International Journal of Grid and Distributed Computing
Vol.6, No.5 (2013), pp.51-60
http://dx.doi.org/10.14257/ijgdc.2013.6.5.05

Cost-aware Workload Dispatching and Server Provisioning for Distributed Cloud Data Centers

Weiwei Fang[1], Quan Zhou[1], Yuan An[2], Yangchun Li[3] and Huijing Zhang[3]

[1]*School of Computer and Information Technology, Beijing Jiaotong University, Beijing, 100044, China*
[2]*State Key Laboratory of Aerospace Dynamics, Xi'an Satellite Control Center, Xi'an, 710043, China*
[3]*Beijing University of Posts and Telecommunications, Beijing, 100876, China*

wwfang@bjtu.edu.cn

Abstract

As the demand on online services and cloud computing has kept increasing in recent years, the power usage and cost associated with cloud data centers' operation have been uprising significantly. Most existing research focuses on reducing power consumption of data centers. However, the ultimate goal of cloud service operators is to reduce the total operating cost of data centers while guaranteeing the quality of service such as service delay to the end users. This paper exploits both the workload dispatching and the service provisioning to address the total electricity cost minimization problem. This problem is formulated as a hierarchical capacitated median model based on mixed integer linear programming (MILP) technique. Extensive evaluations based on real-life electricity price data for multiple data centers show the efficiency and efficacy of our approach.

Keywords: *Cloud Computing, Cost Minimization, Data Center, Smart Grid*

1. Introduction

The past decade witnessed tremendous growth of online cloud services. Generally, such services are expected to scale well, to guarantee performance, and to be highly available. To achieve these goals, these services are typically built atop geographically distributed infrastructures, *i.e.*, cloud data centers located in different regions around the world. However, along with the increasing demand of online cloud services and the deployment of cloud data centers, the operating cost on such data centers is increasing significantly. It has been reported that data center power usage in U.S. doubled between year 2000 and 2006 to about 61 billion kWh (1.5% of U.S. electricity consumption), and is projected to double again by 2011 to more than 100 billion kWh [1]. Therefore, the demand for power management of cloud data centers is becoming even more urgent and important.

Generally, most of the existing power management schemes focus on how to reduce the power demands of data center. However, the ultimate goal of data center operators is to reduce the total operating cost, which depends on not only the consumed power but also the electricity price. In the present multi-electricity-market environment, the price of electricity may exhibit temporal and spatial diversities [2], providing opportunities to data center operators to satisfy the needs of electricity while minimizing their energy

ISSN: 2005-4262 IJGDC

cost [3]. On the other hand, the operators need to guarantee that the cost optimization won't degrade the level of QoS (Quality-of-Service) provided to cloud service users.

In this paper, we propose a new approach to model and solve the total electricity cost minimization problem for distributed data centers. We formulate such a problem as a hierarchical version of the capacitated median problem [4], considering several levels of service demands and corresponding computing resources. The constraints in our model ensure that a given level of service demands can be satisfied by the computing resource of equal or higher level within the requirement on QoS (*i.e.*, service delay). The effectiveness and practicality of our approach is evaluated using the real-world electricity data, and prove that it is useful to reduce the total electricity cost for cloud data centers.

The rest of this paper is organized as follows. Section 2 discusses related work. Section 3 proposes the design and presents the problem modeling and formulation. Numerical results are provided in Section 4. Finally, we conclude the paper in Section 5.

2. Related Work

Along with the tremendous energy usage of cloud data centers, power management has become an important and active research area. While most of the existing works have been developed to reduce the total power consumptions of data centers [5, 6, 7, 8], recently some researchers studied the power management problem in the viewpoint of data center operators by focusing on energy cost other than pure power usage. Qureshi *et al.*, [2] introduced an intuitive approach of utilizing the location diversity of electricity spot price to intelligently direct requests to data centers with lower prices. Rao *et al.*, [3] proposed a joint load balancing and power control approach for data centers to exploit the time and location diversity of electricity prices. Xu *et al.*, [9] presented a general optimization framework for cost efficient data center selection that takes into account both electricity and bandwidth costs. Le *et al.*, [10] proposed VM-migration based load distribution policies that consider all electricity-related costs as well as transient cooling effects. These works, however, do not consider the user-side demands in their problem formulations.

