Analysis of Energy Aware Data Center using Green Cloud Simulator in Cloud Computing

Beant Singh Gill[*1], Satinderjit kaur Gill[*2], Prince Jain[#3]

* Eternal University
Baru Sahib, HP, India
#Malwa Polytechnic College
Faridkot, PB, India

***Abstract*-Cloud computing data centers are becoming popular for the provisioning of computing resources. The cost and operating expenses of data centers have increased with the increase in computing capacity. Several researches and surveys indicate that the energy utilized by computing and communication units within a data center contributes to a considerable portion of the data center operational costs. In this paper, we present a simulation environment for energy aware cloud computing data centers. Along with the workload distribution, the simulator is designed to capture details of the energy consumed by data center components as well as packet-level communication patterns between them. The simulation results obtained for two-tier, three-tier, and three-tier high-speed data center architectures which demonstrate the effectiveness of the simulator in utilizing different power management schema, such as voltage scaling, frequency scaling, and dynamic shutdown.**

Keywords-Data center, Green Cloud, Energy Efficiency, Simulation

I. Introduction

Clouds are a large pool of easily usable and accessible virtualized resources such as hardware, development platform and services. Cloud is an internetwork of different types of server which share resources. A cloud may be public, private, community or hybrid type and managed by the organizations or a third-party. When a cloud is made available in a pay-per-use manner to the public and service being sold as Utility Computing, it is known as Public Cloud. The Private Cloud refers to internal datacenters of an organization that are not made available to the public. The cloud computing is emerging as a model that use "everything as a service" and which in turn provided as cloud services. Virtualized physical resource, virtualized infrastructure, virtualized middleware platforms and business applications are being provided and consumed as service in the cloud. For example, a business solution model is either being built by using cloud service or being provided as a cloud service. The cloud service has stack of services which is arranged from top to bottom on the three layers that are, Hardware, system and application layer. Each component in this stack provides different types of service to cloud. These cloud computing services has much better than the traditional service provisions in context of reduced upfront investment, expected performance, high availability, infinite scalability and tremendous fault tolerance capability [1-4].

A. *Internet Data Center*

Internet Data Center (IDC) is a common form to host cloud computing. An IDC usually deploys hundreds or thousands of blade servers, densely packed to maximize the space utilization. Running services in consolidated servers in IDCs provides customers an alternative to running their software or operating their computer services in-house [5, 6]. The major benefits of IDCs include the usage of economies of scale to amortize the cost of ownership and the cost of system maintenance over a large number of machines. With the rapid growth of IDCs in both quantity and scale, the energy consumed by IDCs, directly related to the number of hosted servers and their workload has been increased. The rated power consumptions of servers have increased by 10 times over the past ten years. This surging demand calls for the urgent need of designing and deployment of energy-efficient Internet data centers [7, 8].

A modern state-of-the-art data center has three main components-data storage, servers, and a local area network (LAN). The data center connects to the rest of the network through a gateway router. The power consumption data for each server was obtained by first calculating the maximum power using HP's power calculator, then following the convention that average power use for midrange/high-end servers is 66% of maximum power. In the following, I outline the functionality of this equipment as well as some of the efficiency improvements in cloud computing data centers over traditional data centers. Long-term storage of data in a data center is provided by hard disk arrays, together with associated equipment [5, 8].

B. *Green cloud computing*

Even though there is a great concern in the community that Cloud computing can result in higher energy usage by the datacenters, the Cloud computing has a green lining. There are several technologies and concepts employed by Cloud providers to achieve better utilization and efficiency than traditional computing. Therefore, comparatively lower carbon emission is expected in Cloud computing due to highly energy efficient infrastructure and reduction in the IT infrastructure itself by multi-tenancy. The key driver technology for energy efficient Clouds is "Virtualization," which allows significant improvement in energy efficiency of Cloud providers by leveraging the economies of scale associated with large number of organizations sharing the same infrastructure. Virtualization is the process of presenting a logical grouping or subset of computing resources so that they can be accessed in ways that give benefits over the original configuration [9, 10]. By consolidation of underutilized servers in the form of multiple virtual machines sharing same physical server at higher utilization, companies can gain high savings in the form of space, management, and energy [10].

