

Models for Efficient Data Replication in Cloud Computing Datacenters

Dejene Boru¹, Dzmityr Kliazovich², Fabrizio Granelli³, Pascal Bouvry², Albert Y. Zomaya⁴

¹ CREATE-NET

Via alla Cascata 56/D, Trento, Italy
dejene.oljira@create-net.org

² University of Luxembourg

6 rue Coudenhove Kalergi, Luxembourg
dzmitry.kliazovich@uni.lu, pascal.bouvry@uni.lu

³ DISI - University of Trento

Via Sommarive 14, Trento, Italy
granelli@disi.unitn.it

⁴ School of Information Technologies

University of Sydney, Australia
albert.zomaya@sydney.edu.au

Abstract: Cloud computing is a computing model where users access ICT services and resources without regard to where the services are hosted. Communication resources often become a bottleneck in service provisioning for many cloud applications. Therefore, data replication which brings data (e.g., databases) closer to data consumers (e.g., cloud applications) is seen as a promising solution. In this paper we present models for energy consumption and bandwidth demand of database access in cloud computing datacenter. In addition we propose an energy efficient replication strategy based on the proposed models which results in improved Quality of Service (QoS) with reduced communication delays. The evaluation results obtained with extensive simulations help to unveil performance and energy efficiency tradeoffs as well as guide the design of future data replication solutions.

Keywords: Cloud computing, Data replication, Energy efficiency

I. INTRODUCTION

Cloud computing is a computing model, where users access ICT services based on their requirements without regard to where the services are hosted. It offers a dynamic provisioning of computing resources supported by datacenters employing virtualization technology. Users can benefit from the convenience of accessing computing services globally from centrally managed backups, high computational capacity and flexible billing strategies [1]. Cloud computing leverages the efficient utilization of servers, data center power planning, virtualization, and optimized software stacks. Nevertheless, cloud computing datacenters consume huge amounts of electricity resulting in higher total cost of operation and carbon dioxide (CO₂) emission to the environment [2].

The growth of Internet services at an unprecedented rate requires the development of novel optimization techniques at all levels to cope with escalation in energy consumption.

Data center resources such as computing, storage, power distribution and cooling are over provisioned to ensure high levels of reliability [3]. Cooling and power distribution

systems consume around 45% and 15% of the total energy respectively, while leaving roughly 40% to the IT equipment [4]. These 40% are shared between computing servers and networking equipment. Depending on the data center load the communication network consumes 30 to 50% of the total power used by the IT equipment [5].

A wide range of datacenter energy efficiency solutions rely on the fact that datacenter infrastructures are underutilized [10] and over provisioned [3]. In essence, the Dynamic Power Management (DPM) method puts idle components into sleep mode [10]. The other method called Dynamic Voltage and Frequency Scaling (DVFS) [11] exploits the relation between power consumption P , supplied voltage V , and operating frequency f : $P = V^2 * f$. Reducing voltage and frequency reduces the power consumption. In computing servers, the effect of DVFS is limited, as power reduction applies only to the CPU, while the other components continue consuming at their peak rates. For communication devices, Dynamic Voltage Scaling is combined with Dynamic Network Shutdown (DNS) for power consumption optimization.

The performance of cloud computing applications depends largely on the availability and efficiency of high-performance communication resources. For better reliability and high performance low latency service provisioning, data resources can be brought closer (replicated) to the physical infrastructure, where the cloud applications are running. However, datacenter infrastructures, such as storage and network devices are required to maintain replicas. Moreover, new replicas need to be synchronized and changes made at one of the sites need to be reflected at the other locations. This incurs an underlying communication costs both in terms of energy and in terms of bandwidth. In this regard, several data replication models for optimization of datacenter energy and bandwidth have been proposed in the literature [3], [6],[7],[8]. These strategies optimize system bandwidth and data availability between geographically distributed data centers. However, none of them focuses on energy efficiency and replication techniques inside data centers.

To address this gap, we propose a data replication technique for cloud computing data centers which optimizes energy consumption, network bandwidth and communication delay both between geographically distributed data centers as well as inside each datacenter. Specifically, our contributions can be summarized as follows.

