Energy-aware High Performance Computing – A Taxonomy Study

Chang Cai[2], Lizhe Wang[1], Samee U. Khan[3], Jie Tao[4]
1 Pervasive Technology Institute, Indiana University, USA
2 School of Computer Science, China University of Geosciences, P. R. China
3 Department of Electrical and Computer Engineering, North Dakota State University, USA
4 Steinbuch Centre for Computing, Karlsruhe Institute of Technology, Germany
† Corresponding author: Lizhe Wang, Lizhe.Wang@gmail.com

***Abstract*—To reduce the energy consumption and build a sustainable computer infrastructure now becomes a major goal of the high performance community. A number of research projects have been carried out in the field of energy-aware high performance computing. This paper is devoted to categorize energy-aware computing methods for the high-end computing infrastructures, such as servers, clusters, data centers, and Grids/Clouds. Based on a taxonomy of methods and system scales, this paper reviews the current status of energy-aware HPC research and summarizes open questions and research directions of software architecture for future energy-aware HPC studies.**

Keywords: High Performance Computing, Energy-aware Computing, Green Computing.

I. Introduction

The research topic of "energy-aware high performance computing" exhibits multiple faces, for example, research methodologies, system scales, and techniques. To understand the context of energy-aware high performance computing and identify the roadmap of future research, a taxonomy of existing research projects and research activities is desired and helpful. In this paper, we develop a taxonomy study of energy-aware high performance computing, where related research effort is categorized in several dimensions (see also Figure 1):

- Compute system scales: based on the system scales, the compute systems under inspection can be classified as: compute servers, compute clusters, distributed virtualized infrastructures, data centers, computational Grids & Clouds.
- System goals: to reduce energy consumption of investigated systems, we try to can reduce its power consumption directly. To decrease the heat emission of compute systems, thus reduce the energy consumption of air-condition components, can also serve as a good solution to the energy-aware computing.
- Research methodologies: a number of approaches and methodologies are adopted in the research of energy-aware parallel computing, for example, hardware re-configuration and operation, virtual machine migration and resource consolidation, programming language and runtime support, and adaptive workload distribution and system management.
- Viewpoints and objectives The research of energy-aware parallel computing can be taken in various contexts and has different objectives.
  - For an application user, typically there is limitation of power consumption for their applications. In general application users are not mandatory to go to "green". Therefore users choose energy-aware solutions voluntarily or with some economical benefit. An application user prefer to save energy consumption with tolerant performance loss or QoS degrade. There is always a tradeoff between application performance & QoS and reduced power consumption.
  - For compute resource owners, power consumption sometime is a economical consideration as the power consumption is a significant portion of the total cost ownership of resources. However, under certain circumstance, compute resources are forced to go "energy-aware" as they reach some limitations for power usage, resource temperature and spaces. Additional work to go energy-aware may affect compute resource performance, like throughput, response time, serviceability, maintainability, availability and reliability.
  - To develop energy aware applications brings further burdens for application developers, which might bring some degrade developing related qualities of applications, for example, buildability, testability, scalability, reusability, portability, security, extensibility and availability.

II. Compute servers

Significant research efforts have been done to reduce power consumption by improving the energy-efficiency for computer servers. This section summarizes energy-aware research in a computer server scale.

A. System level software management

Software control on the operating system level plays a key role in the high performance computing of computer servers. For example, several aspects in the OS are suggested to take

