
Power Management Techniques for Data Centers: A Survey

Sparsh Mittal

Future Technologies Group

Oak Ridge National Laboratory

Oak Ridge, Tennessee, USA 37830

mittals@ornl.gov

Abstract

With growing use of internet and exponential growth in amount of data to be stored and processed (known as “big data”), the size of data centers has greatly increased. This, however, has resulted in significant increase in the power consumption of the data centers. For this reason, managing power consumption of data centers has become essential. In this paper, we highlight the need of achieving energy efficiency in data centers and survey several recent architectural techniques designed for power management of data centers. We also present a classification of these techniques based on their characteristics. This paper aims to provide insights into the techniques for improving energy efficiency of data centers and encourage the designers to invent novel solutions for managing the large power dissipation of data centers.

Keywords

Data Centers, Power Management, Low-power Design, Energy Efficiency, Green Computing, DVFS, Server Consolidation

I. INTRODUCTION

As we move to the exascale era, the requirements of data storage have grown exponentially and hence, the power consumption of data-centers has also increased [1, 2]. It has been reported that in year 2006 alone, the datacenters and servers in U.S. consumed 61 billion kilowatt hours of electricity, which is 1.5% of all US electricity consumption and has a monetary cost of \$4.5 billion [3]. Many data centers which support cloud services employ tens of thousands of servers which draw tens of Mega-Watts of power at peak [4]. It has been estimated that such data centers draw power worth \$9.3 million per year [4]. Moreover, the worldwide expenditure on enterprise power supply and cooling has been estimated to be more than \$30 billion [5]. Hence, achieving energy efficiency in data centers has become an issue of paramount importance.

Recently, several techniques have been proposed for managing power consumption of data centers (e.g. [6–10]). Due to large power consumption levels of data centers, use of these techniques has become essential to maintain both energy efficiency and cost efficiency. In this paper, we highlight the need of achieving energy efficiency in data centers and survey several techniques which have been proposed for ensuring green operation of data centers. Since it is not possible to include detailed discussion of a large

number of techniques, we take the following approach to limit the scope of this paper. We focus on architectural-level and operation-level techniques for improving energy efficiency and not on circuit-level (device-level) techniques. Moreover, we only discuss the key design principles of different techniques and do not discuss the quantitative results, since different techniques have been tested using different evaluation platforms. We only include those techniques which have been evaluated for improving energy efficiency, and do not include the techniques aimed at improving performance, although they are also expected to improve energy efficiency.

The rest of the paper is organized as follows. In the next section, we present the factors which lead to increased power consumption. We also discuss their impact along with the need of managing data center power consumption. We then present a classification of power management techniques based on their characteristics, and then discuss some of these power management techniques in detail. Finally, we provide a discussion of future research directions and concluding remarks.

II. BACKGROUND

A. *Reasons behind Increase in Power Consumption*

Due to recent growth in use of internet, the demand placed on data-centers has increased. Modern data centers typically contain up to tens of thousands of servers and provide 24×7 services to hundreds or thousands of users. As an example, YouTube serves as much as 100 million videos a day, while Facebook has nearly 400 million active users and 3 billion photos uploaded every month [11]. These images, videos and text data are processed using computation-intensive software programs; and stored and accessed from data centers [12, 13]. Further, in recent years, use of high-performance computing techniques has also increased, which require trading off energy consumption for obtaining increased performance. For this reason, data centers have grown in sizes; which has led to increase in their power consumption.

B. *Power Consumption Levels*

Recent years have witnessed a huge increase in the power consumption of servers, data centers and supercomputers [14–16]. The total energy consumption of data centers as a percentage of total US energy consumption has doubled between 2000 and 2007 [2]. As an example, peak power consumption of the most powerful supercomputers in the TOP500 list of supercomputers ranges in tens of megawatts [17, 18]. This amount of power is sufficient to fulfill the demands of a city of 40,000 residents. Research has also shown that the buildings with data centers can be as much as 40 times more energy intensive than the conventional office buildings [19].

The high levels of power consumption demand a costly cooling infrastructure [20]. For example, a 30,000 square feet data center with 10 megawatts power consumption consumes on the order of \$5 million for cooling in a year [21]. Moreover, for every watt of power consumed in the computing equipment, an additional 0.5 to 1W of power is required to operate the cooling system itself [21], which further adds to the cost.

C. Need of Power Management

It has been shown that most of the time, the modern servers operate between 10% to 50% of maximum possible utilization [22, 23]. Further, at these utilization levels, the server energy efficiency also becomes very low [24]. Thus, despite the fact that the average utilization remains low; there exist frequent, brief bursts of activity, and to meet the requirements of service-level-agreements (SLAs), operators are forced to allocate high amount of resources, which leads to poor energy efficiency [25].

Power management is also important from an economic point of view, since effective power management also improves operational efficiencies and increases compaction. Electricity costs for powering servers forms a major cost of operation in data centers and it has been estimated that in near future, energy costs may contribute even more than the cost of IT ([26, 27]). Further, a high ratio of cooling power to computing power restricts the compaction and consolidation possible in data centers, which results in increased operation costs. For example, the high power density poses significant challenges in routing the large amounts of power needed per rack. Currently the power delivery in typical data centers is near 60 Amps per rack and it is expected to reach the limit of power delivery, which will severely affect the operation of servers [23].

The large power consumption and high concentration of nodes in data centers leads to increased node failures. It has been observed that a 15 degree Celsius rise increases the failure rates in hard-disks by a factor of two [28]. Hence, maintaining the computer systems at proper temperature is important for ensuring maximum reliability, longevity, and large return on investment. Finally, large power consumption also has adverse environmental impact, e.g. large carbon emission [29–33]. For these reasons, the design of green solutions for modern data centers has become a topic of paramount importance. Hence, in this paper, we survey techniques for managing power consumption in data centers.