II. METHODOLOGY USED

A. *Problem Formulation*

Recently, cloud computing services have become increasingly popular due to the evolving data centers and parallel computing paradigms. The operation of large geographically distributed data centers requires considerable amount of energy that accounts for a large slice of the total operational costs for cloud data centers for up to 10% of the current data center operational expenses. High power consumption generates heat and requires an accompanying cooling system that costs per year for classical data centers which drastically decreases hardware reliability and may potentially violate the Service Level Agreement with the customers.

The first power saving solutions focused on making the data center hardware components power efficient. Technologies, such as Dynamic Voltage and Frequency Scaling (DVFS), and Dynamic Power Management (DPM) were extensively studied and widely deployed but their power-down and power-off methodologies, the efficiency of these techniques is at best limited. However, achieving the above requires central coordination and energy-aware workload scheduling techniques. Typical energy-aware scheduling solutions attempt to concentrate the workload in a minimum set of the computing resources.

B. *Methodology*

The simulations of an energy-aware data center for two-tier (2T), three-tier (3T), and three-tier high-speed (3Ths) architectures. For comparison reasons, I fixed the number of computing nodes to 1536 for all three topologies, while the number and interconnection of network switches varied. In contrast with other architectures, a 2T data center does not include aggregation switches. The core switches are connected to the access network directly using 1 GE links and interconnected between them using 10 GE links.

The 3Ths architecture mainly improves the 3T architecture with providing more bandwidth in the core and aggregation parts of the network. The bandwidth of the C1– C2 and C2–C3 links in the 3Ths architecture is ten times of that in 3T and corresponds to 100 GE and 10 GE, respectively. The availability of 100 GE links allows keeping the number of core switches as well as the number of paths in the ECMP routing limited to 2 serving the same amount switches in the access.

The workload generation events and the size of the workloads are exponentially distributed. The average size of the workload and its computing requirement depends on the type of task. For CIW workloads, the relation between computing and data transfer parts is chosen to be 1/10, meaning that with a maximum load of the data center its servers will be occupied for 100% while the communication network will be loaded for 10% of its maximum capacity. For DIW workloads the relation is reverse. Under the maximum load, the communication network is loaded for 100% while computing servers for only 10%. Balanced workloads load computing servers and data center network proportionally.

III. DATA CENTRE ARCHITECTURE

The pool of servers in today's data centers overcomes 100,000 hosts with around 70% of all communications performed internally. This creates a challenge in the design of interconnected network architecture and the set of communication protocols. Given the scale of a data center, the conventional hierarchical network infrastructure often becomes a bottleneck due to the physical and cost driven limitations of the used networking equipment. Specifically, the availability of 10 Gigabit Ethernet (GE) components and their price defined the way the data center architectures evolved. The 10 GE transceivers are still too expensive and probably offer more capacity than needed for connecting individual servers. However, their penetration level keeps increasing in the backbone networks, metro area networks, and data centers [11].

A. *Two-tier data center architectures*

In this example, computing Servers physically arranged into racks form the tier-one network. At the tier-two network, Layer-3 (L3) switches provide full mesh connectivity using 10 GE links. The Equal Cost Multi-Path (ECMP) routing is used as a load balancing technology to optimize data flows across multiple paths. It applies load balancing on TCP and UDP packets on a per-flow basis using express hashing techniques requiring almost no processing from a switch's CPU. Other traffic, such as ICMP, is typically not processed by ECMP and forwarded on a single predefined path. The two-tier architecture worked well for early data centers with a limited number of computing servers. Depending on the type of switches used in the access network, the two-tier data centers may support up to 5500 nodes. The number of core switches and capacity of the core links defines the maximum network bandwidth allocated per computing server [4].

B. *Three-tier data center architectures*

Three-tier data center architectures are the most common nowadays. They include: Access, Aggregation and Core layers. The availability of the aggregation layer facilitates the increase in the number of server nodes while keeping inexpensive Layer-2 (L2) switches in the access network, which provides a loop-free topology. Because the maximum number of ECMP paths allowed is eight, a typical three tier architecture consists of eight core switches. Such architecture implements an 8-way ECMP that includes 10 GE Line Aggregation Groups (LAGs), which allow a network client to address several links and network ports with a single MAC address [12-15].

