Energy-Aware Multi-Agent Server Consolidation in Federated Clouds

Alessandro Ferreira Leite and Alba Cristina Magalhaes Alves de Melo

Department of Computer Science
University of Brasilia, Brasilia, Brazil

Abstract. In this paper, we propose and evaluate a server consolidation approach for efficient power management in virtualized federated Data Centers. The main goal of our approach is to reduce power consumption, trying to meet QoS requirements with limited energy defined by a third party agent. In our model, we address application workload considering the costs due to turning servers on/off and Virtual Machine migrations in same Data Center and between different Data Centers. Our simulation results with 2 data centers and 400 simultaneous Virtual Machines show that our approach is able to reduce more than 50% of energy consumption, while still meeting the QoS requirements.

1 Introduction

Cloud Computing is a recent paradigm for provision of computing infrastructure, platform and/or software. This paradigm shifts the location of these components to the Internet in order to reduce costs associated with resource management (hardware and software) [10].

Cloud Computing is gaining popularity since it helps companies to reduce costs and carbon footprint. Usually, services are executed in big Data Centers containing a large number of computing nodes. The energy requirements of the whole Data Center have a high impact on the total operation costs [11], which can be over 60% of the peak load [8, 15]. Therefore, reducing the energy consumption without sacrificing Quality-of-Service (QoS) requirements is an important issue.

In Cloud Computing, Data Centers usually employ virtualization techniques to provide computing resources as utilities and Virtual Machine (VM) technologies for server consolidation. Server consolidation is the process of gathering several virtual machines into a single physical server. It aims at minimizing the number of physical servers required to host a group of Virtual Machines.

Many studies have been conducted to provide power reduction and some of them are based in server consolidation [1]. However, server consolidation in Cloud Computing can introduce some difficulties such as: (i) the Cloud environment must usually provide Quality of Service (QoS) guarantees, normally defined in terms of Service Level Agreements (SLA); (ii) it is common to occur dynamic changes of the incoming requests rate; (iii) the usage pattern of the resources is often unpredictable and (iv) different users have distinct preferences.

M. Yousif, L. Schubert (Eds.): Cloudcomp 2012, LNICST 112, pp. 72–81, 2013.

Currently, with the energy costs increasing, the focus shifts from optimizing Data Center resource management for pure performance to optimize it for energy efficiency while maintaining the level of services [2].

In a Cloud Computing environment, there are distinct participants with distinct objectives, preferences and disposition to pay for services. In this scenario, a Multi-agent System (MAS) can be used where each participant is an autonomous agent that incorporates market and negotiation capabilities.

In this work, we propose a Federated Application Provision (FAP) strategy which uses multiple agents and server consolidation techniques to achieve power-aware resource allocation, by taking into account SLAs, energy consumption and carbon footprint. In our approach, the user should pay according to the efficiency of his/her applications in terms of resource utilization and power consumption. Therefore, we propose that the price paid by the users should increase according to the whole energy consumption of the Data Center, especially when the user does not accept to negotiate QoS requirements.

Experimental results for our FAP consolidation strategy were obtained in the CloudSim [3] simulator, with 2 Data Centers, each one belonging to a different Cloud, and 400 simultaneous virtual machines show that our approach is able to reduce an average of 53.57% of energy consumption, while meeting the SLA requirements.

The remainder of this paper is organized as follows. Section 2 presents concepts of Cloud Computing. Section 3 discusses energy green performance indicators. The proposed strategy for Federated Cloud server consolidation is presented in Section 4. In Section 5, experimental results are discussed. Section 6 presents related work. Finally, Section 7 presents the conclusion and future work.

2 Cloud Computing

There are many definitions of cloud computing in literature. Most of these definitions state that a Cloud Computing system should have (i) pay-per-use capability, (ii) elastic capacity and the illusion of infinite resources, (iii) self-service, (iv) virtualized resources and (v) QoS enhancement functionality. The cloud service models are divided in three classes, according to the abstraction level and the service model of the providers: Infrastructure-as-a-Service (IaaS), Plataform-as-a-Service (PaaS), and Software-as-a-Services (SaaS) [17].