III. OVERVIEW OF POWER MANAGEMENT TECHNIQUES

In recent years, researchers have proposed several techniques for managing power consumption in data centers. While it is very difficult to draw sharp boundaries of classification; for the purpose of study, we broadly classify the techniques in the following four types.

- 1) DVFS (dynamic voltage/frequency scaling) based techniques [34–43].
- 2) Techniques which transition the server/node to low-power or turned-off state; or use server-consolidation based approach to allocate only required amount of server resources [23, 43–55].
- 3) Workload management or task scheduling based techniques [56–63].
- 4) Thermal-aware or thermal-management techniques which take into account the thermal properties [56, 59, 61, 64–70]. Further, some techniques address the issues related to cooling in data centers [21, 59, 71].

Based on other characteristics/parameters, the techniques can be further classified. While most techniques aim to reduce average power (energy), a few techniques aim to reduce peak power consumption [72], or limit power (called power capping) [73, 74]. A few techniques aim to save energy while keeping the performance degradation bounded or maintaining QoS (quality-of-service) guarantee [39, 41, 43, 50, 75–79].

Several techniques use analytical approach and offer control-theoretical algorithms with provable guarantees [38, 74, 79–84], while most other techniques use system approach and focus only on implementation. Some researchers propose use of renewable energy sources [85, 86]. Some researchers aim to reduce disk energy [87–89], while others focus on saving main memory energy [35, 90–93] in data centers.

IV. POWER MANAGEMENT TECHNIQUES

A. DVFS Based Techniques

Dynamic voltage and frequency scaling (DVFS) is a widely-used power-management technique where the clock frequency of a processor is dynamically adjusted to allow a corresponding reduction in the supply voltage to achieve power saving [94]. DVFS is especially useful for memory-bound workloads. The power consumption of a CMOS (Complementary metal oxide semiconductor) circuit is given by

$$P = P_{static} + CFV^2 \quad (1)$$

Here C shows the capacitance of the transistor gates, F shows the operating frequency and V shows the supply voltage. The frequency at which circuit is clocked determines the voltage required for stable operation; and hence, by intelligently reducing the frequency, the supply voltage can also be reduced; which leads to significant power savings due to V^2 relationship shown above. The limitation of DVFS, however, is that a reduction in frequency also reduces the performance of the circuit and hence, DVFS may adversely affect the performance of the processor. For this reason, DVFS needs to be intelligently applied, to maintain high performance. In data centers, DVFS technique is applied to manage the power consumption of multicore processors, DRAM memories and other components.

Sharma et al. [41] propose adaptive algorithms for dynamic voltage scaling in QoS-enabled web servers. Their algorithms aim to minimize energy consumption subject to service level agreements (SLAs). The algorithms are implemented inside the Linux kernel. Their algorithms minimize the energy consumption using a feedback loop which regulates the frequency and voltage levels to keep the instantaneous utilization bounded.

Hsu and Feng [39] propose an algorithm for the dynamically varying voltage and frequency of the processor to realize energy saving while keeping the performance loss bounded. The algorithm takes decisions at the end of fixed time period. Their algorithm uses an estimation model to relate the intensity level of off-chip accesses to total execution time. Using this, the algorithm computes the lowest CPU frequency which keeps the performance loss bounded, while saving largest possible amount of energy. For multiprocessor environment, the same algorithm is repeated for each processor to set the frequency of the processor individually.

Xu et al. [43] propose a technique for saving energy in embedded clusters. Their technique adjusts the number of active nodes based on the system load. In their technique, each node in the cluster performs dynamic voltage scaling independently and runs at the lowest frequency at which it can keep up with the request arrival rate. To prevent the system from reacting to short-term changes in workload, at a time, their technique changes only one node from active to inactive and vice-versa in each interval.

Horvath et al. [38] present methods to dynamically adjust the server voltages to minimize the total system power consumption, while also meeting end-to-end delay constraints in a multi-tier web service environment. In their architecture, server machines use DVS-capable processors. They propose coordinated distributed voltage scaling policy, where decisions on frequency adjustments are made on each locally while minimizing overall power consumption. Compared to other methods such as server on/off, use of dynamic voltage scaling incurs less overhead which facilitates aggressive energy saving.

Deng et al. [91] use DVFS mechanism to save memory energy. Their technique lowers the frequency of DRAM devices, memory channels and memory controllers at the time of low memory activity. This reduces the memory power consumption. They have also extended their technique for coordinating DVFS across multiple memory controllers, memory devices and channels to reduce the overall system power consumption [35].

B. Server Consolidation and Power State Transitioning Based Techniques

As discussed before, modern servers typically operate at low utilization levels. Moreover, to cater to the peak demand and service-level-agreements and ensure reliability, high amount of server resources needs to be allocated which leads to poor energy efficiency [24]. To address this challenge, many approaches have been proposed. Server consolidation is one such approach aimed to ensure efficient usage of server resources by reducing the total number of servers required by a data center, while still delivering same throughput. In this approach, the existing applications are consolidated onto fewer servers, such that unused servers can be transitioned into low-power (or turned-off) state and the used servers can be operated at full utilization levels. Another approach is transitioning the server resources into low-power mode during periods of low activity. These approaches have been widely used to improve energy efficiency of data centers.

Chun et al. [95] propose a hybrid data center design which uses heterogeneous platforms to save power. Under low utilization levels, their technique transfers the running tasks from a high-power, high performance system to a low-power, low-performance system and turns off the higher power servers. Thus, server virtualization along with task migration reduces the number of active servers by consolidating many servers with low-average utilization to a few servers operating at high utilization. The limitation of server consolidation, however, is that it may lead to slow response times and high transition costs.

