

## Towards Dynamic Voltage/Frequency Scaling for Power Reduction in Data Centers

Congfeng Jiang<sup>1, a</sup>, Yinghui Zhao<sup>2, b</sup> and Jian Wan<sup>1, c</sup>

<sup>1</sup>Grid and Service Computing Technology Lab, Hangzhou Dianzi University, Hangzhou 310037, China

<sup>2</sup> Department of Hydraulic Engineering and Agriculture, Zhejiang Tongji Vocational College of Science and Technology, Hangzhou, 311231, China

<sup>a</sup>cjiang@hdu.edu.cn, <sup>b</sup>wanjian@hdu.edu.cn, <sup>c</sup>zhaoyinghuihust@gmail.com

**Keywords:** Power consumption, Dynamic voltage scaling, Dynamic frequency scaling, virtualization, Data centers

**Abstract.** Higher power consumption in data centers results in more heat dissipation, cooling costs and degrades the system reliability. Conventional power reduction techniques such as dynamic voltage/frequency scaling (DVS/DFS) have disadvantages when they are ported to current data centers with virtualization deployments. In this paper, we give a short survey and discussion on some issues and aspects of DVS/DFS in data centers. This paper also presents a simple comparison of four power management schemes in virtualization environments.

### Introduction

The average power requirement per rack continues to grow dramatically, along with the density of the equipment packed into each rack with high-density blade server implementations in typical DCs. Moreover, power densities in various data center types have grown remarkably in recent years [1]. Not only VLSI and CMOS designers, computer manufacturers, but also governments and environmental agencies have recognized and begun to demand new initiatives and solutions for the power consumption problems [2,3]. As for power consumptions and cooling costs, a notable example of this issue which gained much public attention is the ever-increasing of energy costs on Google's server centers [4]. Moreover, power reduction is particularly valuable for data centers with thermal emergencies, cooling constraints, and power-supply constraints. For example, large scale data centers using single power supply source are more power-safe than those with multiple power supplies, where a partial failure of the power supply subsystem can result in performance loss and Quality of Service (QoS) degradation in order to meet a lower power constraint or power budget [5,6]. Consequently, power consumption and induced heat dissipation have become the most important design consideration for data centers.

Since data centers experience large periods of low utilization, they present opportunities for power reduction with minimal performance losses. Virtualization provides flexible construction of numerous virtual machines with almost no hardware limitations, and consequently reduces the total cost of ownership. It also enables data center users to migrate data transparently and seamlessly across multiple server sites or applications without disrupting the transaction state and affecting the application performance and QoS. In summary, virtualization provides transparency, isolation, and security, which makes it a promising approach for service consolidation in DCs, e.g., to reuse and to better exhaust available hardware resources. However, in a virtualization environment, various virtual machines usually require different operating frequencies and voltages. A global cooperative DVS/DFS mechanism for Virtual Machine Monitor (VMM) or Hypervisor to coordinate between numerous Virtual Machines (VMs) or Guest OSes to agree on one single operating frequency or voltage is not available, and usually impossible. Therefore, moving the power consumption control mechanism from hardware-level to fine-grained software-level is indispensable and more promising

for data centers with virtual machine environments, and has the potential to save more power consumption.

The remainder of this paper is organized as follows: In section 2 we formulate the power and power consumption model of processors and review some pre-requirements for DVS/DFS techniques, including power consumption estimation, workload characterization, and OS support power-aware management. Section 3 presents an overview of classic DVS/DFS techniques. We outline the power management challenges in data centers in section 4, including server virtualization and QoS constraints. Finally, we summarize this paper and make some remarks on further research topics in Section 5.

### Formulations and Basic Concepts of Power Consumption

Power is the rate at which the target computing system consumes electrical energy (or dissipates it in the form of heat) while performing application activities, and energy is the total electrical energy the system consumes (or dissipates as heat) over time. The definition is listed below:

$$P = \frac{W}{T} \quad (1)$$

$$E = P \times T \quad (2)$$

Where P is power, E is energy, T is a specific time interval, and W is the total work performed in the corresponding time interval [7]. Usually, energy is measured in *joules*, while power is measured in *watts*. This distinction between power and energy is very important because techniques that reduce power do not reduce energy accordingly. For example, the power consumed by a computer can be reduced by down-scaling the clock frequency, but if the computer then takes much longer time to run the same programs, the total energy consumed will be similar, sometimes even more.

