
loop68  7:68  0      0 loop
loop69  7:69  0      0 loop
loop70  7:70  0      0 loop
loop71  7:71  0      0 loop
loop72  7:72  0      0 loop
loop73  7:73  0      0 loop
loop74  7:74  0      0 loop
loop75  7:75  0      0 loop
loop76  7:76  0      0 loop
loop77  7:77  0      0 loop
loop78  7:78  0      0 loop
loop79  7:79  0      0 loop
loop80  7:80  0      0 loop
loop81  7:81  0      0 loop
loop82  7:82  0      0 loop
loop83  7:83  0      0 loop
loop84  7:84  0      0 loop
loop85  7:85  0      0 loop
loop86  7:86  0      0 loop
loop87  7:87  0      0 loop
loop88  7:88  0      0 loop
loop89  7:89  0      0 loop
loop90  7:90  0      0 loop
loop91  7:91  0      0 loop
loop92  7:92  0      0 loop
loop93  7:93  0      0 loop
loop94  7:94  0      0 loop
loop95  7:95  0      0 loop
loop96  7:96  0      0 loop
loop97  7:97  0      0 loop
loop98  7:98  0      0 loop

```

```

loop99  7:99  0      0 loop
loop100 7:100 0      0 loop
loop101 7:101 0      0 loop
loop102 7:102 0      0 loop
loop103 7:103 0      0 loop
loop104 7:104 0      0 loop
loop105 7:105 0      0 loop
loop106 7:106 0      0 loop
loop107 7:107 0      0 loop
loop108 7:108 0      0 loop
loop109 7:109 0      0 loop
loop110 7:110 0      0 loop
loop111 7:111 0      0 loop
loop112 7:112 0      0 loop
loop113 7:113 0      0 loop
loop114 7:114 0      0 loop
loop115 7:115 0      0 loop
loop116 7:116 0      0 loop
loop117 7:117 0      0 loop
loop118 7:118 0      0 loop
loop119 7:119 0      0 loop
loop120 7:120 0      0 loop
loop121 7:121 0      0 loop
loop122 7:122 0      0 loop
loop123 7:123 0      0 loop
loop124 7:124 0      0 loop
loop125 7:125 0      0 loop
loop126 7:126 0      0 loop
loop127 7:127 0      0 loop

```

NVIDIA-SMI has failed because it couldn't communicate
↪ with the NVIDIA driver. Make sure that the latest
↪ NVIDIA driver is installed and running.

Theta:

LSB Version: n/a

Distributor ID: SUSE

Description: SUSE Linux Enterprise Server 12 SP3

Release: 12.3

Codename: n/a

Linux thetalogin6 4.4.162-94.72-default #1 SMP Mon

↪ Nov 12 18:57:45 UTC 2018 (9de753f) x86_64 x86_64

↪ x86_64 GNU/Linux

Architecture: x86_64

CPU op-mode(s): 32-bit, 64-bit

Byte Order: Little Endian

CPU(s): 32

On-line CPU(s) list: 0-31

Thread(s) per core: 1

Core(s) per socket: 16

Socket(s): 2

NUMA node(s): 2

Vendor ID: GenuineIntel

CPU family: 6

Model: 63

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

```

Model name:          Intel(R) Xeon(R) CPU E5-2698 v3
↳ @ 2.30GHz
Stepping:            2
CPU MHz:              2823.588
CPU max MHz:          3600.0000
CPU min MHz:          1200.0000
BogoMIPS:             4600.18
Virtualization:       VT-x
L1d cache:            32K
L1i cache:            32K
L2 cache:              256K
L3 cache:              40960K
NUMA node0 CPU(s):
↳ 0,2,4,6,8,10,12,14,16,18,20,22,24,26,28,30
NUMA node1 CPU(s):
↳ 1,3,5,7,9,11,13,15,17,19,21,23,25,27,29,31
Flags:                fpu vme de pse tsc msr pae mce
↳ cx8 apic sep mtrr pge mca cmov pat pse36 clflush
↳ dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx
↳ pdpe1gb rdtscp lm constant_tsc arch_perfmon pebs
↳ bts rep_good nopl xtopology nonstop_tsc
↳ aperfmperf eagerfpu pni pclmulqdq dtes64 monitor
↳ ds_cpl vmx smx est tm2 ssse3 sdbg fma cx16 xtpr
↳ pdcm pcid dca sse4_1 sse4_2 x2apic movbe popcnt
↳ tsc_deadline_timer aes xsave avx f16c rdrand
↳ lahf_lm abm ida arat epb invpcid_single pln pts
↳ dtherm kaiser tpr_shadow vnmi flexpriority ept
↳ vpid fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms
↳ invpcid cqm xsaveopt cqm_llc cqm_occup_llc
MemTotal:             263704472 kB
MemFree:               224069692 kB
MemAvailable:          245442064 kB
Buffers:               1290208 kB
Cached:                21929808 kB
SwapCached:            0 kB
Active:                9890932 kB
Inactive:              14074220 kB
Active(anon):          1393268 kB
Inactive(anon):         88880 kB
Active(file):           8497664 kB
Inactive(file):         13985340 kB
Unevictable:           8388688 kB
Mlocked:               8388688 kB
SwapTotal:             134217724 kB
SwapFree:              134217724 kB
Dirty:                 23700 kB
Writeback:              0 kB
AnonPages:             9134600 kB
Mapped:                719096 kB
Shmem:                 736708 kB
Slab:                  4942124 kB
SReclaimable:          738060 kB
SUnreclaim:            4204064 kB
KernelStack:           20064 kB
PageTables:            88544 kB
NFS_Unstable:          0 kB
Bounce:                0 kB
WritebackTmp:          0 kB
CommitLimit:           266069960 kB
Committed_AS:          14294840 kB
VmallocTotal:          34359738367 kB
VmallocUsed:            0 kB
VmallocChunk:           0 kB
HardwareCorrupted:     0 kB
AnonHugePages:          8390656 kB
HugePages_Total:        0
HugePages_Free:         0
HugePages_Rsvd:         0
HugePages_Surp:         0
Hugepagesize:           2048 kB
DirectMap4k:            1129172 kB
DirectMap2M:            151736320 kB
DirectMap1G:            117440512 kB
LESSKEY=/etc/lesskey.bin
MODULE_VERSION_STACK=3.2.11.1
KSH_AUTOLOAD=1
MKLR00T=/opt/intel/compilers_and_libraries_2018.0.12
↳ 8/linux/mkl
PE_LIBSCI_VOLATILE_PRGENV=CRAY GNU INTEL
PE_SMA_DEFAULT_PKGCONFIG_VARIABLES=PE_SMA_COMPFLAG_@
↳ prgenv@
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_mic_knl=160
MANPATH=/opt/cray/pe/mpt/7.7.3/gni/man/mpich:/opt/cr
↳ ay/pe/perftools/7.0.4/man:/opt/cray/pe/papi/5.6.
↳ 0.4/share/pdoc/man:/opt/cray/pe/atp/2.1.3/man:/o
↳ pt/cray/alps/6.6.1-6.0.6.1_4.1__ga6396bb.ari/man
↳ :/opt/cray/job/2.2.3-6.0.6.0_9.47__g6c4e934.ari/
↳ man:/opt/cray/pe/pmi/5.0.14/man:/opt/cray/pe/lib
↳ sci/18.07.1/man:/opt/cray/pe/man/csmlversion:/op
↳ t/cray/pe/craype/2.5.15/man:/opt/intel/man/commo
↳ n:/opt/intel/compilers_and_libraries_2018.0.128/
↳ linux/mpi/man:/opt/intel/documentation_2017/en/d
↳ ebugger//gdb-ia/man:/opt/intel/documentation_20
↳ 17/en/debugger//gdb-mic/man:/opt/intel/document
↳ ation_2017/en/debugger//gdb-igfx/man:/opt/cray/
↳ pe/modules/3.2.11.1/share/man:/usr/local/man:/us
↳ r/share/man:/usr/man:/opt/cray/share/man:/opt/cr
↳ ay/pe/man
DB2INSTANCE=db2cat
NNTPSERVER=news
PE_PAPI_DEFAULT_ACCEL_FAMILY_LIBS_nvidia=-lcupti,-l
↳ cudart,-lcuda
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
PE_PETSC_DEFAULT_GENCOMPS_CRAY_skylake=86
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_CXX_PKGCONFIG_LIBS=mpichcxx
XDG_SESSION_ID=564
HOSTNAME=thetalogin6
IBM_DB_LIB=/dbhome/db2cat/sql/lib
CRAY_UDREG_INCLUDE_OPTS=-I/opt/cray/udreg/2.3.2-6.0.
↳ 6.0_15.18__g5196236.ari/include
PE_FFTW_DEFAULT_TARGET_mic_knl=mic_knl

```