In the Infrastructure-as-a-Service (IaaS) model, the user can request processing power, storage, network and other fundamental computing resources such as the operating system, for a period of time and pay only what he/she uses. Plataform-as-a-Service (PaaS) are development platforms that allow the creation of applications with supported programming languages and tools hosted in the cloud and accessed through a browser. This model can slash development time, offering readily available tools and services. In the Software-as-a-Service (SaaS) model, applications run on the Cloud infrastructure and are accessible from various client devices. From the user view, the SaaS model allows him/her to save money in servers and software licenses.

Cloud Computing can be viewed as a combination of many preexisting technologies such as virtualization and server consolidation, among others.

The term virtualization refers to the abstraction of compute resources (CPU, memory, I/O) from the applications, aiming to improve sharing and utilization of computer systems. One immediate benefit of virtualization is the option to run multiple operating systems and software stacks on a single physical platform.

Server consolidation is the process of gathering several Virtual Machines (VMs) into a single physical server. It is often used by Data Centers to increase resource utilization and reduce electric power consumption costs [11]. For example, consider a set of VMs $\{vm_1, vm_2, vm_3, vm_4\}$ and a set of hosts $\{h_1, h_2, h_3\}$, each of them with a quad-core processor, where each processor is capable of executing one VM. A power efficient allocation schedule could initially assign all the VMs to the same host in such a way that the other hosts could be put in the power-saving state or turned off. A possible solution to this problem would be, therefore, to pack the maximum workload in the smallest number of servers, keeping each resource (CPU, disk, network, among others) on every server at 100% utilization and put the idle servers in power-saving state.

The consolidation process can be performed in a single step using the peak load demands, known as static consolidation, or in a dynamic manner, by reevaluating periodically the workload demand in each VM. In static consolidation, VMs stay in the same physical server during their whole lifetime. The utilization of the peak load demand should ensure that the VM does not overload. However, in a dynamic environment with different access patterns, one or more VMes can become idle, resulting in an inefficient power allocation.Dynamic consolidation aims to tackle this problem by taking into account the current workload demands. Dynamic consolidation may require migrating VMs between physical servers in order to [7]: (i) pull out physical servers from an overload state or (ii) turn off a physical server when it is idle or when the VMs mapped to it can be moved to another physical server.

Server consolidation in a Cloud Computing environment presents some additional difficulties since the Cloud must also provide reliable QoS, normally defined in terms of Service Level Agreements (SLA). An SLA describes characteristics such as maximum throughput and minimum response time, that must be delivered by the deployed systems. If an SLA is violated, economical penalties usually apply.

3 Energy-Aware Computing

Cloud Computing solutions may have a potential impact on green house gas (GHC), which include CO_2 emissions. Saving energy of a Data Center with acceptable QoS requirements is an economical incentive for data center operators, as well as a significant contribution to the environment. This requires the design of energy-aware solutions. Energy-awareness can be characterized by taking into account the amount of resources and QoS requirements required by the applications and also the energy requirements along their life cycle [9].

Several indicators have been introduced to measure Data Center efficiency under the vision of achieving economical, environmental and technological sustainability [13]. In this context, Green Performance Indicators (GPI) are defined as the driving policies for data collection and analysis related do energy consumption. The idea of GPIs is interesting because it can be adapted as part of Service Level Agreements (SLA), where requirements about energy efficiency versus the expected quality of services are specified and need to be satisfied. The GPIs are classified in four clusters (IT Resource Usage, Application Lifecycle, Energy Impact and Organizational). In this work, we consider only the IT Resource Usage and the Energy Impact GPI clusters.

The IT Resource Usage GPIs characterize the energy consumption of an application as a function of the energy consumed by its resources. Examples of metrics are *CPU usage*, *Memory usage* and *I/O activity*.

