An Optimization Framework for Data Centers to Minimize Electric Bill under Day-Ahead Dynamic Energy Prices While Providing Regulation Services

Majid Ghasemi-Gol, Yanzhi Wang, and Massoud Pedram
University of Southern California
Ming Hsieh Department of Electrical Engineering
Los Angeles, United States
{*ghasemig, yanzhiwa, pedram*}*@usc.edu*

Abstract*—Considering the growing number of Internet and cloud computing data centers in operation today and the high, yet flexible data center electric load, data centers can be good candidates to offer ancillary services and respond to regulation signals in a smart grid. This paper considers a problem whereby the smart grid employs both day-ahead dynamic energy prices and regulation signals to incentivize (cloud) data centers to simultaneously reduce their energy consumptions and participate in an ancillary service market. A data center controller schedules task dispatch and performs resource allocation in order to minimize the overall cost, which is the total electricity cost based on time-of-use energy prices minus any monetary compensations that data center may receive due to offering ancillary services. Moreover, the data center must satisfy average latency requirements in processing requests as specified in service-level agreements with clients. A two-tier hierarchical solution is presented for the data center controller, which achieves optimality in minimizing the overall cost with polynomial time complexity. Experimental results on Google trace demonstrate the effectiveness of the proposed solution in minimizing the overall cost in the data center.

I. INTRODUCTION

Information Technology (IT) plays an important part in people's lives at both personal and communal levels. IT services are increasingly used in different areas, e.g., business, commerce, education, manufacturing, and communication services. This leads to the demand for more and more computing resources to satisfy the computational and storage needs of IT customers and clients. As a result, there is a growing demand for Internet and cloud computing services, which in turn brings about rapid increase in the electric power consumption associated with IT infrastructure. This growing power demand precipitates adverse environmental impacts and imposes a significant strain on the Power System Infrastructure [1]-[5].

Large data centers—such as those owned by Google, Microsoft, Facebook, and Amazon—include tens of thousands of computing servers, tens of Petabyte of data storage, various cooling equipment, and power transformers [6], [7]. Although data centers try to make use of renewable energy to reduce their power drawn from the Power Grid, the intermittency of renewable energy sources prevents them from being a reliable source of energy [8]. As a matter of fact, we have witnessed a 56% increase in electricity used by data centers worldwide from 2005 to 2010, and today, 3% of the total electrical power consumption in United States is due to data centers [8]. As a result, electrical energy consumption is a major concern, and plenty of research has been devoted to reducing power dissipation of data centers by employing a variety of techniques e.g., CPU consolidation [9], [10], dynamic voltage and frequency scaling (DVFS) [11], [12], mechanisms to eliminate idle power waste [13], energy aware virtual machine replication and migration in data centers [14], power-aware geographical load balancing [15], and use of *electrical energy storage* systems [14], [15].

Due to monetary costs and environmental impacts, the power market is moving toward integrating less costly and environmentally friendly renewable energy sources (green energy) into the power grid facilities. For example, according to the US Department of Energy [18], wind energy should provide 20% of whole US needs for electricity. Due to the fact that green energy is an uncertain and intermittent source of energy, power market *Independent System Operators* (ISO) are now offering regulation service contracts with their customers. This incentive policy helps the power market customers to reduce their electricity cost, and helps the ISO to make the best use of renewable energy. Basically ISO sends a regulation signal reserve contract to every participant in the regulation market at the beginning of a billing period, specifying the desired power consumption trend for the following billing period (from several minutes to one day). By accepting this contract, the participant is required to adjust its power consumption to the regulation signal received from the ISO. During a billing period, ISO regularly sends regulation signals to market participants, marginally updating the initial regulation trend. This regulation signal can be sent every hour, every 30 minutes, or even every few minutes [19]. If the participant is able to satisfy the regulation signal demands, she will be charged a relatively lower energy price, otherwise, she has to pay the regular price or even pay a higher price for that period, depending on the specific contract [20].

Data centers are considered to be favorable facilities in a power regulation market, because (1) they are large consumers of power in the local Grid, and (2) they have some ability to shape their power consumption profile. The only question is whether data centers will actually benefit from participating in the regulation market. The issue is that although data centers can reduce their energy cost by participating in the power market, they can also incur additional penalties by failing to meet their service level agreements (SLAs) with their own clients. Data center service providers must therefore be convinced that the former benefit outweighs the latter cost.

In this paper, we consider the problem of maximizing the profit in a data center comprising a large number of potentially heterogeneous servers. The data center participates in a power regulation program in which it receives day-ahead energy prices and a regulation signal which dynamically sets regulation target and constraints for power consumption in the data center. The ISO sends this regulation signal at the beginning of each billing period (which is considered to be a single day) as the desired power dissipation profile during the current billing period. The data center attempts to meet the target power consumption profile (with some acceptable variance). Each billing period is divided into multiple time slots. For each time slot, the data center will be rewarded with a higher discount in energy price if it achieves a closer match with the desirable power dissipation value.