```

PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_TRILINOS_DEFAULT_GENCOMPS_CRAY_x86_64=87
XKEYSYMDB=/usr/X11R6/lib/X11/XKeysymDB
CRAY_SITE_LIST_DIR=/etc/opt/cray/pe/modules
LIBRARYMODULES=acml:alps:cray-dwarf:cray-fftw:cray-g
↳ a:cray-hdf5:cray-hdf5-parallel:cray-libsci:cray-
↳ libsci_acc:cray-mpich:cray-mpich2:cray-mpich-abi
↳ :cray-netcdf:cray-netcdf-hdf5parallel:cray-paral
↳ lel-netcdf:cray-petsc:cray-petsc-complex:cray-sh
↳ mem:cray-tpsl:cray-trilinos:cudatoolkit:fftw:ga:
↳ hdf5:hdf5-parallel:iobuf:libfast:netcdf:netcdf-h
↳ df5parallel:ntk:onesided:papi:petsc:petsc-comple
↳ x:pmi:tpsl:trilinos:xt-libsci:xt-mpich2:xt-mpt:x
↳ t-papi
PE_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/
↳ pe/netcdf/4.6.1.3/@PRGENV@/PE_NETCDF_DEFAULT_GE
↳ NCOMPS@/lib/pkgconfig
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/
↳ opt/cray/pe/parallel-netcdf/1.8.1.3/@PRGENV@/PE
↳ _PARALLEL_NETCDF_DEFAULT_GENCOMPS@/lib/pkgconfig
PE_SMA_DEFAULT_COMPFLAG_GNU=-fcray-pointer
PE_TRILINOS_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cra
↳ y/pe/trilinos/12.12.1.1/@PRGENV@/PE_TRILINOS_DE
↳ FAULT_GENCOMPS@/PE_TRILINOS_DEFAULT_TARGET@/lib
↳ /pkgconfig
INTEL_LICENSE_FILE=/opt/intel/compilers_and_librarye
↳ s_2018.0.128/linux/licenses:/opt/intel/licenses:
↳ /intel/licenses
IPPROOT=/opt/intel/compilers_and_libraries_2018.0.12
↳ 8/linux/ipp
PE_ENV=INTEL
PE_HDF5_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
PE_MPICH2_ALTERNATE_LIBS_dpm=-dpm
PE_SMA_DEFAULT_COMPFLAG=
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
SHELL=/bin/bash
TERM=xterm-256color
HOST=thetalingo6
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
PKGCONFIG_ENABLED=1
HISTSIZE=1000
PROFILEREAD=true
INTEL_MINOR_VERSION=0.128
PE_PETSC_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2 7.1
↳ 6.1
SSH_CLIENT=130.202.141.135 56759 22
CRAYPE_DIR=/opt/cray/pe/craype/2.5.15
CRAY_UGNI_POST_LINK_OPTS=-L/opt/cray/ugni/6.0.14-6.0
↳ .6.0.18.12_g777707d.ari/lib64
CRAY_XPMEM_POST_LINK_OPTS=-L/opt/cray/xpmem/2.2.14-6
↳ .0.6.1.5.14_g34333c9.ari/lib64
GDBSERVER_MIC=/opt/intel/debugger_2017/gdb/targets/m
↳ ic/bin/gdbserver
PE_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU

```

```

PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
PE_PETSC_DEFAULT_GENCOMPS_GNU_haswell=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_INTEL_haswell=160
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_TPSSL_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53 49
PE_TPSSL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSCI
PE_TRILINOS_DEFAULT_VOLATILE_PRGENV=CRAY GNU INTEL
LIBRARY_PATH=/opt/intel/compilers_and_libraries_2018
↳ .0.128/linux/ipp/lib/intel64:/opt/intel/compiler
↳ s_and_libraries_2018.0.128/linux/compiler/lib/in
↳ tel64:/opt/intel/compilers_and_libraries_2018.0.
↳ 128/linux/mkl/lib/intel64:/opt/intel/compilers_a
↳ nd_libraries_2018.0.128/linux/tbb/lib/intel64/gc
↳ c4.4:/opt/intel/compilers_and_libraries_2018.0.1
↳ 28/linux/daal/lib/intel64_lin:/opt/intel/compile
↳ rs_and_libraries_2018.0.128/linux/daal/./compil
↳ er/lib/intel64_lin
PE_FFTW_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ /fftw/3.3.8.1/@PE_FFTW_DEFAULT_TARGET@/lib/pkgco
↳ nfig
PE_HDF5_DEFAULT_VOLATILE_PRGENV=GNU
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/op
↳ t/cray/pe/hdf5-parallel/1.10.2.0/@PRGENV@/PE_HD
↳ F5_PARALLEL_DEFAULT_GENCOMPS@/lib/pkgconfig
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PA
↳ TH=/opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/@PRG
↳ ENV@/PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS@/l
↳ ib/pkgconfig
PE_PETSC_DEFAULT_GENCOMPS_CRAY_interlagos=86
CRAY_MPICH2_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-int
↳ el/16.0
INTEL_PATH=/opt/intel/compilers_and_libraries_2018.0
↳ .128
PE_GA_DEFAULT_VOLATILE_PRGENV=GNU
PE_LIBSCI_DEFAULT_GENCOMPS_GNU_x86_64=71 61 51 49
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_mic_knl=86
LD_PRELOAD=/soft/buildtools/trackdeps/$LIB/trackdeps
↳ .so
MORE=--sl
IBM_DB_DIR=/dbhome/db2cat/sqllib
FPATH=/opt/cray/pe/modules/3.2.11.1/init/sh_funcs/n
↳ o_redirect
PERFTOOLS_VERSION=7.0.4
PE_MPICH_DEFAULT_GENCOMPILERS_GNU=7.1 5.1 4.9
PE_PKGCONFIG_PRODUCTS=PE_MPICH:PE_LIBSCI
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_MPICH_GENCOMPS_GNU=71 51 49
DARSHAN_PRELOAD=/soft/perftools/darshan/darshan-3.1.
↳ 5/lib/libdarshan.so
PE_PAPI_DEFAULT_ACCEL_LIBS_nvidia35=-lcupti,-lcudar
↳ t,-lcuda
PE_PETSC_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSC
↳ I:PE_HDF5_PARALLEL:PE_TPSSL
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_haswell=86

```