The Energy Impact GPIs describe the impact of Data Centers and applications on the environment, considering power supply, consumed materials, emissions, and other energy related factors. The most important Energy Impact GPI metrics are: a) application performance indicators, which measure the *energy consumption per computing unit*, using typically FLOPS/kWh or Number of Transactions/kWh; b) *Data Center Infrastructure Efficiency (DCiE)*, which is used to determine the energy efficiency of a Data Center as a whole; and c) *Compute Power Efficiency (CPE)*, which computes the data center power. In this metric, the power consumed by idle servers is computed as overhead.

4 Design of the Multi-Agent Consolidation Mechanism

The main goal of our approach is to meet the QoS requirements of the applications, while keeping the power consumption of the Data Centers below a given energy threshold defined by a third party agent. To achieve this goal, we propose a Multi-Agent strategy to negotiate the resource allocations among Clouds.

We consider a federated Cloud environment with four distinct agents: Cloud Service Provider (CSP), Cloud User (CLU), Energy Power Provider (EPP) and Carbon Emissions Agency Regulator (CEAR) as shown in Figure 4(a). In our design, the CEAR determines the amount of carbon emissions that both the CSP and the EPP can emit in a period of time.

In each Data Center, there is one coordinator responsible for monitoring data center metrics, negotiating with the other agents. There are also sensors to monitor energy consumption, resource usage and SLA violation as shown in Figure 4(b). The scenario proposed is a set of Data Centers composed of a set of Virtual Machines, which are mapped to a set of physical servers that are interconnected and deployed in a hybrid cloud model.

Let R $\{r_1, r_2, \cdots, r_n\}$ be the set of resources in Data Center i with a capacity c_i^k, where $k \in R$. The Energy Consumption for the Data Center (E_i) is defined as $E_i = (p_{max} - p_{min}) * U_i + p_{min}$ [14], where p_{max} is the power consumption at the peak load, p_{min} is the minimum power consumption in active mode, and U is the resource utilization of Data Center i as defined in $U_i = \sum_{j=1}^{n} u_{i,j}$ [14], where $u_{i,j}$ is the resource usage of resource j in Data Center i.

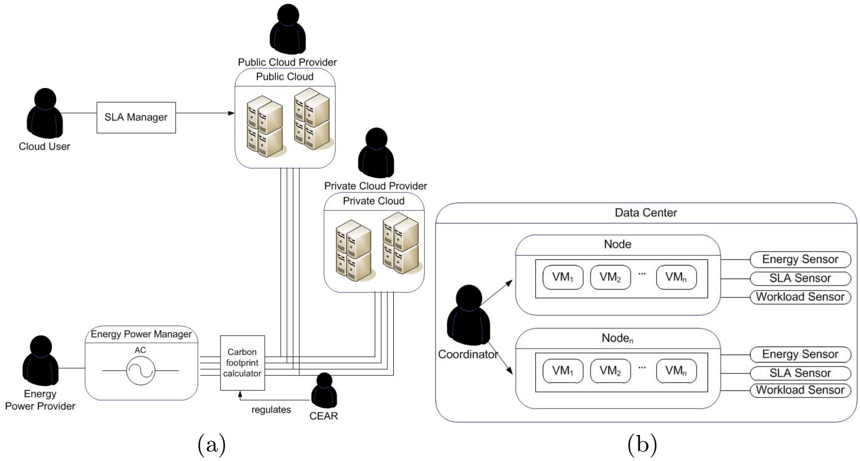

Fig. 1. (a) Agents of the cloud market (b) Detail view of a Data Center

Resources R are managed by the Cloud Service Providers (CSP), which are used by a set of Cloud Users (CLU). The energy power is provided by an Energy Power Provider (EPP). The relation between the CSPs and a CLU is determined by a set of QoS metrics described in SLAs. The Data Center is subject to an energy consumption threshold agreed among the CSP, the EPP and CEAR. When the energy consumption threshold is violated, this implies in additional costs. To calculate the carbon footprint of the CSP and the EPP, the CEAR uses the following metrics: *CPU usage, Memory usage, I/O activity* and *CPE* (Section 3).