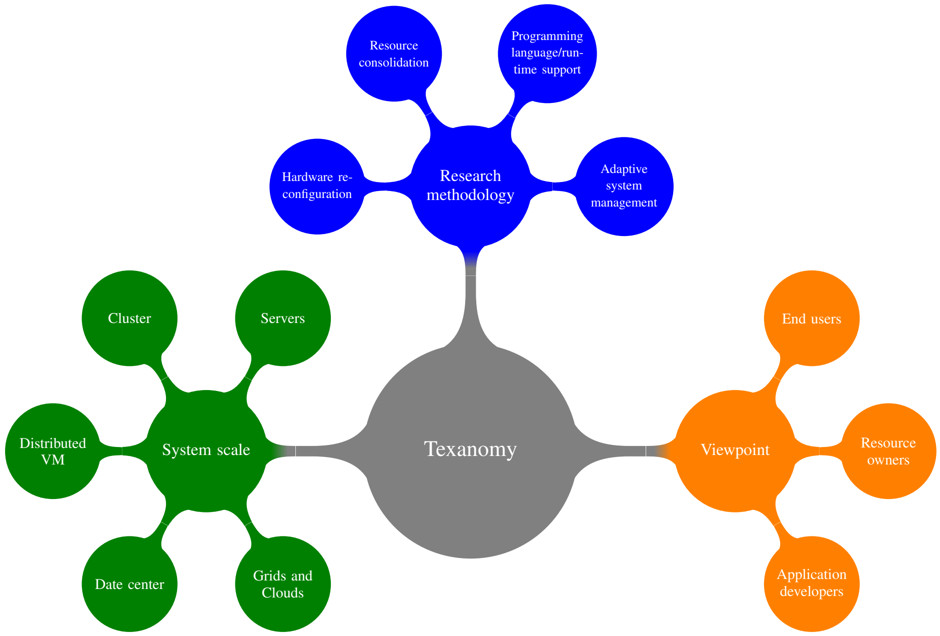

Fig. 1. Taxonomy of energy-aware high performance computing

into account [48] when running on batteries to avoid using energy unnecessarily, such as performing computation remotely and turning off unnecessary devices (including portions of the memory subsystem).

In a later study, Lebeck et al. [27] develop a study on the memory subsystem in OS. This study directs memory accesses to certain memory banks and proposes to turn off unused memory banks. In [3], Benini et al. propose that the operating system is required to monitor the resource usage and locate possibilities to turn off unused components more accurately. Flinn et al. [18] develop a user-level middleware to filter and transcode data that applications fetch. Transcoding changes data quality in order for applications to use the minimum amount of energy when processing it.

The Power and Thermal Management group of the GreenLight [7] aims to develop a system level research of energy-aware computing. It includes 1) technologies to effectively characterize power and thermal related features of workloads and applications, 2) event-driven methods that optimize energy and thermal management, and 3) feedback loops to evaluate the management policies via using different types of sensors to monitor power consumption and temperature.

B. Power management via feedback control

Some work uses the feedback control theory for server power control, which has been proven to be an effective way in improving performance and robustness of computing systems [22]. For example, [58] manages the system power using dynamic voltage scaling by controlling the synchronizing queues in multi-clock-domain processors. To prolong battery lifetime of a laptop, a feedback controller is studied in [31] to reduce the average power consumption.

To manage and fulfill application-level QoS requirements, the control theory is applied in computer server management [43] . Chen et al. [11] also develop a controller to manage the response time in a server cluster. Felter et al. [16] employ the open-loop control to manage sever power budget by by balancing the power between processor and memory components.

III. Compute Clusters

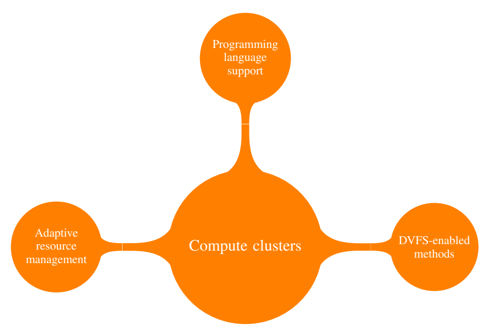

Fig. 2. Taxonomy of energy-aware cluster computing

A. Adaptive resource management

[40] develops algorithms to shutdown unused servers to reduce the average electricity consumption of individual servers.

The β algorithm [24] is developed by Hsu and Feng to control a cluster's power usage. It uses existing workload information (MPIS) to make DVFS decisions on CPU. It has been argued that while appropriate for some applications, the β algorithm is not suitable for many other applications as MIPS information is not so accurate for power consumption.