- Modeling of energy consumption characteristics of data center IT infrastructures.
- Development of a data replication approach for joint optimization of energy consumption and bandwidth capacity of data centers
- Optimization of communication delay to provide quality of user experience for cloud applications.
- Performance evaluation of the developed replication strategy through mathematical modeling and using a packet-level cloud computing simulator, GreenCloud [9].
- Analysis of the tradeoff between performance, serviceability, reliability and energy consumption.

The rest of the paper is organized as follows: In Section II we present a mathematical model for energy consumption, bandwidth demand and delay of cloud applications. Section III provides evaluation of the model outlining theoretical limits for the proposed replication scenarios. Section IV concludes the paper and provides an outline for the future work on the topic.

II. MODELS OF ENERGY CONSUMPTION AND BANDWIDTH

In this section, we present a model of geographically distributed cloud computing system which supports the replication of data. The model focuses on the performance of cloud applications, utilization of communication resources and energy efficiency.

Most cloud applications rely on tight interaction with databases. Data queries can be fulfilled either locally or from the remote location. Data replication is used to ensure availability and reduce access delays.

Fig. 1 presents the timeline of a workload execution in data center. It begins with the user request arrival at the datacenter gateway. After being scheduled, it is forwarded through the data center network to the selected computing resource for execution. At the server, the workload can request data item if it is needed for its execution. For this, it queries a database and waits for the database reply to arrive. The database querying delay corresponds to the round-trip time and depends on the database location. As soon as the database reply is received, the workload execution is started. At the end of the execution, some workloads will send a modified data item back to the database for the update. As a result, the total delay associated with the workload execution in datacenters can be computed as follows:

$$d_{dc} = d_{req} + 2 \cdot d_{db} + d_{exec} + d_{update}, \quad (1)$$

where d_{req} is a time required for the workload description to arrive at the computing server, d_{db} is a one-way communication delay between the server and the database,

d_{exec} is a workload execution time which is defined by the size of the computing work of the workload and computing speed of the server, and d_{update} is the time required to update database.




Fig. 1. Workload execution timeline.

A more detailed description of communication-aware models for cloud computing workloads is available in [23].

A. Cloud Computing System Architecture

Large-scale cloud computing systems are composed of geographically distributed datacenters across the globe (see Fig. 2). The most widely used data center topology is the three tier fat tree [16], which consists of three layers of network switches: core, aggregation and access. The core layer provides packet switching backplane for all the flows going in and outside datacenter. The aggregation layer integrates connections and traffic flows from multiple racks. The access layer is where the computing servers, physically attached to the network, are arranged in racks.

Central database (Central DB) is located in the wide-area network and hosts all the data required by the cloud applications. To speed up database access and reduce access latency, each data center hosts a local database, called datacenter database (Datacenter DB), which is used to replicate the most frequently used data items from the central database. In addition, each rack hosts at least one server capable of running local rack-level database (Rack DB), which is used for subsequent replication from the datacenter database.

When data is queried, the information about requesting server, the rack, and the datacenter is stored. In addition, the statistics showing the number of accesses and updates are maintained for each data item.

A module called Replica Manager is located at the central database and periodically analyzes data access statistics to identify which data items are the most suitable for replication and at which replication sites. The availability of access and update statistics makes it possible to project data center bandwidth usage and energy consumption.

The following subsections present a model of the considered cloud computing system in terms of energy consumption, usage of network bandwidth and communication delays. The objective is to (a) minimize system-level energy consumption, (b) minimize utilization of network bandwidth and (c) minimize communication delays encountered in the data center network.




Fig. 2. Cloud computing datacenter.

B. Energy Consumption of Computing Servers

The power consumption of a server depends on its CPU utilization. As reported in [17] an idle server consumes about two-thirds of its peak power consumption. This is because servers must keep memory modules, disks, I/O resources and other peripherals operational even when no computations are performed. Then, the power consumption scales with offered CPU load according to the following equation [17]:

$$P_s(l) = P_{fixed} + \frac{(P_{peak} - P_{fixed})}{2} (1 + l - e^{-\frac{l}{a}}), \quad (1)$$

where P_{fixed} is an idle power consumption, P_{peak} power consumed at the peak load, l is a server load, and a is a utilization level at which the server attains asymptotic, i.e. close to linear power consumption versus the offered load. For most of the CPUs, $a \in [0.2, 0.5]$.