We adopt the *Generalized Process Sharing* (GPS) modeling, which is a generic stochastic model based on queuing theory [21], to describe the data center behavior. To be realistic, we use the Google data center trace [22] as workload for the data center. Also, we use the average response time for requests from each client as a measure of the *quality of service* (QoS). In other words, the SLA is interpreted as keeping the average response time for client requests below a pre-specified value. This figure of merit is usually used in web service data centers.

Based on the data center modeling and SLA constraints, we present a two-tier hierarchical solution for the data center controller, which performs optimal request dispatch and resource allocation in the data center. The proposed solution achieves optimality in minimizing the overall cost with polynomial time complexity. Experimental results demonstrate that data centers can significantly reduce their overall cost by exploiting the proposed approach.

The rest of this paper is organized as follows. In section II we introduce some related works done in this area. In section III, we formulate the problem and introduce the parameters. In section IV, we introduce our models and methodologies for solving the problem. In section V, our experimental results are presented. Finally in section VI, we conclude and introduce some future works in this area.

II. RELATED WORK

Power regulation market has received much interest in the recent past. A number of studies have focused on investigating the impact of regulation market on data centers, and answering the question of whether or not it is profitable for data centers to participate in the regulation market. In [8], Chen et al. investigated the ability of server clusters to meet the regulation market demands while meeting acceptable *quality of service* (QoS) for users. They proposed a dynamic server power regulation policy using dynamically arriving regulation signal requests from ISO, randomly arriving workload and probabilistic QoS constraints. Ghamkhari et al. proposed an optimization-based strategy to maximize the total profit in data centers [7]. Aksanli and Rosing developed a battery-based framework that helps with providing ancillary services as well as limiting peak power costs in data centers [23]. In [25], Chen et al. used server power capping techniques as well as various power states for servers in order to propose a dynamic power control policy for modulating power consumption in data centers in response to regulation signal requests from ISO.

III. SYSTEM MODEL

We consider a data center comprising N servers. Each server processes service requests based on a *first come first serviced* (FCFS) policy. This data center is equipped with a cooling system whose *coefficient of performance* (COP) is denoted by $COP_{cooling}$. Also we assume that the required electrical energy of the data center is supplied by the power grid, i.e., there are no on-site electrical energy production facilities. Fig. 1 provides an overview of various components in the system. In the following, we explain our system model in terms of power consumption, electrical energy pricing, workload management, and quality of service measurement.

As illustrated by (1), the total power dissipation in a data center consists of two components: the *processing power* ($P_{processing}$) and the *cooling power* ($P_{cooling}$). The processing power refers to the power consumed by servers which are the processing units in the data center. The servers produce a great amount of heat, which necessitates a robust cooling system in the data center. The power consumed by the cooling system is referred to as cooling power. The cooling power is assumed to be proportional to the processing power as shown in (2) [26]. Moreover, there is a coefficient $\frac{1}{eff}$ in (1) which accounts for the effectiveness of the power distribution network in the data center. The $COP_{cooling}$ and $\frac{1}{eff}$ coefficients are extracted from electrical and physical characteristics of the data center.

$$P_{datacenter} = \frac{1}{eff}\left(P_{processing} + P_{cooling}\right) \tag{1}$$

$$P_{cooling} = \frac{1}{COP_{cooling}} \cdot P_{processing} \tag{2}$$

The power consumption of a Blade server (employed in the data center) consists of two parts: the static and dynamic power. The static power in a server (P^{s}_{server}) refers to the power consumption of the server in idle state, when the server is not processing any user requests. P^{s}_{server} is given by $P^{s,max}_{server} \times \phi_{server}$ in which $P^{s,max}_{server}$ is the maximum static power consumption of the server depending on server characteristics, and ϕ_{server} denotes the portion of resources in the server that have been allocated for request processing (please note that the resources that are not allocated for request processing are assumed to be power gated.) The dynamic power consumption refers to the amount of power consumed by a server for processing requests. Dynamic power of a data center server can be modeled as $P^{d,max}_{server} \times U_{server}$ in which $P^{d,max}_{server}$ is the maximum dynamic power in the server depending on server characteristics, and U_{server} is the