Our work relies on the idea of a discrete hierarchical scheduling model for service demand processing, considering several levels of service demands and several types of computing resources. Similar idea has already been proposed in some recent works on the planning for common public facilities, *e.g.*, school network [4], waste bins [11] and healthcare facilities [12]. In this work, we extend the basic model in [4] by taking some characteristics of cloud services into consideration, such as service delay.

3. Formulation and Modeling

Current end users of cloud may have different types (or levels) of service demands, which must be separately assigned to computing resources capable of supporting them. We consider a nested (or successively inclusive) hierarchy of resources where a level-s (s = 1, 2, ..., ns) resource can serve demands of level 1, ..., s. Such a rule has already been adopted by real-world cloud operators through the virtualization technology [7]. For example, Amazon EC2 (Elastic Compute Cloud) provides multiple server instances with different characteristics (*i.e.*, CPU power, memory, disk, OS and software) and different prices [13]. The instances with more resources are capable to support more types of cloud applications and services. Typical cloud users like e-commerce websites will purchase different amounts and types of instances to satisfy their demands. Since these server instances may locate at the data centers in different regions, the operators can dynamically adjust the computing resource (*i.e.*, server instances) usage in each data center according to service demands and electricity prices so as to minimizing the total cost for serving cloud services.

3.1. Notations

We first summarize the notations which will be used throughout this paper in Table 1.

Table 1. Notations used in the paper

Notation	Meaning
I	the set of users
J	the set of data centers
S	the set of demand levels and of facility types
x_{isjk}	1 if the level-s service requests from user i is handled by the level-k server instances at data center j, 0 otherwise
y_{jk}	1 if type-k server instances are available at data center j, 0 otherwise
u_{is}	the level-s service requests from user i
e_{jk}	the unit electricity consumption for type-k server instances at data center j
A_{jk}	server power consumption specified parameter for type-k server instances at data center j
C_{jk}	server power consumption specified parameter for type-k server instances at data center j
f_{jk}	CPU frequency for type-k server instances at data center j
η_{jk}	the processing capability of a type-k server instance at data center j
$P_j(t)$	the electricity price for data center j at time t
m_{jk}	the number of type-k server instances at data center j
B_{jk}	the maximum capacities of type-k server instances at data center j
b_{jk}	the minimum capacities of type-k server instances at data center j
J_k	the set of data centers with existing type-k server instances
p_k	the maximum number of data centers that can start type-k server instances
q_k	the maximum number of data centers that can stop existing type-k server instances
D_{is}	the minimum service delay constraint for level-s service imposed by user i
$D_{jk}(t)$	the actual service delay for level-k server instances to handle the requests of level 1,…, k

3.2. Cost and Constraints Modeling

1) Total Electricity Cost Modeling: The objective of the power management problem is to minimize the total electricity cost T_{cost} for data centers in a multi-electricity-market environment. Therefore, we use the total electricity cost as the objective function to be minimized. The total electricity cost function can be given as:

$$T_{cost} = \sum_{i \in I} \sum_{j \in J} \sum_{s \in S} \frac{x_{isjk} u_{is}}{\eta_{jk}} * e_{jk} * P_j(t) \tag{1}$$

where $x_{isjk}\ u_{is}/\eta_{jk}$ represents the number of the level-k server instances at data center j that is active for handling the level-s service requests from user i. In order to simplify the problem, we assume that the server instances of the same level have the same utilization, the same frequency, and the same request arrival rate [3]. We approximate server power consumption e_{jk} by the following function:

$$e_{jk} = A_{jk} * f_{jk}^{p} + C_{jk} \tag{2}$$

where A_{jk} and B_{jk} are positive constants. In realistic systems, p varies between 2.5 and 3. The value of A_{jk}, B_{jk} and f_{jk} can be well obtained by curve fitting against empirical measurements when profiling the system offline [3].