Let T represent the set of tasks to be executed in a resource r_i which is subject to a set of QoS constraints. The following steps are executed:

1. When a task t_i is submitted, the Cloud Provider calculates the price of t_i's execution.
2. The Cloud Provider tries to place t_i in an appropriate resource, using consolidation techniques to reduce the number of physical servers.
3. If the Cloud Provider does not have enough available resources or the energy threshold will be violated, the Cloud Provider first contacts another Cloud Provider and negotiates with it the execution of this task. In this case, the price of this execution (P_t) is defined as shown in Equation (1).

$$P_t = E_t + \epsilon_t + \lambda_t \qquad (1)$$

   where ϵ_t is the cost Energy Impact of task t on the environment, and λ_t is the cost to transfer task t to another Cloud Provider.
4. If it does not succeed, the Cloud Provider tries to consolidate its VMs considering the tasks SLAs.
5. If not possible, it tries to negotiate the energy threshold with the CEAR and with the EPP agents.

6. If all negotiations fail, the Cloud Provider finds the SLA whose violation implies in lower cost and execute the task. In this case, the price to execute the tasks is defined as shown in Equation (2).

$$V_t = P_t + \gamma + \delta \tag{2}$$

   where γ is the cost of violate the QoS requirements of other tasks and δ is the cost associated with energy consumption violation.

To control task allocation, each Cloud Provider has a matrix representing tasks $t_i \in T$, virtual machines $vm_j \in VM$ and physical servers $r_z \in R$, where: 1 represents that t_i is allocated at vm_j in resource r_z; 0 indicates that t_i can be allocated at vm_j; and -1 represents that this allocation is impossible.

In order to illustrate our strategy, consider a federated Cloud with 2 Data Centers (DC1 and DC2) and a user that contracts DC1 to execute his applications. Consider that DC1 is overloaded and that the QoS requirement described in the SLAs is response time. In this scenario, when the user submits tasks to execute, the DC1 Cloud Provider first tries to execute them locally, considering energy consumption and the available resources. Since DC1 is overloaded, its Cloud Provider contacts DC2 and negotiates with it the execution of the tasks. If DC2 accepts it, the cost of the tasks execution is calculated with Equation (1). If DC2 does not accept, then DC1 tries to consolidate its virtual machines and, if not possible, it tries to negotiate the energy threshold with the CEAR and the EPP agents. If all negotiations fail, then DC1 finds the SLA whose violations implies in lower cost and terminates the execution of its associated task. In this case, the cost to execute the tasks is calculated with Equation (2).

5 Experimental Results

In this section, we present the evaluation of the strategy proposed in Section 4. We used CloudSim [3], which is a well-established Cloud simulator that has been used in many previous works [16], [18], among others. It is a simulation toolkit that enables modeling and simulation of Clouds and application provisioning environments, with support to Data Centers, Virtual Machines and resource provisioning policies.

Since we are dealing with federated environments, we extended CloudSim by adding four classes (*CloudEnergyReg, DCEnergySensor, FedPowerVMAllocPolicy, CustomerDCBroker*) to it, as well as isolation of queue events and support for concurrent execution.

The *CloudEnergyReg* class represents the behavior of the CEAR agent. This agent communicates with the Data Center cloud coordinator to inform the energy consumption threshold. The *DCEnergySensor* class implements the Sensor interface that monitors the energy consumption of the Data Center and informs the coordinator. When the energy consumption is close to the limit, this sensor creates an event and notifies the coordinator, that can take actions. The *FedPowerVmAllocPolicy* class extends the *VmAllocPolicy* class to implement the proposed federated server consolidation mechanism Virtual Machine allocation to hosts that can belong to different Data Centers. Finally, the *CustomerDCBroker* class models the QoS requirements customer behavior, negotiates with the cloud coordinator and requests computations.