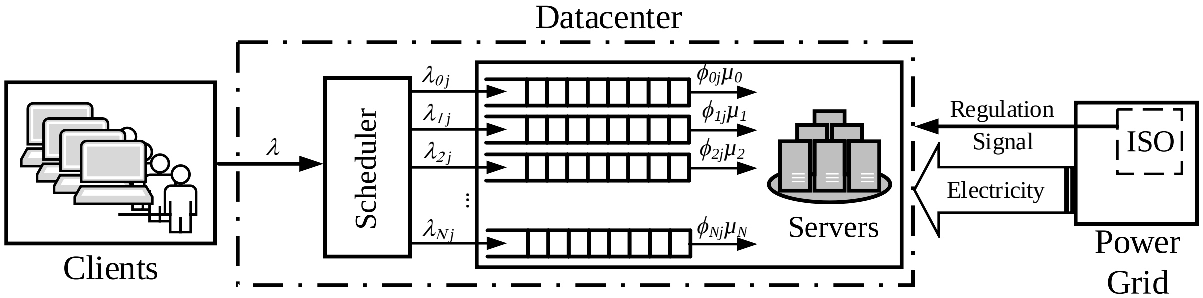

Fig. 1 The overview of the data center system.

utilization level of the server (will be derived later). The total power consumption in a server is given by:

$$P_{server} = P_{server}^{d,max} \cdot U_{server} + P_{server}^{s,max} \cdot \phi_{server} \tag{3}$$

where U_{server} and ϕ_{server} range from 0 to 1. The total processing-related power consumption in the data center is then derived as the summation of the power consumption of each server. As a result, the processing power of the data center is derived as (4).

$$P_{processing} = \sum_{i=1}^{N} P_{server}^{i} \tag{4}$$

In this paper we use a time-of-use (TOU) energy pricing model in which the data center receives the energy price function at the beginning of each billing period, which is one day. The energy price may be different at different times during the billing period. This energy price function will be valid until the end of current billing period. We also assume that the data center is participating in an ancillary service market, and receives the regulation requirements from the ISO in the beginning of each billing period. These regulation requirements will be valid during the current billing period.

An important part of our model is the regulation requirements. By participating in the ancillary service market, the data center controller should try to adjust the data center power consumption with the power regulation target. For each billing period, the data center should establish a power management policy to keep its power consumption close to the desired target levels (and trend) as determined by the regulation signal. As the power consumption in the data center is a function of utilization level of its servers, this power management policy can be interpreted as power-aware service request assignment and processing in the data center. The service request assignment and processing also depend on the arrival rate of service requests, and quality service guarantees. We divide the billing period into time slots. In each time slot, we assign service requests to servers based on energy price, regulation requirements, and quality of service requirements.

As stated earlier we adopt a GPS model for the data center service request management. In the GPS model, each server is assumed to maintain a service request queue for the service requests assigned to it, and process service requests in a FCFS basis. We model the service request queues using M/M/1 queueing model [27]. The processing rate of service requests in each server depends on the amount of allocated resources. Let μ_i denote the request processing rate of server i when all of its resources have been allocated. With a properly chosen time slot length, we can assume that server characteristics do not change during a time slot, and thus, we use ϕ_{ij} to represent the portion of allocated resources in server i during time slot j, which is defined as in (5) (ϕ_{ij} corresponds to ϕ_{server} as defined earlier). In this equation, ψ_i^{total} denotes the total processing resources in server i, and ψ_i^j denotes the amount of allocated resources in server i during time slot j.

$$\phi_{ij} \triangleq \frac{\psi_i^j}{\psi_i^{total}} \tag{5}$$

We assume that user requests arrive at the data center with rate λ_j during time slot j, with an exponentially distributed inter-arrival time (i.e., Poisson arrival model). Each server i in the data center receives a portion of the user requests arriving at the data center in every time slot j, denoted by λ_{ij}, as calculated in (6). In this equation, p_{ij} is the probability of assigning service requests to server i in time slot j.

$$\lambda_{ij} = p_{ij} \cdot \lambda_j \tag{6}$$

According to queuing theory principles [27], we can derive the busy time of a data center server during any time slot. The ratio of the busy time of server i during time slot j to the time slot length is $\frac{\lambda_{ij}}{\phi_{ij}\mu_i}$. The utilization level of the server during the time slot is proportional to this ratio as well as the portion of allocated resources ϕ_{ij}[1]. Hence, the utilization level of server i during time slot j is given by

$$U_{ij} = \frac{p_{ij} \cdot \lambda_j}{\mu_i} \tag{7}$$

[1] This is because the dynamic power consumption depends on both the amount of allocated resources and the portion of time that those available resources that are busy.