CPU power consumption is proportional to $V^2 f$, where V is voltage and f is an operating frequency. Voltage reduction requires frequency downshift. This implies a cubic relation from f . To account of it, Eq. (2) can be rewritten as follows:

$$P_s(l) = P_{fixed} + \frac{(P_{peak} - P_{fixed})}{2} (1 + l^3 - e^{-\frac{l^3}{a}}), \quad (2)$$

Eq. (3) forms the basis for DVFS power management, which can adjust operating frequency when server is underutilized to conserve operational power consumption [5].

C. Energy Consumption of Network Switches

Network switches are hardware devices that consist of the port transceivers, line cards, and switch chassis. All these components contribute to the switch energy consumption. According to [18] and [19], the power consumption of switch chassis and line cards remain constant over time, while the consumption of network ports can scale with the volume of the forwarded traffic as follows:

$$P_{switch} = P_{chassis} + n_c * P_{linecard} + \sum_{r=1}^R n_p^r * P_p^r * u_p^r, \quad (3)$$

where $P_{chassis}$ is a power related to switch chassis, $P_{linecard}$ is the power consumed by a single line card, n_c is number of line cards plugged into switch, P_p^r is a power drawn by a port running at rate r , n_p^r is number of ports operating at rate r and $u_p^r \in [0,1]$ is a port utilization which can be defined as follows:

$$u_p = \frac{1}{T} \int_t^{t+T} \frac{B_p(t)}{C_p} dt = \frac{1}{T * C_p} \int_t^{t+T} B_p(t) dt, \quad (4)$$

where $B_p(t)$ is an instantaneous throughput at the port's link at the time t , C_p is the link capacity, and T is a measurement interval.

D. Bandwidth Model

In this section, we analyze network capacity of data centers and bandwidth requirements of cloud applications that access database for different replication strategies.

An availability of per-server bandwidth is one of the core requirements affecting design of modern data centers. The most widely used three-tier fat tree topology (see Fig. 2) imposes strict limits on the number of hosted core, aggregation, and access switches as well as the number of servers per rack. For example, a rack switch serving 48 servers each connected with 1 Gb/s link has only two 10 Gb/s links in the uplink. As a result, its uplink bandwidth appears to be oversubscribed by a factor of $48 \cdot 1G / 20G = 2.4$, which also limits the per server available bandwidth to 416 Mb/s. Another bandwidth multiplexing occurs at the aggregation layer. An aggregation switch offers 12 ports to the access layer and is connected to all the core layer switches. For the three-tier architecture with 8-way Equal Cost Multipath Routing (ECMP), the oversubscription ratio at the aggregation layer is 1.5. This further reduces the per server bandwidth down to 277 Mbps for fully loaded connections.

Communications inside the datacenter can be broadly categorized to the uplink and downlink. The uplink flows are those directed from the computing servers towards the core switches. Conversely, the downlink flows are those from the core switches to the computing servers.

In the uplink, network bandwidth is used for propagating database requests and when applications need to update modified data items:

$$B_{ul} = N_{serv} (R_a S_{req} + R_u S_{data}), \quad (5)$$

where N_{serv} is the number of computing servers, S_{req} is the size of data request, and S_{data} is the size of the updated data item. R_a and R_u are data access and update rates respectively.

In the downlink, the bandwidth is used for sending job descriptions to computing servers for execution, receiving database objects and propagating data item updates between data replicas:

$$B_{dl} = N_{serv} \cdot R_a \cdot (S_{job} + S_{data}) + B_{rep}, \quad (6)$$

where S_{job} is the size of the job description, S_{data} is the size of the requested data object in bits, and B_{rep} is the bandwidth required to update all the replicas.

B_{rep} is different on different segments of the downlink. For the wide-area network it corresponds to the update between Central DB and Datacenter DBs

$$B_{rep,wan} = N_{serv} \cdot N_{dc} \cdot R_u \cdot S_{data}, \quad (7)$$

while for the network inside data center it corresponds to the update between Datacenter DBs and Rack DBs

$$B_{rep,dc} = N_{serv} \cdot N_{rack} \cdot R_u \cdot S_{data}, \quad (8)$$

where N_{dc} is the number of Datacenter DBs and N_{rack} is the number of Rack DBs in each data center.

Now, having computed the bandwidth required by running applications and their data base interactions, we can obtain residual bandwidth by subtracting it from the network capacity. It will be different for every tier of the data center network due to bandwidth oversubscription involved.

