INL/CON-24-77004-Revision-1

High Performance Computing Peak Shaving for Microreactor Operation

July 2024

Changing the World's Energy Future

Matthew William Anderson, Matthew R Sgambati

Idaho National Laboratory

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/CON-24-77004-Revision-1

High Performance Computing Peak Shaving for Microreactor Operation

Matthew William Anderson, Matthew R Sgambati

July 2024

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

High Performance Computing Peak Shaving for Microreactor Operation

Matthew Anderson
Matthew Sgambati
matthew.anderson2@inl.gov
matthew.sgambati@inl.gov
Idaho National Laboratory
Idaho Falls, Idaho, USA

ABSTRACT

There are multiple nuclear microreactors currently under development that are designed to provide autonomous power for as many as ten or more years without refueling and are designed to power high performance computing (HPC) datacenters. But the load-follow speeds for a nuclear microreactor will be much slower than grid power and slower than the power variance typical of a HPC system. HPC datacenters experience peak power load variance driven by several factors ranging from the operation of cooling systems to remove heat from the servers to supporting a wide range of user application workflows and architectures each with different power signatures. One mechanism to support the limited load-follow of a microreactor is peak shaving where an energy storage mechanism is used to shed peak load and reduce significant power variance. This work explores peak electrical load shaving using uninterruptible power supply (UPS) systems designed for HPC support in the context of peak shaving when operating using a nuclear microreactor with a load-follow limited to 10% of load per minute. Using a self contained HPC datacenter complete with stand-alone cooling system and provisioned with an x86 cluster, an ARM cluster, and a graphics processing unit (GPU) cluster, peak shaving for microreactor operation using the UPS battery backup is explored while running two classes of typical HPC user applications. HPC architecture suitability for microreactor operation under this type of peak shaving is examined.

KEYWORDS

high performance computing, peak shaving, power efficient computing, microreactor integration

1 INTRODUCTION

High performance computing (HPC) datacenters can require as much power as thousands of homes just to support the compute nodes and the top 100 supercomputers in the Top500 list from November 2023 use on average 4 MW without including the power required to cool these systems [1]. An HPC datacenter can also experience periods of significant power variance resulting from different power signatures in user application workloads such as machine learning training versus inference and even from changing weather conditions. For instance, datacenters using evaporative cooling to provide chilled water to remove heat from the datacenter will experience higher power load when there are high humidity atmospheric conditions. The different types of HPC architectures used in the datacenter also contribute to the frequency and size of the datacenter power variance. For example, Figure 1 shows the total power used as a function of time when running the HPL benchmark on a 560 core x86 cluster while also powering an ARM and GPU cluster that are idle. In this example, the power variance exceeds 30% load within a single minute.

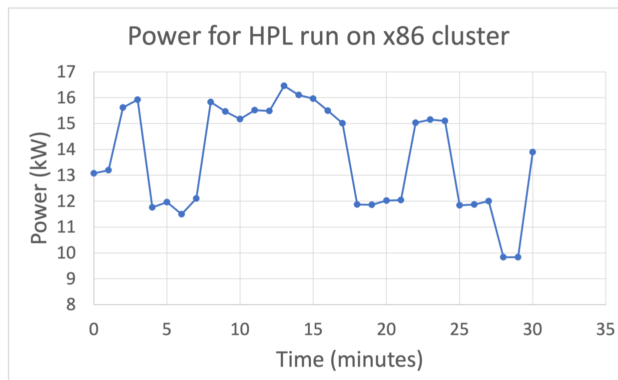

Figure 1: Total power usage including cooling costs as a function of time for running HPL on the 560 core x86 cluster detailed in Table 1. The simulation began at the 2 minute mark and finished at the 27 minute mark. Cooling accounts for most of the variation in power usage. The ARM and the GPU clusters in Table 1 were also on but idle during this run.

