Uncore Power Scavenger: A Runtime for Uncore Power Conservation on HPC Systems

Neha Gholkar¹, Frank Mueller¹, Barry Rountree²

¹North Carolina State University, USA, ngholka@ncsu.edu, mueller@cs.ncsu.edu

²Lawrence Livermore National Laboratory, USA, rountree1@llnl.gov

ABSTRACT

The US Department of Energy (DOE) has set a power target of 20-30MW on the first exascale machines. To achieve one exaflop under this power constraint, it is necessary to minimize wasteful consumption of power while striving to improve performance.

Toward this end, we investigate uncore frequency scaling (UFS) as a potential knob for reducing the power footprint of HPC jobs. We propose Uncore Power Scavenger (UPSCavenger), a runtime system that dynamically detects phase changes and automatically sets the best uncore frequency for every phase to save power without significant impact on performance. Our experimental evaluations on a cluster show that UPSCavenger achieves up to 10% energy savings with under 1% slowdown. It achieves 14% energy savings with the worst case slowdown of 5.5%. We also show that UPSCavenger achieves up to 20% speedup and proportional energy savings compared to Intel's RAPL with equivalent power usage making it a viable solution even for power-constrained computing.

CCS CONCEPTS

• Software and its engineering → Power management;

KEYWORDS

high-performance computing; power awareness

ACM Reference Format:

Neha Gholkar, Frank Mueller, Barry Rountree. 2019. Uncore Power Scavenger: A Runtime for Uncore Power Conservation on HPC Systems. In *SC '19: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, November 17–22, 2019, Denver, CO. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3295500.3356150>

1 INTRODUCTION

The Department of Energy (DOE) has set a power budget of 20-30MW on each of the exascale machine sites to maintain a feasible power demand that can be met by operational budget constraints

and power plants. Today's fastest supercomputer, Sunway Taihu-Light, consumes 15.3MW to deliver 93PFlop/s [5]. Exascale supercomputers, which are slated to arrive in 2020, are expected to achieve an exaflop under 20-30MW of power. To achieve this goal, we need to achieve at least an order of magnitude improvement in performance with respect to the state of art systems without exceeding the DOE's power constraint.

Chip manufactures have provided various knobs such as Dynamic Frequency and Voltage Scaling (DVFS), Intel's Running Average Power Limit (RAPL) [11, 28], and software controlled clock modulation [28] that can be used by system software to improve power efficiency of systems. Various solutions [6, 7, 16, 18, 26, 33, 41, 42] have been proposed that use these knobs to achieve power conservation without impacting performance. While these solutions focused on the power efficiency of cores, they were oblivious of the uncore, which is expected to be a growing component in the future generations of processors [34].

Fig. 1 depicts the architecture of a typical Intel server processor. A chip or a package consists of two main components, core and uncore. A core typically consists of the computation units (e.g., ALU, FPU) and the upper levels of caches (L1 and L2) while the uncore contains the last level cache (LLC), the quick path interconnect (QPI) controllers and the integrated memory controllers. With increasing core count and size of LLC and more intelligent integrated memory controllers on newer generations of processors, the uncore occupies as much as 30% of the die area [25], significantly contributing to the processor's power consumption [8, 23, 51]. The uncore's power consumption is a function of its utilization, which varies not only across applications but it also varies dynamically within a single application with multiple phases. We observed that Intel's firmware sets the uncore frequency to its maximum on detecting even the slightest of uncore activity resulting in high power consumption. We can save power replacing such a naive scheme with a more intelligent uncore frequency modulation algorithm.

Figure 1: A processor is a chip consisting of core and uncore. The uncore consists of memory controllers (MC), Quick Path Interconnect (QPI) and the last level cache (LLC).

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.

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/10.1145/3295500.3356150>

Toward this end, we propose Uncore Power Scavenger (UPSCavenger), a runtime system that automatically modulates the uncore frequency to conserve power without significant performance degradation. For applications with multiple phases, it automatically detects phase changes and dynamically resets the uncore frequency for each phase. To the best of our knowledge, UPSCavenger is the first runtime system that focuses on the uncore to improve power efficiency of the system.

To this end, the paper makes the following contributions:

- We explore uncore frequency scaling and its impact on power consumption and performance of various applications.
- We propose Uncore Power Scavenger (UPSCavenger), a runtime system that reduces the power consumption and the energy footprint of applications by automatically modulating the uncore frequency.
- We propose a new light-weight approach for dynamic phase detection that leverages real-time measurements and does not require any prior information about the application.
- We evaluate UPSCavenger with multiple HPC benchmarks and applications on a cluster over 286 cores to show that it conserves power and achieves significant energy savings compared to the default configuration. Our evaluations also show that UPSCavenger achieves performance improvement over Intel's RAPL under equivalent power constraints.

The paper is organized as follows. Section 2 introduces uncore frequency scaling and analyze its impact on performance. Section 3 discusses the design of UPSCavenger, the proposed runtime system. Section 4 describes the implementation and the experimental setup. Section 5 presents the evaluation of UPSCavenger with multiple applications and benchmarks. To assess energy savings we compare UPSCavenger with the default configuration, i.e. without explicitly setting the uncore frequency, while for performance we compare UPSCavenger with Intel's RAPL mechanism. Section 7 summarizes the contributions.

2 OVERVIEW

From Haswell processors onward, Intel has introduced uncore frequency scaling (UFS) that can be used to modulate the frequency of the uncore independent of the core's operating state. UFS is exposed to software via a model-specific register (MSR) addressed at 0x620. Bits 15 through 8 of this MSR encode the minimum uncore frequency multiplier while bits 7 to 0 encode the maximum uncore frequency multiplier. The product (frequency multiplier $\times$ 100MHz) gives us the frequency. For our architecture the minimum and the maximum uncore frequency multipliers are 12 and 27, respectively. We use the msr-safe [48] kernel module to read from and write to the UFS MSR.

2.1 Single Socket Performance Analysis of UFS

We conducted our experiments on Broadwell nodes. Each node has two sockets. Each socket hosts one Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz processor. Each processor has 8 cores. Each node has 128GB memory. While we conducted performance analysis of a wide range of workloads, in this section we present datasets for three representative workloads that have distinct performance

characteristics, viz., Embarrassingly Parallel (EP), the Block Tri-diagonal solver (BT) and the Multi-grid solver (MG) from the NAS parallel benchmark [4] suite. EP is a compute-intensive benchmark, BT is a last level cache-bound benchmark and MG is a memory-bound benchmark. We used the MPI version of these codes in our experiments. The uncore frequency was modulated from 2.7GHz to 1.2GHz in steps of 0.1GHz. We report performance in terms of job completion time in seconds, average package power and DRAM power in Watts, memory bandwidth and LLC misses per second. Each experiment was repeated five times. We report averages across five runs.

Figures 2(a), 2(b) and 2(c) depict the effects of uncore frequency scaling on EP, BT and MG, respectively. The x-axis represents uncore frequency (from high to low). The default uncore frequency is 2.7GHz. Each figure presents five datasets, viz., completion time in seconds, package power and DRAM power reported by Intel's Running Average Power Limit (RAPL) [28] registers, memory bandwidth and LLC misses per second. The y-axis represents data normalized with respect to the default configuration (UCF=2.7GHz). Table 1 shows the raw data for EP, BT and MG at 2.7GHz.

Table 1: Metrics for EP, BT and MG at 2.7GHz

Metric	EP	BT	MG
PKG power	40W	46W	50W
DRAM power	19W	23W	54W
Memory Bandwidth [MB/s]	8	3983	37327
LLC misses per second	0.003M	3M	33M

Observation 1: UFS leads to (package) power savings without significant performance degradation for CPU-bound applications.

This is illustrated by Figure 2(a) showing data for EP, a representative CPU-bound benchmark. EP uses L1 and L2 caches aggressively but otherwise has a negligible memory footprint. For EP, package power decreases linearly with the uncore frequency. However, completion time and DRAM power remain constant. This is because EP neither uses the LLC nor the memory controllers heavily. The reduction in package power can be attributed to the uncore's idle power consumption.

Observation 2: UFS reduces power for cache-bound applications but aggressive frequency reduction may lead to performance degradation.

This is illustrated by Figure 2(b) showing data for BT, a cache-bound application. It uses LLC aggressively. With the default configuration, the LLC cache miss rate is only 3M per second as most of its LLC accesses are hits and memory bandwidth consumption is 3.9GB/s. Package power decreases linearly with the uncore frequency. However, completion times increase slightly at lower uncore frequencies. Reducing the uncore frequency leads to fewer LLC misses per second and lower memory bandwidth which leads to nominal performance degradation (slowdown) only beyond 1.5GHz. It is important to note that the DRAM power remains constant.

Observation 3: UFS leads to significant performance degradation for memory-bound applications.

This is illustrated by Figure 2(c) showing data for MG, a memory-bound application. With the default configuration, its LLC miss rate is 33M per second while it consumes 37GB/s ($\approx 10\times$ that of BT) of memory bandwidth. Consistent with previous observations,

Figure 2: Effects of UFS on EP, BT and MG.

even for MG package power decreases linearly with the uncore frequency. The completion time increases linearly with the uncore frequency under 2.1GHz. Reducing the uncore frequency leads to slower LLC and slower memory controllers. Slower memory controllers send memory requests at a lower rate to DRAM leading to lower DRAM power consumption as shown in the figure. Unlike EP and BT, DRAM power reduces with the uncore frequency beyond 2.1GHz for MG. This decline in DRAM power is also coupled with significant performance degradation beyond 2.1GHz reaffirming our previously stated inference.

From observations 1-3, we conclude the following:

- The decline in DRAM power as a result of lowering the uncore frequency is an indication of potentially significant performance degradation.
- The uncore frequency can be reduced safely without a significant impact on performance only up to the point until which DRAM power remains unchanged.

Observation 4: Applications consist of multiple distinct phases spanning across the spectrum from compute-intensive to memory-intensive.

This is illustrated by Figure 3 depicting the DRAM power profile and the IPC profile of MiniAMR [44]. The x-axis represents the timeline while the y-axis represents data normalized to each metric at maximum frequency. Five distinctly identifiable phases, P1-P5, are annotated in the plot. P1-P5 are short compute-intensive phases characterized by low DRAM power and high IPC occurring between relatively long memory-intensive phases characterized by high DRAM power and low IPC. Hence, the phase transitions can be detected as follows:

- The transition from a compute-intensive phase to a memory-intensive phase can be identified by a rise in DRAM power.
- The transition from a memory-intensive phase to a compute-intensive phase can be identified by a decline in DRAM power.

Figure 3: Phases in MiniAMR

PHASE CHANGE DETECTION

Let ΔIPC be the change in IPC and $\Delta DRAM_power$ be the change in DRAM power during a fixed time interval. Algorithm 1 depicts the pseudo code for phase change detection.

Hypothesis: Static UFS is sub-optimal for applications with distinct phases.

By combining observations 1-4, we hypothesize that for an application consisting of multiple phases with distinct uncore utilization, DRAM access rates, and CPU intensity, statically setting a single uncore frequency for the entire execution is sub-optimal. Hence, we need a dynamic approach that automatically detects phase changes and determines the best uncore frequency for each phase.

While techniques like power gating can be leveraged to save power for phases with zero uncore utilization [19, 35, 37, 43], i.e., while all processes are suspended, uncore frequency scaling achieves power conservation during phases with non-zero but low uncore utilization, which is often the case for HPC applications. Using techniques like sleep states requires prediction of the length of the zero utilization phases to evaluate the trade-offs of the overheads of entering and exiting from a sleep state. Uncore frequency scaling is free of such overheads. Hence, we use uncore frequency scaling in our proposed approach.

ALGORITHM 1. *Phase Change Detection*

```

1: procedure DETECT_PHASE_CHANGE
2:   if  $\Delta \text{DRAM\_power} == 0$  then
3:     no phase change detected
4:   else if  $\Delta \text{DRAM\_power} > 0$  then
5:     phase change detected
6:     compute-intensive to memory-intensive
7:   else  $\triangleright \Delta \text{DRAM\_power} < 0$ 
8:     phase change detected
9:     memory-intensive to compute-intensive
10:  end if
11: end procedure

```

end

3 UNCORE POWER SCAVENGER (UPSCAVENGER)

We propose UPSCavenger, a power-aware runtime system that aims at conserving power by dialing down the uncore's operating frequency opportunistically without significant performance degradation. Figure 4 depicts the high-level architecture of UPSCavenger and the mapping of its components to the components of an HPC system. Figure 4(a) shows the architecture of a typical HPC cluster. It consists of multiple compute server nodes. Each server node hosts two sockets with a single 12-core processor. Multiple parallel jobs spanning across one or more nodes can run on a cluster, e.g., job1 runs on the top 4 nodes while job2 runs on bottom 4 ones. An instance of UPSCavenger runs alongside each job.