For a three-tier data center with N_{serv} servers, N_{acc} access, N_{agg} aggregation and N_{core} core switches, the corresponding network capacities at each tier can be obtained as follows:

$$BC_{access} = N_{serv} \cdot C_{access}, \quad (11)$$

$$BC_{agg} = 2 \cdot N_{access} \cdot C_{agg}, \quad (12)$$

$$BC_{core} = N_{agg} \cdot N_{core} \cdot C_{core}, \quad (13)$$

where C_{access} , C_{agg} and C_{core} are the capacities at the access, aggregation and core tiers respectively. Commonly, C_{access} is equal to 1 Gb/s, while C_{agg} and C_{core} correspond to 10 Gb/s links in modern datacenters.

The uplink capacity is always limited due to over subscription at lower layers. Therefore, the residual bandwidth in the downlink R_{dl}^l and in the uplink R_{ul}^l available at each tier of the network can be obtained as follows:

$$R_{dl}^l = BC_{dl}^l - B_{dl}, \quad (14)$$

$$R_{ul}^l = BC_{ul}^{l+1} - B_{ul},$$

where $l \in (access, agg, core)$ is an index indicating a tier level. The expression $l + 1$ refers to the tier located above the tier l .

At any moment of time the residual bandwidth left not in use in the data center can be computed as follows:

$$R_{dl} = \min(R_{dl}^{core}, R_{dl}^{agg}, R_{dl}^{access}), \quad (15)$$

$$R_{up} = \min(R_{ul}^{core}, R_{ul}^{agg}, R_{ul}^{access}). \quad (16)$$

E. Database Access and Energy Consumption

Having the model of energy consumption for computing servers (Section II-C) and network switches (Section II-D),

we can obtain total energy consumption of data center IT equipment as follows:

$$E_{dc} = \sum_{s=1}^S E_s + \sum_{k=1}^K E_k^{core} + \sum_{l=1}^L E_l^{agg} + \sum_{m=1}^M E_m^{access}, \quad (17)$$

where E_s is the energy consumed by a computing server s , while E_k^{core} , E_l^{agg} , E_m^{access} are the energy consumptions of k core, l aggregation, and m access switches respectively.

The load of individual servers is proportional to the workload execution and database query delays, therefore, energy consumption of server can be obtained as follows:

$$E_s = P_s(l) \cdot (2 \cdot d_{db} + d_{exec}) \cdot R_a \cdot T, \quad (18)$$

where $P_s(l)$ is a power consumed by the server executing a workload obtained according to Eq. (3), d_{db} is the time required to query and receive a data item from the database, d_{exec} is the workload execution time, R_a is an average database access rate, and T is a total time of the workload execution. The delay d_{db} depends on the database location and employed replication strategy. If data query is satisfied from replica databases, d_{db} becomes smaller, as propagation delay inside datacenter is in the order of microseconds. The delay associated with the database update is not included, as it becomes a job of the network to deliver the update after computing server becomes available for executing other tasks.

For network switches, energy consumption depends on the amount of traversing traffic and utilization of network ports (see Eq. (3)). Port utilization and traffic volumes are proportional to the size of job descriptions, data requests, data traffic, and data updates. Eqs. (5) and (6) allow computing traffic requirements in the uplink and the downlink respectively, while Eqs. (11), (12), and (13) define bandwidth capacity for each segment (access, aggregation, and core) of the network. Based on the aforementioned and by adapting Eq. (3), the energy consumption of the access switches can be computed as follows:

$$E_{access} = \left(P_{fixed} + \frac{N_{serv}}{N_{access}} \cdot P_p^{access} \cdot \frac{B_{dl}}{BC_{access}} + 2 \cdot P_p^{agg} \cdot \frac{B_{ul}}{BC_{access}} \cdot \frac{N_{access}}{N_{serv}} \right) \cdot T, \quad (19)$$

where P_{fixed} corresponds to the power consumption of the switch chassis and line cards, N_{serv}/N_{access} is the number of servers per rack, P_p^{access} and B_{dl}/BC_{access} are power consumption and port utilization of an access link, while P_p^{agg} and B_{ul}/BC_{access} are power consumption and port utilization of an aggregation network link.

