Energy and Power Aware Job Scheduling and Resource Management: Global Survey — An In-Depth Analysis

Gregory A. Koenig*, Matthias Maiterth†‡, Siddhartha Jana†, Natalie Bates*,
Kevin Pedretti§, Milos Puzovic¶, Andrea Borghesi‖, Andrea Bartolini‖, David Montoya**

*EE HPC WG, †Intel Corporation, ‡LMU Munich, §Sandia National Laboratories,
¶The Hartree Centre, ‖University of Bologna, **Los Alamos National Laboratory

***Abstract*—This paper presents a detailed analysis of a first-of-kind global survey of high-performance computing centers that actively employ techniques within job scheduling and resource management middleware layers for managing energy and power on production supercomputers. Our group conducted a series of comprehensive interviews of leading-edge supercomputing centers during 2016 and 2017. The group presented the motivation of the survey center selection, questionnaire details, and a preliminary analysis of the survey results in a previous publication. This paper presents a more detailed analysis of the survey results. The goal is to find commonalities, approaches that are to be developed and hints and guidelines for other centers to move towards a more energy efficient and power aware management of their compute resources.**

***Index Terms*—power, energy, performance, power-aware, computing, scheduling**

I. INTRODUCTION

The Energy Efficient High-Performance Computing Working Group (EE HPC WG) team for Energy and Power Aware Job Scheduling and Resource Management (EPA JSRM) conducted a survey and subsequent series of comprehensive interviews of leading-edge computing centers during 2016 and 2017. The purpose of the survey and interviews was to assess the state of the art in using job scheduling and resource management techniques to manage the energy and power consumption of large-scale supercomputers. An initial paper that describes the survey modalities, the motivation behind the survey, the center selection process, and an overview of the questionnaire and site responses was published earlier this year [28].

The current paper presents a more detailed analysis of the survey results, intended as a follow-up to the initial paper. Section II briefly reviews the overview of the survey, and readers are referred to [28] for a more in-depth overview. The in-depth analysis of the survey is presented in Section III. Based on this analysis, Section V presents insights and forward-looking recommendations. A brief review of Related Work appears in Section IV.

II. SURVEY SUMMARY

The EE HPC WG EPA JSRM team identified all Top500 supercomputing centers that we were aware of as using job scheduling and resource management techniques to manage energy and power of large-scale supercomputers. Of the eleven centers that we initially identified, nine centers agreed to participate in the survey described in this paper while two centers declined due to security-related concerns.

The team was motivated by the desire to identify characteristics of leading-edge centers that are using EPA JSRM techniques, details af the techniques, how the techniques are applied, and the impacts to center operations. To address these motivations, the team developed a questionnaire consisting of eight questions focused on general site description, the site's approach to energy and power aware job scheduling and resource management, and a qualitative self-assessment of the efficacy of the approach. The responding sites provided varying level of detail and follow-up interviews were conducted with each site to supplement the responses for better comparability. Although other papers have discussed these types of techniques previously, we believe our work is valuable because it is, to the best of our knowledge, the only effort to approach these centers directly and in a cohesive manner.

The centers that participated in the survey were:

1) RIKEN Center for Computational Science (RIKEN)
2) Tokyo Institute of Technology (Tokyo Tech)
3) Alternative Energies and Atomic Energy Commission (CEA)
4) King Abdullah University of Science and Technology (Kaust)
5) Leibniz Supercomputing Centre (LRZ)
6) Science and Technology Facilities Council (STFC)
7) Los Alamos and Sandia National Laboratories (Trinity)
8) Consorzio Interuniversitario del Nord est Italiano Per il Calcolo Automatico (CINECA)
9) Joint Center for Advanced High Performance Computing (JCAHPC)

Broadly, the survey results reflect a variety of solutions that cater to center-specific characteristics. For example, one site was motivated to develop an EPA JSRM solution that incorporates a gas turbine co-generation capability that can be invoked at times when the cost of running the turbine is favorable in comparison to the cost of purchasing power from the center's electricity service provider. Another site was motivated by the desire to integrate fine-grained power management runtime systems into scheduling decisions. Overall, a common thread among the approaches reported in the survey was that centers are typically developing their solutions in collaboration with system integrators or HPC-oriented software vendors. In several cases, these efforts are augmented by including various open source software packages in order to leverage work taking place within the HPC community.