The high-level design of the UPSCavenger runtime is depicted in Figure 4(b). The UPSCavenger runtime system allocates one UPSCavenger agent to each socket in a job. Each UPSCavenger agent monitors its socket's performance and DRAM power. It dynamically and automatically manages a socket's uncore frequency to conserve power. UPSCavenger agents make asynchronous decisions based only on the local information pertaining to their respective sockets. Hence, UPSCavenger does not require any global synchronization across the individual UPSCavenger agents. The UPSCavenger agent is a closed-loop feedback controller [49]. The general idea of a closed-loop feedback controller is depicted in Figure 5.

A feedback-control loop consists of three stages that are repeated periodically, viz., sensor, control signal calculator, and actuator.

System is the component whose characteristics are to be monitored and actuated by the controller. The current state of the system is defined in terms of the *process variable* (PV). The desired state of the system, $PV=K$, where K is a constant, is called the setpoint (SP). It is input to the feedback controller. A feedback-control consists of three main modules, viz., sensor, control signal calculator, and actuator.

Sensor

The sensor periodically monitors the state of the system defined in terms of process variable. This is indicated by READ_SENSOR.

Control Signal Calculator

The error (err) in the system state is calculated as the difference between the setpoint and the measured process variable. The feedback controller determines the new value for the actuator as a function of this error to drive the system closer to the setpoint.

Actuator

Figure 4: UPSCavenger Overview

Figure 5: Closed-loop Feedback Controller

The actuator applies the calculated signal to the system. This is indicated by ACTUATE_SIGNAL.

3.1 UPSCavenger Agent

Each UPSCavenger agent is a closed-loop feedback controller that performs three tasks. First, it periodically monitors the socket by measuring its performance in terms of instructions per cycle (IPC) and its DRAM power consumption. Second, depending on the measured power and IPC, it determines the uncore frequency that would drive the system closer to the setpoint and actuates. Third, it detects phase changes at runtime and resets the setpoint for every new phase. The architecture of a UPSCavenger agent is depicted in Figure 6.

Figure 6: UPSCavenger Agent

System, Process Variable and Setpoint

The system of a UPSCavenger is a socket consisting of a processor and its local DRAM. The process variable of a UPSCavenger agent is its socket's DRAM power consumption, i.e., $PV = \text{DRAM power}$, measure of socket's DRAM utilization, which again, correlates with the uncore utilization. Prior work [19, 37, 56] has used models that rely on multiple performance counters to determine application characteristics. Our approach has lower overhead than this approach. UPSCavenger relies on a single register (DRAM RAPL MSR) and hence, issues just one MSR request to quantify/determine the state of the system. UPSCavenger automatically determines the setpoint, which is the maximum DRAM power consumption for every phase of the application, i.e., For an application with multiple phases, UPSCavenger agent automatically resets the setpoint on detecting a phase change.

Sensor

The sensor periodically monitors DRAM power and IPC of the socket. This is indicated by `READ_SENSOR`.

Control Signal Calculator

The control signal calculator takes as inputs the socket's DRAM power consumption and IPC measured at runtime. It then determines the current state of the system to decide the next course of action. Table 2 describes the variables used by the control signal calculator. Figure 7 presents the control logic.

During its first invocation (INITIALIZATION), the UPSCavenger agent sets the setpoint, SP_DRAM_power , to the observed DRAM power, PV_DRAM_power . After initialization, the system enters into a periodic control loop. IPC and DRAM power are measured for each sampling interval. $\Delta DRAM_power$ is calculated as $PV_DRAM_power - SP_DRAM_power$.

State 1: Candidate for UCF reduction

If $\Delta DRAM_power$ is zero, it implies that the DRAM power has not declined (from the setpoint for the current phase) as a result of the previous actuation. This puts the system in state 1, i.e., it is a candidate for lowering uncore frequency in the current iteration of the control loop. Therefore, the actuator decrements the uncore frequency and control returns to the monitoring block.

Table 2: Variables in Control Signal Calculation

Variable	Description
ucf	uncore frequency
PV_DRAM_power	DRAM power consumption measured during current sampling interval
SP_DRAM_power	setpoint
$err = \Delta DRAM_power$	$PV_DRAM_power - SP_DRAM_power$
previous_IPC	IPC measured during previous sampling interval
current_IPC	IPC measured during current sampling interval
ΔIPC	$current_IPC - previous_IPC$
MAX_UCF	maximum uncore frequency for the architecture

State 2: Increment UCF

If $\Delta DRAM_power$ is less than zero, it is an indication of a drop in DRAM power from the setpoint for the current phase. This

Figure 7: Control Logic: A State Machine

drop could either be an indication of a phase transition from a memory-intensive to a compute-intensive phase or a detrimental effect of excessive uncore frequency lowering in the previous iteration of the control loop. To identify the cause of the observed drop, we compare IPC. ΔIPC is the difference between IPC during the current and the previous sampling intervals. If ΔIPC is less than zero, it is an indication of performance degradation. Hence, this observation puts the system in state 2. In response to this state, the actuator increments the uncore frequency and control returns to the monitoring block.

State 3: Phase Transitions

If $\Delta DRAM_power$ is less than zero and ΔIPC is greater than or equal to zero, it is an indication of a rise in CPU activity with a decline in DRAM power. This implies that a phase transition from memory-intensive to compute-intensive phase occurred. Hence, this observation puts the system in state 3 (3a).

If $\Delta DRAM_power$ greater than zero, it implies that the measured DRAM power is higher than the setpoint (maximum) for the current phase. This implies that there has been a phase transition from compute-intensive to memory-intensive phase. Hence, this observation puts the system in state 3 (3b).

Actuator. The actuator is responsible for taking action in response to the determined state of the system. This is indicated by `ACTUATE_SIGNAL`. The UPSCavenger agent sets the socket's uncore frequency by writing to the MSR (at 0x620).

State 1: The actuator decrements the uncore frequency to conserve power.

State 2: The actuator increments the uncore frequency to revert the detrimental effects of the previous actuation.

State 3: On detecting a transition to a different phase, the actuator sets the uncore frequency to the maximum value (MAX_UCF) and sets the setpoint (say SP_DRAM_power) to the measured DRAM power (PV_DRAM_power).

UPSCavenger does not need prior information about any phases that occur during application execution. UPSCavenger's IPC and DRAM power-driven phase change detection suffices to automatically detect these phases with distinct behavior (e.g., checkpointing

to detect that is more memory-intensive than a blocked matrix multiplication kernel).

Resetting the Setpoint to maximum DRAM power for a new phase: It is important to note here that in response to both the phase transitions (even in case of M->C) we set the uncore frequency to MAX_UCF. As a result of this, the socket is guaranteed to have the opportunity to draw maximum DRAM power (at MAX_UCF) during the subsequent sampling interval of the new phase. This ensures that UPSCavenger doesn't constrain the performance of the newly detected phase based on the setpoint of the previous phase. If this maximum DRAM power (measured as PV_DRAM_power during the subsequent sampling interval) is greater than SP_DRAM_power', the system enters state 3 again. In this case, SP_DRAM_power is now intentionally set to the maximum DRAM power for the new phase.

After actuation, the control returns to the monitoring block of the control logic.

Sample Interval. We invoke the UPSCavenger agents periodically every 200ms in our experiments. Sampling under 200ms lead to high overheads. 200ms was the shortest sample interval that had negligible overheads of obtaining measurements, setting the uncore frequency while letting the change of frequency reflect in the next sampling interval. We conducted a sensitivity study by varying the sample interval from 200ms to 1s. We observed that our runtime system has negligible overheads at 200ms while a longer sampling interval makes it sluggish because of two reasons. First, it has fewer opportunities of detecting phase changes and modulating the uncore frequency. Second, it leads to longer reaction times. A sampling interval of 200ms is sufficiently long to actuate the uncore frequency modulation and observe its effect and it is sufficiently frequent to have a fast reacting runtime system. This threshold is driven by the overhead of UPSCavenger, not by any assumption about the length of phases in an application. Applications with even shorter phases would not benefit from UPSCavenger as they would not allow enough time to modulate the frequency.

4 IMPLEMENTATION AND EXPERIMENTAL FRAMEWORK

We conducted our experiments on a 108 node cluster of heterogeneous processors with Infiniband interconnect using 16 Broadwell nodes. Each node has two Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz processors and each processor has 8 cores. Typical HPC systems, like ours, provide exclusive node access. Each MPI job runs on uniquely allocated nodes. No two jobs (or applications) share nodes. We use all the cores of every socket in a job. We pin one MPI rank to a unique core on every socket. Typical MPI applications are NUMA-aware, i.e., all the cores access memory on the socket-local DRAM via the local memory channel consuming the local memory bandwidth. We observed that the remote memory bandwidth consumption is negligible and constant over time compared to the local memory bandwidth for all applications. Hence, we can use each socket's DRAM power as an approximation of the socket's uncore utilization.

UPSCavenger is a light-weight library that is with the application at the time of compilation. It is implemented leveraging the MPI standard profiling interface [38] (PMPI). The UPSCavenger agents

are initiated in the wrapper function of the MPI_Init MPI call, which is the first MPI function call invoked by an MPI application. We use signals to set timers for periodic invocation of UPSCavenger agents on every socket in a job. The system is restored to the default state (maximum uncore frequency multiplier is set to 27) in the wrapper function of the MPI_Finalize MPI call, which is the final MPI call invoked by an MPI application indicating its completion. As UPSCavenger is a shared library that can be linked to the application and it leverages the PMPI interface for initialization and termination, no modifications to the application's source code are required to use UPSCavenger. Applications only need to be linked to our library by modifying their Makefile at the time of compilation.

At each invocation, a UPSCavenger agent periodically monitors the state of the system by measuring DRAM power consumption and IPC. For the DRAM power measurement, we leverage Intel's Running Average Power Limit (RAPL) [28]. From Sandy Bridge processors onward, Intel supports this interface that allows the programmer to measure (or constrain) the power consumption of a processor/socket/package (PKG) as well as the DRAM by reading from (or writing to) the RAPL MSR. We also use RAPL in some of our experiments to constrain the power consumption of the processors. RAPL has a power capping mechanism implemented in hardware that guarantees that the power consumption does not exceed the power limit specified by the user. The UPSCavenger agent measures IPC by reading the fixed counters [28].

Depending on the control calculations, the UPSCavenger agent modulates the uncore frequency by writing to the bits 0 to 7 of the UFS MSR addressed at 0X620 that encode the maximum uncore frequency multiplier. The minimum uncore frequency multiple is set to 12, as 1.2GHz is the minimum uncore frequency on our platform. The msr-safe kernel module installed on this cluster enabled us to read from and to write to Intel's RAPL and UFS MSRs and fixed counters indirectly via the libmsr library [48].

In Figure 7, we compare Δ DRAM_power and Δ IPC to absolute zero. However, in the implementation we use a realistic error margin of $\pm 5\%$ to avoid triggering the UPSCavenger agent for just minor fluctuations in the measurements.

UPSCavenger agents in a job do not need any global synchronization or communication with other agents. They monitor their respective sockets and make local decisions based on local knowledge about the system. Hence, UPSCavenger is a very light weight runtime system with only nominal overheads of reading and writing to MSRs that are several orders of magnitude lower than the total completion time of a typical HPC job.

5 EXPERIMENTAL EVALUATION ON A MULTI-NODE CLUSTER

For evaluation we used Embarrassingly Parallel (EP), Scalar Pentadiagonal solver (SP), Block Tri-diagonal solver (BT), Multi-Grid (MG) and Fourier Transform (FT) from the NAS parallel benchmark suite [4] and CoMD and MiniAMR from the ECP proxy applications suite [1]. We used input sizes C and D, 200x200x200 and 100X100x100 for NPB benchmarks, CoMD and miniAMR, respectively. We used the MPI version of these codes in our experiments. Our selection of applications ensures that we evaluate UPSCavenger

Figure 8: Package and DRAM power savings achieved with UPSCavenger and the resulting slowdowns and energy savings with respect to the baseline

with applications spanning across the spectrum from compute-intensive to memory-intensive. For example, EP is a CPU-bound application that uses uncore minimally, SP and BT use uncore (LLC) intensively with moderate memory (DRAM) access rates while other applications have high memory (DRAM) access rates. We report average package power savings over the duration of the job and slowdowns caused by the UPSCavenger runtime system with respect to the baseline. We chose the default configuration, where Intel’s firmware modulates the uncore frequency, as the

baseline. We also compare an application’s performance with UPSCavenger to its performance with Intel’s Running Average Power Limit (RAPL) with equal power consumption. Each experiment was repeated five times. We report averages across all repetitions with standard deviation of 1-5%.