Similarly, the energy consumption of the aggregation and core switches can be computed as follows:

$$E_{agg} = \left(P_{fixed} + 2 \cdot \frac{N_{access}}{N_{agg}} \cdot P_p^{agg} \cdot \frac{B_{dl}}{BC_{agg}} + N_{core} \cdot P_p^{core} \cdot \frac{B_{ul}}{BC_{core}} \right) \cdot T, \quad (20)$$

$$E_{core} = \left(P_{fixed}^{core} + N_{agg} \cdot P_p^{core} \cdot \frac{B_{dl}}{BC_{core}} \right) \cdot T, \quad (21)$$

where $2 \cdot N_{access}/N_{agg}$ is the number of aggregation switch links connected to racks, while P_p^{core} and B_{ul}/BC_{core} are the power consumption and port utilization of a core network link.

III. PERFORMANCE EVALUATION

In this section, we perform evaluation of the system model developed in Section II. The main performance indicators are: data center energy consumption, available network bandwidth and communication delay. The results obtained from modeling showed good agreement with simulations results obtained from the GreenCloud simulator [9]. The interested reader may refer to [20] for simulation results.

A. Scenario

Considering three-tier data center architecture presented in Fig. 2, we assume a uniform distribution of jobs among the computing servers as well as traffic in the data center network. Both computing servers and network switches implement DVFS [11] and DPM [10] power management techniques. Power consumption of communication ports can be adjusted in network switches based on the load of the forwarded traffic.

Table I summarizes data center setup parameters. The topology is comprised of 1024 servers arranged into 32 racks interconnected by 4 core and 8 aggregation switches. The network links interconnecting the core and aggregation switches as well as the aggregation and access switches are 10 Gb/s. The bandwidth of the access links connecting computing servers to the top-of-rack switches is 1 Gb/s. The propagation delay of all these links is set to 3.3 μ s. There is only one entry point to the datacenter through a gateway switch, which is connected to all the core layer switches with 100 Gb/s, 50 ms links.

TABLE I. DATACENTER TOPOLOGY

| Parameter | Value |
|--------------------------|----------------------|
| Gateway nodes | 1 |
| Core switches | 4 |
| Aggregation switches | 8 |
| Access (rack) switches | 32 |
| Computing servers | 1024 |
| Gateway link | 100 Gb/s, 50 ms |
| Core network link | 10 Gb/s, 3.3 μ s |
| Aggregation network link | 10 Gb/s, 3.3 μ s |
| Access network link | 1 Gb/s, 3.3 μ s |

Table II presents the power consumption profiles of data center servers and network switches. The server peak energy consumption of 301 W is composed of 130 W allocated for a peak CPU consumption and 171 W consumed by other devices like memory, disks, peripheral slots, mother board, fan, and power supply unit [17]. As the only component that scales with the load is the CPU power, the minimum

consumption of an idle server is bound and corresponds to 198 W.

TABLE II. POWER CONSUMPTION OF DATACENTER HARDWARE

| Parameter | Power Consumption [W] | | |
|-------------------------------------|-----------------------|------------|------|
| | Chassis | Line cards | Port |
| Gateway, core, aggregation switches | 1558 | 1212 | 27 |
| Access switches | 146 | - | 0.42 |
| Computing server | 301 | | |

Energy consumption of network switches is almost constant for different transmission rates as 85-97% of the power is consumed by switches' chassis and line cards, and only a small portion of 3-15% is consumed by the port transceivers. The values for power consumption are derived from [22].

B. Evaluation results

According to the model presented in Section II, the energy consumed by IT equipment is composed of the energy consumed by the computing servers as well as core, aggregation, and access switches. Energy consumption of the computing servers is presented in Fig. 3. The servers execute cloud applications that perform a certain amount of computing job and make a single database query for successful completion. The obtained energy consumption increases with the increase in server load. This is because energy is consumed during both phases, while doing computing work as well as while waiting for database data to arrive. The minimum querying time corresponds to the round-trip communication delay between the computing server and the database (see Fig. 1 for details). However, in real systems communication delays are larger and are the subject to queuing delays on congested links and protocol-related procedures which often delay transmissions while waiting for previously transmitted data to be acknowledged.




Fig. 3. Energy consumption of computing servers.

