Simulating Communication Processes in Energy-Efficient Cloud Computing Systems

Dzmitry Kliazovich and Pascal Bouvry
University of Luxembourg
6 rue Coudenhove Kalergi, Luxembourg
dzmitry.kliazovich@uni.lu, pascal.bouvry@uni.lu

Samee Ullah Khan
North Dakota State University
Fargo, ND 58108-6050
samee.khan@ndsu.edu

Abstract— **Cloud computing data centers are becoming increasingly popular for the provisioning of computing resources. In this paper we present a simulation environment, named GreenCloud, for advanced energy-aware studies of cloud computing data centers. GreenCloud offers a detailed fine-grained modeling of the energy consumed by the elements of the data center, such as servers, switches, and links.**

Keywords—Cloud computing, simulator, energy efficiency, cloud communications, data centers

I. INTRODUCTION

Cloud computing is entering our lives and dramatically changing the way people parse information. Cloud provides platforms enabling a large variety of terminal devices owned by individuals to operate. In such an environment, computing, information storage, and communication becomes a utility. Cloud computing is an effective way to offer manageable and secure infrastructure with reduced cost of operations.

Cloud computing relies on the data centers as their primary backend computing infrastructure. Currently, over 500 thousand data centers are deployed worldwide [1]. The operation of large geographically distributed data centers requires a considerable amount of energy that accounts for a large slice of the total operational costs [2]. The Gartner group estimates that the energy consumption accounts for up to 10% of the current data center operational expenses (OPEX), and this estimate may rise to 50% in the next few years [3]. The cost of energy for running servers may already be greater than the cost of the hardware itself [4]. In 2010, data centers consumed about 1.5% of the world's electricity. In terms of CO_2 emissions, it corresponds to more than 50 million metric tons annually. Currently, roughly 40% of the total energy is consumed by the IT equipment [5]. Other systems contributing to the data center energy consumptions are cooling and power distribution systems that account for approximately 45% and 15% of total energy consumption, respectively.

There are two main alternatives for making data center consume less energy: (a) shutting it down or (b) scaling down its performance. The former method, commonly referred to as Dynamic Power Management (DPM) results in most of the savings as the average workload often stays below 30% in cloud computing systems. The latter corresponds to the Dynamic Voltage and Frequency Scaling (DVFS) technology that can adjust the hardware performance and power consumption to match the corresponding characteristics of the workload.

In this paper, we present an environment for simulation of energy-efficient cloud computing data centers. The main focus is devoted to communication processes in data centers which are simulated with a packet-level precision. Furthermore, the paper presents two scheduling solutions which exploit network awareness to provide efficient resource allocation.

II. THE GREENCLOUD SIMULATOR

A. Data Center Architectures

Three-tier trees of hosts and switches form the most widely used data center architecture. It consists of the core tier at the root of the tree, the aggregation tier that is responsible for routing, and the access tier that holds the pool of computing servers. Earlier data centers used two-tier architectures with no aggregation tier. However, such data centers, depending on the type of switches used and per-host bandwidth requirements, could typically support not more than 5,000 hosts.

In current data centers, rack connectivity is achieved with inexpensive top-of-tack switches. A typical top-of-rack switch shares two 10 GE uplinks with 48 GE links that interconnect computing servers within a rack. The difference between the downlink and the uplink capacities of a switch defines its oversubscription ratio, which in the aforementioned case is equal to 48/20 = 2.4:1. Therefore, under full load, only 416 Mb/s will remain available to each of the individual servers out of their 1 Gigabit links. At the higher layers of hierarchy, oversubscription ratios are around 1.5:1. This further reduces the available bandwidth to 277 Mbps.

Future data center architectures will be based on the principle of modular design. Traditional racks of servers will be replaced with standard shipping containers hosting 10 times as many servers as conventional data center in the same volume. Each container is optimized for power consumption. It integrates a combined water and air cooling system and implements optimized networking solutions. These containers, being easy to ship, can become plug-and-play modules in future roof-less data center facilities. To support the trend of modular design, distributed data center architectures, such as DCell, BCube, FiConn, and DPillar, will be adapted [6].

B. Simulator Structure

GreenCloud is a simulation environment for advanced energy-aware studies of cloud computing data centers [11]. It

978-1-4673-2798-5/12/$31.00

offers a detailed fine-grained modeling of the energy consumed by the elements of the data center, such as servers, switches, and links. Fig. 1 presents the structure of the GreenCloud simulator mapped onto the three-tier data center architecture.

