

Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment

Lei Rao^{1,2}, Xue Liu², Le Xie³, Wenyu Liu¹

¹ Department of EI, WNLO, Huazhong University of Science and Technology, Wuhan, China
Email: raolei@smail.hust.edu.cn, liuwuy@mail.hust.edu.cn

² School of Computer Science, McGill University, Montreal, Canada
Email: {leirao, xueliu}@cs.mcgill.ca

³ Department of ECE, Texas A&M University, College Station, USA
Email: lx@andrew.cmu.edu

Abstract—The study of Cyber-Physical System (CPS) has been an active area of research. Internet Data Center (IDC) is an important emerging Cyber-Physical System. As the demand on Internet services drastically increases in recent years, the power used by IDCs has been skyrocketing. While most existing research focuses on reducing power consumptions of IDCs, the power management problem for minimizing the total electricity cost has been overlooked. This is an important problem faced by service providers, especially in the current multi-electricity market, where the price of electricity may exhibit time and location diversities. Further, for these service providers, guaranteeing quality of service (i.e. service level objectives-SLO) such as service delay guarantees to the end users is of paramount importance. This paper studies the problem of minimizing the total electricity cost under multiple electricity markets environment while guaranteeing quality of service geared to the location diversity and time diversity of electricity price. We model the problem as a constrained mixed-integer programming and propose an efficient solution method. Extensive evaluations based on real-life electricity price data for multiple IDC locations illustrate the efficiency and efficacy of our approach.

I. INTRODUCTION

With the advent of ubiquitous, low-cost computing and anytime connectivity, today's systems have evolved into a complex combination of the world's cyber-infrastructures as well as physical infrastructures. However, to date, most of these systems are comprised of components and subcomponents that are designed and developed independently. As a result, the study of systems in which the cyber and physical components are tightly integrated at all scales and levels, which are referred to as Cyber-Physical Systems (CPSs), is of paramount importance.

Internet Data Center (IDC) is one important emergent CPS. As the demand on Internet services drastically increases in recent years, the energy used by IDCs, directly related to the number of hosted servers and their workload, has been skyrocketing. A recent study estimates that the worldwide spending on enterprise and data center power and cooling to top \$30 billion in year 2008 and is likely to even surpass spending on new server hardware [23]. The trend continues

with the emergence of Cloud Computing, where data and computation hosting are outsourced to IDCs, for reliability, management, and cost benefits [1], [7], [16], [17]. Therefore, the need for power management of IDCs is becoming ever more urgent and important.

Recently both industry and academia have proposed various technologies and schemes to address different aspects. Most of these power management approaches have all been proposed as ways to reduce the power demands of data centers. However, the ultimate goal of IDC service providers such as Google, Microsoft etc. is to reduce the total operating cost: it does not depend on the power consumed by the IDCs, but also depend on the price (which could exhibit location and time diversity) of the power consumed.

In North America, service providers (like Google) build their IDCs in different regions because of various considerations including system reliability and electricity cost. In each region, a Regional Transmission Organization (RTO) manages the power grid. Due to the different power generation profiles, electricity prices are regional too. In those regions of the U.S. with wholesale electricity markets, prices may vary on an hourly or 15-min basis. In contrast, electricity prices remain the same for long-time period in other regions without wholesale markets. To better capture the present electricity price fluctuations for IDC locations, we must consider multi-region electricity markets.

For the IDC service providers, in addition to the concern of minimizing the total operating cost, how to guarantee the quality of service (i.e. service level objectives-SLO) such as service delay guarantee to the end users is equally important. This is because a degraded service such as delayed response may frustrate the client, make them switch to competitors, and result in revenue loss. In this paper, we systematically study the problem of how to minimize the total electricity operating for IDCs cost while guaranteeing quality of service geared to the location diversity and time diversity of electricity price in the multi-electricity-market environment.

To this end, we first model the total electricity cost, work-

load constraint and end-to-end delay constraint for IDCs. We then formulate the electricity minimization problem for IDCs in a multi-electricity-market environment as a constrained mixed-integer programming problem, in which the constraints capture the delay guarantee and workload requirements. For current IDCs in multi-electricity-market environments, both client workloads and electricity prices are time-dependent and may change fast. This renders the optimization problem changes accordingly. As a service provider, hence it is important to design efficient solution methods in order to cut electricity cost dynamically. We approximate the problem through a linear programming formulation and solve the linear programming formulation with an efficient fast polynomial-time method. Through extensive evaluations based on real-life electricity price data of certain main Google IDC locations, we show the efficacy of the designed algorithm as well as the total electricity cost reduction.

In summary, the main contribution of this paper is twofold:

- We investigate an important and novel problem of the total electricity cost for IDC service providers under a multi-electricity-market environment. Our formulation not only can help minimize the total electricity cost of IDCs, but also can guarantee the quality of service experience by end users.
- We approximate the mixed-integer optimization formulation with a linear programming formulation and solve it with Brenner's fast polynomial-time method [31]. Result shows that the total electricity cost is greatly reduced for IDC locations.