5.1 Evaluation in Two Scenarios

In order to evaluate the effectiveness of our Federated Application Provisioning strategy (FAP), we used a simulation setup that is similar to the one used in [3]. The simulation environment included 2 Data Centers (DC1 and DC2), with 100 hosts each. These hosts had one CPU core with 1000 MIPS, 2GB of RAM and 1TB of storage. The workload model included provisioning for 400 VMs, where each VM requested one CPU core, 256 MB of RAM and 1GB of storage. The CPU utilization distribution was set to the Poisson distribution, where task required 150 MIPS or 10 minutes to complete execution. We assumed CPU utilization of 20, 40, 60, 80 and 100% and a global energy consumption threshold of 3 kWh of energy per data center. Initially, the provisioner allocates as many as possible virtual machines on a single host, without violating any constraint of the host. The SLA was defined in terms of response time (10 minutes).

In the first evaluation scenario, there are two Data Centers (DC1 and DC2) and tasks are always submitted to DC1. If DC1 becomes overloaded, VMs are migrated from DC1 to DC2. The simulation was repeated 10 times and the mean values for energy consumption without our mechanism using only DC1 (trivial), and with our Federated Application Provision strategy (FAP) mechanism are presented in Figures 2 (a), (b) and (c).

Figure 2(a) shows that the proposed provision technique is able to reduce the total power consumption of the Data Centers, without SLA violation. In this case, an average reduction of 53.37% in power consumption was achieved since DC1 consumed more than 9kWh with the trivial approach and no more than 4.9 kWh was consumed by both Data Centers with our approach (2.92 kWh for DC1 and 1.98 kWh for DC2). In order to achieve this, DC1 tried first to maximize the usage of its resources and to consume the limit of energy power without violating the SLAs. DC2 was used only when DC1 was overloaded, if DC1 was in the imminence of SLA violation or when the energy consumption was close to the limit.

Figure 2(b) presents the number of VM migrations when our mechanism is used. It can be seen that the number of migrations decreases as the threshold of CPU usage increases. This result was expected since with more CPU capacity, the allocation policy tends to use it and allocate more VMs in the same physical machine. In Figure 2(c), we measured the wallclock time needed to execute 400 tasks, with our mechanism (FAP) and without our mechanism (trivial). It can be seen that FAP increases the whole execution time. This occurs because of the overhead caused by the VMs migrations between data centers, and the negotiations between the CLU and the CSP agents. Nevertheless, this increase in less than 22%, since the wallclock execution time without and with the mechanism is 21.5 min and 27.4 min, respectively, for 100% CPU utilization. We consider

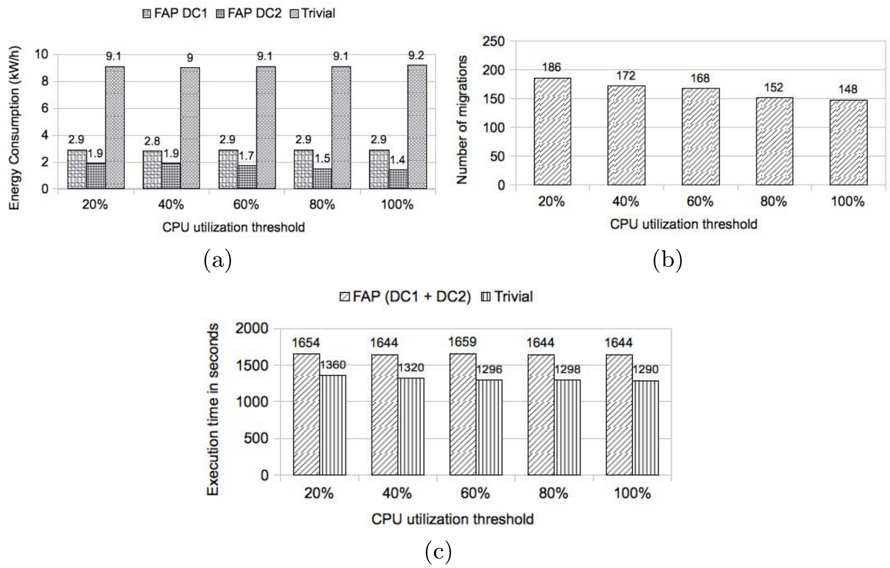

Fig. 2. Case Study 1:(a) Total energy consumption by data centers (b) Number of VM migrations from DC1 to DC2 for the FAP mechanism (c) Execution time of tasks

that this increase in the execution time is very low and it is compensated by the reduction in the power consumption (Figure 2).