For ensuring energy-efficient operation of data centers, Chase et al. [46] discuss a system resource management based approach. Their method controls server allocation and routing of requests to selected servers using a dynamically reconfigurable switch; and thus enables achieving a trade-off between service quality and cost. For optimizing energy consumption, incoming request traffic is continuously monitored and only desired amount of server resources are allocated such that the service level agreements can be fulfilled. Since internet sites observe highly varying usage patterns (e.g. as much as 11:1 peak-to-trough ratio of usage), their method enables intelligent adaptive resource provisioning.

Ranganathan et al. [23] propose a technique for managing server power at the ensemble (i.e. collection of systems) level instead of individual server level. Their design observes the resource-usage trends across multiple systems. By taking advantage of inter-server variations, their technique allows active server to

steal power from the inactive servers. Their technique facilitates reductions in the requirements for power delivery, power consumption, and cooling in the data centers.

To leverage the variation in workload for saving energy, Anagnostopoulou et al. [44] propose a “barely-alive” server design. Their design transitions the servers to a barely-alive power state, where the server can be still accessed, even if many of its other components are turned off. Their design uses a small embedded processor to only keep the memory of idle servers active so that in-memory application code/data remain unaffected, and the free memory space can be used for cooperative application data caching.

To reduce the energy consumption of servers during periods of reduced load, Rusu et al. [77] discuss a cluster-wide QoS-aware technique which uses dynamic reconfiguration based approach. For a given workload, their algorithm dynamically decides the servers which need to remain turned on/off to minimize global power consumption. To minimize the time penalty of server on/off, the algorithm accounts for the booting time of the server and turns on a server before it is actually needed. The authors also use dynamic voltage scaling (DVS) to conduct QoS-aware power management. The authors show the use of their technique in the context of a web server.

Ghosh et al. [49] propose a technique to save energy in data centers by using out-of-band management processors which are typically used for remotely managing a server, to satisfy the I/O requests from a remote server. By transferring the load from the primary server to the management processor, their technique allows the primary server to stay in low-power state for longer time, which improves the energy efficiency.

To utilize the opportunity of energy saving at multiple levels, Da Costa et al. [48] present an integrated framework, called GREEN-NET. GREEN-NET provides a multi-faceted approach for saving energy in clouds and grids. Their approach has three levels. At first level, it increases the awareness of users of their energy consumption. At second level, it involves the users in decisions to trade-off performance for saving energy. Finally, at third level, it conducts adaptive management of grids by techniques such as server turn-off to save energy.

L. Liu et al. [50] use virtual-machine migration approach where a virtual machine (VM) is transferred across physical computers to enable server consolidation and allow more computers to be turned-off. Their technique dynamically decides the time to trigger VM migration, and the alternative physical machines to achieve optimal VM placement. Their technique enables live migration of VMs, such that the users can hardly notice that their applications are being or have been migrated.

Leverich et al. [55] use per-core power-gating (PCPG) approach to manage power consumption of multicore processors. Their technique works by selectively turning on/off the power supply to individual cores of a multicore processor based on the utilization and quality-of-service requirements. The authors have also shown that by synergistically combining their technique with DVFS technique, additional power savings can be obtained.

C. Workload Scheduling Based Techniques

Modern data centers typically have a large number of servers and hence, the decision about placement of workloads on specific servers significantly affects the heat-dissipation and power-consumption. A

poor placement may greatly increase the temperature of the building which will further increase the heat-dissipation of the servers and also increase the cooling requirements. Hence, workload-scheduling techniques have been proposed which intelligently place the workloads on available servers with the goal of saving power, reducing the temperature and the cooling requirements.

Nathuji and Schwan [96] propose a dynamic power management technique to support the isolated and independent operation of VMs running on a virtualized platform, while globally coordinating the diverse power management strategies applied by the VMs to the virtualized resource. Their technique uses a set of virtualized power states to permit guest VMs to run their own, independent power management methods. Further, guest VM-level power management policies act upon these states to minimize power consumption, while meeting application requirements.

A thermal-aware technique for spatial workload placement in data centers is presented by Banerjee et al. [56]. One of their techniques uses the information about dynamic behavior of computer room air conditioned (CRAC) to place the jobs in a manner which reduces the cooling demands from the CRACs. By integrating this technique with a temporal scheduling technique, the decision about when and on which server to execute a job can be taken, leading to a spatio-temporal scheduling technique.

Bradley et al. [57] present a predictive power management technique for saving power in parallel computer systems. In web-based applications where the load on a system varies a lot, minimizing power consumption while meeting the demands of the workload is challenging. For such systems, their technique projects workload ahead of time to allow adequate resources to be powered on and kept ready for work when required to execute the workload. Their algorithm uses CPU utilization data to estimate the workload demand. When the utilization becomes larger than a predefined threshold, additional servers are powered on to reduce utilization on all servers to below that threshold. Conversely, when utilization at all servers is below that threshold and there is adequate capacity in the resulting server group to absorb the load of at least one server without any resource on any server being over-utilized, one or more server are powered off.

D. Thermal-aware Power Management Techniques

Several power management techniques work in thermal-unaware manner, i.e. they do not take the temperature-dependence of server power consumption into account. However, the heat-dissipation of processor components has strong temperature dependency and hence, an increase in operating temperature leads to increased heat-dissipation, which further increases the temperature and so on. This may drastically increase the component failure rates. To address this, several techniques have been proposed which manage the power consumption of the data centers while taking into account the thermal properties.

To minimize the total energy costs of data-center operation while providing a reliable thermal environment, Tang et al. [63] propose a thermal-aware task-scheduling technique. In a blade server, multiple blades are integrated into each chassis, in which blades share a common supply and cooling fan. Each blade may itself have many processors. Operation of a chassis incurs chassis startup power consumption, along with actual power consumption of the blades. Hence, the power consumption cost of adding a task to one chassis may be different depending upon whether it involves waking up an idle

chassis or an idle blade. Based on this observation, they evaluate several thermal-aware task-scheduling techniques, which exercise the trade-off between cost of start-up and power saving coming from turning off the blade servers.

