

Green Approach for Joint Management of Geo-Distributed Data Centers and Interconnection Networks

Amine Barkat^{1,2}  · Mohand-Tahar Kechadi³ · Giacomo Verticale¹ · Ilario Filippini¹ · Antonio Capone¹

Received: 20 December 2016 / Accepted: 6 November 2017 / Published online: 14 November 2017
© Springer Science+Business Media, LLC, part of Springer Nature 2017

Abstract Every time an Internet user downloads a video, shares a picture, or sends an email, his/her device addresses a data center and often several of them. These complex systems feed the web and all Internet applications with their computing power and information storage, but they are very energy hungry. The energy consumed by Information and Communication Technology (ICT) infrastructures is currently more than 4% of the worldwide consumption and it is expected to double in the next few years. Data centers and communication networks are responsible for a large portion of the ICT energy consumption and this has stimulated in the last years a research effort to reduce or mitigate their environmental impact. Most of the approaches proposed tackle the problem by separately optimizing the power consumption of the servers in data centers and of the network. However, the Cloud computing infrastructure of most providers, which includes traditional telcos that are extending their offer, is rapidly evolving toward geographically distributed data

✉ Amine Barkat
amine.barkat@polimi.it

Mohand-Tahar Kechadi
tahar.kechadi@ucd.ie

Giacomo Verticale
giacomo.verticale@polimi.it

Ilario Filippini
ilario.filippini@polimi.it

Antonio Capone
antonio.capone@polimi.it

- ¹ Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB), Politecnico di Milano, Milan, Italy
- ² Department of Computer Science, Faculty of Exact Sciences, University of Bejaia, Bejaia, Algeria
- ³ School of Computer Science and Informatics, University College Dublin, Dublin, Ireland

centers strongly integrated with the network interconnecting them. Distributed data centers do not only bring services closer to users with better quality, but also provide opportunities to improve energy efficiency exploiting the variation of prices in different time zones, the locally generated green energy, and the storage systems that are becoming popular in energy networks. In this paper, we propose an energy aware joint management framework for geo-distributed data centers and their interconnection network. The model is based on virtual machine migration and formulated using mixed integer linear programming. It can be solved using state-of-the-art solvers such as CPLEX in reasonable time. The proposed approach covers various aspects of Cloud computing systems. Alongside, it jointly manages the use of green and brown energies using energy storage technologies. The obtained results show that significant energy cost savings can be achieved compared to a baseline strategy, in which data centers do not collaborate to reduce energy and do not use the power coming from renewable resources.

Keywords Green cloud · Energy consumption · Green energy · VM migration · Energy efficiency · Joint optimization

1 Introduction

In recent years, the wide adoption of Information and Communication Technologies (ICT) and the exponential growth of Internet users have significantly contributed to the increase of the world energy consumption [1, 2], and the impact of the digital economy is expected to increase even more over the next years [1, 3]. Even if ICT is actually helping other sectors of the economy to reduce their environmental impact, the energy consumption of the ICT sector itself cannot be neglected.

The main strategy of the research effort put so far in ICT energy issues is the reduction of energy consumption, with the constraint of guaranteeing the same quality of service. In addition, partially replacing polluting energy by installing green energy plants close to big consumers is also part of the solution to the problem.

In Cloud computing, data centers are well known for being particularly energy hungry. Electricity consumed by global data centers is estimated to be between 1.1 and 1.5% of total electricity use [4]. Typically, data centers are rather inefficient and consume more energy than required [5], leaving room for improvement achievable through intelligent management techniques.

By breaking down the energy consumption of data centers into their components as shown in Fig. 1 [6], we can observe that about 52% of energy is consumed by computing equipments and the remaining 48% are for power equipments and cooling systems.

One of the reasons for energy inefficiency is the underutilization of servers whose consumption is not proportional to computing load. As the statistics show, average server utilization in data centers is around 30% [7], due to capacity over provision based on worst-case scenario in order to ensure high levels of reliability [8].

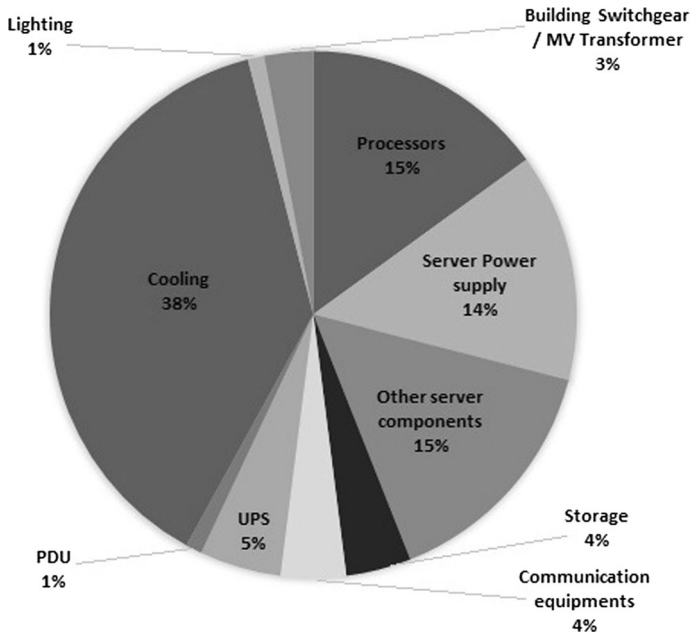


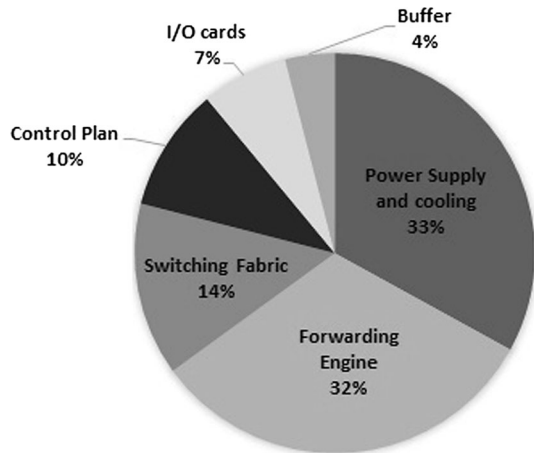
Fig. 1 Data center energy consumption breakdown

To address this problem and improve Cloud systems power efficiency, VM (virtual machine) migration has been proposed and has shown great potential. Migrating a VM consists of changing its physical host without service interruption. This can be done for different purposes, such as fault resilience or for system maintenance. It can be used in power management strategies to move services running on a big set of underutilized servers to a smaller set of optimally loaded servers, so that the others can be switched off for power saving.

Even if VM migration also consumes energy, it has been shown that it is more effective than leaving underutilized servers switched on [9, 10].

Another important component that consumes energy in Cloud services is the communication network. Networks are also typically provisioned for worst-case scenarios such as traffic burst and busy-hours load. Actually, the network has typically a pretty large capacity margin even with respect to peak load for service quality and robustness reasons, and then it usually wastes a lot of energy [11]. In the internal data center network, the main energy consumers are Ethernet switches that are hierarchically interconnected. In the external network based on IP technology, the core network routers dominate energy consumption [12]. The relative contribution to energy consumption of core router components is shown in Fig. 2 [12]. IP networks typically operate at less than 50% utilization, while still consuming almost 100% of maximum power due to an almost flat energy profile (consumption versus load) [11]. For managing network devices in order to consume less energy, two main approaches are used: turning off the nodes or scaling down their performance [13, 14].

Fig. 2 Breakdown of power consumed by a core router



Most of previous work on energy efficiency in Cloud systems focused on managing computing and networking components separately. However, optimizing energy consumption of data centers and their network independently may be significantly inefficient, in particular when dynamic resource management schemes like VM migration are considered. Considering energy consumption of data center servers only may cause traffic congestion and degrade the quality of Cloud services offered to end users, as well as decrease the energy efficiency of the network. On the network side, energy saving techniques are based on estimations of the traffic matrices over time, and if data centers are not considered, large traffic variations due to decisions taken by dynamic resource managers can cause energy waste. Integrating techniques for managing energy consumption of computing and networking components in a new generation of Cloud systems can potentially provide non negligible efficiency gains.

A key aspect that makes some level of integration in services offered by data centers and networks particularly important is the geographic distribution of Cloud systems. Distributing data centers over different locations brings Cloud services closer to end users, and offers the opportunity to better exploit the variation of energy prices in different locations and time zones, as well as the efficient use of the green energy that is locally generated.

Even if there are only a few existing works in literature that have investigated the impact of joint optimization solutions for energy saving in Cloud systems [15–17], the effectiveness of such solutions from energy cost point of view, and their contribution to reducing environmental impact through the use of green energy remain open issues that have motivated our work (see Sect. 2 for more details).

In this paper, we present a holistic approach for jointly managing Cloud data centers and their networks. In the considered scenario, the Cloud system provides Platform as a Service (PaaS) to a variety of users, and data centers are distributed geographically in different locations and interconnected by a network. We propose an optimization model based on mixed integer linear programming (MILP), which has the goal of minimizing the Cloud energy cost and exploit the availability of green

energy sources in different places where data centers are located. The proposed approach covers many aspects of Cloud computing including live migration of VMs, energy storage management, and green energy exploitation. In this model, we consider both energy consumption of data-center servers and their interconnection network, and optimize the use of energy coming from the electrical grid, as well as the energy locally generated using renewable resources exploiting also energy storage [18].¹

The rest of the paper is organized as follows. Section 2 presents related work on energy saving. Section 3 describes the proposed approach and introduces the model formulation. Section 4 presents the tests performed and the obtained results. Finally, concluding remarks are given in Sect. 5.

2 Related Work

As mentioned in the introduction, most of the existing works tackle the problem of energy efficiency in Cloud systems separating the management of data center servers and network nodes.

