

Energy and Carbon-Efficient Placement of Virtual Machines in Distributed Cloud Data Centers

Atefeh Khosravi, Saurabh Kumar Garg*, and Rajkumar Buyya

Cloud Computing and Distributed Systems (CLOUDS) Laboratory
Department of Computing and Information Systems
The University of Melbourne, Australia
atefehk@student.unimelb.edu.au, {sgarg,raj}@csse.unimelb.edu.au

Abstract. Due to the increasing use of Cloud computing services and the amount of energy used by data centers, there is a growing interest in reducing energy consumption and carbon footprint of data centers. Cloud data centers use virtualization technology to host multiple virtual machines (VMs) on a single physical server. By applying efficient VM placement algorithms, Cloud providers are able to enhance energy efficiency and reduce carbon footprint. Previous works have focused on reducing the energy used within a single or multiple data centers without considering their energy sources and Power Usage Effectiveness (PUE). In contrast, this paper proposes a novel VM placement algorithm to increase the environmental sustainability by taking into account distributed data centers with different carbon footprint rates and PUEs. Simulation results show that the proposed algorithm reduces the CO₂ emission and power consumption, while it maintains the same level of quality of service compared to other competitive algorithms.

Keywords: Cloud computing, Data center, Energy efficiency, Carbon footprint, Virtual machine placement.

1 Introduction

The information and communication technology industry (ICT) consumes an increasing amount of energy and most of it is consumed by data centers [8]. A major consequence of this amount of energy consumption by data centers is a significant increase in ecosystem carbon level. According to Gartner, the ICT industry produces 2% of global CO₂ emission, which places it on par with the aviation industry [22]. Therefore, reducing even a small fraction of the energy consumption in ICT, results in considerable savings in financial and carbon emission of the ecosystem.

Cloud computing offers a wide range of services and applications to its users. Three main services that Clouds provide are infrastructure, platform, and software as a service. Infrastructure as a service (IaaS) allows users to run their applications in form of virtual machines (VMs) on a shared infrastructure. Cloud

* Currently working in IBM Research, Melbourne, Australia.

data centers take advantage of virtualization technology [7] to share a physical server's resources among multiple VMs. Each VM has its own characteristics and depending on the resource usage, it consumes energy and leaves carbon footprint. By the arrival of each VM request, the Cloud manager selects the physical resource to instantiate the request. VM placement in Cloud computing system is a complex task and if cannot be done effectively, it leads to high energy usage and high carbon footprint.

Thus, wisely taking into account parameters that affect VM placement and physical server selection result in less energy consumption and less carbon footprint. Distributed Cloud data centers, alongside with bringing high availability and disaster recovery, provide the opportunity to have different energy sources. *Carbon footprint rate* of energy sources is an important parameter, since data centers use electricity driven by these sources to run VMs. By having different energy sources in different data center sites or within a data center site, Cloud providers should increase the use of more clean and off-grid renewable energies [24]. *Power usage effectiveness (PUE)* is coined by the Green Grid consortium [14] and indicates the energy efficiency of a data center. PUE is a ratio of total power consumed by the data center to its power consumed by IT devices. Providers can consider PUE as a parameter to perform VM placement among different data center sites. *Proportional power* is another parameter that can be taken into account for VM placement. Server proportional power has a cubic relation with CPU frequency [17]. Therefore, considering the increase in CPU frequency, which is related to increase in CPU utilization upon new request arrival, will have a great impact on the amount of energy consumption in data centers.

This paper proposes a VM placement algorithm by considering distributed Cloud data centers with the objective of minimizing carbon footprint. Our proposed Cloud computing system, Energy and Carbon-Efficient (ECE) Cloud architecture, benefits from distributed Cloud data centers with different carbon footprint rates, PUE value, and different physical servers' proportional power. ECE Cloud architecture places VM requests in the best suited data center site and physical server. The main contributions of this paper are: an Energy and Carbon-Efficient Cloud architecture, based on distributed Cloud data centers; an efficient VM placement algorithm that integrates energy efficiency and carbon footprint parameters; a comprehensive comparison on carbon footprint and power consumption for different VM placement algorithms with respect to quality of service (number of rejected VMs).

The remainder of the paper is organised as follows. In Section 2 the related work is discussed. Section 3 presents the proposed Cloud architecture with its components, VM placement algorithm, and formulates the objective. The performance evaluation results and the experimental environment are presented in Section 4. Section 5 concludes the paper and presents future works.

