Accounting for Load Variation in Energy-Efficient Data Centers

Dzmitry Kliazovich[1], Sisay T. Arzo[2], Fabrizio Granelli[2], Pascal Bouvry[1], and Samee Ullah Khan[3]

[1] University of Luxembourg
6 rue Coudenhove Kalergi, Luxembourg
dzmitry.kliazovich@uni.lu,
pascal.bouvry@uni.lu

[2] University of Trento
Via Sommarive 14, Trento, Italy
sisay.arzo@studenti.unitn.it,
granelli@disi.unitn.it

[3] North Dakota State University
Fargo, ND 58108-6050
samee.khan@ndsu.edu

Abstract—**The energy consumption in data centers is drastically increasing and becoming a significant portion in the data center operating expenses. Enabling a sleep mode in the idle computing servers and network hardware is the most efficient method to avoid unnecessary power consumption. However, changes in the power modes introduce considerable delays. Moreover, inability to wake up a sleeping server immediately requires an availability of a pool of idle servers able to accommodate incoming load in the short term to prevent QoS degradation.**

In this paper we investigate the amount of computing servers and network hardware needed to accommodate different incoming load patters in the data centers. Furthermore, we propose to build these servers on energy efficient hardware, which is costly but can scale its power consumption with the offered load levels. The evaluation results show that the proposed methodology can save up to $750 per server per year on average.

***Keywords-** cloud computing, load variation, energy efficiency, data center, energy efficient scheduling*

I. Introduction

Cloud computing is an attractive way to deliver services to the end users in the form of utility over the Internet. It provides computation, software, data access, and storage services reducing infrastructure and service deployment costs, ensuring scalability, and responsiveness to changing conditions of the market. Flexibility and elasticity of the clouds foster enterprises to move from massive in-house computer clusters to cloud computing platforms [1].

Cloud computing is typically implemented using a set of geographically distributed data centers, and data centers are energy hungry consuming almost 1.5 percent of the world electricity [2]. On the other hand, data centers turn only around 15 percent of the consumed energy into computing results making it a large territory for potential optimization solutions.

In data centers only 40%[1] of the consumed energy is delivered to IT equipment [2], while the rest is shared between the cooling system (45%) and the power distribution system (15%) [3]. The IT-related consumption can be further divided into the power consumed by computing servers (70%) and communication equipment (30%) [4]. The Gartner Group estimates energy consumptions to account for up to 10% of the current data center operational expenses (OPEX) and can rise up to 50% within next few years [5].

Consequently, the need for energy efficient techniques in data centers is increasing [6], [31]. The two historically popular techniques for power savings in the computing systems are Dynamic Voltage and Frequency Scaling (DVFS) [8] and Dynamic Power Management (DPM) [9]. In DVFS, the supplied voltage and operating frequency are dynamically adjusted in steps based on the required performance level. The aforementioned adjustment helps to conserve the energy for partially loaded hardware due to non-linear relation between the operating frequency and the consumed power. The DPM techniques allow powering off the devices or putting them into a "sleep" mode. To make DPM scheme efficient, a scheduler should attempt to concentrate data center jobs on a smallest set of computing resources maximizing the amount of idle servers available for powering down [22]. Moreover, to be efficient, the scheduler should take into account the following two features: (a) the physical characteristics (e.g., energy) of enabling of a sleep mode and waking up the device and (b) a delay associated with the process [11].

In this paper we: (a) evaluate the energy-efficient scheduler with a job concentration policy, (b) estimate the amount of hardware resources required to accommodate incoming jobs for the time sufficient to wake up additional hardware resources, (c) evaluate the possibility of using energy-efficient hardware for these hardware resources, which is costly, but leads to a fewer power consumption, and (d) calculate economic effect of using energy-efficient hardware. The proposed contribution is focused on addressing the problem of data center energy efficiency in the top-to-bottom fashion where the design of job scheduling policies is guided by particularities of power management implemented in the hardware components.