Moore et al. [59] present techniques to control the heat generation in servers using temperature-aware workload placement. Their techniques take into account the fact that since the heat may travel a large distance inside the data center before arriving at a temperature sensor, cooling inefficiencies can often arise in locations which are spatially uncorrelated from the place where the heat originates. Based on this, their technique uses the information about steady-state hot-spots and cold-spots in the data center to develop scheduling algorithms.

As discussed above, in data centers, work-scheduling algorithms have a significant influence on temperature distribution. Li et al. [97] present a thermal forecasting model to predict the temperatures near the servers in a data center. Their model is based on continuous streams of temperature and airflow measurements. Their model integrates both physical laws and sensor observations in the data center. It uses the data obtained from the sensors to learn the parameters of the cyber-physical system of the data center.

E. Other Techniques

While most techniques aim to reduce energy consumption, a few techniques have also been proposed which aim to control peak power consumption. Lefurgy et al. [74] present a technique for controlling the peak power consumption of a high-density server. Their technique uses a feedback controller to periodically select the highest performance state while keeping the system within a fixed power constraint. The control loop is designed using the control theoretic methodology which helps in obtaining analytical guarantee on system stability and controller performance, despite variation in workloads. Their technique allows controlling power at different time-granularity, e.g. one second, eight seconds etc. The authors have shown that their technique performs better than the ad-hoc and open-loop techniques.

To fulfill the needs of several performance-critical and data-intensive applications that execute on many data center platforms [98–101], a large amount of main memory resources need to be provisioned. Hence, improving the energy efficiency of main memories is important to achieving data center energy efficiency. Yoon et al. [92] propose a technique for saving memory power consumption in data centers by intelligently utilizing low power mobile DRAM components. By using buffering mechanism to aggregate the data outputs from multiple ranks of low frequency mobile DRAM devices (such as 400MHz LPDDR2), their technique enables achieving high bandwidth and high storage capacity equal to server-class DRAM devices (such as 1600MHz DDR3).

Chatterjee et al. [90] propose a technique to take advantage of heterogeneity in DRAM memories to accelerate critical word access. In recent years, processors use DRAM chips with vastly differing latency and energy characteristics. Their technique recognizes the critical word in a cache line which is placed in a low-latency region of the main memory. The remaining non-critical words of the cache line are placed in a low-energy consuming region. Thus, their approach enables achieving high performance while reducing memory energy.

A technique to save the energy spent in unbeneficial refreshes in DRAM is proposed by J. Liu et al. [93]. In a DRAM, only a small number of cells need to be refreshed at the minimum conservative refresh rate and the remaining cells can be refreshed at a much higher rate, while still maintaining their charge. Using this observation, their technique groups DRAM rows in multiple bins and uses different refresh interval for different bins. Further, by refreshing most of the cells with lower frequency than the leaky cells, they achieve reduction in the number of refresh operations required which leads to reduction in memory power consumption.

Keys et al. [102] characterize a wide variety of clusters to find energy-efficient data center building blocks which are useful for data-intensive applications. The authors build homogeneous clusters using embedded, mobile and server systems and execute DryadLINQ applications on these clusters in order to study their energy efficiency for different application types. The authors have observed that their cluster was significantly more energy-efficient than ones with low-power server processors or embedded processors.

To maximize the benefits from the provided quality of computational services, while maintaining energy efficiency, Parolini et al. [82] present an energy saving technique for data centers. They model the data center in terms of mutually coupled cyber and physical systems. Here cyber component refers to the computational network representing the distribution and flow of computational tasks. The physical network refers to the thermal network which characterizes the distribution of thermal energy. The authors have shown that compared to the traditional approach in which the cyber and physical resources are controlled independently, their coordinated control strategy achieves more effective management of power consumption.

V. FUTURE RESEARCH DIRECTIONS

With CMOS scaling, the leakage energy consumption is increasing and hence the dynamic range of energy consumption that DVFS can utilize has reduced. Further, the design complexity of multicore processors also hinders straightforward use of DVFS for saving energy. For this reason, the effectiveness of DVFS based techniques has been reducing. Instead, the state-of-the-art in power management in data centers has focused on achieving energy-proportional computing by addressing hardware-platform heterogeneity [103]. Since modern data centers use tens to hundreds of servers with possibly different configurations, there exists significant heterogeneity in the computing servers. This heterogeneity is likely to increase further with increasing sizes of the data centers. Since workloads are sensitive to hardware platforms, a heterogeneity-oblivious scheduling of workloads on servers may lead to significant performance degradation. Thus, novel techniques have been proposed which take heterogeneity into account for scheduling workloads. Also, as discussed above, research efforts are also being directed towards utilizing renewable energy sources and dynamic power-state transitioning techniques. Also, researchers are exploring low-leakage high-density technologies such as non-volatile memory for reducing the energy consumption of computing systems [104]. These techniques reduce the carbon footprint of data centers, while also improving their energy efficiency.

It is expected that in the near future, several of the above mentioned solutions will be deployed in a synergistic manner for providing even larger savings. A key challenge for this is designing a coordination

framework which is flexible and allows seamless integration of different techniques.

VI. CONCLUSION

With the growing use of internet and requirement of data-storage and processing, the size of modern data centers has greatly increased. This has led to significant increase in the power consumption levels of the data centers. Moreover, the power consumption of data centers is approaching the limit imposed by thermal limitations of cooling solutions and power delivery. Also, since data centers are already consuming tens of Mega Watts, they are also stressing the capabilities of power generation facilities.

As the complexity of operation of data centers increases, power management techniques which also ensure high-performance and low-costs are expected to become a crucial part of future enterprise architectures. In this paper, we highlighted the need for power management in data centers. We reviewed several techniques which have been proposed for reducing power consumption of data centers and classified them based on their characteristics. We believe that our survey will enable the researchers to gain insights into the state-of-the-art in power management of data centers and motivate them to propose innovative solutions for architecting future green data centers.