2 Related Work

There is a growing body of literature that aims to reduce the amount of carbon dioxide of Cloud services in data centers. Most of the works in this area focus

on reducing the energy consumption in a single data center or considering the data center hardware aspects [6] [5]. Well-known technologies that data centers benefit from by applying virtualization technology [7] are VMs migration [15] and consolidation [23]. The main problem with migration and consolidation is that they are complex and, due to the need for resuming and suspending VMs cause overhead to the system [10]. Moreover, these technologies act reactive whereas applying preventive technologies are more efficient.

As idle servers consume almost half of the power when they are in the peak power state [4], work by Lin et al. [18] uses a dynamic right-sizing on-line algorithm to predict the number of active servers that are needed for the arriving workload to the data center. Based on their experiments, dynamic right-sizing can achieve significant energy savings in the data center, but it requires servers to have different power levels and be able to transit to different states. A similar work done by Lefevre et al. [16] proposes Green Open Cloud (GOC) architecture, with advance resource reservation for users to improve the prediction of the arrival requests.

The above mentioned technologies are adopted within a data center and intend to reduce the energy consumption, whilst they do not particularly consider carbon emission. Reducing data center energy consumption does not necessarily lead to reduce in carbon footprint. Works by Aksanli et al. [3] and Goiri et al. [12] consider the availability of both non-polluting (green) and polluting (brown) energy sources in a single data center. They use prediction-based scheduling algorithms to increase usage of green energy sources.

Liu et al. [19] consider reducing the carbon footprint of data centers by considering multiple data center sites. They proposed an algorithm to efficiently use the renewable energies, such as wind and solar, in different places. This algorithm uses the idea of geographic diversity of data center sites and unpredictability of renewable energies to find the optimal percentage of wind/solar energies in order to reduce the brown energy consumption. Garg et al. [11] also consider reducing carbon footprint of Cloud data center sites. They proposed a novel carbon-aware green Cloud architecture, which uses two directories for Cloud providers to register their offered services.

Our work is different from the previous works, since we address the problem of increase in carbon footprint of the Cloud data centers by performing efficient VM placement. Our proposed method accommodates VM requests by considering distributed data center sites of a Cloud provider, with various energy sources and carbon footprint rates. Moreover, we consider data centers' PUE, physical servers' proportional power usage, and user VM requests of different types. Finally, we present an energy and carbon-efficient algorithm that uses two level decision making for VM placement.

3 System Model

In this section, Energy and Carbon-Efficient (ECE) Cloud architecture is described. This architecture assures system quality of service, while minimizing

the Cloud carbon footprint by applying an energy and carbon-efficient heuristic for VM placement.

3.1 ECE Cloud Architecture

The proposed architecture is represented in Figure 1. The system consists of the following components and symbols used in this paper are presented in Table 1:

Table 1. Description of Symbols

Symbol	Description	Symbol	Description
d	Number of Data Center Sites	P	Proportional Power
c	Number of Clusters at each Data Center	P_{fixed}	Server Power Consumption in Idle State
h	Number of Hosts at each Cluster	P_f	Server Power Consumption at Frequency f
cf	Data Center/Cluster Carbon Footprint Rate	f_u	CPU Operating Frequency at Utilization u
CF	Cloud Total Carbon Footprint	$type$	Virtual Machine Instance Type
ht	Virtual Machine Holding Time	$core, pu$	CPU Cores and Total Processing Unit
$ram, storage$	RAM and Storage	bw	Network Bandwidth

Cloud Users: Cloud Users send their VM requests based on predefined requirements to the Cloud provider. Virtual machine types and configurations are inspired by Amazon Elastic Compute Cloud (EC2) [1]. The expected requirements for each VM are specified by its predefined configurations in terms of required number of cores, processing unit of each core, storage, RAM, and network bandwidth. In addition, holding time of a VM depends on the application runs on that VM. We consider two types of applications in this paper: bag-of-tasks and web-requests. Every requested VM by users has the following requirements: $(ht, type)$, where each $type$ consists of the following components: $\{cores, pu, ram, storage, bw\}$. Cloud computing system load at time t , according to the running VMs, is represented as:

$$load = \sum_{i=1}^d \sum_{j=1}^c \sum_{k=1}^h vm_{(i,j,k,t)}.$$

Cloud Provider: A Cloud provider has several geographically distributed data center sites. Each data center is composed of several clusters with various heterogeneous physical servers. Physical servers are characterised by CPU cores, CPU processing unit, amount of RAM, storage, and network bandwidth. In addition to the physical servers configuration, each data center has its own energy-related parameters, shown by PUE and proportional power. Moreover, each data center can have one or more energy sources with different carbon footprint rates.