CPU MISER [21] is a run-time system for performance-directed, power-aware cluster computing. The CPU MISER contains a performance model, workload predicator and DVFS scheduler to provide system-wide, application-independent, fine-grained, DVFS-based power management for generic power-aware clusters.

Productional cluster schedulers, like LSF and Moab, have already provide adapted energy efficient techniques. SLURM provides a solution for energy saving through *CPUfreq* command for changing the processor frequency.

B. Programming language and runtime support

Some work develops compiler and language support for power aware scientific applications [19], [26]. This method provide fine-grained power management of applications, for example, in a function level, a loop level or a Building Block (BB) level. The language level power management for applications can return plausible results. However, this method brings afford for programers and application users to study new APIs and re-write, at least, re-compile their applications.

Some runtime systems are developed for power management of scientific applications. [25], [30]. These implementations develop some online system information logging, analysis and application prediction. With the help of developed system support, application behaviors can be controlled by dynamically scaling the supplied processor voltages.

[19] discusses a method which divides a MPI application into multiple phases with different voltage levels to reduce energy consumption.

[29] discusses power prediction techniques for MPI applications based on task aggregation. [28] proposes power-aware Hybrid MPI/OpenMP programming pattern based on the interaction between communication and computation without using DVFS.

C. DVFS-enabled methods

The Dynamic Voltage Frequency Scaling (DVFS) technology, which an effective technique to reduce processor power dissipation [23], [24], is widely used in the research of power-aware scheduling. The DVFS technique levels down processor clock speed and supply voltage during when CPU is idle or under-used, thus power consumption can be reduced. The DVFS tuning typically requires a proling of the application execution time for all available frequency/voltage settings. With the pre-knowledge of power consumption of the application at each processor operating frequency/voltage, a set of optimal processor operating frequency/voltage is selected for the application. Some example work are [17], [20].

The DVFS approach is limited by the structure of the applications. Only configurable applications can be inserted DVFS instructions for power control.

IV. Distributed virtual machines

Virtualization technologies such as VMware, Xen and Microsoft Virtual Servers can consolidate applications previously running on multiple physical servers onto a single physical server, and so effectively reduce the energy consumption of a data center by shutting down unused servers. It is thus an interesting topic to manage virtualized infrastructure aiming to reducing power consumption.

[57] proposes an architecture of two-layer control to guarantee the response time for virtualized enterprise servers while keeping low power consumption. The VirtualPower approach [36] is an online power management, which supports the isolated and independent operation assumed by guest virtual machines running on virtualized platforms and make it possible to control and globally coordinate the effects of the diverse power management policies applied by these VMs to virtualized resources. [50] implements a middleware which can dynamically schedule virtual machine with a power cost aware model. [51] proposes an scheduling algorithm for dynamically distributing virtual machines in a DVFS-enabled cluster environment. Usher [32] is an implementation of a virtual cluster management middleware in the GreenLight project [7]. Usher exposes dynamic levels of parallelism, and supports the distinctive needs of the GreenLight Instrument's computing environments and it applications, for example, large-scale scientific computing as well as network services with its resource allocation mechanisms and higher level policies for cluster resources.

V. Data center

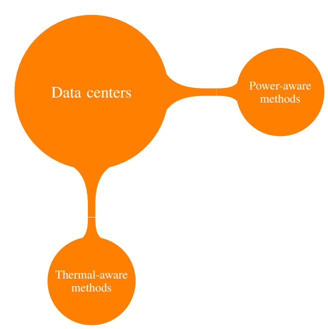

Fig. 3. Taxonomy of energy-aware data center computing

A. Power-aware methods

Some work focuses on the energy aware management in a cluster, server farms or a data center [4], [6], [8], [9], [14], [15], [39]. For example, [5], [10] develop the policies for shutdown unused compute nodes in a data center. [13] present a two-level control framework for managing the cluster-wide power. [41] locates statistical properties of concurrent resource usage in a large cluster or a data center and uses such properties to develop power management policies.