Figure 8 depicts power savings achieved with UPSCavenger, resulting slowdowns and energy savings with respect to the baseline. We show results for 64, 128 and 256 core experiments as shown in the legend. X-axes represent the benchmarks and applications. Y-axes represent (top to bottom) package and DRAM power savings, slowdown and energy saving with respect to the baseline.

Observation 1: UPSCavenger achieves up to 11% package power savings.

For applications like EP, MiniAMR, BT and SP that have longer CPU-bound phases than others (such as MG, FT, CoMD), UPSCavenger achieves higher package power savings. These codes keep the cores busy and incur fewer LLC misses per second. As a result, they are not as dependent on long latency memory accesses and the uncore as the rest. Hence, the uncore frequency can be lowered for the CPU-intensive phases to save power.

Observation 2: UPSCavenger has a marginal impact on DRAM power.

This is a desired effect. UPSCavenger is designed to reduce uncore frequency so long as it does not impact DRAM utilization. We measure DRAM utilization in terms of DRAM power. Hence, a substantial decline in DRAM power is an indication of decline in DRAM utilization (unless a phase change occurs) which can lead to performance degradation.

Observation 3: UPSCavenger causes worst-case slowdown of 5.5%. In some cases, UPSCavenger achieves power savings at the cost of marginal slowdown for a few applications. This includes the overhead of the runtime system, i.e., measuring DRAM and PKG power, IPC, calculating the feedback and setting appropriate uncore frequency. The observed performance degradation is under 5.5% across all codes with some codes experiencing no slowdown. The reported overhead includes the effect of throttling uncore frequency momentarily below the phase’s optimal frequency before UPSCavenger detects performance degradation and recovers by incrementing it.

Observation 4: UPSCavenger achieves up to 14% energy savings. UPSCavenger achieves energy savings of at least 5% across all applications and core counts, despite the slowdown incurred in some cases. UPSCavenger conserves energy by simply reducing the uncore’s energy usage when it is underutilized, thus reducing energy wastage and shifting towards more energy efficient execution.

Observation 5: Intel’s default scheme consists of a naive algorithm that sets the uncore frequency to its maximum.

Intel’s default algorithm sets uncore frequency to 2.7GHz on detecting even the slightest uncore activity. Figure 9 depicts the uncore frequency profile for runs of EP with default configuration (top) and UPSCavenger (bottom). X-axes represent the timeline while y-axes represent the uncore frequency. EP is a CPU-intensive embarrassingly parallel benchmark that does not use the uncore at all. It operates on core-local cache and incurs nominal LLC misses per second. Even in case of such a workload, the default scheme sets the uncore frequency to 2.7GHz, which is the maximum for our specific architecture.

Figure 9: Uncore frequency profiles for EP with UPSCavenger and the default configuration.

Unlike this scheme, UPSCavenger intelligently modulates uncore frequency by making measurement-driven decisions at runtime based on DRAM power and IPC achieved by the socket. Our automatic and dynamic algorithm leads to a constant uncore frequency of 2GHz for the entire duration of EP. For other applications like CoMD and MiniAMR, the uncore frequency continuously changes.

Figure 10 and 11 depict job profiles of MiniAMR and CoMD, respectively, for 4 node (8 socket) runs. We show two sets of plots (left and right). Plots on the left are the job profiles with the default configuration (baseline) while those on the right are with UPSCavenger. X-axes represent the timeline while y-axes represent the uncore frequency, DRAM power, package power (top to bottom) for each of the sockets of the job. Each socket is represented by a unique color as depicted in the legend. With the default configuration, all sockets operate at the maximum uncore frequency (2.7GHz) constantly for both MiniAMR and CoMD. Hence, the uncore frequency profiles of all the sockets coincide. However, with UPSCavenger, the UPSCavenger agents dynamically modulate the uncore frequency of sockets.

Observation 6: UPSCavenger detects phase change and increases the uncore frequency to UCF_MAX leading to spikes in package power. UPSCavenger detects phase changes based on deflection in DRAM power and IPC. Figure 10 shows the DRAM power profile for MiniAMR. At least five phase changes indicated by dips in DRAM power can be observed. When each of these phase changes occurs, UPSCavenger agents automatically detect it and reset the uncore frequency of their corresponding sockets to UCF_MAX. This leads to a spike in the package power consumption as the uncore consumes higher power when operating at a higher uncore frequency. Even with the default configuration, the power consumption varies across packages in a job. This is a result of manufacturing variation across processors [20, 27].

Observation 7: With UPSCavenger, packages spend more time at lower power than with the default configuration. After the initial spike in package power for each newly detected phase, UPSCavenger gradually lowers the uncore frequency to the lowest level before it starts affecting DRAM power. At this low uncore frequency, the packages consume less power than at UCF_MAX (depicted in PKG power plots in Figure 10 and Figure 11), thus leading to power savings.

Observation 8: Even in case of codes with frequent phase changes, UPSCavenger saves power.

For applications such as CoMD depicted in Figure 11 that incur frequent phase changes with short-lived phases, UPSCavenger does not stabilize at a constant lower uncore frequency for an extended duration. Instead, it intelligently and automatically modulates the uncore frequency between several (higher and lower) levels. As a result, packages spend some time at frequencies lower than that selected by the default scheme. This leads to overall power savings over baseline as shown in Figure 8.

Observation 9: UPSCavenger outperforms Intel's RAPL.

Intel's RAPL is a power capping mechanism implemented in hardware that restricts the power consumption of a processor to a value specified in the RAPL MSR. Figure 12 depicts the comparison of UPSCavenger with RAPL. We first recorded the completion time and the resulting power consumption (say P Watts) of an application with UPSCavenger. We then measured the application's performance without UPSCavenger under a P Watt power bound enforced by restricting power consumption of the processors in the job using RAPL. Figure 12 depicts the total (package and DRAM) power savings obtained by UPSCavenger, resulting performance improvements and energy savings with respect to RAPL. We show results for 64, 128 and 256 core experiments (see legend). X-axes represent the benchmarks and applications. Y-axes represent (top to bottom) total power savings, speedup and energy saving with respect to RAPL. The power consumption with UPSCavenger is within $\pm 2\%$ of RAPL. Interestingly, UPSCavenger achieves up to 20% speedup over RAPL leading to proportional energy savings. This indicates that even for a power-constrained computing paradigm that aims at maximizing performance under a strict power budget, UPSCavenger outperforms RAPL.

Observation 10: In a power-constrained environment, UPSCavenger deducts power from the uncore without slowing the cores when possible while Intel's RAPL always deducts power from the cores.

This is illustrated in Figure 13, which depicts the effective per core frequency achieved with RAPL and with UPSCavenger for equivalent power consumption of BT. We first recorded the effective core frequency measured via Intel's APERF and MPERF MSRs [28] and the resulting power consumption (say P Watts) of an application with UPSCavenger. The effective frequency is calculated as $\frac{\Delta APERF * BASE_FREQUENCY}{\Delta MPERF}$, where the base frequency on our architecture is 2.1GHz. We then measured the effective core frequency without UPSCavenger under a P Watt power bound enforced by restricting power consumption of the processors in the job using RAPL. The x-axis represents the timeline while the y-axis represents the effective core frequency.

RAPL lowers the core frequency from 2.3GHz down to 1.8GHz to save power, which leads to performance degradation. UPSCavenger achieves equivalent power savings by opportunistically reducing the uncore frequencies without affecting the core frequencies allowing the cores to operate at the maximum frequency of 2.3GHz.

6 RELATED WORK

Chip manufactures have provided various knobs to modulate power directly, e.g., power capping, or indirectly, e.g., dynamic voltage

Figure 10: Power and uncore frequency profiles for 8 socket runs of MiniAMR with default configuration (left) and UPSCavenger (right).

Figure 11: Uncore frequency profiles for 8 socket runs of CoMD with default configuration (left) and UPSCavenger (right).

frequency scaling (DVFS), of various components of a server. Conventionally, DVFS has been used to modulate the frequency of cores [28]. Intel processors prior to Haswell maintained all cores at a common operating state, i.e., voltage and frequency. Starting with Haswell, Intel introduced per core DVFS. From Sandy Bridge processors onward, Intel introduced running average power limit (RAPL) [11, 28], a power capping mechanism implemented in hardware that constrains the average collective power consumption of the cores and the uncore of a processor. A plethora of solutions [6, 16–18, 20, 21, 26, 33, 41, 42, 50] have been proposed that leverage these knobs to improve the power efficiency of a system. DVFS-based solutions were oblivious of the power consumption of the uncore, which is expected to be a growing component in

the future generations of processors [34]. RAPL-based solutions off-loaded the management of the uncore’s power to Intel’s RAPL implementation, which we show is sub-optimal.

Early work exploited DVFS to reduce CPU frequencies during idle time, e.g., due to early arrival at MPI barriers and collectives [30]. An ILP-based approach to model energy [41] was demonstrated to be effective during the runtime of MPI codes to determine optimal power levels for the cores [42]. These works were trying to conserve energy without sacrificing performance by much when utilizing just one core of a node resulting in underutilization of the system. Memory DVFS [10] was also proposed as a technique to conserve energy for workloads that are not constrained in memory bandwidth. UPSCavenger differs in that its foremost objective is

Figure 12: Package and DRAM power savings, speedups and energy savings achieved by UPSCavenger with respect to RAPL

to reduce the inefficient power consumption of the uncore of each processor of an HPC job without reducing system utilization, i.e., using all cores of a processor. It conserves power and achieves energy savings by dynamically modulating the frequency of the uncore based on real-time feedback from the system.

Prior work [9, 32] leveraged the effect of concurrency throttling and thread locality to save power and increase performance. An ILP-based runtime approach has been shown to determine how many cores an application should be run on to stay within a given power budget [31, 55]. Marathe et al. [36] proposed Conductor, a runtime system that speeds up an application's critical path through an adaptive socket power-allocation algorithm that periodically performs dynamic voltage frequency scaling (DVFS) and dynamic concurrency throttling (DCT) based on application behavior and

Figure 13: Effective core frequency profiles for BT with RAPL and UPSCavenger for equal power consumption.

power usage. Our work differs in terms of granularity and adaptivity. First, we reduce uncore power across resources at processor chip level, exploiting dynamically adaptive feedback methods. Second, UPSCavenger does not modify the concurrency or the thread locality, which can be tightly coupled with the application's input.

In recent work, Gholkar et al. [20] presented PTune, a process variation-aware power tuner that uses performance characterization data for all sockets on a cluster and application execution characteristics to minimize the runtime of a job under its power budget. It assumes that the jobs are moldable, i.e., the number of MPI ranks can be varied at the time of job scheduling. UPSCavenger does not require jobs to be moldable. It also does not require any prior characterization data to make uncore frequency decisions at runtime. In [21], they proposed PShifter, a feedback-based mechanism that shifts power from processors idling at barriers and collectives to the processors on the critical path to improve performance under a power constraint. PShifter does not strive to save power. Its control calculations depend on the utilization of the cores while being oblivious of the uncore's utilization. Our approach is a dynamic closed-loop feedback controller that specifically monitors the uncore and depending on its utilization reduces any wasteful power consumption by modulating the uncore frequency at runtime. Ellsworth et al. proposed Pow, a system-wide power manager that re-distributed the power budget across a system via scheduling while adhering to a global machine-level power cap [13]. This approach salvaged the unused power but did not address inefficient power usage. Our approach goes one step further by detecting and reducing the wasteful power consumption within a job.

Capacity-improving schemes that increase job throughput have been developed under power limitations by exploiting "hardware overprovisioning", i.e., by deploying more nodes that will be powered at a time [14, 15, 39, 45, 46]. In such a system, the characteristics of a code under strong scaling were used to calculate the optimal number of processors considering core and memory power [47]. Modifications to the batch scheduler in how small jobs are backfilled

depending on their power profile can further increase job capacity [40]. While these solutions aimed at achieving performance or throughput improvement while utilizing as much of the power budget as possible, our approach aims at conserving power by reducing wasteful power consumption with marginal impact on performance. We do not explicitly aim at *maximizing* the machine's throughput, which is beyond the scope of this paper, but our runtime reduces the overall power and energy footprint of the jobs making more power available for more new jobs to be scheduled on an overprovisioned system.