The rest of the paper is organized as follows: Section II discusses the state-of-the-art in power management and provides a background on typical data center load patterns; Section III focuses on the proposed solution presenting how scheduling should be performed; Section IV presents the details of the evaluation scenario and discusses the obtained results; finally, Section V draws conclusions and outlines the directions for the future work on the topic.

[1] The values of energy consumption in data centers reported in this paper are approximate and may vary depending on data center architecture, hardware used, and running applications.

978-1-4673-3122-7/13/$31.00 ©2013 IEEE

II. Background and State-of-the-Art

A. Power Management Solutions

The challenge of making data center IT equipment energy-efficient can be addressed at two different levels: (a) by using "greener" technologies while building data center hardware or (b) by coordinating execution of user workloads in a way to allow the most of the computing and communications hardware trigger their power saving modes.

All power management solutions in data centers can be broadly categorized into static power management and dynamic power management solutions. Static power management is primarily achieved using low-power components that typically tradeoff the system performance and the energy consumed by the data center components [12]. In contrast, dynamic power management techniques focus on runtime optimization.

Cubic relation between the CPU's operational frequency and power consumption makes it attractive to gradually scale the operating frequency when no computing power is needed during runtime. These methods are commonly referred as DVFS [8] and can be performed in both computing servers and network switches. Unfortunately, DVFS benefits are limited, as computing components contribute to only a portion of total energy consumption. Other system components like memory modules, power supply units, and peripheral devices cannot scale with the load and account for up to two thirds of total energy consumption in computing servers and up to 80% in network switches [13]. As a result, most of the power savings can be achieved by addressing both the load-dependent and load-independent energy components. To achieve it, the DPM can power down or place in "sleep" mode not only individual component, but also an entire device. To make DPM scheme efficient, a scheduler must consolidate incoming jobs on a minimum set of computing resources to maximize the amount of idle servers [22]. The gains from consolidation scheduling will depend on the average data centers load level that typically stays around 30% [4].

In the next section we provide a detailed outlook on the typical load variations in data centers on different time scales. Further details on existing energy efficiency technologies are available in [14].

B. Load Variation in Data Centers

The data centers are designed to provide a required level of quality of service, defined in Service Level Agreements (SLAs), even at peak loads. Therefore, they tend to overprovision computing and communication resources. In fact, on average, datacenters are loaded only at 30% of their capacity [4]. Being biased to a single or just a few applications the load in data centers has a high correlation with the daytime hours of user activity and regional settings. During day time the, number of users is almost doubled as compared to the night time. Moreover, user arrival rate is not constant, but can spike due to the crowd effect. While peak loads may reach up to 90%, most of the time almost 70% of data center servers, switches, and links remain idle. However, idle servers still need to run OS software, maintain virtual machines, and power on peripheral devices and memory. As a result, even when being idle, the servers consume around two thirds of its peak power consumption [16]. In switches, this ratio is even higher [13]. The energy consumed is shared between the switch's chassis, the line cards, and the transceiver ports. Moreover, several Ethernet standards require uninterrupted transmissions of synchronization symbols at the physical layer to maintain synchronization preventing such switches of scaling their power consumption down, even when no user traffic is transmitted [17]. In [18], the authors introduce an optimized signal encoding scheme to tackle the aforementioned problem.

III. Proposed Solution

Concentrating or groupping the execution of the data center workloads in a minimum amount of computing and communication resources is one of the most widely used scheduling strategies [19], [29] for increasing energy efficiency. It is motivated by the fact that idle servers consume around two thirds of its peak energy consumption and remaining one third scales with the offered load linearly [21]. Moreover, the scheduling strategy of concentration can maximize the number of hardware resources that can be put into sleep mode [22].

Figure 1 shows typical workload distribution among data center servers using the concentration-based scheduler. With the typical data center load of 30% [4], approximately one third of all servers are fully loaded (Segment I of the chart), while remaining two thirds (Segment III of the chart) are left idle and can be put into sleep. The servers in between (correspond to Segment II of the chart) account for the incoming load variation and need to execute workloads even being not fully utilized on average.