C. *Three-tier high-speed data center architectures*

Three-tier high-speed data center architectures are designed to optimize the number of nodes, capacity of core, and aggregation networks that are currently a bottleneck, which limit the maximum number of nodes in a data center or a per-node bandwidth [9]. With the availability of 100 GE links between the core and aggregation switches, reduces the number of the core switches, avoids the shortcomings of LAG technology, reduces cablings, and considerably increases the maximum size of the data center due to physical limitations. Fewer ECMP paths will lead to the flexibility and increased network performance [12, 13].

IV. SIMULATION OF DATA CENTRE

Only a part of the energy consumed by the data center gets delivered to the computing servers directly. A major portion of the energy is utilized to maintain interconnection links and network equipment operations. The rest of the electricity is wasted in the power distribution system, dissipates as heat energy, and used up by air-conditioning systems. It distinguishes three energy consumption components: computing energy, communicational energy and Energy component related to the physical infrastructure of a data center. The efficiency of a data center can be defined in terms of the performance delivered per watt, which may be quantified by the following two metrics: Power Usage Effectiveness (PUE) and Data Center Infrastructure Efficiency (DCiE). Both PUE and DCiE describe which portion of the totally consumed energy gets delivered to the computing servers.

A. *Structure of the simulator*

GreenCloud is an extension to the network simulator Ns2 which we developed for the study of cloud computing environments. The GreenCloud offers users a detailed fine-grained modeling of the energy consumed by the elements of the data center, such as servers, switches, and links. Moreover, GreenCloud offers a thorough investigation of workload distributions. Furthermore, a specific focus is devoted on the packet-level simulations of communications in the data center infrastructure, which provide the finest-grain control and is not present in any cloud computing simulation environment. The structure of the GreenCloud extension mapped onto the three-tier data center architecture [3, 8, 12].

1) Servers

Servers are the staple of a data center that are responsible for task execution. In GreenCloud, the server components implement single core nodes that have a preset on a processing power limit in MIPS or FLOPS, associated size of the memory resources, and contain different task scheduling mechanisms ranging from the simple round-robin to the sophisticated DVFS- and DNS-enabled. The servers are arranged into racks with a Top-of-Rack (ToR) switch connecting it to the access part of the network. The power model followed by server components is dependent on the server state and its CPU utilization. An idle server consumes about 66% of energy compared to its fully loaded configuration. This is due to the fact that servers must manage memory modules, disks, I/O resources, and other peripherals in an acceptable state. Then, the power consumption linearly increases with the level of CPU load. As a result, the aforementioned model allows implementation of power saving in a centralized scheduler that can provision the consolidation of workloads in a minimum possible amount of the computing servers.

Another option for power management is Dynamic Voltage/Frequency Scaling (DVFS) which introduces a tradeoff between computing performance and the energy consumed by the server. The DVFS is based on the fact that switching power in a chip decreases proportionally to V2*f, where V is voltage and f is the switching frequency. Moreover, voltage reduction requires frequency downshift. This implies a cubic relationship from f in the CPU power consumption. Note that server components, such as bus, memory, and disks, do not depend on the CPU frequency. Therefore, the power consumption of an average server can be expressed as follows:

$$P = P_{fixed} + P_f \cdot f^3 \ldots\ldots (1)$$

Where Pfixed is accounts for the portion of the consumed power which does not scale with the operating frequency f, while Pf is a frequency-dependent CPU power consumption [2, 12]. The curve is built for a typical server running an Intel Xeon processor. It consumes 301 W of energy with around 130 W allocated for peak CPU power consumption and around 171 W allocated for other peripheral devices. The scheduling depends on the server load level and operating frequency, and aims at capturing the effects of both of the DVFS and DPM techniques [17].

2) Switches and Links

Switches and Links form the interconnection fabric that delivers workload to any of the computing servers for execution in a timely manner. The interconnection of switches and servers requires different cabling solutions depending on the supported bandwidth, physical and quality characteristics of the link. The quality of signal transmission in a given cable determines a tradeoff between the transmission rate and the link distance, which are the factors defining the cost and energy consumption of the transceivers.