Peak load shaving is where an energy storage device such as a bank of batteries is used for load shedding during periods of higher demand. Peak shaving is typically considered during times in the day when electricity costs are higher and then during times of the day or night when electricity costs are lower the peak shaving energy storage device is recharged. The motivation for this type of peak shaving has been explored extensively for HPC datacenters already [2]. The arrival of nuclear microreactors has now provided a second motivation for exploring peak shaving as the load-follow speed for a microreactor will be constrained.

Nuclear microreactors are much smaller in size than a typical gigawatt-scale reactor, may be mobile, produce a relatively small amount of power, typically less than 10 MW electric, and operate autonomously and without requiring refueling for as many as 10+ years [3, 4]. They are under development by at least 12 nuclear energy manufacturers with HPC datacenters expected to be among the end users [5]. The MARVEL nuclear microreactor currently under construction at Idaho National Laboratory is one prototype of this type of reactor [6]. Unlike typical power sources for an HPC datacenter, load-follow performance for a microreactor is generally constrained. For example, load-follow experiments on the Holos-quad microreactor examined a load-follow benchmark of 10%/minute [7] which is still much slower than the power variance experienced in a datacenter, especially when factoring in cooling considerations. Without peak shaving capability, a microreactor operator will be forced to significantly overprovision power to the datacenter to handle frequent peak load variance and thereby increase the cost of the deployment. For some deployments such as space-based deployments of microreactors, mass constraints for heavy-lift rocketry will limit the maximum power yield achievable with a microreactor intended to power a space-based HPC system [8].

To explore HPC peak shaving in the context of microreactor integration, a self-contained mobile HPC datacenter with integrated cooling has been provisioned with three typical types of HPC architectures: a 560 core x86 cluster using the AMD EPYC chipset, a 1536 core ARM cluster using the Fujitsu A64FX chipset, and a cluster of 20 Nvidia A100 GPUs. A 10 kVa UPS system with eight battery units provides the energy storage mechanism for both peak shaving and UPS operations. Two types of applications are deployed: a finite volume Navier-Stokes incompressible fluid solver, and a machine learning training workflow. Each of these applications as well as the architectures on which they are deployed provide a different power and heating profile for exploring peak shaving impact. This work makes the following contributions:

- A mechanism for peak load shaving using HPC datacenter UPS systems appropriate for microreactor integration is proposed and demonstrated.
- Peak load shaving behavior on three architectures with two classes of typical HPC application workflows is quantified and presented.
- Comparative energy storage requirements are given for peak load shaving for each of the three HPC architectures explored.

2 RELATED WORK

A review of peak load shaving strategies that includes a discussion on renewable integration was conducted by Uddin et al. [9]. This extensive review paper explores energy storage sizing solutions for peak load shaving. Rana et al. [10] explore peak load shaving in a microgrid system as well as incorporate peak shaving using electric vehicles. Neither of these studies explored the specific case of HPC datacenters as end users for peak shaving. Govindan et al. [11] first suggested the use of UPS batteries for peak shaving in datacenters and explored the question if existing UPS could function in the role of peak shaving via a power cap enforcer by looking at lengths of power peaks and expected battery life. Kontorinis et al. [12] explore this same idea but using distributed per server UPS batteries. Aksalini et al. [13] also examine a distributed UPS approach for peak shaving in a datacenter but attempt to take into account physical characteristics of the batteries to improve model accuracy. Liu et al. [14] incorporate super capacitors into the UPS energy storage solution in the context of datacenters. Dabbagh et al. [2] explore a peak shaving strategy in the context of a datacenter using power traces from a Google facility and incorporate workload uncertainty and heterogeneity in their analysis. Those studies aimed to address peak shaving strategies in general without consideration of the HPC system architectures hosted in the datacenter. This study aims to incorporate the electrical peaks associated with cooling the datacenter in addition to the server power usage while exploring different HPC system architectures and typical HPC workloads. It complements those previous efforts by exploring a peak shaving strategy for nuclear microreactor integration rather than just the economic savings mechanism.