```

PE_TPSL_64_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray
↳ /pe/tpsl/18.06.1/@PRGENV@64/@PE_TPSL_64_DEFAULT_
↳ GENCOMPS@/@PE_TPSL_64_DEFAULT_TARGET@/lib/pkgcon
↳ fig
PE_CRAY_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/pa
↳ rallel-netcdf/1.8.1.3/CRAY/8.6/lib/pkgconfig:/op
↳ t/cray/pe/netcdf-hdf5parallel/4.6.1.3/CRAY/8.6/l
↳ ib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/CRAY/8.
↳ 6/lib/pkgconfig:/opt/cray/pe/hdf5-parallel/1.10.
↳ 2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/hdf5/1.1
↳ 0.2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/ga/5.3
↳ .0.8/CRAY/8.6/lib/pkgconfig
PE_TRILINOS_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.7
SSH_TTY=/dev/pts/38
MIC_LD_LIBRARY_PATH=/opt/intel/compilers_and_librari
↳ es_2018.0.128/linux/mpi/mic/lib:/opt/intel/comp
↳ ilers_and_libraries_2018.0.128/linux/compiler/lib
↳ /mic:/opt/intel/compilers_and_libraries_2018.0.1
↳ 28/linux/ipp/lib/mic:/opt/intel/compilers_and_li
↳ braries_2018.0.128/linux/compiler/lib/mic:/opt/i
↳ ntel/compilers_and_libraries_2018.0.128/linux/mk
↳ l/lib/mic:/opt/intel/compilers_and_libraries_201
↳ 8.0.128/linux/tbb/lib/mic
PE_LIBSCI_DEFAULT_OMP_REQUIRES_openmp=_mp
PE_PETSC_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
PE_FORTRAN_PKGCONFIG_LIBS=mpichf90
CRAYPAT_ALPS_COMPONENT=/opt/cray/pe/perftools/7.0.4/
↳ sbin/pat_alps
CRAYPAT_LD_LIBRARY_PATH=/opt/cray/pe/gcc-libs:/opt/c
↳ ray/gcc-libs:/opt/cray/pe/perftools/7.0.4/lib64
PE_SMA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/
↳ mpt/7.7.3/gni/sma@PE_SMA_DEFAULT_DIR_DEFAULT64@/
↳ lib64/pkgconfig
ALLINEA_QUEUE_DLL=/opt/cray/pe/mpt/7.7.3/gni/mpich-i
↳ ntel/16.0/lib/libtvpmpich.so.3.0.1
PE_TRILINOS_DEFAULT_GENCOMPS_INTEL_x86_64=160
CRAY_MPICH_BASEDIR=/opt/cray/pe/mpt/7.7.3/gni
JRE_HOME=/usr/lib64/jvm/java/jre
PE_HDF5_PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1
↳ 5.3 4.9
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2
↳ 7.1 6.1 5.3 4.9
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_haswell=160

```

```

LD_LIBRARY_PATH=/soft/perftools/darshan/darshan-3.1.
↳ 5/lib:/opt/cray/pe/papi/5.6.0.4/lib64:/opt/cray/
↳ job/2.2.3-6.0.6.0_9.47__g6c4e934.ari/lib64:/opt/
↳ intel/compilers_and_libraries_2018.0.128/linux/c
↳ ompiler/lib/intel64:/opt/intel/compilers_and_lib
↳ raries_2018.0.128/linux/compiler/lib/intel64_lin
↳ :/opt/intel/compilers_and_libraries_2018.0.128/l
↳ inux/mpi/intel64/lib:/opt/intel/compilers_and_li
↳ braries_2018.0.128/linux/mpi/mic/lib:/opt/intel/
↳ compilers_and_libraries_2018.0.128/linux/ipp/lib
↳ /intel64:/opt/intel/compilers_and_libraries_2018
↳ .0.128/linux/compiler/lib/intel64:/opt/intel/com
↳ pilers_and_libraries_2018.0.128/linux/mkl/lib/in
↳ tel64:/opt/intel/compilers_and_libraries_2018.0.
↳ 128/linux/tbb/lib/intel64/gcc4.4:/opt/intel/debu
↳ gger_2017/libipt/intel64/lib:/opt/intel/compiler
↳ s_and_libraries_2018.0.128/linux/daal/lib/intel6
↳ 4_lin:/opt/intel/compilers_and_libraries_2018.0.
↳ 128/linux/daal/./compiler/lib/intel64_lin:/dbho
↳ me/db2cat/sql/lib64:/dbhome/db2cat/sql/lib/lib
↳ 64/gskit:/dbhome/db2cat/sql/lib/lib32
LS_COLORS=no=00:fi=00:di=01;34:ln=00;36:pi=40;33:so=
↳ 01;35:do=01;35:bd=40;33;01:cd=40;33;01:or=41;33;
↳ 01:ex=00;32:*.cmd=00;32:*.exe=01;32:*.com=01;32:
↳ *.bat=01;32:*.btm=01;32:*.dll=01;32:*.tar=00;31:
↳ *.tbz=00;31:*.tgz=00;31:*.rpm=00;31:*.deb=00;31:
↳ *.arj=00;31:*.taz=00;31:*.lzh=00;31:*.lzm=00;31
↳ :*.zip=00;31:*.zoo=00;31:*.z=00;31:*.Z=00;31:*.g
↳ z=00;31:*.bz2=00;31:*.tb2=00;31:*.tz2=00;31:*.tb
↳ z2=00;31:*.xz=00;31:*.avi=01;35:*.bmp=01;35:*.fl
↳ i=01;35:*.gif=01;35:*.jpg=01;35:*.jpeg=01;35:*.m
↳ ng=01;35:*.mov=01;35:*.mpg=01;35:*.pcx=01;35:*.p
↳ bm=01;35:*.pgm=01;35:*.png=01;35:*.ppm=01;35:*.t
↳ ga=01;35:*.tif=01;35:*.xbm=01;35:*.xpm=01;35:*.d
↳ l=01;35:*.gl=01;35:*.wmv=01;35:*.aiff=00;32:*.au
↳ =00;32:*.mid=00;32:*.mp3=00;32:*.ogg=00;32:*.voc
↳ =00;32:*.wav=00;32:
PE_FFTW_DEFAULT_TARGET_interlagos=interlagos
PE_LIBSCI_DEFAULT_VOLATILE_PRGENV=CRAY GNU INTEL
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_TPSL_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TRILINOS_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.3
↳ 5.1 4.9
PE_TRILINOS_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
CRAY_RCA_POST_LINK_OPTS=-L/opt/cray/rca/2.2.18-6.0.6
↳ .0_19.14__g2aa4f39.ari/lib64
↳ -lrca
PE_LIBSCI_PKGCONFIG_VARIABLES=PE_LIBSCI_OMP_REQUIRES
↳ _@openmp@:PE_SCI_EXT_LIBPATH:PE_SCI_EXT_LIBNAME
PE_PETSC_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↳ GNU64 INTEL INTEL64
PE_PKGCONFIG_LIBS=darshan-runtime:mpich:AtSigHandle
↳ r:cray-rca:libsci_mpi:libsci
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2
↳ 7.1 5.3 4.9

```