For the data centers, there is a large body of work on energy management of computing resources. We can categorize existing approaches into two classes: server consolidation with power state management and workload scheduling.

Server consolidation consists of efficiently using the available computing resources with the view to reduce the total number of active servers and thus saving energy by turning off unused ones. Entropy is a resource manager proposed in [19] that is based on constraint programming, and it is able to consolidate applications running on a number of underutilized servers to a smaller number of highly utilized servers using live migration of VMs. The adopted scheme does not take into account heterogeneity in application requirements and servers, which is rather common in multiple cloud provider environments. A similar approach able to cope with heterogeneous environments is named pMapper [20], an application placement controller based on continuous optimization. For more complex environments, with a combination of service level agreements (SLAs), different power models and energy policies, a VM consolidation engine named Plug4Green has been proposed in [21].

The workload placement in modern data centers with a large number of servers significantly affects their operating temperature in addition to energy consumption. A smart placement using workload scheduling techniques may reduce cooling requirements and save even more energy. A good example of schemes based on this scenario is EnaCloud [22]. EnaCloud is an energy-aware heuristic-based approach that chooses the most appropriate scheme for dynamic application placement based on their arrivals, departures or resizing events. The approach in [23] proposes an

¹ This paper is an extension to our work presented in [18]. The main difference is that now we consider more complex scenario, we minimize the cost by solving VM migration problem instead of only redirecting requests between data centers.

integer linear programming (ILP) model that combines job allocation and VM migration.

None of the above solutions considers network requirements or the geographic distribution of data centers.

As for previous work on Cloud networks, available contributions are mainly focused on designing and operating the communication infrastructure in order to achieve fault tolerance, scalability, high utilization and cost efficiency [24, 25]. For energy efficiency, beside the hardware improvements, many contributions related to protocols and network architecture aim at achieving a better trade-off between performance and energy consumption.

In [26], Green Route (G-Route), a service routing protocol for achieving energy-efficiency and collaboration among cloud providers, is proposed. It is a routing scheme that creates autonomous energy-efficient paths between different providers before running a specific service. It has been implemented and tested on Amazon EC2 cloud infrastructure, and shown quite significant energy and cost savings per service request. A drawback of this approach is that it needs a trusted third party to control the energy profiling process. Other energy-aware routing solutions can be found in [27–34].

Switching off network nodes and rerouting traffic on other paths has a significant impact on saving energy. In [35] the authors proposed an integer linear programming model and some heuristic algorithms that minimize energy by finding the set of routers and links that must remain powered on for a given traffic level while switching off the others. This model is based on the knowledge of the traffic profile exchanged between source/destination nodes, and the maximum link utilization. Unlike most similar works where the objective is to minimize cost or to maximize performance, the authors minimize the total power consumption of the network. Other works that consider switching off network components and sleep mode for saving energy can be found in [36–43].

PCube [44] is an elastic data center scheme that conserves energy by adjusting the network topology and varying bandwidth availability based on traffic demand. It is designed to be able to dynamically adjust network structure depending on different traffic volumes, and to turn off a set of switches to save energy. Similar solutions are Bcube [45] and ElasticTree [46].

The above solutions focus on network only and do not consider the energy optimization of servers in data centers together with network nodes.

Relatively few papers consider joint management of data centers and network. In [16], the authors proposed an optimization approach to jointly minimize the energy consumption in data center hosts and network. The basic idea is to consider both VM placement and traffic routing for energy saving. To avoid the complexity of the problem, a unified representation method is proposed and the optimization model of VM placement is made similar to a routing problem; then the placement and routing problems are solved as a single one. A similar approach is PowerNetS [15], a power optimization strategy based on workload and traffic correlation analysis. The problem is formulated using constrained programming with the goal of consolidating VMs which are not positively correlated with the same physical machines. At the same time, the model takes into account the network by consolidating VMs that

are linked through traffic flows onto the same server or servers close to each other. In [47] the authors proposed an energy saving scheme for VM placement considering both physical servers and network resources. The problem is modeled as a combination of bin packing and quadratic assignment problems with multi-objective optimization and solved with a greedy algorithm that combines hierarchical clustering with best fit scheme. The above approaches focus on the network internal to the data center and do not consider geographically distributed Cloud systems, as well as the availability of green resources and energy storage.

The approach in [48] jointly minimizes the cost in big data processing based on three factors: task assignment, data placement, and data routing. The cost minimization problem is formulated as a mixed-integer nonlinear programming (MINLP) model, that is then linearized to make it tractable. This work considers geographically distributed data centers but it does not explicitly model the external network for energy consumption, it does not exploit the difference of energy prices in various locations, and it does not take into account green energy usage.

In [49], the authors jointly optimize three problems: VM placement, the distribution of requests, and data center resizing. The geo-distribution of data centers is explored, by considering variation in energy prices within different locations. The problem was formulated using MILP, and solved using a two phase heuristic. However, the energy consumption of the network was not included neither the exploitation of green energy resources.

In this paper, we study the energy cost minimization problem of Cloud systems by managing data centers and Network as a whole. Differently from existing works, where the focus is only on one aspect of Cloud computing like VM placement/migration, in this paper we aggregate multiple Cloud computing aspects into one approach, and propose a holistic energy aware solution for managing Cloud data centers and their interconnection network.

3 Global Green Cloud Management Framework

3.1 Problem Description

We consider a PaaS (Platform as a Service) scenario, where the provider operates on a virtualized infrastructure composed by multiple data centers distributed over different geographical locations. Each data center is equipped with thousands of physical servers. This scenario is rather common nowadays even if the number of locations and servers vary with the size of the provider. For example, Google data centers are distributed among various locations in the world: 19 in the US, 12 in Europe, one in Russia, one in South America, and three in Asia. [50] While Amazon Web Services Cloud operates 42 Zones within 16 geographic Regions around the world [51].

Let \mathcal{I} be the set of available data centers. We assume they are fully connected by a backbone network (mesh topology), where in each path between two data centers i and j , the number of routers and the available bandwidth capacity are known.

We assume that the Cloud Provider is able to host different user applications by offering a set \mathcal{L} of heterogeneous types of VMs. Each type of VMs executes a specific service application that is capable to serve a set \mathcal{K} of user request classes as shown in the system model (Fig. 3).

We consider a 1-day horizon, divided into 24 time periods, in which the duration of each time period is 1 h, and we solve the problem in advance for each day. We consider predictions of the application workload based on historical traffic information [52–54], thus, an estimation of the incoming traffic for each application is provided. We denote by λ_{ik}^t the arrival rate of requests of class $k \in \mathcal{K}$ to data center i at time $t \in \mathcal{T}$.

Based on the traffic profile, our goal is to minimize the total energy cost in the cloud system by allocating VMs to servers and if necessary, migrating them between data centers, considering the fact that migration itself costs energy on both source and destination. Depending on the location of a data center and its time zone (day/night), the price of energy varies. We exploit different energy regions by migrating VMs to data centers where the price of energy is cheaper. We also consider the availability of energy from renewable resources for reducing environmental impact of the cloud system. Therefore, migration of VMs to data centers with more available green energy is an opportunity that can be exploited to optimize costs. Basically, using VM migration we actually reduce the load of expensive and polluting data centers, while we exploit cheap and green energy when available.

We consider that VMs are live migrated between data centers (DCs) using post-copy live migration scheme [55], in which the VM is suspended immediately upon beginning of the migration process. First CPU state is transferred to the destination DC, while the memory is still at the source DC. Then the destination DC requests fault pages from the source DC, while the latter is transferring the memory state to the destination DC. Post-copy live migration decreases the migration traffic therefore reduces the total migration time, since the VM page is transferred only once over the network unlike in pre-copy live migration [56]. Indeed, in this case memory pages are copied from the source DC to the destination DC without interrupting the execution of the virtual machine, which implies a succession of iterations of memory transfer before stopping the VM execution on the source DC and starting it again at the destination DC.

By migrating VMs from one data center to another, we consider migration energy cost of the destination data center, the source data center, and the network. For the network, its energy is assumed to be proportional to the amount of exchanged traffic, which includes bandwidth consumed by migrated VMs including there memory, disk images of VMs, and users traffic.

Moreover, we jointly manage the use of green and brown energy. We assume that all the energy coming from the electrical grid is brown energy, while we consider the energy generated in data centers as green and available for free if it is consumed locally, to privilege the use of on-site generated green energy. However, it is straightforward to modify the formulations to include also an amount of green energy coming from the electrical grid (with a cost depending on location).

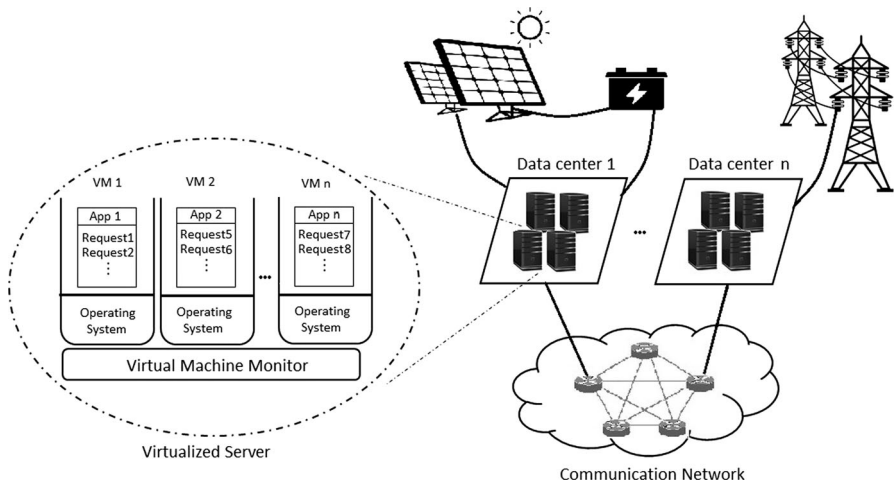


Fig. 3 Cloud system model

We assume that data centers are able to generate on-site an amount of green energy using renewable sources, such as solar energy, wind energy or geothermal energy. The use of this smart electricity sourcing strategies on-site is increasing, e.g., Facebook’s solar-powered data center [57], and Green House Data wind-powered data center [58]. Since our work focuses only on system management, we do not include capital expenditures for renewable sources. The costs of green energy generation are significantly declining over the last few years. Depending on the technology, installation prices vary, e.g., Parabolic trough plants used to generate Concentrating Solar Power (CSP) have capital costs as low as 4600\$/kW in USA market, while wind power technologies tend to be more competitive, between 1800 and 2200\$/kW [59].