ECE Cloud Information Service: Each data center site registers its characteristics in the ECE Cloud information service (ECE-CIS) and they keep their information updated. This information includes available physical resources and




Fig. 1. Energy and Carbon-Efficient (ECE) Cloud Architecture

energy related parameters; such as data center PUE, energy source(s), carbon footprint rate, and physical servers’ current utilization and power consumption. Cloud broker uses this information to perform ECE VM placement in Cloud computing environment.

ECE Cloud Broker: ECE Cloud Broker is the Cloud provider’s interface with Cloud users. It receives user requests and schedules them based on their predefined requirements. Despite users request scheduling, broker should also ensure energy efficient data centers with minimum carbon footprint for Cloud providers.

Resources on the Cloud provider are physical servers in the clusters within each data center. The broker receives the current status of data centers’ physical resources and their energy information from ECE-CIS, and based on this information, assigns the VM to a physical server in a data center site. Based on [25], in today’s Internet and core networks design, average number of hops a packet passes from source to destination is between 12-14 hops. Therefore, we can have data center site selection without considering network distance; especially for sites that are located in a region, such as different states in USA, as we considered in this paper.

3.2 Placement Decision

As stated before, the broker makes the placement decision based on the data centers’ power usage effectiveness (PUE), energy sources carbon footprint rate, and proportional power.

The PUE indicates the energy efficiency of a data center and is a metric to compare different data center designs in terms of electricity consumption. Data center’s PUE is calculated as:

$$PUE_i = \frac{\text{Datacenter}_i \text{TotalPowerConsumption}}{\text{Datacenter}_i \text{ITDevicesPowerConsumption}},$$

where the total power consumption is sum of power drawn by cooling, lightening, and IT equipment. PUE is a value larger than 1 ($PUE \geq 1$). PUE of 1.0 means 100% of the data center's electricity goes to the IT part and is ideal for any data center, but is unattainable pragmatically. In other words, the smaller the PUE, the more energy efficient the data center.

Data center proportional power is the next important metric in physical server selection. According to the experiments by Lien et al. [17] server's power consumption depends on the system base power and the CPU frequency, and the CPU frequency itself depends on the CPU utilization. The data center proportional power, also known as dynamic power, is calculated as: $P = P_{fixed} + P_f \times f_u^3$. The power consumption for a VM on physical server k in cluster j of data center i at time t is modeled as: $P(vm_{(i,j,k,t)})$.

According to the above mentioned metrics the objective is to minimize total carbon footprint of the Cloud provider, CF , for time interval $[0, T]$, and is computed as follows:

$$CF = \sum_{t=1}^T \sum_{i=1}^d (PUE_i \times \sum_{j=1}^c (cf_j \times \sum_{k=1}^h (P(vm_{(i,j,k,t)}) \times ht))),$$

subject to following constraints:

$$\begin{aligned} \sum_{i=1}^d \sum_{j=1}^c \sum_{k=1}^h vm_{(i,j,k)}^{core} &\leq host_{(i,j,k)}^{core}, & \sum_{i=1}^d \sum_{j=1}^c \sum_{k=1}^h vm_{(i,j,k)}^{pu} &\leq host_{(i,j,k)}^{pu}, \\ \sum_{i=1}^d \sum_{j=1}^c \sum_{k=1}^h vm_{(i,j,k)}^{ram} &\leq host_{(i,j,k)}^{ram}, & \sum_{i=1}^d \sum_{j=1}^c \sum_{k=1}^h vm_{(i,j,k)}^{storage} &\leq host_{(i,j,k)}^{storage}. \end{aligned}$$

The above mentioned constraints ensure that allocated resources to the VMs on a physical server do not exceed the total capacity of the server.

3.3 Energy and Carbon-Efficient (ECE) Heuristic for VM Placement

By the arrival of each VM request the broker has $(d \times c \times h)$ different VM placement options. The VM placement problem can be seen as a bin-packing problem with different bin sizes (physical servers). Therefore, we propose the Energy and Carbon-Efficient (ECE) VM placement algorithm (Algorithm 1), which is a derivation of the best-fit heuristic to place the VMs in the data center, cluster, and host with the minimum carbon footprint, PUE, and minimum increase in physical server's power consumption.