In the second scenario, we consider two users, with distinct SLAs and each user makes 400 task execution requests to a different data center (DC1 and DC2). Our goal is to observe the rate of SLA violation when the workload of both Data Centers is high (Figures 3 (a), (b), (c) and (d)).

In Figure 3(a), we can see that, even in a scenario with overloaded Data Centers, our mechanism is able to maintain the power consumption below the threshold (3 kWh) for each Data Center. With the CPU utilization threshold of 80%, the power consumption decreased from 9.13 kWh to 5.65 kWh (DC1 + DC2), reaching 38.2% of reduction in power consumption.

The number of SLA violations with two overloaded Data Centers was lower than the one obtained with one overloaded Data Center (DC1) (Figure 3(d)). With the CPU utilization threshold of 80%, the SLA violation decreased from 43.94% (DC1) to 31.48% (DC1 + DC2), reaching 12.46% of reduction in SLA violations. This shows the appropriateness of VM migrations between different Data Centers in an overloaded scenario.

6 Related Work

Table 1 summarizes the main characteristics of 6 papers that propose server consolidation strategies for distributed environments.

As can be seen in this Table, three approaches [4, 6, 5] use multi-agent systems to reduce power consumption and costs. One of them is targeted to Clouds and

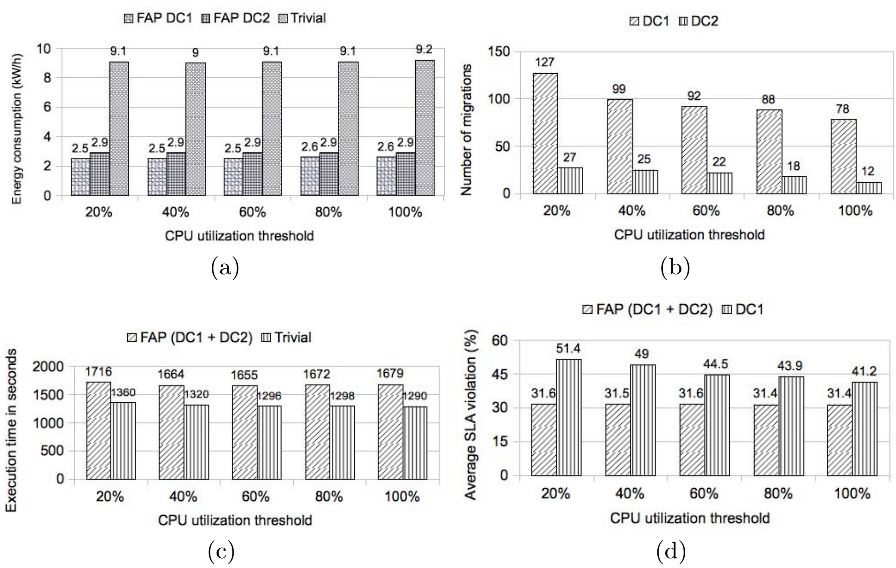

Fig. 3. Case Study 2: (a) Total energy consumption by data centers (b) Number of VM migrations (c) Execution time of tasks with 2 overloaded Data Centers (d) Average SLA Violation for the federated approach

Table 1. Comparative summary of Cloud server consolidation strategies

Paper	Target	Power-Aware	Federated	Multi-agent	Migration	SLA
[4]	Cloud	No	No	Yes	No	No
[6]	Cluster	No	No	Yes	No	No
[5]	Cluster	Yes	No	Yes	No	Yes
[12]	Cluster	Yes	No	No	No	Yes
[14]	Cloud	Yes	No	No	No	No
[3]	Cloud	Yes	No	No	Same DC	Yes
This work	Cloud	Yes	Yes	Yes	Among DCs	Yes

two execute in cluster computing environments. Three approaches [12, 14, 3] reduce power consumption in cloud computing considering SLAs. None of the analyzed proposals consider federation cloud environment nor VM migration among data centers.