The reward model for participating in the ancillary service market in time slot j is given in (8)[2]. In this equation, $P_{datacenter}^{j}$ and P_{target}^{j}, respectively, denote the power consumption of data center and target power consumption during this time slot. ϵ_j and ζ_j are reward coefficients defined in the regulation signal. The reward will be maximized when $P_{datacenter}^{j}$ is equal to P_{target}^{j}, i.e. when the data center exactly meets the target power consumption set by the regulation signal.

$$R_{regulation} = Max\left\{0, \left(\epsilon_j^2 - \left(P_{datacenter}^{j} - P_{target}^{j}\right)^2\right) \cdot \zeta_j\right\} \quad (8)$$

The data center is responsible for providing an acceptable quality of service for its users. In this work, we use the average response time during the whole billing period (a day) to represent the quality of service. The data center needs to satisfy an upper limit on the average response time. The average response time depends on the average service time of requests processed in each server, which depends on the arrival rate of requests and characteristics of the server. The average response time in the data center over the billing period is given by:

$$Resp_{avg} = \frac{1}{M \cdot N} \sum_{j=1}^{M} \sum_{i=1}^{N} \frac{p_{ij}}{\phi_{ij} \cdot \mu_i - p_{ij} \cdot \lambda_j} \quad (9)$$

in which M denotes the number of time slots in the billing period and N denotes the number of servers in the data center.

IV. Problem Formulation

As shown in Fig. 2, we divide the billing period into M time slots, with equal length. Power consumption of the data center during time slot j can be written as in (10). In this equation, $P_{processing}^{j}$ denotes the processing power in the data center during time slot j.

$$P_{datacenter}^{j} = \frac{1}{eff}\left(1 + \frac{1}{COP_{cooling}}\right) \cdot P_{processing}^{j} \quad (10)$$

Service request arrival rate changes over time and usually there are high-load periods as well as near-idle periods in the data center during the billing period. Hence, we consider workload consolidation strategy in our formulation, in which we may decide to turn off/on a number of servers in the data center depending on the service request arrival rate in the beginning of each time slot. Please note that power gated servers can be turned on/off very rapidly, so, the delay to this process is much lower than the time-slot length which is 10 minutes.

Based on (3), (4), and (7), we calculate $P_{processing}^{j}$ using (11). In this equation, Y_{ij} denotes the on/off state of the server i during time slot j. The Y_{ij} value is 1 when the server is on, and

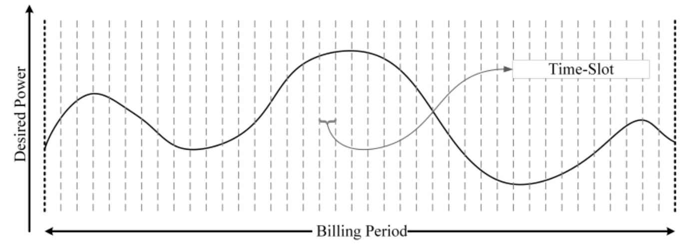

Fig. 2 Regulation signal over a billing period.

0 otherwise. We may change the state of the servers in the beginning of each time slot.

$$P_{processing}^{j} = \sum_{i=1}^{N} Y_{ij} \times \left(P_i^{d,max} \cdot \left(\frac{p_{ij}\lambda_j}{\mu_i}\right) + P_i^{s,max} \cdot \phi_{ij}\right) \quad (11)$$

We calculate the total energy cost of the data center in the billing period as:

$$EC_{DS} = \sum_{j=1}^{M} (EP_j \cdot P_{datacenter}^{j} - R_{regulation}^{j}) \quad (12)$$

where EP_j is the energy price in time slot j, and $R_{regulation}^{j}$ is derived in (8). We also have the equation for average response time $Resp_{avg}$ in data center from (9). The objective is to minimize the energy cost in data center (EC_{DS}) while keeping the average response time less than the acceptable threshold (δ)[3]. The optimization problem is formally described as follows:

$$\min_{Y_{ij},\phi_{ij},p_{ij}} \sum_{j=1}^{M} (EP_j \cdot P_{datacenter}^{j} - R_{regulation}^{j})$$

subject to:

$$\frac{1}{M \cdot N} \sum_{j=1}^{M} \sum_{i=1}^{N} \frac{p_{ij}}{\phi_{ij} \cdot \mu_i - p_{ij} \cdot \lambda_j} \le \delta$$

where:

$$P_{datacenter}^{j} = \frac{1}{eff} \times \left(1 + \frac{1}{COP_{cooling}}\right) \times \sum_{i=1}^{N} Y_{ij} \cdot \left(P_i^{d,max} \cdot \left(\frac{p_{ij} \cdot \lambda_j}{\mu_i}\right) + P_i^{s,max} \cdot \phi_{ij}\right)$$

$$R_{regulation}^{j} = Max\left\{0, \left(\epsilon_j^2 - \left(P_{datacenter}^{j} - P_{target}^{j}\right)^2\right) \cdot \zeta_j\right\}$$

$$Y_{ij} \in \{0, 1\}$$

$$\phi_{ij} \;\&\; p_{ij} \in [0, 1] \;\; and \;\; \sum_{i=1}^{N} p_{ij} = 1$$

$$\phi_{ij}\mu_i - p_{ij}\lambda_j > 0$$

$$\frac{p_{ij} \cdot \lambda_j}{\mu_i} < 1$$

According to aforementioned formulation, we aim at finding proper values for the variables Y_{ij}, p_{ij}, and ϕ_{ij} for every time slot j and every server i during the billing period.

