Not All GPUs Are Created Equal: Characterizing Variability in Large-Scale, Accelerator-Rich Systems

Prasoon Sinha, Akhil Guliani, Rutwik Jain, Brandon Tran, Matthew D. Sinclair and Shivaram Venkataraman
Computer Sciences Department, University of Wisconsin-Madison
Madison, United States of America
Email: {psinha9, guliani, rnjain, bqtran2}@wisc.edu, {sinclair, shivaram}@cs.wisc.edu

Abstract—Scientists are increasingly exploring and utilizing the massive parallelism of general-purpose accelerators such as GPUs for scientific breakthroughs. As a result, datacenters, hyperscalers, national computing centers, and supercomputers have procured hardware to support this evolving application paradigm. These systems contain hundreds to tens of thousands of accelerators, enabling peta- and exa-scale levels of compute for scientific workloads. Recent work demonstrated that power management (PM) can impact application performance in CPU-based HPC systems, even when machines have the same architecture and SKU (stock keeping unit). This variation occurs due to manufacturing variability and the chip's PM. However, while modern HPC systems widely employ accelerators such as GPUs, it is unclear how much this variability affects applications. Accordingly, we seek to characterize the extent of variation due to GPU PM in modern HPC and supercomputing systems. We study a variety of applications that stress different GPU components on five large-scale computing centers with modern GPUs: Oak Ridge's Summit, Sandia's Vortex, TACC's Frontera and Longhorn, and Livermore's Corona. These clusters use a variety of cooling methods and GPU vendors. In total, we collect over 18,800 hours of data across more than 90% of the GPUs in these clusters. Regardless of the application, cluster, GPU vendor, and cooling method, our results show significant variation: 8% (max 22%) average performance variation even though the GPU architecture and vendor SKU are identical within each cluster, with outliers up to $1.5\times$ slower than the median GPU. These results highlight the difficulty in efficiently using existing GPU clusters for modern HPC and scientific workloads, and the need to embrace variability in future accelerator-based systems.

Index Terms—Accelerator Architectures; Dynamic Voltage Scaling; Power Measurement; Temperature Measurement; Time Measurement

I. Introduction

Recently, domain scientists have leveraged the massive parallelism of accelerators for scientific discovery. Some of these discoveries use machine learning (ML) or deep learning (DL) for image recognition [1], [2], speech recognition [3]–[5], and machine translation [6], [7]. Scientists have also used accelerators in other areas, including molecular dynamics, material science, and quantum chemistry. Since these applications often require peta- or exascale levels of compute, running them on massively parallel systems has yielded promising results in areas including protein folding [8], plasma reactor status prediction [9], material science [10], and SARS-CoV-2 [11].

Accordingly, supercomputers, datacenters, and computing centers have procured hardware to accommodate this evolving application paradigm. To reach exascale levels of compute, many of these systems use many accelerators, which offer greater power efficiency and thus support emerging AI


Fig. 1: Normalized runtime across the five compute clusters for SGEMM. All clusters exhibit significant performance variability and contain several outliers.

and HPC workloads within a constrained power budget. For example, nearly all of the top 10 supercomputers leverage GPUs [12] and the second-ranked supercomputer, Oak Ridge National Lab's (ORNL's) Summit, has over 27000 GPUs. Compute centers such as NCSA Delta (840 GPUs), SDSC Expanse (216 GPUs), and Texas Advanced Computing Center (TACC, 744 GPUs) also utilize many GPUs. Similarly, Microsoft, Tesla, and others have deployed accelerator-based supercomputers [13], [14]. The upcoming Aurora, El Capitan, and Frontier supercomputers [15] are expected to have even more GPUs. Thus, current exascale computing systems contain hundreds to tens of thousands of GPUs, and future systems will likely be comprised of a large variety of accelerators and customized chips [13], [14], [16]–[23].

Despite their power efficiency, exascale systems have an enormous power footprint, making it important to consider the hardware's *power management* (PM) algorithms. PM can lead to power and frequency variations across nodes. Such dynamic behavior makes it challenging for repeatable, high performance and can lead to resource underutilization. For example, recent work studying CPU-based supercomputers showed that PM impacts application performance by up to 20%, even for CPUs with the same architecture and vendor SKU (Stock-Keeping Unit) [24]–[29]. This variation occurs due to the manufacturing process and the chip's power constraints [26], [30]. However, despite their increasingly widespread use in modern HPC systems, there is limited work that examines how accelerator PM and manufacturing variability affects application performance. For example, while Scogland, et al. [30] conducted an AMD GPU variability study in 2015, they only used one benchmark and called for a more in-depth study. We discuss related work further in Section II.

In this paper we perform a rigorous study to understand

GPU variability in large scale, accelerator-rich computing clusters. While some systems [13], [14] utilize FPGAs or other accelerators, we focus on GPUs since the majority of the Top-500 systems utilize GPUs. To ensure coverage across GPU workloads that are run on modern systems and to stress different GPU components, we select five applications from ML, sparse graph analytics, and molecular dynamics. Furthermore, we study multiple computing clusters to examine how different scales (hundreds to tens of thousands of GPUs), cooling (air, mineral oil, and water), and GPU vendors (AMD and NVIDIA) impact variability. Specifically, we study five modern compute clusters: ORNL's Summit, Sandia National Lab's (SNL's) Vortex, TACC's Frontera & Longhorn, and Lawrence Livermore National Lab's (LLNL's) Corona. We perform measurements on over 90% of the GPUs on each cluster and verify that our methodology is statistically significant [31]. In total, we collected and analyzed over 18,800 hours of data to study variability in performance, temperature, GPU frequency, and power consumption. Our findings include:

- All clusters exhibit significant performance variability when running SGEMM (Figure 1), which is normalized to a median runtime of 1. On average, SGEMM performance varies between 7%-9% across GPUs in the same cluster, with outliers up to $1.5\times$ slower than the median GPU.

- Performance variability is larger for compute intensive workloads such as matrix multiplication (SGEMM) or training ML models like ResNet-50 (22%, max $3.5\times$ slower than median GPU) than memory intensive workloads like PageRank (1%). BERT pre-training, which has a mix of compute and memory-bound operations, has 8% performance variation. This variation impacts resource utilization, increases runtime, and hurts responsiveness, especially for multi-GPU experiments.

- Air cooled clusters (Longhorn, Corona) have a large temperature range $\geq 30°C$ and very high temperature nodes can have severe power throttling (>15% of TDP). Water and mineral oil cooling reduce temperature variation, but do not reduce performance and power variation.

- On the largest GPU cluster we study (Summit), there are a number of power outliers that are not correlated with temperature or frequency variation.

- Our results are consistent across different days of the week and times of day, indicating that the variability is not transient. Moreover, we also validated that similar or higher variability occurs when limiting the GPU power to be lower than the TDP (Section VI-A).

Overall our results indicate that significant power and performance variation exists in GPU-rich systems. Given there are no standards to expose PM information to administrators, it is difficult today to identify aberrations within the cluster that might require maintenance, modifications, or even replacement of GPUs. Further, unlike CPU-based systems, GPU nodes cannot be assigned as a desktop or a submit node to minimize impact. Thus, our results motivate systematic benchmarking across nodes to provide an early-warning for system administrators to perform maintenance or investigate bad GPUs, without hurting long-term cluster performance. Finally, our work also highlights the need to improve visibility of accelerator variability and design mitigation techniques in runtime frameworks for both users and system administrators.

II. BACKGROUND AND RELATED WORK

A. Hardware Performance Variability

As large scale datacenters and supercomputers are typically power-constrained, manufacturing variability between two identical architectures translates to varying performance and causes load imbalance even for perfectly balanced workloads. PM and temperature variations can also cause thermal throttling, possibly degrading performance further. To mitigate variability, CPU-based HPC systems have adopted various techniques including: dynamic load balancing algorithms [32], [33], cooling mechanisms and temperature-aware job placement [34]–[36], and intelligent adaptive runtimes [24]–[26].

However, less work has examined how GPU PM and manufacturing variability affect applications. Coplin, et al. [37] and Jiao, et al. [38] demonstrated that memory- and compute-bound GPU workloads exhibit significant differences in performance and energy for a variety of configurations. Jiao, et al. also found that increasing frequency often increased both power consumption and performance in Fermi-class GPUs. Moreover, they demonstrated that GPUs utilize dynamic voltage and frequency scaling (DVFS) to help stay within their power limit. However, both studies were conducted on GPUs several generations older than the ones in use today and focus on performance variability for a single-GPU workstation. Similarly, Scogland, et al. examined how LINPACK's performance varied for CPU clusters and AMD GPUs [30]. However, this study was done in 2015, only used a single benchmark, and motivates the need for a more in-depth study of GPU variability. Thus, while these papers are a useful foundation, characterization is needed to determine the impact of PM on modern GPU-rich systems for representative workloads.

Other recent work has also examined how the Titan supercomputer's thousands of GPUs behaved over the machine's lifetime [39]. However, this work focused on issues like reliability, and did not investigate the impact of PM.

B. GPU Power Management

Current many-GPU systems use a local-only PM setup: each GPU has a given thermal limit it must stay within (its thermal design power, TDP), which is 300W for NVIDIA V100 GPUs and AMD MI60 GPUs [40], [41] and 230W for the NVIDIA Quadro RTX 5000 GPUs [42] we study. The GPU PM controller varies the GPU's voltage and frequency using DVFS to avoid exceeding its TDP [43], [44]. Although GPU vendors such as AMD and NVIDIA have not disclosed details about their DVFS schemes (or PM controllers), prior work has shown that, similar to multi-core CPUs, DVFS adjusts the GPU's streaming multiprocessor (SM) and memory frequencies and voltages to stay within the TDP and reduce energy consumption. Sometimes DVFS inadvertently compromises

Cluster	GPU	# GPUs	# Nodes	Cooling
CloudLab [46]	NVIDIA V100	12	4	air
Longhorn [47]	V100	416	104	air
Frontera [47]	RTX 5000	360	90	mineral oil
Vortex [48]	V100	216	54	water
Summit [49]	V100	27648	4608	water
Corona [50]	MI60	328	82	air

TABLE I: Summary of clusters studied.

Benchmark	Input Size	Clusters Observed	Collection Duration
SGEMM [57]	25536×25536	Longhorn	6 Weeks
	25536×25536	Summit	8 Weeks
	24576×24576	Corona	2 Weeks
ResNet-50 [58]	Train. Set: 1.2M images Batch size: 64	Longhorn	2 Weeks
BERT [59]	Train. Set: 30K words Batch size: 64	Longhorn	1 Week
LAMMPS [60]	$(x, y, z) = (8, 16, 16)$	Longhorn	2 Weeks
PageRank [61]	643994×643994	Longhorn	2 Weeks

TABLE II: Summary of applications studied on HPC clusters.

performance [43]–[45]. Thus, PM can affect the performance of GPU applications. Moreover, there is no global PM strategy across GPUs in a cluster. As a result, GPUs can be also affected by the behavior of other GPUs on the same node.

III. METHODOLOGY

To ensure that our results were representative across a variety of GPU clusters with different properties, we sought to run experiments across clusters of different sizes, cooling approaches, and GPU vendors. Moreover, we also selected benchmarks that are representative of how modern systems are used and which stress a variety of GPU components.

Cluster Parameters: Table I summarizes the unique HPC clusters we studied. First, we used a small cluster in Cloud-Lab [46], which contains 12 air-cooled GPUs split across 3 nodes, each with 4 NVIDIA V100-SXM2 GPUs. We also studied five larger HPC clusters. TACC's air-cooled Longhorn cluster [47] has 416 GPUs split across 104 nodes, each with 4 NVIDIA V100 GPUs. SNL's water-cooled Vortex cluster [48] contains 216 NVIDIA V100 GPUs split across 54 nodes. To examine variability in a larger cluster, we also studied ORNL's water-cooled Summit supercomputer [49] (27648 NVIDIA V100 GPUs). Finally, to examine variability across GPU vendors, we studied AMD MI60 GPUs in Livermore's Corona cluster [50]. Although we do not have access to Vortex's temperature setup, the Summit and TACC GPU's shutdown, slowdown, max operating, and max memory operating temperatures are: 90°C, 87°C, 83°C, and 85°C, respectively. The Corona GPU's shutdown, slowdown, and max memory operating temperatures: are 105°C, 100°C, and 99°C, respectively [51], [52]. Finally, the Frontera GPU's shutdown, slowdown, and max operating temperatures are: 96°C, 93°C, and 89°C, respectively. Throughout our experiments, we did not observe GPUs exceed these thresholds.

GPUs: Although AMD and NVIDIA have recently released newer GPUs than the ones available in the systems we study, we chose to study MI60's, RTX 5000's, and V100's because they are widely used in modern HPC systems. Prior work recommended pinning frequency, power limits, fan speed, and voltage ID to ensure that all GPUs have the same initial state (from the software and user perspective) [31]. However, Volta V100 GPUs do not have fans [53]. Instead, like many server-class GPUs, they are attached to large heatsinks and rely on server chassis fans to pull air across them for cooling. As a result, it is not possible to adjust their fan speed (e.g., using nvidia-smi) like desktop-class GPUs. Moreover, as we do not have administrative privileges on the clusters, we could not pin the frequency, power limit, or voltage ID. Instead, we verified that all GPUs were configured to the maximum frequency and

power limit: 1530MHz and 300W for the V100s and 1800MHz and 300W for the MI60s, respectively. We believe this setup is still relevant because it demonstrates how much variability a common HPC user, who does not have these privileges, would see. Moreover, we observed similar variability on a smaller CloudLab cluster where we had administrator privileges and could pin frequency and power (Section VI-B).

Workloads: We selected four applications from parts of CORAL-2 [54], a collection of benchmarks used to measure exascale workload performance. To ensure that our workloads stress different GPU components and thus provide a holistic view of variability, we chose applications that are compute-bound, memory-bound, or balanced [37], [55], [56]. Table II summarizes the applications. SGEMM is a compute-intensive matrix-multiply kernel that is widely used in a number of workloads, including ML. ResNet-50 is a popular, multi-GPU, compute-intensive ML training workload. BERT is a popular, multi-GPU Transformer-based model that is part of the of the widely used MLPerf training benchmark suite. LAMMPS is a popular scalable scientific computing application that is memory-bound our selected configuration. Finally, we tested another memory-bound workload, PageRank, which is used in the Havoq graph analytics benchmark. To ensure we study both single GPU and multi-GPU runs, we run LAMMPS, PageRank, and SGEMM on a single GPU while ResNet and BERT are run across multiple GPUs. We also analyzed ResNet when running as a single GPU application with batch size scaled down appropriately to enable comparisons with the multi-GPU ResNet experiments. To evaluate variability across clusters, we ran SGEMM on all clusters; we ran all other workloads on TACC's Longhorn cluster. We provide further details about each application's setup in their corresponding Results & Analysis sections. Collectively, analyzing these workloads helps us gauge whether GPU variability is application-specific or not in HPC systems.

