Cost-efficient live VM migration based on varying electricity cost in optical cloud networks

Abhishek Gupta[1] · **Uttam Mandal**[1] · **Pulak Chowdhury**[1] · **Massimo Tornatore**[1,2] · **Biswanath Mukherjee**[1]

Received: 14 May 2015 / Accepted: 7 August 2015 / Published online: 4 September 2015

1 Introduction

Cloud-based services have become the norm for enterprises for reasons of scalability and cost-effectiveness. The emergence of cloud-based multimedia services such as Youtube, Pandora, etc. has created a spurt in user demand. This has led cloud-service providers to deploy more large-scale computing infrastructure (DCs) to satisfy user demand. DCs have huge energy usage and cloud infrastructure; and service providers such as Amazon, Microsoft, and Google which have multiple DCs spread geographically across the globe and run thousands of servers, spend a significant part of their operating expenditure on energy costs.

Physical servers in a DC account for upto 52 % of DC energy consumption [2]. Since maximizing the utilization of physical servers is of utmost importance to reduce the energy demand of a DC, VMs (i.e., virtual servers, used to serve different services, that can be hosted in the same physical server) have been successfully introduced in DCs for higher utilization of server resources.

By utilizing VMs, we create virtualized workloads which can be migrated using live VM migration [3]. This provides a technique to move VMs to DCs having cheaper electricity prices, which in turn reduces the energy costs incurred in

✉ Abhishek Gupta
abgupta@ucdavis.edu

Uttam Mandal
umandal@ucdavis.edu

Pulak Chowdhury
pchowdhury@ucdavis.edu

Massimo Tornatore
mtornatore@ucdavis.edu

Biswanath Mukherjee
bmukherjee@ucdavis.edu

1 Department of Computer Science, University of California, Davis, CA, USA

2 Politecnico di Milano, Milan, Italy

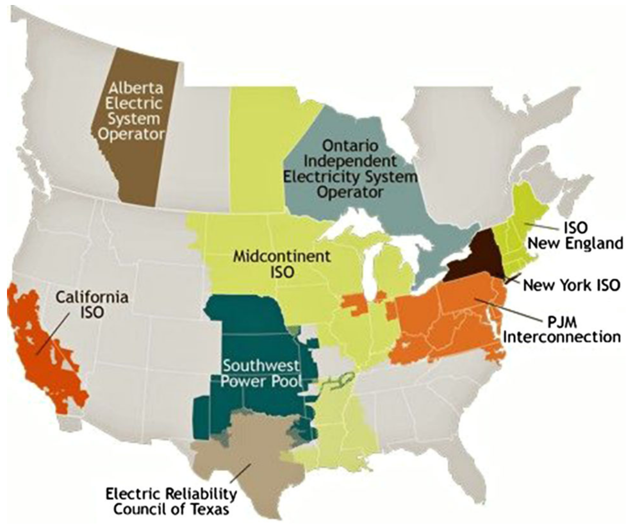

Fig. 1 ISOs/RTOs in North America [1]

operation. Electricity prices are available hourly through the Independent System Operators/Regional transmission Organizations (ISOs/RTOs) shown in Fig. 1. So, we can migrate VMs at the beginning of each hour to the cheapest DC. However, it should be noted that migration done on a hourly basis may be costlier when analyzed with a multi-hour perspective, i.e., although cost at a DC currently may be higher, it may be the cheapest for the remaining period of time. Without a multi-hour perspective, we could be performing unnecessary migrations which would reduce the effectiveness of our scheme.

In this work, we take a multi-hour approach while considering the price of electricity at other candidate DC locations, cost of migration, and cost related to turning servers and racks on/off, to make VM migration decisions. We find that migrations performed using our approach are cost-efficient for the multi-hour period in consideration.

2 Dynamic electricity pricing

In this work, we study how energy cost for DC operation can be reduced by considering the *spatial and temporal variation in electricity cost* that can be observed across a large country, such as the USA. Differences in electricity prices among different regions in USA can be as large as 11 times [4], due to factors such as type of fuels, power plants, transmission and distribution lines, supply and demand, weather conditions, and regulations. These electricity prices are regionally regulated by organizations called Independent System Operator (ISO)/Regional Transmission Organization (RTO). Electricity prices can be flatly marked or other schemes such as dynamic electricity pricing are employed. Dynamic electricity pricing is a model where the cost of electricity changes every hour. It has to be mentioned that not all grid markets in the USA have dynamic electricity pricing, and the ISOs/RTOs shown in Fig. 1 are the major ones in North America which offer dynamic pricing of electricity.

There are seven major ISOs/RTOs in the USA that offer dynamic pricing, namely California ISO (CAISO), Southwest Power Pool ISO (SPP), Midcontinent ISO (MISO), Electric Reliability Council of Texas (ERCOT), New York ISO (NYISO), New England ISO, and PJM Interconnection (PJM). These ISOs/RTOs are responsible for regulating wholesale market prices and managing the electricity grid in the region among other functions (here we consider that ISOs/RTOs have the same functions though subtlety exists in defining them). Figure 2 shows the variation in electricity prices over a 24-h period across ISOs/RTOs in the USA for July 17, 2014. The prices have been synchronized according to Eastern Daylight Time (EDT) to get the perspective of variation in electricity prices across regions and time.