Matching exactly the energy consumption with green energy generation is difficult and can potentially generate inefficiencies when produced energy cannot be consumed immediately. Therefore, we relax this problem by considering the use of energy storage technologies. In our scenario, data centers are equipped with rechargeable battery systems that are able to store the locally generated green energy. Balancing the use of green energy produced, between immediate usage and storage in batteries for later consumption allows for green energy to be available when the price of brown energy is high, as well as to solve the problem of discontinuous availability of renewable resources.

For each time period, the model defines how to allocate the load in each data center. In other words, how many VMs are kept active or off. The same thing for the network, where we define for each time period which links are to be turned on and which should be off depending on the number of routers in each link and their capacity. Note that, even if we associate energy consumption to links, the real energy consumers in the network are line cards connected to the links in the routers on both sides.

3.2 Problem Formulation

In this section, we first introduce decision variables. We then formulate the cost minimization objective function and the problem constraints.

3.2.1 Decision Variables

The goal of our optimization model is twofold: (1) finding how many VMs to be migrated among DCs, and the source and destination of the migrations, (2) managing the usage of the available green energy sources. These decisions allow to move the load and the energy consumption among DCs during the day in order to exploit the availability of green energy and the differences in brown energy prices of various DC locations.

We formulate the problem using several sets of decision variables. The first pair of main decision variables refer to VMs migration and requests forward. The integer variable v_{ijl}^t represents the number of VMs of type l to be migrated from DC i to data center j during time period t . This variable depends on the number of received requests by each data center. We use a continuous positive variable x_{ijkl}^t to represent the arrival rate of class k requests received by VM of type l in data center i and then served in data center j after live migration.

The second pair of main decision variables are related to the green energy management. In particular, variables sge_i^t and dge_i^t indicate the sources of the green energy supplied to the DC i during the time period t . sge_i^t is the amount of energy coming from DC batteries, while dge_i^t is the green energy produced at the DC i directly supplied to the DC for its operations.

Together with these main decision variables, there are several other secondary variables that, depending on the values of the main ones, are used to model the behavior of the system. We can group them in the three domains, they refer to: VM and Migration, Networking, Battery and Green Energy Management.

3.2.2 VM and Migration

Integer variable \bar{w}_{il}^t denotes the number of active VMs type l that are originally running on data center i , while we use the integer variable w_{il}^t to refer to the number of VMs executing on a data center i after all live migrations took place. The variables won_{il}^t and $woff_{il}^t$ depict the number of VMs to be turned on and off respectively at time period t . The energy consumption associated to the migration of type- l VMs, including both current VMs at DC i migrated to other DCs and new VMs migrated to DC i from other DCs, is captured by the variable mig_{il}^t . Figure 4 describes a small illustrative scenario of the migration process, in which a DC i receive a number v_{ijl}^{t-1} of migrated VMs from a DC j during the time slot $t - 1$, then, during time t it migrates a number v_{ijl}^t to a DC j , we mention the space between time bands is just to show the number of turned off VMs $woff_{il}^t$.

Networking As for the network, we assume that the energy consumption of a link is proportional to its load, expressed in terms of the ratio of used bandwidth over




Fig. 4 VM variables

available bandwidth. We rely on variable b_{ij}^t to express the bandwidth used at the link (i, j) connecting DC i with DC j during the time period t , which is determined by the traffic volume exchanged by the two DCs due to all migration processes among them. In addition, we let unused links to be switched off. We capture this behavior using binary variables z_{ij}^t , equal to 1 if the link (i, j) is active, and 0 otherwise. Similarly to data centers, zon_{ij}^t and $zoff_{ij}^t$ indicate whether each link has to be turned on or off at the beginning of time period t according to its status during the previous time period $(t - 1)$. Figure 5 illustrates a small example of networking variables changes during a small scenario for transmissions in the link between a DC i and a DC j . In this small example, the link is used for transmissions from time band 5 until time band 8, after that, during time band 9 there are no transmissions, therefore the link is turned off, and we indicate that by assigning 1 to the variable $zoff_{ij}^t$. The link has to be turned on again during time band 11, though, the variable zon_{ij}^t takes the value 1.

Battery and Green Energy Management Concerning green energy and batteries, the variable c_i^t represents the amount of energy charged in a battery i at time t from renewable energy sources installed at DC i . It is related to the main decision variables sg_e^t and dge_i^t as described by the Fig. 6, where basically the generated green energy not immediately provided in dge_i^t is used to recharge the batteries.

In addition to the above-mentioned variables, we have a set of variables to model the energy charging and discharging phases at DC i . We suppose that the energy charged at a time t cannot be used in the same time period. To define the energy




Fig. 5 Networking variables

level of a battery we use two different variables, one refers to the level of energy at the beginning of a time period t denoted by \bar{s}_i^t , and the other one s_i^t refers to the level of energy at the end of the time period t . Figure 7 shows how variable sg_e^t , which indicates the amount of the batteries' energy consumed during time period t to run the DC, is connected to other energy-related variables. Finally, variables sg_e^t and dge_i^t define the total amount of green energy provided to the DC during the time period t , expressed by the variable g_i^t .

All defined variables are summarized in Table 1.

3.2.3 Objective Function

The objective of our model is to minimize the energy cost, which consists of two components: DC energy consumption and networking energy consumption. Therefore, we design the following objective function:

$$\min \sum_{t \in T} \sum_{i \in \mathcal{I}} \left\{ M_i^t \left[\rho_i \sum_{l \in \mathcal{L}} (\alpha_{il} w_{il}^t + \eta_{il} won_{il}^t + \theta_{il} woff_{il}^t + mig_{il}^t) - g_i^t \right] \right\} + \sum_{t \in T} \sum_{i, j \in \mathcal{I}} E_{ij}^t R_{ij} \left[(\gamma_{ij} - \delta_{ij}) \frac{b_{ij}^t}{Q_{ij}} + \delta_{ij} z_{ij}^t + \tau_{ij} zon_{ij}^t + \zeta_{ij} zoff_{ij}^t \right] \tag{1}$$

The first term accounts for the cost of data centers consumption. It considers the costs of all data centers over all time periods, where for each data center we multiply the specific site cost of brown energy, M_i^t , and PUE (Power Usage efficiency), ρ_i , for the total energy consumed by the servers. The consumed energy consists of:

- the total consumption of running VMs, where α_{il} is energy needed for running a type l VM in DC i (e.g., Wh)
- the energy needed for turning on and off the servers, where η_{il} and θ_{il} are, respectively the energy needed for turning on or off a type l VM in DC i (e.g., Wh)




Fig. 6 Energy generation, storage and consumption




Fig. 7 Batteries functioning

Table 1 Decision variables

Main decision variables	
v_{ijl}^t	Number of migrated VM type l from DC i to DC j at time t
x_{ijkl}^t	Number of requests class k received by i and treated in j
sge_i^t	Energy supplied to DC i from batteries at time t
dge_i^t	Green energy directly supplied to DC i at time t
VMs related variables	
\bar{w}_{il}^t	Number of VMs type l at DC i before live migration
w_{il}^t	Number of VMs type l at DC i after live migration
$migs_{il}^t$	Migrations energy consumption of VMs class l in DC i during t
won_{il}^t	Number of VMs type l to be turned on
$woff_{il}^t$	Number of VMs of type l to be turned off
Network variables	
b_{ij}^t	Bandwidth used in link (i, j) at time t
z_{ij}^t	Link (i, j) status at time t (binary)
zon_{ij}^t	Whether the link (i, j) has to be turned on (binary)
$zoff_{ij}^t$	Whether the link (i, j) has to be turned off (binary)
Batteries and green energy variables	
c_i^t	Energy charger in battery i at time t
\bar{s}_i^t	Energy level in battery i at the beginning of time period t
\underline{s}_i^t	Energy level in battery i at the end of time period t
g_i^t	All green energy used by DC i during t

- the energy consumed by DCs to migrate VMs, captured by variable $migs_{il}^t$

In addition, the consumed energy is discounted by the amount of green energy provided by renewable energy sources installed at DCs represented by g_i^t . As we consider the green energy produced locally, the operating cost of green plants is set to zero in the model.

The second term of the objective function accounts for the network consumption. It is computed as a sum of each path cost, which in turn consists of two components:

- the energy required to operate each router on link (i, j) during the time period t , considering both active and idle state energy consumption, respectively indicated by γ_{ij} and δ_{ij}
- the energy required to turn on and off a router at the beginning of the time period, respectively, τ_{ij} and ξ_{ij}

The model assumes that all the routers of the link (i, j) are identical, therefore, the energy consumption of each router multiplied by the number of routers along the link, R_{ij} . This assumption can be easily modified by considering individualized energy consumption values, which have been omitted here for sake of ease in presentation. Finally, in order to compute the energy cost of this second term, the total energy consumption is multiplied by the average energy price along the link during time period t , E_{ij}^t .

3.2.4 Constraints

In this section we present different groups of constraints used to model the Vm migration, the operations of data centers, the network, and battery and green energy management features.