Measurement: We collected four metrics for the duration of each application: kernel runtime (iteration runtime for ResNet-50 and BERT) in milliseconds (ms), GPU CU/SM temperature (°C), GPU CU/SM power consumption (Watts), and GPU CU/SM frequency (MHz). To collect this data we use the GPU vendor's profilers [62], [63]. As in prior work [31], we computed the recommended sample size (number of GPUs) for each cluster to obtain $\lambda = 0.5\%$ accuracy for average power within a 95% confidence interval. Given that we sample measurements from almost all GPUs in each cluster, our sample size is 2.9× larger than the worst-case

recommendations. Thus, our measurements are statistically significant. Unless otherwise specified, we use the median of each measurement to avoid one-off outliers. Since 1ms is the minimum sampling interval for these profilers, we configured our input sizes to ensure kernel durations were larger than 1ms. To characterize applications as compute- or memory-intensive, we also collected profiler metrics related to functional unit (FU) utilization (the utilization level of arithmetic functional units measured by nvprof on a scale of 0 to 10), DRAM utilization, and stalls. As all the clusters we study are shared, we collected measurements while other machines were in use. However, for consistent measurement we ensured there was no timesharing of our allocated nodes or GPUs during data collection. By using exclusive allocations and staggered run times, we eliminated spatial and temporal effects on variability for all applications. We discuss spatial effects further in Section VII. Finally, to protect against transient effects we collected data for multiple runs on the same machine over multiple days or weeks.

IQR & Variability: We use box and whiskers plots and inter-quartile regression (IQR) to help determine variation and categorize statistically significant outliers. Each box plot represents the spread from quartile 1 to 3 ($Q1$ to $Q3$). The center of the box represents the median ($Q2$). With $IQR = Q3 - Q1$, the upper and lower whiskers in the plot represent $Q3 + 1.5IQR$ and $Q1 - 1.5IQR$, respectively. The IQR captures 99.3% of the Gaussian distribution within the box-and-whisker format. We define *range* as the difference between the upper and lower whiskers. We denote the variation of a metric as $\frac{range}{Q2}$ and **denote all data points that fall outside the whiskers as** *outliers*. **Thus, outliers are** *not* **included in our variance calculations**.

IV. VARIATION ACROSS CLUSTERS

A. Methodology for SGEMM Application

We use SGEMM to study variability across clusters. Our SGEMM application is a single SGEMM kernel that performs matrix multiplication on two matrices containing single-precision floats. To do this, we use optimized SGEMM implementations from NVIDIA's cuBLAS [57] and AMD's hipBLAS [64] libraries for the respective GPUs. To study the effects of the GPU's PM controller, it is important that all streaming multiprocessors (SMs) or compute units (CUs) are fully occupied and the work is evenly distributed across SMs/CUs. Moreover, because the PM controller relies on DVFS to maintain the power limit for safe operation [44], [65], the kernel run must be long enough for the DVFS controller to reach a stable state [37]. Thus, we carefully tuned the matrix size (Table II) for both NVIDIA V100s and AMD MI60s to (i) achieve a sufficient runtime (Section III), (ii) ensure high performance, and (iii) provide high SM/CU occupancy.

We define 1 run of our experiment as 100 repetitions of the SGEMM kernel. The repetitions help avoid statistical bias and transient effects. Before collecting data, we run one warm-up run to avoid counting cuDNN startup overheads [66], [67]. We provide the exact same matrix inputs to every GPU, since the data inputs themselves are not important in our study.

B. SGEMM on TACC Longhorn

Figure 2 shows aggregated box plots for kernel duration, frequency, power, and temperature for SGEMM on Longhorn. Overall, SGEMM has 9% performance variation on Longhorn. Moreover, despite verifying that the GPUs are configured to run at the highest frequency (1530MHz), Figure 2 shows that the GPUs run at lower frequencies (1300-1440MHz). The overall frequency variation is 140MHz or 11%. We also observe a wide spread in temperatures (33°C between $Q1$ and $Q3$). Furthermore, certain GPUs also run at power levels far below 300W: at around 250W. This significant variation in performance, power, and temperature motivated us to further investigate the relationship between these metrics.

Figure 3 presents scatter plots between the distinct measurement pairs. In Figure 3, although sometimes the slowest kernels run on the GPUs at the highest temperatures – and the fastest kernels run on the lowest temperature GPUs – overall performance and temperature are not strongly correlated ($\rho = 0.46$). For example, multiple GPUs running at the same temperature have up to a 200ms (10%) performance difference. Additionally, the top-right of the performance vs temperature plot shows a cluster of GPUs with both high temperature and longer runtimes. This implies that if a GPU is running at a high temperature (>78°C), it is likely to have worse performance. However, this is not always the case, because other GPUs run at similar temperatures but complete much faster (e.g., c004). This is surprising because as frequency increases, temperature is expected to increase. Moreover, usually higher frequencies improve SGEMM's performance [30], [44] and our results agree with this trend. Figure 3c shows that performance and frequency are strongly correlated on Longhorn ($\rho = -0.97$). Further, Figure 3b shows that certain GPUs consume less than 290W, and these GPUs usually have lower performance and higher temperatures. This indicates that higher temperature levels may hurt performance. OVerall, these variations in trends highlighted the need to examine these GPUs more closely. So, we reported our findings to Longhorn system administrators, which helped them perform early identification of a poorly performing node and investigate it in greater detail.

Takeaway 1: *Despite running the same kernel on similarly configured GPUs, we see 9% performance variance across GPUs in the Longhorn cluster, with temperature range $\geq 30°C$ and some power outliers at 250W.*

C. SGEMM on ORNL Summit

We next examine the larger Summit to see if similar patterns are seen. As shown in Table I, Summit contains 27648 GPUs and is water cooled. We collect the same measurements for SGEMM as Longhorn. However, since Summit has so many GPUs, we further breakdown our measurements based on the particular row and column where a machine resides (using node location references derived from ORNL's layout [68]).

Figure 4 shows aggregated kernel duration, frequency, power, and temperature box plots, grouped by rows. Similar to Longhorn, Summit has 8% performance variation across
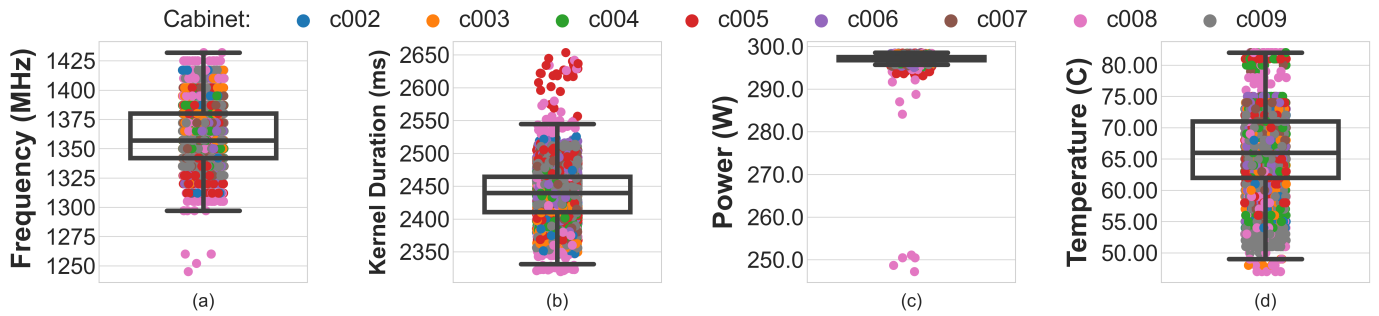
Fig. 2: Summary results for SGEMM on Longhorn cluster presented as box plots of (a) frequency, (b) kernel duration (performance), (c) power, and (d) temperature. The color indicates the cabinet of the GPU.
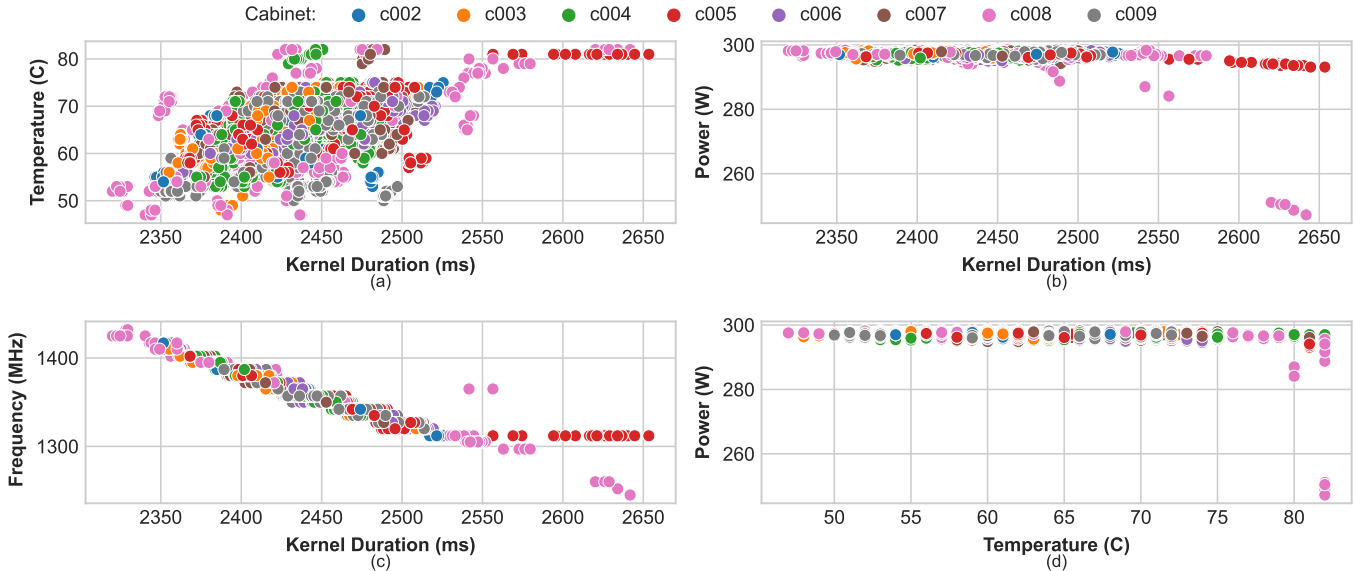

Fig. 3: Scatter plots for SGEMM on Longhorn showing that (a) performance and temperature have a weak positive correlation (Pearson correlation coefficient ρ =0.46), (b) power and performance have a weak inverse correlation ($\rho = -0.35$), (c) performance and frequency are strongly correlated ($\rho = -0.97$) and (d) power and temperature are uncorrelated ($\rho = -0.1$). The color indicates the cabinet of the GPU.

all rows, with rows D and F having the most outliers (Figure 4a). Likewise, frequency variation (Figure 4b) is again near 100MHz across all rows, although rows D and F have some outliers below 1300MHz. Moreover, although all row's IQRs (Figure 4c) range from 295-300W, a number of GPUs consume less than 290W, especially in rows A and H. However, unlike Longhorn, Summit's temperature range is narrow: 40°-62°C. This shows the benefits of water cooling [39] compared to Longhorn, although water cooling does not significantly impact frequency or performance variation.

Figure 5 shows Summit's scatter plots. Similar to **Takeaway 1**, performance and frequency were directly correlated ($\rho = -0.99$) on Summit (Figure 5a). However, unlike Longhorn, Summit has a string of power outliers below 290W with 2510ms runtime (Figure 5b). This is unexpected, since GPUs with a wide power range typically vary in their performance.

To understand these outliers, we analyzed row H, column 36 because of its variance in power consumption.[1] We found

that 7 of its nodes exhibited power outliers, with power as low as 255W, while the remaining 9 nodes did not have any outliers. However, despite the power outliers, nodes 10 and 11 did not have any temperature outliers, unlike Longhorn. Hence, these two nodes show that although water cooling decreases temperature variation, it does not prevent nodes from having large performance and power variation. Moreover, the variance in power consumption seen in row H, column 36 highlights the difficulty in drawing conclusions from cluster-wide summaries and provides another example of how our study can be used to flag underperforming nodes early for system administrators to investigate further, similar to Longhorn.

Takeaway 2: *Summit and Longhorn have similar performance and frequency variation trends, but Summit has severe power outliers that are concentrated in a few rows.*

Takeaway 3: *While Summit and Longhorn have the same GPU temperature setups, water cooling reduces temperature variation, but does not improve performance and power variation.*

[1]See Appendix B graphs for this row-column pair.

Fig. 4: Summary results for SGEMM on Summit, showing variation in (a) kernel duration (performance), (b) frequency, (c) power, and (d) temperature. We breakdown box plots by row because of the large scale of the cluster.
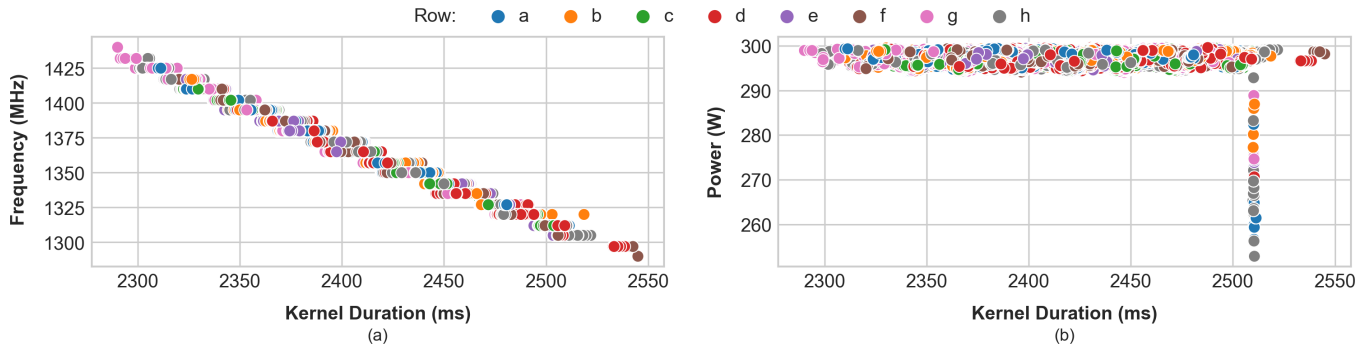

Fig. 5: Scatter plots of Summit data showing (a) performance and frequency have strong negative correlation ($\rho = -0.99$) and (b) performance and power have almost no correlation ($\rho = -0.09$). The color indicates the row for the GPU.
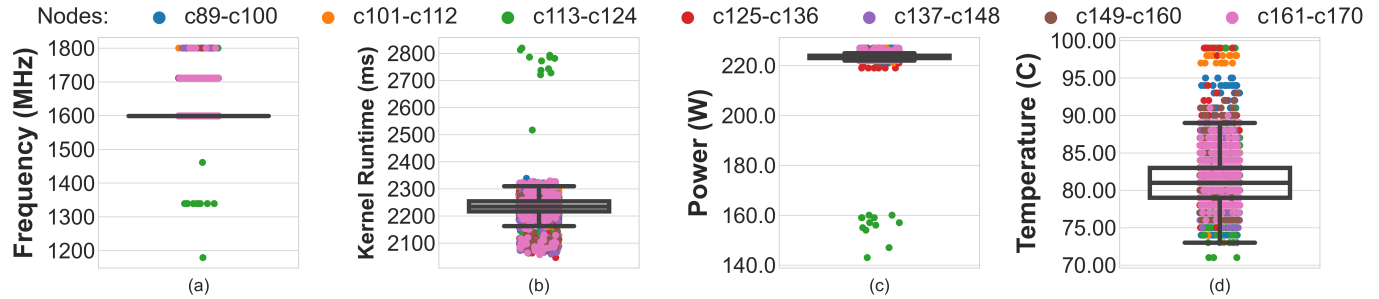

Fig. 6: Summary results for SGEMM on Corona presented as box plots of (a) frequency, (b) mean kernel duration (performance), (c) power, and (d) temperature. Node c115 (green) is the outlier.