In order to maximize bisection bandwidth the servers performing computations (Segment I) are spread between racks. Such that each rack hosts approximately equal number of fully loaded and idle servers. This allows keeping operating costs almost unchanged while maximizing bisection bandwidth available to each server.

Figure 1 shows that around 10% of the data center resources are allocated for the fluctuation in the incoming load, which typically corresponds to the short-term incoming workload rate variation on a second or minute time scales. However, on a daily time scale the average data center load can vary between 10% and 90% [4]. As a result, even the resources dedicated to account for incoming load variation should adapt voltage using the DVFS scheme or be put into sleep by DPM scheme.

The exact number of hardware resources kept operational depends on the speed of incoming load variation and hardware wakeup time. In addition, the energy consumption levels during power state transitions should also be taken into account. When a virtual machine changes its power state, it will consume in the order of 14 W for a wake-up and 17 W for a shutdown [25], in addition to its idle consumption. These values are approximate and may vary depending on the virtual machine, running applications, and the amount of memory used and needed to be stored on the hard disk.

Due to high energy and time overheads, the number of virtual machine wake-ups and shutdowns should be minimized.

Therefore, the number of servers assigned to compensate incoming load variation should be carefully selected and adjusted during runtime by predicting incoming workload variations based on the recent history. Leaving more hardware resources idle will lead to unnecessary energy consumption while underestimating the load will introduce additional delay in the workload execution and may confront with a customer SLA.

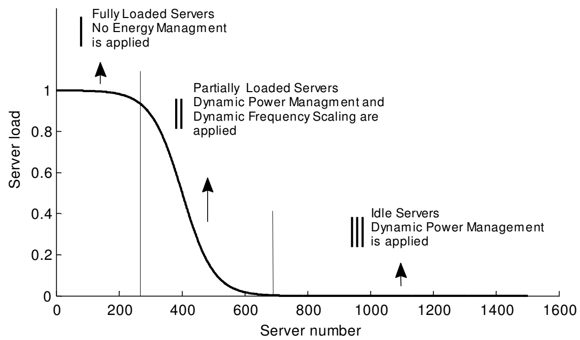

Figure 1. Load distribution amoung data center servers with concentration-based scheduling.

As we may see from Fig. 1, the scheduling policy based on workload concentration as a method to achieve energy efficiency allows to divide data center hardware into the following three categories according to the level of their utilization: (a) fully utilized (more than 95% of utilization), (b) partially utilized (95% of utilization and less), and (c) idle (zero utilization). In homogeneous data centers the same hardware can be used for all the three types. However, in this paper we would like to explore the benefits of using energy efficient hardware.

There are two main alternatives for building data center server hardware: (a) rack servers and (b) blade servers. Rack servers occupy the space of one rack unit (1RU). Each rack can host up to 42 of such servers providing them common power and networking solutions. Blade servers are designed to minimize space by eliminating redundant components. Being not limited to 1RU, a rack of blades can contain more than a hundred of servers. While a single blade server can be more expensive than a rack server, it becomes significantly less expensive in large quantities [27]. Blade servers being initially attractive for their price also have several drawbacks. For example, high server density in a rack increases a rack weight and becomes a heavy heat producer requiring an increased performance from the cooling system. But a more important drawback of blade servers is in their increased energy consumption [27]. This defines a tradeoff between Capital Expenditure (CAPEX) invested into data center hardware and Operating Expenses (OPEX) spent on the energy to power up data center hardware. In each particular case the choice of rack or blade servers should be guided by the projected load, data center lifespan, and regional cost of electricity.

Energy efficient low power architectures, like Intel Atom processor, were initially designed for mobile devices. Typically, they are more expensive and less computationally powerful, if compared with legacy server platforms. However, they allow so-called energy-proportional computing, when the amount of the consumed power scales almost linearly with the delivered computing performance. As a result, energy efficient platforms become an ideal hardware in data centers when used for only partially utilized or idle computing servers and network equipment.