3 EXPERIMENTAL CONFIGURATION

All experiments were performed in a World Wide Technology HPC Mobile Datacenter (MDC) as shown in Figures 2–3. The MDC

Figure 2: Self-contained datacenter hosting the three classes of HPC hardware, stand-alone cooling system, and UPS battery backup.

uses 208V three-phase power and is fed from utility power through a 100A circuit breaker. This system contains a fire suppression system, UPS, HVAC, and three 42U racks for computer equipment. The HVAC system is a Bard B-410A/W36H Wall-Mount Heat Pump with a heating/cooling capacity of 3 Tons. The UPS is a Schneider Electric Symmetra 20K with a capacity of 10 kVA and contains eight out of 12 V66 batteries for storage capacity as shown in Figure 5. The three 42U compute racks have a total of 6 L630 20A power distribution units between them, which all systems used for testing are plugged into. The HVAC system has been directly wired into the UPS to allow it to run off the V66 batteries for peak shaving experiments, while everything else is connected to the utility power of the MDC, as shown in Figure 4. This strategy for UPS based peak shaving was selected since it does not impact or inhibit the server or

Figure 3: The datacenter is 208 volt three-phase and single panel fed for all HPC systems and cooling.

benchmark performance. However, it also does not eliminate power variance originating from different computational phases of the applications running on the servers themselves. As a consequence, architectures like ARM that tend to have very little power variance through different computational phases will be more suited for this type of strategy than architectures like GPUs which tend to have a large power variance for different phases.

All power measurements were captured using a Fluke power meter connected to the MDC's main breakers utility power feed. This allowed the most accurate power measurements possible since there is no loss from measuring further upstream. This approach also captured the power draw for the entire HPC datacenter, including the cooling system. All computational experiments were performed on the clusters outlined in Table 1. There is an x86 based system, an ARM based system, and a GPU based system with chipsets representative of those typically found in current HPC datacenters. The x86 base system consists of 5 x 2U SuperMicro nodes with 2 x AMD EPYC 7663 64 core CPUs, 1 TB RAM, and 100 Gb/s EDR InfiniBand running Rocky Linux 8.8. The ARM based system is a cluster of 4 x HPE Apollo 80 2U chassis with each chassis housing 4 x HPE Apollo blades, and each blade containing 2 x nodes running RHEL 8.6. The nodes have 1 x A64FX 48 core CPU with 32 GB High Bandwidth Memory (HBM) and 100 Gb/s EDR InfiniBand. The GPU based system is 3 x HPE Apollo 6500 chassis with 2 x nodes each running Rocky Linux 8.8. The nodes have a single AMD EPYC 7543P 32 core CPU, 1 TB RAM, 1 Gb/s Ethernet, and 4 x NVIDIA A100 40 GB GPUs using the SMX2 Redstone platform with NVLink. The x86 based system has a total of 560 cores and 5 TB of memory. The ARM based system has a total of 1,536 cores and 1 TB of HBM. Finally, the GPU based system has 160 total cores, 5 TB of memory, and 20 Nvidia A100 GPUs.

To test peak shaving using this configuration, jobs were started on one cluster at a time and while they were running, the main utility feed into the UPS was turned off. This forced the HVAC system to run entirely off of the UPS systems batteries, while the rest of the experiment ran off the main utility feed of the MDC. This test was then done again for each of the remaining clusters.

For both the ARM and the x86 cluster, the benchmark test was the MOOSE framework Navier-Stokes solver [15] running a low viscosity fluid through a butterfly valve. For the GPU cluster, the benchmark test used Horovod [16] to train on the MNIST dataset [17].

4 RESULTS AND ANALYSIS

Full system power usage for the x86, ARM, and GPU architectures with and without peak shaving are shown in Figures 6–8. The results for the x86 cluster running the MOOSE Navier-Stokes solver are shown in Figure 6 while for ARM they are in Figure 7; the results for the GPU cluster running the Horovod training on the MNIST dataset are shown in Figure 8. The UPS battery charge at the beginning and end of each peak shaving experiment is shown in Table 2.