2) Dispatching Constraint Modeling: Firstly, we need to ensure that all demands of all levels from all users are satisfied. Therefore, we have:

$$\sum_{j\in J} x_{isjk} = 1, \forall i \in I, s \in S, k \in S \tag{3}$$

Secondly, it must be imposed that a given level of demands can only be satisfied by available server instances of equal or higher level. Therefore, we have:

$$x_{isjk} \le y_{jk}, \forall i \in I, j \in J, s \in S, k \in S \tag{4}$$

$$x_{isjk} \le \frac{k}{s}, \forall i \in I, j \in J, s \in S, k \in S \tag{5}$$

3) Workload Constraint Modeling:

$$\sum_{i\in I} \sum_{s\in S|s\le k} x_{isjk} u_{is} \ge b_{jk} y_{jk}, \forall j \in J, k \in S \tag{6}$$

$$\sum_{i\in I} \sum_{s\in S|s\le k} x_{isjk} u_{is} \le B_{jk} y_{jk}, \forall j \in J, k \in S \tag{7}$$

Constraints (6) and (7) impose maximum and minimum limits on capacity, according to the type of server instance. With this formulation, capacity is shared by all demand levels.

4) Delay Constraint Modeling: To quantify the service level agreement (SLA), we use the steady-state results of the *M/M/n* queuing model to model each server instance in a data center [3]. In this model, the average service delay is given as:

$$D_{jk}(t) = \frac{1}{m_{jk}\eta_{jk} - \sum_{i\in I} \sum_{s\in S|s\le k} x_{isjk} u_{is}}, \forall j \in J, k \in S \tag{8}$$

To meet the SLA requirements of end users, there is a delay constraint D_{is} for the average service delay at data center j, and $D_{jk}(t) \le D_{is}$ ($s \le k$). Therefore, we have:

$$\sum_{i\in I} \sum_{s\in S|s\le k} x_{isjk} u_{is} \le m_{jk}\eta_{jk} - \frac{1}{D_{is}}, \forall j \in J, k \in S \tag{9}$$

5) Planning Constraint Modeling: Sometimes, the operator may plan for start or stop to provide some levels of server instances at some data centers. Generally, there are limitations for controlling the number of levels.

$$\sum_{j\in J\setminus J_k} y_{jk} \le p_k, \forall k \in S \tag{10}$$

$$\sum_{j\in J_k} y_{jk} \ge |J_k| - q_k, \forall k \in S \tag{11}$$

Constraints (10) and (11) limit the number of new level of server instances to start or existing levels to close.

3.3. Problem Formulation

The goal of the power management problem in this paper is to minimize the total electricity cost T_{cost} for cloud data centers in a multi-electricity-market environment. Hence, total electricity cost is the objective function. The decision variables are $x_{isjk} \in \{0, 1\}$ and $y_{jk} \in \{0, 1\}$. (Noted that y_{jk} can be known variables. If all y_{jk} are already known, this solution can provide us the final optimized results on workload dispatching for cost minimization. If some of the y_{jk} are unknown, this solution can be used for computing facility planning, *i.e.*, how many levels of server instances should be changed (start/stop) for minimizing the electricity cost. Examples are on these two cases in Section 4.) The constraints on dispatching, workload, delay and planning are discussed in Section 3.2. In conclusion, we have the following optimization problem:

$$
\begin{aligned}
& \min \sum_{i\in I}\sum_{j\in J}\sum_{s\in S}\frac{x_{isjk}u_{is}}{\eta_{jk}} * (A_{jk} * f_{jk}^{p} + C_{jk}) * P_j(t) \\
& \textit{subject to} \\
& \sum_{j\in J} x_{isjk} = 1, \forall i\in I, s,k\in S \\
& x_{isjk} \le y_{jk}, \forall i\in I, j\in J, s\in S, k\in S \\
& x_{isjk} \le \frac{k}{s}, \forall i\in I, j\in J, s\in S, k\in S \\
& \sum_{i\in I}\sum_{s\in S|s\le k} x_{isjk}u_{is} \ge b_{jk}y_{jk}, \forall j\in J, k\in S \\
& \sum_{i\in I}\sum_{s\in S|s\le k} x_{isjk}u_{is} \le B_{jk}y_{jk}, \forall j\in J, k\in S \\
& \sum_{i\in I}\sum_{s\in S|s\le k} x_{isjk}u_{is} \le m_{jk}\eta_{jk} - \frac{1}{D_{is}}, \forall j\in J, k\in S \\
& \sum_{j\in J\setminus J_k} y_{jk} \le p_k, \forall k\in S \\
& \sum_{j\in J_k} y_{jk} \ge |J_k| - q_k, \forall k\in S \\
& x_{isjk} \in \{0,1\}, y_{jk} \in \{0,1\}, \forall i\in I, j\in J, s\in S, k\in S
\end{aligned}
\tag{12}
$$

We implement the models in (12) based on the *linprog* solver in Matlab. In particular, *linprog* uses a simplex method, which has been proven to have a low complexity in practice.

4. Performance Evaluation

4.1. Experiment setup

The electricity system in the US is organized in ten cross-state regions such as CAISO, MISO and New England [3, 9]. To capture the location diversity of the cloud infrastructure and electricity market, we assume the data centers are deployed across the continental USA. The locations include Mountain View, CA, Fort Wayne, IN and Oxford, CT, which belongs to electricity region CAISO, MISO and New England respectively. From the raw data obtained from publicly available government agencies [14, 15], we calculate the average daily electricity price for all these locations in Feb 2012 and present the results in Figure 1 (a). We also calculate the hourly electricity price for all these locations on Feb 1, 2012 and present the results in Figure 1 (b). From Figure 1, we can know that the electricity prices for different regions may significantly different from others, and electricity prices are dynamically judged by the market with no obvious predictable pattern.

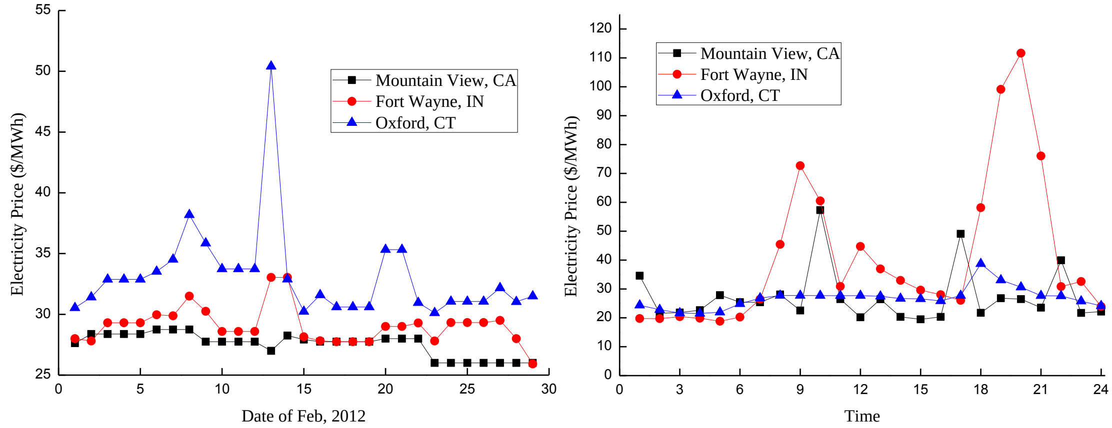

(a) (b)

Figure 1. Electricity price data for three major locations of cloud data centers: (a) daily data for Feb, 2012; (b) hourly data for Feb 1, 2012

To emulate data center server power consumption, we assume that workload comes from five enterprise user-groups to the three data centers. The setting on request workload, server instance are illustrated in Table 2 and 3, respectively. We assume that the server instances of the same level but at different data centers are homogeneous. The delay constraint D_{is} is 1ms [3]. Besides, $B_{jk} \leq m_{jk}\eta_{jk}$.