D. SGEMM on Corona LLNL

Since compute clusters use GPUs from multiple vendors, it is important that we also examine how AMD GPUs vary. Corona uses air-cooling for its AMD MI60 GPUs. We col-

lected the same measurements as Longhorn and Summit. However, unlike Longhorn and Summit we were unable to obtain a system map and thus analyze our data by node. Nevertheless, to improve readability we group nodes in "cabinets" of 12 GPUs, like the similarly-sized Longhorn cluster.
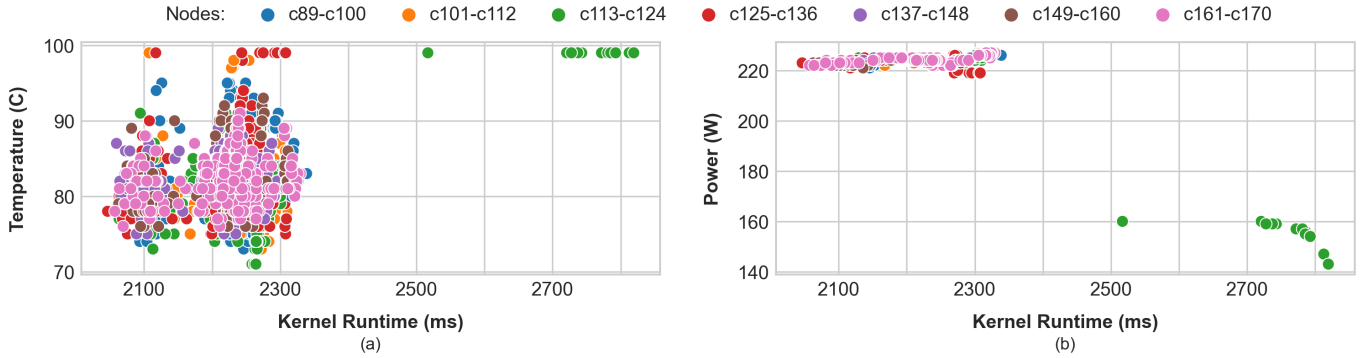
Fig. 7: Scatter plots for SGEMM on Corona showing correlations between (a) kernel duration (performance) and temperature ($\rho = 0.20$) and (b) kernel duration and power ($\rho = -0.48$). Node c115 (green) is the outlier.

Figure 6 presents aggregated box plots for our 4 metrics. Individual data points are distinguished by our node groupings. Overall, Corona exhibits 7% runtime variation (Figure 6b), similar to Longhorn (9%) and Summit (8%). Interestingly, frequency and performance are not as strongly correlated on Corona ($\rho = -0.76$ on Corona, versus Longhorn's -0.93 and Summit's -0.99), as Figure 6a shows much less variability than Longhorn and Summit (**Takeaways 1-2**). Moreover, the MI60s have coarser frequency levels than the NVIDIA V100s, suggesting that the GPU vendor's DVFS schemes vary significantly. Finally, the power (2%, Figure 6c) and temperature (20%, Figure 6d) variability are similar to Longhorn. However, whereas Longhorn has more power outliers consuming as little as 250W, Corona only has one outlier node (c115) which consumes 165W. Surprisingly Corona's nodes never reach the max power of 300W. Since the temperature are as high as 99°C (near the slowdown temperature), we believe the DVFS controller heavily throttles the frequency (most frequencies also do not reach the peak 1800MHz), thereby decreasing the power since $P_{dynamic} \propto frequency$.

To further analyze the outlier node, c115, Figure 7 correlates our metrics. Corona's kernel runtime and power relationship (Figure 7b) is similar to Longhorn and matches expectations: slower GPUs tend to consume less power. However, Summit does not exhibit this trend, as similarly performing GPUs in Summit exhibit a wide power range. The performance-power correlation coefficients are -0.48 and -0.35 on the similarly-sized Corona and Longhorn clusters, respectively, but -0.09 on Summit. However, since the cluster sizes differ, we also compared Summit against a scaled normal distribution of Longhorn's performance numbers to determine cluster size's impact. The scaled normal performance distribution projects that the Longhorn data would have 9.4% variability on a Summit-sized cluster. Since our actual Summit measurements (Section IV-C) had 8% performance variability, this suggests the cluster size may impact the severity of variability. Also, while c115 runs very hot compared to most nodes (Figure 7a), there is not a clear relationship between kernel runtime and temperature ($\rho = 0.20$). Typically higher temperatures occur when GPUs run at higher frequencies and achieve better
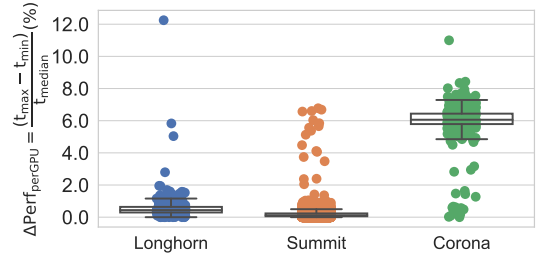

Fig. 8: Normalized performance variation within a GPU for SGEMM on Longhorn, Summit, and Corona clusters.

performance. However, there are several nodes that run as hot as c115 but which do not perform worse than the median GPU, which likely rules out that this outlier was due to inadequate air cooling. Moreover, all of these GPU's temperatures are near the slowdown temperature. Additionally, Summit nodes that perform similarly also exhibit a wide temperature range. For example, Corona GPUs performing around 2210ms have a 25°C temperature range. Thus, c115 appears to be an under-performing GPU and another candidate for further inspection and potential replacement by the system administrators.

Lastly, a single GPU exhibits little performance variation across independent application runs. Figure 8 shows the normalized per-GPU performance variance for SGEMM on various clusters. The median variance values are 0.44%, 0.12% and 6.06% on Longhorn, Summit, and Corona, respectively. Thus, our results are repeatable and any transient effects cause only minor variations. Although there are a few outliers as high as 12% (Longhorn node c002-010), we found that these do not correspond to the worst performing GPUs. Instead, they fall between the median and $Q3$, which suggests that ill-performing GPUs are consistently ill-performing.

Takeaway 4: *Corona's AMD GPUs exhibit similar behavior to the like-sized Longhorn cluster with NVIDIA GPUs. However, the performance-power relationship in the larger Summit differs from Corona and Longhorn, suggesting that cluster size affects variability degree.*

E. SGEMM on SNL Vortex

We also evaluated SGEMM's behavior on SNL's Vortex cluster [48]. Vortex is similarly sized to Longhorn and Corona,
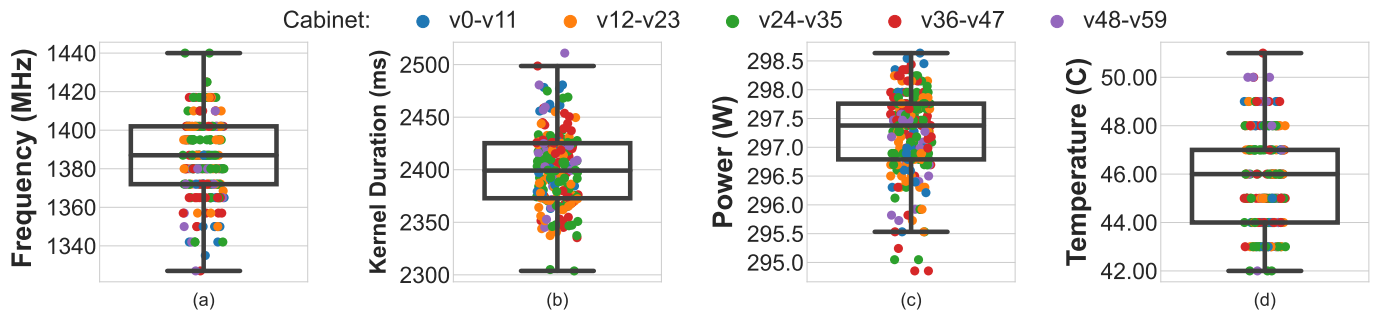
Fig. 9: SGEMM box plot summary results on SNL Vortex for (a) frequency, (b) performance, (c) power, and (d) temperature.

but is water-cooled whereas Longhorn and Corona are air-cooled. We gathered data from 184 GPUs on Vortex. However, we were unable to obtain information about their location within the cluster. Thus, similar to Corona, we simplify readability of plots by grouping nodes into "cabinets" of 12 GPUs.

Figure 9 shows aggregated box plots for kernel duration, frequency, power, and temperature when the SGEMM kernel is run on Vortex. We observed a performance variation (9%) similar to Longhorn (9%) and Summit (8%) clusters. Despite the GPUs being configured to run at the highest frequency level (1530MHz), Figure 9 shows that they often run at lower frequencies (1330-1442 MHz), similar to other clusters. These frequencies vary by about 100 MHz or 10% between the fastest and slowest GPUs observed. The temperature difference between the GPUs has a narrower spread (10°C between $Q1$ and $Q3$) on Vortex than Longhorn. This is likely due to Vortex's use of water-cooling. From Figure 9c, we also observe that all GPUs operated within 5W of the 300W power limit, unlike Longhorn where some outliers consuming significantly lower power (Figure 2).

Figure 10 presents correlations between the metrics. Similar to **Takeaway 1**, on Vortex we again see a strong inverse correlation between frequency and performance. Figure 10b shows the correlation between temperature and kernel duration. Similar to Longhorn (Figure 3), multiple GPUs running at the same temperature have up to a 200 ms (10%) difference in performance. However, unlike Longhorn we do not observe GPUs with high temperature having longer runtimes in Vortex. Thus, temperature and performance apper to be weakly correlated in water-cooled systems. This may be because all observed GPUs are running relatively cooler (median 46°C) when compared to Longhorn (median 66°C).

From these observations, we find a consistent trend across Longhorn, Summit, Corona, and Vortex of performance variations being inversely correlated with changes in frequency. To verify that these frequency changes occur due to DVFS mechanisms employed by PM infrastructure, we examine the timelines of frequency and power for two GPUs at the extremes of kernel performance (Figure 11). We examine a 10s slice in the middle of the overall application run, during which 4 SGEMM kernels are launched and run to completion. Vertical lines indicate the start time of a new kernel. Figure 11a shows that an SGEMM kernel is launched at 430.6s and the

GPU reports a frequency of 1340MHz. Similarly, a kernel is launched on GPU-2 at 432s, with the GPU reporting an initial frequency slightly under 1440MHz. As the kernels begin executing, the frequency starts to rise along with an increase in power draw (Figure 11b). However, as soon as power reaches and exceeds the 300W TDP, the DVFS mechanism triggers a frequency drop until the GPU reaches a state where the power draw remains below 300W. Note that DVFS affects both GPUs but GPU-1 crosses the power limit at much lower frequencies than GPU-2 (GPU-1 runs at a median frequency of 1327 MHz, whereas GPU-2 is running at 1440 MHz). This difference in frequency translates into a 173 ms or 8% difference in performance between the two GPUs, even when the GPUs have the same temperature (44°C on average) and nearly the same average power draw. These timelines demonstrate that (1) differences in performance between GPUs is correlated to changes in frequency, and that (2) these frequency changes are driven by the PM infrastructure trying to keep power consumption under the GPU's TDP.

Overall, the findings on Vortex also reinforce **Takeaway 3** – water cooling reduces temperature variation but does not diminish performance or power variability.

F. SGEMM on TACC Frontera

Finally, we study TACC's Frontera cluster. As specified in Table I, Frontera is a mineral oil-cooled system with Turing-class RTX 5000 GPUs [42], [47]. We compare results on Frontera with that of the similarly sized Longhorn (air-cooled V100 GPUs) and Vortex (water-cooled V100 GPUs) clusters.

Figure 12 shows aggregated box plots for our metrics, grouped by cabinet number. Overall, Frontera shows 5% performance variation and 7% frequency variation. Since Quadro RTX GPUs have a faster boost clock [42] than Volta V100s, Frontera's range of operating frequencies is higher than other clusters. Turing-class GPUs also have a lower TDP (230W [42]), and almost all GPUs on Frontera operate within 5W of this limit, similar to Vortex and Longhorn. However, two GPUs in cabinet c197 are outliers – they run 1100-1600ms slower than the median kernel duration, 16°C cooler than the median temperature, and consume 59W lesser than the median power consumption. These outliers are similar to those on Longhorn, Summit, and Corona. Moreover, similar to Corona, presenting this data to system administrators led them to flag the the pump that stirs the mineral oil in this cabinet for