Hackenberg et al. [24] observed that the uncore frequency depends on the core frequency even when there is no uncore activity. Unlike UPSCavenger, which dynamically modulates the uncore frequency using the measurement-driven closed-loop feedback controller, Sundriya et al. [52, 53] propose a runtime approach that relies on a static model for estimating the frequency. More recent work [54] proposed a dynamic core and uncore frequency scaling strategy that partially relied on regression models using prior data to make dynamic decisions. UPSCavenger is independent of such dependencies and overheads of prior characterization. Von et al. [57] proposed a neural network-based uncore frequency scaling approach that was simulated for chip multi-core (CMP) power management. Da-Cheng et al. [29] simulated a semi-supervised reinforcement learning based approach for dynamic frequency scaling to maximize performance under a power budget. Unlike this work, we implement and evaluate our runtime system on real hardware, i.e., a Broadwell cluster. Our runtime is suitable for power management of a multi-node and a multi-core system.

Our UPSCavenger work is unique in that it is the first of its kind that targets the otherwise neglected uncore component of a chip. It dynamically modulates the frequency of the uncore depending on its utilization measured at runtime. It automatically detects new phases within an application and resets the uncore frequency for each phase. Proposed frameworks like Redfish, the PowerAPI, and Intel's GEOPM [2, 3, 12, 22] can integrate UPSCavenger as a unique closed-loop feedback-based policy for job power management on a cluster. UPSCavenger will also relieve the application developers of the burden to explicitly indicate phase changes as required by the APIs (like GEOPM) as UPSCavenger automatically detects phases without any explicit information from the developer.

7 SUMMARY

We explored uncore frequency scaling and its impact on performance, power and energy consumption of various HPC applications. To the best of our knowledge, this is the first study of uncore frequency scaling conducted on Broadwell processors. We proposed UPSCavenger, a runtime system that automatically modulates the uncore frequency to conserve power without significant performance degradation. As a part of UPSCavenger, we also introduced an algorithm that automatically detects phase changes at runtime. Our evaluations indicate that UPSCavenger achieves up to 10% energy reduction with less than 1% slowdown. It achieves up to 14% energy savings with a worst case slowdown of 5.5% compared to the default configuration. We also show that UPSCavenger achieves up to 20% speedup with proportional energy savings compared to Intel's RAPL with equivalent power usage. While UPSCavenger

can be easily deployed as a runtime on any HPC system, we intend to collaborate with semiconductor manufacturers to incorporate this into their firmware.

8 ACKNOWLEDGEMENTS

We extend our thanks to Livermore Computing for facilitating Dedicated Access Time on the LLNL cluster. This work was supported in part by NSF grants 1525609, 1058779, 1217748, 0958311, and by Consortium for Advanced Simulation of Light Water Reactors (CASL). This material is also based upon work supported by the U.S. Department of Energy's Lawrence Livermore National Laboratory, Office of Science, under Award number DE-AC52-07NA27344 and supported by Office of Science, Office of Advanced Scientific Computing Research (LLNL-CONF-656877).

REFERENCES

- [1] Ecp proxy apps suite. <https://proxyapps.exascaleproject.org/ecp-proxy-apps-suite/>.
- [2] Panel on power api/redfish/geopm. https://eehpcwg.llnl.gov/documents/conference/sc16/SC16_Laros_Eastep_Benson_API.pdf.
- [3] Redfish api. <https://www.dmtf.org/standards/redfish/>.
- [4] NASA Advanced Supercomputing Division, NAS Parallel Benchmark Suite v3.3. 2006. <http://www.nas.nasa.gov/Resources/Software/npb.html>.
- [5] Top500 Supercomputer Sites. 2017. <https://www.top500.org/lists/2017/11/>.
- [6] P. E. Bailey, A. Marathe, D. K. Lowenthal, B. Rountree, and M. Schulz. Finding the limits of power-constrained application performance. In SC, 2015.
- [7] S. Bhalachandra, A. Porterfield, S. L. Olivier, and J. F. Prins. An adaptive core-specific runtime for energy efficiency. In *IPDPS*, pages 947–956, May 2017.
- [8] H.-Y. Cheng, J. Zhan, J. Zhao, Y. Xie, J. Sampson, and M. J. Irwin. Core vs. uncore: The heart of darkness. In *Proceedings of the 52Nd Annual Design Automation Conference, DAC '15*, pages 121:1–121:6, New York, NY, USA, 2015. ACM.
- [9] M. Curtis-Maury, J. Dzierwa, C. D. Antonopoulos, and D. S. Nikolopoulos. Online power-performance adaptation of multithreaded programs using hardware event-based prediction. In *Supercomputing*, pages 157–166, 2006.
- [10] H. David, C. Fallin, E. Gorbato, U. R. Hanebutte, and O. Mutlu. Memory power management via dynamic voltage/frequency scaling. In *Proceedings of the 8th ACM International Conference on Autonomic Computing, ICAC '11*, pages 31–40, New York, NY, USA, 2011. ACM.
- [11] H. David, E. Gorbato, U. R. Hanebutte, R. Khanna, and C. Le. RAPL: Memory Power Estimation and Capping. In *Proceedings of the 16th ACM/IEEE international symposium on Low power electronics and design, ISLPED '10*, pages 189–194, New York, NY, USA, 2010. ACM.
- [12] J. Eastep, S. Sylvester, C. Cantalupo, B. Geltz, F. Ardanaz, A. Al-Rawi, K. Livingston, F. Keceli, M. Maiterth, and S. Jana. Global extensible open power manager: A vehicle for hpc community collaboration on co-designed energy management solutions. In J. M. Kunkel, R. Yokota, P. Balaji, and D. Keyes, editors, *High Performance Computing*, pages 394–412, Cham, 2017. Springer International Publishing.
- [13] D. A. Ellsworth, A. D. Malony, B. Rountree, and M. Schulz. Pow: System-wide dynamic reallocation of limited power in hpc. In *High-Performance Parallel and Distributed Computing*, pages 145–148, 2015.
- [14] M. E. Femal and V. Freeh. Boosting data center performance through non-uniform power allocation. In *International Conference on Autonomic Computing*, 2005.
- [15] M. E. Femal and V. W. Freeh. Safe overprovisioning: using power limits to increase aggregate throughput. In *International Conference on Power-Aware Computer Systems*, December 2005.
- [16] V. Freeh, F. Pan, N. Kappiah, and D. K. Lowenthal. Using multiple energy gears in mpi programs on a power-scalable cluster. In *PPoPP*, pages 164–173, June 2005.
- [17] V. Freeh, F. Pan, N. Kappiah, D. K. Lowenthal, and R. Springer. Exploring the energy-time tradeoff in mpi programs on a power-scalable cluster. In *IPDPS*, May 2005.
- [18] R. Ge, X. Feng, W. c. Feng, and K. W. Cameron. Cpu miser: A performance-directed, run-time system for power-aware clusters. In *2007 International Conference on Parallel Processing (ICPP 2007)*, pages 18–18, Sept 2007.
- [19] P. Gepner and M. F. Kowalik. Multi-core processors: New way to achieve high system performance. In *International Symposium on Parallel Computing in Electrical Engineering (PARELEC'06)*, pages 9–13, Sept 2006.
- [20] N. Gholkar, F. Mueller, and B. Rountree. Power tuning hpc jobs on power-constrained systems. In *PACT*. ACM, 2016.
- [21] N. Gholkar, F. Mueller, B. Rountree, and A. Marathe. Pshifter: Feedback-based dynamic power shifting within hpc jobs for performance. In *HPDC*. ACM, 2018.

- [22] R. E. Grant, M. Levenhagen, S. L. Olivier, D. DeBonis, K. T. Pedretti, and J. H. L. III. Standardizing power monitoring and control at exascale. *Computer*, 49(10):38–46, Oct 2016.
- [23] V. Gupta, P. Brett, D. Koufaty, D. Reddy, S. Hahn, K. Schwan, and G. Srinivasa. The forgotten 'uncore': On the energy-efficiency of heterogeneous cores. In *Proceedings of the 2012 USENIX conference on Annual Technical Conference*, June 2012.
- [24] D. Hackenberg, R. SchÄune, T. Ilsche, D. Molka, J. Schuchart, and R. Geyer. An energy efficiency feature survey of the intel haswell processor. In *2015 IEEE International Parallel and Distributed Processing Symposium Workshop*, pages 896–904, May 2015.
- [25] D. L. Hill, D. Bachand, S. Bilgin, R. Greiner, P. Hammarlund, T. Huff, S. Kulick, and R. Safranek. The uncore: A modular approach to feeding the high-performance cores. In *Intel Technology Journal*, volume 14, 2010.
- [26] C.-h. Hsu and W.-c. Feng. A power-aware run-time system for high-performance computing. In *Supercomputing*, page 1, 2005.
- [27] Y. Inadomi, T. Patki, K. Inoue, M. Aoyagi, B. Rountree, M. Schulz, D. Lowenthal, Y. Wada, K. Fukazawa, M. Ueda, M. Kondo, and I. Miyoshi. Analyzing and mitigating the impact of manufacturing variability in power-constrained supercomputing. In *SC*, pages 78:1–78:12, 2015.
- [28] Intel. Intel-64 and IA-32 Architectures Software Developer's Manual, Volumes 3A and 3B: System Programming Guide. 2011.
- [29] D.-C. Juan and D. Marculescu. Power-aware performance increase via core/uncore reinforcement control for chip-multiprocessors. In *Proceedings of the 2012 ACM/IEEE International Symposium on Low Power Electronics and Design*, pages 97–102, 2012.
- [30] N. Kappiah, V. Freeh, and D. K. Lowenthal. Just in time dynamic voltage scaling: Exploiting inter-node slack to save energy in mpi programs. In *SC*, Nov 2005.
- [31] A. Langer, E. Toton, U. S. Palekar, and L. V. Kalé. Energy-efficient computing for hpc workloads on heterogeneous manycore chips. In *Proceedings of Programming Models and Applications on Multicores and Manycores*. ACM, 2015.
- [32] D. Li, B. R. de Supinski, M. Schulz, K. W. Cameron, and D. S. Nikolopoulos. Hybrid mpi/openmp power-aware computing. In *IPDPS*, volume 10, pages 1–12, 2010.
- [33] M. Y. Lim, V. Freeh, and D. K. Lowenthal. Adaptive, transparent frequency and voltage scaling of communication phases in mpi programs. In *SC*, 2006.
- [34] G. H. Loh. The cost of uncore in throughput-oriented many-core processors. In *In Proc. of Workshop on Architectures and Languages for Troughput Applications (ALTA)*.
- [35] K. Ma and X. Wang. Pgcapping: Exploiting power gating for power capping and core lifetime balancing in cmps. In *Proceedings of the 21st International Conference on Parallel Architectures and Compilation Techniques*, PACT '12, pages 13–22, New York, NY, USA, 2012. ACM.
- [36] A. Marathe, P. E. Bailey, D. K. Lowenthal, B. Rountree, M. Schulz, and B. R. de Supinski. A run-time system for power-constrained hpc applications. In J. M. Kunkel and T. Ludwig, editors, *High Performance Computing*, pages 394–408, Cham, 2015. Springer International Publishing.
- [37] D. Meisner, C. M. Sadler, L. A. Barroso, W.-D. Weber, and T. F. Wenisch. Power management of online data-intensive services. In *Proceedings of the 38th Annual International Symposium on Computer Architecture*, ISCA '11, pages 319–330, New York, NY, USA, 2011. ACM.
- [38] S. Mintchev and V. Getov. Pmpi: High-level message passing in fortran77 and c. In B. Hertzberger and P. Sloot, editors, *High-Performance Computing and Networking*, pages 601–614, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.
- [39] T. Patki, D. K. Lowenthal, B. Rountree, M. Schulz, and B. R. de Supinski. Exploring Hardware Overprovisioning in Power-constrained, High Performance Computing. In *International Conference on Supercomputing*, pages 173–182, 2013.
- [40] T. Patki, D. K. Lowenthal, A. Sasidharan, M. Maiterth, B. Rountree, M. Schulz, and B. R. de Supinski. Practical Resource Management in Power-Constrained, High Performance Computing. In *HPDC*, 2015.
- [41] B. Rountree, D. K. Lowenthal, S. Funk, V. Freeh, B. R. de Supinski, and M. Schulz. Bounding energy consumption in large-scale mpi programs. In *SC*, Nov 2007.
- [42] B. Rountree, D. K. Lowenthal, M. Schulz, V. Freeh, and T. Bletsch. Adagio: Making dvs practical for complex hpc applications. In *ICS*, Nov 2009.
- [43] A. Samih, R. Wang, A. Krishna, C. Maciocco, C. Tai, and Y. Solihin. Energy-efficient interconnect via router parking. In *2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, pages 508–519, Feb 2013.
- [44] Sandia National Laboratory. Mantevo project home page. <https://software.sandia.gov/mantevo>, June 2011.
- [45] O. Sarood. *Optimizing Performance Under Thermal and Power Constraints for HPC Data Centers*. PhD thesis, University of Illinois, Urbana-Champaign, Dec 2013.
- [46] O. Sarood, A. Langer, A. Gupta, and L. V. Kale. Maximizing throughput of overprovisioned hpc data centers under a strict power budget. In *SC*, 2014.
- [47] O. Sarood, A. Langer, L. V. Kale, B. Rountree, and B. de Supinski. Optimizing Power Allocation to CPU and Memory Subsystems in Overprovisioned HPC Systems. In *Proceedings of IEEE Cluster 2013*, Sept 2013.
- [48] K. Shoga, B. Rountree, M. Schulz, and J. Shafer. Whitelisting msrs with msr-safe. In *3rd Workshop on Extreme-Scale Programming Tools at SC*, Nov. 2014.
- [49] C. A. Smith and A. Corripio. Chapter 6: Control systems and their basic components. In *Principles and Practice of Automatic Process Control*, pages 154–192.
- [50] R. Springer, D. K. Lowenthal, B. Rountree, and V. Freeh. Minimizing execution time in MPI programs on an energy-constrained, power-scalable cluster. In *PPoPP*, May 2006.
- [51] B. Subramaniam and W.-c. Feng. Towards energy-proportional computing for enterprise-class server workloads. In *Proceedings of the 4th ACM/SPEC International Conference on Performance Engineering*, ICPE '13, pages 14–26, NY, USA, 2013. ACM.
- [52] V. Sundriyal, M. Sosonkina, B. M. Westheimer, and M. Gordon. Uncore frequency scaling vs dynamic voltage and frequency scaling: A quantitative comparison. In *Society for Modeling & Simulation International*, SpringSim-HPC, 2018.
- [53] V. Sundriyal, M. Sosonkina, B. M. Westheimer, and M. S. Gordon. Comparisons of core and uncore frequency scaling modes in quantum chemistry application GAMESS. In *Proceedings of the High Performance Computing Symposium, SpringSim (HPC) 2018*, Baltimore, MD, USA, April 15-18, 2018, pages 13:1–13:11, 2018.
- [54] V. Sundriyal, M. Sosonkina, B. M. Westheimer, and M. S. Gordon. Core and uncore joint frequency scaling strategy. In *Journal of Computer and Communications*, 6, pages 184–201, 2018.
- [55] E. Toton, A. Langer, J. Torrellas, and L. Kale. Scheduling for hpc systems with process variation heterogeneity. In *Technical Report YCS-2009-443, Department of Computer Science, University of York*, 2014.
- [56] V. M. Weaver and S. A. McKee. Can hardware performance counters be trusted? In *2008 IEEE International Symposium on Workload Characterization*, pages 141–150, Sept 2008.
- [57] J. Y. Won, X. Chen, P. Gratz, J. Hu, and V. Soteriou. Up by their bootstraps: Online learning in artificial neural networks for cmp uncore power management. In *2014 IEEE 20th International Symposium on High Performance Computer Architecture (HPCA)*, pages 308–319, Feb 2014.