REFERENCES

- [1] J. Baliga, R. W. Ayre, K. Hinton, and R. Tucker, "Green cloud computing: Balancing energy in processing, storage, and transport," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 149–167, 2011.
- [2] K. Kant, "Data center evolution: A tutorial on state of the art, issues, and challenges," *Computer Networks*, vol. 53, no. 17, pp. 2939–2965, 2009.
- [3] R. Brown *et al.*, "Report to congress on server and data center energy efficiency: Public law 109-431," 2008.
- [4] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 68–73, 2008.
- [5] C. Patel and P. Ranganathan, "Enterprise power and cooling," *ASPLOS Tutorial*, 2006.
- [6] R. A. Bergamaschi, L. Piga, S. Rigo, R. Azevedo, and G. Araújo, "Data center power and performance optimization through global selection of p-states and utilization rates," *Sustainable Computing: Informatics and Systems*, vol. 2, no. 4, pp. 198–208, 2012.
- [7] G. Cabusao, M. Mochizuki, K. Mashiko, T. Kobayashi, R. Singh, T. Nguyen, and X. P. Wu, "Data center energy conservation utilizing a heat pipe based ice storage system," in *CPMT Symposium Japan, 2010 IEEE*, 2010, pp. 1–4.
- [8] J. Chang, J. Meza, P. Ranganathan, A. Shah, R. Shih, and C. Bash, "Totally green: evaluating and designing servers for lifecycle environmental impact," in *ACM SIGARCH Computer Architecture News*, vol. 40, no. 1, 2012, pp. 25–36.
- [9] H. Chen, Y. Li, and W. Shi, "Fine-grained power management using process-level profiling," *Sustainable Computing: Informatics and Systems*, vol. 2, no. 1, pp. 33–42, 2012.

- [10] G. L. T. Chetsa, L. Lefevre, J. Pierson, P. Stolf, and G. Da Costa, "Beyond cpu frequency scaling for a fine-grained energy control of hpc systems," in *Computer Architecture and High Performance Computing (SBAC-PAD), 2012 IEEE 24th International Symposium on*. IEEE, 2012, pp. 132–138.
- [11] A. Hepburn, "Facebook statistics, stats & facts for 2011," *Digital Buzz*, 2011.
- [12] S. Mittal, S. Gupta, and A. Mittal, "BioinQA: metadata-based multi-document QA system for addressing the issues in biomedical domain," *International Journal of Data Mining, Modelling and Management*, vol. 5, no. 1, pp. 37–56, 2013.
- [13] S. Mittal and A. Mittal, "Versatile question answering systems: seeing in synthesis," *International Journal of Intelligent Information and Database Systems*, vol. 5, no. 2, pp. 119–142, 2011.
- [14] S. Mittal, "A survey of architectural techniques for dram power management," *International Journal of High Performance Systems Architecture*, vol. 4, no. 2, pp. 110–119, 2012.
- [15] Y. Zhang and N. Ansari, "Green data centers," *Handbook of Green Information and Communication Systems*, p. 331, 2012.
- [16] S. Mittal, Z. Zhang, and Y. Cao, "CASHIER: A Cache Energy Saving Technique for QoS Systems," *26th International Conference on VLSI Design(VLSID)*, pp. 43–48, 2013.
- [17] W.-c. Feng and K. W. Cameron, "The green500 list: Encouraging sustainable supercomputing," *Computer*, vol. 40, no. 12, pp. 50–55, 2007.
- [18] www.top500.org.
- [19] S. Greenberg, E. Mills, B. Tschudi, P. Rumsey, and B. Myatt, "Best practices for data centers: Lessons learned from benchmarking 22 data centers," *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE*, vol. 3, pp. 76–87, 2006.
- [20] S. Mittal, "A survey of techniques for improving energy efficiency in embedded computing systems," *International Journal of Computer Aided Engineering and Technology (IJCAET)*, 2014.
- [21] C. D. Patel, C. E. Bash, R. Sharma, M. Beitelmal, and R. Friedrich, "Smart cooling of data centers," *Pacific RIM/ASME International Electronics Packaging Technical Conference and Exhibition (IPACK03)*, 2003.
- [22] L. A. Barroso and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 4, no. 1, pp. 1–108, 2009.
- [23] P. Ranganathan, P. Leech, D. Irwin, and J. Chase, "Ensemble-level power management for dense blade servers," in *ACM SIGARCH Computer Architecture News*, vol. 34, no. 2, 2006, pp. 66–77.
- [24] L. A. Barroso and U. Hölzle, "The case for energy-proportional computing," *IEEE computer*, vol. 40, no. 12, pp. 33–37, 2007.
- [25] S. Mittal and Z. Zhang, "EnCache: Improving cache energy efficiency using a software-controlled profiling cache," in *IEEE International Conference On Electro/Information Technology*, Indianapolis, USA, May 2012.
- [26] L. A. Barroso, J. Dean, and U. Holzle, "Web search for a planet: The Google cluster architecture," *Micro*, vol. 23, no. 2, pp. 22–28, 2003.