Fig. 10: Scatter plots for SGEMM on SNL Vortex showing (a) kernel duration (performance) and frequency have a strong negative correlation ($\rho = -0.98$) and (b) kernel duration and temperature are uncorrelated ($\rho = 0.04$).
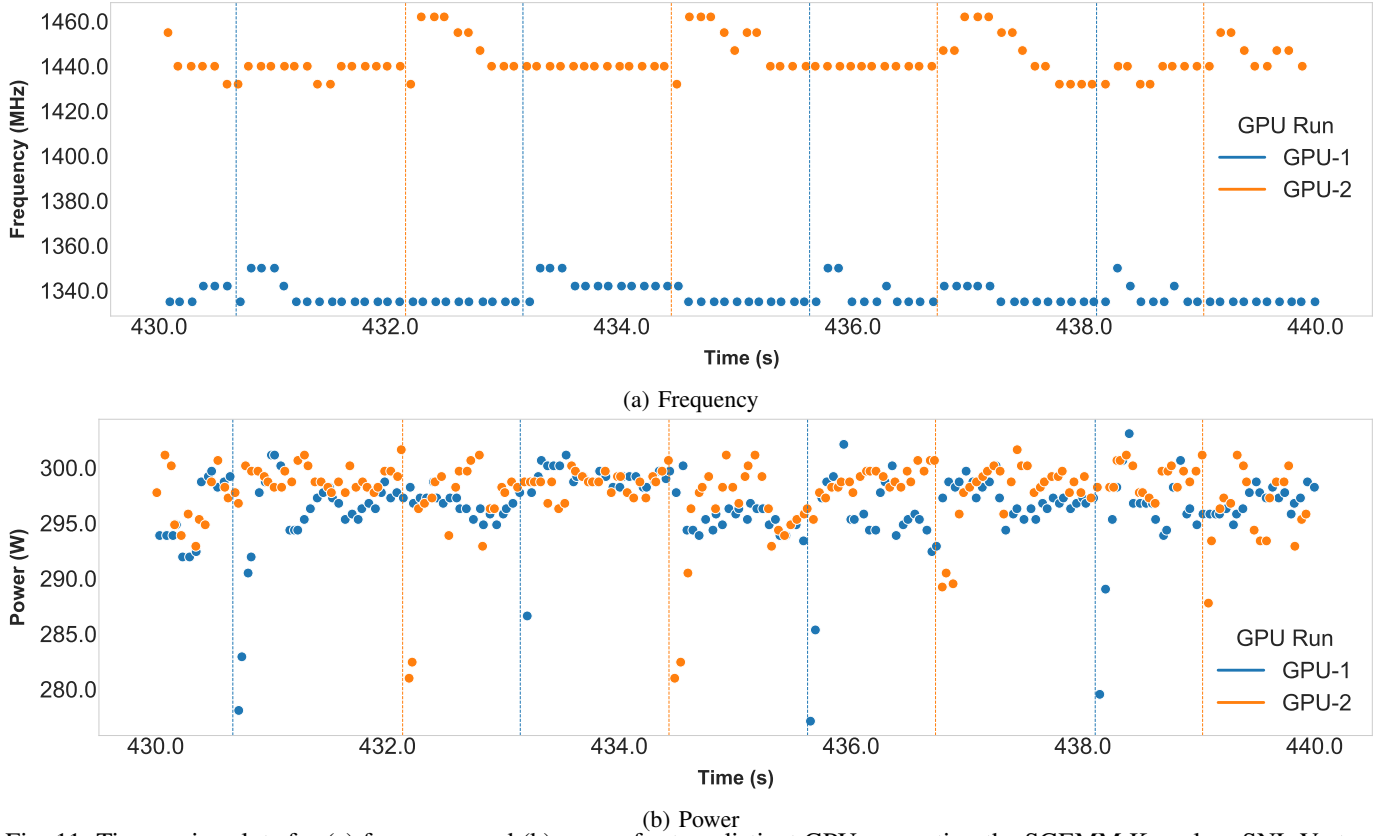

(a) Frequency

(b) Power

Fig. 11: Time-series plots for (a) frequency and (b) power for two distinct GPUs executing the SGEMM Kernel on SNL Vortex. Vertical lines indicate start of a new kernel.

additional investigation – further demonstrating the benefits of our approach in helping operators manage cluster operation (discussed further in Section VII). Similar to the water-cooled Vortex, Frontera has a narrow temperature spread ($Q3 - Q1 = 4°C$), but the median temperature ($76°C$) is much higher than Vortex ($46°C$). These temperature results indicate that mineral-oil cooling lies somewhere between air and water-cooling in effectiveness at reducing thermal variation.

Finally, Figure 13 presents scatterplots that correlate our metrics. Despite the c197 outliers, there is a strong correlation between performance and power/frequency ($\rho = -0.96$). However, Figure 13b shows that GPUs consuming similar

amounts of power can have widely varying temperatures ($\rho = -0.1$). This plot, together with observations from Vortex, suggests that temperature correlation with other metrics is weaker in systems cooled by water or mineral oil, as opposed to air-cooled clusters. Finally, the outlier on Frontera is concentrated to one cabinet and performs significantly worse, similar to outliers we observed on Summit (**Takeaway 2**).

V. VARIATION ACROSS APPLICATIONS

Overall, our study with SGEMM demonstrated that we can see significant variation across different scales, GPU vendors, and cooling approaches. However, SGEMM is a single compu-
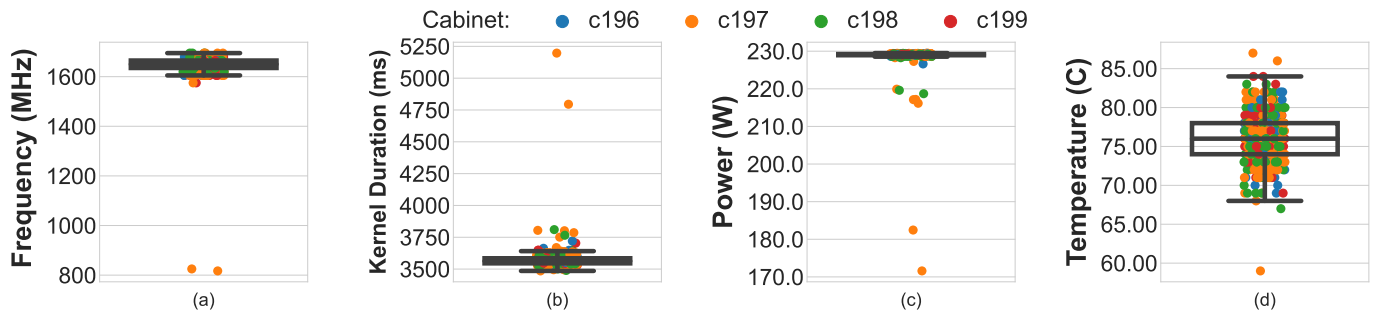
Fig. 12: SGEMM box plot summary results on Frontera for (a) frequency, (b) performance, (c) power, and (d) temperature.
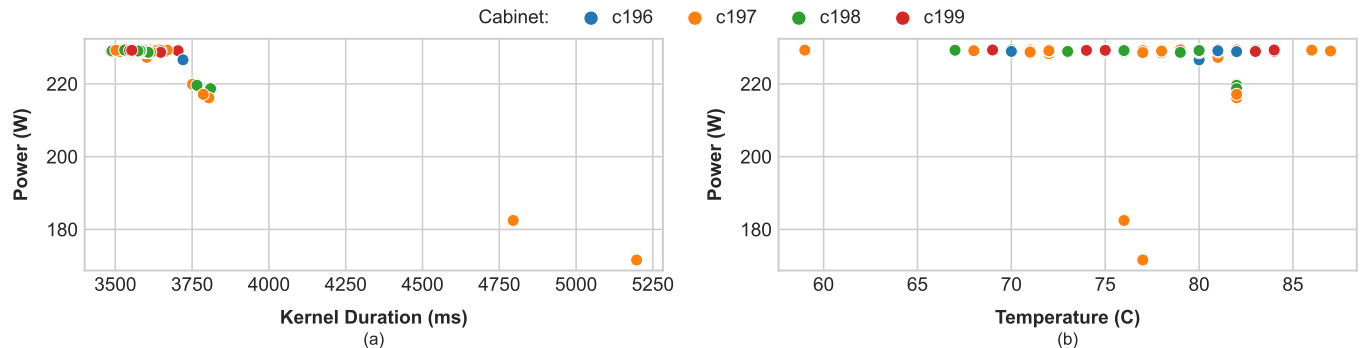

Fig. 13: Scatter plots for SGEMM on Frontera showing that (a) kernel duration (performance) and power have a strong negative correlation ($\rho = -0.96$) and (b) power and temperature are almost uncorrelated ($\rho = -0.1$)

tationally intensive kernel. Thus, although SGEMM is widely used in modern GPU applications, it may not be representative of larger applications. Accordingly, we next study different applications to determine whether variability is application-specific or not.

A. ResNet-50 on TACC Longhorn

Given their widespread use on GPU-rich clusters, it is important to understand performance variability for ML applications. Thus, we studied the ResNet-50 CNN for the most intensive training phase. We chose the 50-layer version because it is a stable, commonly used benchmark in the HPC community [69], [70]. In general, ResNet-50 is computationally intensive, especially for its kernels that perform convolution. However, ResNet-50 also performs other, more memory intensive operations, which reduce its overall computational intensity compared to SGEMM. This is confirmed by examining FU utilization: averaging across kernels, ResNet-50's was 5.4 (on a $1-10$ scale) while SGEMM's FU utilization was 10. This also aligns with observations from prior work on classifying GPU applications [71]. Our training set uses 1.2 million images from ImageNet with batch size 64. We define one training run as 500 iterations. Note that we did not complete training on our entire training set; 500 training iterations was sufficiently long to collect profiling data while training was stable. Finally, since ResNet-50 is commonly trained using multiple GPUs, we trained across four GPUs on one node and trained $3-4$ times per node. Like SGEMM, before collecting data we perform one warm up run.

Unlike SGEMM, ResNet-50 has ≈85 unique kernels and over 1.3 million kernels per run. Since 75% of these kernels run for less than 2ms, accurately measuring them is challenging (Section III). Thus, averaging kernel duration would not fairly characterize its overall performance. Hence, for ResNet-50 we use iteration duration instead of kernel duration as our performance metric; iteration duration is also more informative for HPC users training ML models. All other metrics are the same as SGEMM. However, we ignore the initialization kernels (e.g., NCCL) to avoid one-time startup costs.

Figure 14 presents ResNet-50's aggregated box plots on Longhorn for our 4 metrics. Unlike our SGEMM experiments on Longhorn (Section IV-B), ResNet-50 has little frequency variation and most nodes run at the max 1530MHz. However, ResNet-50 has 22% performance variation (Figure 14b), our largest observed variation. The temperature variation (Figure 14d) is similar to SGEMM's: a 30°C range. Surprisingly, unlike SGEMM there is significant power variation (104%, Figure 14c). We believe this reflects ResNet-50's more varied behavior across its different kernel types – although its convolution kernels are similar to SGEMM, ResNet-50's other, less computationally intense kernels have less onerous power demands. Consequently, the DVFS algorithm does not need to reduce voltage or frequency to remain beneath the TDP.

Figure 15 presents scatter plots of our distinct metrics. Figure 15a shows that most iterations complete within 100-150ms and GPUs have frequencies around 1530MHz. However, surprisingly, there are tails on both axes. In particular, several c008 runs have much lower frequency yet perform similarly to
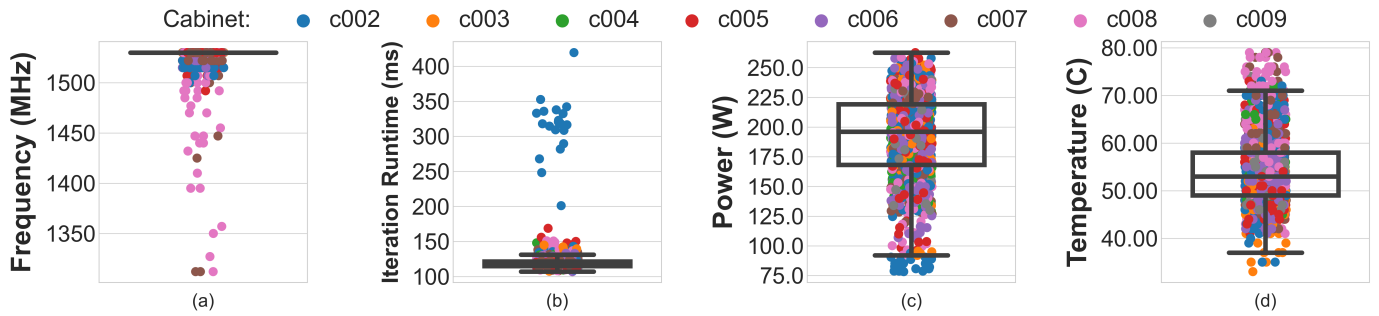
Fig. 14: Multi-GPU ResNet-50 Longhorn box plot summaries for (a) frequency, (b) performance, (c) power, and (d) temperature.
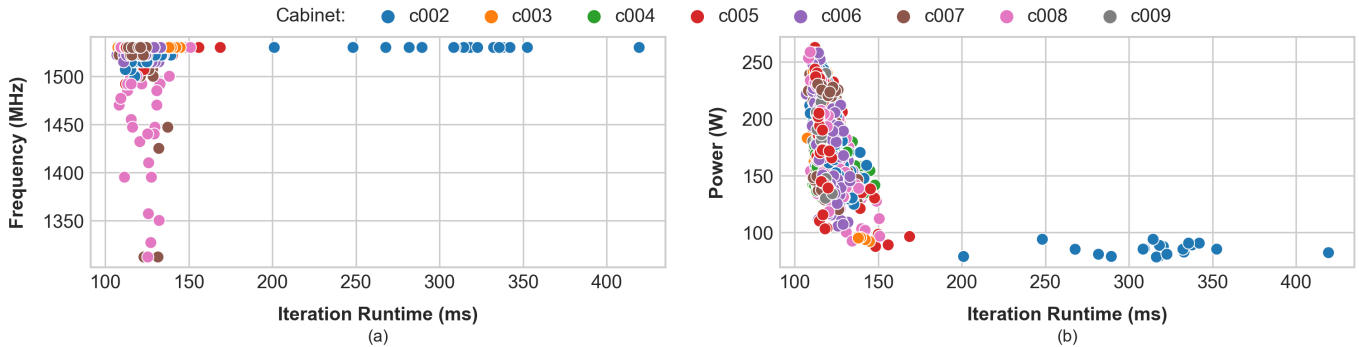

Fig. 15: Scatter plots for multi-GPU ResNet-50 experiments on Longhorn showing (a) iteration duration (performance) and frequency are almost uncorrelated ($\rho = -0.01$) and (b) iteration duration and power have a negative correlation ($\rho = -0.48$).

those running at the max 1530MHz. Meanwhile, several GPUs in c002 run at 1530MHz (Figure 15a), but perform poorly and consume much less power (Figure 15b). We expect c002's stragglers to consume more power than the tail of c008's runs, since $P_{dynamic} \propto frequency$. However, we observe the opposite: c002's worst performing runs consume much less power (as low as 76W) than c008 and other cabinets. This differs from Longhorn's SGEMM runs. Thus, ResNet-50 and SGEMM's variability and correlation differences on Longhorn suggest that variability is application-specific. Moreover, 8 of the 10 worst-performing GPUs for SGEMM were also outliers for ResNet, highlighting that variability is not transient and some GPUs consistently do not perform well.

We also ran ResNet as a single-GPU experiment with the same training set but proportionally scale the batch size to 16. Figure 16 presents the aggregated boxplots for the single-GPU experiments. Similar to the multi-GPU runs, power consumption stays well within TDP limits for all GPUs, hence they run at the max frequency of 1530MHz and observe little PM interference. The absolute power consumption values are lower than multi-GPU experiments, but we still see a significant median power variation (24%) across GPUs. Similarly, the absolute iteration duration values are also lower than the multi-GPU results, but single-GPU ResNet continues to demonstrate significant performance variability (14%). While the c002 outliers observed for SGEMM and multi-GPU ResNet still run at high temperatures (Figure 16(d)), they demonstrate lower degradation in performance compared to multi-GPU ResNet. This shows that multi-GPU jobs with a bulk synchronous pattern end up running as fast as the slowest GPU, leading

to larger performance degradation.