[2] Reward calculation is usually based on power reserve as in [24], [7]. In this paper we assume that the regulation signal is available for the whole day. So, the we can write the reward function as in (8)

[3] Please note that our optimization framework is general in that the percentile of response time that exceeds a limit value can be considered as a metric in SLA, as shown in reference [7].

V. PROBLEM SOLUTION

Solving the overall optimization problem in one shot will have a high computation time because of the large number of optimization variables, and complexity of convex optimization algorithms. In order to reduce the computation complexity, we employ a hierarchical solution for the overall problem. More specifically, we first focus on a single time slot and minimize the energy cost during that time slot with different response time constraints, denoted by δ'. In other words, we derive a relationship between the response time constraint δ' and the minimum energy cost in the data center during each time slot. Based on this relationship in each time slot, we apply a dynamic programming approach to solve the overall optimization problem during the billing period.

We can write the optimization problem for time slot *j* as:

$$\min_{\phi_{ij},p_{ij}} EP_j \cdot P^j_{datacenter} - R^j_{regulation}$$

Subject to: $$\frac{1}{N}\sum_{i=1}^{N} \frac{p_{ij}}{\phi_{ij}\cdot\mu_i - p_{ij}\cdot\lambda_j} \le \delta'$$

Where: $$P^j_{datacenter} = \frac{1}{eff}\cdot\left(1+\frac{1}{COP_{cooling}}\right) \times \sum_{i=1}^{N}\left(P_i^{d,max}\cdot\left(\frac{p_{ij}\cdot\lambda_j}{\mu_i}\right) + P_i^{s,max}\cdot\phi_{ij}\right)$$

$$R^j_{regulation} = Max\left\{0, \left(\epsilon_j^2 - \left(P^j_{datacenter} - P^j_{target}\right)^2\right)\cdot\zeta_j\right\}$$

$$\phi_{ij} \;\&\; p_{ij} \in [0,1] \;\; and \;\; \sum_{i=1}^{N} p_{ij} = 1$$

$$\frac{p_{ij}\cdot\lambda_j}{\mu_{ij}} \le 1$$

$$\phi_{ij}\mu_i - p_{ij}\lambda_j > 0$$

$$\delta_{min} \le \delta' \le \delta_{max}$$

$$\frac{p_{ij}\cdot\lambda_j}{\mu_i} < 1$$

In this formulation, we define maximum and minimum values for response time constraints (constraint on δ'), denoted by δ_{min} and δ_{max}, respectively. δ_{min} depends on the service request arrival rate at data center and the processing rate of the servers. On the other hand, δ_{max} is set by the data center manager and determines the maximum acceptable response time to service requests during the time slot. This value should be greater than δ.

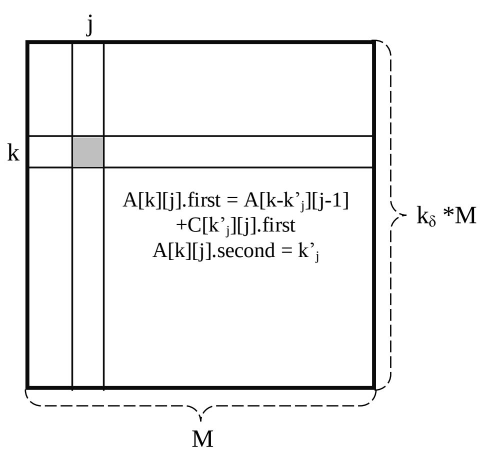

Fig. 3 Dynamic programming matrix (Matrix *A*) constructed to solve the problem over the billing period.