In the peak shaving experiments for the each of the HPC architectures, the steady-state power variance was easily brought to within the 10%/minute limitation needed for operation using a nuclear microreactor when following the UPS peak shaving strategy from Figure 4. For the GPU and x86 architectures, the power usage did increase in the first minute of running as the benchmark application started running when compared to the idle state. In contrast, for the ARM architecture, there was no significant difference in power usage between when the system was idle and when the benchmark application was running. Without peak shaving, the power variance exceeded the 10%/minute load-follow constraint in every case.

The most pronounced peak shaving was seen in the GPU results, where the shaved peak load had a width of nearly 12 minutes. The shortest peak was seen in the x86 case with multiple peaks with widths of only two minutes. The ARM case only had one shaved peak throughout the 20 minute tests with a width of four minutes. The need for peak shaving to meet the nuclear microreactor load-follow constraint was most easily met using the ARM system and closely followed by the x86 system. The GPU system had the widest peak that required load shedding in order to meet the microreactor load-follow constraint.

Architecture	UPS charge at start	UPS charge at end
x86	84	58
ARM	100	46
GPU	84	28

Table 2: UPS Battery charge level at the start and end of each peak shaving experiment.

As seen in Table 2, the GPU cluster had the most expensive impact on the UPS energy storage charge level during the 20 minutes test while the x86 cluster had the least. Because the ARM system has almost three times as many cores as the x86 system, it is beneficial to compute the per core impact on the energy storage system's charge level over the course of the test. When scaling the battery impact to a per-core basis the ARM system showed a 30% lower impact per core on the energy storage system's charge level than the x86 system.

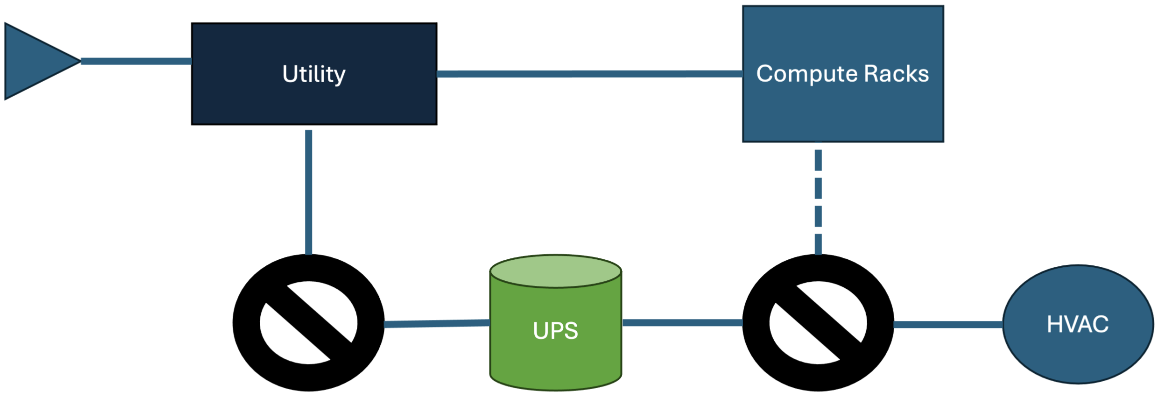

Figure 4: The UPS Configuration implemented for the microreactor peak shaving strategy. In this approach, the HVAC is wired directly into the UPS where everything else is connected to utility power. The dashed line indicates the potential to also power the compute racks from UPS in the event of utility failure, which was not tested in this study.

Architecture	Node Count	CPU	Cores	Memory	Network
x86	5	AMD EPYC 7663	112	1 TB	EDR InfiniBand
ARM	32	FUJITSU A64FX	48	32 GB HBM	EDR InfiniBand
GPU	5	AMD EPYC 7543P	32	1 TB	1 Gb Ethernet

Table 1: Details for the three clusters used in all computational experiments.

Figure 5: The eight UPS batteries used for peak shaving.