Table 2. Workload u_{is} Setting

i \ s	1	2	3
1	20000	10000	2000
2	15000	15000	2000
3	20000	15000	2000
4	15000	20000	2000
5	10000	15000	2000

Table 3. Server Setting for data center j

k	m_{jk}	η_{jk}	f_{jk}	A_{jk}	C_{jk}
1	30000	1.25	2.4	4.485	53
2	20000	1.75	3	2.370	70
3	10000	2	3.4	3.206	68

4.2. Experiment results

In the first set of experiments, we assume that all y_{jk} are known as in Table 4. Note that only the data center in Fort Wayne can provide server instances of level-3.

Table 4. y_{jk} setting

j \ k	1	2	3
1 (Mountain View)	1	1	0
2 (Fort Wayne)	1	1	1
3 (Oxford)	1	1	0

Our approach is compared to the total power minimization approach [3], in which the objective is as follows

$$T_{cost} = \sum_{i\in I}\sum_{j\in J}\sum_{s\in S}\frac{x_{isjk}u_{is}}{\eta_{jk}} * e_{jk} \qquad (13)$$

Figure 2 shows the comparison results of hourly total electricity cost for these two methods. It can be found that our approach is more capable for helping data center operators to save a lot in their electricity bill, especially when the differences between the prices of different regions are relatively higher (*e.g.*, hour 9 and hour 20). From Figure 3(a), we could notice that in our approach the optimal workload assignment changes with time and price for different locations, and the data center at the location with high electricity price (*e.g.*, data center at Fort Wayne, which even supports more levels than the other two locations) will generally receive much fewer service requests than others. We also check the workload dispatching for a single level (*i.e.*, level-1) workload. The results in Figure 3(b) exhibit a similar variation trend as those of total workload in Figure 3(a). In all, it is clear that there is a coherent relationship between the variations of the hourly workload assignment show in Figure 3 and the dynamics of the hourly electricity price shown in Figure 1(b). This indicates that the electricity prices at different locations affect the amount of optimal assigned workload.

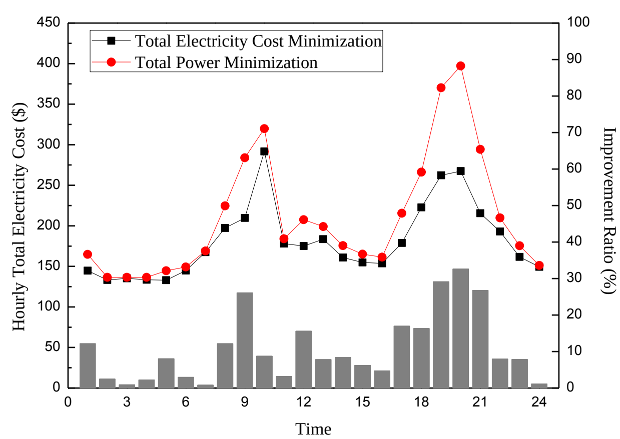

Figure 2. Comparison on hourly total electricity cost

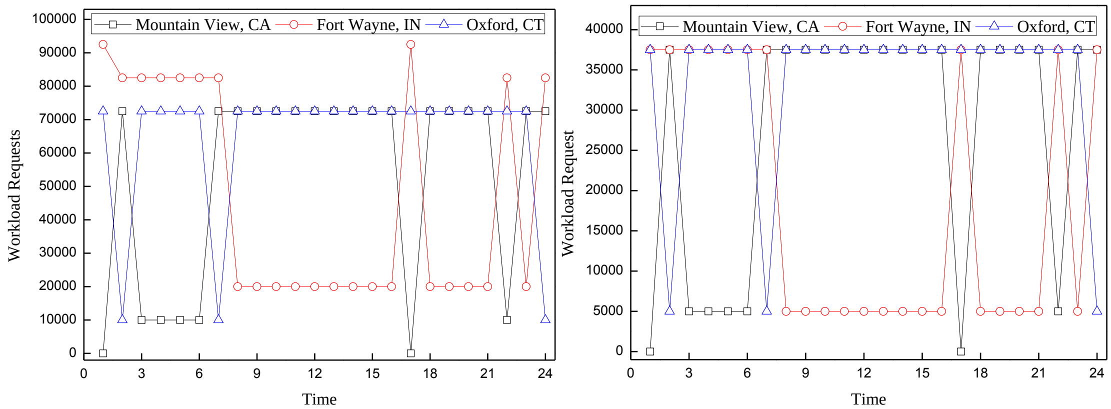

Figure 3. Assigned workloads for the three data centers: (a) total workload; (b) level-1 workload