```

PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_MPICH_FIXED_PRGENV=INTEL
XNLSPATH=/usr/share/X11/nls
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_skylake=16.0
PE_PETSC_DEFAULT_GENCOMPS_GNU_interlagos=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_GNU_sandybridge=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_PETSC_DEFAULT_GENCOMPS_INTEL_sandybridge=160
PE_TPSL_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
INTEL_VERSION=18.0.0.128
MIC_LIBRARY_PATH=/opt/intel/compilers_and_libraries_
↪ 2018.0.128/linux/mpi/mic/lib:/opt/intel/compiler_
↪ s_and_libraries_2018.0.128/linux/compiler/lib/mi
↪ c:/opt/intel/compilers_and_libraries_2018.0.128/
↪ linux/tbb/lib/mic
MPICH_ABORT_ON_ERROR=1
PE_LIBSCI_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_PAPI_DEFAULT_PKGCONFIG_VARIABLES=PE_PAPI_ACCEL_LI
↪ BS@accelerator@
↪ PE_PAPI_ACCEL_FAMILY_LIBS@accelerator_family@
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_PETSC_DEFAULT_GENCOMPS_GNU_mic_knl=53
PE_PETSC_DEFAULT_GENCOMPS_INTEL_mic_knl=160
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2
↪ 7.1 5.3 4.9
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_sandybridge=160
MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-intel/16.0
HOSTTYPE=x86_64
ATP_POST_LINK_OPTS=-Wl,-L/opt/cray/pe/atp/2.1.3/libA
↪ pp/
CPATH=/opt/intel/compilers_and_libraries_2018.0.128/
↪ linux/ipp/include:/opt/intel/compilers_and_libra
↪ ries_2018.0.128/linux/mkl/include:/opt/intel/com
↪ pilers_and_libraries_2018.0.128/linux/tbb/includ
↪ e:/opt/intel/compilers_and_libraries_2018.0.128/
↪ linux/daal/include
PE_FFTW_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_FFTW_DEFAULT_TARGET_sandybridge=sandybridge
PE_HDF5_PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_NETCDF_HDF5PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_
↪ HDF5_PARALLEL
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_MPICH_FORTRAN_PKGCONFIG_LIBS=mpichf90
IBM_DB_HOME=/dbhome/db2cat/sqllib
RCLocal_PRGENV=true
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_mic_knl=5.3
FROM_HEADER=
CHPL_CG_CPP_LINES=1
OFFLOAD_INIT=on_start
PE_LIBSCI_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_LIBSCI_GENCOMPS_INTEL_x86_64=160

```

```

PE_PRODUCT_LIST=CRAYPE_MIC-KNL:CRAY_RCA:CRAY_ALPS:DV
↪ S:CRAY_XPMEM:CRAY_DMAPP:CRAY_PMI:CRAY_UGNI:CRAY_
↪ UDREG:CRAY_LIBSCI:CRAYPE:INTEL:PERFTOOLS:CRAYPAT
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_TPSL_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53 49
PAGER=less
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 5.3 4.9
PE_TPSL_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
CRAY_MPICH_ROOTDIR=/opt/cray/pe/mpt/7.7.3
CSHEDIT=emacs
PE_LIBSCI_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1 4.9
PE_PETSC_DEFAULT_GENCOMPS_GNU_skylake=61
PE_PETSC_DEFAULT_GENCOMPS_INTEL_skylake=160
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_MPICH_GENCOMPILERS_CRAY=8.6
PE_MPICH_MODULE_NAME=cray-mpich
XDG_CONFIG_DIRS=/etc/xdg
CRAYPAT_ROOT=/opt/cray/pe/perftools/7.0.4
INTEL_MAJOR_VERSION=18.0
PE_LIBSCI_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_LIBSCI_GENCOMPS_CRAY_x86_64=86
PE_MPICH_DEFAULT_VOLATILE_PRGENV=CRAY_GNU
PE_MPICH_TARGET_VAR_nvidia20=-lcudart
PE_TPSL_64_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIB
↪ SCI
PE_TPSL_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_TPSL_DEFAULT_GENCOMPS_CRAY_sandybridge=86
MINICOM=-c on
LIBGL_DEBUG=quiet
USERMODULES=acml:alps:apprentice:apprentice2:atp:blc
↪ r:cce:chapel:cray-ccdb:cray-fftw:cray-ga:cray-hd
↪ f5:cray-hdf5-parallel:cray-lgdb:cray-libsci:cray
↪ -libsci_acc:cray-mpich:cray-mpich2:cray-mpich-com
↪ pat:cray-netcdf:cray-netcdf-hdf5parallel:cray-pa
↪ rallel-netcdf:craypat:craype:cray-petsc:cray-pet
↪ sc-complex:craypkg-gen:cray-shmem:cray-snplaunch
↪ er:cray-tpsl:cray-trilinos:cudatoolkit:ddt:fftw:
↪ ga:gcc:hdf5:hdf5-parallel:intel:iobuf:java:lgdb:
↪ libfast:libsci_acc:mpich1:netcdf:netcdf-hdf5para
↪ llel:netcdf-nofsync:netcdf-nofsync-hdf5parallel:
↪ ntk:onesided:papi:parallel-netcdf:pathscale:perf
↪ tools:perftools-lite:petsc:petsc-complex:pgi:pmi
↪ :PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:PrgEnv-path
↪ scale:PrgEnv-pgi:stat:totalview:tpsl:trilinos:xt
↪ -asynce:xt-craypat:xt-lgdb:xt-libsci:xt-mpich2:x
↪ t-mpt:xt-papi:xt-shmem:xt-totalview
CRAY_DMAPP_INCLUDE_OPTS=-I/opt/cray/dmapp/7.1.1-6.0.
↪ 6.0_51.37__g5a674e0.ari/include
↪ -I/opt/cray/gni-headers/5.0.12-6.0.6.0_3.26__g52
↪ 7b6e1.ari/include
CRAY_LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
CRAY_LIBSCI_DIR=/opt/cray/pe/libsci/18.07.1
DVS_VERSION=0.9.0

```

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

```
NLSPATH=/opt/intel/compilers_and_libraries_2018.0.12
↳ 8/linux/compiler/lib/intel64/locale/%l_%t/%N:/opt
↳ t/intel/compilers_and_libraries_2018.0.128/linux
↳ /mk1/lib/intel64/locale/%l_%t/%N:/opt/intel/debu
↳ gger_2017/gdb/intel64_mic/share/locale/%l_%t/%N:
↳ /opt/intel/debugger_2017/gdb/intel64/share/local
↳ e/%l_%t/%N
PE_LIBSCI_PKGCONFIG_LIBS=libsci_mpi:libsci
PE_NETCDF_DEFAULT_GENCOMPS_GNU=
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPS_GNU=51 49
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49
PATH=/opt/xalt/bin:/soft/buildtools/trackdeps/bin:/s
↳ oft/perftools/darshan/darshan-3.1.5/bin:/opt/mod
↳ ulefiles/nompirun/bin:/opt/cray/pe/mpt/7.7.3/gni
↳ /bin:/opt/cray/pe/perftools/7.0.4/bin:/opt/cray/
↳ pe/papi/5.6.0.4/bin:/opt/cray/rca/2.2.18-6.0.6.0
↳ _19.14_g2aa4f39.ari/bin:/opt/cray/alps/6.6.1-6.
↳ 0.6.1_4.1_ga6396bb.ari/sbin:/opt/cray/job/2.2.3
↳ -6.0.6.0_9.47_g6c4e934.ari/bin:/opt/cray/pe/cray
↳ pe/2.5.15/bin:/opt/intel/compilers_and_libraries
↳ _2018.0.128/linux/bin/intel64:/opt/intel/compile
↳ rs_and_libraries_2018.0.128/linux/mpi/intel64/bi
↳ n:/opt/intel/debugger_2017/gdb/intel64_mic/bin:/
↳ opt/cray/pe/modules/3.2.11.1/bin:/usr/local/bin:
↳ /usr/bin:/bin:/usr/bin/X11:/usr/games:/dbhome/db
↳ 2cat/sqllib/bin:/dbhome/db2cat/sqllib/adm:/dbhom
↳ e/db2cat/sqllib/misc:/usr/lib/mit/bin:/usr/lib/m
↳ it/sbin:/opt/cray/pe/bin
MODULE_VERSION=3.2.11.1
PAT_REPORT_PRUNE_NAME=_cray$mt_execute_,_cray$mt_sta
↳ rt_,_cray_hwpc_,f_cray_hwpc_,cstart_,_pat_,pat_
↳ region_,PAT_,OMP.slave_loop,slave_entry,_new_sla
↳ ve_entry,_thread_pool_slave_entry,THREAD_POOL_jo
↳ in_,_libc_start_main,_start_,_start_,start_thread
↳ ,_wrap_,UPC_ADIO_,_upc_,upc_,_caf_,_pgas_,sys
↳ call_,_device_stub
PE_HDF5_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ /hdf5/1.10.2.0/@PRGENV@/@PE_HDF5_DEFAULT_GENCOMP
↳ S@/lib/pkgconfig
PE_PKGCONFIG_DEFAULT_PRODUCTS=PE_TRILINOS:PE_TPSSL_64
↳ :PE_TPSSL:PE_PETSC:PE_PARALLEL_NETCDF:PE_NETCDF_H
↳ DF5PARALLEL:PE_NETCDF:PE_MPICH:PE_LIBSCI:PE_HDF5
↳ _PARALLEL:PE_HDF5:PE_GA:PE_FFTW
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1 5.3
↳ 4.9
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_interlagos=86
PE_MPICH_GENCOMPILERS_GNU=7.1 5.1 4.9
CPU=x86_64
XTPE_NETWORK_TARGET=aries
ATP_IGNORE_SIGTERM=1
PE_FFTW_DEFAULT_TARGET_abudhabi=abudhabi
PE_NETCDF_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPILERS_GNU=5.1 4.9
PE_PETSC_DEFAULT_GENCOMPS_CRAY_mic_knl=86
```