B. Thermal-aware methods

Some research develop thermal-aware methods to reduce energy-consumption of air-condition systems in data centers. [37] develops an adaptive resource control system that dynamically adjusts the resource shares to individual tiers of multi-tier application in order to meet application-level QoS goals while achieving high resource utilization in the data center.

[47] proposes a "Thermal-Aware Job Scheduling" and produces a number of results for thermal aware task scheduling in a data center. [35] proposes a software architecture for thermal aware data center management. [45], [46] exploits the abstract heat recirculation model to formulate the problem of minimizing the heat recirculation by appropriately assigning the incoming tasks around the servers.

[33] describes Weatherman, which is an automated, online, predictive thermal management for data centers. ANN techniques is used to learn and predict the complexities of the thermal topology in a data center.

[34], [42] develop temperature-aware workload placement algorithms and present the first comprehensive exploration of the benefits from these policies.

[2], [42] make a study on the methods and mechanisms to control data center temperatures with CFD based thermal model. [38] proposes a blueprint of Energy Aware Grid to provide a global utility infrastructure explicitly incorporating energy efficiency and thermal management among data centers.

[49] proposes the concept of task-temperature profile, which is the temperature increase along with the task execution. Wang et al. develops a thermal aware task scheduling algorithm [53], [55] by predicating resource temperatures based on online task-temperature profiles. In their algorithm, an online task-temperature profile is calculated with the pre-knowledge of task-temperature profile and RC-thermal model [44]. Another method to predict resource temperatures is using artificial intelligence techniques, such as support vector machine, neural network, generic algorithm [54], [56]. [52] develop a taxonomy study of performance metrics for green data centers.

VI. Computational Clouds and Grids

Some research are carried out to reach energy-efficiency in large-scale distributed systems, for example, Grids and Clouds. The GreenLight project and Green-Net project are typical examples.

The GreenLight project [7] develops research to measures, monitors, and optimizes the energy consumption of large-scale scientific applications from many different areas. The GreenLight project uses experimental and ethnographic methods to examine how access to energy costs influences researchers' behavior in using shared computational resources, by recording and analyzing data of their real-world behavior.

The Green-Net [1], [12] project, a cooperative research action developed at INRIA, aims to design energy-aware software frameworks dedicated to large scale distributed systems. The Green-Net framework collects information of energy usage and provide them to resources managers and schedulers. The Green-Net project also develops large scale experimental validations on the Grid5000 and DSLLAB testbeds.

VII. Summary and discussion

For the research to reduce energy consumption for high performance computing applications, especially in a large scale distributed computing system, there are several open questions unaddressed for this research:

- How to monitor the "energy-aware" aspects of a parallel computing system in multiple scales? There are multiple interesting system indexes, for example, power, temperature, and humidity. This monitoring process will be taken in multiple scales, for example, threads, processes, services, compute nodes, racks, compute clusters and compute Clouds.
- What is the energy usage profile and compute cost model for high performance applications in a large scale distributed compute system? There is still no effective methods to estimate energy usage and cost for users. For example, how much energy does a parallel application cost? How much additional energy will be used if a program changes all its float variable from single precision to double precision?
- How to develop tools, middleware and software framework to support energy efficient high performance applications in a parallel computing system? A single software solution a certain level may not effective as it may increase power consumption in another level or in another system component. Some time, the power consumption maybe reduced with tolerant compute performance loss. This result maybe unacceptable as application QoS, such as reliability and availability, is significantly degraded.
- Current "energy-aware" policies and solutions are not clearly defined by roles involved. An application user and a data center owner certainly have different interests and goals. Therefore, in the "green" computing context, the business model for users, application developers and resource owners are unclear. Software tools and service are missed to support the interaction and negotiation between them.

To address above open questions, A coordinated multi-level, multi-approach and multi-role methodology should be adopted. The solution follows the paradigm of "monitoring $\rightarrow$ analyzing $\rightarrow$ execution":

- computing systems and applications will be monitored,
- their energy-aware related behaviors and patterns will be analyzed, and
- Energy-aware policies and software framework will be implemented.