Algorithm 1. Energy and Carbon-Efficient (ECE) VM Placement Algorithm

Input: *datacenterList*, *clusterList*, *hostList*
Output: *destination*

```

1 while vmRequest do
2   Get data centers' Information from ECE-CIS;
3   foreach datacenter in datacenterList do
4     Add clusterList into aggregateClusterList;
5   Sort aggregateClusterList in an ascending order of  $(PUE \times cf)$ ;
6   foreach cluster in aggregateClusterList do
7     foreach host in hostList do
8        $P_1 \leftarrow$  Get current hostDynamicPower;
9        $P_2 \leftarrow$  Calculate hostDynamicPower after initiating the vm;
10       $\Delta P \leftarrow P_2 - P_1$ ;
11     Sort hostList in an ascending order of  $\Delta P$ ;
12     foreach host in hostList do
13       if host is suitable for vm then
14          $destination \leftarrow (datacenter, cluster, host)$ ;
15         return destination;
16    $destination \leftarrow null$ ; //rejection of request;
17   return destination;
```

The broker receives a VM request and selects the best physical server for the VM. Its objective is to minimize the data centers' carbon footprint and accordingly power consumption. Therefore, ECE placement algorithm gets data centers' resources and energy status from ECE-CIS, upon the arrival of a new VM request (Line 2). According to the received information, ECE adds the clusters of all the data centers into an aggregated cluster list (Lines 3-4), and sorts the new list based on the minimum $(PUE \times cf)$ (Line 5). By receiving the data centers and clusters status, ECE calculates the amount of power consumption that will be added to each host after initiating the new VM (Lines 8-10). Afterwards, ECE sorts the hosts list based on the estimated ΔP (Line 11), and if the host has enough resources for the VM (Line 13), it submits the VM to the selected data center, cluster, and host.

In order to show the time complexity of Algorithm 1, we consider v VM requests. Line 3-4 take $O(d)$, and the sort function at Line 5 can be done in $O(dc \log(dc))$. Lines 7-9 need $O(h)$ time, and the sort function for hosts at Line 11 needs $O(h \log(h))$ to be done. Lines 12-15 take $O(h)$, in the worst case. Thus, the total running time of the algorithm is $O(v(d + dc \log(dc) + dc(h + h \log(h) + h)))$. Since there are small number of data center sites and clusters (dc) for a Cloud provider, the complexity of this algorithm is dominated by the number of VM requests and hosts sort function. The total time complexity of the algorithm is $O(vdch \log(h))$.

4 Performance Evaluation

We use the CloudSim toolkit [9] to evaluate the Cloud computing virtualized environment. We have extended CloudSim to enable energy and carbon-efficient VM placement simulations. Apart from being aware of data center's PUE, carbon footprint rate, and dynamic power, the extended package has the ability to simulate dynamic VM requests with different instance types.

In order to evaluate the proposed algorithm, we modeled an IaaS provider with 4 data center sites, and each site with 90 heterogeneous physical servers. Each data center has a unique PUE value and 2 clusters with different carbon footprint rates. Table 2 shows data centers' PUE value and carbon footprint rate for different group of clusters. The PUE value is based on the work by Greenberg et al. [13], and is in the range [1.56, 2.1]. Data centers' carbon footprint rates, the last column of Table 2, are extracted based on the information from US Department of Energy, Appendix F, Electricity Emission Factors [2]. In this simulation, we use 5 different physical servers whose characteristics are given in Table 3. According to the linear relationship between CPU utilization and frequency, and the cubic relation between CPU frequency and system proportional power, the following is the power models for the platforms:

$$\begin{aligned} \text{CPU Frequency(in GHz): } \{f(u) : (1.4, 1.57, 1.74, 1.91, 2.08, 2.25, 2.42, 2.6, 2.77, 2.94, 3.11)\} \\ \text{Power Model1(in Watt): } \{P_f : (60, 63, 66.8, 71.3, 76.8, 83.2, 90.7, 100, 111.5, 125.4, 140.7)\} \\ \text{Power Model2(in Watt): } \{P_f : (41.6, 46.7, 52.3, 57.9, 65.4, 73, 80.7, 89.5, 99.6, 105, 113)\} \end{aligned}$$

VM characteristics are inspired by Amazon EC2 instance types given in Table 4. The physical resources to the VMs are allocated based on the VM resource requirements and all the VMs are considered to perform at the maximum utilization during their lifetime. The VM type and the number of VMs requested by users depend on the user type (bag-of-tasks or web-requests), and are based on the related probabilities. The VM type related probability is shown in the last column of Table 4 and is derived from the work by Mills et al. [21].