III. SURVEY ANALYSIS

This section presents an in-depth analysis as well as insights about the survey responses. These contents extend the comprehensive overview given in Table 1 of our previous work [28]. Our goal of this analysis is to identify commonalities among approaches taken by the surveyed computing centers as well as to identify particularly noteworthy approaches.

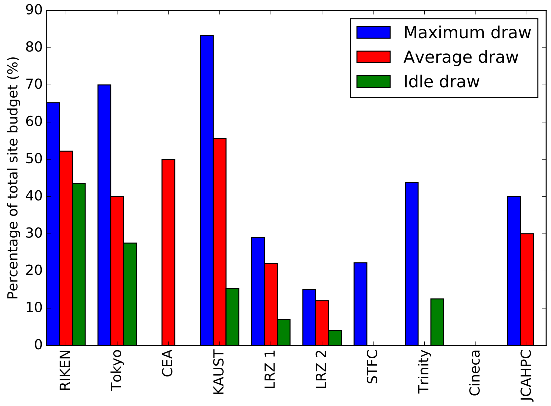

Fig. 1. Maximum, average and idle power consumption as percentage of total power budget

Table I in the current paper shows a high-level summary of responses from the sites regarding power and cooling capacity as well as maximum, average, and idle power draws from each site. Following from the data presented in the Table, Figure 1 presents a graph of the percentage of total power capacity that is reached for each of maximum, average, and idle power draw at each site.

In analyzing the responses to the survey, one way of dividing the responses is into two groups: (1) techniques that primarily focus on managing power, such as by limiting the impact of idle resources on a center's current power draw, and (2) techniques that primarily focus on managing energy, such as by evaluating the performance-per-Watt profile of each job and attempting to determine a way of scheduling jobs to optimize this metric. The following subsections explore these delineations in greater detail.

A. Power-Oriented Approaches

The first broad category of EPA JSRM techniques focuses on managing power consumption. In many ways, power-oriented techniques are more straightforward because power consumption tends to be a more short-term goal and, accordingly, the telemetry necessary to achieve power-oriented goals can be taken from more ephemeral sources. That said, the power draw of a supercomputing center is in many cases the biggest limiting factor to the maximimum size of computer that a center can run, so managing power is a critical success factor for many centers.

In the simplest approach, the batch job scheduler combines information about idle resources and upcoming jobs and uses the resource management software to *shut down the idle resources*. This simple approach has two advantages. First, powered off nodes consume no (or essentially no) power. Second, most job schedulers and resource managers (e.g., Moab [1], [2] and SLURM [31], [32]) support this functionality. However, despite the advantages related to simplicity and effectiveness, shutting down idle resources is often not possible for the simple fact that supercomputers typically have consistently high utilization except during scheduled maintenance windows. High purchase costs and limited lifespans, typically around 60 months, mean that supercomputing centers strive to keep all computing resources highly utilized with few periods of idleness.

One survey site that uses the technique of shutting down idle resources to significant effect is the Tokyo Institute of Technology. The benefits of the approach can be observed in Figure 1 which shows the percentage of maximum, average and, idle power draw. Tokyo Institute of Technology sees high idle power draws, thus this technique outweighs the disadvantages the approach can have and the center achieves the highest average power draw compared to its maximum power draw.

Another power-oriented approach is *dynamic job termination*. Dynamic job termination is used to directly respond to power constraint situations by terminating one or more running jobs in order to keep the total system power draw under a defined critical value. The power limit may be either a hard limit (e.g., an actual hardware limits) or a soft limit (e.g., due to exceptional situation-based power costs or due to a buffered hardware safeguard). Dynamic job termination approaches may vary in complexity from fully manual to fully autonomous selection and cancellation of jobs. Candidate jobs for termination are usually selected based on each job's contribution to overall power conumption. Additional factors that may be considered include the progress that each job has made toward its overall requested wallclock time, priority or some time-critical designation of each job, and whether each job is a piece of a longer-term multi-job workflow. Generally, terminated jobs are re-queued for execution at a future time when power demands are lower.

RIKEN is an example of a site that employs dynamic job termination. When RIKEN's resource management system detects an excessive power usage in relation to power available to the center, the resource management system begins stopping jobs by canceling (not suspending) the job using the largest number of nodes, repeating the process until power consumption falls below the critical threshold. Research at the center

TABLE I

TOTAL POWER AND COOLING BUDGET WITH MAXIMUM, AVERAGE AND IDLE POWER DRAWS (UNITS ARE IN MW)

Center	Power Budget	Cooling Budget	Maximum draw	Average draw	Idle draw
RIKEN	23	36	15	12	10
TokyoTech	2	2	1.4	0.8	0.55
CEA	10	7.5		5	
KAUST	3.6	2.9	3	2	0.55
LRZ 1	10	10	2.9	2.2	0.7
LRZ 2	10	10	1.5	1.2	0.4
STFC	4.5	2	1		
Trinity	19.2	27	8.4		2.4
Cineca	6				
JCAHPC	8	4.2	3.2	2.4	

is currently underway to develop a more sophisticated policy based around some of the factors described in the previous paragraph.