Reference to Fig. 1, fully loaded servers can be implemented using either rack or blade servers. Similarly, there is no need to use energy-efficient hardware for idle servers (rightmost part of the chart) if it can be put into sleep mode that can reduce power consumption close to zero. However, for partially loaded servers energy-efficient hardware would be the most beneficial. Extra investments in energy-efficient hardware will return back with operational savings on energy. To understand the order of these savings and time required to breakeven, we analyzed typical hardware costs for all the three types of server hardware (rack, blade, and energy-efficient) and verified their average power consumption levels.

Table I summarizes our hardware selections. The values for power consumption and hardware costs are averaged for the following types of the hardware. The HP blade system c-Class, PowerEdge M610x Blade and IBM BladeCenter HS23 Express are selected as representative blade server systems. On the average they consume 1027 W of energy and cost $2,323. The HP DL380, PowerEdge R420, and IBM System x3550 M3 are selected as representative rack server hardware. With 225 W, they consume about a quarter of the rack servers consumption, but cost almost 79% more. The Seamicro SM10000-64 [20] is selected for energy-efficient hardware. It is a solution integrating 512 64-bit X86 Intel Atom processors in a single rack-like module offering common Ethernet switching, server management, and application load balancing solutions. For comparative studies the hardware cost and average power consumption reported in Table I are normalized to the equivalent of the rack server performance.

TABLE I. TYPICAL DATA CENTER SERVER HARDWARE AND ITS POWER CONSUMPTION

Server Type	Server Hardware	Average Power Consumption (W)	Hardware Average Cost ($)
Blade	HP blade system c-Class/ PowerEdge M610x Blade / IBM BladeCenter HS23	1027	2323
Rack	HP DL380/ PowerEdge R420/ IBM System x3550	225	2933
Energy-Efficient	Seamicro SM10000-64	55	3700

Figure 2 reports the costs of hardware and power consumed during operation for a typical data center life span of 10 years [26]. The costs are calculated for a single server with an average load of 50%. It can be observed that operating expenses largely prevail on the initial hardware investments. For example, a rack server consumes almost three times its initial costs over ten years. A blade server, being initially attractive with inexpensive hardware, turns to require the largest investment over ten years. In contrast to rack and blade servers, the power consumed by the energy-efficient server forms just a portion (around 40%) of the hardware cost. As a result, rack servers becomes more economically feasible after 1.7 years of operation, surpassing blade servers. However, after 5.3 years energy-efficient server overcomes high hardware costs with low energy consumption surpassing the rack server. Finally, on a ten years life span each rack server replaced with energy-efficient server saves $117 per year. Similarly, almost $742 is saved for every blade server replaced with energy-efficient servers.

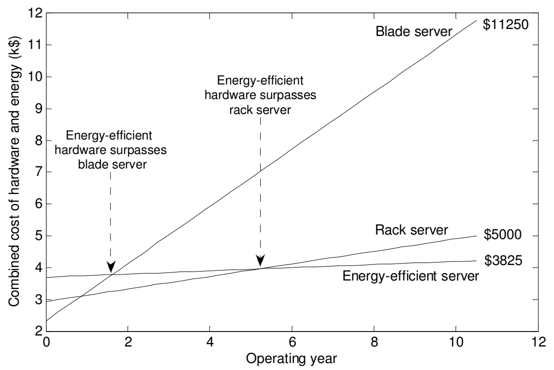

Figure 2. Costs of server hardware and power consumtion.

The costs reported in Fig. 2 correspond to the case of constantly running server hardware under partial (below 50%) load level. Therefore, they are appropriate to apply to partially loaded servers represented by Segment II of the chart in Fig. 1. However, these results should not be applied to fully loaded server hardware (Segment I) as well as for idle servers (Segment III) where sleep mode can be applied due to significant differences in energy consumption profiles.

Another important factor leading to power saving comes from the fact that two or more energy-efficient servers are needed to replace a single rack or a blade server due to their lower computing capacity. The aforementioned brings finer granularity to sleep mode employment. With even a minor load it will be required to keep the entire rack of servers operational. Alternatively, only one energy-efficient server from two (or more) needed for equivalent rack server replacement will stay operational while the rest will be put into sleep mode.