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

Hardware : We conducted our experiments on a cluster of 16 Broadwell nodes. Each node has two Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz processors as mentioned in the paper.

Software: Linux operating system (kernel release 4.10.13-1.el7.elrepo.x86_64) was deployed on the cluster. For evaluation, we used NAS Parallel Benchmarks v3.3.1 benchmarks and ECP proxy applications that were compiled using MPI compiler with gcc version 5.4.0.

ARTIFACT AVAILABILITY

Software Artifact Availability: Some author-created software artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

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

Data Artifact Availability: Some author-created data artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

Proprietary Artifacts: No author-created artifacts are proprietary.

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

Software is available on our private git repository
↳ that can be made available upon request.

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz

Operating systems and versions: Linux operating system (kernel release 4.10.13-1.el7.elrepo.x86_64)

Compilers and versions: MPI compiler with gcc version 5.4.0.

Applications and versions: NPB 3.3.1 and ECP Proxy Applications

Libraries and versions: MPICH

Paper Modifications: #msr kernel module was loaded modprobe msr

Output from scripts that gathers execution environment information.

```
MANPATH=/opt/ohpc/pub/mpi/mvapich2-gnu/2.2/man:/opt/
↳ ohpc/pub/compiler/gcc/5.4.0/share/man:/opt/ohpc/
↳ pub/autotools/share/man:/usr/local/share/man:/us
↳ r/share/man/overrides:/usr/share/man/en:/usr/sha
↳ re/man
XDG_SESSION_ID=28875
LC_MONETARY=en_US.UTF-8
HOSTNAME=c94
LMOD_FAMILY_CUDA=cuda
```

```
_ModuleTable003_=LFsiZnVsbE5hbWUiXT0ibXZhcGljaDIvMi4
↳ yIixbImxvYWRPcmRlciJdPTQscHJvcFQ9e30sWyJzaG9ydCJ
↳ dPSJtdmFwaWNoMiIsWyJzdGF0dXMiXT0iYWN0aXZlIix9LG9
↳ ocGM9e1siRk4iXT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWx
↳ lcy9vaHBjIixbImRlZmF1bHQiXT0wLFsiZnVsbE5hbWUiXT0
↳ ib2hwYyIsWyJsb2FkT3JkZXIiXT02LHBByb3BUPXt9LFsic2h
↳ vcnQiXT0ib2hwYyIsWyJzdGF0dXMiXT0iYWN0aXZlIix9LHB
↳ ydW49e1siRk4iXT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWx
↳ lcy9wcnVulzEuMSIsWyJkZWZhdWx0Il09MSxbImZ1bGx0YWY
↳ lI109InBydW4vMS4xIixbImxvYWRPcmRlciJdPTIscHJvcFQ
↳ 9e30sWyJzaG9ydCJdPSJwcnVuIixbInN0YXR1cyJdPSJh
SHELL=/bin/bash
TERM=xterm-256color
HISTSIZE=1000
SSH_CLIENT=10.1.1.2 38196 22
LMOD_DEFAULT_MODULEPATH=/opt/ohpc/admin/modulefiles:
↳ /opt/ohpc/pub/modulefiles
LC_NUMERIC=en_US.UTF-8
LMOD_PKG=/opt/ohpc/admin/lmod/lmod
LMOD_VERSION=6.5.11
SSH_TTY=/dev/pts/0
USER=USER
CUDA_HOME=/usr/local/cuda
```

```

LS_COLORS=r-s=0:di=38;5;27:ln=38;5;51:mh=44;38;5;15:p
i=40;38;5;11:so=38;5;13:do=38;5;5:bd=48;5;232;38
;5;11:cd=48;5;232;38;5;3:or=48;5;232;38;5;9:mi=0
5;48;5;232;38;5;15:su=48;5;196;38;5;15:sg=48;5;1
1;38;5;16:ca=48;5;196;38;5;226:tw=48;5;10;38;5;1
6:ow=48;5;10;38;5;21:st=48;5;21;38;5;15:ex=38;5;
34:*.tar=38;5;9:*.tgz=38;5;9:*.arc=38;5;9:*.arj=
38;5;9:*.taz=38;5;9:*.lha=38;5;9:*.lz4=38;5;9:*.
lzh=38;5;9:*.lzma=38;5;9:*.tlz=38;5;9:*.txz=38;5
;9:*.tzo=38;5;9:*.t7z=38;5;9:*.zip=38;5;9:*.z=38
;5;9:*.Z=38;5;9:*.diz=38;5;9:*.gz=38;5;9:*.lrz=38
;5;9:*.lz=38;5;9:*.lzo=38;5;9:*.xz=38;5;9:*.bz2=
38;5;9:*.bz=38;5;9:*.tbz=38;5;9:*.tbz2=38;5;9:*.
tz=38;5;9:*.deb=38;5;9:*.rpm=38;5;9:*.jar=38;5;9
:*.war=38;5;9:*.ear=38;5;9:*.sar=38;5;9:*.rar=38
;5;9:*.alz=38;5;9:*.ace=38;5;9:*.zoo=38;5;9:*.cp
io=38;5;9:*.7z=38;5;9:*.rz=38;5;9:*.cab=38;5;9:*.
jpg=38;5;13:*.jpeg=38;5;13:*.gif=38;5;13:*.bmp=
38;5;13:*.pbm=38;5;13:*.pgm=38;5;13:*.ppm=38;5;1
3:*.tga=38;5;13:*.xbm=38;5;13:*.xpm=38;5;13:*.ti
f=38;5;13:*.tiff=38;5;13:*.png=38;5;13:*.svg=38;
5;13:*.svgz=38;5;13:*.mng=38;5;13:*.pcx=38;5;13
:*.mov=38;5;13:*.mpg=38;5;13:*.mpeg=38;5;13:*.m2v
=38;5;13:*.mkv=38;5;13:*.webm=38;5;13:*.ogm=38;5
;13:*.mp4=38;5;13:*.m4v=38;5;13:*.mp4v=38;5;13:*.
vob=38;5;13:*.qt=38;5;13:*.nuv=38;5;13:*.wmv=38
;5;13:*.asf=38;5;13:*.rm=38;5;13:*.rmvb=38;5;13
:*.flc=38;5;13:*.avi=38;5;13:*.fli=38;5;13:*.flv=
38;5;13:*.gl=38;5;13:*.dl=38;5;13:*.xcf=38;5;13
:*.xwd=38;5;13:*.yuv=38;5;13:*.cgm=38;5;13:*.emf=
38;5;13:*.axv=38;5;13:*.anx=38;5;13:*.ogv=38;5;1
3:*.ogx=38;5;13:*.aac=38;5;45:*.au=38;5;45:*.fla
c=38;5;45:*.mid=38;5;45:*.midi=38;5;45:*.mka=38;
5;45:*.mp3=38;5;45:*.mpc=38;5;45:*.ogg=38;5;45:*.
ra=38;5;45:*.wav=38;5;45:*.axa=38;5;45:*.oga=38
;5;45:*.spx=38;5;45:*.xspf=38;5;45:

LD_LIBRARY_PATH=/usr/lib64/nvidia:/usr/local/cuda/li
b64:/usr/local/cuda/targets/x86_64-linux/lib:/op
t/ohpc/pub/mpi/mvapich2-gnu/2.2/lib:/opt/ohpc/pu
b/compiler/gcc/5.4.0/lib64
_ModuleTable004=Y3RpdmlU1H0sfSxtcGF0aEE9eyIvb3B0L29
ocGMvcHViL21vZHVzZWRLcHMvZ251LW12YXBpY2gyIiw1L29
wdC9vaHBjL2B1Yi9tb2R1bGVkZXBzL2dudSIsIi9vcHQvb2h
wYy9hZG1pbj9tb2R1bGVmaWxlcysIi9vcHQvb2hwYy9wdWI
vbW9kdWx1ZmlsZXMiLH0sWyJzeXN0ZW1CYXNlTVBBVEgiXT0
iL29wdC9vaHBjL2FkbWluL21vZHVzZWZpbGVzOi9vcHQvb2h
wYy9wdWVibW9kdWx1ZmlsZXMiLFsidmVyc2lubiJdPTIsfQ==
LMOD_PREPEND_BLOCK=normal

PATH=/usr/local/cuda/bin:/opt/ohpc/pub/mpi/mvapich2-
gnu/2.2/bin:/opt/ohpc/pub/compiler/gcc/5.4.0/bin
:/opt/ohpc/pub/prun/1.1:/opt/ohpc/pub/autotools/
bin:/opt/ohpc/pub/bin:/usr/local/sbin:/usr/local
/bin:/usr/sbin:/usr/bin:/USER/.local/bin:/USER/b
in

MAIL=/var/spool/mail/USER