Unlike in the case of computing servers, the energy consumption of network switches is less sensitive to

variation of the forwarded traffic. It is mainly because only port level power consumption scales with the traffic load under DVFS power saving, while other hardware components, such as switch chassis and line cards, remain always active. Fig. 4, reports the obtained energy consumption levels of network equipment. The result suggests that devising power saving modes that shut down entire hardware components of a switch would allow substantial savings. However, it has to be noted that applying such kind of approaches will affect network connectivity and may result in system performance degradation, as datacenter load cannot be accurately predicted.




Fig. 4. Energy consumption of network switches.

Fig. 5 presents the system bandwidth requirements in downlink without database updates. Since the bandwidth is proportional to both the size of data item and access rate, it grows fast and becomes higher than the corresponding capacities at the core, aggregation, and the access segments of the datacenter network necessitating replication. Having only 100 Gb/s at the gateway link would trigger replication even for the small data items of less than 12 MB (or 8 Ethernet packets) for the access rate of 1 Hz requiring data replication from Central DB to the Datacenter DB in order to avoid the bottleneck. The bandwidth provided by the core network of 320 Gb/s will be exceeded with data items larger than 40 MB for the access rate of 1 Hz. Similarly, the bandwidth of the aggregation network of 640 Gb/s will be exceeded after 78 MB and will require additional data replication from Datacenter DB to Rack DBs. Finally, data size larger than 125 MB will cause congestion in the access segment of the network clearly indicating the limits.

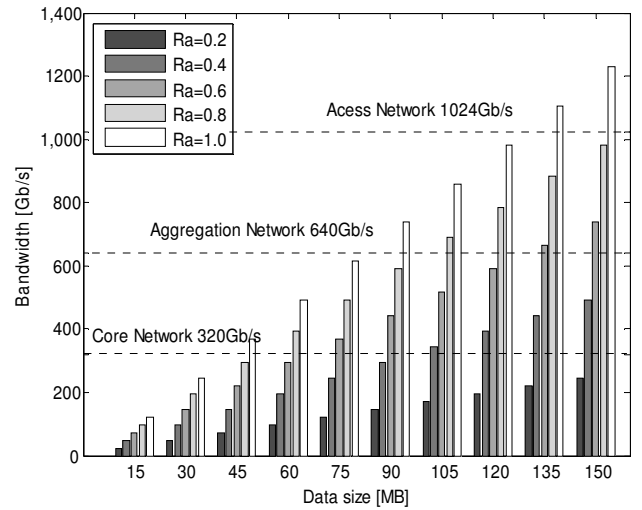


Fig. 5. Downlink bandwidth demand.

Fig. 6, reports data access delays measured as an average time elapsed from the moment of sending data request and having the requested data arrived. As expected, access delay becomes smaller for replicas located closer to servers and for all the replication scenarios an increase in the size of data objects increases data access delay.




Fig. 6. Data access delay.

IV. CONCLUSIONS AND FUTURE WORK

This paper reviews the topic of data replication in geographically distributed cloud computing data centers and presents models of energy consumption and bandwidth that can be used for efficient data replication strategy. In addition, optimization of communication delays leads to improvements in quality of user experience of cloud applications.

The evaluation of the proposed replication solution is based on the developed mathematical model and simulations using GreenCloud [9]. The obtained results confirm that replicating data closer to data consumers, i.e., cloud applications, can reduce energy consumption, bandwidth usage and communication delays substantially.

Future work on the topic will focus on developing a testbed implementation of the proposed solution.

ACKNOWLEDGEMENT

The authors would like to acknowledge the funding from National Research Fund, Luxembourg in the framework of ECO-CLOUD project (C12/IS/3977641).