The rest of this paper is organized as follows. Section II discusses the related work. Section III gives the modeling and formulation of the total electricity cost minimization problem. Section IV solves the optimization problem by first converting the optimization problem to a minimum cost flow problem and then adopting a fast polynomial-time algorithm to get the solutions. Section V discusses a multi-electricity market environment for some of main Google IDC locations. Section VI first shows the effectiveness of the designed method based on two groups of real-life electricity price data and then shows that the proposed method greatly reduces the electricity cost. Finally, Section VII concludes the paper.

II. RELATED WORK

Cyber-Physical System (CPS) is an important and new active research area. The position paper published in the NSF workshop on CPSs [20] gives a good overview of different aspects of CPS research. Internet Data Center (IDC) is an important CPS. In this paper, we study the problem of minimizing the total electricity cost in a multi-electricity-market environment for IDCs. Due to the space limitation, in this section we only review the work most relevant to our work.

A. Power Management in IDCs

Due to the growing demand on Internet services in recent years, the energy used by IDCs grows rapidly. Liu et al. [12]

present a latest overview of challenges toward power management in IDCs. Efforts such as the Climate Savers Computing Initiative (www.climatesaverscomputing.org) intended to help lower worldwide computer energy consumption by promoting widespread adoption of high-efficiency power supplies and encouraging the use of power-savings features already present in users equipment. The adoption of low power chipset with chip multiprocessing [13], virtualization technologies [24], [21], power control theories [34], and other energy saving solutions [23], [14] have also contributed positively toward more energy-efficient servers.

While most of the new technologies have been developed aim to reduce the total power consumptions of IDCs[5], [8], [10], [18], [25], [28], [35], we are lacking a holistic approach to help better reduce the total electricity cost of IDC service providers. In this paper, we study the novel problem of minimizing IDC total electricity cost. In the new scheme, we coordinate the workload assigned from each front-end Web portal server to individual IDC locations in order to achieve the goal of minimizing the total electricity cost. Parolini et al. [15] also consider the load management for IDCs, however their goal is to minimize the integrated weighted sum of power consumption and computational performance, which is different from ours.

B. Multi-Region Electricity Markets

Research on modeling of electric power prices began about 15 years ago due to the deregulation of electric power industry [26]. The electric power grid in North America is operated on a regional basis by regional transmission organizations (RTOs). Due to different power generation profiles in different regions, electricity price exhibits location diversities. Electricity price varies on an hourly or 15-minute basis in regions of the U.S with wholesale electricity markets, whereas electricity prices remain the same for regions with regulated industry structure. While in [3] Qureshi et al. study the problem of reducing the electricity cost in a wholesale market environment, we investigate the problem in multi-region electricity markets to better capture the present electricity price situation in IDC locations.

C. Minimum Cost Flow: Problem and Solution

The minimum cost flow problem is to find the minimum cost way of sending a certain amount of flow through a flow network [30]. This problem has been well studied and many approaches were proposed to solve the minimum cost flow with linear programming method. In this paper, we convert the approximated linear programming to a minimum cost flow problem for deriving fast and efficient solution.

A fast strongly polynomial-time solution algorithm for general uncapacitated minimum cost flow problems is described by Orlin [11]. Other algorithms exploit the special structure of the minimum cost flow problem [9], [19], [29]. Our solution leverages a fast polynomial-time algorithm proposed by Brenner [31]. Brenner's algorithm has a worst case running time of $O(xy^2(\log x + y \log y))$, where x is the number of sources and

y is the number of sinks in the minimum cost flow problem, respectively.

III. FORMULATION AND MODELING

In this section, we give the detailed modeling of the minimization of total electricity cost for distributed Internet Data Centers (IDCs) under multi-electricity-market environments. To this end, we first give the modeling of the total electricity cost, the workload constraint and client end-to-end delay constraint respectively. Then we formulate the total electricity cost minimization problem as a constrained mixed-integer programming problem.

A. Notations

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

TABLE I
NOTATIONS

i	IDC location i
j	Front-end Web portal server j
t	Time t
N	Total number of IDC locations
C	Total number of front-end Web portal servers
L_j	Total workload at front-end Web portal server j
$Pr_i(t)$	Spot price of electricity at i
PO_i	Power consumption for one server at i
λ_{ji}	The request rate from j to i
D_i	The end-to-end delay constraint at i
m_i	The number of turned on servers at i

B. Total Electricity Cost Modeling for IDC

The electric power grid in North America is operated on a regional basis by regional transmission organizations (RTOs). Regions in which most large-scale data centers are built are under competitive electricity market structure. The spot price of electricity is determined by the clearing of supply and demand functions, while transmission limits are observed. A bottom-up bid-based stochastic price model [26] is adopted here to characterize the hourly spot price dynamics $Pr_i(t)$.

In this paper, we focus on the discussion of the impact of price diversities on total electricity cost. In order to simplify the problem, we make two important assumptions. First, we assume that the power consumption profile is constant for each server at every location. This is usually true when each data center use homogenous servers and configurations.

The real power consumption of a single server at each data center location i can be estimated based on the current server frequency, f_i , and CPU utilization, U_i [27]. So the total power consumption PO_i at each data center i can be obtained by multiplying the total number of servers at location i with the power consumption value for each server.