Takeaway 5: *ResNet-50 exhibits the highest performance variability (22% for multi-GPU, 14% in single-GPU) across all our benchmarks. Moreover, the difference in ResNet-50's and SGEMM's compute intensities, and the significant difference in performance-frequency correlations, suggests that variability is application-specific.*

B. BERT on TACC Longhorn

Given the increasing adoption of Transformer-based models [7], [19], [72]–[79], we also study another multi-GPU ML workload: pre-training of Bidirectional Encoder Representations from Transformers (BERT) [59]. BERT is widely used in Natural Language Processing (NLP) and is in the popular MLPerf suite [69], [70]. We used $BERT_{LARGE}$, which has 24 encoders with 16 bidirectional self-attention heads. Our training set was 30522 words and uses a batch size of 64. Like ResNet, we limit each training run to 250 iterations and run across all 4 GPUs in a node. We performed one warm up run and then performed measurements 5 times on each node. Overall we observed 53 nodes in Longhorn for BERT and collect the same metrics as ResNet. Similar to ResNet, we omit the initialization kernels to avoid one-time startup costs.

Figure 17 shows aggregated box plots on Longhorn for our 4 metrics. Compared to ResNet, BERT has less overall power consumption (Figure 17c): BERT's median power consumption is around 40W lower. This is expected because ResNet uses more compute intensive GEMMs, which increase its power consumption. In comparison, BERT's GEMMs are much less computationally intensive: although GEMMs make

Fig. 16: Single-GPU ResNet-50 box plot summaries for (a) frequency, (b) performance, (c) power, and (d) temperature.
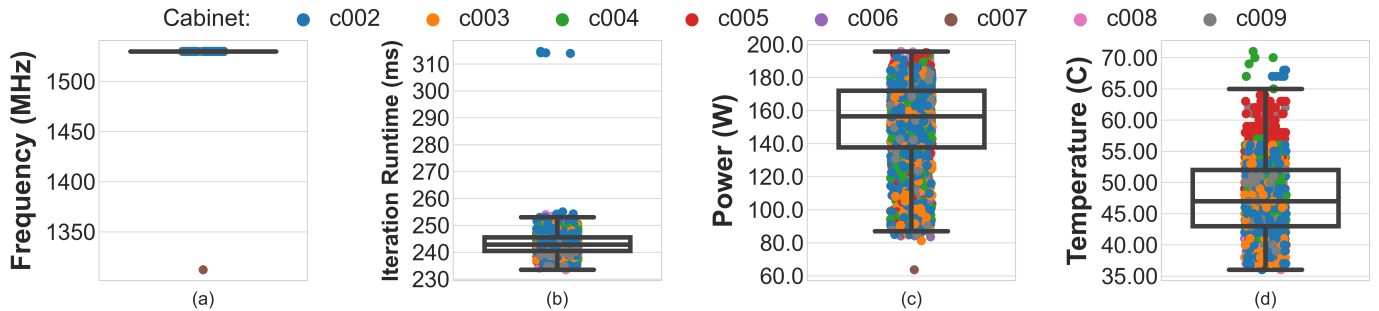

Fig. 17: Multi-GPU BERT box plot summaries on Longhorn for (a) frequency, (b) performance, (c) power, and (d) temperature.

up 30-65% of its total runtime, they only utilize 40-50% of the GPU [78], [80]–[82] – which decreases BERT's overall power consumption. Nonetheless, like ResNet there is large power variability ($\approx 87\%$). However, because BERT is less computationally intensive than ResNet, it has less performance variability ($\approx 8\%$) and fewer frequency outliers (Figure 17b and Figure 17a, respectively). Perhaps unsurprisingly, the performance outliers in Figure 17 in cabinet c002 on Longhorn are also outlier nodes for ResNet in Figure 14b. This suggests that the same GPUs perform poorly for both ML applications.

Takeaway 6: *Like ResNet-50, BERT exhibits large power variability. However, BERT's less computationally intensive GEMMs reduce performance variability (8%). Moreover, BERT's and ResNet-50's outlier nodes are the same.*

C. LAMMPS on TACC Longhorn

LAMMPS is a popular molecular dynamics (MD) application that simulates a variety of atomic and molecular systems [60]. Whether LAMMPS is compute or memory-bound depends on the selected settings and hardware [54], [83], [84]. We chose to use the REAXC setting [85] to simulate a chemical reaction. When REAXC is run in a distributed multi-GPU setup with large simulation sizes, LAMMPS is compute-intensive [54]. However, to compare LAMMPS variability with SGEMM, we ran it as a single-GPU experiment, with input configuration parameters (x, y, z) that determine GPU occupancy and are limited by the 16GB device memory on NVIDIA V100s. After careful tuning, we selected the (8,16,16) input size, which lead to high GPU utilization, while still remaining within device memory limits. Profiling shows that LAMMPS has $42\times$ higher DRAM utilization than ResNet, while ResNet kernels utilized FUs $4.3\times$ more

than LAMMPS. Thus, LAMMPS is memory-bound in our experiments.

Each LAMMPS run is composed broadly of 2 types of kernels, short-running ($\leq 60 \mu s$ long) and long-running kernels (4 unique kernels, 20-200ms long). Long-running kernels make up 98% of the total runtime of a LAMMPS job but there are 4 unique long kernels interspersed with short ones. Thus, similar to ResNet-50, median kernel duration is not a good performance measure. We therefore use the sum of all large kernel durations as our performance metric. All other metrics are measured as specified in Section III.

Figure 18 presents LAMMPS' aggregated box plots on Longhorn for our 4 metrics. Interestingly, median power for all LAMMPS jobs was ≤ 180 W (Figure 18c), in contrast to more compute-intensive applications such as SGEMM, which often touched V100 TDP. Moreover, similar to BERT and ResNet-50, frequency quickly saturates to the maximum value of 1530MHz (Figure 18a) and does not change throughout the course of the job. Additionally, performance varies by less than 1% in Figure 18b, which completely differs from BERT (8%), SGEMM (9%), and ResNet-50 (22%). However, we still observe power variability of 20% and temperature variability of 8°C between $Q1$ and $Q3$ (Figure 18d). These observations are in line with prior work [83] and emphasize that high energy consumption is undesirable in a memory-bound application because it is not accompanied by any significant increase in performance. In part, this happens because the GPU's memory frequency is lower than the compute frequency, and even if the application is memory-bound, this does not stress the TDP as much as the compute-heavy applications. Overall, this suggests that (i) SM frequency gets pinned in applications that are less compute-intensive (ii) performance of such applications
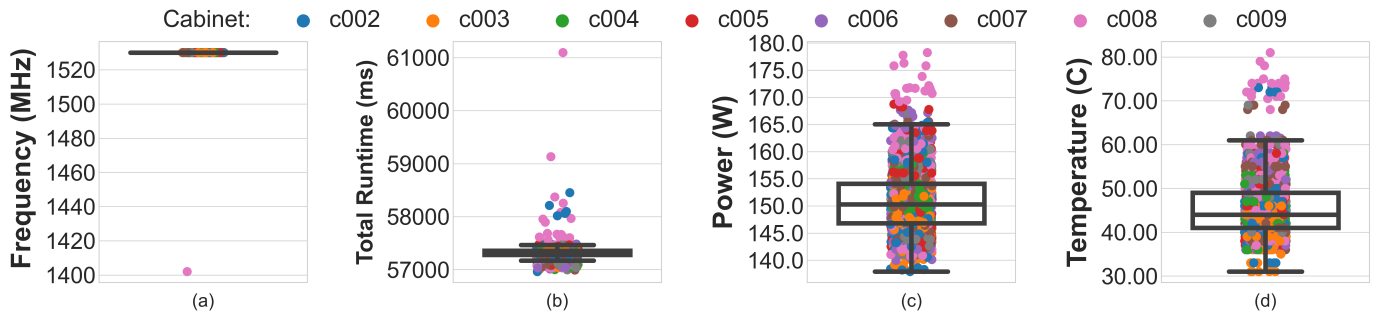
Fig. 18: LAMMPS box plot summary results on Longhorn for (a) frequency, (b) performance, (c) power, and (d) temperature.
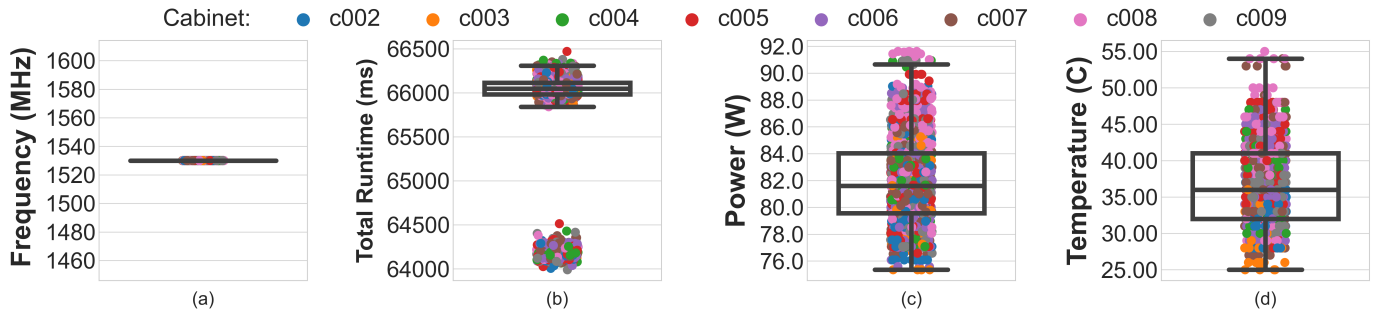

Fig. 19: PageRank box plot summary results on Longhorn for (a) frequency, (b) performance, (c) power, and (d) temperature.

is relatively more predictable with very low variability in application runtime and (iii) significant temperature and power variability are still observed across GPUs.

Takeaway 7: *Similar to **Takeaway 5**, LAMMPS' results show that variability is application-specific with memory-intensive applications seeing lower performance variance, but still having significant power and temperature variability.*

D. PageRank on TACC Longhorn

PageRank is a popular graph analytics algorithm that is used in recommender systems, search engines, social network analysis, and bibliometrics [86]. Although neither push- nor pull-based graph analytics algorithms always provide the best performance on GPUs [87]–[89], we focus on pull-based algorithms because they are more widely used. The most common PageRank algorithm utilizes Sparse matrix-Vector (SPMV) computation [61], [86]. Since PageRank's sparsity is input graph dependent, its memory access pattern can be highly irregular and often dwarfs the amount of compute, making it both memory bandwidth-bound and highly irregular. We chose an input graph that fully utilizes the SMs of a V100 GPU and provides sufficiently long kernels runtimes (Section III). We ran PageRank with *rajat30*, an undirected graph for circuit simulation [90]. The other configuration parameters are the same as SGEMM on Longhorn. Like LAMMPS, SPMV computations are also memory bound, but irregular. Consequently, they do not stress memory as much: LAMMPS has 4.24× higher DRAM utilization than PageRank. However, PageRank kernels had 61% memory dependency stalls, versus 7% for LAMMPS and 3% for SGEMM. PageRank also had negligible FU execution dependency stalls (12× less than SGEMM), showing that PageRank is not compute-bound.

Figure 19 summarizes Longhorn variability for PageRank. Similar to LAMMPS (Section V-C), PageRank has little frequency variation and very little performance variability (1%) across GPUs. However, we observed 22% variability in median power across runs and a temperature variation of 8°C between $Q1$ and $Q3$, as seen with LAMMPS.

Takeaway 8: *Like LAMMPS (**Takeaway 6**), PageRank has little frequency and performance variation. This suggests that memory-bound workloads can use worse-performing nodes without significant performance penalty.*

VI. VARIATION ACROSS TIME, POWER BUDGET

Next, we analyze how variability changes across days of the week and how the GPU power limit affects variability.

A. Variation Across Days of the Week

To determine if the variations we observed on these clusters hold over time, we ran SGEMM on Summit on each day of the week across a period of eight weeks. Figure 20 shows a consistent trend across each day of the week: around 8% variation. On Mondays, Wednesdays, and Fridays there is a higher concentration of GPUs which consume below 290W. However, even with more power outliers, performance on these days is very similar to the rest of the week and **Takeaways 1-3**, as is also shown in Figure 5.

Longhorn also shows similar trends, as shown in Figure 21. There is consistent performance variability across all days of the week. Compared to Summit, however, the variation is lower - around 3% each day. We also observe more outliers on Mondays, Thursdays and Fridays than other days of the week.
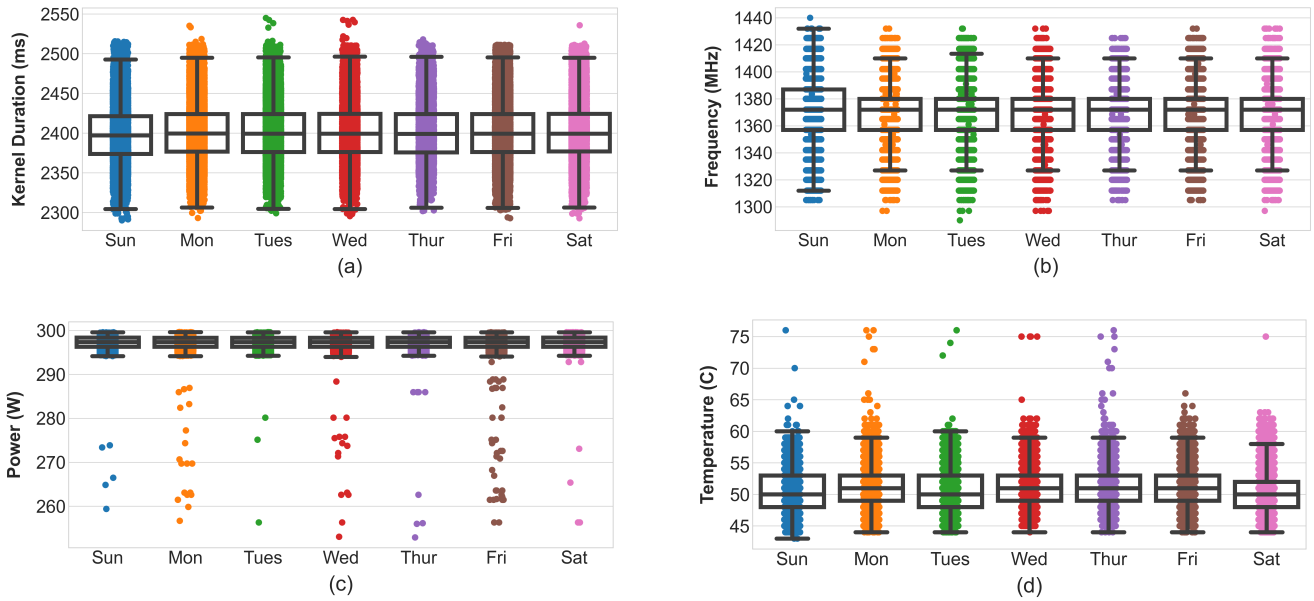
Fig. 20: Day of the week summary results for Summit. We see more power outliers on Mondays, Wednesdays, and Fridays.


Fig. 21: Day of the week summary results for Longhorn. We observed performance outliers on Mondays, Thursdays and Sundays.