The following equation shows the dynamic power consumption of processors, which depends on supply voltage (V), clock frequency (f), and an activity factor ( $\alpha$ ):

$$P_{dynamic} \propto \alpha \cdot V^2 \cdot f \quad (3)$$

Eq.3 suggests that there are three ways to reduce dynamic power consumption: to reduce the clock frequency, to reduce the supply voltage, and to reduce the factor  $\alpha$ . Bansal et al [8] analyzes the relationship between the processor speed and its power consumption with a more accurate algorithm, named Arbitrary Power Function.

Although the above three ways can ideally reduce dynamic power dissipation cubically, each of them has different tradeoffs, disadvantages, and not all of them reduce the total energy consumed accordingly. They also have limitations and cannot always be applied. For example, reducing the clock frequency or supply voltage, i.e., lowering fixed frequencies and voltages, will worsen the target system performance, Quality of Service, and does not always reduce the total energy consumed. Moreover, reducing the supply voltage increases gate delays, additional overheads and usually it also requires reducing the clock frequency to allow the circuit to work properly. This is not feasible for heterogeneous multi-core architectures and DCs.

**Power-Performance Metrics and Tradeoffs.** In order to reduce power consumption, the first step is to measure power consumption before DVS/DFS implementation. On the other hand, the power-efficiency evaluation of implemented DVS/DFS algorithm is also based on the practical and real-time power consumption measurement. Accurate power consumption characteristics and measurement is vital to help activate DVS/DFS mechanism to make right decisions according with preset selection criteria. Although some computing system vendors report some power-performance figures, but they are often not directly comparable due to differences in workload, configuration, test environment, etc. The SPECpower\_ssjs2008[9] only addresses one subset of server workloads: the performance of server side Java.

Power Usage Effectiveness (PUE) metric is often used to evaluate the efficiency of data centers [10]. PUE is the ratio of the total power consumed by a data center to the power consumed by the IT equipment that operates the facility:

$$PUE = \frac{TotalFacilityPower}{ITEquipmentPower} \quad (4)$$

PUE puts emphasis on maximizing the power supplied to the equipment running customer applications and minimizing the power consumed by cooling and power distribution. For example, a PUE of 2.0 indicates that for every watt of IT power, an additional watt is consumed to cool and distribute power to the IT equipment.

**Workload Characterization.** DVS/DFS can dynamically allocate computing resources to meet the workload requirements more effectively and tolerate acceptable performance loss or degradation. However, the main weakness resided in DVS/DFS or adaptive power management (APM) is the highly expensive workload characterization problem. Moreover, increasing demand for power-efficient, high-performance computing requires tuning applications and/or the underlying hardware to improve the mapping between workload heterogeneity and computational resources. To support power management, approaches to profile the power usage of applications play a main role in power management and software design decisions [11, 12].

Power-driven statistical sampling and thread-level power estimation are main techniques to estimate power consumption, measurement and monitoring in computing systems. In practical implementation, power estimation techniques for power management should not add to additional overheads or impede prediction accuracy and runtime performance under various workloads. Huang [13] proposed a novel approach to reduce power consumption while better maintaining high performance requirements via a novel behavioral quantification of workload. In order to automate the collection and analysis of workload traces and performance data from data centers, methods and toolkits for repeatable and verifiable experiments and workload emulation is needed. Moreover, the toolkits or software packages should provide functions to allow detailed and flexible emulation of enterprise-class workloads.

In summary, the feasibility of energy management will be decided by assessing benefits and costs, i.e., power performance tradeoff. However, how to find a system-wide balancing point between potential power reduction benefits and performance loss in DCs has long been difficult for designers and researchers due to the challenging dimensionality of the analysis.

### Existing Dynamic Voltage/Frequency Scaling Algorithms

Power consumption is always one of main concerns to the designers of processors. Circuit-level techniques have been proven effective for managing power consumption and heat dissipations. As transistor density increases, designers have to implement adaptive structures to provide the promised energy efficiency. Except for the circuit-level and architectural-level adaption, various excellent task-specific dynamic voltage/frequency scaling (DVS/DFS) algorithms have been proposed for power reduction and system reliabilities. Within DVS/DFS schemes, when system workloads are under average level, considerable power consumption can be reduced through CPU voltage or frequency scaling, e.g. slowing down CPUs or switching CPUs to low-power modes.