Second, we also assume each server at the same location receives the same traffic rate in steady-state. This is usually achieved through load balancing at each location. Thus, the




Fig. 1. M/M/n Queueing Model

total electricity cost ¹ for N data center locations is given as:

$$T_{total} = \sum_{i=1}^N m_i Pr_i(t) PO_i.$$

C. Workload Constraint Modeling

In this subsection, we model the workload constraints in order to capture the requests allocations among all the data center locations for each front-end Web portal. A client request is first handled by a front-end Web portal server, then it is forwarded to one server at a specific location of the IDC to be processed. We denote the requests demand at each front-end Web portal server j as L_j ($j = 1, \dots, C$). The request arrival rate from front-end Web portal server j to location i is λ_{ji} . Therefore we have:

$$\sum_{i=1}^N \lambda_{ji} = L_j, \forall j = 1, \dots, C. \quad (1)$$

At each location of an IDC, there are usually hundreds or even thousands servers so as to be able to afford large number of requests. However, there is a limitation M_i on the number of servers m_i at each location i of IDCs. Therefore, we have

$$m_i \leq M_i, \forall i = 1, \dots, N. \quad (2)$$

D. Delay Constraint Modeling

In this subsection, we use the M/M/n queueing model shown in Figure 1 to model each server in the data center. In the M/M/n queueing model, the average delay D , is expressed as $\frac{1}{n\mu - \lambda} P_Q$, given the number of servers n , the service rate μ , the arrival rate λ and the probability of customers waiting in queue P_Q . In a data center, without loss of generality, we can assume that the servers are always busy, i.e. there are always requests waiting in queue. Hence, we have P_Q equals 1.

¹Note that energy is the time integral of power consumption, and electricity cost is the time integral of the product of electricity price and power consumption. In this paper we mainly discuss the optimization at a specific time, hence we omit the time integral in the formulation. So we use power consumption directly instead of energy consumption.

At location i with m_i servers, when each sever has the service rate μ_i and the total arrival rate is $\sum_{j=1}^C \lambda_{ji}$, the average delay D_i^{CPU} is given as

$$D_i^{CPU} = \frac{1}{m_i \mu_i - \sum_{j=1}^C \lambda_{ji}}.$$

To meet the requirements of end users, there is a delay constraint D_i for the average delay at each location i :

$$\frac{1}{m_i \mu_i - \sum_{j=1}^C \lambda_{ji}} \leq D_i, \forall i = 1, \dots, N. \quad (3)$$

E. Formulation of Total Electricity Cost Minimization Problem

The goal of the power management problem in this paper is to minimize the total electricity cost T_{total} for IDCs in a multi-electricity-market environment. Hence, total electricity cost is the objective function. The decision variables are the number of machines m_i and the assigned work rate λ_{ji} from each front-end Web portal server j to each location i . The constraints are the workload constraints discussed in Subsection III-C and the delay constraints discussed in Subsection III-D. In conclusion, we have the following optimization problem named **Problem One**:

$$\min_{m_i, \lambda_{ji}} \sum_{i=1}^N m_i Pr_i(t) P o_i, \quad (4)$$

subject to

$$\frac{1}{m_i \mu_i - \sum_{j=1}^C \lambda_{ji}} \leq D_i, \forall i = 1, \dots, N, \quad (5)$$

$$\sum_{i=1}^N \lambda_{ji} = L_j, \forall j = 1, \dots, C, \quad (6)$$

$$m_i \leq M_i, \forall i = 1, \dots, N, \quad (7)$$

$$m_i \in \mathcal{N}, \forall i = 1, \dots, N. \quad (8)$$

IV. SOLUTION METHOD DESIGN

In this section, we present a solution method for **Problem One** discussed in section III. In order to solve the problem, we approximate the problem by a linear programming formulation. We then convert the problem to a minimum cost flow problem and solve it with a fast polynomial-time algorithm.

A. Mixed Integer Linear Programming Formulation

We first transform (5) in **Problem One** to $m_i \mu_i - \sum_{j=1}^C \lambda_{ji} \geq \frac{1}{D_i}, \forall i = 1, \dots, N$, and then we have $m_i \geq \frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i}, \forall i = 1, \dots, N$. Hence **Problem One** in Section III can be rewritten as

$$\min_{m_i, \lambda_{ji}} \sum_{i=1}^N m_i Pr_i(t) P o_i, \quad (9)$$

subject to

$$m_i \geq \frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i}, \forall i = 1, \dots, N, \quad (10)$$

$$\sum_{i=1}^N \lambda_{ji} = L_j, \forall j = 1, \dots, C, \quad (11)$$

$$m_i \leq M_i, \forall i = 1, \dots, N, \quad (12)$$

$$m_i \in \mathcal{N}, \forall i = 1, \dots, N. \quad (13)$$

This is a mixed integer linear programming problem that we will solve. Since m_i is an integer, from (11) we have

$$m_i = \left\lceil \frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i} \right\rceil. \text{ Therefore, we have}$$

$$\frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i} \leq m_i \leq \frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i} + 1. \quad (14)$$