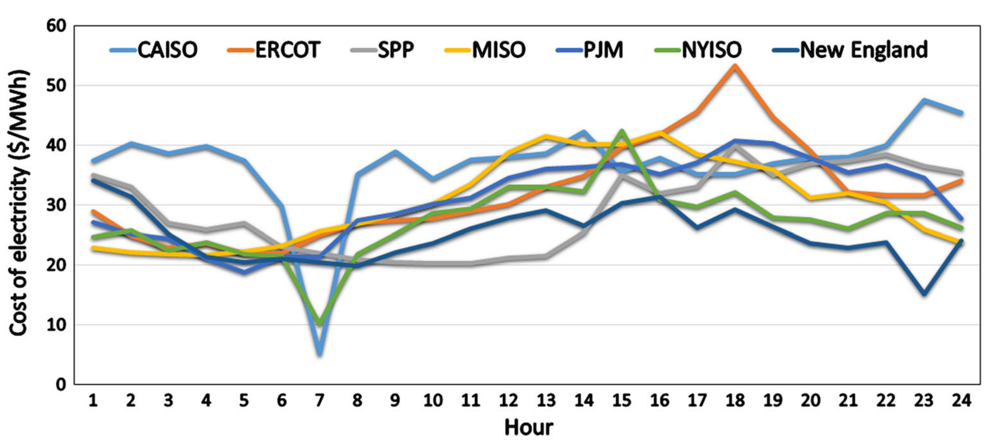

Fig. 2 Twentyfour hour electricity prices synchronized to EDT across ISOs/RTOs in United States for July 17, 2014

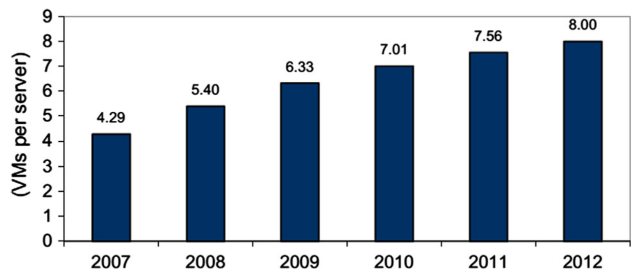

Fig. 3 Server virtualization densities [5]

3 Virtual machine migration

Cloud infrastructure providers virtualize physical servers to create virtual machines (VMs) for easy and secure resource allocation in a multi-tenant DC environment. VMs help in better utilization of physical resources, thereby reducing server deployment rates, which is essential since DC infrastructure build-up is capital intensive.

VMs being virtual servers are used for deploying services. As service requests increase, more VMs may need to be instantiated. The number of VMs that can be hosted on a physical server depends on the servers's physical configuration. Figure 3 shows data on the virtualization densities of a physical server in recent years. This trend reinforces the notion of virtualization of physical resources using VMs. One of the major reasons for such an increase in VM density per server is better physical server configuration.

In this work, we intend to move VMs to DCs with cheaper electricity prices without disrupting the existing connections. *Live VM migration* has been used to move VMs to achieve cloud bursting, geographical load balancing, IT consolidation within a single DC, etc. While VM migration over a local area network (LAN), e.g., inside a DC, is already widely adopted, migration over a wide area network (WAN) has started to receive a great deal of research and industry attention.

VM migration over a WAN is carried out by establishing a network connection between the source and destination locations of the VM, then transferring the memory and disk storage states, and finally re-configuring the VM and resuming the activity at the destination site. Memory (RAM) transfer occurs in three phases. First, a full memory copy is performed. To account for the modified or dirtied memory during this period and successive transfer periods, an iterative copy phase is started, as shown in Fig. 4. In this phase, the dirtied memory after the last iteration is copied to the destination. Transferred data at each iteration depend on the provisioned network bandwidth. Service is always available at the source location during this phase. When a certain predefined stop condition is met, the iterative copy phase is stopped, and the VM is halted for a short duration for the final memory copy. Before resuming the VM at the destination site, the network is reconfigured and existing connections are migrated to the destination location. The VM is down during this last phase. Disk storage state is migrated similarly as memory states. However, because disk-dirtying rates are generally lower than memory-dirtying rates, initial disk-copy phase consumes most resources and time. From the network's resource point of view, a VM migration requires a high-bandwidth connection between the source and destination DCs. Migration time and service downtime are reduced if higher network bandwidth is provided. Note that, in this work, we employ the same bandwidth for VM migration for the whole duration of migration in order to create a simple and effective migration model. The final memory size (when significant) of the VM may require additional bandwidth to be deployed to meet latency constraints.

Figure 5 shows the variation in migration duration as the bandwidth allocated for migration is increased. The duration of a live migration depends nonlinearly on the bandwidth allocated for migration (for lower bandwidths, more memory at the source node will be *dirtied* during migration, and this penalizes the overall migration time). For this reason, we utilize the function-by-point model which depicts VM migration characteristics taken from [7]. In this function-by-point model, each function point is a tuple of bandwidth and migration time, i.e., it gives the migration time for the band-

Fig. 4 Iterative phases of VM migration [3]