As processors consume a large portion of power in computer systems, it is instinctive to reduce power consumptions of processors. AbouGhazaleh et al [14] proposed an integrated DVS approach to synergistically manage the energy of chip components in different clock domains. Lin et al [15] proposed User-Driven Frequency Scaling (UDFS) and Process-Driven Voltage Scaling (PDVS) that can be readily employed independently or together for energy management. Isci et al [16] proposed a system for predicting phases of applications at runtime using performance counters. They use a global phase history table predictor leveraged from a common branch predictor technique. These

runtime phase predictions are used to guide DVS/DFS as the underlying dynamic management technique.

DVS/DFS techniques are effective at reducing average power consumptions, however, these techniques have the undesired side-effect of increasing both the variability of power dissipation and the variability of current drawn by the processor. This in turn increase in current variability, often referred to as the  $dI/dt$  problem, can cause supply voltage fluctuations. Joseph et al [17] proposed a methodology for augmenting packaging techniques for  $dI/dt$  with micro-architecture control mechanisms using resonant frequencies most relevant to current microprocessor packages.

Most processors designed today include power management features that provide processor operating points which can be used in energy management algorithms. However, existing power management algorithms implicitly assume that lower performance points are more energy efficient than higher performance points. Miyoshi et al [18] observed that for many systems, this assumption is not valid. They introduced a new concept called critical power slope to explain and capture the power-performance characteristics of systems with power management features. However, DVS/DFS techniques may not provide enough granularities and can hurt runtime performance. Actually, because of recent hardware advances in platform and CPU architectures such as sophisticated memory subsystem design, DVS/DFS gives much more limited energy savings with relatively high performance loss as compared to running workloads at high speed and then switching into low power mode in modern computing systems[19].

To reduce the considerable core-to-core variations in power and performance caused by manufacturing process variations, the recent trend towards chip-multiprocessors (CMP) executing multi-threaded workloads with heterogeneous behavior motivates the need for per-core DVS/DFS control mechanisms. Fine-grained dynamic voltage/frequency scaling is an important tool in managing the balance between power and performance in chip-multiprocessors [20, 21]. Kim et al [22] explored the potential system-wide energy savings offered by implementing both fine-grained and per-core DVS/DFS in a 4-core CMP system.

Although process migration techniques can provide some relief, currently there are no solutions for multi-core or CMP architectures that suggest process migration schedules to improve thermals when all cores or threads are active. Therefore, DVS/DFS techniques should be a last resort to protect the processor from energy or thermal emergencies. Moreover, DVS/DFS techniques do not gracefully handle scenarios where processors are fully loaded or active, i.e. there are no free threads or cores for process scheduling. Since DVS/DFS operates on the core scale and it is very effective in decreasing temperature, it may not provide satisfied granularity for thermal management since most current systems do not support independent frequency scaling of cores or threads.

### **Power Management Challenges in Data Centers**

In data centers various applications usually run on shared heterogeneous hardware, including virtual machines, virtual local area networks, and virtual storage. These challenges have motivated recent work in frameworks and implementations for coordinated monitoring and control of power consumption in large scale data centers.

Parolini et al [23] proposed a coordinated approach to data center energy management and the simulation results show the potential for a coordinated control strategy to achieve better energy management than traditional schemes that control the computational and cooling subsystems separately. Currently, control theory is often utilized for power consumption monitoring and control. A classic closed-loop feedback control flowchart used in DCs includes measurement, monitor, analyzer, actuator, etc. The input and output is the real-time power consumption data. The measurement block also monitors and measure server dynamic activity information in DCs, including power consumption, workload status, network utilization, service status, etc. All the

DVS/DFS mechanism is realized by a system administrator or management software through electrical or mechanical actuator.

Control-theoretic techniques have recently shown their potentials on power management due to their better control performance and theoretical guarantees on control precision and system robustness. However, in DCs many servers are correlated with each other by common workloads and share common power supplies. Conventional servers use simple open-loop policies to set a safe performance level in order to limit peak power consumption, despite workloads and runtime environments, usually controlling a single server independently from others. Lefurgy et al [24] propose a technique which controls the peak power consumption of a high-density server through a feedback controller that uses precise, system-level power measurement to periodically select the highest performance state while keeping the system within a fixed power constraint. Wang et al [25], the authors use an optimal multi-input-multi-output (MIMO) control theory in power control and propose a cluster-level power controller which regulates power among servers according to their performance requirements, while controlling the total power of the cluster to be lower than a constraint. However, it maybe unfair to use DVS/DFS mechanism to reduce power consumptions in typical situations because it may unfairly penalize cooler running processes by reducing their runtime performance.