In order to solve the problem, we leverage the rounding technique in [2] to approximate the mixed integer linear programming formulation with a safe bound. Since our paper does not focus on approximation of mixed integer programming problem, we do not discuss the technique details here. For detailed discussions, refer to [2]. According to [2], we transform the optimization problem as following, named **Problem Two**:

$$\min_{\lambda_{ji}} \sum_{i=1}^N Pr_i(t) P o_i \left\{ \frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i} + 1 \right\}, \quad (15)$$

subject to

$$\sum_{i=1}^N \lambda_{ji} = L_j, \forall j = 1, \dots, C, \quad (16)$$

$$\sum_{j=1}^C \lambda_{ji} \leq \mu_i (M_i - \frac{1}{\mu_i D_i}), \forall i = 1, \dots, N, \quad (17)$$

$$m_i = \left\lceil \frac{\sum_{j=1}^C \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i} \right\rceil, \forall i = 1, \dots, N. \quad (18)$$

We first solve (15), (16) and (17) to decide the workload λ_{ji} to assign from each front-end Web portal server j to each location i , and then calculate (18) to decide the number of servers m_i to run for each location i .

B. Polynomial-Time Solution for Approximated Total Electricity Cost Minimization Problem

For current IDCs in multi-electricity-market environments, both client workloads and electricity prices are time-dependent and may change fast, which renders the optimization problem changes accordingly. It is important for service providers to design efficient solution methods in order to cut electricity cost dynamically. In this subsection, we show that **Problem**

Two can be converted to a minimum cost flow problem. We leverage a fast polynomial-time algorithm to solve this problem.

To illustrate the polynomial-time solution for **Problem Two**, without loss of generality, we consider a simple case that $N = 3$ and $C = 5$. The optimization problem in (15), (16) and (17) can be written as:

$$\min_{\lambda_{ji}} \sum_{i=1}^3 Pr_i(t) Po_i \left\{ \frac{\sum_{j=1}^5 \lambda_{ji}}{\mu_i} + \frac{1}{\mu_i D_i} + 1 \right\}, \quad (19)$$

subject to

$$\lambda_{11} + \lambda_{12} + \lambda_{13} = L_1, \quad (20)$$

$$\lambda_{21} + \lambda_{22} + \lambda_{23} = L_2, \quad (21)$$

$$\lambda_{31} + \lambda_{32} + \lambda_{33} = L_3, \quad (22)$$

$$\lambda_{41} + \lambda_{42} + \lambda_{43} = L_4, \quad (23)$$

$$\lambda_{51} + \lambda_{52} + \lambda_{53} = L_5, \quad (24)$$

$$\lambda_{11} + \lambda_{21} + \lambda_{31} + \lambda_{41} + \lambda_{51} \leq \mu_1 \left(M_1 - \frac{1}{\mu_1 D_1} \right), \quad (25)$$

$$\lambda_{12} + \lambda_{22} + \lambda_{32} + \lambda_{42} + \lambda_{52} \leq \mu_2 \left(M_2 - \frac{1}{\mu_2 D_2} \right), \quad (26)$$

$$\lambda_{13} + \lambda_{23} + \lambda_{33} + \lambda_{43} + \lambda_{53} \leq \mu_3 \left(M_3 - \frac{1}{\mu_3 D_3} \right), \quad (27)$$

where $\sum_{j=1}^C L_j = L$ is the total workload from five Web servers,

and $\sum_{i=1}^N \mu_i \left(M_i - \frac{1}{\mu_i D_i} \right)$ is the maximum workload that all locations can afford.

Constraints (20), (21), (22), (23) and (24) are known as requirement constraints. These constraints ensure that each front-end Web portal server obtains its workload requirement. Constraints (25), (26) and (27) are known as availability constraints. There is one such constraint for each of the three locations. These constraints ensure that the total workload of a location does not exceed its capacity. If the sum total of the requirements exactly matched the sum total of the availabilities then constraints, constraints (25), (26) and (27) can be treated as '=' instead of ' \leq '. However, this is not so for most of the case.

The above problem can be looked at graphically as illustrated in Figure 2. In the network of Figure 2, we have the locations T_1 , T_2 and T_3 and the six front-end Web portal servers S_1 , S_2 , S_3 , S_4 , S_5 and S_6 (including the dummy server). T_i and S_j provide the nodes of the network to which we have attached the (positive) capacities or (negative) requests. The possible supply patterns T_i to S_j provide the arcs of the network to which we have attached the unit supply costs. Our problem can now be regarded as one where we wish to obtain the minimum cost flow through the network. The T_i nodes are 'sources' for the flow entering the system and the S_j nodes are 'sinks' where flow leaves the system.

Minimum cost flow problem has been well studied and there are several methods that could solve this problem in polynomial time. We utilize a fast polynomial-time algorithm proposed by Brenner [31] to solve our problem. On

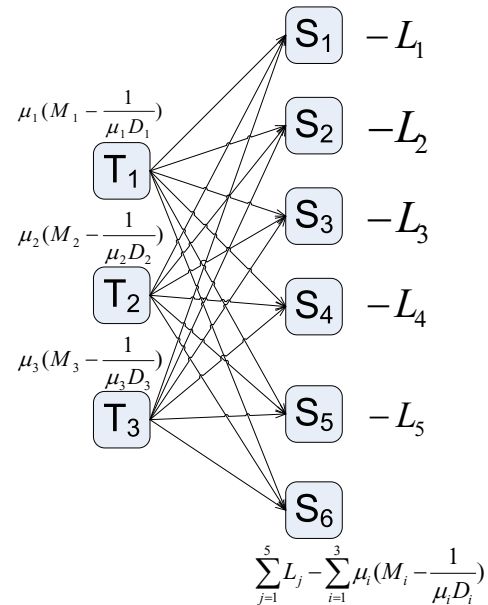


Fig. 2. Illustration of the Simple Case in Subsection IV-B

instances with N sources and $C+1$ sinks, Brenner's algorithm has a worst case running time of $O(N(C+1)^2(\log N + (C+1)\log(C+1)))$. This algorithm closes a gap between algorithms with running time linear in N but exponential in C and a polynomial-time algorithm with running time $O(N(C+1)^2 \log^2 N)$.