In order to generate the workload, we need VM requests arrival rate and holding time. The Lublin-Feitelson [20] workload model is employed to generate the bag-of-tasks VM requests. We take benefit of Lublin to set arrival request parameters, including simulation duration, number of requests, requests arrival time, and request holding time. We consider each generated request in Lublin as a VM request. In order to generate VMs with longer holding time, we increased the first parameter of the Gamma distribution and left other Lublin parameters with their default value. To generate the web-requests, we use the same arrival time model of bag-of-tasks requests, and for the holding time we use a hyper gamma distribution with expectation value 73 and variance 165. For both workloads, we omit 5% of created requests at the start (warm-up period) and end (cool-down period) of the simulation to get a steady environment. We apply 240-hour long workload with different number of requests. Finally, for the purpose of accuracy, each experiment is repeated 30 times and the mean is reported for measured values for experimental results.

Table 2. Data Centers Characteristics

Data Center Site	PUE	Carbon Footprint Rate (Tons/MWh)
DC1 -Oregon, USA	1.56	0.124, 0.147
DC2 -California, USA	1.7	0.350, 0.658
DC3 -Virginia, USA	1.9	0.466, 0.782
DC4 -Dallas, USA	2.1	0.678, 0.730

Table 3. Platform Types Characteristics

Platform Type	Number of Cores	Core Speed (GHz)	Memory (GB)	Storage (GB)	Network Bandwidth (Mbps)	Bits	Power Model
Platform1	2	2	16	2000	1000	B32	PowerModel1
Platform2	4	4	32	6000	1000	B64	PowerModel1
Platform3	8	4	32	7000	2000	B64	PowerModel2
Platform4	8	8	64	7000	4000	B64	PowerModel2
Platform5	8	16	128	9000	4000	B64	PowerModel2

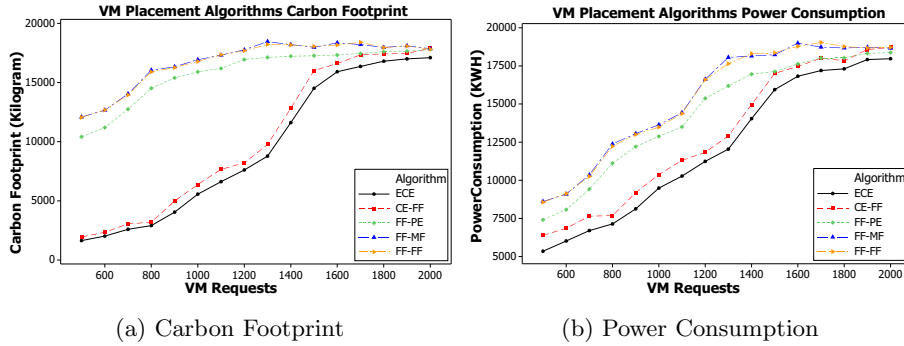
4.1 Results

We use the described workload data to compare the proposed VM placement algorithm with respect to carbon and power efficiency with four competing algorithms. The first algorithm is a version of ECE, that its data center and cluster selection is same as ECE, and uses first-fit bin-packing for host selection. We refer to this algorithm as Carbon-Efficient First-Fit (CE-FF). The other group of algorithms are three bin-packing heuristics that use first-fit heuristic for data center/cluster selection, without considering carbon footprint parameters. First-Fit Power-Efficient (FF-PE) uses power-efficient policy for host selection (same as ECE host selection). First-Fit Most-Full First (FF-MF) selects the physical server with least available resources. Finally, the last algorithm uses first-fit heuristic for data center, cluster, and host selection (FF-FF).