We can employ a convex optimization method to solve the aforesaid problem in each time slot *j* as explained next. In the optimization problem, whenever a server needs to be turned off (its power consumption is 0, and it is not involved in response time calculation), we can either assign an binary (0 or 1) coefficient Y_{ij} and set it to 0 (as shown in the formulation in the previous section), or set p_{ij} and ϕ_{ij} to 0. The latter one is more desirable since we do not want to add more complexity to the optimization problem by adding integer variables. So, we do not need to consider Y_{ij} in this optimization procedure. We observe that when either the p_{ij} values or the ϕ_{ij} values are fixed, the optimization problem becomes a convex optimization problem with convex objective function and convex inequality constraints. For the objective function, $P^j_{datacenter}$ is a linear function of ϕ_{ij}'s (or p_{ij}'s), and $R^j_{regulation}$ is a concave function of ϕ_{ij}'s (or p_{ij}'s) when $\left|P^j_{datacenter} - P^j_{target}\right| \le \epsilon_j$. Thus the objective function is a convex function of ϕ_{ij}'s with given p_{ij} values and vice versa. On the other hand, the constraints are either linear equality constraints, or convex inequality constraints of ϕ_{ij}'s with given p_{ij} values and vice versa. Hence, we adopt an iterative solution to solve the optimization problem in each time slot *j*. In an iterative manner, we derive the optimal ϕ_{ij} values with given p_{ij}'s, and derive the optimal p_{ij} values with given ϕ_{ij}'s, using standard convex optimization techniques in polynomial time complexity [28]. This is a near-optimal solution according to the theoretical analysis in [21].

We use Matlab to solve the problem (*fmincon* function). We have to solve the problem by using different values for δ' during different time slots. In this paper, we use *K*=100 different values for δ' uniformly distributed between δ_{min} and δ_{max} for each time slot.

Up to now, we have derived the minimum energy cost in each time-slot for each δ' value. We adopt a dynamic programming approach to find the minimum total energy cost during the billing period as explained next. We pick one of the $(\delta', EC^{j,min}_{datacenter})$ pairs from each time slot j. We pick these pairs so that the summation of the δ' values over all time-slots is less than $M\cdot\delta$. In order to make this problem solvable using a dynamic programming approach (i.e., solvable by filling tables), we map each δ' value (between δ_{min} and δ_{max}) to an index k (note that we have K different δ' values whose indices are from 1 to K). We build a matrix *C* in a way that every element $C[k][j]$ stores both the minimum energy cost ($C[k][j].first$) and the corresponding values for the optimization variables ($C[k][j].second$) in time slot *j* under the k^{th} value of response time constraint δ' (δ'_k). As shown in Fig. 3 we create a matrix *A* comprised of $k_\delta \cdot M$ rows and M columns (M is the number of time slots and $k_\delta \cdot M < K$) which will be used to derive the minimum energy cost in data center during the billing period. We calculate the elements in matrix A following the steps explained in Algorithm 1:

Algorithm 1: Creating matrix A.

Input: Matrix C, containing the energy cost values of data center during each time slot for different response time constraints and the corresponding values for optimization variables.

Output: Matrix A.

FOR $j = 1{:}\,M$
  FOR $k = 1{:}\,k_\delta \cdot M$
    IF j *equals to* 1 *THEN*
      $A[k][j].first \leftarrow \min\limits_{1 \le k_j \le k} \{C[k_j][j].first\}$
      $A[k][j].second \leftarrow \operatorname{Argmin}\limits_{1 \le k_j \le k} \{C[k_j][j].first\}$
    ELSE
      $k'_j \leftarrow \operatorname{Argmin}\limits_{1 \le k_j \le k} \{C[k_j][j].first + A[k-k_j][j-1]\}$
      $A[k][j].first \leftarrow A[k-k'_j][j-1] + C[k'_j][j]$
      $A[k][j].second \leftarrow k'_j$
    ENDIF
  ENDFOR
ENDFOR
RETURN A

In this algorithm, we assume that k_δ^{th} δ' value is equal to δ. This assumption will be valid by proper choosing of δ' values. Each element $A[k][j]$ maintains (i) a minimal energy cost value and (ii) the corresponding optimization variable values picking which assures that the summation of response times for the first j time slots will be less than or equal to δ'_k.

Using matrix A we can derive the minimum energy cost value in the data center during the billing period by following the steps in Algorithm 2.

Algorithm 2: Finding minimum energy cost value in data center during the billing period and proper values for optimization variables in each time slot which results in minimizing the energy cost during the billing period.

Input: Matrix C, containing the energy cost values of data center during each time slot for different response time constraints and the corresponding values for optimization variables.

Output: $EC_{datacente}^{min}$ and values for p_{ij} and ϕ_{ij} for all the servers during all the time slots which results in minimizing the energy cost in data center (vector B).