- [27] R. Bianchini and R. Rajamony, "Power and energy management for server systems," *IEEE Computer*, vol. 37, no. 11, pp. 68–74, 2004.
- [28] D. Anderson, J. Dykes, and E. Riedel, "More than an interface-scsi vs. ata." in *FAST*, vol. 2, no. 00, 2003, p. 3.
- [29] T. Daim, J. Justice, M. Krampits, M. Letts, G. Subramanian, and M. Thirumalai, "Data center metrics: An energy efficiency model for information technology managers," *Management of Environmental Quality: An International Journal*, vol. 20, no. 6, pp. 712–731, 2009.
- [30] "Energy star program," www.energystar.gov/.
- [31] W.-C. Feng, "Making a case for efficient supercomputing," *Queue*, vol. 1, no. 7, p. 54, 2003.
- [32] M. Mazzucco and I. Mitrani, "Empirical evaluation of power saving policies for data centers," *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, no. 3, pp. 18–22, 2012.
- [33] L. H. Seigo, A. Márquez, A. Rawson, T. Cader, K. Fox, W. I. Gustafson Jr, and C. J. Mundy, "Implementing the data center energy productivity metric," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 8, no. 4, p. 30, 2012.
- [34] P. Bohrer, E. N. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, C. McDowell, and R. Rajamony, "The case for power management in web servers," in *Power aware computing*. Springer, 2002, pp. 261–289.
- [35] Q. Deng, D. Meisner, A. Bhattacharjee, T. F. Wenisch, and R. Bianchini, "Multiscale: memory system dvfs with multiple memory controllers," in *international symposium on Low power electronics and design*, 2012, pp. 297–302.
- [36] 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," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 18, no. 6, pp. 835–848, 2007.
- [37] A. Gandhi, M. Harchol-Balter, R. Das, and C. Lefurgy, "Optimal power allocation in server farms," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 37, no. 1, 2009, pp. 157–168.
- [38] T. Horvath, T. Abdelzaher, K. Skadron, and X. Liu, "Dynamic voltage scaling in multitier web servers with end-to-end delay control," *Computers, IEEE Transactions on*, vol. 56, no. 4, pp. 444–458, 2007.
- [39] 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*. IEEE Computer Society, 2005, p. 1.
- [40] L. Rao, X. Liu, L. Xie, and W. Liu, "Minimizing electricity cost: optimization of distributed internet data centers in a multi-electricity-market environment," in *IEEE INFOCOM*, 2010, pp. 1–9.
- [41] V. Sharma, A. Thomas, T. Abdelzaher, K. Skadron, and Z. Lu, "Power-aware QoS management in web servers," in *24th IEEE Real-Time Systems Symposium (RTSS)*, 2003, pp. 63–72.
- [42] B. Shi and A. Srivastava, "Unified datacenter power management considering on-chip and air temperature constraints," *Sustainable Computing: Informatics and Systems*, vol. 1, no. 2, pp. 91–98, 2011.

- [43] R. Xu, D. Zhu, C. Rusu, R. Melhem, and D. Mossé, “Energy-efficient policies for embedded clusters,” in *ACM SIGPLAN Notices*, vol. 40, no. 7, 2005, pp. 1–10.
- [44] V. Anagnostopoulou, S. Biswas, A. Savage, R. Bianchini, T. Yang, and F. T. Chong, “Energy conservation in datacenters through cluster memory management and barely-alive memory servers,” in *Proceedings of the 2009 Workshop on Energy Efficient Design*, 2009.
- [45] P. Bohrer, D. Cohn, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, R. Rajamony, F. Rawson, and E. Hensbergen, “Energy conservation for servers,” in *IEEE Workshop on Power Management for Real-time and Embedded Systems*, 2001.
- [46] J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, and R. P. Doyle, “Managing energy and server resources in hosting centers,” in *ACM SIGOPS Operating Systems Review*, vol. 35, no. 5. ACM, 2001, pp. 103–116.
- [47] H. Chen, M. Kesavan, K. Schwan, A. Gavrilovska, P. Kumar, and Y. Joshi, “Spatially-aware optimization of energy consumption in consolidated data center systems,” in *ASME 2011 Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems*. American Society of Mechanical Engineers, 2011, pp. 461–470.
- [48] G. Da Costa, M. D. De Assuncao, J.-P. Gelas, Y. Georgiou, L. Lefèvre, A.-C. Orgerie, J.-M. Pierson, O. Richard, and A. Sayah, “Multi-facet approach to reduce energy consumption in clouds and grids: the green-net framework,” in *Proceedings of the 1st international conference on energy-efficient computing and networking*. ACM, 2010, pp. 95–104.
- [49] S. Ghosh, M. Redekopp, and M. Annavaram, “Knightshift: shifting the i/o burden in datacenters to management processor for energy efficiency,” in *Computer Architecture*. Springer, 2012, pp. 183–197.
- [50] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, “GreenCloud: a new architecture for green data center,” in *6th international conference industry session on Autonomic computing and communications industry session*. ACM, 2009, pp. 29–38.
- [51] E. Pinheiro, R. Bianchini, E. V. Carrera, and T. Heath, “Dynamic cluster reconfiguration for power and performance,” in *Compilers and operating systems for low power*. Springer, 2003, pp. 75–93.
- [52] K. Rajamani and C. Lefurgy, “On evaluating request-distribution schemes for saving energy in server clusters,” in *Performance Analysis of Systems and Software, 2003. ISPASS. 2003 IEEE International Symposium on*. IEEE, 2003, pp. 111–122.
- [53] A. Verma, G. Dasgupta, T. K. Nayak, P. De, and R. Kothari, “Server workload analysis for power minimization using consolidation,” in *Proceedings of the 2009 conference on USENIX Annual technical conference*. USENIX Association, 2009, pp. 28–28.
- [54] W. Zheng, R. Bianchini, G. J. Janakiraman, J. R. Santos, and Y. Turner, “JustRunIt: Experiment-based management of virtualized data centers,” in *Proc. USENIX Annual Technical Conference*, 2009.
- [55] J. Leverich, M. Monchiero, V. Talwar, P. Ranganathan, and C. Kozyrakis, “Power management of datacenter workloads using per-core power gating,” *Computer Architecture Letters*, vol. 8, no. 2,

pp. 48–51, 2009.