```
PE_TPSSL_64_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2
↳ 7.1 6.1
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1 5.3
↳ 4.9
SSH_SENDS_LOCALE=yes
JAVA_BINDIR=/usr/lib64/jvm/java/bin
PE_HDF5_PARALLEL_DEFAULT_FIXED_PRGENV=CRAY INTEL
PE_HDF5_PARALLEL_DEFAULT_GENCOMPS_GNU=
PE_NETCDF_HDF5PARALLEL_DEFAULT_FIXED_PRGENV=CRAY
↳ INTEL
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS_GNU=
PE_SMA_DEFAULT_DIR_CRAY_DEFAULT64=64
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
CRAY_UDREG_POST_LINK_OPTS=-L/opt/cray/udreg/2.3.2-6.
↳ 0.6.0_15.18_g5196236.ari/lib64
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_TPSSL_64_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↳ GNU64 INTEL INTEL64
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_interlagos=160
TBBROOT=/opt/intel/compilers_and_libraries_2018.0.12
↳ 8/linux/tbb
INPUTRC=/etc/inputrc
CRAYPE_VERSION=2.5.15
CRAY_ALPS_POST_LINK_OPTS=-L/opt/cray/alps/6.6.1-6.0.
↳ 6.1_4.1_ga6396bb.ari/lib64
PE_TPSSL_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
PE_MPICH_VOLATILE_PRGENV=CRAY GNU
JAVA_HOME=/usr/lib64/jvm/java
TARGETMODULES=craype-abudhabi:craype-abudhabi-cu:cray
↳ ype-accel-host:craype-accel-nvidia20:craype-acce
↳ l-nvidia30:craype-accel-nvidia35:craype-barcelon
↳ a:craype-broadwell:craype-haswell:craype-hugepag
↳ es128K:craype-hugepages128M:craype-hugepages16M:
↳ craype-hugepages256M:craype-hugepages2M:craype-h
↳ ugepages32M:craype-hugepages4M:craype-hugepages5
↳ 12K:craype-hugepages512M:craype-hugepages64M:cra
↳ ype-hugepages8M:craype-intel-knc:craype-interlag
↳ os:craype-interlagos-cu:craype-istanbul:craype-i
↳ vybridge:craype-mc12:craype-mc8:craype-mic-knl:c
↳ raype-network-aries:craype-network-gemini:craype
↳ -network-infiniband:craype-network-none:craype-ne
↳ twork-seastar:craype-sandybridge:craype-shanghai
↳ :craype-target-compute_node:craype-target-local_
↳ host:craype-target-native:craype-xeon:xtpe-barce
↳ lona:xtpe-interlagos:xtpe-interlagos-cu:xtpe-ist
↳ anbul:xtpe-mc12:xtpe-mc8:xtpe-network-gemini:xtp
↳ e-network-seastar:xtpe-shanghai:xtpe-target-nati
↳ ve:xtpe-xeon
```

```

_LMFILES=/opt/cray/pe/modulefiles/modules/3.2.11.1:
↳ /opt/modulefiles/intel/18.0.0.128:/opt/cray/pe/c
raype/2.5.15/modulefiles/craype-network-aries:/o
pt/cray/pe/modulefiles/craype/2.5.15:/opt/cray/p
e/modulefiles/cray-libsci/18.07.1:/opt/cray/ari/
modulefiles/udreg/2.3.2-6.0.6.0_15.18__g5196236.
ari:/opt/cray/ari/modulefiles/ugni/6.0.14-6.0.6.
0_18.12__g777707d.ari:/opt/cray/pe/modulefiles/p
mi/5.0.14:/opt/cray/ari/modulefiles/dmapp/7.1.1-
6.0.6.0_51.37__g5a674e0.ari:/opt/cray/ari/module
files/gni-headers/5.0.12-6.0.6.0_3.26__g527b6e1.
ari:/opt/cray/ari/modulefiles/xpmem/2.2.14-6.0.6.
1_5.14__g34333c9.ari:/opt/cray/ari/modulefiles/
job/2.2.3-6.0.6.0_9.47__g6c4e934.ari:/opt/cray/a
ri/modulefiles/dvs/2.7_2.2.100-6.0.6.1_12.6__g54
2256c:/opt/cray/ari/modulefiles/alps/6.6.1-6.0.6.
1_4.1__ga6396bb.ari:/opt/cray/ari/modulefiles/r
ca/2.2.18-6.0.6.0_19.14__g2aa4f39.ari:/opt/cray/
pe/modulefiles/atp/2.1.3:/opt/cray/pe/modulefile
s/perftools-base/7.0.4:/opt/cray/pe/modulefiles/
PrgEnv-intel/6.0.4:/opt/cray/pe/craype/2.5.15/mo
dulefiles/craype-mic-kl:/opt/cray/pe/modulefile
s/cray-mpich/7.7.3:/opt/modulefiles/nompirun/nom
pirun:/soft/environment/modules/modulefiles/dars
han/3.1.5:/soft/environment/modules/modulefiles/
trackdeps:/soft/environment/modules/modulefiles/
xalt
PE_LIBSCI_DEFAULT_OMP_REQUIRES=
PE_MPICH_DEFAULT_GENCOMPS_CRAY=86
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_sandybridge=7.1
↳ 5.3 4.9
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_LIBSCI_MODULE_NAME=cray-libsci/18.07.1
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_skylake=8.6
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
LANG=en_US.UTF-8
GDB_CROSS=/opt/intel/debugger_2017/gdb/intel64_mic/b
↳ in/gdb-mic
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
PE_INTEL_FIXED_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.7.3
↳ /gni/mpich-intel/16.0/lib/pkgconfig
MODULEPATH=/opt/cray/pe/perftools/7.0.4/modulefiles:
↳ /opt/cray/pe/craype/2.5.15/modulefiles:/opt/cray
↳ /pe/modulefiles:/opt/cray/modulefiles:/opt/modul
↳ efiles:/soft/environment/modules/modulefiles:/op
↳ t/cray/ari/modulefiles
PYTHONSTARTUP=/etc/pythonstart
PE_LIBSCI_GENCOMPILERS_CRAY_x86_64=8.6
PE_MPICH_NV_LIBS_nvidia20=-lcudart
PE_MPICH_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.
↳ 7.3/gni/mpich-@PRGENV@PE_MPICH_DIR_DEFAULT64@/@
↳ PE_MPICH_GENCOMPS@lib/pkgconfig
SDK_HOME=/usr/lib64/jvm/java
TZ=UTC

```