```

```

_ModuleTable001_=X01vZHVszVRhYmxlXz17WyJhY3RpdmdVTAxpJ
↳ 1lI09NixiYXNlTXBhdGhBPXsiL29wdC9vaHBjL2FkbWluL2l
↳ vZHVszZWZpbGVzIiwiL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWx
↳ lcyIsfSxbImNfcmlldWlsZFRpbWUiXT1mYWxzZSxbImNfc2h
↳ vcnRUaW1lI09ZmFsc2UsZmFtaWxz5PXTbIk1QSSJdPSJtdmF
↳ waWNoMiIsWyJjb21waWxlcjJdPSJnbniUiLFsiY3VkySjYdPSJ
↳ jdWRhIix9LGlueYWN0aXZlPxt9LGlUPXthdXRvdG9vbHM9e1s
↳ iRk4iXT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9hdXR
↳ vdG9vbHMiLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJsZSd
↳ hdXRvdG9vbHMiLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJs
↳ bInNob3J0Ii09ImF1dG90b29scyIsWyJzdGF0dXMiXT0i
_=/usr/bin/env
LC_COLLATE=en_US.UTF-8
LMOD_SETTARG_CMD=:
PWD=/USER
XMODIFIERS=@im=none
_LMFILES=/opt/ohpc/pub/modulefiles/autotools:/opt/o
↳ hpc/pub/modulefiles/prun/1.1:/opt/ohpc/pub/modul
↳ efiles/gnu/5.4.0:/opt/ohpc/pub/moduledeps/gnu/mv
↳ apich2/2.2:/opt/ohpc/pub/modulefiles/cuda.lua:/o
↳ pt/ohpc/pub/modulefiles/ohpc
LANG=en_US.UTF-8
MODULEPATH=/opt/ohpc/pub/moduledeps/gnu-mvapich2:/op
↳ t/ohpc/pub/moduledeps/gnu:/opt/ohpc/admin/module
↳ files:/opt/ohpc/pub/modulefiles
LC_MEASUREMENT=en_US.UTF-8
LOADEDMODULES=autotools:prun/1.1:gnu/5.4.0:mvapich2/
↳ 2.2:cuda:ohpc
_ModuleTable_Sz_=4
LMOD_CMD=/opt/ohpc/admin/lmod/lmod/libexec/lmod
HISTCONTROL=ignoredups
HOME=/USER
SHLVL=2
LANGUAGE=
MPI_DIR=/opt/ohpc/pub/mpi/mvapich2-gnu/2.2
_ModuleTable002_=YWN0aXZlIix9LGN1ZGE9e1siRk4iXT0iL29
↳ wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9jdWRhLmx1YSIsWyJ
↳ kZWZhdWx0Ii09MkxibG90YXV1Ii09ImN1ZGEiLFsiZGVmYXV
↳ hZE9yZGVyIi09NSxwcm9wVD17fSxbInNob3J0Ii09ImN1ZGE
↳ iLFsic3RhdHVzIi09ImFjdGZlZSIsfSxbImNfcmlldWlsZFR
↳ iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9nbniUvNS40LjA
↳ iLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJsZSdPSJnbniU
↳ vNS40LjAaIi09ImF1dG90YXV1Ii09Myxwcm9wVD17fSxbInNob3J
↳ 0Ii09ImdudSIsWyJzdGF0dXMiXT0iYWN0aXZlIix9LGl12YXB
↳ pY2gyPXTbIkZ0Ii09Ii9vcHhvdW9wY3RvdWVibW9kdWxlcY9
↳ wcy9nbniUvNXZhcGljaDIvMi4yIixbImRlZmF1bHQiXT0x
BASH_ENV=/opt/ohpc/admin/lmod/lmod/init/bash
LOGNAME=USER
CVS_RSH=ssh
SSH_CONNECTION=10.1.1.2 38196 10.4.1.95 22
MODULESHOME=/opt/ohpc/admin/lmod/lmod
LESSOPEN=||/usr/bin/lesspipe.sh %s
PKG_CONFIG_PATH=/opt/ohpc/pub/mpi/mvapich2-gnu/2.2/l
↳ ib/pkgconfig
LMOD_FULL_SETTARG_SUPPORT=no
LMOD_FAMILY_COMPILER=gnu

```

Uncore Power Scavenger: A Runtime for Uncore Power Conservation on HPC Systems

```
XDG_RUNTIME_DIR=/run/user/0
LMOD_DIR=/opt/ohpc/admin/lmod/lmod/libexec
INCLUDE=/usr/local/cuda/include:/opt/ohpc/pub/compil_
↳ er/gcc/5.4.0/include
LC_TIME=en_US.UTF-8
LMOD_COLORIZE=no
LMOD_FAMILY_MPI=mvapich2
BASH_FUNC_module()=() { eval "$($LMOD_CMD bash "$@");
[ $? = 0 ] && eval "$(${LMOD_SETTARG_CMD:-} -s sh)
}
BASH_FUNC_ml()=() { eval "$($LMOD_DIR/ml_cmd "$@")
}
+ lsb_release -a
collect_environment.sh: line 10: lsb_release: command
↳ not found
+ uname -a
Linux c94 4.10.13-1.el7.elrepo.x86_64 #1 SMP Thu Apr
↳ 27 12:06:06 EDT 2017 x86_64 x86_64 x86_64
↳ GNU/Linux
+ lscpu
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: 2
Core(s) per socket: 8
Socket(s):         2
NUMA node(s):      2
Vendor ID:         GenuineIntel
CPU family:        6
Model:            79
Model name:        Intel(R) Xeon(R) CPU E5-2620 v4
↳ @ 2.10GHz
Stepping:          1
CPU MHz:           2100.000
CPU max MHz:       2101.0000
CPU min MHz:       1200.0000
BogoMIPS:          4199.93
Virtualization:    VT-x
L1d cache:         32K
L1i cache:         32K
L2 cache:          256K
L3 cache:          20480K
NUMA node0 CPU(s): 0-7,16-23
NUMA node1 CPU(s): 8-15,24-31
```

```
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 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 3dnowprefetch epb cat_l3 cdp_l3
↳ intel_ppin intel_pt tpr_shadow vnmi flexpriority
↳ ept vpid fsgsbase tsc_adjust bmi1 hle avx2 smep
↳ bmi2 erms invpcid rtm cqm rdt_a rdseed adx smap
↳ xsaveopt cqm_llc cqm_occup_llc cqm_mbm_total
↳ cqm_mbm_local dtherm ida arat pln pts
+ cat /proc/meminfo
MemTotal:           65858832 kB
MemFree:            55725052 kB
MemAvailable:       55457320 kB
Buffers:            8488 kB
Cached:             290800 kB
SwapCached:         0 kB
Active:             9281064 kB
Inactive:           232584 kB
Active(anon):       9216412 kB
Inactive(anon):     155824 kB
Active(file):        64652 kB
Inactive(file):     76760 kB
Unevictable:        0 kB
Mlocked:            0 kB
SwapTotal:          0 kB
SwapFree:           0 kB
Dirty:              48 kB
Writeback:          0 kB
AnonPages:          9214512 kB
Mapped:             69396 kB
Shmem:              157860 kB
Slab:               147848 kB
SReclaimable:       39992 kB
SUnreclaim:        107856 kB
KernelStack:        8320 kB
PageTables:         22188 kB
NFS_Unstable:       0 kB
Bounce:             0 kB
WritebackTmp:       0 kB
CommitLimit:        32929416 kB
Committed_AS:       10799664 kB
VmallocTotal:       34359738367 kB
VmallocUsed:        0 kB
VmallocChunk:       0 kB
HardwareCorrupted:  0 kB
AnonHugePages:      9168896 kB
ShmemHugePages:     0 kB
ShmemPmdMapped:     0 kB
HugePages_Total:    0
HugePages_Free:     0
HugePages_Rsvd:     0
```



```
HugePages_Surp:      0
Hugepagesize:        2048 kB
DirectMap4k:         1128200 kB
DirectMap2M:         39651328 kB
DirectMap1G:         28311552 kB
+ inxi -F -c0
collect_environment.sh: line 14: inxi: command not
↳ found
+ lsbblk -a
NAME      MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sdb       8:16    0 256.2G  0 disk 
sda       8:0     0 931.5G  0 disk 
└─sda2    8:2     0 927.7G  0 part /
└─sda1    8:1     0   3.8G  0 part /boot
+ lsscsi -s
collect_environment.sh: line 16: lsscsi: command not
↳ found
+ module list
++ /opt/ohpc/admin/lmod/lmod/libexec/lmod bash list
+ eval 'LMOD_DEFAULT_MODULEPATH="/opt/ohpc/admin/modu
↳ ulefiles:/opt/ohpc/pub/modulefiles";' export
↳ 'LMOD_DEFAULT_MODULEPATH;'
↳ 'MODULEPATH="/opt/ohpc/pub/moduledeps/gnu-mvapic
↳ h2:/opt/ohpc/pub/moduledeps/gnu:/opt/ohpc/admin/
↳ modulefiles:/opt/ohpc/pub/modulefiles";' export
↳ 'MODULEPATH;' echo '\'\''\''';' echo
↳ '\'\''Currently'\'' '\'\''Loaded'\''
↳ '\'\''Modules:\'\''';' echo '\'\''\'''\'' '\'\''\'''\''
↳ '\'\''1)\'\'' '\'\''autotools'\'' '\'\''\'''\''
↳ '\'\''\'''\'' '\'\''3)\'\'' '\'\''gnu/5.4.0'\''
↳ '\'\''\'''\'' '\'\''\'''\'' '\'\''\'''\'' '\'\''\'''\''
↳ '\'\''\'''\'' '\'\''5)\'\'' '\'\''cuda'\''';' echo
↳ '\'\''\'''\'' '\'\''\'''\'' '\'\''2)\'\''
↳ '\'\''prun/1.1'\'' '\'\''\'''\'' '\'\''\'''\'' '\'\''\'''\''
↳ '\'\''\'''\'' '\'\''4)\'\'' '\'\''mvapich2/2.2'\''
↳ '\'\''\'''\'' '\'\''\'''\'' '\'\''6)\'\''
↳ '\'\''ohpc'\''';' echo '\'\''\'''\''';'
↳ 'LMOD_DEFAULT_MODULEPATH="/opt/ohpc/admin/module
↳ files:/opt/ohpc/pub/modulefiles";' export
↳ 'LMOD_DEFAULT_MODULEPATH;'
↳ 'MODULEPATH="/opt/ohpc/pub/moduledeps/gnu-mvapic
↳ h2:/opt/ohpc/pub/moduledeps/gnu:/opt/ohpc/admin/
↳ modulefiles:/opt/ohpc/pub/modulefiles";' export
↳ 'MODULEPATH:'
```

```

_ModuleTable001_="X01vZHVszVRHymxLz17WjJhY3RpdmVtA
↳ XplI09NixiYXNlTXBhdGhBPXsiL29wdC9vaHBjL2FkbWlu
↳ 21vZHVsZWZpbGVziIiwL29wdC9vaHBjL3B1Yi9tb2R1bGVma
↳ WxlcYIsfSxbImNfcMvidWlsZFrbpWUiXT1mYWxzZSxbImNfc
↳ 2hvcnRUaw1l1l09ZmfSc2UsZmfTawX5PXtbIk1QSSJdPSJtd
↳ mFwaWNNoMiSwWyJjb21waWxlcjJdPSJnbnuUilFsiY3VkYSJdP
↳ SJjdWRHiIX9LGluYWN0aXZlPxT9LG1UPXthdXRvdG9vbHM9e
↳ 1sirK4iT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9hd
↳ XRvdG9vbHMiLFsiZGVmYXVsdCJDPTAsWyJmdWxsTmfTZSJdP
↳ SjhdxRvdG9vbHMiLfisibG9hZE9yZGVyIl09MSxwcm9wVD17f
↳ SbInNob3J0Il09ImF1dG90b28scysIsWyJzdGF0dXMiT0i"
↳ ;' export
↳ '_ModuleTable001_';

'_ModuleTable002_'="YWN0aXZlIiwL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9jdWRHLmx1YSIsW
↳ yJkZWZhWx0l109MCxbImZlbgOYw1l1l09ImN1ZGEiLFsic3
↳ G9hZE9yZGVyIl09NSxwcm9wVD17fSxbInNob3J0Il09ImN1Z
↳ GEiLFsic3RhddHVzIl09ImFjdGl2ZSIIsfSxbnbnU9e1sirK4iX
↳ T0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9nbnuVns40L
↳ jAilFsiZGVmYXVsdCJDPTesWyJmdWxsTmfTZSJdPSJnbnuVN
↳ S40ljAiLFsic3G9hZE9yZGVyIl09Myxwcm9wVD17fSxbInNob
↳ 3J0Il09ImduOSYhWyJzdGF0dXMiT0iYWN0aXZlIiwL2Yy
↳ XBpY2gyPXtbIkZOIl09Iiw9cHQvb2hwYy9wdWIvbW9kdWxlZ
↳ GVwcy9nbnuUvbXZhcGljaDIvMi4yIiw9cHQvb2hwYy9wdWI
↳ ;' export
↳ '_ModuleTable002_';

'_ModuleTable003_'="LFsiZnVsbE5hbWUiXT0ibXZhcGljaDIvM
↳ i4yIiw9cHQvb2hwYy9wdWIvbW9kdWxlZmlsZXMiLFsic3
↳ CJdPSJtdmFwaWNNoMiSwWyJzdGF0dXMiT0iYWN0aXZlIiwL
↳ G9ocGM9e1sirK4iT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVma
↳ WxlcY9vaHBjIiw9cHQvb2hwYy9wdWIvbW9kdWxlZmlsZXMi
↳ T0ib2hwYyIsWyJjb2I2fSkZjKXIiXT0LHBYyb3BUPXt9LFsic
↳ 2hvcnRUaw1l1l09ZmfSc2UsZmfTawX5PXtbIk1QSSJdPSJtd
↳ HBydW49e1sirK4iT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVma
↳ WxlcY9wcnVuLzEuMSIsWyJkZWZhWx0l109MSxbImZlbgOY
↳ W1l1l09InBydW4vMS4xIiw9cHQvb2hwYy9wdWIvbW9kdWxl
↳ FQ9e30swYjZaG9ydCJDPSJwcnVuIiw9cHQvb2hwYy9wdWI
↳ ;' export '_ModuleTable003_';

'_ModuleTable004_'="Y3RpdmUiLH0sfSxtcGF0aEE9eyIvb
↳ 3B0L29ocGMvchViL21vZHVsZWRLCHMVZ251LW12YXBpY2gyI
↳ iwIL29wdC9vaHBjL3B1Yi9tb2R1bGVkbXBzL2dudSIsIiw9c
↳ HQvb2hwYy9hZG1pbj9tb2R1bGVmaWxlcYIsIiw9cHQvb2hwY
↳ y9wdWIvbW9kdWxlZmlsZXMiLH0swYjZeXNOZW1CYXNlTVBBV
↳ EgiXT0iL29wdC9vaHBjL2FkbWluL21vZHVsZWZpbGVzOi9vc
↳ HQvb2hwYy9wdWIvbW9kdWxlZmlsZXMiLFsidmVyc2l2bviJdP
↳ TIIsfQ=";' export '_ModuleTable004_';

'_ModuleTable_Sz_'="4";' export
↳ '_ModuleTable_Sz_';

++ LMOD_DEFAULT_MODULEPATH=/opt/ohpc/admin/modulefil
↳ es:/opt/ohpc/pub/modulefiles
++ export LMOD_DEFAULT_MODULEPATH
++ MODULEPATH=/opt/ohpc/pub/moduledeps/gnu-mvapich2:
↳ /opt/ohpc/pub/moduledeps/gnu:/opt/ohpc/admin/mod
↳ ulefiles:/opt/ohpc/pub/modulefiles
++ export MODULEPATH
++ echo ''