Desired research will be implemented and work coordinately in multiple levels of computing systems: a compute node, a rack, a data center and a compute Cloud. The software tools and services implemented will be built over programming model, application service and high level policy. Three roles will be involved in the energy-aware computing activities: application developers, users, and compute resource owners. Multiple system metrics and qualities and be investigated, for example, compute system performance such as response time, throughout, FLOPS, application QoS like reliability, availability, portability, security, development qualities, e.g., buildability, testability and understandability. system operational metrics, e.g., maintainability, serviceability. In specific, software toolkits and services are demanded:

- a monitoring or auditing service for a large scale distributed computing system,
- a software tool for workload characterization and model. This research will find the correlation between workload patterns and their energy usage,
- adaptive energy aware workload placement algorithms and services,
- energy aware programming models and middleware for data intensive applications,
- energy aware virtual machine management system in a large-scale distributed system,
- business models and SLA for public users, resource owners, and publicity.

REFERENCES

[1] INRIA Green-Net project. Website, 2009. http://www.ens-lyon.fr/LIP/RESO/Projects/GREEN-NET/.

[2] Abdlmonem H. Beitelmal and Chandrakant D. Patel. Thermo-Fluids Provisioning of a High Performance High Density Data Center. *Distributed and Parallel Databases*, 21(2-3):227–238, 2007.

[3] Luca Benini, Alessandro Bogliolo, Stefano Cavallucci, and Bruno Riccó. Monitoring system activity for os-directed dynamic power management. In *ISLPED '98: Proceedings of the 1998 international symposium on Low power electronics and design*, pages 185–190, New York, NY, USA, 1998. ACM.

[4] Ricardo Bianchini and Ram Rajamony. Power and energy management for server systems. *IEEE Computer*, 37(11):68–74, November 2004.

[5] P. Bohrer, D. Cohn, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, R. Rajamony, F. Rawson, and E. Van Hensbergen. Energy conservation for servers. In *In IEEE Workshop on Power Management for Real-Time and Embedded Systems*, pages 1–4, 2001.

[6] Pat Bohrer, Mootaz Elnozahy, Mike Kistler, Charles Lefurgy, Chandler McDowell, and Ramakrishnan Rajamony. The case for power management in web servers. In Robert Graybill and Rami Melhem, editors, *Power Aware Computing*. Kluwer Academic Publishers, 2002.

[7] Calit2. the GreenLight project. Web Page.

[8] Jeff Chase, Darrell Anderson, Prachi Thakur, and Amin Vahdat. Managing energy and server resources in hosting centers. In *Proceedings of the Eighteenth Symposium on Operating System Principles (SOSP'01)*, October 2001.

[9] Jeff Chase and Ron Doyle. Balance of power: Energy management for server clusters. In *Proceedings of the Eighth Workshop on Hot Topics in Operating Systems (HotOS'01)*, May 2001.

[10] Jeffrey S. Chase, Darrell C. Anderson, Prachi N. Thakar, and Amin M. Vahdat. Managing energy and server resources in hosting centers. In *In Proceedings of the 18th ACM Symposium on Operating System Principles (SOSP*, pages 103–116, 2001.

[11] Yiyu Chen, Amitayu Das, Wubi Qin, Anand Sivasubramaniam, Qian Wang, and Natarajan Gautam. Managing server energy and operational costs in hosting centers. In *SIGMETRICS '05: Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 303–314, New York, NY, USA, 2005. ACM.

[12] Georges Da Costa, Jean-Patrick Gelas, Yiannis Georgiou, Laurent Lefevre, Anne-Cecile Orgerie, Jean-Marc Pierson, Olivier Richard, and Kamal Sharma. The green-net framework: Energy efficiency in large scale distributed systems. In *IPDPS '09: Proceedings of the 2009 IEEE International Symposium on Parallel&Distributed Processing*, pages 1–8, Washington, DC, USA, 2009. IEEE Computer Society.