The twisted pair is the most commonly used medium for Ethernet networks that allows organizing Gigabit Ethernet transmissions for up to 100 meters with the consumed transceiver power of around 0.4Wor 10 GE links for up to 30 meters with the transceiver power of 6W. The twisted pair cabling is a low cost solution. However, for the organization of 10 GE links it is common to use optical multimode fibers. The multimode fibers allow transmissions for up to 300 meters with the transceiver power of 1 W. On the other hand, the fact that multimode fibers cost almost 50 times of the twisted pair cost motivates the trend to limit the usage of 10 GE links to the core and aggregation networks as spending for the networking infrastructure may top 10–20% of the overall data center budget [9, 11, 12].

The number of switches installed depends on the implemented data center architecture. However, as the computing servers are usually arranged into racks, the most common switch in a data center is the Top-of-Rack (ToR) switch. The ToR switch is typically placed at the top unit of the rack unit (1RU) to reduce the amount of cables and the heat produced. The ToR switches can support either gigabit or 10 gigabit speeds. However, taking into account that 10 GE switches are more expensive and that current capacity of aggregation and core networks is limited, gigabit rates are more common for racks.

Similar to the computing servers early power optimization proposals for interconnection network were based on DVS links. The DVS introduced a control element at each port of the switch that depending on the traffic pattern and current levels of link utilization could downgrade the transmission rate. Due to the comparability requirements, only few standard link transmission rates are allowed, such as for GE links 10 Mb/s, 100 Mb/s, and 1 GB/s are the only options. On the other hand, the power efficiency of DVS links is limited as only a portion of the consumed power scales linearly with the link rate [16, 12]. As demonstrated by the experiments, the energy consumed by a switch and all its transceivers can be defined as:

$P_{switch} = P_{chassis} + n_{linecards} + P_{linecard} + \sum_{i=0}^{r} n\ ports, r + P_r$……….. (2)

where Pchassis is related to the power consumed by the switch hardware, Plinecard is the power consumed by any active network line card, Pr corresponds to the power consumed by a port running at the rate r. only the last component appears to be dependent on the link rate while other components, such as Pchassis and Plinecard remain fixed for all the duration of switch operation. Therefore, Pchassis and Plinecard can be avoided by turning the switch hardware off or putting it into sleep mode. The proposed GreenCloud simulator implements energy model of switches and links with the values of power consumption for different elements taken in accordance. The implemented powers saving schemes are: DVS only, DNS only and DVS with DNS [12, 17].

3) Workloads

Workloads are the objects designed for universal modeling of various cloud user services, such as social networking, instant messaging, and content delivery. In grid computing, the workloads are typically modeled as a sequence of jobs that can be divided into a set of tasks. The tasks can be dependent, requiring an output from other tasks to start execution, or independent. Moreover, due to the nature of grid computing applications, the number of jobs available prevail the number of computing resources available.

While the main goal is the minimization of the time required for the computing of all jobs which may take weeks or months, the individual jobs do not have a strict completion deadline. In cloud computing, incoming requests are typically generated for such applications like web browsing, instant messaging, or various content delivery applications. The jobs tend to be more independent, less computationally intensive, but have a strict completion deadline specified in SLA. To cover the vast majority of cloud computing applications, define three types of jobs. The execution of each workload object in GreenCloud requires a successful completion of its two main components: computing and communicational [12, 18].

The computing component defines the amount of computing that has to be executed before a given deadline on a time scale. The deadline aims at introducing QoS constraints specified in SLA. The communicational component of the workload defines the amount and the size of data transfers that must be performed prior, during, and after the workload execution. It is composed of three parts: size of the workload, size of internal and size of external to the data center communications [12, 19].

V. RESULTS AND DISCUSSION

The simulations of an energy-aware data center for two-tier, three-tier and three-tier high-speed architectures. For comparison reasons, fixed the number of computing nodes to 1536 for all three topologies, while the number and interconnection of network switches varied. In contrast with other architectures, a 2T data center does not include aggregation switches. The core switches are connected to the access network directly using 1 GE links and interconnected between them using 10 GE links. The 3Ths architecture mainly improves the 3T architecture with providing more bandwidth in the core and aggregation parts of the network. The bandwidth of the C1–C2 and C2–C3 links in the 3Ths architecture is ten times of that in 3T and corresponds to 100 GE and 10 GE, respectively. The availability of 100 GE links allows keeping the number of core switches as well as the number of paths in the ECMP routing limited to 2 serving the same amount switches in the access. The propagation delay of all the links is set to 10 ns.