- [56] A. Banerjee, T. Mukherjee, G. Varsamopoulos, and S. K. Gupta, “Integrating cooling awareness with thermal aware workload placement for hpc data centers,” *Sustainable Computing: Informatics and Systems*, vol. 1, no. 2, pp. 134–150, 2011.
- [57] D. J. Bradley, R. E. Harper, and S. W. Hunter, “Workload-based power management for parallel computer systems,” *IBM Journal of Research and Development*, vol. 47, no. 5.6, pp. 703–718, 2003.
- [58] K. Le, R. Bianchini, M. Martonosi, and T. D. Nguyen, “Cost-and energy-aware load distribution across data centers,” *Proceedings of HotPower*, pp. 1–5, 2009.
- [59] J. D. Moore, J. S. Chase, P. Ranganathan, and R. K. Sharma, “Making scheduling cool: Temperature-aware workload placement in data centers.” in *USENIX annual technical conference, General Track*, 2005, pp. 61–75.
- [60] R. Nathuji, C. Isci, and E. Gorbato, “Exploiting platform heterogeneity for power efficient data centers,” in *Fourth International Conference on Autonomic Computing (ICAC)*. IEEE, 2007, pp. 5–5.
- [61] L. Wang, G. Von Laszewski, J. Dayal, X. He, A. J. Younge, and T. R. Furlani, “Towards thermal aware workload scheduling in a data center,” in *10th International Symposium on Pervasive Systems, Algorithms, and Networks (ISPAN)*. IEEE, 2009, pp. 116–122.
- [62] M. Zapater, J. L. Ayala, and J. M. Moya, “Leveraging heterogeneity for energy minimization in data centers,” in *Cluster, Cloud and Grid Computing (CCGrid), 2012 12th IEEE/ACM International Symposium on*, 2012, pp. 752–757.
- [63] Q. Tang, S. K. Gupta, D. Stanzione, and P. Cayton, “Thermal-aware task scheduling to minimize energy usage of blade server based datacenters,” in *Dependable, Autonomic and Secure Computing, 2nd IEEE International Symposium on*. IEEE, 2006, pp. 195–202.
- [64] C. E. Bash, C. D. Patel, and R. K. Sharma, “Dynamic thermal management of air cooled data centers,” in *Thermal and Thermomechanical Phenomena in Electronics Systems, 2006. IThERM’06. The Tenth Intersociety Conference on*. IEEE, 2006, pp. 8–pp.
- [65] W. Huang, M. Allen-Ware, J. B. Carter, E. Elnozahy, H. Hamann, T. Keller, C. Lefurgy, J. Li, K. Rajamani, and J. Rubio, “TAPO: Thermal-aware power optimization techniques for servers and data centers,” in *Green Computing Conference and Workshops (IGCC), 2011 International*. IEEE, 2011, pp. 1–8.
- [66] D. Kudithipudi, Q. Qu, and A. K. Coskun, “Thermal management in many core systems,” in *Evolutionary Based Solutions for Green Computing*. Springer, 2013, pp. 161–185.
- [67] A. Mirtar, S. Dey, and A. Raghunathan, “Adaptation of video encoding to address dynamic thermal management effects,” in *Green Computing Conference (IGCC)*. IEEE, 2012, pp. 1–10.
- [68] J. Moore, R. Sharma, R. Shih, J. Chase, C. Patel, P. Ranganathan *et al.*, “Going beyond CPUs: The potential of temperature-aware solutions for the data center,” in *Proc. 2004 First Workshop Temperature-Aware Computer Systems (TACS-1) Held in Conjunction with ISCA-31*, 2004.
- [69] E. Pakbaznia, M. Ghasemazar, and M. Pedram, “Temperature-aware dynamic resource provisioning

- in a power-optimized datacenter,” in *Proceedings of the Conference on Design, Automation and Test in Europe*. European Design and Automation Association, 2010, pp. 124–129.
- [70] L. Ramos and R. Bianchini, “C-Oracle: Predictive thermal management for data centers,” in *IEEE 14th International Symposium on High Performance Computer Architecture (HPCA)*, 2008, pp. 111–122.
- [71] J. Kim, M. Ruggiero, and D. Atienza, “Free cooling-aware dynamic power management for green datacenters,” in *High Performance Computing and Simulation (HPCS), 2012 International Conference on*. IEEE, 2012, pp. 140–146.
- [72] W. Felter, K. Rajamani, T. Keller, and C. Rusu, “A performance-conserving approach for reducing peak power consumption in server systems,” in *19th annual international conference on Supercomputing*, 2005, pp. 293–302.
- [73] J. Choi, S. Govindan, B. Urgaonkar, and A. Sivasubramaniam, “Profiling, prediction, and capping of power consumption in consolidated environments,” in *Modeling, Analysis and Simulation of Computers and Telecommunication Systems, 2008. MASCOTS 2008. IEEE International Symposium on*. IEEE, 2008, pp. 1–10.
- [74] C. Lefurgy, X. Wang, and M. Ware, “Server-level power control,” in *Autonomic Computing, 2007. ICAC’07. Fourth International Conference on*. IEEE, 2007, pp. 4–4.
- [75] J. O. Kephart, H. Chan, R. Das, D. W. Levine, G. Tesauro, F. L. Rawson III, and C. Lefurgy, “Coordinating multiple autonomic managers to achieve specified power-performance tradeoffs.” in *ICAC*, vol. 7, 2007, p. 24.
- [76] D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang, “Power and performance management of virtualized computing environments via lookahead control,” *Cluster computing*, vol. 12, no. 1, pp. 1–15, 2009.
- [77] C. Rusu, A. Ferreira, C. Scordino, and A. Watson, “Energy-efficient real-time heterogeneous server clusters,” in *Proceedings of the 12th IEEE Real-Time and Embedded Technology and Applications Symposium*, 2006, pp. 418–428.
- [78] K. Sudan, S. Srinivasan, R. Balasubramonian, and R. Iyer, “Optimizing datacenter power with memory system levers for guaranteed quality-of-service,” in *Proceedings of the 21st international conference on Parallel architectures and compilation techniques*. ACM, 2012, pp. 117–126.
- [79] X. Wang and Y. Wang, “Co-con: Coordinated control of power and application performance for virtualized server clusters,” in *Quality of Service, 2009. IWQoS. 17th International Workshop on*. IEEE, 2009, pp. 1–9.
- [80] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wang, and N. Gautam, “Managing server energy and operational costs in hosting centers,” in *ACM SIGMETRICS Performance Evaluation Review*, vol. 33, no. 1, 2005, pp. 303–314.
- [81] T. Horvath and K. Skadron, “Multi-mode energy management for multi-tier server clusters,” in *17th international conference on Parallel architectures and compilation techniques*. ACM, 2008, pp. 270–279.