[13] Mark E. Femal and Vincent W. Freeh. Boosting data center performance through non-uniform power allocation. In *ICAC '05: Proceedings of the Second International Conference on Automatic Computing*, pages 250–261, Washington, DC, USA, 2005. IEEE Computer Society.

[14] Mootaz Elnozahy, Michael Kistler, and Ramakrishnan Rajamony. Energy conservation policies for web servers. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems USITS'03*, March 2003.

[15] W. M. Felter, T. W. Keller, M. D. Kistler, C. Lefurgy, K. Rajamani, R. Rajamony, F. L. Rawson, B. A. Smith, and E. Van Hensbergen. On the performance and use of dense servers. *IBM Journal of Research and Development*, 47(5):671–688, 2003.

[16] Wes Felter, Karthick Rajamani, Tom Keller, and Cosmin Rusu. A performance-conserving approach for reducing peak power consumption in server systems. In *ICS '05: Proceedings of the 19th annual international conference on Supercomputing*, pages 293–302, New York, NY, USA, 2005. ACM.

[17] Xizhou Feng, Rong Ge, and Kirk W. Cameron. Power and energy profiling of scientific applications on distributed systems. In *IPDPS*, 2005.

[18] Jason Flinn and M. Satyanarayanan. Energy-aware adaptation for mobile applications. In *SOSP '99: Proceedings of the seventeenth ACM symposium on Operating systems principles*, pages 48–63, New York, NY, USA, 1999. ACM.

[19] Vincent W. Freeh and David K. Lowenthal. Using multiple energy gears in mpi programs on a power-scalable cluster. In Keshav Pingali, Katherine A. Yelick, and Andrew S. Grimshaw, editors, *PPOPP*, pages 164–173. ACM, 2005.

[20] Vincent W. Freeh, Feng Pan, Nandini Kappiah, David K. Lowenthal, and Robert Springer. Exploring the energy-time tradeoff in mpi programs on a power-scalable cluster. In *IPDPS*. IEEE Computer Society, 2005.

[21] R. Ge, Feng X., Feng W., and K.W. Cameron. CPU MISER: A Performance-Directed, Run-Time System for Power-Aware Clusters. In *Proceedings of the 2007 International Conference on Parallel Processing*. IEEE Computer Society Washington, DC, USA, 2007.

[22] Joseph L. Hellerstein, Yixin Diao, Sujay Parekh, and Dawn M. Tilbury. *Feedback Control of Computing Systems*. John Wiley & Sons, 2004.

[23] Chung hsing Hsu and Wu chun Feng. A feasibility analysis of power awareness in commodity-based high-performance clusters. In *Proceedings of IEEE International Conference on Cluster Computing*, pages 1–10, Los Alamitos, CA, USA, 2005. IEEE Computer Society.

[24] C. Hsu and W. Feng. A power-aware run-time system for high-performance computing. In *Proceedings of the 2005 ACM/IEEE conference on Supercomputing*. IEEE Computer Society Washington, DC, USA, 2005.

[25] Chung-Hsing Hsu and Wu chun Feng. A power-aware run-time system for high-performance computing. In *SC*, page 1. IEEE Computer Society, 2005.

[26] Chung-Hsing Hsu and Ulrich Kremer. The design, implementation, and evaluation of a compiler algorithm for cpu energy reduction. In *PLDI*, pages 38–48. ACM, 2003.

[27] Alvin R. Lebeck, Xiaobo Fan, Heng Zeng, and Carla Ellis. Power aware page allocation. *SIGPLAN Not.*, 35(11):105–116, 2000.

[28] Dong Li, Bronis R. de Supinski, Martin Schulz, Kirk W. Cameron, and Dimitrios S. Nikolopoulos. Hybrid mpi/openmp power-aware computing. In *IPDPS*, pages 1–12, 2010.