Takeaway 9: *The variability we observe is consistent throughout the week, suggesting that regardless of when the experiments are run, our observations hold.*

B. Varying Power Limit

We also studied how varying a GPU's power limit affects variability. Varying the power limit (e.g., using `nvidia-smi`)

requires administrative access. Since we did not have administrator (root) access on the large compute clusters, we used the smaller CloudLab [46]. As discussed in Section III, CloudLab has 12 NVIDIA V100 GPUs. Figure 22 shows the performance variation as the power limit varies from 100-300W while running SGEMM. Interestingly, unlike prior
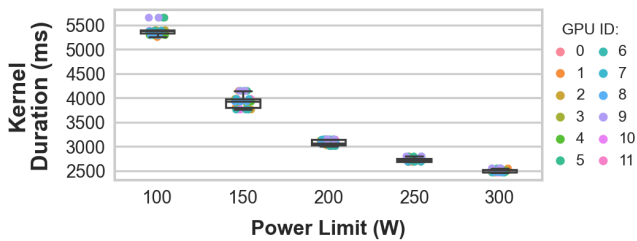
Fig. 22: Performance variation of SGEMM on NSF CloudLab GPUs when varying the power limit.

work [30] performance variation was similar when pinning and not pinning (Section IV-B). Rather, our variability results for a 300W limit are in line with those from the large clusters: 9% variation. Moreover, as expected, the kernel durations increase with lower power limits. However, variability and the number of outliers also increase with lower power limits. For example, at a 150W limit there is 18% variability versus 9% at a 300W limit. Our conversations with GPU manufacturers indicate that this is potentially occurring because GPU DVFS algorithms are less optimized for extremely low power budgets, but variability under power limits would become important if future exascale machines are operating under a varying power budget [91].

VII. Conclusions and Takeaways

Modern compute clusters are embracing accelerator-rich systems, especially GPUs. Although prior work has identified how variability in CPUs affects these systems, it is unclear how much GPU variability affects these systems. Thus, we conducted a detailed study and analysis of how GPU variability affects a wide range of modern HPC and scientific workloads across five computing centers of varying scales, cooling methods, and GPU vendors. Our results show that there is significant variability in these systems: up to 22% for the applications, with outliers up to **1.5×** slower than the median GPU. This demonstrates the need to embrace variability in future accelerator-based systems. We conclude by highlighting several implications and mitigation strategies that practitioners and researchers can pursue in the future:

Impact on Users: In addition to the performance variation, inter-experiment variability is also important. For example, when running SGEMM on Longhorn, 18% of the GPUs are 6-7% (or about 150ms) slower than the fastest GPUs. Thus, users running SGEMM-like single GPU experiments would have an 18% chance of being assigned a slower GPU. In Summit, again with SGEMM, 9% of the GPUs are 6-7% (or about 160ms) slower than the fastest GPUs. Thus, users running a single GPU experiment on Summit have a 9% chance of being assigned a slower GPU. Like Longhorn, Summit also has variance within specific GPUs. Users running multi-GPU experiments are even more likely to be assigned a slower GPU: if a user requests 4 GPUs on Longhorn, 40%-50% of the time they will be assigned a slower GPU. We also observed this in our multi-GPU BERT and ResNet-50 experiments.

Spatial Effects: For all of our experiments, we obtained exclusive access to a machine and all its GPUs. This eliminated any spatial effects from jobs running on neighboring GPUs. While this is the typical allocation mode in modern supercomputing clusters, spatial effects would be relevant for other scenarios like cloud computing or enterprise clusters [92] where GPUs are allocated individually. We plan to study both spatial and temporal (i.e., variability due to a preceding job run on the same GPU) effects in the future.

Blacklisting, Maintenance: Cluster operators can use our study to improve the cluster's operation and help develop strategies for better maintenance. For example, our study has already helped TACC's operators identify and perform targeted maintenance on problematic nodes with clearly underperforming GPUs in Frontera and Longhorn. Performing periodic variability benchmarking can help automate this.

Application-aware Frameworks: Since performance variation is application dependent, the next generation of HPC allocation frameworks should take application properties into account to mitigate variance. Similar to prior work [71], our profiling indicates that metrics like FU utilization, DRAM utilization, and memory stalls can be used by operators to classify applications and modify schedulers to assign medium- and high-compute intensity workloads on nodes with less variation. Similarly, memory-bound applications can be run on higher-variation nodes without incurring significant performance loss. However, this may change in future as thermal performance degrades below 14nm [93].

New Hardware and System Design: A major limiter to further improving efficiency is the lack of standards for exposing power information in modern accelerators. Thus, for future systems, designers can build on the insights generated by our benchmarks for current systems and apply co-design that makes the hardware, software, and runtime layers aware of the variance in the systems. To do this, we will need to design a standard for accelerators to expose PM information from the hardware to the software and runtime. Using this information, we can develop techniques for **global power management** that can enable optimal PM decisions across accelerators and further reduce performance variability.

REFERENCES

[1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conference on Computer Vision and Pattern Recognition*, ser. CVPR. IEEE, 2009, pp. 248–255.

[2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[3] D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. Chen, M. Chrzanowski, A. Coates, G. Diamos, E. Elsen, J. Engel, L. Fan, C. Fougner, A. Y. Hannun, B. Jun, T. Han, P. LeGresley, X. Li, L. Lin, S. Narang, A. Y. Ng, S. Ozair, R. Prenger, S. Qian, J. Raiman, S. Satheesh, D. Seetapun, S. Sengupta, C. Wang, Y. Wang, Z. Wang, B. Xiao, Y. Xie, D. Yogatama, J. Zhan, and Z. Zhu, "Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin," in *Proceedings of the 33nd International Conference on Machine Learning*, ser. ICML, 2016, pp. 173–182.

[4] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates *et al.*, "Deep speech: Scaling up end-to-end speech recognition," *arXiv preprint arXiv:1412.5567*, 2014.

[5] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangguan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S. yiin Chang, K. Rao, and A. Gruenstein, "Streaming end-to-end speech recognition for mobile devices," 2018.

[6] M. Ott, S. Edunov, D. Grangier, and M. Auli, "Scaling Neural Machine Translation," 2018.

[7] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean, "Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation," *CoRR*, vol. abs/1609.08144, 2016. [Online]. Available: http://arxiv.org/abs/1609.08144

[8] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko *et al.*, "Highly accurate protein structure prediction with alphafold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.

[9] J. Kates-Harbeck, A. Svyatkovskiy, and W. Tang, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning," *Nature*, vol. 568, no. 7753, pp. 526–531, 2019.

[10] Z. Fan and E. Ma, "Predicting orientation-dependent plastic susceptibility from static structure in amorphous solids via deep learning," *Nature communications*, vol. 12, no. 1, pp. 1–13, 2021.

[11] L. Casalino, A. C. Dommer, Z. Gaieb, E. P. Barros, T. Sztain, S.-H. Ahn, A. Trifan, A. Brace, A. T. Bogetti, A. Clyde, H. Ma, H. Lee, M. Turilli, S. Khalid, L. T. Chong, C. Simmerling, D. J. Hardy, J. D. Maia, J. C. Phillips, T. Kurth, A. C. Stern, L. Huang, J. D. McCalpin, M. Tatineni, T. Gibbs, J. E. Stone, S. Jha, A. Ramanathan, and R. E. Amaro, "AI-driven multiscale simulations illuminate mechanisms of SARS-CoV-2 spike dynamics," *The International Journal of High Performance Computing Applications*, p. 10943420211006452, 2021.

[12] top500.org, "TOP500 List," 2021, https://www.top500.org/lists/top500/list/2021/11/.

[13] J. Fowers, K. Ovtcharov, M. Papamichael, T. Massengill, M. Liu, D. Lo, S. Alkalay, M. Haselman, L. Adams, M. Ghandi, S. Heil, P. Patel, A. Sapek, G. Weisz, L. Woods, S. Lanka, S. K. Reinhardt, A. M. Caulfield, E. S. Chung, and D. Burger, "A Configurable Cloud-scale DNN Processor for Real-time AI," in *Proceedings of the 45th Annual International Symposium on Computer Architecture*, ser. ISCA. Piscataway, NJ, USA: IEEE Press, 2018, pp. 1–14. [Online]. Available: https://doi.org/10.1109/ISCA.2018.00012

[14] Tesla, "Ahead of 'Dojo,' Tesla Reveals Its Massive Precursor Supercomputer," 2021, https://www.hpcwire.com/2021/06/22/ahead-of-dojo-tesla-reveals-its-massive-precursor-supercomputer/.

[15] Oak Ridge National Laboratory, "Frontier, Direction of Discovery," https://www.olcf.ornl.gov/frontier/, 2021.

[16] D. Abts, J. Ross, J. Sparling, M. Wong-VanHaren, M. Baker, T. Hawkins, A. Bell, J. Thompson, T. Kahsai, G. Kimmell, J. Hwang, R. Leslie-Hurd, M. Bye, E. Creswick, M. Boyd, M. Venigalla, E. Laforge, J. Purdy, P. Kamath, D. Maheshwari, M. Beidler, G. Rosseel, O. Ahmad, G. Gagarin, R. Czekalski, A. Rane, S. Parmar, J. Werner, J. Sproch, A. Macias, and B. Kurtz, "Think Fast: A Tensor Streaming Processor (TSP) for Accelerating Deep Learning Workloads," in *Proceedings of ACM/IEEE 47th Annual International Symposium on Computer Architecture*, ser. ISCA, 2020, pp. 145–158.

[17] P. Alcorn, "Hot Chips 2017: A Closer Look At Google's TPU v2," http://www.tomshardware.com/news/tpu-v2-google-machine-learning,35370.html, September 2017.

[18] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, R. Boyle, P.-l. Cantin, C. Chao, C. Clark, J. Coriell, M. Daley, M. Dau, J. Dean, B. Gelb, T. V. Ghaemmaghami, R. Gottipati, W. Gulland, R. Hagmann, C. R. Ho, D. Hogberg, J. Hu, R. Hundt, D. Hurt, J. Ibarz, A. Jaffey, A. Jaworski, A. Kaplan, H. Khaitan, D. Killebrew, A. Koch, N. Kumar, S. Lacy, J. Laudon, J. Law, D. Le, C. Leary, Z. Liu, K. Lucke, A. Lundin, G. MacKean, A. Maggiore, M. Mahony, K. Miller, R. Nagarajan, R. Narayanaswami, R. Ni, K. Nix, T. Norrie, M. Omernick, N. Penukonda, A. Phelps, J. Ross, M. Ross, A. Salek, E. Samadiani, C. Severn, G. Sizikov, M. Snelham, J. Souter, D. Steinberg, A. Swing, M. Tan, G. Thorson, B. Tian, H. Toma, E. Tuttle, V. Vasudevan, R. Walter, W. Wang, E. Wilcox, and D. H. Yoon, "In-Datacenter Performance Analysis of a Tensor Processing Unit," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, ser. ISCA. New York, NY, USA: ACM, 2017, pp. 1–12. [Online]. Available: http://doi.acm.org/10.1145/3079856.3080246

[19] N. P. Jouppi, D. H. Yoon, M. Ashcraft, M. Gottscho, T. B. Jablin, G. Kurian, J. Laudon, S. Li, P. Ma, X. Ma, N. Patil, S. Prasad, C. Young, Z. Zhou, and D. Patterson, "Ten Lessons from Three Generations Shaped Google's TPUv4i," in *Proceedings of the 48th Annual International Symposium on Computer Architecture*, ser. ISCA, 2021.

[20] M. La and A. Chien, "Cerebras Systems: Journey to the Wafer-Scale Engine," University of Chicago, Tech. Rep., 2020.

[21] S. Lie, "Multi-million core, multi-wafer ai cluster," in *2021 IEEE Hot Chips 33 Symposium (HCS)*. IEEE Computer Society, 2021, pp. 1–41.

[22] L. R. M. Mohan, A. Marshall, S. Maddrell-Mander, D. O'Hanlon, K. Petridis, J. Rademacker, V. Rege, and A. Titterton, "Studying the potential of Graphcore IPUs for applications in particle physics," *arXiv preprint arXiv:2008.09210*, 2020.

[23] Z. Jia, B. Tillman, M. Maggioni, and D. P. Scarpazza, "Dissecting the Graphcore IPU Architecture via Microbenchmarking," 2019. [Online]. Available: https://arxiv.org/abs/1912.03413

[24] B. Acun, A. Langer, E. Meneses, H. Menon, O. Sarood, E. Totoni, and L. V. Kalé, "Power, Reliability, and Performance: One System to Rule them All," *Computer*, vol. 49, no. 10, pp. 30–37, 2016.

[25] D. Chasapis, M. Casas, M. Moretó, M. Schulz, E. Ayguadé, J. Labarta, and M. Valero, "Runtime-Guided Mitigation of Manufacturing Variability in Power-Constrained Multi-Socket NUMA Nodes," in *Proceedings of the 2016 International Conference on Supercomputing*, ser. ICS '16, 2016.

[26] D. Chasapis, M. Moretó, M. Schulz, B. Rountree, M. Valero, and M. Casas, "Power Efficient Job Scheduling by Predicting the Impact of Processor Manufacturing Variability," in *Proceedings of the ACM International Conference on Supercomputing*, ser. ICS '19, 2019, p. 296–307. [Online]. Available: https://doi.org/10.1145/3330345.3330372

[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 *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC, 2015. [Online]. Available: https://doi.org/10.1145/2807591.2807638

[28] T. Patel, A. Wagenhäuser, C. Eibel, T. Hönig, T. Zeiser, and D. Tiwari, "What does Power Consumption Behavior of HPC Jobs Reveal? : Demystifying, Quantifying, and Predicting Power Consumption Characteristics," in *IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS, 2020, pp. 799–809.

[29] D. Skinner and W. Kramer, "Understanding the causes of performance

variability in HPC workloads," in *Proceedings of the IEEE Workload Characterization Symposium*, ser. IISWC, 2005, pp. 137–149.

[30] T. Scogland, J. Azose, D. Rohr, S. Rivoire, N. Bates, and D. Hackenberg, "Node Variability in Large-Scale Power Measurements: Perspectives from the Green500, Top500 and EEHPCWG," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '15. New York, NY, USA: Association for Computing Machinery, 2015. [Online]. Available: https://doi.org/10.1145/2807591.2807653

[31] T. R. Scogland, C. P. Steffen, T. Wilde, F. Parent, S. Coghlan, N. Bates, W.-c. Feng, and E. Strohmaier, "A Power-Measurement Methodology for Large-Scale, High-Performance Computing," in *Proceedings of the 5th ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 149–159. [Online]. Available: https://doi.org/10.1145/2568088.2576795

[32] B. Acun, P. Miller, and L. V. Kale, "Variation Among Processors Under Turbo Boost in HPC Systems," in *Proceedings of the 2016 International Conference on Supercomputing*, ser. ICS '16. New York, NY, USA: Association for Computing Machinery, 2016. [Online]. Available: https://doi.org/10.1145/2925426.2926289

[33] B. Acun and L. V. Kale, "Mitigating Processor Variation through Dynamic Load Balancing," in *IEEE International Parallel and Distributed Processing Symposium Workshops*, ser. IPDPSW, 2016, pp. 1073–1076.