```

```

++ echo Currently Loaded Modules:
Currently Loaded Modules:
++ echo ' ' '1)' autotools ' ' '3)' gnu/5.4.0 ' '
  ↳ ' ' ' ' '5)' cuda
    1) autotools    3) gnu/5.4.0    5) cuda
++ echo ' ' '2)' prun/1.1 ' ' ' '4)' mvapich2/2.2
  ↳ ' ' ' '6)' ohpc
    2) prun/1.1    4) mvapich2/2.2    6) ohpc
++ echo ' '

++ LMOD_DEFAULT_MODULEPATH=/opt/ohpc/admin/modulefiles
  ↳ es:/opt/ohpc/pub/modulefiles
++ export LMOD_DEFAULT_MODULEPATH
++ MODULEPATH=/opt/ohpc/pub/moduledeps/gnu-mvapich2:
  ↳ /opt/ohpc/pub/moduledeps/gnu:/opt/ohpc/admin/mod
  ↳ ulefiles:/opt/ohpc/pub/modulefiles
++ export MODULEPATH
++ _ModuleTable001=_X01vZHVszVRHymxlXz17WyJhY3RpdmVT
  ↳ aXplI109NixiYXNlTXBhdGhBPXsiL29wdC9vaHBjL2FkbWlu
  ↳ L21vZHVszWZpbGVZiwiL29wdC9vaHBjL3B1Yi9tb2R1bGVm
  ↳ aWxlcysfSxbImNfcmVidWlsZFRpbWUiXT1mYWxzZXsxbImNf
  ↳ c2hvcnRUaW1lI09MZFsc2UsZmFtaWx5PXBtIk1QSSJdPSJt
  ↳ dmFwaWNoMiIsWyJjb2IwXWxlcjJdPSJnbmUiLFsiY3VkySjJ
  ↳ PSJjdWRhIix9LGlUyWN0aXZlPXt9LG1UPXthdXRvdG9vbHM9
  ↳ e1siRk4iXT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9h
  ↳ dXRvdG9vbHMiLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJd
  ↳ PSJhdXRvdG9vbHMiLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJd
  ↳ fSxbInNob3J0I109ImF1dG90b29scysIswYJzdGF0dXMiXT0i
++ export _ModuleTable001_
++ _ModuleTable002_=YWN0aXZlIix9LGN1ZGE9e1sirK4iXT0i
  ↳ L29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9jdWRhLmx1YSIs
  ↳ WyJkZWZhdWx0I109MCCxbImZ1bGx0YW1lI109ImN1ZGEiLFsi
  ↳ bG9hZE9yZGVyI109NSxwcm9wVD17fSxbInNob3J0I109ImN1
  ↳ ZGEiLFsic3RhdHVzI109ImFjdG12ZSIsfSxnbU9e1siRk4i
  ↳ XT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVmaWxlcY9nbmUvNS40
  ↳ LjAiLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJdPSJnbmUv
  ↳ NS40LjAiLFsiZGVmYXVsdCJdPTAsWyJmdWxsTmFtZSJdPSJnbmUv
  ↳ b3J0I109ImdudSIsWyJzdGF0dXMiXT0iYWN0aXZlIix9LGI2
  ↳ YXBpY2gyPXBtIkZ0I109Ii9vcHQvb2hwYy9wdWIvbW9kdWxl
  ↳ ZGVwcy9nbmUvbXZhcGljaDViMi4yIixbImRlZmF1bHQiXT0x
++ export _ModuleTable002_
++ _ModuleTable003_=LFsiZnVsEb5hbWUiXT0ibXZhcGljaDIv
  ↳ Mi4yIixbImxvYWRPcmRlciJdPTQschJvcFQ9e30sWyJzaG9y
  ↳ dCJdPSJtdmFwaWNoMiIsWyJzdGF0dXMiXT0iYWN0aXZlIix9
  ↳ LG9ocGM9e1siRk4iXT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVm
  ↳ aWxlcY9vaHBjIixbImRlZmF1bHQiXT0wLFsiZnVsEb5hbWUi
  ↳ XT0ib2hwYyIsWyJsb2FkT3JkZXIiXT02LHBHyb3BUPXt9LFsi
  ↳ c2hvcnQiXT0ib2hwYyIsWyJzdGF0dXMiXT0iYWN0aXZlIix9
  ↳ LHBYdW49e1siRk4iXT0iL29wdC9vaHBjL3B1Yi9tb2R1bGVm
  ↳ aWxlcY9wcnVuLzEuMSIsWyJkZWZhdWx0I109MSxbImZ1bGx0
  ↳ YW1lI109InBydW4vMS4xIixbImxvYWRPcmRlciJdPTIschJv
  ↳ cFQ9e30sWyJzaG9ydcJdPSJwcnVuIixbInN0YXR1cyJdPSJh
++ export _ModuleTable003_

```

```
++ _ModuleTable004_=Y3RpdmiUjLH0sfSxtcGF0aEE9eyIvb3B0
└─ L29ocGMvchVjL21vZHVhZSWZRlchMvZ251LW12YXBpY2gyIiwj
└─ L29wdC9vaHBjL3B1Yi9tb2R1bGVkZXBzL2dudSIsIi9vcHQv
└─ b2hwYy9hZG1pbj9tb2R1bGVmaWxlcysIsIi9vcHQvb2hwYy9w
└─ dWIvbW9kdWx1ZmIsZXMiLH0sWyJzeXN0ZW1CYXNlTVBBVEgi
└─ XT0iL29wdC9vaHBjL2FkbWluL21vZHVhZSWZpbGVzOi9vcHQv
└─ b2hwYy9wdWIvbW9kdWx1ZmIsZXMiLfsidmVyc2lubiJdTPTIs
└─ fQ==
++ export _ModuleTable004_
++ _ModuleTable_Sz_=4
++ export _ModuleTable_Sz_
+ '[' 0 = 0 ']'
++ : -s sh
+ eval
+ lshw -short -quiet -sanitize
+ cat
collect_environment.sh: line 19: lshw: command not
└─ found
└─ lspci
00:00.0 Host bridge: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DMI2 (rev 01)
00:01.0 PCI bridge: Intel Corporation Xeon E7 v4/Xeon
└─ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 1
└─ (rev 01)
00:03.0 PCI bridge: Intel Corporation Xeon E7 v4/Xeon
└─ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 3
└─ (rev 01)
00:04.0 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 0 (rev 01)
00:04.1 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 1 (rev 01)
00:04.2 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 2 (rev 01)
00:04.3 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 3 (rev 01)
00:04.4 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 4 (rev 01)
00:04.5 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 5 (rev 01)
00:04.6 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 6 (rev 01)
00:04.7 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
└─ Channel 7 (rev 01)
00:05.0 System peripheral: Intel Corporation Xeon E7
└─ v4/Xeon E5 v4/Xeon E3 v4/Xeon D
└─ Map/VTd_Misc/System Management (rev 01)
```

```

00:05.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D IIO Hot Plug (rev
↳ 01)
00:05.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D IIO RAS/Control
↳ Status/Global Errors (rev 01)
00:05.4 PIC: Intel Corporation Xeon E7 v4/Xeon E5
↳ v4/Xeon E3 v4/Xeon D I/O APIC (rev 01)
00:11.0 Unassigned class [ff00]: Intel Corporation
↳ C610/X99 series chipset SPSR (rev 05)
00:11.4 SATA controller: Intel Corporation C610/X99
↳ series chipset sSATA Controller [AHCI mode] (rev
↳ 05)
00:14.0 USB controller: Intel Corporation C610/X99
↳ series chipset USB xHCI Host Controller (rev 05)
00:16.0 Communication controller: Intel Corporation
↳ C610/X99 series chipset MEI Controller #1 (rev 05)
00:16.1 Communication controller: Intel Corporation
↳ C610/X99 series chipset MEI Controller #2 (rev 05)
00:1a.0 USB controller: Intel Corporation C610/X99
↳ series chipset USB Enhanced Host Controller #2
↳ (rev 05)
00:1c.0 PCI bridge: Intel Corporation C610/X99 series
↳ chipset PCI Express Root Port #1 (rev d5)
00:1c.4 PCI bridge: Intel Corporation C610/X99 series
↳ chipset PCI Express Root Port #5 (rev d5)
00:1d.0 USB controller: Intel Corporation C610/X99
↳ series chipset USB Enhanced Host Controller #1
↳ (rev 05)
00:1f.0 ISA bridge: Intel Corporation C610/X99 series
↳ chipset LPC Controller (rev 05)
00:1f.2 SATA controller: Intel Corporation C610/X99
↳ series chipset 6-Port SATA Controller [AHCI mode]
↳ (rev 05)
00:1f.3 SMBus: Intel Corporation C610/X99 series
↳ chipset SMBus Controller (rev 05)
01:00.0 Ethernet controller: Intel Corporation
↳ Ethernet Controller 10-Gigabit X540-AT2 (rev 01)
01:00.1 Ethernet controller: Intel Corporation
↳ Ethernet Controller 10-Gigabit X540-AT2 (rev 01)
04:00.0 PCI bridge: ASPEED Technology, Inc. AST1150
↳ PCI-to-PCI Bridge (rev 03)
05:00.0 VGA compatible controller: ASPEED Technology,
↳ Inc. ASPEED Graphics Family (rev 30)
7f:08.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 0 (rev
↳ 01)
7f:08.2 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 0 (rev
↳ 01)
7f:08.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 0 (rev
↳ 01)
7f:09.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 1 (rev
↳ 01)
7f:09.2 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 1 (rev
↳ 01)
7f:09.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 1 (rev
↳ 01)
7f:0b.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
↳ (rev 01)
7f:0b.1 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
↳ (rev 01)
7f:0b.2 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
↳ (rev 01)
7f:0b.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link Debug
↳ (rev 01)
7f:0c.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0c.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0f.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0f.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0f.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)
7f:0f.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
↳ 01)

```

Uncore Power Scavenger: A Runtime for Uncore Power Conservation on HPC Systems

7f:0f.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

7f:10.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D R2PCIE Agent (rev 01)

7f:10.1 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D R2PCIE Agent (rev 01)

7f:10.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Ubox (rev 01)

7f:10.6 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D Ubox (rev 01)

7f:10.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Ubox (rev 01)

7f:12.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Home Agent 0 (rev 01)

7f:12.1 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D Home Agent 0 (rev 01)

7f:13.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Target Address/Thermal/RAS (rev 01)

7f:13.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Target Address/Thermal/RAS (rev 01)

7f:13.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

7f:13.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

7f:13.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

7f:13.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

7f:13.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Broadcast (rev 01)

7f:13.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Global
↳ Broadcast (rev 01)

7f:14.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 0 Thermal Control (rev 01)

7f:14.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 1 Thermal Control (rev 01)

7f:14.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 0 Error (rev 01)

7f:14.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 1 Error (rev 01)