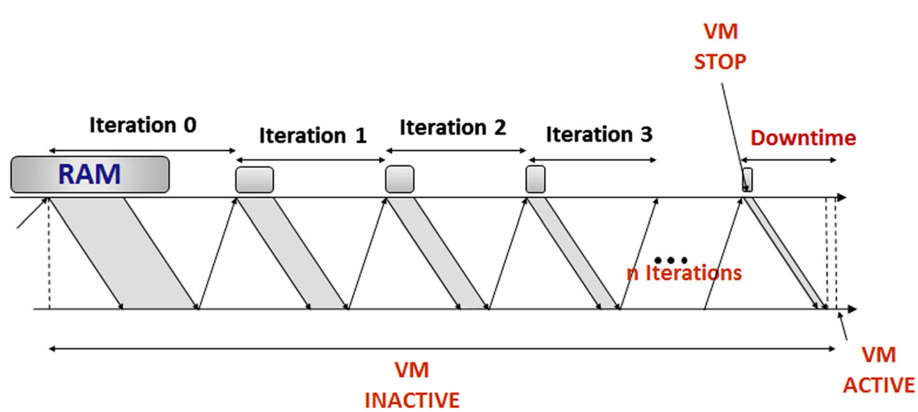

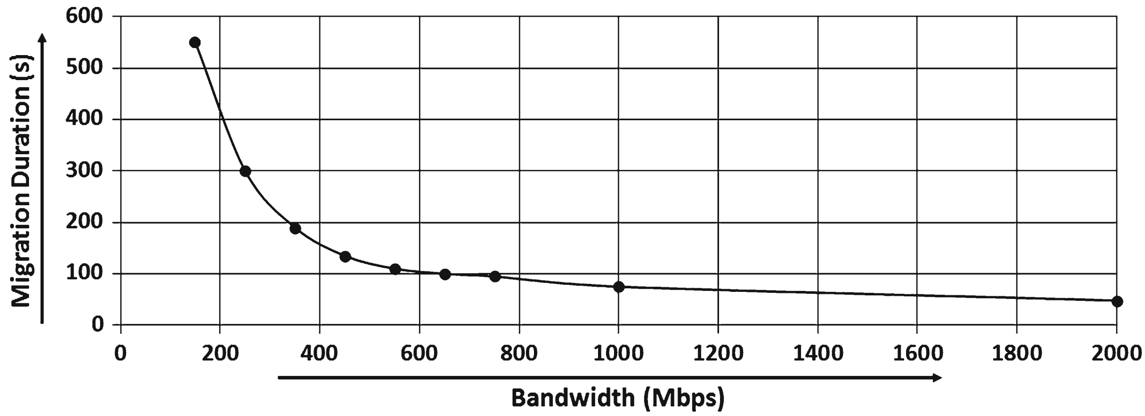

Fig. 5 VM migration time as a function of provided bandwidth [6]

width deployed for VM migration. Note that this migration time also includes the down time for the VM migration.

An important aspect of live VM migration is to keep existing client connections to the VM alive during and after migration, i.e., the migration is seamless and does not impact client's perceptions. This requires the existing connections not to be taken down and is ensured by creating a virtual LAN between the source and destination DCs [6,8]. Using this VLAN, the existing connections are kept alive by routing through the following path: client to the source DC, and then to the destination DC. This is achieved by appending to the existing client connection an additional path from source DC to the destination DC.

4 Related work

In this section, we review related studies which have investigated minimizing the cost of cloud operation as well as those which have utilized variation in electricity cost for reducing the carbon footprint among broader works which can be included under Green Cloud Networks.

Energy efficiency in cloud networks has been the focus of a number of previous works. The authors in [9] present a short survey of the current technologies and tools used to achieve energy efficiency in cloud networks, while Ref. [10] analyzes the energy consumption in transmission and switching networks which connect users to the cloud. With respect to switching networks, Ref. [11] provides a detailed survey of the techniques and approaches for energy efficiency in optical networks.

Authors in [4] use load balancing to reduce energy cost in ICT. However, Ref. [12] shows that it can actually lead to an increase in total energy use and simultaneously investigates the environmental gains that can be exploited from the system. Algorithms developed in [13] to dynamically route optical paths to transfer jobs to renewable-energy-rich DCs can be used to reduce energy cost while increasing the environmental gain. Ref. [14] reduces the energy consumption by focusing on reducing ISP power consumption cost. This work employs various strategies such as concentrating traffic on a minimal set of resources, shutting off nodes, and increasing link utilizations, among others, to achieve this. The authors in [15] go further and aim to reduce the green house gas (GHG) emissions associated with a network node along with the energy consumption using an improved Manycast Drop at Member Node (MA-DMN) overlay algorithm.

While reduction in energy consumption is important, the trade-off must not adversely affect the user's Quality of Experience (QoE). Ref. [16] considers minimizing the energy consumption by reducing the operating cost of DCs but also guarantee Quality of Service (QoS). This is achieved by employing a Mixed Integer Linear Program (MILP). Similarly, Ref. [17] maintains service-level objectives (SLOs) while providing an optimization-based framework to reduce the brown energy (energy obtained from carbon-rich and non-renewable sources such as coal and nuclear energy) consumption while improving green energy (energy obtained from renewable sources such as solar, wind, and hydro) usage in a multi-DC environment.

5 Problem description

The operating cost of VMs in a DC will vary depending on the ISO/RTO operating the grid where the DC is located. This is so because VMs are hosted on physical servers which consume power, depending on their power rating (amount of electrical power required for a particular device) and utilization. The cost of the power consumed will depend on the local cost of electricity, so cheaper electricity prices result in lower operating costs. To exploit this variation in prices, we use live VM Migration to migrate VMs to DCs in regions with cheaper electricity prices.

We model the problem of minimizing the operating cost of VMs by employing live VM migration over a multi-hour period and using a mixed-integer linear program (MILP) formulation. The idea of using the variation in electricity prices to migrate DC load toward cheaper DCs was originally proposed in [4], but in this study we consider live VM migration as a novel means to migrate load. In previous related work [18], a partial and dynamic version of this problem was solved, with a specific focus on the maximization of renewable energy usage.

5.1 Power model

We look at the power consumption, both at the source and destination DCs, as well as the power consumed in the network during migration.

5.1.1 Datacenter power consumption

VMs are hosted on physical servers in DCs. DC space is expensive as it also needs power and cooling. To best utilize this space as well as to help the cooling requirements of physical servers, racks are used as shown in Fig. 6. The standard size for racks is 42U (where 'U' is a metric unit equal to 1.75 in.). Physical servers can have variable size (in integral multiples of U), depending on their configuration and occupancy of rack slots. The number of occupied rack slots depends on the space available for server storage (space is also required for cooling and power equipment) and the size of each physical server.