Figure 2a illustrates the carbon footprint of ECE in comparison with other placement algorithms under different number of VM requests. Based on the experiments, as the number of VMs increases, the system utilization increases as well to the point that system performs with highest utilization and reaches to the saturation point. Therefore, increase in system load leads to increase in the total carbon footprint in data centers. Based on the Figure 2a, ECE in comparison to CE-FF (carbon-efficient) and other heuristics (non carbon-efficient) reduces carbon footprint with an average of 10% and 45% respectively. The same behaviour can be seen for the data centers' power consumption in Figure 2b. The ECE algorithm has lower power consumption in comparison to the other algorithms and consumes on average 8% and 20% less power than CE-FF and other heuristics placement algorithms respectively. Considering the differences between algorithms behaviour in both figures, we can infer that just considering power-efficient parameters is not enough to reduce the data centers' carbon footprint. However, taking into consideration data centers' energy and carbon rate parameters, at the same time, leads to significant reduction in terms of Cloud computing system carbon footprint and consumed power.

Table 4. VM Types and Simulated User Types; (Bag-of-Task Users (BT) and Web-Request Users (WR))

VM Type		Number of Cores	Core Speed (GHz)	Memory (MB)	Storage (GB)	Network Bandwidth (Mbps)	Bits	Probability and User Type
Standard Instances	M1Small	1	1	1740	160	500	B32	0.25-BT
	M1Large	2	4	7680	850	500	B64	0.12-WR
	M1XLarge	4	8	15360	1690	1000	B64	0.08-WR
High Memory Instances	M2XLarge	2	6.5	17510	420	1000	B64	0.12-WR
	M22XLarge	4	13	35020	850	1000	B64	0.08-WR
High CPU Instances	C1Medium	2	5	1740	320	500	B32	0.1-BT

**Fig. 2.** Comparison of ECE Algorithm with other VM Placement Algorithms**Table 5.** SLA Violation for Different VM Placement Algorithms

VM Placement Algorithm	SLA Violation Under Different VM Requests					
	1000	1200	1400	1600	1800	2000
ECE	0.0%	0.05%	0.4%	2.9%	8.6%	13.0%
CE-FF	0.0%	0.0%	0.3%	0.7%	6.0%	11.4%
FF-PE	0.0%	0.0%	0.3%	2.5%	9.4%	15.3%
FF-MF	0.0%	0.0%	0.2%	2.5%	9.7%	15.3%
FF-FF	0.0%	0.0%	0.1%	2.6%	9.7%	15.3%

Table 5 shows the SLA violation under different system loads for different VM placement algorithms. It shows that, the SLA violation (number of rejected VMs) under low system load for ECE is slightly higher than the other algorithms. However, by increasing system load, ECE will have lower SLA violation. Overall, all the VM placement algorithms have close values for violation, while ECE considerably reduces carbon footprint and power consumption.

5 Conclusion and Future Work

In this paper, the problem of VM placement to reduce Cloud computing energy consumption and carbon footprint is investigated. We used ECE Cloud

information service (ECE-CIS), as part of next generation Cloud computing environment. ECE-CIS obtains energy and carbon related information from data centers and enables the broker to carry out carbon and power-efficient VM placement. We introduced the energy and carbon-efficient (ECE) VM placement algorithm, and compared it with a carbon-efficient algorithm (CE-FF) and three other heuristic algorithms (FF-PE, FF-MF, FF-FF). We performed the simulations by extending CloudSim and used different VM instance types with different holding times for the system workload. Based on the experiment results, ECE can on average save up to 10% and 45% carbon footprint in the ecosystem in comparison to CE-FF and three other heuristics respectively, while keeping SLA violation level as the same. Moreover, ECE reduces the power consumption in data centers by an average of 8% and 20% in comparison to CE-FF and other three algorithms respectively; which illustrates the importance of considering data centers' carbon footprint rate and PUE to reduce Cloud computing carbon footprint.

In the future, we plan to study the impact of different user applications and VM holding times on the VM placement policies. Moreover, we want to explore the effect of inter-data centers network distance and data locality on the Cloud computing system carbon footprint.

Acknowledgments. We would like to thank Amir Vahid Dastjerdi, Adel Nadjaran Toosi, Rodrigo N. Calheiros, and Deepak C Poola for their help and constructive suggestions.