VM Migration First, we must ensure that all the requests received from cloud users are processed by the data centers. The requests of different classes have to be processed by suitable type of VMs. For this purpose we use the following constraints:

$$\sum_{l \in \mathcal{L}_k} \sum_{j \in \mathcal{I}} x'_{ijkl} = \lambda'_{ik} \quad \forall i \in \mathcal{I}, \forall k \in \mathcal{K}, \forall t \in \mathcal{T} \quad (2)$$

In particular, constraint (2) ensures that all the incoming traffic is processed in the Cloud, by any of the DCs with an appropriate VM. Note that we consider the set \mathcal{L}_k , which is the set of VM classes that can process class- k requests.

In addition, a migration plan requires to define the number of VMs to migrate, as well as their source and destination DCs. For this purpose we use the following constraints:

$$\sum_{j \in \mathcal{I}} v'_{ijl} \geq \sum_{k \in \mathcal{K}} \left[\left(\sum_{j \in \mathcal{I}} x'_{ijkl} \right) - x'_{iikl} \right] / \mu_l \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (3)$$

$$\sum_{j \in \mathcal{I}} v'_{jil} \geq \sum_{k \in \mathcal{K}} \left[\left(\sum_{j \in \mathcal{I}} x'_{jikl} \right) - x'_{iikl} \right] / \mu_l \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (4)$$

$$v'_{iil} = 0 \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (5)$$

Equations (3) and (4) compute the number of VM that were sent and received, respectively, by each DC. They are proportional to the rate of request, respectively, redirected to and received from other DCs. Note that the term x_{ikl}^t represents the requests arrived to DC i and locally served. The number of migrated VMs must be sufficient to serve the rate of forwarded requests, this is captured by the parameter μ_l expressing the maximum service rate for a type l VM. The last constraint (5) guarantees that a DC does not migrate VMs to itself.

The following constraint (6) determines the number of active VMs in a data center i after making the necessary migrations. Starting from the number of VMs that was created in a data center (\bar{w}_{il}^t), we subtract the number of VMs that migrated and add the ones that arrived from other locations. In most cases, a data center makes one of these operations: sending or receiving VMs but not both at the same time, and that depends on its capacity and the energy constraints.

$$\underline{w}_{il}^t = \bar{w}_{il}^t - \sum_{j \in \mathcal{I}} v_{ijl}^t + \sum_{j \in \mathcal{I}} v_{jil}^t \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \tag{6}$$

Finally, constraint (7) calculates the energy consumed by migration operations.

$$mig_{il}^t = EH_l \sum_{j \in \mathcal{I}} v_{jil}^t + ES_l \sum_{j \in \mathcal{I}} v_{ijl}^t \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \tag{7}$$

Parameters ES_l and EH_l represent the energy consumption for migrating a VM type l consumed, respectively, at source and destination DC.

DC Behavior Together with the migration plan, we need to model the DC behavior. We assume each DC has a maximum number of requests that can process per time period. This number depends on the number of VMs that a data center is able to handle simultaneously. Hence, we use the following constraints to ensure that the capacity requirements of data centers are not exceeded:

$$\underline{w}_{il}^t \leq P_{il} \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \tag{8}$$

$$\bar{w}_{il}^t \geq \sum_{j \in \mathcal{I}} \sum_{k \in \mathcal{K}} \frac{x_{ijk}^t}{\mu_l} \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \tag{9}$$

$$\underline{w}_{il}^t \geq \sum_{j \in \mathcal{I}} \sum_{k \in \mathcal{K}} \frac{x_{jik}^t}{\mu_l} \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \tag{10}$$

where, constraint (8) ensures that the number of running VMs after all migration operations took place does not exceed the capacity of the system resources. In other words, the overall utilization of resources dedicated to run class- l VMs is below a planned threshold in each DC i , P_{il} . Constraint (9) defines the number of VMs of type l originated in DC i , while constraint (10) defines the number of VMs of type l running on DC i after making all the necessary live migrations. Both numbers depend on outgoing and ingoing rates of migrated requests.

In order to ensure time continuity in the number of running VMs at each DC, we need the following constraints.

$$won_{il}^t \geq \bar{w}_{il}^t - \underline{w}_{il}^{t-1} \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (11)$$

$$woff_{il}^t \geq \underline{w}_{il}^{t-1} - \bar{w}_{il}^t \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (12)$$

Constraints (11) and (12) determine the number of VMs to be turned on and off at the beginning of time period t , according to their number at the end of time period $t - 1$.

Networking For the network, the following constraints are defined to ensure that we do not exceed bandwidth capacity and to guarantee the proper operation of network links:

$$b_{ij}^t = \sum_{l \in \mathcal{L}} \phi_{ijl}^t (VMsize_l + DI_l) + \sum_{k \in \mathcal{K}} \left(B_k \sum_{l \in \mathcal{L}} x_{ijkl}^t \right) \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{I}, \forall t \in \mathcal{T} \quad (13)$$

$$b_{ij}^t + b_{ji}^t \leq Q_{ij} z_{ij} \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{I}, \forall t \in \mathcal{T} \quad (14)$$

$$z_{ij}^t = z_{ji}^t \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{I}, \forall t \in \mathcal{T} \quad (15)$$

Constraint (13) computes the portion of bandwidth used for transferring data between different DCs. In our scenario, all the exchanged data is related to migration operations. Basically, we consider the size of different types of migrated VMs, in terms of memory state and the content of CPU registers, then, we associate to each type of VMs an amount of users traffic in terms of requests. Parameters $VMsize_l$ and DI_l indicates the bandwidth consumed to migrate a type- l VM and the size of disk images of type- l VMs, while B_k is the bandwidth required for a type- k request. In order to consider the effect of large latency between data centers, a scale factor can be added to Eq. (13) to take into account the increase of bandwidth needed to manage memory consistency during VM migration. However, a detailed investigation of the effects of the latency on VMs' performance is out of the scope of this paper, which mainly focuses on energy efficiency aspects. Therefore, in the rest of the paper we assume this scale factor can be approximatively set to 1.

Constraint (14) guarantees that the VM exchanges do not exceed the link capacity, Q_{ij} , which is forced to 0 when the link is switched off ($z_{ij} = 0$). Finally, Constraint (15) ensures that if a link is active in one direction, it is also active in the other one.

Similarly to the case of the number of VMs at DCs, we need to ensure the time continuity of the link status. Constraints (16) and (17) define which links have to be turned on or off at each time period transition.

$$zon_{ij}^t \geq z_{ij}^t - z_{ij}^{t-1} \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (16)$$

$$zoff_{ij}^t \geq z_{ij}^{t-1} - z_{ij}^t \quad \forall i \in \mathcal{I}, \forall l \in \mathcal{L}, \forall t \in \mathcal{T} \quad (17)$$

Green Energy and Batteries Management Green energy management and storage play an important role in our model. The constraints below guarantee the appropriate behavior of the available renewable resources and batteries.

$$g_i^t \leq \rho_i \sum_{l \in \mathcal{L}} (\alpha_{il} w_{il}^t + \eta_{il} w_{on_{il}}^t + \theta_{il} w_{off_{il}}^t + mig_{il}^t) \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{18}$$

$$g_i^t = sge_i^t \beta_i + dge_i^t \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{19}$$

$$c_i^t + dge_i^t \leq \Gamma_i^t \quad \forall i \in \mathcal{I}, \quad \forall t \in \mathcal{T} \tag{20}$$

$$\sum_{i \in \mathcal{I}} g_i^t \leq \sum_{i \in \mathcal{I}} \Gamma_i^t \quad \forall i \in \mathcal{I} \tag{21}$$

In particular, constraint (18) ensures that the consumed green energy during a time period is less than the required amount to run the corresponding DC. Constraint (19) states that the green energy at time period t can be provided directly from the renewable source (dge_i^t) or from batteries (sge_i^t), taking into account the energy loss rate during a time period due to its storage, denoted by β_i . Constraint (20) guarantees that the amount of energy charged in batteries and the one directly supplied are less than the total amount generated in one time period, denoted by Γ_i^t , while constraint (21) ensures that the total green energy consumed during a day in a data center ($\sum_{i \in \mathcal{I}} g_i^t$) does not exceed the total amount generated ($\sum_{i \in \mathcal{I}} \Gamma_i^t$). This makes the daily repetition of the plan sustainable.

$$\bar{s}_i^t = c_i^{t-1} \psi_i + \underline{s}_i^{t-1} \zeta_i \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \setminus \{1\} \tag{22}$$

$$sge_i^t = \bar{s}_i^t - \underline{s}_i^t \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{23}$$

$$sge_i^t \leq \bar{s}_i^t \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{24}$$

Constraints (22), (23), and (24) are related to batteries. Constraint (22) states that the energy level in a battery at the beginning of the period t is given by the energy remaining at the end of time period $t - 1$ (considering the energy discharging efficiency of the battery ζ_i) and the energy charged during $(t - 1)$ (considering the energy charging efficiency of the battery ψ_i). Constraint (23) forces the amount of discharged energy, sge_i^t , to be equal to the difference between the level of energy at the beginning and at the end of time period t . Constraint (24) makes sure that the energy discharged from a battery is less than the available energy in that battery at the beginning of the time period.

The following final constraints are related to physical limitations of the batteries. Constraints (25), (26) and (27) ensure that the model does not exceed the energetic capacity and charging and discharging rate limits of a battery.

$$\bar{s}_i^t \leq Smax_i \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{25}$$

$$c_i^t \leq Cmax_i \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{26}$$

$$sge_i^t \leq Dmax_i \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \tag{27}$$

Parameters $Smax_i$, $Cmax_i$, and $Dmax_i$ refer, respectively, to the maximum energy storage capacity, energy charging in 1 h, and energy discharging in 1 h of the considered batteries (which is the duration of one time period). We also assume the energy charged into a battery during time t can not be used until the next time period, therefore, we add the constraints (28), (29) and (30) to initialize batteries' status in $t = 1$.

$$sge_i^1 = 0 \quad \forall i \in \mathcal{I} \quad (28)$$

$$\bar{s}_i^1 = 0 \quad \forall i \in \mathcal{I} \quad (29)$$

$$\underline{s}_i^1 = 0 \quad \forall i \in \mathcal{I} \quad (30)$$

Table 2 summarizes all model parameters.