[34] B. Acun, E. K. Lee, Y. Park, and L. V. Kale, "Support for Power Efficient Proactive Cooling Mechanisms," in *IEEE 24th International Conference on High Performance Computing*, ser. HiPC, 2017, pp. 94–103.

[35] H. Menon, B. Acun, S. G. De Gonzalo, O. Sarood, and L. Kalé, "Thermal aware automated load balancing for HPC applications," in *IEEE International Conference on Cluster Computing*, ser. CLUSTER, 2013, pp. 1–8.

[36] K. Zhang, S. Ogrenci-Memik, G. Memik, K. Yoshii, R. Sankaran, and P. Beckman, "Minimizing Thermal Variation Across System Components," in *IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS, 2015, pp. 1139–1148.

[37] J. Coplin and M. Burtscher, "Energy, Power, and Performance Characterization of GPGPU Benchmark Programs," in *IEEE International Parallel and Distributed Processing Symposium Workshops*, ser. IPDPSW, May 2016, pp. 1190–1199. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/IPDPSW.2016.164

[38] Y. Jiao, H. Lin, P. Balaji, and W. Feng, "Power and Performance Characterization of Computational Kernels on the GPU," in *IEEE/ACM International Conference on Green Computing and Communications & International Conference on Cyber, Physical and Social Computing*, 2010, pp. 221–228.

[39] G. Ostrouchov, D. Maxwell, R. A. Ashraf, C. Engelmann, M. Shankar, and J. H. Rogers, "GPU Lifetimes on Titan Supercomputer: Survival Analysis and Reliability," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '20. IEEE Press, 2020.

[40] NVIDIA, "NVIDIA Tesla V100 GPU Architecture The World's Most Advanced Data Center GPU," http://www.nvidia.com/object/volta-architecture-whitepaper.html , NVIDIA, 2017.

[41] AMD, "AMD Radeon Instinct MI60," https://www.amd.com/system/files/documents/radeon-instinct-mi60-datasheet.pdf, 2018.

[42] NVIDIA, "NVIDIA Turing GPU Architecture Graphics Reinvented," https://images.nvidia.com/aem-dam/en-zz/Solutions/design-visualization/technologies/turing-architecture/NVIDIA-Turing-Architecture-Whitepaper.pdf, 2018.

[43] S. Bharadwaj, S. Das, K. Mazumdar, B. Beckmann, and S. Kosonocky, "Predict; Do not React for Enabling Efficient Fine Grain DVFS in GPUs," 2022. [Online]. Available: https://arxiv.org/abs/2205.00121

[44] R. Ge, R. Vogt, J. Majumder, A. Alam, M. Burtscher, and Z. Zong, "Effects of Dynamic Voltage and Frequency Scaling on a K20 GPU," in *42nd International Conference on Parallel Processing*, ser. ICPP, 2013, pp. 826–833.

[45] P. Meinerzhagen, C. Tokunaga, A. Malavasi, V. Vaidya, A. Mendon, D. Mathaikutty, J. Kulkarni, C. Augustine, M. Cho, S. Kim, G. Matthew, R. Jain, J. Ryan, C.-C. Peng, S. Paul, S. Vangal, B. P. Esparza, L. Cuellar, M. Woodman, B. Iyer, S. Maiyuran, G. Chinya, C. Zou, Y. Liao, K. Ravichandran, H. Wang, M. Khellah, J. Tschanz, and V. De, "An energy-efficient graphics processor featuring fine-grain DVFS with integrated voltage regulators, execution-unit turbo, and retentive sleep in 14nm tri-gate CMOS," in *IEEE International Solid - State Circuits Conference*, ser. ISSCC, 2018, pp. 38–40.

[46] D. Duplyakin, R. Ricci, A. Maricq, G. Wong, J. Duerig, E. Eide, L. Stoller, M. Hibler, D. Johnson, K. Webb, A. Akella, K. Wang, G. Ricart, L. Landweber, C. Elliott, M. Zink, E. Cecchet, S. Kar, and P. Mishra, "The Design and Operation of Cloudlab," in *Proceedings of the 2019 USENIX Conference on Usenix Annual Technical Conference*, ser. USENIX ATC, 2019, p. 1–14.

[47] TACC, "Texas Advanced Computing Center," https://www.tacc.utexas.edu/, 2021.

[48] SNL, "Sandia National Laboratories," https://www.sandia.gov/asc/advanced-simulation-and-computing/computational-systems/haaps/, 2022.

[49] ORNL, "Oak Ridge National Laboratory," https://www.olcf.ornl.gov/summit/, 2022.

[50] LLNL, "Lawrence Livermore National Laboratory," https://hpc.llnl.gov/hardware/compute-platforms/corona, 2022.

[51] M. Chandrasekhar, "AMD Radeon Community Update: More Control Over GPU Power and Performance, Enhanced Thermal Monitoring, Maximized Performance," https://community.amd.com/t5/gaming/amd-radeon-community-update-more-control-over-gpu-power-and/ba-p/418629, 2019.

[52] SafeTemp, "AMD Radeon RX Vega 11 Max Temp," https://safetemp.blogspot.com/2021/11/amd-radeon-rx-vega-11-max-temp.html, November 2021.

[53] "GPU Fan Speed Measurements," https://groups.google.com/g/cloudlab-users/c/mBhDK-MoWs8/m/KkijNWgeAAAJ.

[54] Lawrence Livermore National Labs, "CORAL-2 Benchmarks," https://asc.llnl.gov/coral-2-benchmarks, 2020.

[55] T. Vijayaraghavan, Y. Eckert, G. H. Loh, M. J. Schulte, M. Ignatowski, B. M. Beckmann, W. C. Brantley, J. L. Greathouse, W. Huang, A. Karunanithi, O. Kayiran, M. Meswani, I. Paul, M. Poremba, S. Raasch, S. K. Reinhardt, G. Sadowski, and V. Sridharan, "Design and analysis of an apu for exascale computing," in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2017, pp. 85–96.

[56] O. Villa, D. R. Johnson, M. Oconnor, E. Bolotin, D. Nellans, J. Luitjens, N. Sakharnykh, P. Wang, P. Micikevicius, A. Scudiero, S. W. Keckler, and W. J. Dally, "Scaling the power wall: A path to exascale," in *SC '14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2014, pp. 830–841.

[57] NVIDIA, "cuBLAS," https://developer.nvidia.com/cublas.

[58] "ResNet: Deep residual networks pre-trained on ImageNet," https://pytorch.org/hub/pytorch_vision_resnet/.

[59] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, ser. NAACL-HLT, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 4171–4186. [Online]. Available: https://doi.org/10.18653/v1/n19-1423

[60] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, "LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales," *Comp. Phys. Comm.*, vol. 271, p. 108171, 2022.

[61] S. Che, B. M. Beckmann, S. K. Reinhardt, and K. Skadron, "Pannotia: Understanding Irregular GPGPU Graph Applications," in *IEEE International Symposium on Workload Characterization*, ser. IISWC, Sept 2013, pp. 185–195.

[62] NVIDIA Corp, "Profiler User's Guide," https://docs.nvidia.com/cuda/profiler-users-guide/index.html, 2018.

[63] Advanced Micro Devices, Inc., "System Management Interface (SMI)," https://rocmdocs.amd.com/en/latest/ROCm_System_Managment/ROCm-System-Managment.html, 2022.

[64] AMD, "hipBLAS Library," https://github.com/ROCmSoftwarePlatform/hipBLAS, 2020.

[65] NVIDIA, "NVIDIA Tesla V100 GPU Architecture The World's Most Advanced Data Center GPU," http://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf, 2017.

[66] Y. H. Chou, C. Ng, S. Cattell, J. Intan, M. D. Sinclair, J. Devietti, T. G. Rogers, and T. M. Aamodt, "Deterministic Atomic Buffering," in *53rd*

Annual IEEE/ACM International Symposium on Microarchitecture, ser. MICRO, Oct 2020, pp. 981–995.

[67] NVIDIA, "cuDNN Frontend API," https://github.com/NVIDIA/cudnn-frontend#execution-plan-caching, 2021.

[68] "Job Step Viewer - Summit ORNL," https://jobstepviewer.olcf.ornl.gov/summit/871957-1.

[69] P. Mattson, V. J. Reddi, C. Cheng, C. Coleman, G. Diamos, D. Kanter, P. Micikevicius, D. Patterson, G. Schmuelling, H. Tang, G.-Y. Wei, and C.-J. Wu, "MLPerf: An Industry Standard Benchmark Suite for Machine Learning Performance," *IEEE Micro*, vol. 40, no. 2, pp. 8–16, 2020.

[70] P. Mattson, C. Cheng, G. Diamos, C. Coleman, P. Micikevicius, D. Patterson, H. Tang, G.-Y. Wei, P. Bailis, V. Bittorf, D. Brooks, D. Chen, D. Dutta, U. Gupta, K. Hazelwood, A. Hock, X. Huang, D. Kang, D. Kanter, N. Kumar, J. Liao, D. Narayanan, T. Oguntebi, G. Pekhimenko, L. Pentecost, V. Janapa Reddi, T. Robie, T. St John, C.-J. Wu, L. Xu, C. Young, and M. Zaharia, "MLPerf Training Benchmark," in *Proceedings of Machine Learning and Systems*, I. Dhillon, D. Papailiopoulos, and V. Sze, Eds., vol. 2, 2020, pp. 336–349. [Online]. Available: https://proceedings.mlsys.org/paper/2020/file/02522a2b2726fb0a03bb19f2d8d9524d-Paper.pdf

[71] J. Guerreiro, A. Ilic, N. Roma, and P. Tomás, "DVFS-aware application classification to improve GPGPUs energy efficiency," *Parallel Computing*, vol. 83, pp. 93–117, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167819118300243

[72] D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," in *Proceedings of the Third International Conference on Learning Representation*, ser. ICLR, 2015.

[73] G. M. Correia, V. Niculae, and A. F. T. Martins, "Adaptively Sparse Transformers," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, ser. EMNLP-IJCNP, 2019.

[74] N. Kitaev, Łukasz Kaiser, and A. Levskaya, "Reformer: The Efficient Transformer," in *Proceedings of the Eighth International Conference on Learning Representation*, ser. ICLR, 2020.

[75] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A Lite BERT for Self-supervised Learning of Language Representations," in *Proceedings of the Seventh International Conference on Learning Representation*, ser. ICLR, 2019.

[76] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, ser. EMNLP. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 1412–1421. [Online]. Available: https://www.aclweb.org/anthology/D15-1166

[77] T. Shen, T. Zhou, G. Long, J. Jiang, and C. Zhang, "Bi-Directional Block Self-Attention for Fast and Memory-Efficient Sequence Modeling," in *Proceedings of the Sixth International Conference on Learning Representation*, ser. ICLR, 2018.

[78] M. Shoeybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, "Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism," 2019.

[79] J. Zhu, Y. Xia, L. Wu, T. He, Di amd Qin, W. Zhou, H. Li, and T. Liu, "Incorporating BERT into Neural Machine Translation," in *Proceedings of the Eighth International Conference on Learning Representation*, ser. ICLR, 2020.

[80] A. Ivanov, N. Dryden, T. Ben-Nun, S. Li, and T. Hoefler, "Data Movement Is All You Need: A Case Study on Optimizing Transformers," in *Proceedings of Machine Learning and Systems*, A. Smola, A. Dimakis, and I. Stoica, Eds., vol. 3, 2021, pp. 711–732. [Online]. Available: https://proceedings.mlsys.org/paper/2021/file/c9e1074f5b3f9fc8ea15d152add07294-Paper.pdf

[81] S. Pati, S. Aga, N. Jayasena, and M. D. Sinclair, "Demystifying BERT: Implications for Accelerator Design," 2021.

[82] A. H. Zadeh, Z. Poulos, and A. Moshovos, "Deep Learning Language Modeling Workloads: Where Time Goes on Graphics Processors," in *IEEE International Symposium on Workload Characterization*, ser. IISWC, IEEE. Washington, DC, USA: IEEE Computer Society, 2019, pp. 131–142.

[83] V. Vecher, V. Nikolskii, and V. Stegailov, "GPU-Accelerated Molecular Dynamics: Energy Consumption and Performance," in *Supercomputing*, V. Voevodin and S. Sobolev, Eds. Cham: Springer International Publishing, 2016, pp. 78–90.

[84] H. Aktulga, J. Fogarty, S. Pandit, and A. Grama, "Parallel reactive molecular dynamics: Numerical methods and algorithmic techniques," *Parallel Computing*, vol. 38, no. 4, pp. 245–259, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167819111001074

[85] T.-R. Shan, R. Wixom, C. Yarrington, and A. Thompson, "Atomistic Simulation of Initiation in Hexanitrostilbene," in *APS Shock Compression of Condensed Matter Meeting Abstracts*, ser. APS Meeting Abstracts, Jun. 2015, p. Y5.003.

[86] L. Page, S. Brin, R. Motwani, and T. Winograd, "The PageRank Citation Ranking: Bringing Order to the Web." November 1999, previous number = SIDL-WP-1999-0120. [Online]. Available: http://ilpubs.stanford.edu:8090/422/

[87] M. Besta, M. Podstawski, L. Groner, E. Solomonik, and T. Hoefler, "To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations," in *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 93–104. [Online]. Available: https://doi.org/10.1145/3078597.3078616

[88] G. Salvador, W. H. Darvin, M. Huzaifa, J. Alsop, M. D. Sinclair, and S. V. Adve, "Specializing Coherence, Consistency, and Push/Pull for GPU Graph Analytics," in *IEEE International Symposium on Performance Analysis of Systems and Software*, ser. ISPASS, 2020.

[89] T. Sorensen, S. Pai, and A. F. Donaldson, "One Size Doesn't Fit All: Quantifying Performance Portability of Graph Applications on GPUs," in *IEEE International Symposium on Workload Characterization*, ser. IISWC, November 2019.

[90] SuiteSparse Matrix Collection, "Rajat/rajat30 circuit simulation matrix," https://sparse.tamu.edu/Rajat/rajat30, 2006.

[91] Office of Advanced Scientific Computing Research, "Energy Efficient Hardware and Software Design Constrained to a Power Envelope of 20 Megawatts," https://science.osti.gov/-/media/grants/pdf/lab-announcements/2013/LAB_13-02.pdf, 2021.

[92] M. Jeon, S. Venkataraman, A. Phanishayee, J. Qian, W. Xiao, and F. Yang, "Analysis of Large-Scale Multi-Tenant GPU Clusters for DNN Training Workloads," in *2019 USENIX Annual Technical Conference*, ser. USENIX ATC, 2019, pp. 947–960.

[93] K. Heyman, "DRAM Thermal Issues Reach Crisis Point," https://semiengineering.com/dram-thermal-issues-reach-crisis-point/, 2022.