REFERENCES

- [1] B. Hayes, "Cloud computing," *Magazine Communications of the ACM*, vol. 51, no. 7, pp. 9-11, 2008.
- [2] R. H. Katz, "Tech Titans Building Boom," *IEEE Spectrum*, vol. 46, no. 2, pp. 40-54, Feb. 2009.
- [3] Ruay-Shiung Chang, Hui-Ping Chang, and Yun-Ting Wang, "A dynamic weighted data replication strategy in data grids," *IEEE/ACS International Conference on Computer Systems and Applications*, pp. 414-421, April 2008.
- [4] Brown, R., et al.: Report to congress on server and data center energy efficiency: public law 109-431. Lawrence Berkeley National Laboratory, Berkeley, 2008.
- [5] Li Shang, Li-Shiuan Peh, and N. K. Jha, "Dynamic voltage scaling with links for power optimization of interconnection networks," *Ninth International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 91-102, February 2003.
- [6] Bin Lin, Shanshan Li, Xiangke Liao, Qingbo Wu, and Shazhou Yang, "eStor: Energy efficient and resilient data center storage," *2011 International Conference on Cloud and Service Computing (CSC)*, pp. 366-371, December 2011.
- [7] Xiaowen Dong, T. El-Gorashi, and J. M. H. Elmirghani, "Green IP Over WDM Networks With Data Centers," *Journal of Lightwave Technology*, vol. 29, no. 12, pp. 1861-1880, June 2011.
- [8] Fan Ping, Xiaohu Li, C. McConnell, R. Vabbalareddy, Jeong-Hyon Hwang, "Towards Optimal Data Replication Across Data Centers," *International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pp. 66-71, June 2011.
- [9] D. Kliazovich, P. Bouvry, and S. U. Khan, "GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers," *Journal of Supercomputing*, vol. 62, no. 3, pp. 1263-1283, 2012.
- [10] Shengquan Wang, Jun Liu, Jian-Jia Chen, and Xue Liu, "PowerSleep: A Smart Power-Saving Scheme With Sleep for Servers Under Response Time Constraint," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 1, no. 3, pp. 289-298, September 2011.
- [11] T. Horvath, T. Abdelzaher, K. Skadron, and X. Liu, "Dynamic voltage scaling in multitier web servers with end-to-end delay control," *IEEE Transactions on Computing*, vol. 56, no. 4, pp. 444-458, 2007.
- [12] Li Shang, Li-Shiuan Peh, and N. K. Jha, "Dynamic voltage scaling with links for power optimization of interconnection networks," *Ninth International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 91-102, February 2003.
- [13] Xiaowen Dong, T. El-Gorashi, and J. M. H. Elmirghani, "Green IP Over WDM Networks With Data Centers," *Journal of Lightwave Technology*, vol. 29, no. 12, pp. 1861-1880, June 2011.
- [14] Fan Ping, Xiaohu Li, C. McConnell, R. Vabbalareddy, Jeong-Hyon Hwang, "Towards Optimal Data Replication Across Data Centers," *International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pp. 66-71, June 2011.
- [15] S. Asur, B. A. Huberman, G. Szabo, and C. Wang, "Trends in social media: Persistence and decay," *5th International AAAI Conference on Weblogs and Social Media*, 2011.
- [16] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," *ACM SIGCOMM*, pp. 63-74, 2008.
- [17] X. Fan, W.-D. Weber, and L. A. Barroso, "Power Provisioning for a Warehouse-sized Computer", *ACM International Symposium on Computer Architecture*, pp. 13-23, San Diego, CA, June 2007.
- [18] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A power benchmarking framework for network devices," *8th International IFIP-TC 6 Networking Conference*, pp. 795 - 808, Aachen, Germany, May 2009.
- [19] P. Reviriego, V. Sivaraman, Z. Zhao, J. A. Maestro, A. Vishwanath, A. Sanchez-Macian, and C. Russell, "An energy consumption model for Energy Efficient Ethernet switches," *International Conference on High Performance Computing and Simulation (HPCS)*, pp. 98-104, July 2012.
- [20] D. Boru, D. Kliazovich, F. Granelli, P. Bouvry, A. Y. Zomaya, "Energy-Efficient Data Replication in Cloud Computing Datacenters," *Springer Cluster Computing*, 2014.
- [21] Intel Inc. (2010) Intel® Xeon® Processor 5000 Sequence, available at:http://www.intel.com/p/en_US/products/server/processor/xeon5000
- [22] N. Farrington, E. Rubow, and A. Vahdat, "Data center switch architecture in the age of merchant silicon," *17th IEEE symposium on high performance interconnects (HOTI '09)*, pp. 93-102, 2009.
- [23] D. Kliazovich, J. E. Pecero, A. Tchernykh, P. Bouvry, S. U. Khan, and A. Y. Zomaya, "CA-DAG: Communication-Aware Directed Acyclic Graphs for Modeling Cloud Computing Applications," *IEEE 6th International Conference on Cloud Computing (CLOUD)*, Santa Clara, CA, USA, 2013.