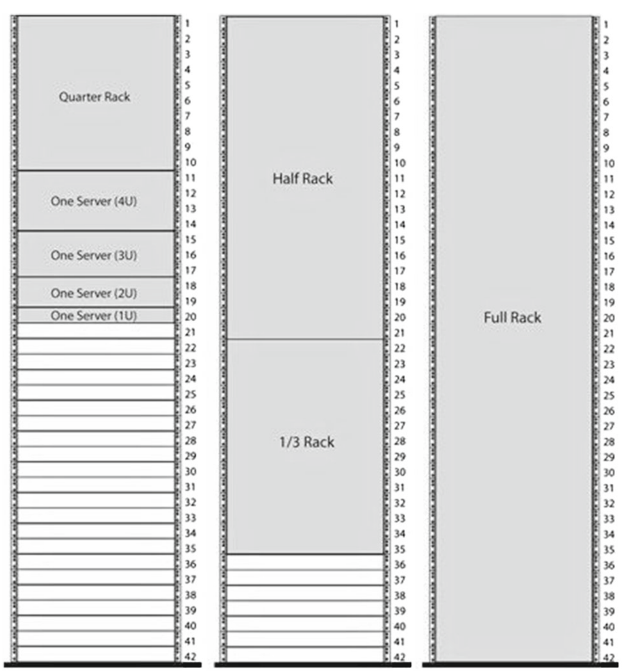

Fig. 6 Rack and physical server deployments [19]

The base power consumption of a physical server is the power consumed by it in idle state. The power utilized by a server is directly proportional to the load on the server until it reaches the maximum power consumption. To reduce the complexity of the problem, we assume a server to be in an idle state or in an active state (when it is serving requests). Rack power is calculated as the power required to keep all the servers switched on when in an idle state. For example, typical values of power consumed by a rack and a server are 3300 and 275 W [20], respectively. As a result, we know the base power consumed by a rack; and, as the number of active servers in the rack increases, the power consumption also increases until it reaches full utilization. This is done to consolidate VMs in a server and its corresponding rack. By consolidating VMs in a single rack, we can switch off any idle racks and avoid the overhead of switching on more racks.

We denote the power consumed in a DC by P_{DC}, power consumed by an active rack as P_{R} which is the power consumed in turning the rack on, power consumed by an active server as P_{S}, number of active racks as y, maximum number of servers per rack as β, and number of active servers as x ($x \leq \beta * y$). By the definition of our power model, the following relation holds:

$$P_{\mathrm{DC}} = y * P_{\mathrm{R}} + x * P_{\mathrm{S}}$$

5.1.2 Migration power consumption

Backbone networks, which connect DCs, consume power when migrating a VM from a source DC to a destination DC. This power consumption is calculated as a product of the total number of bits transferred (bandwidth-duration product) in a VM migration, the power consumed by a core router in transmitting one bit,[1] and the cost of electricity at the core router.

5.2 Mathematical formulation

VM migration time is modeled as a set of function points derived from [6]. The function-by-point model is a set of bandwidth and migration duration tuples. The function gives the migration duration for the bandwidth used and vice versa while giving the MILP the flexibility to choose the bandwidth for VM migration. The energy cost (or *migration cost*) to migrate a VM depends on the migration power consumption and the cost of electricity used at the intermediate core routers. A VM is enabled at the destination DC when migration begins. During migration, the VM is kept switched on at the source DC. After migration completion, the racks and

[1] Core routers such as a Cisco CRS-1 multi-shelf system can give a switching capacity of 92 Tb/s full duplex while consuming 1020 kW [21], so the power consumed in transmitting one bit is 1.108×10^{-5} W.

servers at the source DC which were hosting the migrated VMs can be powered off. We refer to the power consumed by racks and servers at the source node during migration (i.e., energy consumed before they can be switched off) as *residual power*. Residual power is estimated by accounting for the maximum migration time possible by a VM in order to reduce the complexity of the problem.

Our model decides on VM migration based on the trade-off in the operating costs (at the source and possible destination DCs) and the migration cost/residual power cost. As each migration entails a penalty due to residual power and migration cost, the decision to migrate is taken only if the savings from migration significantly exceed the penalties.

The resulting formal optimization-problem statement is as follows: Given an optical backbone network topology, a set of DC nodes, initial locations of each VM, hourly prices of electricity at each node, link capacities, maximum number of VMs a DC can host, and a multi-hour period, our objective is to minimize the operating cost of the VMs over this period by deciding whether (including when and where) or not to migrate the VMs. We formulate the problem as an MILP as follows:

5.2.1 Input parameters

- $G(V, E)$: Physical topology of the network; V is set of nodes and E is set of links.
- $D \in V$: Set of candidate DC locations.
- F: Set of function points containing VM migration time.
- b^f: Migration bandwidth of function point f.
- d^f: Migration duration of function point f.
- H: Given multi-hour time period.
- $h \in H$: A hour duration in given multi-hour time period.
- ζ_h^i: Cost of electricity at node $i \in V$ for hour 'h'.
- VM: Set of VMs in the system.
- ϑ: Maximum VM capacity for a DC.
- δ: Maximum migration duration for a VM.
- $A^v_{zh_{\text{previous}}} \in \{0, 1\}$: Position of a VM v at DC $i \in D$ in hour h_{previous} before H.
- P_{R}: Power consumed by a rack (3300 W).
- θ: Number of VMs per rack.
- P_{S}: Power consumed by a physical server (275 W).
- κ: Number of VMs per server.
- P_{b}: Transmission power per bit of a router.

5.2.2 Variables

- $X_{ij}^{szhvf} \in \{0, 1\}$: 1 if link $ij \in E$ is utilized for migrating VM $v \in VM$ from source DC $s \in D$ to destination DC $z \in D$ in hour $h \in H$ using VM migration characteristics of function point $f \in F$.
- $t_f^{szvh} \in \{0, 1\}$: 1 if VM migration characteristics of f are used in VM v's migration from source DC s to destination DC z in hour h.
- $Y_h^{szv} \in \{0, 1\}$: 1 if migration of VM v occurs from source DC s to destination DC z in hour h.
- $A_{hz}^v \in \{0, 1\}$: 1 if VM v is located at DC z for hour h.
- $Q_h^{sv} \in \{0, 1\}$: 1 if VM v does not migrate from DC s for hour h.
- R_h^z: Number of racks active during h at DC z.
- S_h^z: Number of servers active during h at DC z.
- ρ_h^z: Number of racks switched off in $h \in H$ at DC $z \in D$.
- σ_h^z: Number of servers switched off in $h \in H$ at DC $z \in D$.