5 CONCLUSIONS

The development of nuclear microreactors for use as power sources for HPC datacenters repesents a significant departure from how these datacenters have been powered in the past. There are significant differences in load-follow speed constraints when operating with grid power versus operating solely from a microreactor where it is anticipated that load-follow will be substantially constrained. This study has used the 10%/minute load-follow constraint for the Holos-quad microreactor as a benchmark for evaluating HPC datacenters as an end user for nuclear microreactors and evaluated the power variance characteristics for three distinct HPC cluster architectures installed in a MDC with its own self-contained cooling for a complete evaluation of consumed power. Two separate benchmarks typical of HPC operations were examined in terms of power variance and load-follow: an incompressible Navier-Stokes solver and a machine learning training workflow. These two benchmarks and three HPC cluster architectures were used to illustrate a peak shaving strategy to support nuclear microreactor integration that leverages the energy storage of existing UPS batteries typically used in a datacenter in the event of grid power loss. Each cluster architecture was evaluated independently to quantify the impact on the energy storage system in the context of the peak shaving strategy introduced in this study.

Out of the three architectures explored, the ARM architecture was the best suited for nuclear microreactor integration and showed the best peak shaving characteristics and lowest impact per core on the energy storage system's charge level. The peak shaving

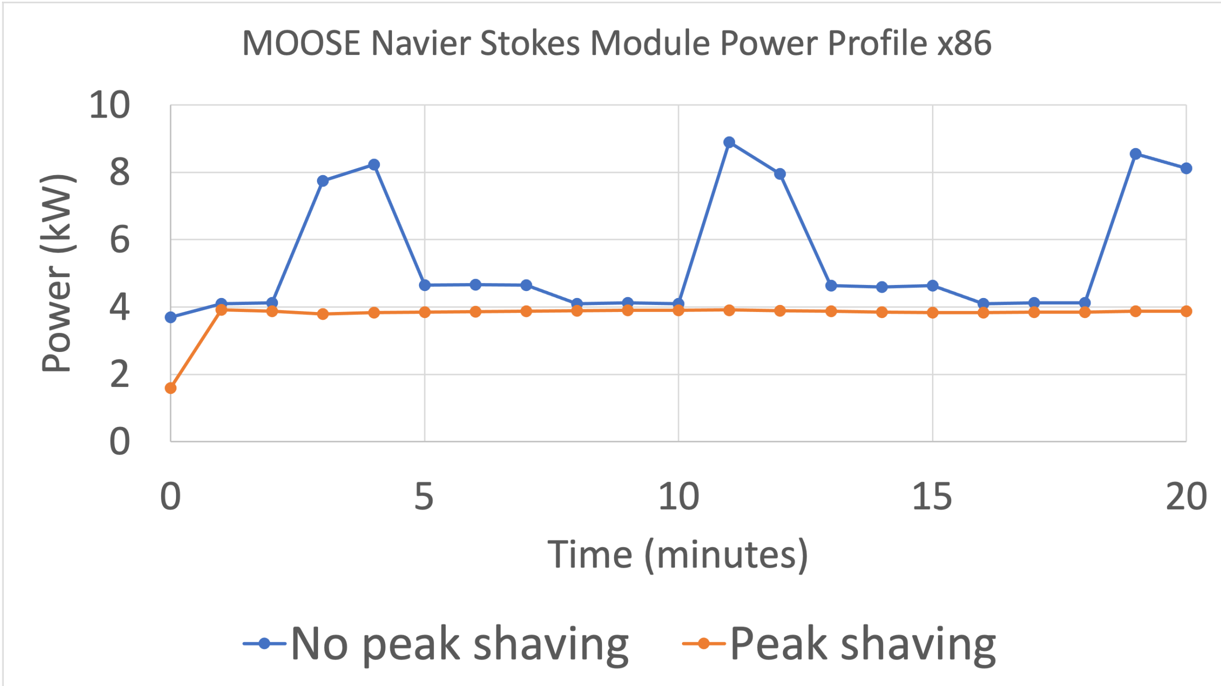

Figure 6: Total power usage including cooling costs as a function of time for running the Navier-Stokes incompressible fluid solver on the 560 core x86 AMD EPYC cluster with and without peak shaving. In the peak shaving case, the energy storage system was the system UPS. In the first minute of operation, there is a power demand change as the system goes from idle to busy.