```

LOADED_MODULES=modules/3.2.11.1:intel/18.0.0.128:cray
↳ pe-network-aries:craype/2.5.15:cray-libsci/18.07
↳ .1:udreg/2.3.2-6.0.6.0_15.18__g5196236.ari:ugni/
↳ 6.0.14-6.0.6.0_18.12__g777707d.ari:pmi/5.0.14:dm
↳ app/7.1.1-6.0.6.0_51.37__g5a674e0.ari:gni-header
↳ s/5.0.12-6.0.6.0_3.26__g527b6e1.ari:xpmem/2.2.14
↳ -6.0.6.1_5.14__g34333c9.ari:job/2.2.3-6.0.6.0_9.4
↳ 7__g6c4e934.ari:dvs/2.7_2.2.100-6.0.6.1_12.6__g5
↳ 42256c:alps/6.6.1-6.0.6.1_4.1__ga6396bb.ari:rca/
↳ 2.2.18-6.0.6.0_19.14__g2aa4f39.ari:atp/2.1.3:per
↳ ftools-base/7.0.4:PrgEnv-intel/6.0.4:craype-mic-
↳ knl:cray-mpich/7.7.3:nompirun/nompirun:darshan/3
↳ .1.5:trackdeps:xalt
SHMEM_ABORT_ON_ERROR=1
CRAY_DMAPP_POST_LINK_OPTS=-L/opt/cray/dmapp/7.1.1-6.
↳ 0.6.0_51.37__g5a674e0.ari/lib64
PE_FFTW_DEFAULT_TARGET_ivybridge=ivybridge
PE_FFTW_DEFAULT_TARGET_share=share
PE_FFTW_DEFAULT_TARGET_x86_skylake=x86_skylake
PE_PKG_CONFIG_PATH=/opt/cray/pe/cti/1.0.7/lib/pkgcon
↳ fig:/opt/cray/pe/cti/1.0.6/lib/pkgconfig:/opt/cr
↳ ay/pe/cti/1.0.4/lib/pkgconfig:/opt/cray/pe/cti/1
↳ .0.5/lib/pkgconfig
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53
↳ 49
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
CRAY_RCA_INCLUDE_OPTS=-I/opt/cray/rca/2.2.18-6.0.6.0
↳ _19.14__g2aa4f39.ari/include
↳ -I/opt/cray/krca/2.2.4-6.0.6.1_5.30__g8505b97.ar
↳ i/include
↳ -I/opt/cray-hss-devel/8.0.0/include
PAT_BUILD_PAPI_BASEDIR=/opt/cray/pe/papi/5.6.0.4
PE_LIBSCI_OMP_REQUIRES_openmp=_mp
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_skylake=6.1
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_mic_knl=86
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
CRAY_MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-inte
↳ l/16.0
PE_MPICH_CXX_PKGCONFIG_LIBS=mpichcxx
DB2_HOME=/dbhome/db2cat/sqllib
DAALROOT=/opt/intel/compilers_and_libraries_2018.0.1
↳ 28/linux/daal
PE_LIBSCI_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_MPICH_PKGCONFIG_VARIABLES=PE_MPICH_NV_LIBS_@accel
↳ erator@:PE_MPICH_ALTERNATE_LIBS_@multithreaded@:
↳ PE_MPICH_ALTERNATE_LIBS_@dpm@
APP2_STATE=7.0.4
CRAY_PMI_POST_LINK_OPTS=-L/opt/cray/pe/pmi/5.0.14/li
↳ b64
MPM_LAUNCHER=/opt/intel/debugger_2017/mpm/mic/bin/st
↳ art_mpm.sh
PE_HDF5_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16.0

```

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks

```

PE_TPSL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ /tpsl/18.06.1/@PRGENV@/PE_TPSL_DEFAULT_GENCOMPS
↳ @/PE_TPSL_DEFAULT_TARGET@/lib/pkgconfig
CRAY_MPICH2_VER=7.7.3
PE_MPICH_PKGCONFIG_LIBS=mpich
GPG_TTY=/dev/pts/38
INTEL_PYTHONHOME=/opt/intel/debugger_2017/python/intel
↳ el64/
PE_GA_DEFAULT_GENCOMPILERS_GNU=5.3 4.9
PE_LIBSCI_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/libsci
↳ i/18.07.1/@PRGENV@/PE_LIBSCI_GENCOMPS@/PE_LIBS
↳ CI_TARGET@/lib/pkgconfig
PE_MPICH_ALTERNATE_LIBS_multithreaded=_mt
PE_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_PARALLEL_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
SHLVL=2
JDK_HOME=/usr/lib64/jvm/java
QT_SYSTEM_DIR=/usr/share/desktop-data
CRAY_LIBSCI_VERSION=18.07.1
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
PE_MPICH_TARGET_VAR_nvidia35=-lcudart
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
PE_PKGCONFIG_PRODUCTS_DEFAULT=PE_PAPI
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_haswell=8.2 7.1 5.3 4.9
OSTYPE=linux
LESS_ADVANCED_PREPROCESSOR=no
IBM_DB_INCLUDE=/dbhome/db2cat/sql/lib/include
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
PE_MPICH_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ e/mpt/7.7.3/gni/mpich-@PRGENV@/PE_MPICH_DEFAULT_
↳ DIR_DEFAULT64@/PE_MPICH_DEFAULT_GENCOMPS@/lib/pkgconfig
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_TPSL_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU GNU64
↳ INTEL_INTEL64
XCURSOR_THEME=DMZ
LS_OPTIONS=-N --color=none -T 0
CRAY_PMI_INCLUDE_OPTS=-I/opt/cray/pe/pmi/5.0.14/include
↳ ude
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_interlagos=8.6
PE_TPSL_DEFAULT_GENCOMPS_INTEL_sandybridge=16.0
WINDOWMANAGER=
PRGENVMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:PrgEnv-pgi
CRAYPE_NETWORK_TARGET=aries
ATP_MRNET_COMM_PATH=/opt/cray/pe/atp/2.1.3/libexec/atp
↳ tp_mrnet_commnnode_wrapper
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PKGCONFIG_PATH_DEFAULT=/opt/cray/pe/papi/5.6.0.4/lib64/pkgconfig
PE_MPICH_DIR_CRAY_DEFAULT64=64
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_haswell=7.1 5.3 4.9
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2 7.1
↳ 5.3 4.9
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0

```