5.2.3 Problem formulation

$$\text{Minimize: } \sum_{h\in H}\sum_{z\in D} R_h^z * P_{\text{R}} * \zeta_h^z + \sum_{h\in H}\sum_{z\in D} S_h^z * P_{\text{S}} * \zeta_h^z + \sum_{h\in H}\sum_{z\in D} \rho_h^z * \delta * P_{\text{R}} * \zeta_h^z + \sum_{h\in H}\sum_{z\in D} \sigma_h^z * \delta * P_{\text{S}} * \zeta_h^z + \sum_{v\in VM}\sum_{s\in D}\sum_{z\in D}\sum_{h\in H}\sum_{f\in F} b^f * d^f * P_{\text{b}} * \sum_{ij} X_{ij}^{szvhf} * \zeta_h^i \tag{1}$$

$$\text{Subject to: } \sum_{i\epsilon A_j} X_{ij}^{szhvf} - \sum_{i\epsilon A_j} X_{ji}^{szhvf} = \begin{cases} -t_f^{szvh} & j = s \\ t_f^{szvh} & j = z \\ 0 & \text{otherwise} \end{cases} \quad \forall s, \forall z, \forall h, \forall v, \forall f \tag{2}$$

$$\sum_{f\in F} t_f^{szvh} = Y_h^{szv} \quad \forall s, \forall z, \forall v, \forall h \tag{3}$$

$$\vartheta_h^s - \sum_{z\in D}\sum_{v\in VM} Y_h^{szv} + \sum_{z\in D}\sum_{v\in VM} Y_h^{zsv} = \vartheta_{h+1}^s \quad \forall s, \forall h \tag{4}$$

$$\vartheta_{h+1}^s \leq \vartheta \quad \forall s, \forall h \tag{5}$$

$$\sum_{s\in D}\sum_{z\in D}\sum_{v\in VM}\sum_{f\in F} X_{ij}^{szhvf} * b^f \leq c_{ij} \quad \forall (i, j), \forall h \tag{6}$$

$$Q_h^{sv} + \sum_{z\in D} Y_h^{szv} = A_{h-1s}^v \quad \forall s, \forall v, \forall h \tag{7}$$

$$A_{hs}^v = Q_h^{sv} + \sum_{z\in D} Y_h^{zsv} \quad \forall z, \forall h, \forall v \tag{8}$$

$$R_h^z \geq \sum_{v\in VM} A_{hz}^v/\theta \quad \forall z, \forall h \tag{9}$$

$$S_h^z >= \sum_{v\in VM} A_{hz}^v/\kappa \quad \forall z, \forall h \tag{10}$$

$$\rho_h^s \geq \sum_{v\in VM}\sum_{z\in D} Y_h^{szv}/\theta \quad \forall s, \forall h \tag{11}$$

$$\sigma_h^s \geq \sum_{v\in VM}\sum_{z\in D} Y_h^{szv}/\kappa \quad \forall s, \forall h \tag{12}$$

The objective function in Eq. (1) evaluates the dollar cost of operating the VMs across all the DCs for the multi-hour period and includes the residual power cost and migration cost. Eqs. (2) and (3) establish whether a VM is migrated in a particular hour using a given function point. Eqs. (4) and (5) constrain the number of VMs hosted at a DC for placement and migration. Equation (6) sets the capacity constraint for link i, j. Equations (7) and (8) give location of the VM during an hour. Equations (9), (10), (11), and (12) calculate the number of racks and servers along with racks and servers containing migrating VMs at each DC on an hourly basis.

6 Illustrative numerical examples

We conduct simulation experiments on a 10-node topology as shown in Fig. 7. We first run simulations on a 4-h period beginning 0400 EDT and ending 0800 EDT. Three nodes, namely nodes 2, 5 and 6, are considered as DC locations. The DC nodes fall in different time zones and are geographically distributed across the US to set up an environment where the spatial and temporal variations of electricity prices can be exploited. We associated ISOs/RTOs with the nodes based on [1]. These nodes fall under three different ISOs/RTOs, i.e., node 2 falls under California ISO (CAISO), node 5 falls under ERCOT, and node 6 falls under PJM Interconnection. We consider the number of physical servers per rack to be 12. The maximum number of VMs that can be hosted on a server is 4. This is different from our previous assumption of 1 VM per server in [22].

The experiments are run for 100, 200, and 300 VMs in the system. It is assumed that we are given the initial distribution of VMs across the three DCs for each of these cases as shown in Table 1.

6.1 4-h simulations

We compare our MILP approach with the cost of running the VMs according to their given initial placements for the 4-h period, 0400-0800 EDT (Eastern Daylight Time). This

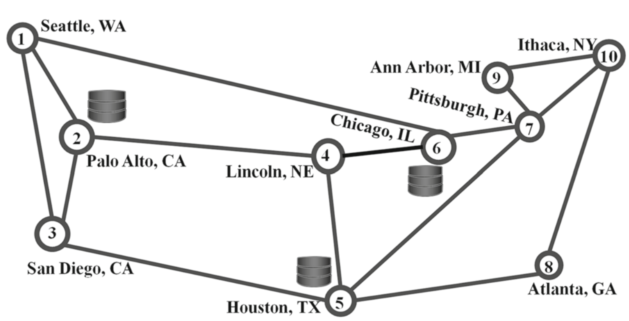

Fig. 7 Ten-node network

Table 1 Given initial VM distribution

VMs	Node 2	Node 5	Node 6
100	36	29	35
200	71	57	72
300	106	96	98

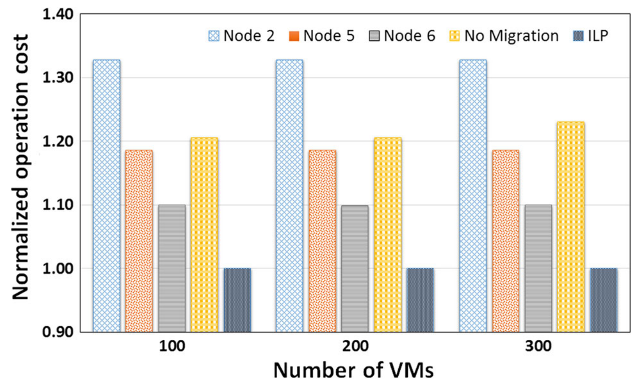

Fig. 8 Operational cost versus number of VMs (4 h)

cost is represented as “No Migration” in Fig. 8. In the “No Migration” scenario, the VMs are run in their given location, for the duration of 4 h. Since DC capacity is not constrained, we also consider the case in which VMs are all placed in a single node (2, 5, or 6), i.e., all the VMs are hosted in node 2, 5, or 6 for the duration of 4 h.