**Virtualization.** Server virtualization consolidates multiple under-utilized servers into a single physical server, thereby reducing the physical and environmental space and power supplies required to host them. Therefore, virtualization can reduce power consumptions significantly. In DCs with virtualization deployment, simultaneous on-demand provisioning and accesses of shared physical infrastructures to virtual machines (VMs) makes it harder for the capacities management and to fulfill the Service Level Agreements (SLAs) with least cost. Therefore, managing server resources in data centers where content providers/website owners share the hardware platform is a challenge for system designer and operators. Moreover, there also exist some disadvantages associated with individual DVS/DFS algorithms when they are ported to DCs, where virtualization technologies and emerging multi-core processors are widely deployed, such as additional overheads during mode selection and transitions, latencies, performance degradations, conflicts and coordination, etc. Virtualization is more and more deployed in modern data centers, which cause problematic issues for DVS/DFS because it is hard and impossible for Virtual Machine Monitor(VMM) or Hypervisor to coordinate between numerous Virtual Machines (VMs) or Guest OSes to agree with a single voltage or frequency. It is same with the emerging multi-core or chip multi processor (CMP) architectures.

Table 1 shows a simple comparison of four power management schemes in virtualization environments. We observe from Table 1 that DVS/DFS schemes are not utilized in most of the existing solutions for power management in virtualization environments due to the complexity of multiple scaling modes coordination among various VMs. In contrast, they use VM migration for consolidation and to save more power consumptions since it is much easier to implement VM migration for the underlying VMM or hypervisor. However, this does not mean that VMs can randomly be migrated among all nodes. Actually, the potential overheads caused by live migrations of VMs can not be ignored in many cases because they may have severe negative impacts on global system performance including QoS and SLA satisfactions.

Moreover, in DCs with legacy hardware and software applications, it is impossible to implement global DVS/DFS schemes among all machines due to the heterogeneity. Therefore, more power reduction can be achieved if fine grained DVS/DFS schemes are used elegantly in virtualization environments.

**Service Level Agreement Guarantees.** The Service Level Agreement (SLA) defines service quality like online time, response time, failure percentage, etc. From a general-purpose viewpoint, performance can be defined by QoS constraints for corresponding underlying workload heterogeneity [31]. Due to the high-density of service consolidation and development of new applications and the increasing number of users with heterogeneous requests, providing users with

QoS guarantees while executing applications and saving power consumption has become a crucial problem that needs to be addressed. This problem is referred as the QoS-constrained DVS/DFS problem and it is typically NP-hard.

Table 1. Comparison of power management schemes in virtualization environments

Evaluation items	VirtualPower[26]	Magnet[27]	ClientVisor[28]	Ref.[29]	VPM Tokens[30]
Testbed	Multiple PCs	A 64-hosts cluster	Desktop virtualization environment	PC	Cluster
Hardware Heterogeneity	Homogeneous +Heterogeneous	Homogeneous	Homogeneous	Homogeneous	homogeneous +heterogeneous
CPU Vendor	Intel	AMD	Intel	Intel	Intel
CPU Type	Dual Core Pentium 4	Athlon 3500+	Core2 Duo T9400	Pentium D	Pentium 4/ Dual Core
Using DVS/DFS	Yes	N/A	N/A	N/A	Yes
VMM	Xen	Xen	Xen	L4 micro-kernel	Xen
Number of VMs	$\geq 4$	N/A	3	N/A	$\geq 6$
Online/Offline	online	online	online	online	online
Max. Power savings	34%	74.8%	22%	N/A	Fixed power budgets
Performance degradations	Little performance penalties	Adjustably acceptable	2%~3%.	N/A	26%
With QoS/SLA guarantees	Yes	Yes	N/A	N/A	Yes
VM migration	Yes	Yes	N/A	N/A	Yes
Workload	RUBiS	bit-r, m-sort, m-m, t-sim, metis, r-sphere, and r-wing	SPECpower_ssj	DAQ/bzip2 application	SPEC CPU2000 for Batch/Transactional Workload and Nutch for Web Service Workload

N/A: stands for not mentioned in the corresponding paper.