```

MACHTYPE=x86_64-suse-linux
LESS=-M -I -R
G_FILENAME_ENCODING=@locale,UTF-8,ISO-8859-15,CP1252
CRAY_GNI_HEADERS_INCLUDE_OPTS=-I/opt/cray/gni-header
↳ s/5.0.12-6.0.6.0.3.26_g527b6e1.ari/include
CRAY_LIBSCI_PREFIX_DIR=/opt/cray/pe/libsci/18.07.1/INTEL
↳ INTEL/16.0/x86_64
PE_HDF5_DEFAULT_GENCOMPS_GNU=
PE_MPICH_NV_LIBS=
PE_NETCDF_DEFAULT_REQUIRED_PRODUCTS=PE_HDF5
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1
↳ 5.3 4.9
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16
↳ .0
PE_TPSL_DEFAULT_GENCOMPS_GNU_x86_64=8.2 7.1 5.3 4.9
PE_TRILINOS_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_HDF5
↳ F5_PARALLEL:PE_NETCDF_HDF5PARALLEL:PE_LIBSCI:PE_
↳ TPSL
CVS_RSH=ssh
DMAPP_ABORT_ON_ERROR=1
PE_LIBSCI_OMP_REQUIRES=
PE_MPICH_DEFAULT_GENCOMPILERS_CRAY=8.6
PE_TRILINOS_DEFAULT_GENCOMPS_GNU_x86_64=8.2 7.1 5.3 4.9
PE_MPICH_GENCOMPS_CRAY=8.6
SSH_CONNECTION=130.202.141.135 56759 140.221.69.22 22
CLASSPATH=/opt/intel/compilers_and_libraries_2018.0.
↳ 128/linux/daal/lib/daal.jar:/dbhome/db2cat/sql/lib
↳ b/java/db2java.zip:/dbhome/db2cat/sql/lib/java/sql
↳ lj.zip:/dbhome/db2cat/sql/lib/function:/dbhome/db
↳ 2cat/sql/lib/java/db2jcc_license_cu.jar:/dbhome/db
↳ b2cat/sql/lib/tools/clplplus.jar:/dbhome/db2cat/sql
↳ lib/tools/jline-0.9.93.jar:/dbhome/db2cat/sql/lib
↳ b/java/db2jcc.jar:.
XDG_DATA_DIRS=/usr/share
TOOLMODULES=apprentice:apprentice2:atp:chapel:cray-l
↳ gdb:craypat:craypkg-gen:cray-snp-launcher:ddt:gdb
↳ :iobuf:papi:perftools:perftools-lite:stat:totalview
↳ iew:xt:craypat:xt-lgdb:xt-papi:xt-totalview
DVS_INCLUDE_OPTS=-I/opt/cray/dvs/2.7.2.2.100-6.0.6.1
↳ _12.6_g542256c/include
PE_LIBSCI_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_MPICH_DEFAULT_FIXED_PRGENV=INTEL
PE_MPICH_DEFAULT_GENCOMPS_GNU=7.1 5.1 4.9
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
TRACKDEPS_CMDS=cc1 cc1plus f951 clang clang-3.9
↳ clang-4.0 clang-5.0 clang-6.0 clang-7.0 cpp as
↳ mcpcom fortcom ccfe ftnfe driver ar ld.bfd
↳ ld.gold ld
MODULESHOME=/opt/cray/pe/modules/3.2.11.1
CRAY_PRGENVINTEL=loaded
PE_GA_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_LIBSCI_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ pe/libsci/18.07.1/@PRGENV@/PE_LIBSCI_DEFAULT_GENCOMPS
↳ @/PE_LIBSCI_DEFAULT_TARGET@/lib/pkgconfig

```

```

PE_TPSL_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2 7.1
↳ 5.3 4.9
TRACKDEPS_LOG_PATH=/lus/theta-fs0/logs/trackdeps
LESSOPEN=lessopen.sh %s
PE_MPICH_NV_LIBS_nvidia35=-lcudart
PE_PETSC_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/p
↳ e/petsc/3.8.4.0/complex/@PRGENV@/PE_PETSC_DEFAU
↳ LT_GENCOMPS@/PE_PETSC_DEFAULT_TARGET@/lib/pkgco
↳ nfig
PKG_CONFIG_PATH=/soft/perftools/darshan/darshan-3.1.
↳ 5/lib/pkgconfig:/opt/cray/rca/2.2.18-6.0.6.0_19.
↳ 14__g2aa4f39.ari/lib64/pkgconfig:/opt/cray/alps/
↳ 6.6.1-6.0.6.1_4.1__ga6396bb.ari/lib64/pkgconfig:
↳ /opt/cray/xpmem/2.2.14-6.0.6.1_5.14__g34333c9.ar
↳ i/lib64/pkgconfig:/opt/cray/gni-headers/5.0.12-6
↳ .0.6.0_3.26__g527b6e1.ari/lib64/pkgconfig:/opt/c
↳ ray/dmapp/7.1.1-6.0.6.0_51.37__g5a674e0.ari/lib6
↳ 4/pkgconfig:/opt/cray/pe/pmi/5.0.14/lib64/pkgcon
↳ fig:/opt/cray/ugni/6.0.14-6.0.6.0_18.12__g77707
↳ d.ari/lib64/pkgconfig:/opt/cray/udreg/2.3.2-6.0.
↳ 6.0_15.18__g5196236.ari/lib64/pkgconfig:/opt/cra
↳ y/pe/craype/2.5.15/pkg-config:/opt/cray/pe/iobuf
↳ /2.0.8/lib/pkgconfig:/opt/cray/pe/atp/2.1.3/lib/
↳ pkgconfig
COMPILER_PATH=/opt/xalt/bin
PELOCAL_PRGENV=true
CRAYPAT_OPTS_EXECUTABLE=sbin/pat-opts
LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_x86_64=160
INFOPATH=/opt/intel/documentation_2017/en/debugger//
↳ gdb-ia/info:/opt/intel/documentation_2017/en/de
↳ bugger//gdb-mic/info:/opt/intel/documentation_2
↳ 017/en/debugger//gdb-igfx/info/
LIBSCI_VERSION=18.07.1
PE_LIBSCI_DEFAULT_PKGCONFIG_VARIABLES=PE_LIBSCI_DEFA
↳ ULT_OMP_REQUIRES_@openmp@:PE_SCI_EXT_LIBPATH:PE_
↳ SCI_EXT_LIBNAME
PE_MPICH_NV_LIBS_nvidia60=-lcudart
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53
↳ 49
PE_TPSL_DEFAULT_GENCOMPS_INTEL_mic_knl=160
XDG_RUNTIME_DIR=/run/user/32451
CRAY_PRE_COMPILE_OPTS=-hnetwork=aries
CRAY_ALPS_INCLUDE_OPTS=-I/opt/cray/alps/6.6.1-6.0.6.
↳ 1_4.1__ga6396bb.ari/include
PE_FFTW_DEFAULT_TARGET_broadwell=broadwell
PE_LIBSCI_GENCOMPILERS_INTEL_x86_64=16.0
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1
↳ 5.3 4.9
CRAY_CPU_TARGET=mic-knl
CRAY_UGNI_INCLUDE_OPTS=-I/opt/cray/ugni/6.0.14-6.0.6
↳ .0_18.12__g77707d.ari/include
CRAY_XPMEM_INCLUDE_OPTS=-I/opt/cray/xpmem/2.2.14-6.0
↳ .6.1_5.14__g34333c9.ari/include
PE_LIBSCI_REQUIRED_PRODUCTS=PE_MPICH
PE_PAPI_DEFAULT_ACCELL_FAMILY_LIBS=

```

```

PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_x86_64=86
craype_already_loaded=0
PE_LIBSCI_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1
↳ 4.9
PE_LIBSCI_GENCOMPS_GNU_x86_64=71 61 51 49
PE_TPSL_DEFAULT_GENCOMPS_INTEL_haswell=160
LESSCLOSE=lessclose.sh %s %s
DB2LIB=/dbhome/db2cat/sql/lib/lib
ATP_HOME=/opt/cray/pe/atp/2.1.3
PE_FFTW_DEFAULT_TARGET_x86_64=x86_64
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
G_BROKEN_FILENAMES=1
CRAY_LD_LIBRARY_PATH=/opt/cray/pe/mpt/7.7.3/gni/mpic
↳ h-intel/16.0/lib:/opt/cray/pe/perftools/7.0.4/li
↳ b64:/opt/cray/rca/2.2.18-6.0.6.0_19.14__g2aa4f39
↳ .ari/lib64:/opt/cray/alps/6.6.1-6.0.6.1_4.1__ga6
↳ 396bb.ari/lib64:/opt/cray/xpmem/2.2.14-6.0.6.1_5
↳ .14__g34333c9.ari/lib64:/opt/cray/dmapp/7.1.1-6.
↳ 0.6.0_51.37__g5a674e0.ari/lib64:/opt/cray/pe/pmi
↳ /5.0.14/lib64:/opt/cray/ugni/6.0.14-6.0.6.0_18.1
↳ 2__g77707d.ari/lib64:/opt/cray/udreg/2.3.2-6.0.
↳ 6.0_15.18__g5196236.ari/lib64:/opt/cray/pe/libsc
↳ i/18.07.1/INTEL/16.0/x86_64/lib
PE_FFTW_DEFAULT_TARGET_haswell=haswell
PE_GA_DEFAULT_GENCOMPS_GNU=53 49
PE_GA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/g
↳ a/5.3.0.8/@PRGENV@/PE_GA_DEFAULT_GENCOMPS@/lib/
↳ pkgconfig
PE_INTEL_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/p
↳ arallel-netcdf/1.8.1.3/INTEL/16.0/lib/pkgconfig:
↳ /opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/INTEL/1
↳ 6.0/lib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/IN
↳ TEL/16.0/lib/pkgconfig:/opt/cray/pe/mpt/7.7.3/gn
↳ i/mpich-intel/16.0/lib/pkgconfig:/opt/cray/pe/hd
↳ f5-parallel/1.10.2.0/INTEL/16.0/lib/pkgconfig:/o
↳ pt/cray/pe/hdf5/1.10.2.0/INTEL/16.0/lib/pkgconfi
↳ g:/opt/cray/pe/ga/5.3.0.8/INTEL/18.0/lib/pkgconf
↳ ig
PE_PAPI_DEFAULT_ACCEL_LIBS=
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_interlagos=7.1 5.3
↳ 4.9
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_SMA_DEFAULT_DIR_PGI_DEFAULT64=64
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16
↳ .0
COLORTERM=1
JAVA_ROOT=/usr/lib64/jvm/java
I_MPI_ROOT=/opt/intel/compilers_and_libraries_2018.0
↳ .128/linux/mpi
PE_MPICH_DEFAULT_DIR_CRAY_DEFAULT64=64
PE_PETSC_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_PETSC_DEFAULT_GENCOMPS_GNU_x86_64=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_INTEL_x86_64=160
intel_already_loaded=0

```