[29] Dong Li, Dimitrios S. Nikolopoulos, Kirk W. Cameron, Bronis R. de Supinski, and Martin Schulz. Power-aware mpi task aggregation prediction for high-end computing systems. In *IPDPS*, pages 1–12, 2010.

[30] T. Li and L. K. John. Run-time modeling and estimation of operating system power consumption. In *SIGMETRICS International conference on Measurement and modeling of computer systems*, pages 160–171, New York, NY, USA, 2003. ACM.

[31] Eui-Young Chung Luca, Luca Benini, and Giovanni De Micheli. Dynamic power management using adaptive learning tree. In *In Proc. International Conference on Computer-Aided Design*, pages 274–279, 1999.

[32] Marvin McNett, Diwaker Gupta, Amin Vahdat, and Geoffrey M. Voelker. Usher: An Extensible Framework for Managing Clusters of Virtual Machines. In *Proceedings of the 21st Large Installation System Administration Conference (LISA)*, November 2007.

[33] Justin Moore, Jeffrey Chase, and Parthasarathy Ranganathan. Weatherman: Automated, Online, and Predictive Thermal Mapping and Management for Data Centers. In *The Third IEEE International Conference on Autonomic Computing*, June 2006.

[34] Justin D. Moore, Jeffrey S. Chase, Parthasarathy Ranganathan, and Ratnesh K. Sharma. Making scheduling "cool": Temperature-aware workload placement in data centers. In *USENIX Annual Technical Conference, General Track*, pages 61–75, 2005.

[35] Tridib Mukherjee, Qinghui Tang, Corbett Ziesman, Sandeep K. S. Gupta, and Phil Cayton. Software Architecture for Dynamic Thermal Management in Datacenters. In *COMSWARE*, 2007.

[36] Ripal Nathuji and Karsten Schwan. Virtualpower: coordinated power management in virtualized enterprise systems. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, pages 265–278, New York, NY, USA, 2007. ACM Press.

[37] Pradeep Padala, Kang G. Shin, Xiaoyun Zhu, Mustafa Uysal, Zhikui Wang, Sharad Singhal, Arif Merchant, and Kenneth Salem. Adaptive

control of virtualized resources in utility computing environments. *SIGOPS Oper. Syst. Rev.*, 41(3):289–302, 2007.

[38] Chandrakant Patel, Ratnesh Sharma, Cullen Bash, and Sven Graupner. Energy Aware Grid: Global Workload Placement based on Energy Efficiency. Technical report, HP Laboratories, 2002.

[39] E. Pinheiro, R. Bianchini, E. V. Carrera, and T. Heath. Dynamic cluster reconfiguration for power and performance. In Luca Benini, Mahmut Kandemir, and J. Ramanujam, editors, *Compilers and Operating Systems for Low Power*. Kluwer Academic Publishers, 2002.

[40] Eduardo Pinheiro, Ricardo Bianchini, Enrique V. Carrera, and Taliver Heath. Load balancing and unbalancing for power and performance in cluster-based systems. In *In Workshop on Compilers and Operating Systems for Low Power*, 2001.

[41] Parthasarathy Ranganathan, Phil Leech, David Irwin, and Jeffrey Chase. Ensemble-level power management for dense blade servers. In *Proceedings of the 33rd annual international symposium on Computer Architecture*, pages 66–77, Washington, DC, USA, 2006. IEEE Computer Society.

[42] Ratnesh K. Sharma, Cullen Bash, Chandrakant D. Patel, Richard J. Friedrich, and Jeffrey S. Chase. Balance of power: Dynamic thermal management for internet data centers. *IEEE Internet Computing*, 9(1):42–49, 2005.

[43] Vivek Sharma, Arun Thomas, Tarek Abdelzaher, Kevin Skadron, and Zhijian Lu. Power-aware qos management in web servers. In *RTSS '03: Proceedings of the 24th IEEE International Real-Time Systems Symposium*, page 63, Washington, DC, USA, 2003. IEEE Computer Society.