Parameters	**Two tier DC**	**Three Tier DC**	**Three Tier High Speed DC**
Core nodes	16	8	2
Aggregation nodes	0	8	4
Access Switches	512	512	512
Servers	1536	1536	1536

Table 1: data center Simulation Parameters

The workload generation events and the size of the workloads are exponentially distributed. The average size of the workload and its computing requirement depends on the type of task. For CIW workloads, the relation between communication network and data transfer parts is chosen to be 1/10. For DIW workloads the relation is reverse. Balanced workloads load computing servers and data center network proportionally. The workloads arrived to the data center are scheduled for execution using energy aware green scheduler. This green scheduler tends to group the workloads on a minimum possible amount of computing servers. In order to account for DIW workloads, the scheduler continuously tracks buffer occupancy of network switches on the path. In case of congestion, the scheduler avoids using congested routes even if they lead to the servers able to satisfy computational requirement of the workloads. The servers left idle are put into sleep mode while on the under loaded servers the supply voltage is reduced. The time required to change the power state in either mode is set to 100 ms.

The whole load of the data center is mapped onto approximately one third of the servers maintaining load at a peak rate. This way, the remaining two thirds of the servers can be shut down using DNS technique. A tiny portion of the approximately 50 out of 1536 servers which load represents a falling slope of the chart are under-utilized on average, and DVFS technique can be applied on them. The server peak energy consumption of 301Wis composed of 130W allocated for a peak CPU consumption and 171 W consumed by other devices. As the only component which scales with the load is the CPU power, the minimum consumption of an idle server is bounded and corresponds to 198W were also a portion of CPU power consumption of 27 W required to keep the operating system running is included [7, 12, 17].

Parameters	Two tier PC	Three Tier PC	Three Tier High Speed PC
Data center	477	503	508
Servers	351	351	351
Switches	126	152	157
Core	51	25	56
Aggregation	0	51	25
Access	75	75	75

Table 2: Data center Architecture Power Consumption

The switches consumption is almost constant for different transmission rates as most of the power is consumed by their chassis and line cards and only a small portion is consumed by their port transceivers. Depending on the employed data center topology, the core and aggregation switches will consume differently. For the 3T topology where the fastest links are 10 G the core and aggregation switches consume a few kilowatts, while in the 3Ths topology where links are of 10 G speed faster switches are needed which consume tens of kilowatts. The simulation results obtained for three evaluated data center topologies with no energy saving management involved for an average load of the data center of 30% [9, 12, 21]. The obtained numbers aim to estimate the scope of the energy-related spending components in modern data centers and define where the energy management schemes would be the most efficient.

On average, the data center consumption is around 432 kWh during an hour of the runtime. The processing servers share around 70% of total data center energy consumption, while the communicational links and switches account for the rest 30%. Furthermore, the consumption of switches breaks with 17% allocated for core switches, 34% for aggregation switches, and 50% for the access switches. It means that after computing servers lowering the power consumption of access switches will have the highest impact. The core and aggregation switches together account for 15% of total energy consumption. However, taking into account the requirements for network performance, load balancing, and communication robustness, the obvious choice is to keep core and aggregation switches constantly running possibly applying communication rate reduction in a distributed manner.

Parameters	No Energy Saving	DVFS	DNS	DVFS+DNS
Data center	503	96	37	35
Servers	351	97	39	37
Switches	152	95	32	31
Energy Cost	$441	$435	$163	$157

Table 3: Comparison of energy efficient scheme

The data center network accounts for the differences between power consumption levels of different data center architectures. With the respect to the 2T architecture, the 3T architecture adds around 25 kW for aggregation layer which enables the data center scale beyond 10,000 nodes. The 3Ths architecture contains fewer core and aggregation switches. However, the availability of 100 G links comes at a price of the increase per-switch energy consumption. As a result, a 3Ths network consumes more than a 3T network. Compares the impact on energy consumption of DVFS, DNS, and DVFS with DNS schemes applied on both computing several and networking equipment. The results are obtained for balanced tasks loading both computing servers and interconnection network equally for an average system load of 30%.