strategy presented in this work also showed success for the x86 cluster architecture although it would require more energy storage capacity to meet the same peak shaving performance as the ARM architecture. The GPU cluster architecture was found to be poorly suited for this peak shaving strategy for nuclear microreactor integration and would likely require either significant overprovisioning of power by a nuclear microreactor due to the width of the load peaks or significantly more energy storage resources than would be necessary for either an x86 or ARM cluster architecture.

This work focused entirely on the use of UPS as the energy storage system due to their ubiquity in HPC datacenters. As part of future work, the integration of nuclear microreactors with HPC datacenters will be explored in the context of microgrids which provide greater flexibility in energy storage and input mechanisms to meet the load-follow constraints for microreactors.

6 ACKNOWLEDGMENTS

This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517. This manuscript has been authored by Battelle Energy Alliance, LLC under Contract No. DE-AC07-05ID14517 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

REFERENCES

[1] *Top500: The List.* https://top500.org. Nov. 2023.

[2] Mehiar Dabbagh et al. "Shaving Data Center Power Demand Peaks Through Energy Storage and Workload Shifting Control". In: *IEEE Transactions on Cloud Computing* 7.4 (2019), pp. 1095–1108. DOI: 10.1109/TCC.2017.2744623.

[3] K. Araújo G. Black D. Shropshire and A. van Heek. "Prospects for Nuclear Microreactors: A Review of the Technology, Economics, and Regulatory Considerations". In: *Nuclear Technology* 209.sup1 (2023), S1–S20. DOI: 10.1080/00295450.2022.2118626.

[4] Raffaella Testoni, Andrea Bersano, and Stefano Segantin. "Review of nuclear microreactors: Status, potentialities and challenges". In: *Progress in Nuclear Energy* 138 (2021), p. 103822. ISSN: 0149-1970. DOI: https://doi.org/10.1016/j.pnucene.2021.103822.

[5] Steven E. Aumeier et al. "Microreactor Applications in U.S. Markets: Evaluation of State-Level Legal, Regulatory, Economic and Technology Implications". In: (Mar. 2023). DOI: 10.2172/1964093.

[6] Jhansi Kandasamy and Elizabeth Brunner. "Idaho national laboratory to demonstrate collaboration first versus competition to accelerate achieving a secure clean energy future by 2031". In: *Nuclear Engineering and Technology* (2023). ISSN: 1738-5733. DOI: https://doi.org/10.1016/j.net.2023.11.050. URL: https://www.sciencedirect.com/science/article/pii/S1738573323005508.

[7] Anton Moisseytsev and Claudio Filippone. "Load Following Analysis of the Holos-Quad 10MWe Micro-Reactor with Plant Dynamics Code". In: (May 2022). DOI: 10.2172/1877020. URL: https://www.osti.gov/biblio/1877020.

[8] D. Watson et al. "SpaceX Falcon Heavy Mass constraints as Design Driver for Practical Heat Pipe Stirling Micro Reactors". In: *Proceedings of Nuclear and Emerging Technologies for Space.* 2023, pp. 305–311. DOI: 10.13182/NETS23-41911.