7 Final Consideration and Future Work

In this paper, we proposed and evaluated a server consolidation approach for efficient power management in virtualized data centers taking into account energy consumption, and QoS requirements. The results obtained in the CloudSim [3] simulator, with 2 data centers and 400 simultaneous virtual machines show that very good energy consumption savings are obtained with our approach, while meeting the QoS requirements. The best gain (53.57%) was obtained when we have one overloaded data center. In this case, we were able to reduce the energy consumption from 9.13 kW/h with the trivial approach to 4.9 kW/h with an increase of less than 22% in the execution time.

As future work, we intend to investigate formal models for power-aware resource allocation in Cloud Computing Systems and propose extensions that take more parameters into consideration.

Acknowledgment. This work is partially supported by FAPDF/Brazil and CNPq/Brazil.

References

[1] Beloglazov, A., et al.: A taxonomy and survey of energy-efficient data centers and cloud computing systems. Advances in Computers 82, 49 (2010)
[2] Buyya, R., et al.: Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility. Future Gener. Comput. Syst. 25, 599–616 (2009)
[3] Calheiros, R.N., et al.: Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. Software: Practice and Experience 41(1), 23–50 (2011)
[4] Chen, Y., Yeh, H.: An implementation of the multiagent system for market-based cloud resource allocation. J. Computing 2(11), 27–33 (2010)
[5] Das, R., et al.: Autonomic multi-agent management of power and performance in data centers. In: AAMAS, pp. 107–114 (2008)
[6] Ejarque, J., Sirvent, R., Badia, R.M.: A multi-agent approach for semantic resource allocation. In: CloudCom, pp. 335–342 (December 2010)
[7] Ferreto, T.C., et al.: Server consolidation with migration control for virtualized data centers. Future Gener. Comput. Syst. 27(8), 1027–1034 (2011)
[8] Fan, X., Weber, W., Barroso, L.A.: Power provisioning for a warehouse-sized computer. In: ISCA, pp. 13–23. ACM (2007)
[9] Ferreira, A., et al.: Energy-aware design of service-based applications. In: ISOCC, pp. 99–114 (2009)
[10] Hayes, B.: Cloud computing. Commun. ACM 51, 9–11 (2008)
[11] Hoelzle, U., Barroso, L.A.: The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines. M. C. Pub. (2009)
[12] Kim, K.H., et al.: Sla-based scheduling of bag-of-tasks applications on power-aware cluster systems. IEICE Trans. on Inf. Sys. E93-D(12), 3194–3201 (2010)
[13] Kipp, A., Jiang, T., Fugini, M., Salomie, I.: Layered green performance indicators. Future Gener. Comput. Syst. 28(2), 478–489 (2012)
[14] Lee, Y.C., Zomaya, A.Y.: Energy efficient utilization of resources in cloud computing systems. The Journal of Super Computing, 1–13 (2010)
[15] Lefurgy, C., Wang, X., Ware, M.: Server-level power control. In: Fourth Autonomic Computing. IEEE Computer Society (2007)
[16] Shi, Y., Jiang, X., Ye, K.: An energy-efficient scheme for cloud resource provisioning based on cloudsim. In: IEEE ICCC, pp. 595–599 (2011)
[17] Vaquero, L.M., et al.: A break in the clouds: towards a cloud definition. SIGCOMM Comput. Commun. Rev. 39, 50–55 (2008)
[18] Zhang, Q., Gürses, E., Boutaba, R., Xiao, J.: Dynamic resource allocation for spot markets in clouds. In: Hot-ICE 2011, pp. 1–6 (2011)