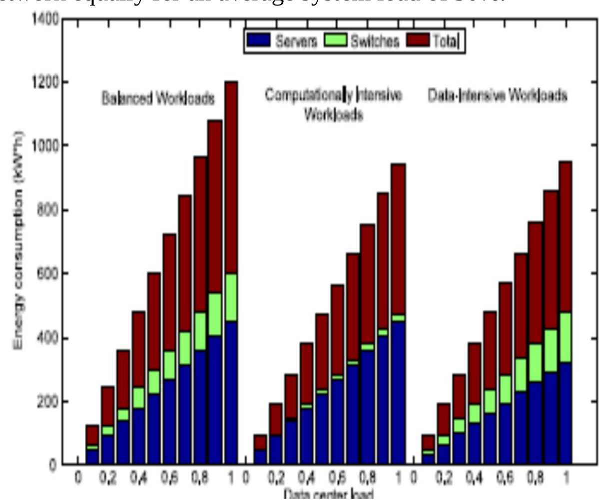

Figure 1: Data center energy consumption for different types of workloads

The DVFS scheme alone reduces power consumption to only 96% from the nominal level. Most of the power saving in servers comes from downshifting CPU voltage on the under-loaded servers. On the other hand, DVFS shows itself ineffective for the switches as only 3% to 15% of the switch's energy is sensitive to the transmission rate variation. The most effective results are obtained by DNS scheme. It is equally effective for both servers and switches as the most of their energy consumed shows no dependency on the operating frequency. However, in order to utilize DNS scheme effectively, its design should be coupled with the data center scheduler positioned to unload the maximum number of the servers.

The bottom of the table provides estimates of the data center energy cost on a yearly basis. Initial energy spending of $441 thousand can be reduced down to almost a third, $157 thousand, by a combination of DVFS and DNS schemes. The curves are presented for balanced type of the workloads and correspond to the total data center consumption as well as the energy consumed by the servers and switches. The DVFS scheme shows itself little sensitive to the input load of the servers and almost insensitive to the load of network switches. On the contrary, the DNS scheme appears to capture load variation precisely adjusting power consumptions of both servers and switches accordingly. The results reported are averaged over 20 runs with the random seed controlling random number generator. The introduced uncertainty affected mainly the way the workloads arrive to the data center slightly impacting the number of servers and network switches required to be powered. The maximum variance of 95% confidence intervals from the mean value accounted for less than 0.2% for the energy consumed by the servers and less than 0.1% for the energy consumed by the network switches.

VI. CONCLUSION AND FUTURE WORK

A simulation environment for energy-aware cloud computing data centers. GreenCloud is designed to capture details of the energy consumed by data center components as well as packet-level communication patterns between them. The simulation results obtained for two-tier, three-tier, and three-tier high-speed data center architectures demonstrate applicability and impact from the application of different power management schemes like voltage scaling or dynamic shutdown applied on the computing as well as on the networking components. The future work will focus on the simulator extension adding storage area network techniques and further refinement of energy models used in the simulated components.

REFERENCES

[1] Raj Kumar Buyya, Rajiv Ranjan, Rodrigo N. Calheiros, “Modeling and simulation of scalable cloud computing environments and the CloudSim toolkit: challenges and opportunities”, proceedings of the 7th high performance computing and simulation conference, Leipzig, Germany, June 21-24, 2009.

[2] Yiyu Chen, Amitayu Das, Wubi Qin, Anand Sivasubramaniam, Qian Wang, Natarajan Gautam, “Managing Server Energy and Operational costs in Hosting centers”, Proceedings of the ACM Sigmetrics international conference on measurement and modeling of computer systems, ACM, New York, pp: 303–314, 2005.

[3] Gong Chen, Wenbo He, Jie Liu, Suman Nath, Leonidas Rigas, Lin Xiao, Feng Zhao, “Energy-aware Server Provisioning and Load Dispatching for Connection-Intensive Internet Services”, in 5th USENIX symposium on networked systems design and implementation”, PP: 337-350, Berkeley, USA 2008.

[4] Xiaobo Fan, Wolf-Dietrich Weber, Luiz André Barroso, “Power provisioning for a warehouse-sized computer”, Proceedings of the ACM international symposium on computer architecture, San Diego, CA, June 2007.

[5] Fan Xiaobo, Weber Wolf-Dietrich, Barroso Luiz Andre, “Power provisioning for warehouse sized computer”, Proceedings of the 34th annual international symposium on computer architecture (ISCA), ACM, pp: 13–23, New York, June 2007.