In order to satisfy the QoS requirements of user applications, the status of the DCs must be monitored and the performance data should be recorded. However, in a large DC, the collection of system performance data, the coordinated optimization of power consumption and QoS requirements will add on a large amount of computation and communication overhead. Thus, efficient system-wide and per-component monitoring and discovering technologies must be developed. In a real DC, asking the users to fully specify their QoS requirements quantitatively is an unreasonable burden. For example, user only need to specify a power consumption level such as low, middle, or high when executing jobs rather than the numerical values. Therefore, how to evaluate the qualitative and quantitative effects is a key factor that impacts the coordinated optimization of power consumption and QoS requirements heavily [32, 33].

Conventional servers respond to power supply constraint situations by using simple open-loop policies to set a safe performance level in order to limit peak power consumption. Due to the various sources of QoS requirements, energy reduction techniques should consider the coordinated optimization of power consumption and QoS requirements together. It is feasible to dynamically change device operation modes under different workloads using DVS/DFS schemes. However, current work mainly focuses on global energy consumption without attacking hotspots, which leads to bad system performance and QoS performances. In some cases the system performance, such as throughput or response time, can benefit significantly from global energy management through coordinated optimization of power consumption and QoS requirements, including hotspot eliminations. However, global energy consumption mode is different in different systems with specific performance-oriented applications. Moreover, energy consumption mode is also different for different DCs with different performance constraints and QoS requirements. It is an open problem to reduce energy consumption, while still meeting performance demands, system loads and reliability.

### Open Problems and Future Work

Power consumption has emerged as a critical issue in large scale DCs. In this paper, we discuss several aspects and characteristics of DVS/DFS algorithms, e.g. performance, overheads, feasibility and usability in Data Centers. This paper also presents the challenges of power reduction in large Data Centers. It can be summarized as follows:

(1) Hardware-level design is the fundamental solution to reduce power consumptions in DCs. However, due to excessive design complexity and high power requirements, it is feasible to save system energy from a system-wide and software-level standpoint. Moreover, in DCs with heterogeneous platforms and legacy hardware and software applications, it is impossible to implement global DVS/DFS schemes among all machines due to the workload heterogeneity. In these cases non-uniform power allocation and workload allocation can save more power consumption than DVS/DFS implementation only [34, 35].

(2) DVS/DFS scheme is a key technique to save energy when energy-efficiency is an important concern or performance metric. However, DVS/DFS techniques do not work better in all cases. For example, server workloads are likely to contain different characteristics in web servers and multimedia streaming applications, especially in DCs with virtualization deployment.

(3) In order to balance workload between various processors or cores, applications, tasks or processes will be migrated from one processor or core to another. However, migration overhead must be considered in both sensor-based or performance counter-based migration techniques. When more and more cores are packaged in one chip, load balancing and energy balancing is becoming more and more critical to provide performance assurances and system reliabilities.

Although the power savings depend on the specific application environment, the hardware, and the energy source's characteristics in large scale DCs, it's worthy of implementing a global and system-wide power management schemes to save power as much as possible. Coordinated optimization of power consumption and QoS requirements is still challenging and active for researchers and engineers. With more focuses shifted from micro architectural and hardware level to software-level and system-wide techniques, it is highly likely that in the future we can overcome this optimization problem and provide higher power-performance ratio for DCs.

### Acknowledgments

The funding support of this work by State Key Development Program of Basic Research of China (Grant No. 2007CB310900), Natural Science Fund of China (NSFC, Grant No. 60873023, 60973029), Science and Technology Research and Development Program of Zhejiang Province, China (Grant No. 2009C31033), and Hangzhou Dianzi University Startup Fund (Grant No. KYS 055608058) are greatly appreciated.