[44] Kevin Skadron, Tarek F. Abdelzaher, and Mircea R. Stan. Control-Theoretic Techniques and Thermal-RC Modeling for Accurate and Localized Dynamic Thermal Management. In *HPCA*, pages 17–28, 2002.

[45] Qinghui Tang, Sandeep K. S. Gupta, and Georgios Varsamopoulos. Thermal-aware task scheduling for data centers through minimizing heat recirculation. In *CLUSTER*, pages 129–138, 2007.

[46] Qinghui Tang, Sandeep K. S. Gupta, and Georgios Varsamopoulos. Energy-Efficient Thermal-Aware Task Scheduling for Homogeneous High-Performance Computing Data Centers: A Cyber-Physical Approach. *IEEE Trans. Parallel Distrib. Syst.*, 19(11):1458–1472, 2008.

[47] Qinghui Tang, Tridib Mukherjee, Sandeep K. S. Gupta, and Phil Cayton. Sensor-Based Fast Thermal Evaluation Model For Energy Efficient High-Performance Datacenters. In *Proceedings of the Fourth International Conference on Intelligent Sensing and Information Processing*, pages 203–208, Oct. 2006.

[48] Amin Vahdat, Alvin Lebeck, and Carla Schlatter Ellis. Every joule is precious: the case for revisiting operating system design for energy efficiency. In *Proceedings of the 9th workshop on ACM SIGOPS European workshop*, pages 31–36, New York, NY, USA, 2000. ACM.

[49] Daniel C. Vanderster, Amirali Baniasadi, and Nikitas J. Dimopoulos. Exploiting task temperature profiling in temperature-aware task scheduling for computational clusters. In *Asia-Pacific Computer Systems Architecture Conference*, pages 175–185, 2007.

[50] Akshat Verma, Puneet Ahuja, and Anindya Neogi. pmapper: power and migration cost aware application placement in virtualized systems. In *Middleware '08: Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware*, pages 243–264, New York, NY, USA, 2008. Springer-Verlag New York, Inc.

[51] Gregor von Laszewski, Lizhe Wang, Andrew J. Younge, and Xi He. Power-aware scheduling of virtual machines in dvfs-enabled clusters. In *CLUSTER*, pages 1–10, 2009.

[52] Lizhe Wang and Samee Khan. Review of performance metrics for green data centers: a taxonomy study. *Journal of Supercomputing*, 2011.

[53] Lizhe Wang, Samee Khan, and Jai Dayal. Thermal aware workload placement with task-temperature profiles in a data center. *The Journal of Supercomputing*, pages 1–24, 2011. 10.1007/s11227-011-0635-z.

[54] Lizhe Wang, Gregor von Laszewski, and Jai Dayal. Thermal Aware Workload Scheduling with Backfilling for Green Data Centers. In *Proceedings of the 28th IEEE International Performance Computing and Communications Conference*, Dec. 2009.

[55] Lizhe Wang, Gregor von Laszewski, Jai Dayal, Xi He, Andrew J. Younge, and Thomas R. Furlani. Towards thermal aware workload scheduling in a data center. In *Proceedings of the 10th International Symposium on Pervasive Systems, Algorithms and Networks*, Dec. 2009.

[56] Lizhe Wang, Gregor von Laszewski, Fang Huang, Jai Dayal, Tom Frulani, and Geoffrey Fox. Task scheduling with ann-based temperature prediction in a data center: a simulation-based study. *Engineering with Computers*, pages 1–11, 2011. 10.1007/s00366-011-0211-4.

[57] Yefu Wang, Xiaorui Wang, Ming Chen, and Xiaoyun Zhu. Power-efficient response time guarantees for virtualized enterprise servers. In *RTSS '08: Proceedings of the 2008 Real-Time Systems Symposium*, pages 303–312, Washington, DC, USA, 2008. IEEE Computer Society.

[58] Qiang Wu, Philo Juang, Margaret Martonosi, Li-Shiuan Peh, and Douglas W. Clark. Formal control techniques for power-performance management. *IEEE Micro*, 25(5):52–62, 2005.