Figure 8 shows the operating cost of the previous approaches normalized with respect to the MILP result. Our MILP approach achieves from a minimum of 10 % to a maximum of 35–36 % reduction in operation costs across the 100, 200, and 300 VM systems. However, results in Fig. 8 are obtained under the assumption that all VMs can be hosted at a single DC (no DC capacity limit). Thus, to investigate more practical cases where DC capacity is limited, we enforce a Max-DC constraint and create a metric to denote the unallocated capacity in the system. We define this metric as the ratio of the number of VMs to the maximum number of VMs that can be allocated in a system and refer to it as ‘Load’ (L). For example, if we have 100 VMs in a 1000-VM capacity system, the load in the system will be 0.1 (100/1000). Similarly, the maximum number of VMs that we can have in our system is 1000, in which case the load in the system will be 1. Thus, the range of values for load will be between 0 and 1. Our aim here is to observe the effectiveness of the MILP as L increases.

To analyze the case of limited capacity, we also develop a Max and Min placement strategy. Min places VMs at the lowest-cost DC node; and, when DC capacity is exhausted, the next cheapest DC node is selected. After the VMs have been placed in this fashion, they are run in this Min placement configuration for the duration of 4 h. VMs do not change their

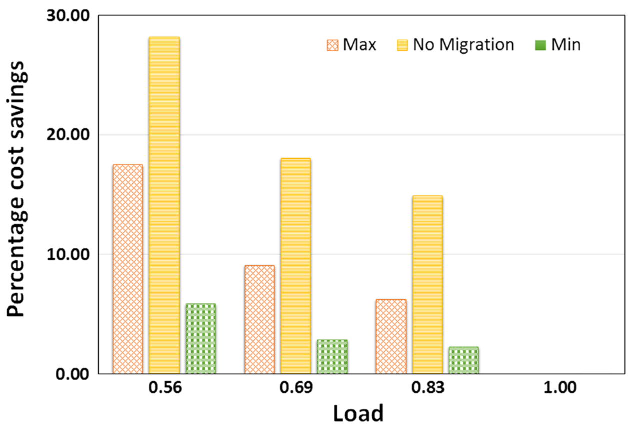

Fig. 9 Percentage cost savings versus load (4 h)

placements. Max places VMs at the costliest node, and on capacity exhaustion, the next-costliest node is selected. VMs are run in this Max placement for the duration of 4 h. VMs do not change their placements.

Figure 9 shows the percentage cost savings for our MILP approach with respect to the other approaches. We find that the cost savings gradually decrease as load L increases as less optimization in VM placements is possible at high loads. We emphasize that this is a fortuitous scenario and that it would not be possible to use a low-price DC at full capacity for a multi-hour period, while other DC capacity remains unallocated.

We demonstrated the reduction in energy costs achieved by our model. However, it will also be interesting to analyze how the migrations are dependent on hourly electricity price and other factors such as VM packing capacity of a rack, bandwidth capacity of the backbone links, etc.

Figure 10 shows the number of migrations occurring in each hour for the 100, 200, and 300 VM systems. Table 2 gives the electricity costs for 0400–0800 EDT across the three DC locations.

Figure 10 gives details of VM migrations through the 4-h period. Figure 10a shows the number of migrations that occur in each hour for the 100, 200, and 300 VM systems.

Table 2 shows DC node 6 to be cheapest for hour 0400. We find that all VMs residing at DC node 2 and DC node 5 migrate toward the cheaper DC node 6 for the 100 and 200 VM systems. However, we find that, for the 300-VM system, not all VMs residing at DC node 5 as shown in Fig. 10b (96