[9] Moslem Uddin et al. "A review on peak load shaving strategies". In: *Renewable and Sustainable Energy Reviews* 82 (2018), pp. 3323–3332. ISSN: 1364-0321. DOI:

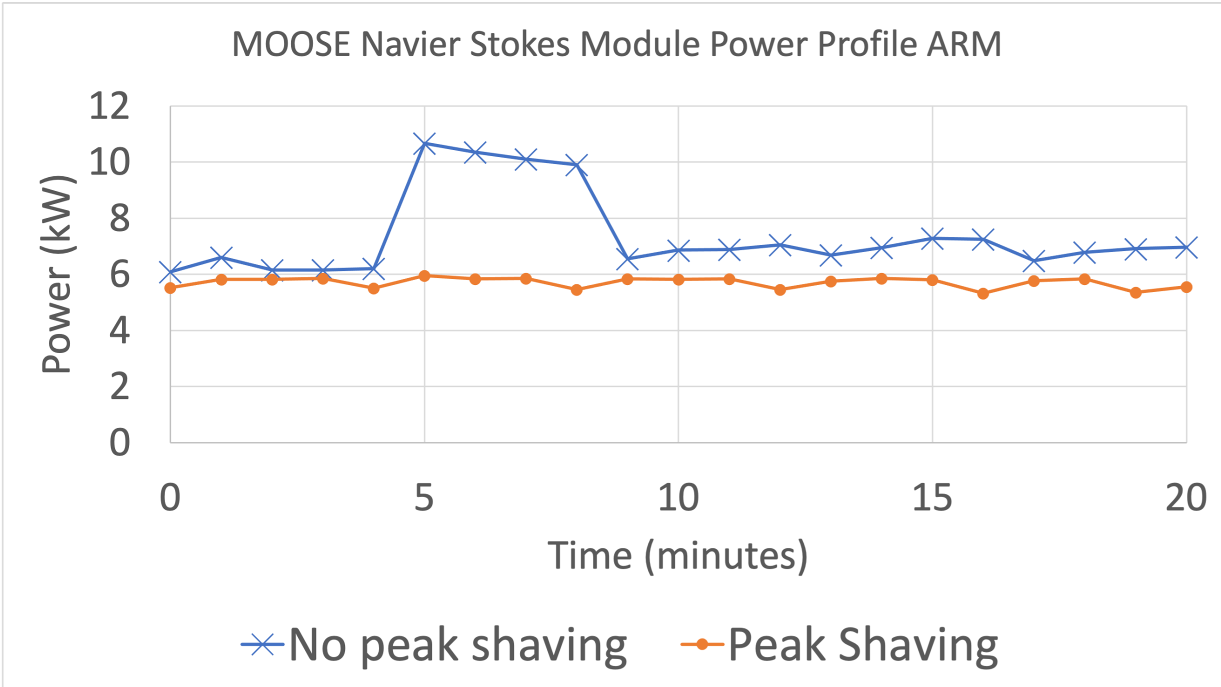

Figure 7: Total power usage including cooling costs as a function of time for running the Navier-Stokes incompressible fluid solver on the 1536 core ARM A64FX cluster with and without peak shaving. In the peak shaving case, the energy storage system was the system UPS. There is no change for this architecture in the power demand as the system goes from idle to busy.

https://doi.org/10.1016/j.rser.2017.10.056. URL: https://www.sciencedirect.com/science/article/pii/S1364032117314272.

[10] Md Masud Rana et al. "A Review on Peak Load Shaving in Microgridm-dash;Potential Benefits, Challenges, and Future Trend". In: *Energies* 15.6 (2022). ISSN: 1996-1073. DOI: 10.3390/en15062278. URL: https://www.mdpi.com/1996-1073/15/6/2278.

[11] Sriram Govindan, Anand Sivasubramaniam, and Bhuvan Urgaonkar. "Benefits and limitations of tapping into stored energy for datacenters". In: *SIGARCH Comput. Archit. News* 39.3 (June 2011), pp. 341–352. ISSN: 0163-5964. DOI: 10.1145/2024723.2000105. URL: https://doi.org/10.1145/2024723.2000105.

[12] Vasileios Kontorinis et al. "Managing distributed UPS energy for effective power capping in data centers". In: *2012 39th Annual International Symposium on Computer Architecture (ISCA)*. 2012, pp. 488–499. DOI: 10.1109/ISCA.2012.6237042.