[6] Farrington N, Rubow E, Vahdat A, “Data center switch architecture in the age of merchant silicon”, Proceedings of the 17th IEEE symposium on high performance interconnects (HOTI’ 09), IEEE Computer Society, Washington, pp: 93–102, 2009.

[7]Albert G. Greenberg, Parantap Lahiri, David A. Maltz, Parveen Patel, Sudipta Sengupta, “Towards next generation data center architecture: scalability and commoditization”, Proceedings of the ACM Sigcomm, Workshop on Programmable Routers for Extensible Services of Tomorrow PRESTO 2008, Seattle, WA, USA, August 22, 2008

[8] Chuanxiong Guo, Haitao Wu, Kun Tan, Lei Shiy, Yongguang Zhang, Songwu Lu, “DCell: A Scalable and Fault Tolerant Network Structure for Data Centers,” In SIGCOMM 2008 on Data communication, PP: 75-86, Seattle, USA, August 2008.

[9] Chuanxiong Guo, Guohan Lu, Dan Li, Haitao Wu, Xuan Zhang, Yunfeng Shi, Chen Tian, Yongguang Zhang, and Songwu Lu, “BCube: A High Performance, Server-centric Network Architecture for Modular Data Centers,” In SIGCOMM 2009 on Data communication, PP: 63-74, Barcelona, Spain, August 2009.

[10] Tibor Horvath, Tarek Abdelzaher, Kevin Skadron and X. Liu, "Dynamic Voltage Scaling in Multi-Tier Web Servers with End-to-End Delay Control", IEEE Transaction *on* Computers, vol. 56, no. 4, pp. 444-458, April 2007.

[11] IEEE 802.3ad-2000, “Amendment to Carrier Sense Multiple Access With Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications-Aggregation of Multiple Link Segments”, IEEE Press, New York, March 2000.

[12] Dzmitry Kliazovich, Pascal Bouvry, Samee Ullah Khan, “GreenCloud: a packet-level simulator of energy-aware cloud computing data centers”, Springer Science+Business Media, LLC, Nov 2010.

[13] IEEE std 802.3ba-2010, “Media Access Control parameters, Physical layers and Management parameters for 40 Gb/s and 100 Gb/s operation”, IEEE press, NY, June, 2010

[14] Liang Zhong, Bo Li, Jianxin Li, Jinpeng Huai, Tianyu Wo, Qin Li and Zhong L, “ EnaCloud: An Energy-saving Application Live Placement Approach for Cloud Computing Environments”, IEEE International Conference on Cloud Computing, PP: 17-24, Bangalore, India, 2009.

[15] Seung-Hwan Lim, Bikash Sharma, Gunwoo Nam, Eun Kyoung Kim, and Chita R. Das, “MDCSim: AMulti-Tier Data center simulation, platform”, IEEE international conference on cluster computing and workshops (CLUSTER), PP: 1-9, Aug-Sept, 2009.

[16] Jie Liu, Feng Zhao, Xue Liu, Wenbo He, “Challenges Towards Elastic Power Management in Internet Data Centers”, Proceedings of the 2nd international workshop on cyber-physical systems (WCPS) in conjunction with ICDCS 2009, Montreal, Canada, June 2009.

[17] Priya Mahadevan, Puneet Sharma, Sujata Banerjee and Parthasarathy Ranganathan, “Energy Aware Network Operations,” IEEE INFOCOM, HP labs, pp: 1 – 6, 2009.

[18] Priya Mahadevan, Puneet Sharma, Sujata Banerjee and Parthasarathy Ranganathan, “A power benchmarking framework for network devices”, Proceedings of the 8th international IFIP-TC 6 networking conference, Aachen, Germany, May, 2009.

[19] Moore J, Chase J, Puneet Sharma and Parthasarathy Ranganathan, “Making scheduling cool: Temperature aware workload placement in data centers”, Proceedings of the Annual conference on USENIX Annual Technical Conference, ATEC '05, USENIX Association, Berkeley, CA, USA, 2005.

[20] Ramya Raghavendra, Parthasarathy Ranganathan, Vanish Talwar, Zhikui Wang, Xiaoyun Zhu, “No power struggles: Coordinated multi-level power management for the data center”, SIGARCH Computer Architecture News, vol: 36(1), PP:48-59, 2008.