4 Model Evaluation

Our model has been evaluated using a state-of-the-art MILP solver and considering various instances and workload configurations. In this section, we present results obtained on a set of scenarios with realistic values for parameters.

4.1 Parameter Setting

We considered 15 data centers distributed geographically all over the world. For their locations, we have taken inspiration from Cloud computing infrastructure of Google [60]. We used four geographical macro areas: West USA, East USA, Europe and Asia. In each area, data centers have the same or close time zone. A detailed view of used data centers location and number of servers is provided in Table 3. We mention here that the number of servers used for each data center is not the real number of Google DCs, but it is generated within a the range of [5000, 16,000].

For each data center, we associate a PUE value in order to include the power facilities that support the IT equipment load, such as cooling systems. According to [61] the global average PUE of the largest data centers is around 1.7, while the average PUE for all Google data centers is 1.12. In our tests, we vary PUE values between 1.1 and 2.

For data centers capacity, we generated a random number of physical servers for each DC, within the range [5000:16,000], and we assume 1:1 ratio for the physical to virtual resources assignment (i.e., 1 physical core is assigned to 1 virtual core of equal capacity).

Regarding technical characteristics of servers in DCs, we consider an HP ProLiant DL370 G6, with a Intel Xeon W5580 processor (8 cores at 3200 Mhz) and 96 GB of total memory. Even if we considered three different classes of VMs (see below), we modeled only a single server type, in order to simplify energy consumption analysis. For this reason, all VMs require the same amount of energy to run at peak load or when idle, while they differ in the class of requests that can be

Table 2 Model parameters

\mathcal{I}	Set of data centers
\mathcal{T}	Set of time bands
\mathcal{L}	Set of VMs types
\mathcal{K}	Set of request classes
Data centers and VMs parameters	
λ_{ik}^t	Incoming class k requests rate at DC i in time t (Req/h)
B_k	Bandwidth required for a type k request
P_{il}	The maximum number of VMs type l in DC i
$VMsize_l$	Bandwidth consumed to migrated a type- l VM type
DI_l	Size of disk images of VMs type l
μ_l	Maximum service rate for a type l VM
EH_l	Destination DC energy consumption for migrating a VM type l
ES_l	Source DC energy consumption for migrating a VM type l
α_{il}	Energy needed for running a type l VM in DC i
η_{il}	Energy needed for turning on a type l VM in DC i
θ_{il}	Energy needed for switching off a type l VM in DC i
ρ_i	PUE (Power Usage efficiency) for DC i
M_i^t	Price of energy in DC i at time t
Network parameters	
R_{ij}	Number of routers in link (i, j)
Q_{ij}	Maximum bandwidth in link (i, j)
γ_{ij}	Energy needed for running a router in (i, j)
δ_{ij}	Energy needed for keeping idle a router in (i, j)
τ_{ij}	Energy needed for turning on a router in link (i, j)
ζ_{ij}	Energy consumption for switching off a router in (i, j)
E_{ij}^t	Price of energy in link (i, j)
Batteries and Green Energy parameters	
Γ_i^t	Green Power that could be generated at DC i in time t (kWh)
ψ_i	Energy charging efficiency in DC i battery
β_i	Energy loss rate per time in DC i battery
ζ_i	Energy discharging efficiency in battery i
$Smax_i$	Maximum energy storage capacity in battery i
$Cmax_i$	Maximum energy charging in 1 h for battery i
$Dmax_i$	Maximum energy discharging in 1 h for battery i

processed and their total number per time period. However, the model is flexible to include more than one type of physical servers, this is possible through assigning each class of VMs to a specific configuration of physical servers, therefore, different classes of VMs can take different values of energy consumption. The values of energy consumed for migrating a VM are taken from an experiment that is designed to estimate servers (host/destination) consumption due to live migration [9]. However, the model is tested on an example of homogeneous data centers, but the

Table 3 Data centers location and number of servers

DC	City	Country	Macro-area	Time zone	Servers
DC1	Mountain View—CA	USA	West America	UTC−08	6000
DC2	Pleasanton—CA	USA	West America	UTC−08	5250
DC3	San Jose—CA	USA	West America	UTC−08	6000
DC4	Atlanta—Georgia	USA	East America	UTC−05	10,500
DC5	Reston—Virginia	USA	East America	UTC−05	6750
DC6	Berlin	Germany	Europe	UTC+01	9750
DC7	Groningen	Netherlands	Europe	UTC+01	12,000
DC8	Mons	Belgium	Europe	UTC+01	13,000
DC9	Paris	France	Europe	UTC+01	15,750
DC10	Dublin	Ireland	Europe	UTC+00	13,500
DC11	Milan	Italy	Europe	UTC+01	9000
DC12	Moscow	Russia	Asia	UTC+03	14,250
DC13	Tokyo	Japan	Asia	UTC+09	11,250
DC14	Hong Kong	China	Asia	UTC+08	8250
DC15	Beijing	China	Asia	UTC+08	11,250

formulations are flexible to include heterogeneity by giving different energy consumption for different classes of VMs.

Another important parameter is the energy cost, which varies over time, with peak hours not simultaneously occurring at different time-zones. In this paper, we have used as input the average energy prices in the day-ahead market in different markets in the world including GME (Gestore dei Mercati Energetici) in Italy, New England Market and PJM in California USA, SEMO in Ireland and many others. Energy prices were collected and averaged during October 2014. Table 4 reports the list of market managers considered. The resulting costs of energy varies between 10 and 65 Euro/MWh. Figure 8 represents energy price trend for each different macro area.

Table 4 Energy market managers considered in the paper

Country	Market manager
USA	New England Market (ISO-NE) [62], California ISO [63], PJM [64]
Canada	Independent Electricity System Operator (IESO) [65]
Italy	Gestore dei Mercati Energetici (GME) [66]
France	Powernext [67]
Netherlands, UK, and Belgium	APX-ENDEX [68]
Germany	European Energy Exchange (EEX) [69]
Ireland	SEMO [70]
Japan	Japan Electric Power Exchange (JEPX) [71]
Russia	Trade System Administrator (ATS) [72]

We assume that data centers are fully connected, thus, we consider that capacity of different links varies between 0.5 and 1 Gbps. Moreover, a typical link connecting data centers is built up both by physical lines (such as optical fiber) and network components (such as routers and switches). Therefore, we estimate the energy cost of each link as the cost of energy consumed by its routers, proportionally to the bandwidth in use. For the number of routers in each link, a traceroute application was used to determine the number of hops between two nodes. We have also considered a single reference router which is a Juniper E320, with a maximum power consumption of 3.84 kW [73]. The other values of parameters related to routers are listed in Table 6.

We built workloads based on a trace of requests registered at a website of a big University. This trace was collected hourly during 1 year, from sessions registered on 100 servers. To generate the workload, we consider the total number of Internet users for each country where a data center is located [74], then, based on the number of Google search done all over the world and percentages of Internet users of the same country, we estimate the requests rate for each data center.

For VMs types, we consider 3 classes of VMs, where each class is able to serve 5 types of requests. The differences between this three classes is mainly the size of the VM and the size of its disk images. While each type of request have different requirements in terms of bandwidth. For VMs size and related parameters, we took inspiration from Amazon EC2 Instance Types [75]. Disk images size varies depending on the type of VMs. In the considered scenario, we assume that the type of VM are not storage intensive therefore we consider that the size of disk images is between 0.5 and 20 Gb. However, latency factor is considered to be 1 in the following tests. The considered values of each class are summarized in Table 5.

In order to estimate the total amount of green energy produced by each data center during a single day, we multiply the average energy produced by a green plant per square meter with the average data center size that we vary between 450 and 10,000 m² [76]. Moreover, we consider that data centers are equipped with Li-ion (lithium-ion) batteries with overall capacity of 1486 Ah. This kind of batteries

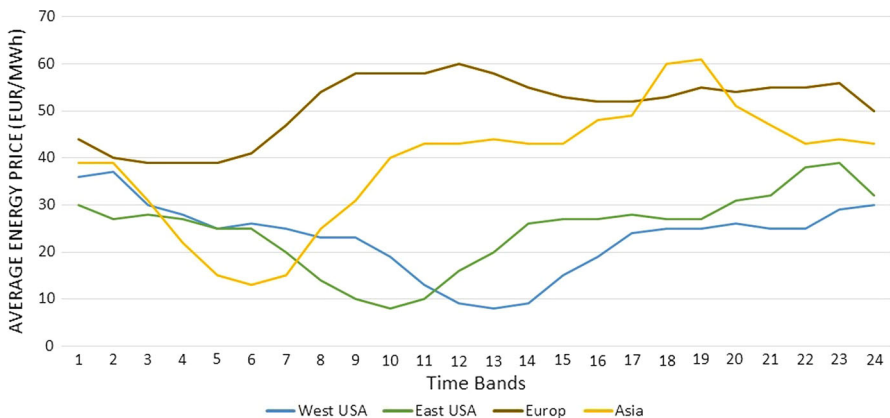


Fig. 8 Average energy prices for different macro areas during 1 day

Table 5 VMs settings

VM class	VM size (GB)	Disk images size (GB)
Class 1	0.5	0.5
Class 2	1	7
Class 3	2	20

have a C-rate around 73 Ah per module, so with a voltage of 14.8 V, a module can charge 1.08 kWh during 1 h, which is almost the full capacity of a module. Energy charging and discharging efficiency are considered equal to 88% [77]. Table 6 summarizes parameter settings used to test the model.