B. Energy-Oriented Approaches

The second broad category of EPA JSRM techniques focuses on managing energy. Because supercomputing centers typically try to keep their computational resources engaged at all times, energy-oriented approaches are in some ways more aligned with overall center goals because energy-oriented approaches tend to involve managing active resources. To this end, energy-oriented approaches generally involve optimizing around some objective involving factors such as utilization, performance, fairness, time to completion, and energy consumption. Making these decisions often involves maintaining historic information about job execution and combining this information with information taken from the currently queued upcoming jobs as well as with information about the center's longer term goals.

It is interesting to notice that optimizing toward energy use is often the same as optimizing toward performance, and that these goals are in many cases *in opposition* to the goal of optimizing toward power consumption. Consider the well-understood case of optimizing a parallel application's performance by load balancing processing elements across the CPUs allocated to the job. The typical outcome of load balancing is that application performance is improved leading to less execution time to complete the job. This objective likely coincides with the objective of improving the energy efficiency of the job due to the job's shorter wallclock time to complete. However, because load balancing typically results in much higher CPU utilization on each CPU used in the job, the peak power of the job may be much higher.

As described above, EPA JSRM techniques that are focused on energy typically need to have some understanding of how a given job configuration (executable, input data, number of nodes, node layout, etc.) is expected to perform. Accordingly, one approach for energy-oriented job scheduling reported by several centers is *application tagging*. With this approach, users provide information to the job scheduler about characteristics of submitted jobs by manually applying one or more "tag values" to each job. The job scheduler and resource management system use this information to configure the allocated nodes accordingly. Example tags can be characteristics, such as CPU, memory, or I/O affinity of a job. In other similar approaches, the tags that a user applies to a job relate the job to historical records of past runs of the job.

LRZ is perhaps the site that makes the most extensive use of application tagging. LRZ's approach was developed in collaboration with IBM and leverages capabilities of the IBM LoadLeveler job scheduler [23] used on LRZ's SuperMUC system. The approach encourages users to tag their jobs at submit time with metadata that allows the scheduler to choose the best DVFS settings for the nodes allocated to a job. The work involved development of models for energy consumption for various applications running on SuperMUC under various DVFS settings. The concepts are explained in the work by Auweter et al. [3] and involve an initial run of each application with a low default frequency and then setting the frequency for subsequent submission of the same application at higher frequencies according to the model if the energy/performance trade-off is favorable. The effectiveness of the technique is shown in Figure 1 where LRZ's maximum and average power draw, taken as a percentage of total power capacity, are very similar. The long-term effect of using application tagging at LRZ is an overall reduction in energy consumption. Follow-on work at LRZ is planned to implement similar functionality in the open source SLURM scheduler for upcoming machines deployed at LRZ.

Approaches involving *predictive models* extend application tagging approaches by attempting to automatically tag jobs according to specific characteristics of jobs that are typically stored in a historical database of job runs. These approaches are challenging due to difficulties related to extrapolating historical experience to new job configurations (e.g., different input datasets to the same executable) and to unseen numbers or configurations of nodes.

IV. RELATED WORK

Energy and power aware job scheduling and resource management has received increasing attention in both academic research and vendor efforts in several ways related to the experiences reported by survey sites.

Research in EPA JSRM involves several topic areas. Multiple surveys regarding power management techniques have been published [10], [21], [27], [29]. These efforts have primarily focused on the research aspects of EPA JSRM rather than on developing capabilities suitable for production deployment. Several efforts focus on optimizing power consumption by applying specific techniques, such as moldable jobs [4], [33], [34], or general autotuning approaches used to improve scheduling [24], [35]. Other work looks at specific changes to CPU operation, e.g., Dynamic Voltage and Frequency Scaling (DVFS) in [3], [15]–[17], [17], [18], [22], and Intel's Running Average Power Limit (RAPL) [12] as well as works using it [6], [14], [14]. Several efforts focus more on the scheduling aspect specific parts thereof such as job co-scheduling [7]–[9], alteration of execution order [3], [5], [25], [26] and usage of energy-tags [3], [7], [36], [37].