Now, having understood the level of cost savings per server, it is important to estimate the amount of partially loaded servers in a typical data center to assess the impact of the proposed scheme. The number of partially loaded servers is constantly changing during runtime following variations in the incoming load patterns guided by the speed of incoming load variation. Visually, it can be represented by the sharpness of the falling slope (Segment II) in Fig. 1. Sharp slopes represent deterministic systems, where user workloads arrive to the data center regularly, and are characterized by a fixed number of fully loaded servers. On the opposite, smooth slopes correspond to a bulky workload arrival which requires large variation in the number of fully loaded servers.

A representative example evaluated in the next session shows a small data center hosting 1,536 servers. Around 3% of the servers appeared to be the candidates for the replacement with energy-efficient hardware. This corresponds to the OPEX savings of up to $300,000.

IV. Performance Evalution

We used GreenCloud simulator [24], [28] to perform evaluation of the proposed methodology under realistic user load levels. GreenCloud, developed as an extension of ns-2 simulation platform [15], is the first cloud computing simulator to capture data center communication patterns at the packet level and calculate energy consumed by data center hardware (servers, switches, and links). GreenCloud is available for download at [7].

A. *Simulation Scenario*

Figure 3 presents a topology of a simulated three-tire data center architecture, selected for being the most widely used nowadays [10]. The Layer-3, which consists of the core switches, ensures tasks dispatching and load balancing among different parts of the data center. The Layer-2, comprised of the aggregation layer switches, serves sets of racks ensuring scalability of the data center. The Layer-1 hosts access switches and computing servers arranged into racks with a single Top-of-Rack (ToR) switch serving all communication demands of the entire rack.

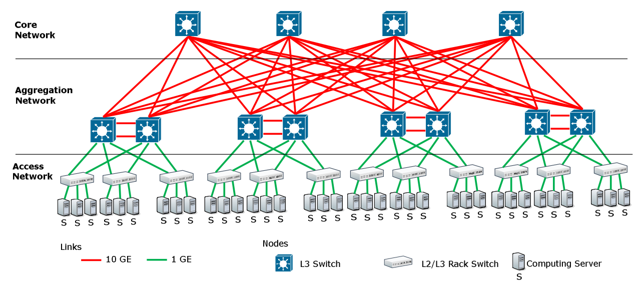

Figure 3. Simulated data center topology.

The simulated data center consists of 1536 servers arranged into 32 racks each holding 48 servers, served by 4 core and 8 aggregation switches. We used 1 Gigabit Ethernet (GE) links for interconnecting servers inside racks, while 10 GE links were used to form a fat-tree topology interconnecting access, aggregation, and core switches.

To simulate typical rate of user arrival to the system incoming workloads were exponentially distributed on a time scale. A scheduling decision is taken for each newly arrived task as from the approach of the runtime scheduling. After the scheduler specifies the exact virtual machine and hardware server, the task gets fragmented into a number of IP packets to be sent over the network. The size of the task is equal to 15 KB, which corresponds to 10 Ethernet packets. During their execution, the workloads produce a constant bit rate stream of 1Mb/s directed out of the data center. Such a stream is designed to mimic the behavior of the most common video sharing applications. To add uncertainties and mimic a database access, during the execution, each workload communicates with another randomly selected workload by sending a 75 KB message internally. The average load of the data center is kept around 30% that is distributed among the servers using energy efficient scheduler. This scheduler always attempts to concentrate execution of active workloads on a minimum amount of hardware maximizing the amount of idle devices for sleep mode.

B. Results

Figure 4 shows incoming load variation in terms of the number of loaded servers (left axis) and average data center load level (right axis) for 600 seconds of simulation time. The average load level fluctuates around 0.3 with a span of 160 servers.