7f:14.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

7f:14.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

7f:14.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

7f:14.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

7f:15.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 2 Thermal Control (rev 01)

7f:15.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 3 Thermal Control (rev 01)

7f:15.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 2 Error (rev 01)

7f:15.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 3 Error (rev 01)

7f:16.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Target
↳ Address/Thermal/RAS (rev 01)

7f:16.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Broadcast (rev 01)

7f:16.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Global
↳ Broadcast (rev 01)

7f:17.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 1 - Channel 0 Thermal Control (rev 01)

7f:17.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)

7f:17.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)

7f:17.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)

7f:17.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)

7f:1e.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)

7f:1e.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)

7f:1e.2 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 ↳ (rev 01)

7f:1e.3 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 ↳ (rev 01)

7f:1e.4 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 ↳ (rev 01)

7f:1f.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 ↳ (rev 01)

7f:1f.2 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 ↳ (rev 01)

80:01.0 PCI bridge: Intel Corporation Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 1
 ↳ (rev 01)

80:02.0 PCI bridge: Intel Corporation Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 2
 ↳ (rev 01)

80:04.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 0 (rev 01)

80:04.1 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 1 (rev 01)

80:04.2 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 2 (rev 01)

80:04.3 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 3 (rev 01)

80:04.4 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 4 (rev 01)

80:04.5 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 5 (rev 01)

80:04.6 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 6 (rev 01)

80:04.7 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA
 ↳ Channel 7 (rev 01)

80:05.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D
 ↳ Map/VTd_Misc/System Management (rev 01)

80:05.1 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D IIO Hot Plug (rev
 ↳ 01)

80:05.2 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D IIO RAS/Control
 ↳ Status/Global Errors (rev 01)

80:05.4 PIC: Intel Corporation Xeon E7 v4/Xeon E5
 ↳ v4/Xeon E3 v4/Xeon D I/O APIC (rev 01)

81:00.0 InfiniBand: Mellanox Technologies
 ↳ MT25408A0-FCC-QI ConnectX, Dual Port 40Gb/s
 ↳ InfiniBand / 10GigE Adapter IC with PCIe 2.0 x8
 ↳ 5.0GT/s In... (rev b0)

82:00.0 VGA compatible controller: NVIDIA Corporation
 ↳ TU106 [GeForce RTX 2070 Rev. A] (rev a1)

82:00.1 Audio device: NVIDIA Corporation TU106 High
 ↳ Definition Audio Controller (rev a1)

82:00.2 USB controller: NVIDIA Corporation TU106 USB
 ↳ 3.1 Host Controller (rev a1)

82:00.3 Serial bus controller [0c80]: NVIDIA
 ↳ Corporation TU106 USB Type-C Port Policy
 ↳ Controller (rev a1)

ff:08.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 0 (rev
 ↳ 01)

ff:08.2 Performance counters: Intel Corporation Xeon
 ↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 0 (rev
 ↳ 01)

ff:08.3 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 0 (rev
 ↳ 01)

ff:09.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 1 (rev
 ↳ 01)

ff:09.2 Performance counters: Intel Corporation Xeon
 ↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 1 (rev
 ↳ 01)

ff:09.3 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D QPI Link 1 (rev
 ↳ 01)

ff:0b.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
 ↳ (rev 01)

ff:0b.1 Performance counters: Intel Corporation Xeon
 ↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
 ↳ (rev 01)

ff:0b.2 Performance counters: Intel Corporation Xeon
 ↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
 ↳ (rev 01)

ff:0b.3 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D R3 QPI Link Debug
 ↳ (rev 01)

ff:0c.0 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
 ↳ 01)

ff:0c.1 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
 ↳ 01)

ff:0c.2 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
 ↳ 01)

ff:0c.3 System peripheral: Intel Corporation Xeon E7
 ↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev
 ↳ 01)

Uncore Power Scavenger: A Runtime for Uncore Power Conservation on HPC Systems

ff:0c.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0c.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0c.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0c.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0f.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0f.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0f.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0f.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:0f.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Caching Agent (rev 01)

ff:10.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D R2PCIe Agent (rev 01)

ff:10.1 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D R2PCIe Agent (rev 01)

ff:10.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Ubox (rev 01)

ff:10.6 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D Ubox (rev 01)

ff:10.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Ubox (rev 01)

ff:12.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Home Agent 0 (rev 01)

ff:12.1 Performance counters: Intel Corporation Xeon
↳ E7 v4/Xeon E5 v4/Xeon E3 v4/Xeon D Home Agent 0 (rev 01)

ff:13.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Target Address/Thermal/RAS (rev 01)

ff:13.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Target Address/Thermal/RAS (rev 01)

ff:13.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

ff:13.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

ff:13.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

ff:13.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel Target Address Decoder (rev 01)

ff:13.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Broadcast (rev 01)

ff:13.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Global
↳ Broadcast (rev 01)

ff:14.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 0 Thermal Control (rev 01)

ff:14.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 1 Thermal Control (rev 01)

ff:14.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 0 Error (rev 01)

ff:14.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 1 Error (rev 01)

ff:14.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

ff:14.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

ff:14.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

ff:14.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface (rev 01)

ff:15.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 2 Thermal Control (rev 01)

ff:15.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 3 Thermal Control (rev 01)

ff:15.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 2 Error (rev 01)

ff:15.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 0 - Channel 3 Error (rev 01)

ff:16.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Target
↳ Address/Thermal/RAS (rev 01)

ff:16.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Broadcast (rev 01)

```

ff:16.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Global
↳ Broadcast (rev 01)
ff:17.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Memory Controller
↳ 1 - Channel 0 Thermal Control (rev 01)
ff:17.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)
ff:17.5 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)
ff:17.6 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)
ff:17.7 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface (rev 01)
ff:1e.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)
ff:1e.1 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)
ff:1e.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)
ff:1e.3 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)
ff:1e.4 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)
ff:1f.0 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)
ff:1f.2 System peripheral: Intel Corporation Xeon E7
↳ v4/Xeon E5 v4/Xeon E3 v4/Xeon D Power Control Unit
↳ (rev 01)

```

+ lsmod

Module	Size	Used by
nvidia_uvm	765952	0
fuse	102400	3
nfs_v3	40960	1
nfs_acl	16384	1 nfs_v3
nfs	249856	2 nfs_v3
lockd	94208	2 nfs,nfs_v3
grace	16384	1 lockd
fscache	65536	1 nfs
8021q	32768	0
garp	16384	1 8021q
mrp	20480	1 8021q
rpcrdma	94208	0
ib_isert	49152	0
iscsi_target_mod	294912	1 ib_isert
ib_iser	49152	0
libiscsi	57344	1 ib_iser

scsi_transport_iscsi	94208	2 ib_iser,libiscsi
ib_srpt	45056	0
target_core_mod	356352	3
↳ iscsi_target_mod,ib_isert,ib_srpt		
ib_srp	49152	0
scsi_transport_srp	20480	1 ib_srp
ib_ipoib	110592	0
rdma_ucm	28672	0
ib_ucm	24576	0
ib_uverbs	65536	2 ib_ucm,rdma_ucm
ib_umad	24576	0
rdma_cm	53248	4
↳ ib_iser,ib_isert,rpcrdma,rdma_ucm		
ib_cm	49152	5
↳ rdma_cm,ib_srp,ib_ipoib,ib_srpt,ib_ucm		
iw_cm	45056	1 rdma_cm
mlx4_ib	163840	0
mlx4_en	118784	0
ib_core	200704	14 ib_iser,ib_cm,rdma_cm,ib_uverbs,ib_umad,ib_srp,ib_isert,ib_uverbs,rpcrdma,ib_ipoib,iw_cm,ib_srpt,ib_ucm,rdma_ucm,mlx4_ib
sunrpc	339968	19
↳ rpcrdma,nfs_acl,lockd,nfs,nfs_v3		
nvidia_drm	45056	0
nvidia_modeset	1040384	1 nvidia_drm
sb_edac	32768	0
edac_core	57344	1 sb_edac
snd_hda_codec_hdmi	45056	1
nvidia	16588800	2
↳ nvidia_modeset,nvidia_uvm		
x86_pkg_temp_thermal	16384	0
snd_hda_intel	36864	0
intel_powerclamp	16384	0
snd_hda_codec	126976	2
↳ snd_hda_intel,snd_hda_codec_hdmi		
coretemp	16384	0
kvm_intel	200704	0
snd_hda_core	81920	3
↳ snd_hda_intel,snd_hda_codec,snd_hda_codec_hdmi		
kvm	589824	1 kvm_intel
snd_hwdep	16384	1 snd_hda_codec
snd_seq	65536	0
snd_seq_device	16384	1 snd_seq
irqbypass	16384	1 kvm
snd_pcm	98304	4 snd_hda_intel,snd_hda_codec,snd_hda_codec_hdmi,snd_hda_core
↳ _codec,snd_hda_core,snd_hda_codec_hdmi		
crct10dif_pclmul	16384	0
crc32_pclmul	16384	0
snd_timer	32768	2 snd_seq,snd_pcm
crc32c_intel	24576	0
snd	77824	8
↳ snd_hda_intel,snd_hwdep,snd_seq,snd_hda_codec,snd_hda_codec_hdmi,snd_timer,snd_hda_codec_hdmi,snd_seq_device,snd_pcm		
ghash_clmulni_intel	16384	0
soundcore	16384	1 snd
i2c_algo_bit	16384	0

pcbc	16384	0	
ttm	98304	0	
aesni_intel	167936	0	
drm_kms_helper	151552	1	nvidia_drm
crypto_simd	16384	1	aesni_intel
glue_helper	16384	1	aesni_intel
cryptd	24576	3	
↪ crypto_simd,ghash_clmulni_intel,aesni_intel			
drm	352256	4	
↪ ttm,nvidia_drm,drm_kms_helper			
intel_cstate	20480	0	
intel_rapl_perf	16384	0	
mlx4_core	294912	2	mlx4_en,mlx4_ib
ixgbe	303104	0	
syscopyarea	16384	1	drm_kms_helper
sysfillrect	16384	1	drm_kms_helper
sysimgblt	16384	1	drm_kms_helper
fb_sys_fops	16384	1	drm_kms_helper
ptp	20480	2	ixgbe,mlx4_en
pps_core	16384	1	ptp
joydev	20480	0	
input_leds	16384	0	
iTCO_wdt	16384	0	
pcspkr	16384	0	
iTCO_vendor_support	16384	1	iTCO_wdt
i2c_i801	28672	0	
ipmi_si	57344	0	
mdio	16384	1	ixgbe
fjes	77824	0	
ipmi_devintf	20480	0	
ipmi_msghandler	45056	3	
↪ nvidia,ipmi_devintf,ipmi_si			
sg	36864	0	
shpchp	36864	0	
mxm_wmi	16384	0	
mei_me	40960	0	
wmi	20480	1	mxm_wmi
mei	102400	1	mei_me
lpc_ich	24576	0	
ioatdma	53248	0	
mfd_core	16384	1	lpc_ich
dca	16384	2	ioatdma,ixgbe
acpi_cpufreq	20480	0	
acpi_power_meter	20480	0	
acpi_pad	180224	0	
binfmt_misc	20480	1	
ip_tables	24576	0	
ext4	581632	2	
jbd2	106496	1	ext4
mbcache	16384	1	ext4
sd_mod	49152	3	
ahci	36864	2	
libahci	32768	1	ahci
libata	241664	2	ahci,libahci

ARTIFACT EVALUATION

Verification and validation studies: Results were reproduced independently by another team member of the same hardware/software setup following these instructions:

```
git clone --no-checkout git@arc:ngholka.git cd ngholka git config core.sparseCheckout true echo "Projects/Project3 - DUFM" >> .git/info/sparse-checkout git checkout -- cd cd Projects/Project3 - DUFM/source #look at README to generate all the libraries and test
```

Misc: 1. You need to have root privileges to access relevant MSRs. 2. If you are running submitting a multinode job as a root user, please make sure that the directory pointed by LIBMSR_NODE_PATH is writable by root. 3. Hyperthreading should be turned off. 4. Run clean_cat/submit_clean_cat before starting experiments or in case of a node(s) being left in dirty CAT state

+++++

To reproduce all the results Step 1 : login as root; Allocate a Broadwell node `srun -N2 -w c80 -pty /bin/bash`

Step 2 : Compile all benchmarks `cd DUFM/benchmarks`

#compile all NAS `make suite`

#compile ecp `cd miniAMR/ref make cd ../../`

#compile CoMD `cd CoMD/src-mpi make cd ../../`

Step 3 : Check the bins `ls bin/* ldd bin/*` #release the node exit

Step 4 : Submit jobs `cd scripts ./submit-jobs ./submit-jobs-rapl`