Vendor support for EPA JSRM includes both software and hardware solutions. Vendor solutions provide capabilities for system wide power capping [31], [32], and administrative tools for node control [30] such as Cray's System Management Workstation [11]. ACPI [20]) is provided by SLURM, PBSPro, Moab, and CAPMC. Tools for ramp-rate control also exist for example Cray's power staging tools. Job Level Power management is provided by [13]. These efforts are enabled by the respective hardware level monitors and controls.

V. RECOMMENDATIONS AND CONCLUSIONS

Our objective in conducting the survey and subsequent interviews described in this paper was to understand EPA JSRM techniques that are being developed within the high-performance computing community was a way of helping to foster the growth of these efforts. Based on our analysis of the survey and interview comments, we are prepared to make a few recommendations for ways the high-performance computing community might further invest into EPA JSRM.

- *User Awareness* — In order to incentivize users toward taking steps such as optimizing their applications for energy efficiency, helping HPC system software understand how to best deploy their applications (e.g., by using tag values when submitting jobs, as described in Section III-B), or making the highest and best use of hardware that provides favorable energy efficient application execution such as GPU accelerators, we recommend investing heavily into software tools that provide users with accurate ways of measuring and reporting application performance, power draw, and energy efficiency. This point was emphasized during our interview with STFC and their description of using tools such as LSF and Allinea Forge (Map) in their datacenter environment.
- *Monitoring Applications and Resources* — Related to the previous item, we recommend that centers identify monitoring and reporting capabilities such as those described in the Green500 *Energy Efficient High Performance Computing Power Measurement Methodologies* report [19]. While the previous item focuses more on users, this item is related to measuring power and energy in the center as a whole. Most or all sites that participated in the survey indicated that measurements such as these are a critical success factor in achieving their power and energy goals.
- *Funding Agency Emphasis* — When energy efficiency is incentivized by a center's funding agency, the center becomes focused on this objective and typically begins exploring ways of driving advances in this area. As noted above, because energy-oriented optimizations often overlap somewhat with performance-oriented optimizations, this focus on energy efficiency may improve operations within a center in many ways. Further, when centers are focused on a particular topic, vendors who sell hardware and software to the centers also become incentivized to explore these topics. In our survey, STFC expressed that their funding agency placed high incentive on energy efficiency as a way of spending more monetary resources on computing rather than on electricity, and LRZ expressed that their funding agency placed emphasis on a drive toward green computing.
- *System Software* — Several centers expressed recognition of the importance of investing in energy and power aware system software stacks such as runtime systems and resource managers. We recommend that centers identify promosing research that is aligned with center goals in this space and engaging the research groups conducting this research. Examples of such efforts include the PowerStack Group [38].
- *HPC Community Outreach* — During the course of interviewing the survey sites, the authors realized that a wide variety of techniques are being developed that may be obvious to the broad high-performance computing community. We recommend that centers engaged in EPA JSRM development efforts to publish their experience in formal peer reviewed papers or in less formal technical reports so the community can benefit from this experience.

Our team within the Energy Efficient High-Performance Computing Working Group is encouraged by the EPA JSRM efforts currently underway within the high-performance computing community. Although many of these efforts are based on an evolutionary approach to address local needs, we believe that the techniques developed in this space will be of increasing interest to more Top500 sites in coming years.

ACKNOWLEDGMENT

We thank the participating HPC Centers for providing us with their time and knowledge for the interviews, and reviews thereof. We would especially like to thank Keiji Yamamoto, Fumiyoshi Shoji, Kazuo Minami, Toshio Endo, Francis Belot, Gilles Wiber, Matthieu Hautreux, Jean-Marc Ducos, Bilel Hadri, Torsten Wilde, Michael Rudgyard, Neil Morgan, Stephen Hill, James Laros, Josip Loncaric, Carlos Cavazzoni, Toshihiro Hanawa, Masaaki Kondo and all others participating in the interviews.

REFERENCES

[1] Adaptive Computing Enterprises Inc. Green Computing. http://docs.adaptivecomputing.com/suite/8-0/basic/help.htm#topics/moabWorkloadManager/topics/ powerManagement/greenTasks/nodeIdlePowerAction.htm, online.

[2] Adaptive Computing Enterprises Inc. Moab Workload Manager Administrator Guide 9.0.3. http://docs.adaptivecomputing.com/9-0-3/MWM/Moab-9.0.3.pdf.