```

BASH_FUNC_module%%=( ) { eval
↪ ` /opt/cray/pe/modules/3.2.11.1/bin/modulecmd
↪ bash $* `
}
_=/usr/bin/env
collect_environment.sh: line 8: inxi: command not
↪ found
NAME      MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
loop0      7:0    0    3G  0 loop
loop1      7:1    0   50.8G  0 loop
↪ /var/opt/cray/imps-image-binding/PE/squash_mount_
↪ s/squashfs_NSX14L_mount_point
loop2      7:2    0   1.6G  0 loop
↪ /var/opt/cray/imps-image-binding/diags/squash_mo_
↪ unts/squashfs_lueY5C_mount_point
loop3      7:3    0        1 loop
loop4      7:4    0        0 loop
loop5      7:5    0        0 loop
loop6      7:6    0        0 loop
loop7      7:7    0        0 loop
sda        8:0    0   700G  0 disk
├─sda1      8:1    0   1007K  0 part
├─sda2      8:2    0     2G  0 part /boot
├─sda3      8:3    0    20G  0 part
├─sda4      8:4    0   256G  0 part /tmp
├─sda5      8:5    0   128G  0 part [SWAP]
sdb        8:16   0  417.3G  0 disk
├─sdb1      8:17   0    10G  0 part /var/crash
├─sdb2      8:18   0  407.3G  0 part
↪ /var/opt/cray/persistent
sr0        11:0    1  1024M  0 rom
mira-home 253:150 0        0 disk /gpfs/mira-home
theta-fs1 253:151 0        0 disk /gpfs/theta-fs1
[0:2:0:0] disk DELL PERC H330 Mini 4.25
↪ /dev/sda 751GB
[0:2:1:0] disk DELL PERC H330 Mini 4.25
↪ /dev/sdb 448GB
[10:0:0:0] cd/dvd PLDS DVD-ROM DU-8D5LH FD51
↪ /dev/sr0 -
Currently Loaded Modulefiles:
  1) modules/3.2.11.1                                13)
    ↪ dvs/2.7_2.2.100-6.0.6.1_12.6__g542256c
  2) intel/18.0.0.128                                14)
    ↪ alps/6.6.1-6.0.6.1_4.1__ga6396bb.ari
  3) craype-network-aries                             15)
    ↪ rca/2.2.18-6.0.6.0_19.14__g2aa4f39.ari
  4) craype/2.5.15                                    16)
    ↪ atp/2.1.3
  5) cray-libsci/18.07.1                             17)
    ↪ perftools-base/7.0.4
  6) udreg/2.3.2-6.0.6.0_15.18__g5196236.ari
    ↪ 18) PrgEnv-intel/6.0.4
  7) ugni/6.0.14-6.0.6.0_18.12__g777707d.ari
    ↪ 19) craype-mic-knl

```

```

8) pmi/5.0.14                                         20)
    ↪ cray-mpich/7.7.3
  9) dmapp/7.1.1-6.0.6.0_51.37__g5a674e0.ari
    ↪ 21) nompirun/nompirun
 10) gni-headers/5.0.12-6.0.6.0_3.26__g527b6e1.ari
    ↪ 22) darshan/3.1.5
 11) xpmem/2.2.14-6.0.6.1_5.14__g34333c9.ari
    ↪ 23) trackdeps
 12) job/2.2.3-6.0.6.0_9.47__g6c4e934.ari
    ↪ 24) xalt
collect_environment.sh: line 12: nvidia-smi: command
↪ not found
collect_environment.sh: line 13: lshw: command not
↪ found
collect_environment.sh: line 13: lspci: command not
↪ found

```

ARTIFACT EVALUATION

Verification and validation studies: A primary goal of this paper is to validate that the new benchmark suite, GPCNeT, achieves the goals of reliably producing and measuring congestion across a range of architectures and run configurations. In addition to the validation study presented in the paper, we performed several verification studies during the development effort. We verified the isolated performance of the communication kernels in GPCNeT against known hardware limitations and against other benchmarks with similar kernels. This includes hardware limitations, such as latency and bandwidth, on Cray Aries and Infiniband networks. We also verified GPCNeT's measurement of Allreduce performance against accepted benchmarks, such as OSU MPI Microbenchmarks. The results presented on the Cray Slingshot network were in part verified against expectations from network simulations as well.

Accuracy and precision of timings: Precision timing is particularly important for GPCNeT. Measurements of latency are on the order of a microsecond, and software timers with significant overhead (e.g., MPI_Wtime) cannot be used in many cases. In Section 3 of the paper, we show pseudo-code of key loop structures in GPCNeT along with details of iteration counts. Two timers are used in all cases. One is a "bulk" timer, used to reproduce the timing techniques in the HPCC benchmark. That timer uses MPI_Wtime, which has sufficient resolution for the inner iteration loops. The shortest wall-time for any of the tests' loops timed with this is 200 microseconds. Another timer is used to produce histograms of every iteration of a communication kernel. This timer uses clock_gettime(), which has sub-microsecond resolution after accounting for its overhead.

The latency and bandwidth values produced from these timing values were verified against hardware limitations and against other benchmarks with similar communication kernels.

Quantified the sensitivity of results to initial conditions and/or parameters of the computational environment: Section 3 and 4 of the paper review our efforts to describe the sensitivity of GPCNeT's measurements to job size, processes per node, MPI library, and network architecture. Although not all included in the paper, the key statistical values for every run is included with the artifacts for this paper. Statistics, such as average and 99% tail, for runs in all of these

dimensions are included. In the paper, Table 2 includes detailed results across several systems and run configurations. This data shows the response of measured values to job size and processes per node.

Controls, statistics, or other steps taken to make the measurements and analyses robust to variability and unknowns in the system. As described in the paper, GPCNeT relies on randomization and a large number of samples to overcome potential biases due to process placement and node topology. For example, the random ring infrastructure discussed in Section 3 generates 30 different random rings for each rank on a node with different random sequences. This produces 30xPPN (processes per node) different rings over which communication performance is measured. An outer loop, shown in Listing 1, repeats measurements over those rings to improve the significance of measurements. GPCNeT also uses warmup iterations to reduce "jitter" from caching data, for example. A congestor warmup, described in Section 3.3, is used before sensitive traffic begins running to remove variability while the network is filled with traffic.