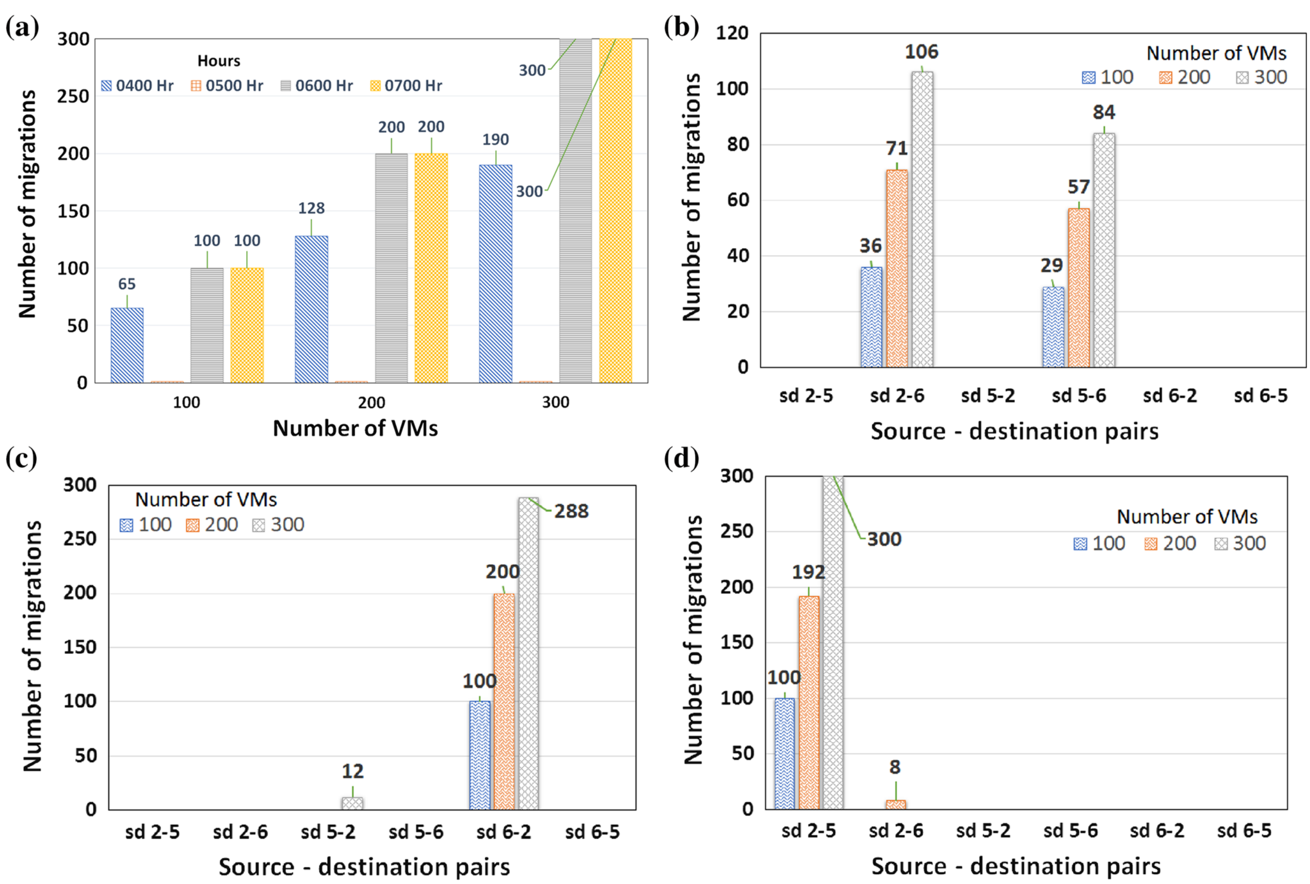

Fig. 10 VM migrations across 0400-0800 EDT for 100, 200, and 300 VM systems, **a** number of migrations in each hour (4 h), **b** number of migrations in hour starting 0400 EDT, **c** number of migrations in hour starting 0600 EDT, **d** number of migrations in hour starting 0700 EDT

Table 2 Electricity prices in \$/MWh

Hour	Node 2	Node 5	Node 6
0400	37.46	21.99	18.97
0500	29.82	22.01	21.50
0600	5.17	25.01	21.37
0700	35.19	27.01	27.51

VMs were hosted at DC node 5 as shown in Table 1) migrate to DC node 6. This comes as a result of VM consolidation in racks. A rack can host 48 VMs (a server can host 4 VMs and a rack can hold 12 physical servers). In the 300-VM system, 190 VMs migrate to DC node 6, and DC node 6 already hosts 98VMs, so in total 288 VMs are hosted at DC node. This means all the VMs are perfectly consolidated into racks.

The model does not migrate the remaining 12 VMs from DC node 5 for a variety of factors. First, migrating them would lead to turning on an underutilized rack at DC node 6 while the underutilized rack at DC node 5 will still be on for the migration period. Second, the electricity price difference between DC node 5 and DC node 6 is minimal for the hour period beginning 0500 EDT which means that running an underutilized rack at DC node 5 will not be that much costlier than DC node 6. Finally, our model provisions the highest possible bandwidth for VM migration, so as to reduce the migration duration. This leads to the cheapest routes getting utilized if there are a high number of migrations. As a result, the 12 remaining VMs may need to be migrated over a costlier route, and this migration cost penalty cannot be compensated by running those 12 VMs at DC node 6. These factors result in the model hosting the 12 VMs at DC node 5.

No migration happens in the hour starting 0500 since DC node 6 still remains the cheapest in electricity price. We find that the numbers of migrations occurring in hours starting 0600 and 0700 are equal to the total number of VMs in the system. This is as expected since we do not have a limit on the capacity of DCs in the system and also because DC node 2 is by far the cheapest in hour starting 0600, while DC node 5 is much cheaper than DC node 2 in hour starting 0700.

In hour starting 0600, DC node 2 becomes the cheapest by a large margin. As a result, we see that all VMs hosted at DC node 5 and DC node 6 migrate to DC node 2. Almost all VMs were hosted at DC node 6 for a 2-h period from 0400 to 0600 across the 100, 200, and 300 VM systems. As seen in Fig. 10c, all the VMs migrate to DC node 2, including the 12 VMs that were hosted at DC node 5.

Figure 10d indicates that all the VMs placed at DC node 2 for the hour period 0600–0700 are migrated in the hour starting 0700 for 100, 200, and 300 VM systems, but all the VMs are migrated to DC node 5 which has the cheapest electricity price. However, 8 VMs of the 200-VM system are migrated to DC node 6. This happens as DC node 5 is only

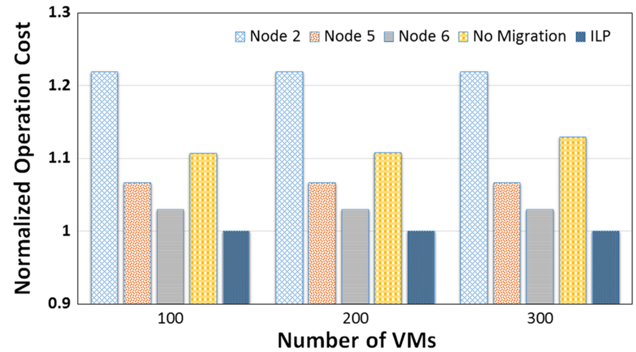

Fig. 11 Operational cost versus number of VMs (24 h)

marginally cheaper than DC node 6. As a result, not all VMs are migrated from DC node 2 to DC node 5 in hour starting 0700, because the cheapest route was already completely utilized for migrations, and the number of VMs migrated to DC node 5 were all efficiently consolidated into racks. The model decides on a trade-off between migrating all VMs to DC node 5 and paying for an underutilized rack, on the one hand, along with the migration cost incurred on the costlier route and running the VMs at DC node 6, on the other hand. The trade-off (which considers the penalty of residual power cost and migration cost vs. the gain from migration) favors not migrating all the VMs to DC node 5, and hence, the VMs are migrated to DC node 6.