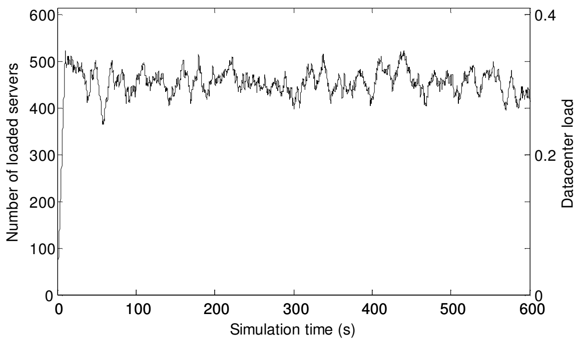

Figure 4. Data center load variation.

Figure 5 provides more insights into data center load distribution among the servers. It can be observed that most of the servers are either fully loaded or idle with only a fraction of servers being partially loaded. These partially loaded servers are the candidates for energy-efficient hardware replacement. A server is considered to be partially loaded when its load level is below 95%. The number of partially loaded servers varies with the intensity of incoming load variation, which is captured by three curves in Figure 5 for low, medium, and high load variation.

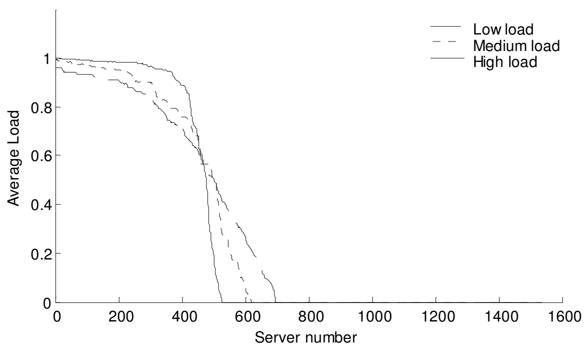

Figure 5. Data center load distribution.

However, not all the partially loaded servers should be replaced with energy-efficient hardware, but only those required to accommodate any incoming load variations for the time needed to wake up more servers. To estimate the number of servers that should be replaced, we take the first derivative from the number of active servers.

Figure 6 shows the rate of change in the number of loaded servers for both fully loaded and partially loaded servers (left axis) scaled with the typical server wake up time (right axis). According to [11], typical wake up time is estimated to be equal to 2 seconds.

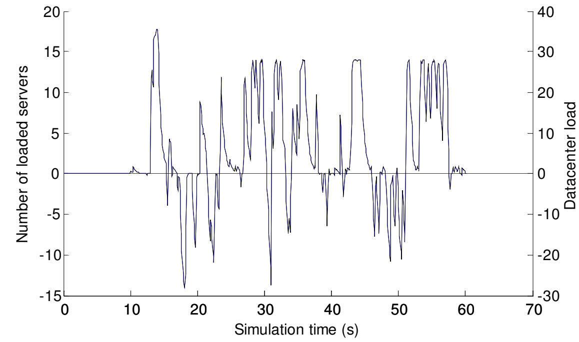

Figure 6. Change in the number of loaded servers.

In the aforementioned example up to 40 servers are needed to compensate incoming data center load for the duration of 2 seconds required to wake up additional servers. This corresponds to 2.6% of total 1,536 servers. Figure 7 calculates investment cost and operating expenses for different hardware profiles used for implementing these 40 servers. Blade server hardware turns to be the most expensive over fifteen years of data center life span with $450,000 followed by rack server hardware with $200,000 and energy-efficient hardware with $153,000. This means, $47,000 can be saved if rack servers are replaced with energy-efficient hardware and $297,000 in savings if blade servers are upgraded with energy-efficient hardware. However, taking into account the time required to break even, which is 1.75 years for the blade server and almost 5.2 years for the blade server, the benefits of energy-efficient hardware are maximized for the long-lived data centers.