In the second set of experiments, we set $p_k = 1$, *i.e.*, one more level-3 server instances can be started at data center 1 or 3. The results are shown in Table 5. It is obvious that our scheme (*i.e.*, optimized setting) can minimize the total electricity cost by optimizing the server provisioning. Although the setting is relatively simple, we believe that our approach is capable and suitable for planning optimization. In conclusion, our method can not only schedule workloads among different data centers, but also optimize the planning of cloud resources, to achieve energy-efficient computing.

Table 5. Comparison on server provisioning

	y_{13}	y_{23}	y_{33}	cost at hour 9	cost at hour 20	average daily cost
original setting	0	1	0	144.74	209.69	4348.03
optimized setting	1	1	0	165.02	188.58	4010.88
another setting	0	1	1	171.56	193.81	4054.97

5. Conclusion

This paper investigates an emergent and important problem of minimizing the total electricity cost for cloud data centers under a multi-electricity-market environment. We propose a scheme based on the hierarchical capacitated median model to minimize the total electricity cost while guaranteeing the QoS to end users. The experiments based on real-life electricity price data demonstrate the effectiveness of the proposed approach as well as total electricity cost reduction. It is worth noting that other online optimization and control approaches [16, 17] can be employed together with our method, which is left as a future work.

Acknowledgements

This work was supported in part by the National Science Foundation of China (No.61202430, 61303245 and 61272483), the Science and Technology Foundation of Beijing Jiaotong University (No. 2012RC040).

References

[1] A. Greenberg, J. Hamilton, D. A. Maltz and P. Patel, "The Cost of a Cloud: Research Problems in Data Center Networks", ACM SIGCOMM Computer Communication Review, vol. 39, no. 1, **(2009)**, pp. 68-73.

[2] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs, "Cutting the Electric Bill for Internet-scale Systems", Proceedings of 15th ACM SIGCOMM Conference, **(2009)**, pp. 123-134.

[3] L. Rao, X. Liu, M. D. Ilic and J. Liu, "Distributed Coordination of Internet Data Centers Under Multiregional Electricity Markets", Proceedings of the IEEE, vol. 100, no. 1, **(2011)**, pp. 269-282.

[4] J. C. Teixeira and A. P. Antunes, "A Hierarchical Location Model for Public Facility Planning", European Journal of Operational Research, vol. 185, no. 1, **(2008)**, pp. 92-104.

[5] L. T. Lee, K. Y. Liu, H. Y. Huang and C. Y. Tseng, "A Dynamic Resource Management with Energy Saving Mechanism for Supporting Cloud Computing", International Journal of Grid and Distributed Computing, vol. 6, no. 1, **(2013)**, pp. 67-76.

[6] J. Liu, F. Zhao, X. Liu and W. He, "Challenges Towards Elastic Power Management in Internet Data Centers", Proceedings of the 29th IEEE International Conference on Distributed Computing Systems Workshops, **(2009)**, pp. 65-72.

[7] W. W. Fang, X. M. Liang, S. X. Li, L. Chiaraviglio and N. X. Xiong, "VMPlanner: Optimizing Virtual Machine Placement and Traffic Flow Routing to Reduce Network Power Costs in Cloud Data Centers", Computer Networks, vol. 57, no. 1, **(2013)**, 179-196.

[8] Z. H. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. K. Wang, M. Marwah and C. Hyser, "Renewable and Cooling Aware Workload Management for Sustainable Data Centers", Proceedings of the 12th ACM SIGMETRICS Conference, **(2012)**, pp. 175-186.