4.2 Numerical Results

We used the commercial solver IBM ILOG CPLEX 12.1 as a MILP optimization solver [78]. The model has run on an 8-core 2.4 GHz Intel Xeon server with 96 Gb RAM. To evaluate the energy saving of joint optimization and compare it to traditional strategies that separate data centers and network energy management, we considered the proposed model denoted by *Global Green* and two different scenarios.

The *Servers Only* scenario, where we optimize only servers side without considering an optimization for the network. In other words, network equipments are considered to be turned on all the time without any energy management strategy. A second scenario, called *Separated*, consists of two separated energy management strategies for both data centers and the network. In this case, data centers collaborate to minimize their energy consumption without considering that their interchangeable traffic can have an impact on the network energy consumption, while the network side turn on and off the routers based on the amount of traffic imposed due to DCs load balancing. A third scenario called *Global Green* represents our proposed approach.

Considering that our model is solved based on 1 day horizon, all of the following results represent the behavior of the system in 1 day, energy expenditures included.

Figure 9 shows the energy costs for the different scenarios for different values of traffic (number of requests per day). We also report the energy savings (in

Table 6 Parameters settings

B_k [200, 450] kb	α_{il} [60, 90] Wh	η_{il} [2, 3] Wh
θ_{il} [0.28, 14] Wh	ρ_i [1.1,2]	γ_{ij} 3.84 kWh
δ_{ij} 0.768 kWh	τ_{ij} 0.128 kWh	ζ_{ij} 0.128 kWh
EH_l [203, 908] Ws	ES_j [203, 908] Ws	PUE [1.1,2]
$Smax_i$ 1486 Ah	$Cmax'_i$ 1.08 kWh/module	ψ_i and $\beta_i = 88\%$
M'_i [10, 65] Euro/MWh	DI_l [0.5, 20] Gb	Q_{ij} [0.5,1] Gbps
R_{ij} [1,36] Routers	Nb. of servers [5250, 157,500]	

percentage) of the *Global Green* scenario with respect the other two. These costs are the objective function values resulting from the optimization described in Sect. 3.2. The joint optimization approach (*Global Green*) can save large amount of energy cost up to 70% compared to the cost without network power management. The reason behind this is the non-negligible energy saved in the network turning on and off routers according to traffic dynamics. The savings can be up to 34% compared to the separate case, and this is due obviously to the separate solution of the two problems that leads to suboptimal solutions, and to the use of fixed traffic values for the optimization of the network. As expected, the savings tend to decrease as traffic increases since we are forced to keep more system nodes (servers and routers) active to accommodate a larger number of requests and we have less room (smaller space of admission solutions) for optimizing energy consumption.

In Fig. 10 we plot the energy consumption of the same three scenarios. Obviously, we observe a significant saving of about 43% of the joint model compared to the *Servers Only* case, because of the network energy consumption. On the other side, the energy consumption of the separate and joint models is comparable. The reason is that in the objective function we considered the energy costs rather than the energy consumption. Since the local generated green energy is available for free, the system tends to exploit it at best using also storage to better match production periods with consumption.

In order to better investigate the behavior of the proposed joint optimization model, we performed a series of other tests aimed at understanding the contribution of different system features like the geographical distribution of data centers with different energy prices and green energy availability. To this purpose we have considered with three additional scenarios: the *Brown Base* scenario, where we do not consider any load balancing between data centers nor the use of green resources. In the scenario called *Green Base*, we introduce the use of green resources locally but without transferring load between data centers. Conversely, in *Global Brown*

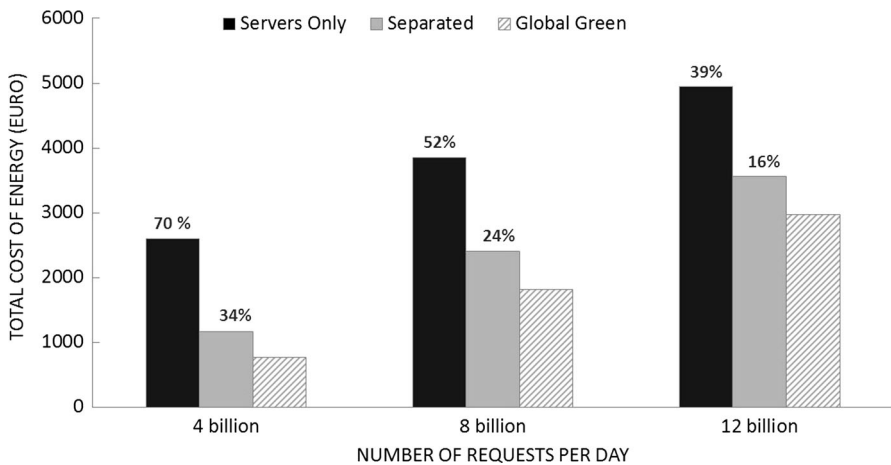


Fig. 9 Comparison of the optimal daily cost of energy for the different optimization approaches




Fig. 10 Comparison of the optimal daily energy consumption for the different optimization approaches

scenario, we exploit traffic distribution among data centers but without using green energy generation. Note that in scenarios where green energy production is considered, we also use storage to optimize its use over time.

Before analyzing the results of the model, it is worth considering their computational complexity. Figure 11 shows the average execution time required by CPLEX to solve the models using different workload configurations. We observe that in all cases it is possible to solve the problem within a few minutes even for a very large number of user requests. For small instances, i.e., 2 Billion requests, the solver takes around 1 min for the Global Green scenario, and half a minute for the Global brown scenario. While for the two base cases (Brown and Green), solution time is always less than half a second, even if the problems become unfeasible for instances with more than 16 billion requests, due to capacity limits of data centers servers. In the worst case, with a very high traffic compared to capacity (40 Billion




Fig. 11 Solving time

requests per day), the solution of the joint optimization problem took 4.52 min, which is very good for a 24 h time horizon of traffic planning.

Figure 12 shows the results obtained for energy costs using the scenarios mentioned above for different traffic levels. On top of each bar, we indicate the percentages of savings of the Global Green model compared the one indicated by the bar. It can be easily noted how significant cost reduction is achieved through collaboration between data centers using VMs migration. Moreover, with the cooperative and jointly optimized schemes the available capacity of the Cloud system is higher than with the non-cooperative schemes, as for high load levels the Base Brown and Green models are not feasible. We notice also that using the cooperative model without green energy (*Global Brown*) still provides non-negligible savings compared with the non cooperative approach (*Base Brown*).

To better understand how the *Global Green* model uses energy in the system, we can analyze the split of energy use in Figs. 13 and 14. The amount of green energy is limited by the capacity of generators considered in our instances, and it is a significant portion of the energy used only at low traffic levels, while it becomes rather small at high load. Taking a closer look to optimal solutions we notice that the system in different time periods tends to saturate the capacity of sites where energy is cheaper and green energy available for migrating VMs, and uses the other sites for the load exceeding capacity until the savings are significant compared to migration costs. As load increases, the capacity of cheap and green sites tends to be saturated by local demand and the cost savings decrease since migration is used mainly for load balancing.

Figure 15, shows the number of migrated VMs in both *Global Green* and *Global Brown* scenarios. As it is expected, it is proportional to the number of received requests. Even though VM migration process itself costs energy, the overall cost saved is more significant. We can notice also that with the presence of green sources




Fig. 12 Overall cost comparison




Fig. 13 Energy consumption split

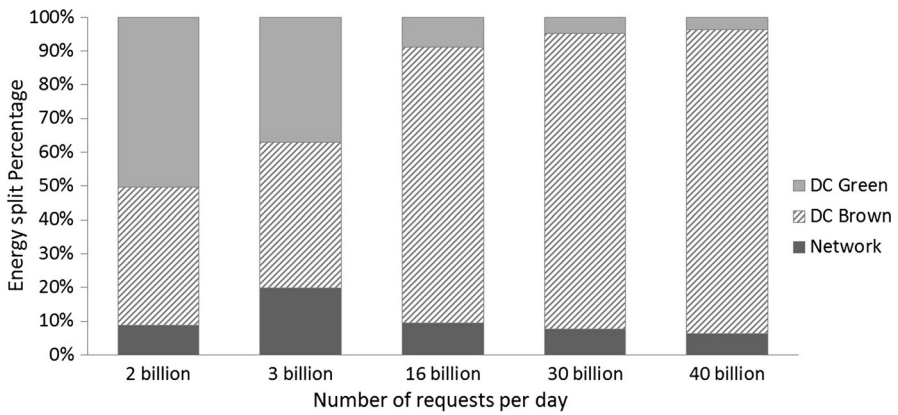


Fig. 14 Energy split percentage




Fig. 15 Number of virtual machines migrated

of energy more VMs are migrated to exploit it, unlike the brown case where we only benefits from the difference of prices between various locations.

While our model can achieve significant savings of the total power cost, it may cause consumption of larger amounts of energy as shows Fig. 16. The reason is that we use beside data center servers, network devices for VMs load balancing and this consume more power. On the other hand, our model performs better in exploiting green energy by migrating VMs as reported in Fig. 17. Therefore, even with the additional amount of energy that we consume, the proposed model is greener because it uses less brown sources of power by replacing them with green resources.




Fig. 16 Total energy consumption




Fig. 17 Green energy usage

5 Conclusion

Most of existing work on energy optimization in Cloud systems manages separately data center servers and their interconnection network. In this paper we presented a new optimization framework based on MILP for jointly management of Cloud data centers and their network.

The proposed model considers a set of data centers geographically distributed over different locations around the world. Data centers collaborate by migrating VMs between them when necessary to exploit different energy prices in various time zones. Another factor that we consider is the availability of green energy resources in some data centers and the possibility to store this energy using batteries.