[3] A. Auweter, A. Bode, M. Brehm, L. Brochard, N. Hammer, H. Huber, R. Panda, F. Thomas, and T. Wilde. A case study of energy aware scheduling on supermuc. In J. M. Kunkel, T. Ludwig, and H. W. Meuer, editors, *Supercomputing*, pages 394–409, Cham, 2014. Springer International Publishing.

[4] P. E. Bailey, A. Marathe, D. K. Lowenthal, B. Rountree, and M. Schulz. Finding the limits of power-constrained application performance. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '15, pages 79:1–79:12, New York, NY, USA, 2015. ACM.

[5] A. A. Bhattacharya, D. Culler, A. Kansal, S. Govindan, and S. Sankar. The need for speed and stability in data center power capping. *Sustainable Computing: Informatics and Systems*, 3(3):183 – 193, 2013. Selected papers from the 2012 IEEE International Green Computing Conference (IGCC 2012).

[6] D. Bodas, J. Song, M. Rajappa, and A. Hoffman. Simple power-aware scheduler to limit power consumption by hpc system within a budget. In *Proceedings of the 2Nd International Workshop on Energy Efficient Supercomputing*, E2SC '14, pages 21–30, Piscataway, NJ, USA, 2014. IEEE Press.

[7] A. Borghesi, A. Bartolini, M. Lombardi, M. Milano, and L. Benini. Predictive modeling for job power consumption in hpc systems. In J. M. Kunkel, P. Balaji, and J. Dongarra, editors, *High Performance Computing*, pages 181–199, Cham, 2016. Springer International Publishing.

[8] A. Borghesi, F. Collina, M. Lombardi, M. Milano, and L. Benini. Power capping in high performance computing systems. In G. Pesant, editor, *Principles and Practice of Constraint Programming*, pages 524–540, Cham, 2015. Springer International Publishing.

[9] A. Borghesi, C. Conficoni, M. Lombardi, and A. Bartolini. Ms3: A mediterranean-stile job scheduler for supercomputers - do less when it's too hot! In *2015 International Conference on High Performance Computing Simulation (HPCS)*, pages 88–95, July 2015.

[10] K. W. Cameron, R. Ge, and X. Feng. High-performance, power-aware distributed computing for scientific applications. *IEEE Computer*, 38(11):40–47, 2005.

[11] CRAY. System Management Workstation (SMW). http://docs.cray.com/books/S-2480-7204a/S-2480-7204a.pdf, online.

[12] H. David, E. Gorbatov, U. R. Hanebutte, R. Khanna, and C. Le. Rapl: Memory power estimation and capping. In *Proceedings of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design*, ISLPED '10, pages 189–194, New York, NY, USA, 2010. ACM.

[13] J. Eastep, S. Sylvester, C. Cantalupo, B. Geltz, F. Ardanaz, A. Al-Rawi, K. Livingston, F. Keceli, M. Maiterth, and S. Jana. Global Extensible Open Power Manager: A Vehicle for HPC Community Collaboration on Co-Designed Energy Management Solutions. In J. M. Kunkel, R. Yokota, P. Balaji, and D. Keyes, editors, *High Performance Computing*, pages 394–412, Cham, 2017. Springer International Publishing.

[14] D. A. Ellsworth, A. D. Malony, B. Rountree, and M. Schulz. Dynamic power sharing for higher job throughput. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '15, pages 80:1–80:11, New York, NY, USA, 2015. ACM.

[15] M. Etinski, J. Corbalan, J. Labarta, and M. Valero. Optimizing job performance under a given power constraint in hpc centers. In *International Conference on Green Computing*, pages 257–267, Aug 2010.

[16] M. Etinski, J. Corbalan, J. Labarta, and M. Valero. Parallel job scheduling for power constrained hpc systems. *Parallel Comput.*, 38(12):615–630, Dec. 2012.

[17] F. Fraternali, A. Bartolini, C. Cavazzoni, and L. Benini. Quantifying the impact of variability and heterogeneity on the energy efficiency for a next-generation ultra-green supercomputer. *IEEE Transactions on Parallel and Distributed Systems*, PP(99):1–1, 2017.

[18] V. W. Freeh, D. K. Lowenthal, F. Pan, N. Kappiah, R. Springer, B. L. Rountree, and M. E. Femal. Analyzing the energy-time trade-off in high-performance computing applications. *IEEE Transactions on Parallel and Distributed Systems*, 18(6):835–848, June 2007.