- [82] L. Parolini, N. Tolia, B. Sinopoli, and B. H. Krogh, "A cyber-physical systems approach to energy management in data centers," in *Proceedings of the 1st ACM/IEEE International Conference on Cyber-Physical Systems*. ACM, 2010, pp. 168–177.
- [83] R. Urgaonkar, U. C. Kozat, K. Igarashi, and M. J. Neely, "Dynamic resource allocation and power management in virtualized data centers," in *Network Operations and Management Symposium (NOMS)*. IEEE, 2010, pp. 479–486.
- [84] X. Wang and M. Chen, "Cluster-level feedback power control for performance optimization," in *IEEE 14th International Symposium on High Performance Computer Architecture (HPCA)*, 2008, pp. 101–110.
- [85] Í. Goiri, R. Beauchea, K. Le, T. D. Nguyen, M. E. Haque, J. Guitart, J. Torres, and R. Bianchini, "Greenslot: scheduling energy consumption in green datacenters," in *International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 2011, p. 20.
- [86] K.-K. Nguyen, M. Cheriet, M. Lemay, M. Savoie, and B. Ho, "Powering a data center network via renewable energy: A green testbed," *Internet Computing, IEEE*, vol. 17, no. 1, pp. 40–49, 2013.
- [87] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, "DRPM: dynamic speed control for power management in server class disks," in *30th Annual International Symposium on Computer Architecture*. IEEE, 2003, pp. 169–179.
- [88] D. Narayanan, A. Donnelly, and A. Rowstron, "Write off-loading: Practical power management for enterprise storage," *ACM Transactions on Storage (TOS)*, vol. 4, no. 3, p. 10, 2008.
- [89] E. Pinheiro, R. Bianchini, and C. Dubnicki, "Exploiting redundancy to conserve energy in storage systems," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 34, no. 1, 2006, pp. 15–26.
- [90] N. Chatterjee, M. Shevgoor, R. Balasubramonian, A. Davis, Z. Fang, R. Illikkal, and R. Iyer, "Leveraging heterogeneity in dram main memories to accelerate critical word access," in *Microarchitecture (MICRO), 2012 45th Annual IEEE/ACM International Symposium on*. IEEE, 2012, pp. 13–24.
- [91] Q. Deng, D. Meisner, L. Ramos, T. F. Wenisch, and R. Bianchini, "Memscale: active low-power modes for main memory," *ACM SIGPLAN Notices*, vol. 46, no. 3, pp. 225–238, 2011.
- [92] D. H. Yoon, J. Chang, N. Muralimanohar, and P. Ranganathan, "BOOM: enabling mobile memory based low-power server DIMMs," in *ACM SIGARCH Computer Architecture News*, vol. 40, no. 3. IEEE Computer Society, 2012, pp. 25–36.
- [93] J. Liu, B. Jaiyen, R. Veras, and O. Mutlu, "RAIDR: Retention-aware intelligent DRAM refresh," in *Computer Architecture (ISCA), 2012 39th Annual International Symposium on*. IEEE, 2012, pp. 1–12.
- [94] S. Mittal, "A survey of architectural techniques for improving cache power efficiency," *Sustainable Computing: Informatics and Systems*, 2013.
- [95] B.-G. Chun, G. Iannaccone, G. Iannaccone, R. Katz, G. Lee, and L. Niccolini, "An energy case for hybrid datacenters," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 1, pp. 76–80, 2010.
- [96] R. Nathuji and K. Schwan, "VirtualPower: coordinated power management in virtualized enterprise systems," in *ACM SIGOPS Operating Systems Review*, vol. 41, no. 6, 2007, pp. 265–278.

- [97] L. Li, C.-J. M. Liang, J. Liu, S. Nath, A. Terzis, and C. Faloutsos, “Thermocast: a cyber-physical forecasting model for datacenters,” in *17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2011, pp. 1370–1378.
- [98] S. Gupta, S. Mittal, and S. Dasgupta, “Guaranteed QoS with MIMO Systems for Scalable Low Motion Video Streaming Over Scarce Resource Wireless Channels,” in *Proceedings Second International Conference On Information Processing*, 2008, pp. 452–466.
- [99] S. Gupta, S. Mittal, S. Dasgupta, and A. Mittal, “MIMO Systems For Ensuring Multimedia QoS Over Scarce Resource Wireless Networks,” in *ACM International Conference On Advance Computing, India*, 2008.
- [100] S. Mittal *et al.*, “Design exploration and implementation of simplex algorithm over reconfigurable computing platforms,” in *IEEE International Conference on Digital Convergence*, 2011, pp. 204–209.
- [101] A. Pande and S. Mittal, “BayWave: BAYesian WAVElet-based image estimation,” *International Journal of Signal and Imaging Systems Engineering*, vol. 2, no. 4, pp. 155–162, 2009.
- [102] L. Keys, S. Rivoire, and J. D. Davis, “The search for energy-efficient building blocks for the data center,” in *Computer Architecture*. Springer, 2012, pp. 172–182.
- [103] C. Delimitrou and C. Kozyrakis, “Paragon: QoS-aware scheduling for heterogeneous datacenters,” in *ASPLOS*. ACM, 2013, pp. 77–88.
- [104] S. Mittal, *Architectural Techniques For Managing Non-volatile Caches*. Lambert Academic Publishing (LAP), 2013.