[9] H. Xu and B. C. Li, "Cost Efficient Datacenter Selection for Cloud Services", Proceedings of the 1st IEEE Conference on Communications in China, **(2012)**, pp. 51-56.

[10] K. Le, J. R. Zhang, J. D. Meng, R. Bianchini, Y. Jaluria and T. D. Nguyen, "Reducing Electricity Cost Through Virtual Machine Placement in High Performance Computing Clouds", Proceedings of the 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, **(2011)**, pp. 1-12.

[11] A. S. Aremu and B. F. Sule, "A Case Study Evaluation of the Impacts of Optimized Waste Bin Locations in a Developing City", Civil Engineering and Environmental Systems, vol. 29, no. 2, **(2012)**, pp. 137-146.

[12] A. R. A. V. Fo and I. S. Mota, "Optimization Models in the Location of Healthcare Facilities: A Real Case in Brazil", vol. 4, no. 1, **(2012)**, pp. 37-50.

[13] G. Juve, E. Deelman, K. Vahi, G. Mehta, B. Berriman, B. P. Berman and P. Maechling, "Data Sharing Options for Scientific Workflows on Amazon EC2", Proceedings of the 2010 International Conference for High Performance Computing, Networking, Storage and Analysis, **(2010)**, pp. 1-9.

[14] U.S. Energy Information Administration (EIA), "Independent Statistics and Analysis", http://www.eia.doe.gov/.

[15] Federal Energy Regulatory Commission (FERC), "Electric Power Markets: National Overview", http://www.ferc.gov/.

[16] M. Salehpur and A. Shahbahrami, "Alleviating Dynamic Resource Allocation for Bag of Tasks Applications in Cloud Computing", International Journal of Grid and Distributed Computing, vol. 5, no. 3, **(2012)**, pp. 95-110.

[17] R. Z. Farahani and M. Hekmatfar, "Facility Location: Concepts, Models, Algorithms and Case Studies", Springer-Verlag, Berlin Heidelberg, **(2009)**.

Authors

Weiwei Fang

Weiwei Fang, BSc (2003), PhD (2010), is now an assistant professor at School of Computer and Information Technology, Beijing Jiaotong University. His current research interests include Distributed Computing, Cloud Data Centers and Wireless Sensor Networks. He has coauthored 1 book chapter and more than 30 papers, and served in the Technique Program Committee of more than 30 conferences and workshops.

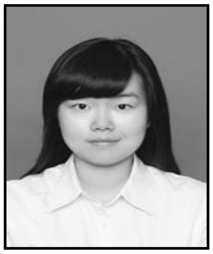

Quan Zhou

Quan Zhou received her BSc degree from School of Management, Tianjin Normal University, in 2011. She is now a MSc candidate at School of Economics and Management, Beijing Jiaotong University. From 2012, she also works as a research assistant at School of Computer and Information Technology, Beijing Jiaotong University. Her current research interests include Cloud Computing and Information Systems.

Yuan An

Yuan An, BSc(2003), PhD (2010), is now a researcher at State Key Laboratory of Aerospace Dynamics, Xi'an Satellite Control Center, China. His current research interests include Distributed Computing and Wireless Communication.

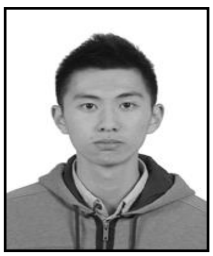

Yangchun Li

Yangchun Li is now a BSc candidate at School of Information and Communication Engineering, Beijing University of Posts and Telecommunications. From 2013, he also works as a research assistant at School of Computer and Information Technology, Beijing Jiaotong University. His current research interests include Distributed Computing and Cloud Data Centers.

Huijing Zhang

Huijing Zhang is now a BSc candidate at School of Computer Science, Beijing University of Posts and Telecommunications. From 2013, she also works as a research assistant at School of Computer and Information Technology, Beijing Jiaotong University. Her current research interests include Distributed Computing and Cloud Data Centers.