The basic idea of solving the total electricity cost minimization problem with the fast polynomial-time algorithm is: according to the electricity prices at different data center locations, Brenner's algorithm is calculated to decide the amount of workload assigned from different Web servers to each location to minimize the total electricity cost for IDCs; according to the assigned workload, each location decides how many servers to run to meet the end-to-end delay constraint.

V. ELECTRICITY PRICE AT CERTAIN LOCATIONS FOR GOOGLE INTERNET DATA CENTERS

In this section, we present the electricity price data at certain locations for Google Internet Data Centers (IDCs), which contains the following locations: Mountain View, California, Houston, Texas, and Atlanta, GA. The data reported here was obtained from publicly available government agencies [32] [33].

The electricity system in the US is organized in cross-state regions such as New England, PJM (primarily covering Pennsylvania, New Jersey, and Maryland), and ERCOT (Texas). Both regulated utility structure and deregulated wholesale market structure exist in different regions. The two locations of Mountain View and Houston are in the deregulated electricity market regions, and the other one (i.e. Atlanta) is located in the regulated utility region (with fixed electricity rate). For the two deregulated market regions, California has an hour-ahead forward market, whereas Texas has a 15-min




Fig. 3. Daily Electricity Price Data for Three Major Locations of Google IDCs on May, 2009




Fig. 4. Hourly Electricity Price Data for Three Major Locations of Google IDCs on May 2nd, 2009





Fig. 5. Comparing Price Variation for Hourly Market and Real-time Market in Houston on May 2nd, 2009

ahead forward market. From the raw data we calculate the average daily electricity price for all these locations on May, 2009 and present the results in Figure 3 (www.caiso.com, www.ercot.com). We also calculate the hourly electricity price for all these locations on May 2nd, 2009 and present the results in Figure 4. Figure 5 illustrates the different dynamics for Houston. Compared to the hourly market, the 15 mins real-time market is more volatile, with more high-frequency variation.

VI. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed method based on real-life electricity price data for for some known Google IDCs in the U.S. as discussed in the previous section. We first show the total electricity cost reduction by two groups of data at rush hour. We then present the performance evaluation based on the data of May 2nd, 2009 and the optimal workload assignments from Web servers to these IDC locations. At the end of this section, we further discuss on some insights about IDC resource planning for service providers based on our model in terms of shadow price.

A. Electricity Cost Reduction

To emulate Google IDC server power consumption, we assume that the total workload from five front-end Web portal servers to three back-end IDCs in the three respective locations (i.e. $C = 5$, $N = 3$) is 10^5 requests per second; the delay constraint is 1 ms; each server is operating at 120 Watts;

the processing speed for each server at three locations are 2.0 requests per second, 1.25 requests per second and 1.75 requests per second, respectively. The other parameters are shown in Table II, Table III and Table IV, where T_{total} represents the total electricity cost with proposed method in this paper; T_{ave} represents the average electricity cost with average workload assignment method; R_{reduce} indicates the reduced electricity cost ratio of optimal workload assignment with respect to the average workload assignment. In this paper, the unit for electricity price Pr_i at location i is $\$/MWh$, and the unit for electricity cost T_{total} and T_{ave} is \$.

As mentioned in Section V, we select three known Google IDC locations to emulate a dynamic multi-electricity-market environment. These locations, form a subset of the whole Google IDC locations, in the U.S.. Without loss of generality, we select two groups of data at GMT-2:00 Time 9:00 in the morning and 16:00 in the afternoon, which are all working hours at three different locations. We solve the optimization problem with $Pr = [42.92566, 20.27, 55.30] \$/MWh$ and $Pr = [77.57629, 29.48, 55.30] \$/MWh$, respectively. The optimal hourly total electricity cost for the first group of data is 219.2845 \$ compared with 285.4376 \$ for the average workload assignment case, while the optimal hourly total electricity cost for the second group of data is 319.3067 \$ compared with 387.1758 \$ for the average workload assignment case. With the designed algorithm for power management, we can reduce the hourly total electricity cost by 30.15% and 17.53%, respectively. This result demonstrates the importance of optimal workload assignment for data centers.