[13] Baris Aksanli, Tajana Rosing, and Eddie Pettis. "Distributed battery control for peak power shaving in datacenters". In: *2013 International Green Computing Conference Proceedings*. 2013, pp. 1–8. DOI: 10.1109/IGCC.2013.6604477.

[14] Longjun Liu et al. "HEB: Deploying and managing hybrid energy buffers for improving datacenter efficiency and economy". In: *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*. 2015, pp. 463–475. DOI: 10.1145/2749469.2750384.

[15] Alexander Lindsay et al. "MOOSE Navier–Stokes module". In: *SoftwareX* 23 (2023), p. 101503.

[16] Alexander Sergeev and Mike Del Balso. "Horovod: fast and easy distributed deep learning in TensorFlow". In: *arXiv preprint arXiv:1802.05799* (2018).

[17] Li Deng. "The mnist database of handwritten digit images for machine learning research". In: *IEEE Signal Processing Magazine* 29.6 (2012), pp. 141–142.

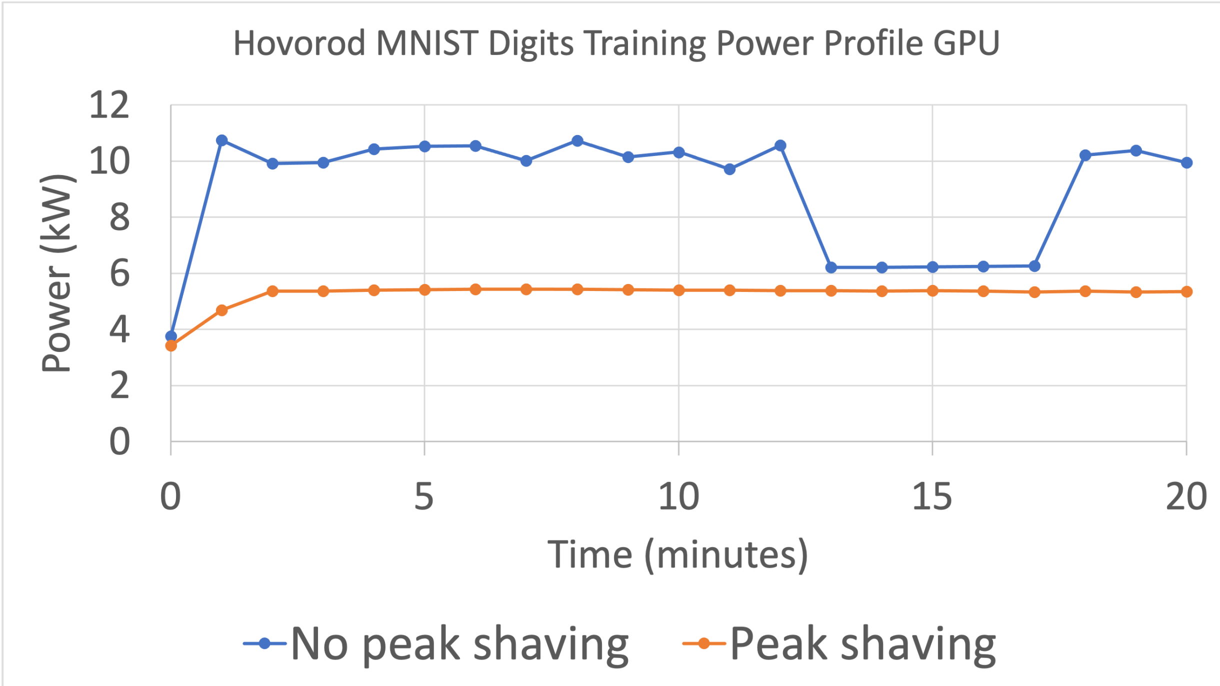

Figure 8: Total power usage including cooling costs as a function of time for using Horovod to train the MNIST dataset on the 20 NVIDIA A100 GPUs with and without peak shaving. In the peak shaving case, the energy storage system was the system UPS. In the first minute of operation, there is a power demand change as the system goes from idle to busy.