[19] Green500, Top500, Green Grid, EE HPC WG. Energy Efficient High Performance Computing Power Measurement Methodology. https://www.top500.org/green500/resources/eehpc-wg-power-measurement-methology/, online.

[20] Hewlett-Packard Corporation, Intel Corporation, Microsoft Corporation, Phoenix Technologies Ltd., and Toshiba Corporation. Advanced configuration and power interface specification. June 2009. Revision 4.0.

[21] C. Hsu, W. Feng, and J. S. Archuleta. Towards efficient supercomputing: A quest for the right metric. In *19th International Parallel and Distributed Processing Symposium (IPDPS 2005), CD-ROM / Abstracts Proceedings, 4-8 April 2005, Denver, CO, USA*. IEEE Computer Society, 2005.

[22] C.-h. Hsu and W.-c. Feng. A power-aware run-time system for high-performance computing. In *Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*, SC '05, pages 1–, Washington, DC, USA, 2005. IEEE Computer Society.

[23] IBM. IBM LoadLeveler User Manual. https://www.ibm.com/support/knowledgecenter/SSFJTW/loadl_welcome.html, online.

[24] 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*, SC '15, pages 78:1–78:12, New York, NY, USA, 2015. ACM.

[25] B. Khemka, R. Friese, S. Pasricha, A. A. Maciejewski, H. J. Siegel, G. A. Koenig, S. Powers, M. Hilton, R. Rambharos, and S. Poole. Utility maximizing dynamic resource management in an oversubscribed energy-constrained heterogeneous computing system. *Sustainable Computing: Informatics and Systems*, 5:14 – 30, 2015.

[26] K. Leal. Energy efficient scheduling strategies in federated grids. *Sustainable Computing: Informatics and Systems*, 9(Complete):33–41, 2016.

[27] Y. Liu and H. Zhu. A survey of the research on power management techniques for high-performance systems. *Softw. Pract. Exper.*, 40(11):943–964, Oct. 2010.

[28] M. Maiterth, G. Koenig, K. Pedretti, S. Jana, N. Bates, A. Borghesi, D. R. Montoya, A. Bartollini, and M. Puzovic. *Energy and Power Aware Job Scheduling and Resource Management: Global Survey Initial Analysis*. Institute of Electrical and Electronics Engineers Inc., Vancouver, Canada, 5 2018.

[29] M. Maiterth, T. Wilde, D. Lowenthal, B. Rountree, M. Schulz, J. Eastep, and D. Kranzlmueller. *Power aware high performance computing: Challenges and opportunities for application and system developers - Survey & tutorial*, pages 3–10. Institute of Electrical and Electronics Engineers Inc., United States, 9 2017.

[30] S. Martin, D. Rush, and M. Kappel. Cray advanced platform monitoring and control (capmc). In *Proc. Cray Users Group Technical Conference (CUG)*, 2015.

[31] S. Morris Jette. Slurm power saving guide. https://slurm.schedmd.com/power_save.html, online.

[32] S. Morris Jette. Technical: Slurm power management support. https://slurm.schedmd.com/SLUG15/Power_mgmt.pdf, technical presentation.

[33] A. W. Mu'alem and D. G. Feitelson. Utilization, predictability, workloads, and user runtime estimates in scheduling the ibm sp2 with backfilling. *IEEE Trans. Parallel Distrib. Syst.*, 12(6):529–543, June 2001.

[34] T. Patki, D. K. Lowenthal, A. Sasidharan, M. Maiterth, B. L. Rountree, M. Schulz, and B. R. de Supinski. Practical resource management in power-constrained, high performance computing. In *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '15, pages 121–132, New York, NY, USA, 2015. ACM.

[35] H. Shoukourian. Adviser for energy consumption management: green energy conservation.

[36] H. Shoukourian, T. Wilde, A. Auweter, and A. Bode. Predicting the energy and power consumption of strong and weak scaling hpc applications. *Supercomput. Front. Innov.: Int. J.*, 1(2):20–41, July 2014.

[37] A. Sîrbu and O. Babaoglu. Power consumption modeling and prediction in a hybrid cpu-gpu-mic supercomputer. In *Proceedings of the 22Nd International Conference on Euro-Par 2016: Parallel Processing - Volume 9833*, pages 117–130, New York, NY, USA, 2016. Springer-Verlag New York, Inc.

[38] The HPC PowerStack Group. The HPC PowerStack. https://powerstack.lrr.in.tum.de/, online.