TABLE II
 PARAMETERS OF THREE GOOGLE IDC LOCATIONS

i	μ_i	Po_i	M_i	D_i
1	2.00	120	30000	0.001
2	1.25	120	60000	0.001
3	1.75	120	25000	0.001

TABLE III
 PARAMETERS OF FIVE GOOGLE FRONT-END WEB PORTAL SERVERS

j	1	2	3	4	5
L_j	30000	15000	15000	20000	20000

B. Performance Evaluation Based on Real-life Electricity Price Data for Google IDCs

Service providers like Google requires large computational resources in order to provide reasonable service. Google has a large number of data centers scattered around the world. At least 12 significant Google IDC installations are located in the United States. In this paper, we evaluate the performance of the proposed method with the electricity prices for three of these locations: Mountain View, California, Houston, Texas, and Atlanta, GA [32] [33]; and five front-end Web portal servers.




Fig. 6. Hourly Running Number of Servers for Three Major Locations of Google IDCs on May 2nd, 2009




Fig. 7. Comparison of Hourly Total Electricity Cost for Optimal Workload Assignment and Average Workload Assignment

Figure 6 shows the number of servers for each location to run to meet the workload requirement and the quality of

TABLE IV
 COMPARISON OF OPTIMAL WORKLOAD ASSIGNMENT AND AVERAGE WORKLOAD ASSIGNMENT FOR TWO GROUPS OF REAL-LIFE DATA

i	t	Pr_i	m_i	T_{total}	T_{ave}	R_{reduce}
1	9:00	42.92566	13501			
2	9:00	20.27	60000	219.2845	285.4376	30.15%
3	9:00	55.30	572			
i	t	Pr_i	m_i	T_{total}	T_{ave}	R_{reduce}
1	16:00	77.57629	501			
2	16:00	29.48	60000	319.3067	387.1758	17.53%
3	16:00	55.30	15429			

service constraints for all the front-end Web portal servers. First we observe it is obvious that the number of servers at each location varies a lot with the electricity price variations, which indicates the optimal workload assignment changes with time for different locations. To this end, it is necessary to calculate the number of servers for different time at different locations. However, as the workloads at different locations may vary, there might be increasing bandwidth cost. In this paper, though we focus on the simple setup of the problem without bandwidth constraints, it is worth noting that our formulation can be easily extended to accommodate the bandwidth constraints. On the other hand, a service provider as large as Google can negotiate contracts with carriers on a nationwide basis.

Figure 7 shows the comparison of electricity cost for optimal workload assignment and average workload assignment. As we can see, the electricity cost is greatly reduced in every single hour, which indicates that the proposed method would help service providers save a lot in their electricity bill.

C. Optimal Workload Assignment from front-end Web portal servers to Internet Data Centers

Figure 8 illustrates the optimal workload assignment for all of the five front-end Web portal servers. As shown in Table III, since front-end Web portal server 2 and Web server 3 have the same workload, we can see in Figure 8 that the optimal workload assignments are the same. This indicates the electricity cost for two front-end Web portal servers are the same, independent of their transmission costs. We can observe similar results for front-end Web portal server 4 and front-end Web portal server 5.

D. Discussion on the IDC Resource Planning in Terms of Shadow Prices

From the evaluation results in Subsection VI-B, we can observe that workloads assignment changes with the prices' variation in the electricity markets in order to achieve the goal of minimizing the total electricity cost. The other important observation is that workloads at each front-end Web portal server may exhibit time diversities. In this case, service providers (i.e. decision makers) may face the problem of how to do electricity cost budget planning in such a dynamic market environment. Meanwhile, to satisfy the increasing Internet services, service providers may expect to expand the scale of IDCs. Therefore, to scale which IDC location is also

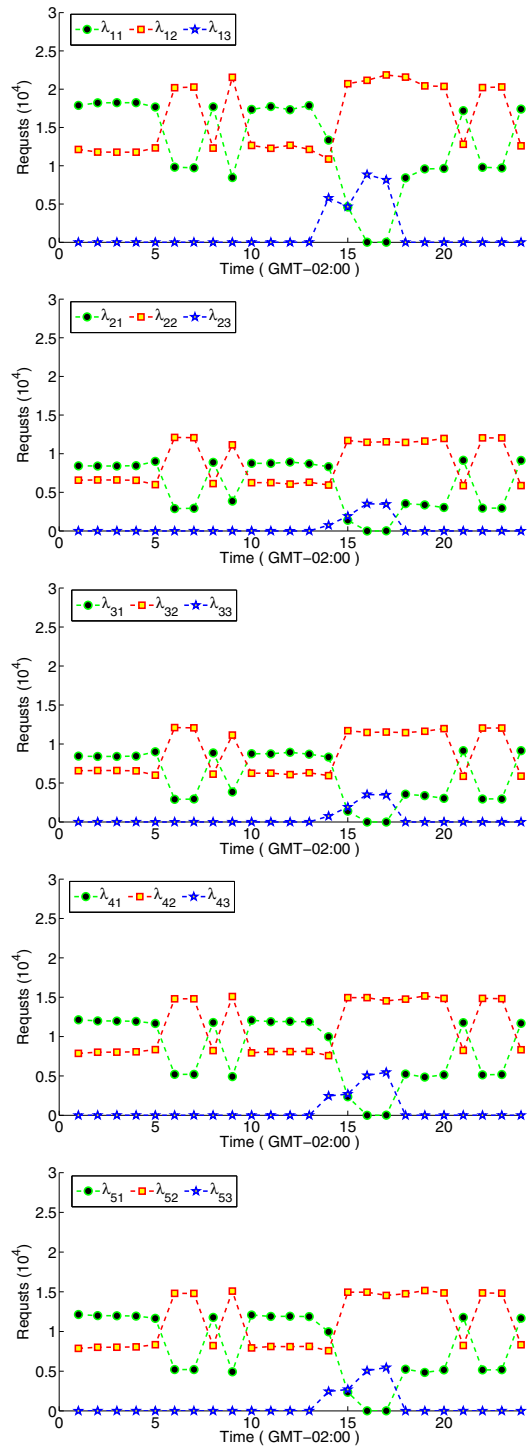


Fig. 8. Hourly Assigned Workloads of front-end Web portal servers to Three Google IDC Locations