$k' \leftarrow \operatorname{Argmin}\limits_{1 \le k \le M \cdot k_\delta} \{A[k][M].first\}$
$EC_{datacente}^{min} \leftarrow A[k'][M].first$
$B[M] \leftarrow C[A[k'][M].second][M].second$
$temp \leftarrow A[k'][M].second$
FOR $j = M{:}\,2$
  $temp \leftarrow A[k' - temp][j-1].second$
  $k' \leftarrow k' - temp$
  $B[j-1] \leftarrow C[temp][j-1].second$
ENDFOR
RETUN $EC_{datacente}^{min}$ *and B*

In this algorithm we first find the minimum energy cost value in the data center during the billing period, satisfying the response time constraints. To do this we simply find the minimum energy cost value in the M^{th} column of matrix A, from row 1 to $M \cdot k_\delta$. Now, we have to find all the corresponding elements in matrix A that were involved in calculating the element we picked from the last column. For this, we use the backtrack method [29]. After finding all these elements, we can use matrix C to find the corresponding values for optimization variables.

The dynamic programming method described above produces an optimal solution to the problem in $O(M \cdot K\text{^}2)$ time complexity and with $O(M \cdot K \cdot N)$ space complexity.

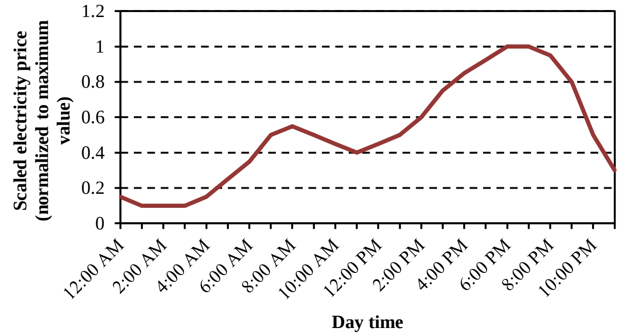

Fig. 4 Normalized time of use energy pricing profile during a day.

VI. Experimental Results

In the experiments we set the billing period to be one day and the time slot to be 10 minutes. Hence, one billing period is comprised of 144 time slots. During this period, we adopt a time-of-use energy price function as shown in Fig. 4. The energy price function is normalized by the maximum price during the billing period, and changes every hour during the billing period. The service request arriving rate during the billing period is shown in Fig. 5. This information is extracted from the google trace. We also set the characteristics of our benchmarking data center as shown in Table 1.

Table 1 Data center characteristics in the experiments.

Number of servers (N)	100
$COP_{cooling}$	5
$\frac{1}{eff}$	1.4

We conduct experiments on a data center with 100 servers, the power consumptions of which are estimated based on their processing resources (CPU and memory). We assume that the static and dynamic power consumption of a server has a linear relation with its processing resources ($\sim \alpha \cdot CPU + \beta \cdot MEM$). In the Google trace, resources of different servers are normalized with respect to the largest amount of resources among them.

Table 2 Different server types in the data center.

CPU	memory	count
1	1	6
0.5	0.75	8
0.5	0.5	55
0.5	0.25	29
0.25	0.25	2

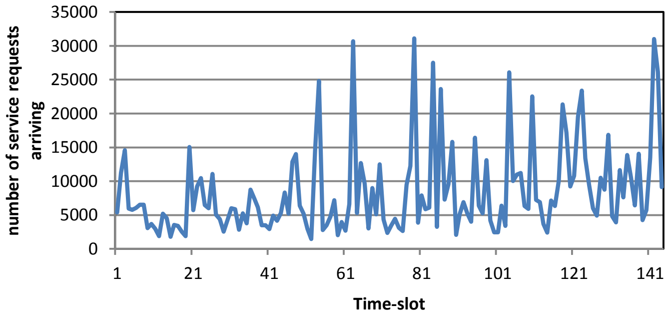

Fig. 5 Service request arrival rate at different time slots in a day.

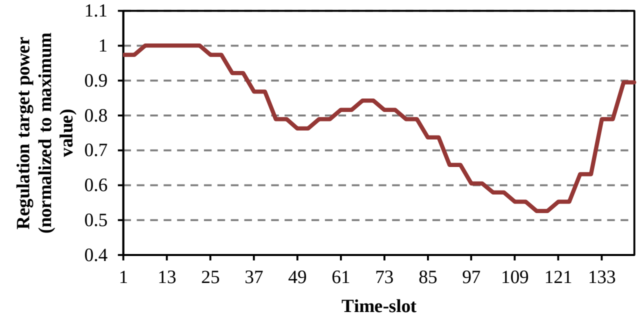

Fig. 6 The target regulation power profile during the billing period.