A. Abstract

This artifact provides container specifications and scripts for reproducing experiments that measure and benchmark variability of GPUs across different applications. The artifact can be used to run different benchmarks across machine learning, molecular dynamics and graph analytics. We also include scripts that use vendor-specific profiling tools to measure power (W), temperature ($^\circ$ C), frequency (MHz) and performance (kernel runtime/iteration duration/total runtime depending on application). The experiments, their corresponding paper sections and the base framework/tools used to run them are summarised in Table III.

Experiment	Section	Base Framework/ Tools Used
SGEMM on NVIDIA GPUs [57]	IV-A	CUDA 10.1, nvcc 10.1 NVProf 10.1, gcc 4.8.5
SGEMM on AMD GPUs [57]	IV-D	ROCm 4.0.1
ResNet-50 [58]	V-A	CUDA 10.1, NVProf 10.1 PyTorch 1.9
BERT [59]	V-B	CUDA 10.1, NVProf 10.1 PyTorch 1.9
LAMMPS [60]	V-C	CUDA 10.1, NVProf 10.1, nvcc 10.1
PageRank [61]	V-D	CUDA 10.1, NVProf 10.1 nvcc 10.1

TABLE III: Summary of experiments in artifact.

B. Artifact Metadata

- **Persistent ID**: DOI: 10.5281/zenodo.7010207
- **GitHub URL**: https://github.com/hal-uw/gpu_variability_sc22_artifact
- **Artifact name**: Artifact: Characterizing Variability in Large-Scale, Accelerator-Rich Systems
- **Run-time environment**: Singularity
- **Dependencies**: See Section A-C1

C. Description

The artifact is organized into directories per-experiment (and GPU vendor) to allow any subset of the paper's experiments to be reproduced independent of one another. We include separate singularity scripts for each experiment. These scripts pull container images, install all dependencies and compile library code into respective application binaries. Directions to run each application using these scripts can be found in the respective application's `README.md` file. At the end of the `README`s, we also provide troubleshooting/FAQ sections for common errors that users could encounter, and how to resolve them.

1) Prerequisites: The following hardware and software dependencies must be satisfied to run the experiments successfully:

- Machine/compute node with an NVIDIA/AMD GPU
- Relevant GPU drivers are installed
- Singularity is installed

All the build and run scripts have been tested with Singularity v3.7.2-4.el7a on either NVIDIA V100 (Volta), RTX 5000 (Turing), or AMD MI60 GPUs.

2) Installation: The artifact repository can be accessed through the persistent DOI https://doi.org/10.5281/zenodo.7010207. The main `README` file in the artifact provides step-wise instructions on cloning the repository and points to per-application instructions for running different experiments.

3) Customization: Each experiment can be customized for different GPU architectures, different input sizes, and different iterations/time steps. For instance, all of the NVIDIA scripts assume a Volta V100 GPU by default, but the scripts can be configured differently if the user has a different GPU. The `Prerequisites` section in each application-specific `README` provides further details on how to customize respective experiments.

APPENDIX B
OUTLIER ANALYSIS: SGEMM ON ORNL SUMMIT

A. Row H Results

Our results in Section IV-C show that outliers exist in Summit. To understand these outliers better, we more deeply examined the results within Summit's row H and found additional interesting trends. Behavior varies significantly across the different columns in row H, despite being located near one another. For example, Figure 23 shows that 7 of the 29 columns in row H have no frequency, performance, power, or temperature outliers. In the remaining columns that do have at least one outlier, the vast majority of the outliers come from columns 13, 14, 28, 33, 36, and 50. Columns 33 and 36 in particular have a very large number of outliers; both columns have frequency, performance, power, and temperature outliers (the remaining columns have outliers in 1-3 of these categories).

Figure 24 shows that despite several columns having outliers in multiple metrics, these outliers are not always well correlated. Note that, since we do not have sole access to Summit and are not guaranteed access to a specific node when enqueuing a job, the number of runs for different nodes in the scatter plots ranges from 1 (column50-node08, column50-node12, and column50-node15) to 16 (column28-node13). Nevertheless, correlating multiple metrics yields further interesting insights. Perhaps unsurprisingly, performance and frequency (Figure 24a) are often well correlated – i.e., as frequency in a given GPU decreases, performance also increases. However, the correlation between performance and temperature (Figure 24b), performance and power (Figure 24c), and power and temperature (Figure 24d) show significantly less correlation. In particular, the power outliers show fascinating trends when correlated with performance: although the outliers all take approximately 2510 ms to complete, their power consumption across the different nodes is very different – ranging from 250 - 285W, despite running for the same amount of time. Interestingly, every GPU in the scatter plots appears to have at least one outlier. For example, rowh-col28-n13-3, shown in purple in Figure 24c, appears to be representative of a GPU

(a) Performance

(b) Frequency
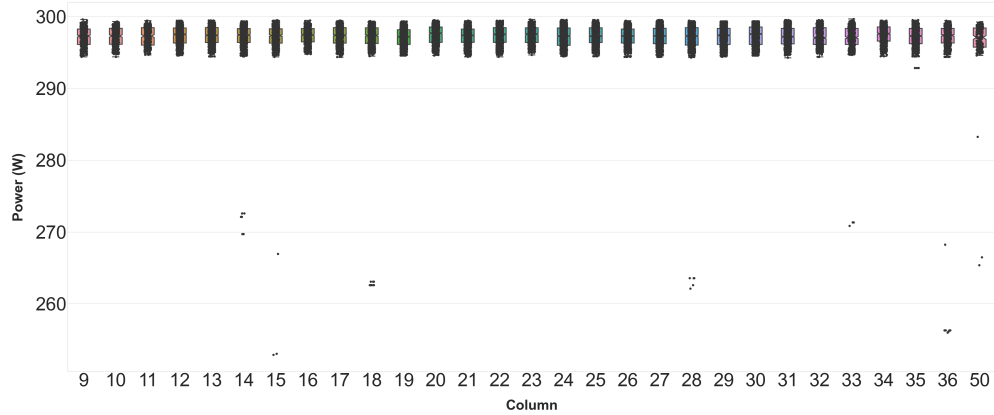
(c) Power

(d) Temperature

Fig. 23: Summary results for Summit Row H, showing variation in performance (kernel duration), frequency, power and temperature as reported by NVIDIA's profiler, when running the GPUs unthrottled at the TDP value of 300W.
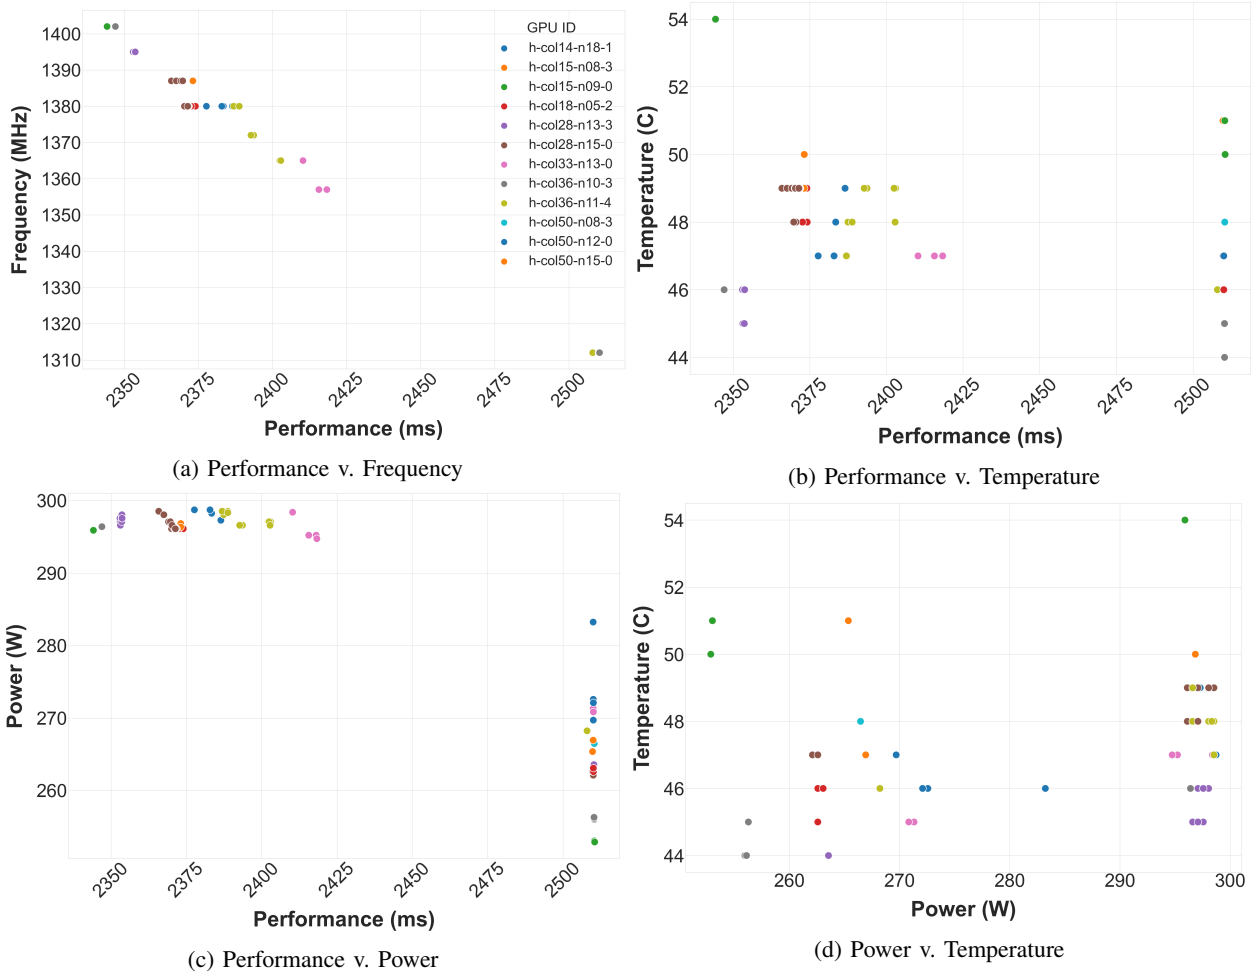
(a) Performance v. Frequency

(b) Performance v. Temperature

(c) Performance v. Power

(d) Power v. Temperature

Fig. 24: Scatter plots for GPUs in Row H of Summit showing possible correlations between (a) temperature and performance (kernel duration), (b) frequency and performance, (c) power and performance, and (d) temperature and power. The color indicates the GPU. Only GPUs with at least one reported power level < 290 W are included.

that consistently does not show outliers: it usually completes in 2353 ms while consuming 295 - 299W. However, it has an outlier where it consumes 264W and completes in 2510 ms. Rowh-col14-n18-1, shown in blue in Figure 24c, is more consistent than rowh-col28-n13-3, except that it consistently has outliers: it completes in 2510 ms while consuming 269 - 273W. Other GPUs show more divergent behavior, including runs that are outliers and runs that are not outliers. For example, rowh-col36-n11-4, shown in lime green in Figure 24c, has runs that consume up to 30W different power while taking between 2385 and 2510 ms to complete (even though the range of temperatures for this GPU is less than 3°C). This is perhaps unsurprising, because rowh-col36-n11-4 has wider variance in frequency (68 MHz, Figure 24a[2]) than other GPUs. In particular, this GPU exhibits large swings in frequency, power, and performance – the water cooling appears to be doing its job since the temperature variation is small, but this

is not preventing the GPU from having large variations in behavior.

These interesting results led us to further examine the behavior across multiple runs on one of the same GPUs. In particular, we selected rowh-col36-node10-3 (shown in gray in Figure 24), because this GPU exhibits significant variation and contains power outliers. Figure 25 also shows a time-series of continuous measurements for power and frequency from two runs on this GPU. We can see that this GPU clearly exhibits outlier behavior, since the maximum power consumption only reaches 259W. Even more interestingly, we see that this GPU is always utilizing the same frequency across both runs – 1312 MHz, despite the instantaneous power consumption rising and falling as the kernels progress. In particular, the startup power consumption is lower, showing that the GPU is able to increase its power consumption as it reaches the steady state in a given kernel invocation. However, the frequency consumption does not rise with it, unlike what we observed in TACC. Nevertheless, given the large number of outliers, across all metrics, row H exemplifies the outliers across Summit. Yet,

---

[2]In Figure 24a, we jitter the yellow point from (2510, 1312) to (2508, 1312) to make the frequency range of 68 MHz for this GPU (rowh-col36-n11-4) easily noticeable.

(a) Power

(b) Frequency

Fig. 25: Time-series plots showing continuous measurements for (a) power and (b) frequency for two distinct runs of one GPU in Row H of Summit with at least one reported power level $< 290W$. The GPU plotted is rowh-col36-n10-3.

because each node in a column of row H contains multiple GPUs, in Appendix B-B we further sub-divide, analyze, and examine the behavior of row H, column 36, since this column exhibits numerous outliers across all metrics.

B. Row H, Column 36 Results

Figure 26 shows the performance, frequency, power, and temperature breakdown across multiple runs for each node in row H, column 36. Similar to Appendix B-A, the breakdown within row H, column 36 shows that the outliers come from a specific subset of the nodes: 7 nodes (nodes 2, 6, 8, 10, 11, 13, 14, and 18) have at least one outlier, while the remaining 9 nodes do not have any. This further reinforces the difficulty in drawing conclusions from the cluster-wide summary, but also highlights the need to intelligently zoom in on specific problematic row-column pairs, given the vast amount of data across the 27,648 GPUs. In particular, nodes 10 and 11 show the most outliers across frequency, performance, and power. This is corroborated by Figure 24, which shows rowh-col36-node10-3 (gray) and rowh-col36-node11-4 (lime green). These GPUs have a wide disparity in results (including points which are outliers and not outliers), as discussed in Appendix B-A. In particular, these nodes have large outliers in performance (up to 2510 ms) and power (255W in several cases). However, neither of these nodes have any temperature outliers, unlike our prior findings on TACC. Interestingly, all temperature outliers occur on rowh-col36-node2 – which does not have any frequency, performance, or power outliers, but which runs

at up to 73°C. Thus, water cooling does not appear to be completely successful in keeping this node's GPUs within the desired range. Overall, these results highlight how several nodes repeatedly cause outlier results across several metrics, and do not always fit the established patterns we observed on other clusters. Like other clusters, frequency is usually a consistent issue for the variations. However, the results for rowh-col36-node2 show that variations can be caused (in Summit) by temperature alone. Compared to Vortex, Summit also showed considerably more variation (potentially due to gathering data across a larger amount of time and larger number of GPUs), despite Summit and Vortex both using water cooling. Thus, relying on water cooling alone to reduce variation does not appear to be sufficient.

(a) Performance

(b) Frequency

(c) Power

(d) Temperature

Fig. 26: Summary results from Summit Row H Column 36, showing variation in performance (kernel duration), measured frequency, power and temperature as reported by NVIDIA's profiler, when running the GPUs unthrottled at the TDP value of 300W.