an important problem. In this section, we make a detailed discussion on these practical problems in terms of shadow prices.

Every linear programming problem, referred to as a primal problem, can be converted into a dual problem [22]. Assume

the primal problem has the following form:

$$\min o^T x, \text{ subject to } Ax = u, x \geq 0.$$

The corresponding dual problem is

$$\max u^T y, \text{ subject to } A^T y \geq o, y \geq 0,$$

where y is used instead of x as the decision variable vector. The strong duality theorem [22] states that if the primal has an optimal solution, x^* , then the dual also has an optimal solution, y^* , such that $o^T x^* = b^T y^*$.

Then for **Problem Two**, the optimal dual values for constraint (16) are as follows:

$$\forall j = 1, \dots, C : q_j = \frac{Pr_i(t)Po_i(t)\lambda_{ji}^*}{\mu_i L_j}, \quad (28)$$

and the optimal dual values for constraint (17) are:

$$\forall i = 1, \dots, N : p_i = \frac{Pr_i(t)Po_i(t)\lambda_{ji}^*}{\mu_i^2 (M_i - \frac{1}{\mu_i D_i})}. \quad (29)$$

In linear programming, the optimal dual values are also called shadow prices [4], [22]. The shadow price captures the change in the objective function arising from an change in the amount of constrain constant. The value of the shadow price can provide decision makers powerful insight into our problems. For instance if there is a constraint that limits the workload arrived to the front-end Web portal server per second, the shadow price will indicate the total electricity cost change for an additional unit of the workload. In this paper, according to (28), for different front-end Web portal server j , the shadow price q_j tells the average electricity cost of the request processing at different front-end Web portal servers. Service providers hence can decide the electricity cost budget as the requests demand increases or decreases at each front-end Web portal server j . Similarly, according to (29), for different location i , the shadow price p_i can tell the average electricity cost of the upper limitation number of servers at each location. Service providers will prefer to increase the number of servers at location i that is with lower shadow price in this case.

VII. CONCLUSION AND FUTURE WORK

Power expenses are becoming an increasingly important fraction of Internet Data Center (IDC) operating cost. In this paper, we investigate an important and emergent problem of minimizing the total electricity cost for IDC under a multi-electricity-market environment. In order to provide a holistic solution, we model the total electricity cost, workload constraint and end-to-end delay constraint respectively and formulate the total electricity cost minimization problem as a constrained mixed-integer programming problem.

In order to obtain efficient solutions, we approximate the optimization problem through a linear programming formulation. We then convert the linear programming formulation to a minimum cost flow problem. We leverage Brenner's fast polynomial-time algorithm [31] to solve the minimum cost flow problem. Through extensive evaluations based on real-life Google IDC data and associated electricity price, we show

the efficacy of the proposed method as well as total electricity cost reduction.

As a next step work, we plan to extend the proposed method to capture bandwidth cost constraints.

ACKNOWLEDGMENT

This work was supported in part by Chinese National 863 project (No.2007AA01Z223), National Science & Technology Pillar Program of China (No. 2006BAH02A24), NSERC Discovery Grant 341823-07, NSERC Strategic Grant STPGP 364910-08, and FQRNT Grant 2010-NC- 131844.