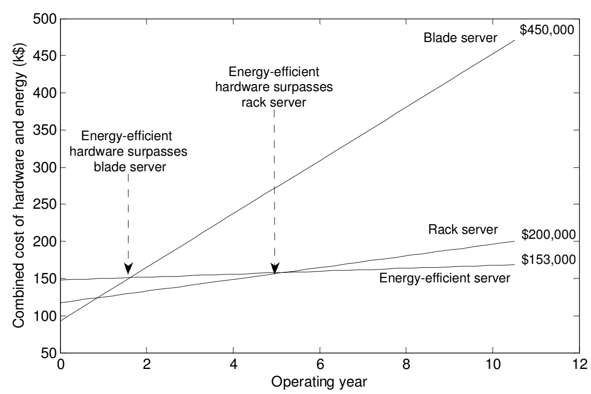

Figure 7. Hardware and energy costs for partially loaded servers.

V. Conclusions

Energy consumption in data centers is drastically increasing and becomes significant portion in data center operating expenses. Enabling sleep mode in idle hardware including computing servers and network equipment is the most efficient method to avoid unnecessary power consumption. However, changes in power modes introduce considerable delays. Moreover, inability to wake up a sleeping server immediately requires the availability of pool of idle servers able to accommodate incoming load in the short term to prevent QoS degradation.

In this paper we (a) evaluate energy-efficient job concentration scheduler, (b) estimate the amount of hardware resources required to accommodate incoming jobs for the time sufficient for waking up additional hardware resources, (c) propose using energy-efficient hardware for these hardware resources, which is costly, but results in a fewer power consumption, and (d) calculate economic effect of using such hardware. The proposed contribution is focused on addressing the problem of data center energy efficiency in the top-to-bottom fashion where the design of job scheduling policies is guided by particularities in power management of the hardware components. The evaluation results show that the proposed methodology can save up to $750 per server per year on average or a total of $300,000 in small data centers.

In the ongoing work we focus on QoS aspects of job execution, use job arrival traces from partner data centers, and consider prototype implementation of the designed methodology.

ACKNOWLEDGMENT

The authors would like to acknowledge the funding from National Research Fund, Luxembourg in the framework of ECOCLOUD project (C12/IS/3977641) and Marie Curie Actions of the European Commission (FP7-COFUND). Samee U. Khan's work was partly supported by the Young International Scientist Fellowship of the Chinese Academy of Sciences, (Grant No. 2011Y2GA01).

REFERENCES

[1] M. Rosoff, “Here's Why Cloud Computing Is So Hot Right Now”, 2011.

[2] “Report to congress on server and data center energy efficiency: Public law 109-431,” Lawrence Berkeley National Laboratory, 2008.

[3] Li Shang, Li-Shiuan Peh, and Niraj K. Jha, “Dynamic Voltage Scaling with Links for Power Optimization of Interconnection Networks,” Proceedings of the 9th International Symposium on High-Performance Computer Architecture table of contents, 2003.

[4] J. Liu, F. Zhao, X. Liu, and W. He, “Challenges Towards Elastic Power Management in Internet Data Centers”, WCPS 2009, in conjunction with ICDCS 2009., Montreal, Quebec, Canada, June 2009.

[5] Gartner Group, available at: http://www.gartner.com/.

[6] L. Wang and S. U. Khan, “Review of Performance Metrics for Green Data Centers: A Taxonomy Study,” Journal of Supercomputing, 2011.

[7] Greencloud - The green cloud simulator, available at http://greencloud.gforge.uni.lu/.

[8] J. Pouwelse, K. Langendoen, and H. Sips, “Energy priority scheduling for variable voltage processors,” International Symposium on Low Power Electronics and Design, pp. 28 – 33, 2001.

[9] L. Benini, A. Bogliolo, and G. De Micheli, “A survey of design techniques for system-level dynamic power management,” IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 8, no. 3, pp. 299 – 316, June 2000.

[10] “Cisco Data Center Infrastructure 2.5 Design Guide,” Cisco press, 2010.

[11] M. Zhao and R. J. Figueiredo, “Experimental study of virtual machine migration in support of reservation of cluster resources,” International Workshop on Virtualization Technology in Distributed Computing, New York, NY, USA, ACM, 2007, pp. 1-8.

[12] M. E. Tolentino, J. Turner, and K. W. Cameron, “Memory-miser: a performance-constrained runtime system for power-scalable clusters,” International Conference Computing Frontiers, pp. 237–246, 2007.