### References

- [1] Datacom Equipment Power Trends and Cooling Applications. American Society of Heating, Refrigerating and Air Conditioning Engineers(ASHRAE) (2005)
- [2] Report to Congress on Server and Data Center Energy Efficiency. U.S. Environmental Protection Agency(2007)
- [3] L. Ramos, R. Bianchini: C-Oracle: Predictive Thermal Management for Data Centers, Proceedings of 14th International Symposium on High-Performance Computer Architecture (HPCA-14) , IEEE Computer Society Press, Los Alamitos,CA (2008)
- [4] L. Barroso: The price of performance,ACM Queue, Vol.3(2005),p.48

- 
- [5] X. Wang, C. Lefurgy, and M. Ware: Managing Peak System-level Power with Feedback Control, IBM Research Technical Report(2005)
  - [6] K. Shen, A. Zhang, T. Kelly, C. Stewart: Operational analysis of processor speed scaling, Proceedings of the twentieth annual ACM Symposium on Parallel Algorithms and Architectures (SPAA08), ACM Press, New York (2008)
  - [7] V. Venkatachalam and M. Franz: Power Reduction Techniques for Microprocessor Systems, ACM Computing Surveys, Vol. 37(2005), p. 195
  - [8] N. Bansal, H. Chan, K. Pruhs: Speed scaling with an arbitrary power function, Proceedings of the Nineteenth Annual ACM -SIAM Symposium on Discrete Algorithms (SODA'09), Society for Industrial and Applied Mathematics (SIAM) Press, Philadelphia, PA (2009)
  - [9] [http://www.spec.org/power\\_ssj2008/](http://www.spec.org/power_ssj2008/)
  - [10] [http://www.thegreengrid.org/gg\\_content/TGG\\_Data\\_Center\\_Power\\_Efficiency\\_Metrics\\_PUE\\_and\\_DCiE.pdf](http://www.thegreengrid.org/gg_content/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf)
  - [11] M. Curtis-Maury, J. Dzierwa, C. D. Antonopoulos, and D. S. Nikolopoulos: Online Power-Performance Adaptation of Multithreaded Programs using Hardware Event-Based Prediction, Proceedings of the 20th Annual International Conference on Supercomputing (ICS'06) ACM Press, New York (2006)
  - [12] S. Gurumurthi, A. Sivasubramaniam, M.J. Irwin, N. Vijaykrishnan, M. Kandemir, T. Li, L.K. John: Using Complete Machine Simulation for Software Power Estimation: The SoftWatt Approach, Proceedings of the 8th International Symposium on High Performance Computer Architecture (HPCA-8), IEEE Computer Society Press, Los Alamitos, CA (2002)
  - [13] S. Huang and W. Feng: A Workload-Aware, Eco-Friendly Daemon for Cluster Computing, Proceedings of the 9th IEEE international Symposium on Cluster Computing and the Grid (CCGrid09), IEEE Computer Society Press, Los Alamitos, CA (2009).
  - [14] N. AbouGhazaleh, B. R. Childers, D. Mossé, R. G. Melhem: Integrated CPU Cache Power Management in Multiple Clock Domain Processors, Lecture Notes in Computer Science, Vol. 4917(2008), p. 209
  - [15] B. Lin, A. Malik, P. A. Dinda, G. Memik, and R. P. Dick: Power Reduction through Measurement and Modeling of Users and CPUs: Summary, Proceedings of ACM SIGMETRICS'07, ACM Press, New York (2007)
  - [16] C. Isci, G. Contreras, and M. Martonosi: Live, Runtime Phase Monitoring and Prediction on Real Systems with Application to Dynamic Power Management, Proceedings of 39th ACM/IEEE International Symposium on Micro-architecture (MICRO-39) IEEE Computer Society Press, Los Alamitos, CA (2006)
  - [17] R. Joseph, D. Brooks, and M. Martonosi: Control Techniques to Eliminate Voltage Emergencies in High Performance Processors, Proceedings of the 9th International Symposium on High-Performance Computer Architecture (HPCA-9), IEEE Computer Society Press, Los Alamitos, CA (2003)
  - [18] A. Miyoshi, C. Lefurgy, E. V. Hensbergen, R. Rajamony, and R. Rajkumar: Critical Power Slope: Understanding the Runtime Effects of Frequency Scaling, Proceedings of the 16th Annual ACM International Conference on Supercomputing (ICS'02), ACM Press, New York (2002)
  - [19] G. Dhiman, K. K. Pusukuri, T. Rosing: Analysis of Dynamic Voltage Scaling for System Level Energy Management, Proceedings of USENIX Workshop on Power Aware Computing and Systems (HotPower08), USENIX Association Press, Berkeley, CA (2008)
  - [20] S. Herbert, D. Marculescu: Variation-Aware Dynamic Voltage Frequency Scaling, Proceedings of the 15th International Symposium on High-Performance Computer Architecture (HPCA-15), IEEE Computer Society Press, Los Alamitos, CA (2009)