In Cloud scenarios, migrating VMs between different sites needs additional network resources due to the size of VMs themselves and their data. Beside managing both data centers and their network, we also manage both the use of brown and green energies. Our strategy consists on redirecting the load to sites with more available green energy. We suppose also that some data centers are able to store the generated energy for later use, therefore we can save the clean power to use it when its generation is not possible or during peak energy price periods, thus we solve the problem of the possibility to be discontinued.

We show that the proposed optimization model can be solved using a state of the art MILP solver (CPLEX) in a reasonable time even for big size instances. The obtained results are very promising and shows that our approach allows significant cost saving compared to the base scenarios used nowadays. Moreover, from an environmental point of view, our model reduces greenhouse gas emission by pushing the Cloud to use more green power resources, alongside with optimizing its use in each data center using local energy storage.

Acknowledgements This work is partially funded by the European Commission under the Erasmus Mundus Green IT project (Green IT for the benefit of civil society. 3772227-1-2012-ES-ERA MUNDUS-EMA21; Grant Agreement No. 2012-2625/001-001-EMA2).

References

1. Van Heddeghem, W., Lambert, S., Lannoo, B., Colle, D., Pickavet, M., Demeester, P.: Trends in worldwide ICT electricity consumption from 2007 to 2012. *Comput. Commun.* **50**, 64–76 (2014)
2. Lannoo, B., Lambert, S., Van Heddeghem, W., Pickavet, M., Kuipers, F., Koutitas, G., Nivais, H., Satsiou, A., Beck, M., Fischer, A.: Overview of ICT energy consumption. *Netw. Excel. Internet Sci. Deliv.* **D8 1**(02), 1–59 (2013)
3. Gelenbe, E., Caseau, Y.: The impact of information technology on energy consumption and carbon emissions. *Ubiquity* **2015**(June), 1 (2015)
4. Koomey, J.: Growth in Data Center Electricity Use 2005 to 2010. A report by Analytical Press, completed at the request of The New York Times, p. 9 (2011)
5. Delforge, P.: Americas data centers consuming and wasting growing amounts of energy. In: *Natural Resource Defence Council* (2014)
6. Power, E.N.: Energy Logic: Reducing Data Center Energy Consumption by Creating Savings that Cascade Across Systems. White paper, Emerson Electric Co, vol. 9 (2009)

7. Liu, J., Zhao, F., Liu, X., He, W.: Challenges towards elastic power management in internet data centers. In: 29th IEEE International Conference on Distributed Computing Systems Workshops, 2009. ICDCS Workshops' 09, pp. 65–72. IEEE (2009)
8. Chang, R.S., Chang, H.P., Wang, Y.T.: A dynamic weighted data replication strategy in data grids. In: IEEE/ACS International Conference on Computer Systems and Applications, 2008. AICCSA 2008, pp. 414–421. IEEE (2008)
9. Strunk, A., Dargie, W.: Does live migration of virtual machines cost energy? In: 2013 IEEE 27th International Conference on Advanced Information Networking and Applications (AINA), pp. 514–521. IEEE (2013)
10. Huang, Q., Gao, F., Wang, R., Qi, Z.: Power consumption of virtual machine live migration in clouds. In: 2011 Third International Conference on Communications and Mobile Computing (CMC), pp. 122–125. IEEE (2011)
11. Baliga, J., Ayre, R.W., Hinton, K., Tucker, R.S.: Green cloud computing: balancing energy in processing, storage, and transport. Proc. IEEE **99**(1), 149–167 (2011)
12. Hinton, K., Baliga, J., Feng, M., Ayre, R., Tucker, R.S.: Power consumption and energy efficiency in the internet. IEEE Netw. **25**(2), 6–12 (2011)
13. Shang, L., Peh, L.S., Jha, N.K.: Dynamic voltage scaling with links for power optimization of interconnection networks. In: The Ninth International Symposium on High-Performance Computer Architecture, 2003. HPCA-9 2003. Proceedings, pp. 91–102. IEEE (2003)
14. Nedeveschi, S., Popa, L., Iannaccone, G., Ratnasamy, S., Wetherall, D.: Reducing network energy consumption via sleeping and rate-adaptation. NSDI **8**, 323–336 (2008)
15. Zheng, K., Wang, X., Li, L., Wang, X.: Joint power optimization of data center network and servers with correlation analysis. In: INFOCOM, 2014 Proceedings IEEE, pp. 2598–2606. IEEE (2014)
16. Jin, H., Cheochnrgngarn, T., Levy, D., Smith, A., Pan, D., Liu, J., Pissinou, N.: Joint host-network optimization for energy-efficient data center networking. In: 2013 IEEE 27th International Symposium on Parallel and Distributed Processing (IPDPS), pp. 623–634. IEEE (2013)
17. Jiang, J.W., Lan, T., Ha, S., Chen, M., Chiang, M.: Joint VM placement and routing for data center traffic engineering. In: INFOCOM, 2012 Proceedings IEEE, pp. 2876–2880. IEEE (2012)
18. Barkat, A., Capone, A.: Effective management of green cloud data centers using energy storage technologies. In: 2015 23rd International Conference on Software, Telecommunications and Computer Networks (SoftCOM), pp. 27–31. IEEE (2015)
19. Hermenier, F., Lorca, X., Menaud, J.M., Muller, G., Lawall, J.: Entropy: a consolidation manager for clusters. In: Proceedings of the 2009 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, pp. 41–50. ACM (2009)
20. Verma, A., Ahuja, P., Neogi, A.: pMapper: power and migration cost aware application placement in virtualized systems. In: Middleware 2008, pp. 243–264. Springer, Berlin (2008)
21. Dupont, C., Hermenier, F., Schulze, T., Basmaadjian, R., Somov, A., Giuliani, G.: Plug4green: a flexible energy-aware vm manager to fit data centre particularities. Ad Hoc Netw. **25**, 505–519 (2015)
22. Li, B., Li, J., Huai, J., Wo, T., Li, Q., Zhong, L.: Enacloud: an energy-saving application live placement approach for cloud computing environments. In: IEEE International Conference on Cloud Computing, 2009. CLOUD'09, pp. 17–24. IEEE (2009)
23. Ghribi, C., Hadji, M., Zeglache, D.: Energy efficient VM scheduling for cloud data centers: exact allocation and migration algorithms. In: 2013 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), pp. 671–678. IEEE (2013)
24. Guo, C., Lu, G., Wang, H.J., Yang, S., Kong, C., Sun, P., Wu, W., Zhang, Y.: Secondnet: a data center network virtualization architecture with bandwidth guarantees. In: Proceedings of the 6th International CConference, p. 15. ACM (2010)
25. Guo, C., Wu, H., Tan, K., Shi, L., Zhang, Y., Lu, S.: Dcell: a scalable and fault-tolerant network structure for data centers. ACM SIGCOMM Comput. Commun. Rev. **38**(4), 75–86 (2008)
26. Itani, W., Ghali, C., Kayssi, A., Chehab, A., Elhadj, I.: G-route: an energy-aware service routing protocol for green cloud computing. Cluster Comput. **18**(2), 889–908 (2015)
27. Baldi, M., Ofek, Y.: Time for a “greener” internet. In: IEEE International Conference on Communications Workshops, 2009. ICC Workshops 2009, pp. 1–6. IEEE (2009)
28. Restrepo, J.C.C., Gruber, C.G., Machuca, C.M.: Energy profile aware routing. In: IEEE International Conference on Communications Workshops, 2009. ICC Workshops 2009, pp. 1–5. IEEE (2009)
29. Bianzino, A.P., Chaudet, C., Larroca, F., Rossi, D., Rougier, J.L.: Energy-aware routing: a reality check. In: GLOBECOM Workshops (GC Wkshps), 2010 IEEE, pp. 1422–1427. IEEE (2010)