6.2 24-h simulations

Ideally, we would optimize the VM migration process over a 24-h period. Unfortunately, due to complexity reasons, our model cannot scale to a 24-h optimization, but we report results for a 24-h period by simulating 4-h instances. The results will be sub-optimal, but will provide an important indicator of the daily gain achievable by our approach.

Figure 11 shows the operating costs of the various approaches in a 24-h period normalized with our MILP result. We find that the trends are similar to results in the 4-h simulation. However, savings are less in relative terms since spatial variations of costs in some periods of the day (e.g., at night) are less remarkable. Nonetheless, it is confirmed that savings on the order of 25 % can be obtained.

7 Conclusion

We investigated the problem of reducing the VM operation cost in backbone cloud networks using an MILP formulation. Live VM migration is used to migrate VMs, depending on the variation in electricity costs. Our model explicitly accounts for residual power cost and migration cost, and it also promotes consolidation of VMs. We analyzed the migration behavior of VMs and how the multiple factors considered by our model affect it. We find that our method gives significant improvement in operating cost in comparison with any initial placement of VMs. To solve this problem in a dynamic scenario is an open problem for future research.

Acknowledgments This work was supported by NSF Grant No. CNS-1217978.

References

[1] ISO/RTO Council. http://www.isorto.org/about/default
[2] Emerson Network Power. White Paper. http://www.emersonnetworkpower.com/documentation/white%20paper Liu,
[3] H., Jin, H., Xu, C., Liao, X.: Performance and energy mod-eling for live migration of virtual machines. Cluster Comput. **16**, 249–264 (2013)
[4] Qureshi, A., Weber, R., Balakrishnan, H., Guttag, J., Maggs, B.: Cutting the electric bill for internet-scale systems. SIGCOMM **39**, 123–134 (2009)
[5] VMware Technical Paper. www.f5.com/pdf/white-papers/cloud-vmotion-f5-wp
[6] Wood, T., Ramakrishnan, K., Van Der Merwe, J., Shenoy, P.: Cloudnet: A platform for optimized wan migration of virtual machines, University of Massachusetts Technical Report, TR-2010-002 (2010)
[7] Mandal, U., Habib, M., Zhang, S., Tornatore, M., Mukherjee, B.: Greening the cloud using renewable-energy-aware service migration. IEEE Netw. **27**, 36–43 (2013)
[8] VMware and F5 Technical Brief. Enabling long distance live migration with f5 and vmware vmotion. https://f5.com/resources/white-papers/enabling-long-distance-live-migration-with-f5-and-vmware-vmotion
[9] Berl, A., Gelenbe, E., Di Girolamo, M., Giuliani, G., De Meer, H., Dang, M., Pentikousis, K.: Energy-efficient cloud computing. Comput. J. **53**, 1045–1051 (2010)
[10] Baliga, J., Ayre, R., Hinton, K., Tucker, R.: Green cloud computing: balancing energy in processing, storage, and transport. Proc. IEEE **99**, 149–167 (2011)
[11] Zhang, Y., Chowdhury, P., Tornatore, M., Mukherjee, B.: Energy efficiency in telecom optical networks. IEEE Commun. Surv. Tutor. **12**, 441–458 (2010)
[12] Liu, Z., Lin, M., Wierman, A., Low, Andrew, S.L.: Greening geographical load balancing. In: Proceedings of ACM SIGMETRICS, pp. 233–244 (2011)
[13] Gattulli, M., Tornatore, M., Fiandra, R., Pattavina, A.: Low-carbon routing algorithms for cloud computing services in IP-over-WDM networks. In: Proceedings of 2012 IEEE International Conference on Communications (ICC), pp. 2999–3003 (2012)
[14] Chiaraviglio, L., Mellia, M., Neri, F.: Minimizing ISP network energy cost: formulation and solutions. IEEE/ACM Trans. Netw. **20**, 463–476 (2012)
[15] Schondienst, T., Davis, D., Plante, J., Vokkarane, V.: Renewable energy-aware manycast overlays, Technical Report, University of Massachusetts Lowell
[16] Rao, L., Liu, X., Xie, L., Liu, W.: Minimizing electricity cost: optimization of distributed internet data centers in a multi-electricity-market environment. In: Proceedings of IEEE INFOCOM (2010)
[17] Le, K., Bianchini, R., Nguyen, T., Bilgir, O., Martonosi, M.: Capping the brown energy consumption of internet services at low cost. In: Proceedings of 2010 IEEE International Green Computing Conference, pp. 3–14 (2010)
[18] Mandal, U., Habib, M., Zhang, S., Tornatore, M., Mukherjee, B.: Bandwidth and routing assignment for virtual machine migration in photonic cloud networks, In: ECOC 2013
[19] http://www.ksc.net/
[20] Hewlett Packard. Technology Brief, Power efficiency and power management in HP proliant servers. http://h10032.www1.hp.com/ctg/manual/c03161908
[21] Baliga, J., Hinton, K., Tucker, R.: Energy Consumption of the Internet. Department of Electrical and Electronic Engineering, University of Melbourne, Melbourne (2011)
[22] Gupta, A., Mandal, U., Chowdhury, P., Tornatore, M., Mukherjee, B.: Cost-efficient live VM migration based on varying electricity cost in optical cloud networks. In: IEEE International Confer-ence on Advanced Networks and Telecommunications Systems (ANTS) (2014)