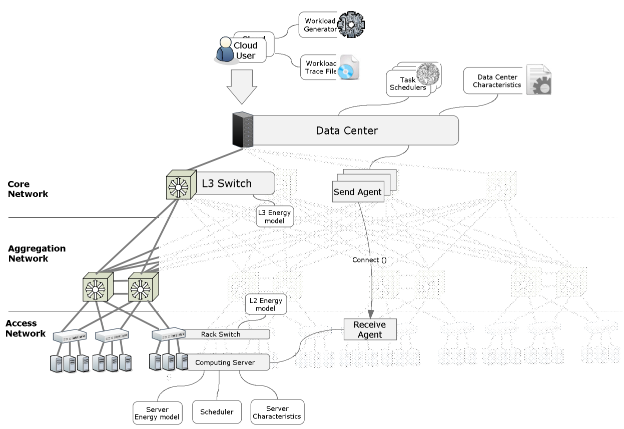

Figure 1. Architecture of the GreenCloud simulator.

C. *Energy Models*

In Greencloud, computing servers are implemented as single core nodes. According to the basic model, an idle server consumes around two-thirds of its peak load to keep memory, disks, and I/O resources running, while the rest of the power is consumed by the CPU and scaled with the offered computing load. In this work, we rely on a more detailed energy consumption model, which is based on the energy-consumption benchmarks for different severs from a number of manufactures [7]:

$$P(l) = P_{idle} + \frac{P_{peak} - P_{idle}}{2}\left(1 + l - e^{-\frac{l}{\tau}}\right), \qquad (1)$$

where $P(l)$ is the energy consumed by the server operating at the load level $l \in [0,1]$, P_{idle} and P_{peak} are a server consumptions at the minimum and the maximum load levels respectively, τ is a scaling coefficient.

Energy consumption of network switches depends on the: (a) type of switch, (b) number of ports, (c) port transmission rates, and (d) employed cabling solutions. The energy consumed by a switch can be generalized by the following [8]:

$$P_{switch} = P_{chassis} + n_{linecards} \cdot P_{linecard} + \sum_{i=0}^{R} n_p \cdot P_r, \qquad (2)$$

where $P_{chassis}$ is the power consumed by the switch base hardware, $P_{linecard}$ is the power consumed by an active linecard, and P_r corresponds to the power consumed by an active port (transmitter) running at the rate r.

III. Network-Aware Scheduling and Resource Allocation

Job scheduling is at the heart of the successful operation of power management in data centers. As most of the energy savings come from the hardware standby mode, job schedulers follow workload consolidation policy by maximizing the load on the operational computing servers and increase the number of idle servers that can be put into the "sleep" mode. Such scheduling policy works well in systems that can be abstracted as a homogenous pool of computing servers. However, being implemented in a data center with cloud applications requiring communication, the scheduler should tradeoff workload concentration with the load balancing of network traffic.

A. *DENS scheduler*

The DENS (Data Center Energy-Efficient Network-aware Scheduling) methodology minimizes the total energy consumption of a data center by selecting the best-fit computing resources for job execution based on the load level and communication potential of data center components. The communicational potential is defined as the amount of end-to-end bandwidth provided to individual servers or group of servers by the data center architecture.

Fig. 2 presents a DENS metric used for server selection. DENS favors servers located in racks with a light load of the network links. It also penalizes the selection of idle servers and increases the probability of selection of servers operating at the moderate load level. This strategy allows maximizing the number of idle servers which can be turned off to save power.

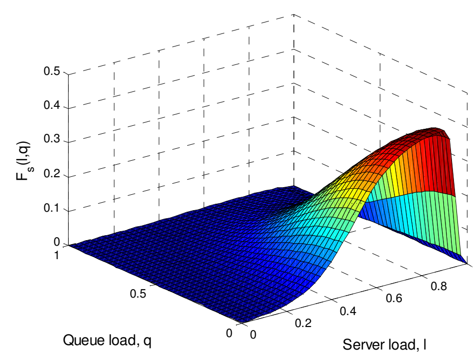

Figure 2. Server selection by DENS metric according to its load and communicational potential.

A detailed description of the DENS scheduler and associated resource allocation policies is provided in [9].

B. *e-STUB scheduler*

The **e**nergy-efficient **s**cheduler for cloud computing applications with **tra**ffic load **b**alancing (e-STAB) treats communicational demands of the jobs equally important to that of the computing requirements. e-STAB is a scheduler aiming to: **(a)** balance communication flows produced by the jobs and **(b)** consolidate jobs on a minimum amount of the computing servers. As network traffic can be highly dynamic and often difficult to predict, the e-STAB scheduler analyses both the load on the network links and the occupancy of outgoing queues at the network switches. e-STAB allocates jobs favoring network resources that offer the most of the available bandwidth and penalizes resources for which the load approaches the available transmission capacity when the traffic queues growing in size. Queuing analysis aids in preventing a buildup of network congestion. Such techniques are already implemented in several transport-layer protocols that estimate buffer occupancy of the network switches and can react before congestion related losses occur.