30. Bolla, R., Davoli, F., Bruschi, R., Christensen, K., Cucchietti, F., Singh, S.: The potential impact of green technologies in next-generation wireline networks: is there room for energy saving optimization? *IEEE Commun. Mag.* **49**(8), 80–86 (2011)
31. Chiaraviglio, L., Ciullo, D., Mellia, M., Meo, M.: Modeling sleep modes gains with random graphs. In: 2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), pp. 355–360. IEEE (2011)
32. Addis, B., Capone, A., Carello, G., Gianoli, L.G., Sanso, B.: Multi-period traffic engineering of resilient networks for energy efficiency. In: Online Conference on Green Communications (Green-Com), 2012 IEEE, pp. 14–19. IEEE (2012)
33. Amaldi, E., Capone, A., Gianoli, L.G., Mascetti, L.: Energy management in IP traffic engineering with shortest path routing. In: 2011 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), pp. 1–6. IEEE (2011)
34. Nam, T.M., Thanh, N.H., Thu, N.Q., Hieu, H.T., Covaci, S.: Energy-aware routing based on power profile of devices in data center networks using SDN. In: 2015 12th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON), pp. 1–6. IEEE (2015)
35. Chiaraviglio, L., Mellia, M., Neri, F.: Energy-aware networks: reducing power consumption by switching off network elements. In: FEDERICA-Phosphorus tutorial and workshop (TNC2008) (2008)
36. Shang, Y., Li, D., Xu, M.: A comparison study of energy proportionality of data center network architectures. In: 2012 32nd International Conference on Distributed Computing Systems Workshops (ICDCSW), pp. 1–7. IEEE (2012)
37. Wang, L., Zhang, F., Aroca, J.A., Vasilakos, A.V., Zheng, K., Hou, C., Li, D., Liu, Z.: A general framework for achieving energy efficiency in data center networks. *arXiv preprint arXiv:1304.3519* (2013)
38. Liu, R., Gu, H., Yu, X., Nian, X.: Distributed flow scheduling in energy-aware data center networks. *IEEE Commun. Lett.* **17**(4), 801–804 (2013)
39. Buysse, J., Georgakilas, K., Tzanakaki, A., De Leenheer, M., Dhoedt, B., Develder, C., Demeester, P.: Calculating the minimum bounds of energy consumption for cloud networks. In: 2011 Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN), pp. 1–7. IEEE (2011)
40. Shirayanagi, H., Yamada, H., Kenji, K.: Honeyguide: a VM migration-aware network topology for saving energy consumption in data center networks. *IEICE Trans. Inf. Syst.* **96**(9), 2055–2064 (2013)
41. Huu, T.N., Ngoc, N.P., Thu, H.T., Ngoc, T.T., Minh, D.N., Tai, H.N., Quynh, T.N., Hock, D., Schwartz, C.: Modeling and experimenting combined smart sleep and power scaling algorithms in energy-aware data center networks. *Simul. Model. Pract. Theory* **39**, 20–40 (2013)
42. Fang, W., Liang, X., Li, S., Chiaraviglio, L., Xiong, N.: Vmplanner: optimizing virtual machine placement and traffic flow routing to reduce network power costs in cloud data centers. *Comput. Netw.* **57**(1), 179–196 (2013)
43. Addis, B., Capone, A., Carello, G., Gianoli, L., Sansó, B.: Energy management through optimized routing and device powering for greener communication networks. *IEEE/ACM Trans. Netw.* **22**(1), 313–325 (2014)
44. Huang, L., Jia, Q., Wang, X., Yang, S., Li, B.: Pcube: improving power efficiency in data center networks. In: 2011 IEEE International Conference on Cloud Computing (CLOUD), pp. 65–72. IEEE (2011)
45. Guo, C., Lu, G., Li, D., Wu, H., Zhang, X., Shi, Y., Tian, C., Zhang, Y., Lu, S.: Bcube: a high performance, server-centric network architecture for modular data centers. *ACM SIGCOMM Comput. Commun. Rev.* **39**(4), 63–74 (2009)
46. Heller, B., Seetharaman, S., Mahadevan, P., Yiakoumis, Y., Sharma, P., Banerjee, S., McKeown, N.: Elastictree: saving energy in data center networks. *NSDI* **10**, 249–264 (2010)
47. Dong, J., Jin, X., Wang, H., Li, Y., Zhang, P., Cheng, S.: Energy-saving virtual machine placement in cloud data centers. In: 2013 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), pp. 618–624. IEEE (2013)
48. Gu, L., Zeng, D., Li, P., Guo, S.: Cost minimization for big data processing in geo-distributed data centers. *IEEE Trans. Emerg. Top. Comput.* **2**(3), 314–323 (2014)
49. Gu, L., Zeng, D., Guo, S., Ye, B.: Joint optimization of VM placement and request distribution for electricity cost cut in geo-distributed data centers. In: 2015 International Conference on Computing, Networking and Communications (ICNC), pp. 717–721. IEEE (2015)

50. Robertazzi, T.: Introduction to Computer Networking. Springer, Berlin (2017)
51. Amazon Web Services: Infrastructure mondiale. <https://aws.amazon.com/fr/about-aws/global-infrastructure/>. Accessed 13 Nov 2017
52. Taheri, M., Ansari, N.: Power-aware admission control and virtual machine allocation for cloud applications. In: 2015 36th IEEE Sarnoff Symposium, pp. 134–139. IEEE (2015)
53. Wolke, A.: Energy efficient capacity management in virtualized data centers. Ph.D. thesis, Universität München (2015)
54. Barati, M., Sharifian, S.: A hybrid heuristic-based tuned support vector regression model for cloud load prediction. *J. Supercomput.* **71**(11), 4235–4259 (2015)
55. Hines, M.R., Gopalan, K.: Post-copy based live virtual machine migration using adaptive pre-paging and dynamic self-ballooning. In: Proceedings of the 2009 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, pp. 51–60. ACM (2009)
56. Clark, C., Fraser, K., Hand, S., Hansen, J.G., Jul, E., Limpach, C., Pratt, I., Warfield, A.: Live migration of virtual machines. In: Proceedings of the 2nd Conference on Symposium on Networked Systems Design and Implementation, vol. 2, pp. 273–286. USENIX Association (2005)
57. Miller, R.: Facebook installs solar panels at new data centers. *Datacenter Knowl.* (2011). <http://www.datacenterknowledge.com/archives/2011/04/16/facebook-installs-solar-panels-at-new-data-center/>. Accessed 13 Nov 2017
58. Miller, R.: Wind-powered data center in wyoming (2014). <http://www.datacenterknowledge.com/archives/2007/11/29/wind-powered-data-center-in-wyoming/>. Accessed 13 Nov 2017
59. International Renewable Energy Agency (IRENA), Summary for policy makers: renewable power generation costs (2012). http://www.irena.org/DocumentDownloads/Publications/Renewable_Power_Generation_Costs.pdf. Accessed 13 Nov 2017
60. Google.com: Data center locations (2016). <https://www.google.com/about/datacenters/inside/locations/index.html>. Accessed 13 Nov 2017
61. Google: Efficiency of data centers. <http://www.google.ca/about/datacenters/efficiency/internal/>. Accessed 13 Nov 2017
62. FERC: Electric power markets—new england (iso-ne). <https://www.ferc.gov/market-oversight/mkt-electric/new-england.asp>. Accessed 13 Nov 2017
63. California ISO. <https://www.caiso.com/market/Pages/default.aspx>. Accessed 13 Nov 2017
64. FERC: Electric Power Markets. <https://www.ferc.gov/market-oversight/mkt-electric/pjm.asp>. Accessed 13 Nov 2017
65. Independent Electricity System Operator (IESO). <http://www.ieso.ca/>. Accessed 13 Nov 2017
66. GME—Gestore dei Mercati Energetici SpA. <http://www.mercatoelettrico.org/it/>. Accessed 13 Nov 2017
67. PEGAS Spot. <http://www.powernext.com/>. Accessed 13 Nov 2017
68. EPEX SPOT. <https://www.apxgroup.com/press-releases/apx-endex-to-separate-into-derivatives-gas-spot-exchange-and-power-spot-exchange/>. Accessed 13 Nov 2017
69. EEX Group. <https://www.eex.com/en/>. Accessed 13 Nov 2017
70. SEMO—Single Electricity Market Operator. <http://www.sem-o.com/Pages/default.aspx>. Accessed 13 Nov 2017
71. JEPX. <http://www.jepx.org/english/>. Accessed 13 Nov 2017
72. Trade System Administrator (ATS). <https://www.interrao.ru/en/activity/trading/>. Accessed 13 Nov 2017
73. Juniper Networks, E320 broadband services router specifications (2011). <https://www.juniper.net/documentation/hardware/erx/junose151/hw-e120-e320-hardware/index.html>. Accessed 13 Nov 2017
74. Internetworldstats: World internet users statistics and 2015 world population stats (2016). <http://www.internetworldstats.com/stats.htm>. Accessed 13 Nov 2017
75. Amazon Web Services: Ec2 instance types amazon web services (aws) (2016). <https://aws.amazon.com/ec2/instance-types/>. Accessed 13 Nov 2017
76. Barroso, L.A., Clidaras, J., Hölzle, U.: The datacenter as a computer: an introduction to the design of warehouse-scale machines. *Synth. Lect. Comput. Archit.* **8**(3), 1–154 (2013)
77. Diaz-González, F., Sumper, A., Gomis-Bellmunt, O., Villafafila-Robles, R.: A review of energy storage technologies for wind power applications. *Renew. Sustain. Energy Rev.* **16**(4), 2154–2171 (2012)
78. IBM: Cplex optimizer. <http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/>. Accessed 13 Nov 2017

Amine Barkat is a Ph.D. student at Politecnico di Milano, Italy. He received B.S. in Computer Science and M.S. in Engineering of Networks and Communications from University of Batna, Algeria. He started the Ph.D. in Cloud Computing at University of Bejaia, Algeria. In 2013, he joined Politecnico di Milano under the European Project GreenIT. His main doctoral research focuses on Cloud computing systems, virtualization and networks energy efficiency.

Mohand-Tahar Kechadi is Professor at School of Computer Science and Informatics, University College Dublin (UCD), Ireland. He received Ph.D. and M.S. in Computer Science from University of Lille 1, France. He is currently funded investigator in the Insight Centre for data analytics, UCD. His research interests include data mining, cloud and grid computing, and forensic computing and cybercrime investigations. He is EiC of the Journal of Computer Science of Science Publication and associate editor of the Journal of Future Generation of Computer Systems, Elsevier. He is a member of the International Knowledge Cloud Consortium (IKCC) and CERN.

Giacomo Verticale is Assistant Professor at Politecnico di Milano, Italy. He obtained his Ph.D. in Telecommunication Engineering in year 2003 from Politecnico di Milano. In the years 1999–2001 he was with the research center CEFRIEL, working on the Voice-over-IP and ADSL technologies. He was involved in several European research projects advancing the Internet technology. His current interests focus on the security issues of the Smart Grid and on Network Function Virtualization. He is a member of IEEE and ACM.

Ilario Filippini is Assistant Professor at Politecnico di Milano, Italy. He received M.S. and Ph.D. from Politecnico di Milano, both in Electrical Engineering. His research interests include planning and optimization of wired and wireless networks, traffic management in software defined networks, and resource management in wireless access networks. He serves as Editor of Wiley/Hindawi Wireless Communications and Mobile Computing and Elsevier Computer Networks.

Antonio Capone is Full Professor at Politecnico di Milano, Italy, where he is the Director of the ANTLab. His main research activities include radio resource management in wireless networks, traffic management in software defined networks, network planning, and optimization. On these topics he has published more than 200 peer-reviewed papers. He was an Editor of the ACM/IEEE Transactions on Networking from 2010 to 2014. He is an Editor of the IEEE Transactions on Mobile Computing, Computer Networks, and Computer Communications.