References

1. Amazon EC2 instance types, <http://aws.amazon.com/ec2/instance-types/>
2. US Department of Energy, Appendix F, Electricity Emission Factors, <http://www.eia.doe.gov/oiaf/1605/pdf/Appendix>
3. Aksanli, B., Venkatesh, J., Zhang, L., Rosing, T.: Utilizing green energy prediction to schedule mixed batch and service jobs in data centers. In: Proc. of the 4th Workshop on Power-Aware Computing and Systems, pp. 5:1–5:5. ACM (2011)
4. Barroso, L., Holzle, U.: The case for energy-proportional computing. *IEEE Computer* 40(12), 33–37 (2007)
5. Beloglazov, A., Buyya, R., Lee, Y., Zomaya, A.: A taxonomy and survey of energy-efficient data centers and cloud computing systems. *Advances in Computers* 82(2), 47–111 (2011)
6. Berl, A., Gelenbe, E., Di Girolamo, M., Giuliani, G., De Meer, H., Dang, M., Pentikousis, K.: Energy-efficient cloud computing. *The Computer Journal* 53(7), 1045–1051 (2010)
7. Brey, T., Lamers, L.: Using virtualization to improve data center efficiency. *The Green Grid, Whitepaper* 19 (2009)
8. Brown, R., et al.: Report to congress on server and data center energy efficiency, pp. 109–431. Public law (2008)
9. Calheiros, R., Ranjan, R., Beloglazov, A., De Rose, C., Buyya, R.: 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)

10. Clark, C., Fraser, K., Hand, S., Hansen, J., Jul, E., Limpach, C., Pratt, I., Warfield, A.: Live migration of virtual machines. In: Proc. of the 2nd Conference on Symposium on Networked Systems Design & Implementation, vol. 2, pp. 273–286. USENIX Association (2005)
11. Garg, S., Yeo, C., Buyya, R.: Green cloud framework for improving carbon efficiency of clouds. In: Proc. of the 17th International Conference on Parallel Processing, pp. 491–502 (2011)
12. Goiri, I., Beauchea, R., Le, K., Nguyen, T., Haque, M., Guitart, J., Torres, J., Bianchini, R.: Greenslot: scheduling energy consumption in green datacenters. In: Proc. of the International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 20:1–20:11. ACM (2011)
13. Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., Myatt, B.: Best practices for data centers: Lessons learned from benchmarking 22 data centers. In: Proc. of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE, August 3, vol. 87, pp. 76–87 (2006)
14. Haas, J., Froedge, J., Pflueger, J., Azevedo, D.: Usage and public reporting guidelines for the green grids infrastructure metrics, pue/dcie (2009)
15. Harney, E., Goasguen, S., Martin, J., Murphy, M., Westall, M.: The efficacy of live virtual machine migrations over the internet. In: Proc. of the 2nd International Workshop on Virtualization Technology in Distributed Computing. ACM (2007)
16. Lefèvre, L., Orgerie, A.: Designing and evaluating an energy efficient cloud. *The Journal of Supercomputing* 51(3), 352–373 (2010)
17. Lien, C., Bai, Y., Lin, M.: Estimation by software for the power consumption of streaming-media servers. *IEEE Transactions on Instrumentation and Measurement* 56(5), 1859–1870 (2007)
18. Lin, M., Wierman, A., Andrew, L., Thereska, E.: Dynamic right-sizing for power-proportional data centers. In: Proc. of the IEEE INFOCOM, pp. 1098–1106. IEEE (2011)
19. Liu, Z., Lin, M., Wierman, A., Low, S., Andrew, L.: Geographical load balancing with renewables. *ACM SIGMETRICS Performance Evaluation Review* 39(3), 62–66 (2011)
20. Lublin, U., Feitelson, D.: The workload on parallel supercomputers: modeling the characteristics of rigid jobs. *Journal of Parallel and Distributed Computing* 63(11), 1105–1122 (2003)
21. Mills, K., Filliben, J., Dabrowski, C.: Comparing vm-placement algorithms for on-demand clouds. In: Proc. of the Third International Conference on Cloud Computing Technology and Science, pp. 91–98. IEEE (2011)
22. Pettey, C.: Gartner estimates ict industry accounts for 2 percent of global co2 emissions (2007)
23. Srikantaiah, S., Kansal, A., Zhao, F.: Energy aware consolidation for cloud computing. In: Proc. of the 2008 Conference on Power Aware Computing and Systems, p. 10. USENIX Association (2008)
24. Stewart, C., Shen, K.: Some joules are more precious than others: Managing renewable energy in the datacenter. In: Proc. of the Workshop on Power Aware Computing and Systems (2009)
25. Van Mieghem, P.: Performance analysis of communications networks and systems. Cambridge University Press (2006)