The e-STAB scheduling policy can be defined with the following two steps executed for every incoming cloud computing data center workload:

- **Step 1:** Select a group of servers S connected to the data center network with the highest available bandwidth, provided that at least one of the servers in S can accommodate the computational demands of the scheduled job. The available bandwidth is defined as an unused capacity of the link or a set of links connecting the group of servers S to the rest of the data center network.
- **Step 2:** Within the selected group of servers S, select a computing server with the smallest available computing capacity, but sufficient to satisfy the computational demands of the scheduled task.

Fig. 3 presents the e-STAB metric used for the selection of server racks based on the observed link load and buffer occupancy.

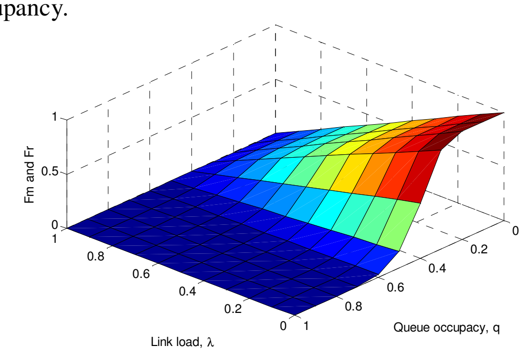

Figure 3. Selection of racks and modules by the STAB scheduler.

A detailed description of the e-STAB scheduler and associated resource allocation policies is provided in [10].

IV. Simulation Results

The main focus of the GreenCloud simulator [11, 12] is on the fine-grained modeling of computing and communication processes in modern data centers. Three types of workloads exist to model typical cloud computing applications, such as web browsing, instant messaging, or content delivery. The computationally intensive workloads model HPC applications. They load computing servers considerably, but require almost no data transfers in the interconnection network of the data center. The data-intensive workloads, on the contrary, produce almost no load at the computing servers, but require heavy data transfers. The balanced workloads aim to model the applications having both computing and data transfer requirements.

System performance and energy efficiency are two main sets of metrics reported after the simulation. Fig. 4 presents a screenshot with an example of GreenCloud simulation results.

A more detailed description of the GreenCloud simulator is available in [11].

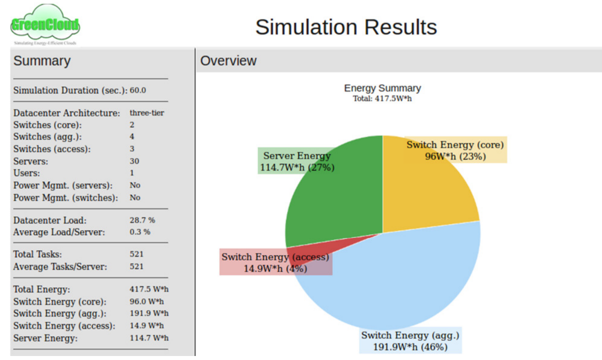

Figure 4. GreenCloud simulation results.

V. Conclusions`

This paper presents a simulation environment for advanced energy-aware studies of cloud computing data centers, named GreenCloud. GreenCloud offers a detailed fine-grained modeling of the energy consumed by the elements of the data center, such as servers, switches, and links, as well as assures a packet-level precision in modeling of communication processes.

Acknowledgments

The authors would like to acknowledge the funding from National Research Fund, Luxembourg in the framework of GreenIT project (C09/IS/05) 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] "State of the Data Center 2011," Emerson Network Power, 2011.

[2] R. Raghavendra at el. "No "Power" Struggles: Coordinated Multi-level Power Management for the Data Center," APLOS 2008.

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

[4] A. Vasan and A. Sivasubramaniam "Worth their watts?-an empirical study of datacenter servers," IEEE 16th International Symposium on High Performance Computer Architecture (HPCA), 2010.

[5] R. Brown et al. "Report to congress on server and data center energy efficiency: public law 109-431," Lawrence Berkeley National Laboratory, Berkeley, 2008.

[6] R. Boutaba at el. "Data Center Network Virtualization: A Survey," IEEE Communications Surveys & Tutorials, vol. PP, no. 99, pp. 1-20, 2012.

[7] Server Power and Performance characteristics, available on: http://www.spec.org/power_ssj2008/

[8] P. Mahadevan at el. "A Power Benchmarking Framework for Network Devices," IFIP-TC 6 Networking Conference, Aachen, Germany, May 11 - 15, 2009.

[9] D. Kliazovich, P. Bouvry, and S. U. Khan, "DENS: Data Center Energy-Efficient Network-Aware Scheduling," Cluster Computing, 2012.

[10] D. Kliazovich, S. T. Arzo, F. Granelli, P. Bouvry, and S. U. Khan, "e-STAB: Energy-Efficient Scheduling for Cloud Computing Applications with Traffic Load Balancing," submitted for publication to Parallel & Distributed Processing Symposium, 2013.

[11] D. Kliazovich, P. Bouvry, and Samee 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.

[12] The GreenCloud simulator, available at http://greencloud.gforge.uni.lu.