- 
- [21] R. Teodorescu, J. Torrellas, Variation-Aware Application Scheduling and Power Management for Chip Multiprocessors, Proceedings of the 35th International Symposium on Computer Architecture (ISCA2008), IEEE Computer Society Press, Los Alamitos, CA (2008)
  - [22] W. Kim, M. Gupta, Gu-Yeon Wei, David Brooks: System Level Analysis of Fast, Per-Core DVS/DFS using On-Chip Switching Regulators, Proceedings of 14th International Symposium on High-Performance Computer Architecture (HPCA-14), IEEE Computer Society Press, Los Alamitos, CA (2008)
  - [23] L. Parolini, B. Sinopoli, B. H. Krogh: Reducing Data Center Energy Consumption via Coordinated Cooling and Load Management, Proceedings of USENIX Workshop on Power Aware Computing and Systems (HotPower08), USENIX Association Press, Berkeley, CA (2008)
  - [24] C. Lefurgy, X. Wang, M. Ware: Server-Level Power Control, Proceedings of the Fourth International Conference on Autonomic Computing (ICAC07), IEEE Computer Society Press, Los Alamitos, CA (2007)
  - [25] X. Wang, M. Chen: Cluster-level feedback power control for performance optimization, Proceedings of IEEE 14th International Symposium on High Performance Computer Architecture (HPCA 2008), IEEE Computer Society Press, Los Alamitos, CA (2008)
  - [26] R. Nathuji, and K. Schwan, Virtualpower: Coordinated Power Management in Virtualized Enterprise Systems, ACM SIGOPS Operating Systems Review, Vol.41(2007), p.265
  - [27] L. Hu, H. Jin, X. Liao, X. Xiong, H. Liu: Magnet: A Novel Scheduling Policy for Power Reduction in Cluster with Virtual Machines. Proceedings of the 2008 IEEE International Conference on Cluster Computing (Cluster 2008), IEEE Computer Society Press, Los Alamitos, CA (2008)
  - [28] H. Chen, H. Jin, Z. Shao, K. Yu, K. Tian: ClientVisor: Leverage COTS OS Functionalities for Power Management in Virtualized Desktop Environment. Proceedings of the 2009 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE09), ACM Press, New York (2009)
  - [29] Stoess, J., Lang, C., Bellosa, F.: Energy Management for Hypervisor-Based Virtual Machines. Proceedings of the 2007 USENIX Annual Technical Conference (USENIX'07), USENIX Association Press, Berkeley, CA (2007)
  - [30] Ripal Nathuji, Karsten Schwan: VPM Tokens: Virtual Machine-Aware Power Budgeting in Datacenters. Proceedings of the 17th international symposium on High Performance Distributed Computing (HPDC2008), IEEE Computer Society Press, Los Alamitos, CA (2008)
  - [31] F. Guo, Y. Solihin, Li Zhao, and R. Iyer: A Framework for Providing Quality of Service in Chip Multi-Processors, Proceedings of the 40th Annual IEEE/ACM Symposium on Microarchitecture (MICRO-2007), IEEE Computer Society Press, Los Alamitos, CA (2007)
  - [32] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, X. Zhu: No Power Struggles: Coordinated Multi-level Power Management for the Data Center, Proceedings of 13th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XIII), ACM Press, New York (2008)
  - [33] R. Nathuji, A. Somani, K. Schwan, and Y. Joshi: CoolIT: Coordinating Facility and IT Management for Efficient Datacenters, Proceedings of USENIX Workshop on Power Aware Computing and Systems (HotPower08), USENIX Association Press, Berkeley, CA (2008)
  - [34] Mark E. Femal, Vincent W. Freeh: Boosting Data Center Performance Through Non-Uniform Power Allocation, Proceedings of the Second International Conference on Automatic Computing (ICAC'05), IEEE Computer Society Press, Los Alamitos, CA (2005)
  - [35] Ripal Nathuji, Canturk Isci, Eugene Gorbatoov, Exploiting Platform Heterogeneity for Power Efficient Data Centers, Proceedings of the Fourth International Conference on Automatic Computing (ICAC'07), IEEE Computer Society Press, Los Alamitos, CA (2007)