REFERENCES

- [1] Amazon Web Services LLC, Amazon elastic computing cloud (amazon ec2), 2008. [Online]. Available: <http://aws.amazon.com/ec2/>
- [2] A. Neumaier and O. Shcherbina, Safe bounds in linear and mixed-integer linear programming, *Mathematical Programming*, vol. 99, Mar. 2004, pp. 283-296.
- [3] A. Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, and B. Maggs, Cutting the Electric Bill for Internet-Scale Systems, *ACM SIGCOMM*, 2009.
- [4] C. Courcoubetis and R. Weber, Pricing communication networks: economics, technology, and modelling, John Wiley & Sons Inc, 2003.
- [5] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, F. Zhao, "Energy-Aware Server Provisioning and Load Dispatching for Connection-Intensive Internet Services", in Proceedings of the 5th USENIX Symposium on Networked Systems Design & Implementation (NSDI'08), San Francisco, CA, April 2008.
- [6] F. D. Galiana, "An application of system identification and state prediction to electric load modeling and forecasting," *Electric Power Systems Laboratory*, MIT, Cambridge, MA, 1971.
- [7] Google, Google apps, 2008. [Online]. Available: <http://www.google.com/apps/intl/en/business/index.html>
- [8] Hewlett-Packard Development Company, HP Enterprise Power Calculators. <http://h30099.www3.hp.com/configurator/powercalcs.asp>, 2006.
- [9] H.P. Kleinschmidt, H. Schannath, A strongly polynomial algorithm for the transportation problem, *Mathematical Programming* 68 (1995) 1C13.
- [10] J. Heo, D. Henriksson, X. Liu, T. Abdelzaher, "Integrating Adaptive Components: An Emerging Challenge in Performance-Adaptive Systems and a Server Farm Case-Study," in Proceedings of the 28th IEEE Real-Time Systems Symposium (RTSS'07), Tucson, Arizona, 2007
- [11] J.B. Orlin, A faster strongly polynomial minimum cost flow algorithm, *Operations Research* 41 (1993) 338C350.
- [12] J. Liu, F. Zhao, X. Liu, W. He, "Challenges Towards Elastic Power Management in Internet Data Centers," *icdcs*, pp.65-72, 2009 29th IEEE International Conference on Distributed Computing Systems Workshops, 2009
- [13] L. A. Barroso and U. Holzle, The case for energyproportional computing. *Computer*, vol. 40, no. 12, pp. 33C 37, 2007.
- [14] L. Liu, H. Wang, X. Liu, X. Jin, W.B. He, Q.B. Wang, and Y. Chen, GreenCloud: a new architecture for green data center, Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session, Barcelona, Spain: ACM, 2009, pp. 29-38.
- [15] L. Parolini, B. Sinopoli, and B. Krogh, Reducing Data Center Energy Consumption via Coordinated Cooling and Load Management, Hot-Power '08: Workshop on Power Aware Computing and Systems, ACM, 2008.
- [16] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, Above the clouds: A Berkeley view of cloud computing, EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28, Feb 2009. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html>
- [17] Microsoft, Azure service platform, 2008. [Online]. Available: <http://www.microsoft.com/azure/default.aspx>
- [18] M. Elnozahy, M. Kistler, R. Rajamony, "Energy-Efficient Server Clusters," *Power-Aware Computer Systems*, 2003, pp. 179-197.
- [19] N. Megiddo, A Tamir, Linear time algorithms for some separable quadratic programming problems, *Operation Research Letters* 13 (1993) 203C211.
- [20] National Science Foundation. Cyber-physical systems. Technical report, NSF Workshop on Cyber-Physical Systems, 2006. <http://varma.ece.cmu.edu/cps/>.
- [21] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, A. Warfield, "Xen and the art of virtualization", *ACM SIGOPS Operating Systems Review*, Volume 37 , Issue 5,December 2003.
- [22] R. Dorfman, P.A. Samuelson, and R.M. Solow, *Linear programming and economic analysis*, Courier Dover Publications, 1987.
- [23] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu, No power struggles: coordinated multi-level power management for the data center, in *ASPLOS XIII: Proceedings of the 13th international conference on Architectural support for programming languages and operating systems*. New York, NY, USA: ACM, 2008, pp. 48C59.
- [24] R. Nathuji, K. Schwan, "VirtualPower: coordinated power management in virtualized enterprise systems", *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, Stevenson, Washington, USA, 2007.
- [25] S. Nedevschi, L. Popa1, G. Iannaccone, S. Ratnasamy, D. Wetherall, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation", In *Proceedings of the 5th USENIX Symposium on Networked Systems Design & Implementations (NSDI'08)*, San Francisco, CA, April 2008.
- [26] P. Skantze, M.D. Ilić and J. Chapman, "Stochastic modeling of electric power prices in a multi-market environment," *IEEE Power Engineering Society Winter Meeting*, 2000.
- [27] T. Horvath and K. Skadron, Multi-mode energy management for multi-tier server clusters, *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*, Toronto, Ontario, Canada: ACM, 2008, pp. 270-279.
- [28] T. Horvath, T. Abdelzaher, K. Skadron, and X. Liu, Dynamic voltage scaling in multitier Web servers with end-to-end delay control, *Computers*, *IEEE Transactions on*, vol. 56, no. 4, pp. 444C458, April 2007.
- [29] T. Tokuyama, J Nakano, Geometric algorithms for a minimum cost assignment problem, in: *Proceedings of the Seventh Symposium on Computational Geometry*, 1991. pp. 262C271.
- [30] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein (2001) [1990]. "29". *Introduction to Algorithms* (2nd edition ed.). MIT Press and McGraw-Hill. pp. 787C788. ISBN 0-262-03293-7.
- [31] U. Brenner, A faster polynomial algorithm for the unbalanced Hitchcock transportation problem, *Operations Research Letters*, vol. 36, Jul. 2008, pp. 408-413.
- [32] United States Energy Information Administration, Department of Energy, available online at www.eia.doe.gov.
- [33] United States Federal Energy Regulatory Commission, available online at www.ferc.gov.
- [34] X. Liu, J. Heo, L. Sha, X. Zhu, "Adaptive Control of Multi-Tiered Web Application Using Queueing Predictor," in *Proceedings of the 10th IEEE/IFIP Network Operations and Management Symposium (NOMS 2006)*, Vancouver, Canada, 2006.
- [35] X. Fan, W. Weber, L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *Proceedings of the 34th annual international symposium on Computer architecture*, San Diego, California, USA, 2007.