[13] Dennis Abts, Mike Marty, Philip Wells, Peter Klausler, and Hong Liu, “Energy Proportional Datacenter Networks”, International Symposium on Computer Architecture, 2010, pp. 338-347.

[14] G. L. Valentini, W. Lassonde, S. U. Khan, N. Min-Allah, S. A. Madani, J. Li, L. Zhang, L. Wang, N. Ghani, J. Kolodziej, H. Li, A. Y. Zomaya, C.-Z. Xu, P. Balaji, A. Vishnu, F. Pinel, J. E. Pecero, D. Kliazovich, and P. Bouvry, "An Overview of Energy Efficiency Techniques in Cluster Computing Systems," Cluster Computing, pp. 1-12, 2011.

[15] The Network Simulator Ns2, avialable at: http://www.isi.edu/nsnam/ns/.

[16] X. Fan, W. D. Weber, L. A. Barroso (2007) “Power provisioning for a warehouse-sized computer.” In: Proceedings of the 34th annual international symposium on computer architecture (ISCA ‘07). ACM, New York, pp 13–23.

[17] IEEE Std 802.3-2008, "Local and Metropolitan Area Networks,” pp.1-315, 2008.

[18] Y. Audzevich, P. Watts, S. Kilmurray, and A. W. Moore, “Efficient photonic coding: a considered revision,” SIGCOMM workshop on Green networking (GreenNets '11), ACM, New York, NY, USA, 2011.

[19] Bo Li, Jianxin Li, Jinpeng Huai, Tianyu Wo, Qin Li, and Liang Zhong, “EnaCloud: An Energy-Saving Application Live Placement Approach for Cloud Computing Environments,” IEEE International Conference on Cloud Computing, Bangalore, India, 2009.

[20] Seamicro, available at: http://www.seamicro.com/

[21] T. Horvath, T. Abdelzaher, K. Skadron, and Xue Liu, “Dynamic Voltage Scaling in Multitier Web Servers with End-to-End Delay Control,” IEEE Transactions on Computers, vol. 56, no. 4, pp. 444 – 458, 2007.

[22] D. Kliazovich, P. Bouvry, and Samee U. Khan, “DENS: Data Center Energy-Efficient Network-Aware Scheduling,” Cluster Computing, special issue on Green Networks, pp. 1-11, 2011.

[23] IEEE std 802.3ba-2010, “Media access control parameters, physical layers and management parameters for 40 Gb/s and 100 Gb/s operation,” June 2010.

[24] D. Kliazovich, P. Bouvry, Y. Audzevich, and S. U. Khan, “GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers,” IEEE Global Communications Conference (GLOBECOM), Miami, FL, USA, December 2010.

[25] A. C. Orgerie, M. D. de Assuncao, and L. Lefevre “Energy Aware Clouds”, Grids, Clouds and Virtualization, pp. 145-170, October 2010.

[26] “Introduction to Data Center,” Cisco Systems, 2009.

[27] “HP Blade vs Rack Servers - what's best for a mid-sized Oracle Infrastructure?”, available at: http://www.veriton.co.uk/roller/fmw/entry/hp_blade_vs_rack_servers.

[28] D. Kliazovich, P. Bouvry, and S. U. Khan, “GreenCloud: a packet-level simulator of energy-aware cloud computing data centers,” Journal of Supercomputing, special issue on Green Networks, vol. 62, no. 3, pp. 1263-1283, 2012.

[29] L. Wang, S. U. Khan, and J. Dayal, "Thermal Aware Workload Placement with Task-Temperature Profiles in a Data Center," Journal of Supercomputing, vol. 61, no. 3, pp. 780-803, 2012.

[30] Jonathan G. Koomy, “Estimating Regional Power Consumption By Servers: A Technical Note” Final report, December 2007.

[31] J. Kolodziej and S. U. Khan, "Data Scheduling in Data Grids and Data Centers: A Short Taxonomy of Problems and Intelligent Resolution Techniques," LNCS Transactions on Computational Collective Intelligence, 2011.