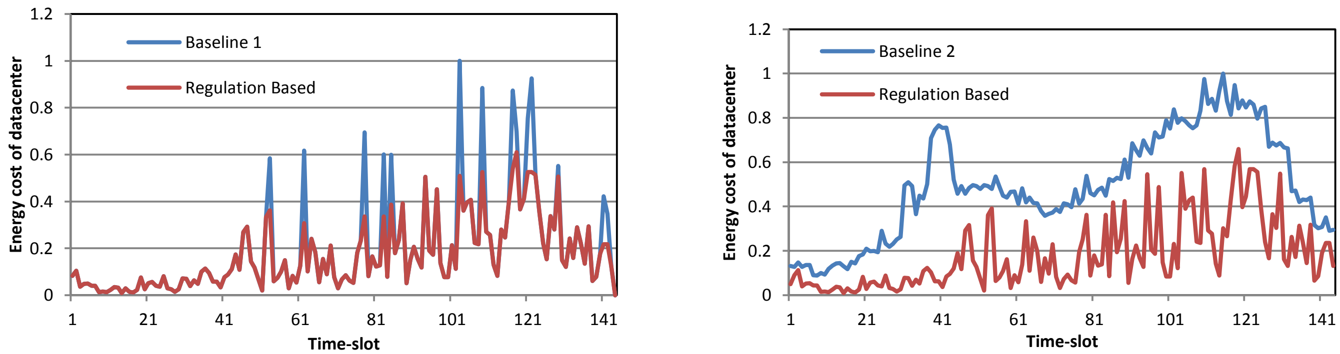

Fig. 7 The energy cost profile during the billing period using different approaches.

Table 2 shows the normalized amounts of resources in data center servers and the number of servers with that amount of resources.

We run our experiment for 144 time slots using the fmincon function in Matlab. We conduct the experiment with 100 different δ' values, which are uniformly distributed between the minimum possible response time of the data center (when the service request rate is the minimum among the Google trace and all the servers have allocated all their resources) and the maximum possible response time of the data center (when the service request rate is the maximum among the Google trace and the minimum number of servers are powered on using the minimum amount of resources). The minimum number of servers and minimum resources in the latter case are achieved by solving the optimization problem without considering the response time constraint, and trying to minimize the power consumption. We derive the optimization variable values that result in the minimum energy cost for different δ' values in each time slot. After that, we employ our dynamic programming approach so that the total energy cost of the data center is minimized in the billing period and the response time constraint is satisfied.

We also investigate two baseline approaches in which the data center is not participating in the regulation market and only tries to minimize its power consumption. In Baseline1 the data center uses the same request dispatch and resource allocation approach as the proposed system except that it does not participate in the regulation market, whereas in Baseline2 the incoming service requests are equally assigned to data center servers, i.e. all p_{ij} values are set to be $\frac{1}{N}$ and we find the proper ϕ_{ij} values to minimize the energy cost in the data center.

Fig. 6 shows the target regulation power profile during the billing period. The values in this figure are normalized to the highest value during the billing period. The regulation reward is set relative to the energy price value, and it ranges from 0 to 50 percent of the energy cost in the data center in each time slot, depending on how close is the power consumption in data center to the target regulation power during the time slot. In our simulations, we assume $\epsilon_j = 0.15 \cdot P^j_{target}$. We also set ζ_j based on the value of EP_j.

Table 3 illustrates the total energy cost values during the billing period for the two baseline approaches and our proposed approach. As one can see, the energy cost in Baseline2 is very high compared to Baseline1 and the proposed approach. Fig. 7 shows the energy cost values in the data center during each time slot. These values are normalized to the maximum energy cost value which happens in the 116th time slot using Baseline2.

Table 3 energy cost in the data center for different approaches.

Approach	Energy Cost
Baseline1	0.42
Baseline2	1
Regulation-based	0.36

As we can see, Baseline2 results in a huge energy cost in the data center. Also, there is a 13% reduction in the energy cost for the data center in the case of participating in the power regulation market as compared to the case in which we only optimize requests dispatch and resource allocation (Baseline1).

VII. CONCLUSION

This paper considers the scenario where the smart grid employs both day-ahead dynamic energy prices and regulation signals to incentivize data centers to simultaneously reduce energy cost and participate in an ancillary service market. The data center controller schedules task dispatch and resource allocation among servers in order to minimize its overall cost, which is the electricity cost based on time-of-use energy prices subtracted by the monetary compensation that the data center may receive due to offering ancillary services. Moreover, the cloud computing system needs to satisfy the average latency requirement in processing requests as specified in its service-level agreements (SLAs). A two-tier hierarchical solution is proposed for the data center controller, which achieves optimality in minimizing the overall cost with polynomial time complexity. Experimental results on Google traces demonstrate the effectiveness of the proposed solution on minimizing the overall cost of the cloud computing system.

Acknowledgment- This work was supported in part by a grant from the National Science Foundation.

VIII. REFERENCES

[1] M. Pedram, “Energy-Efficient Datacenters,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, no. 10, pp. 1465–1484, Oct. 2012.

[2] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, “Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility,” *Future Generation Computer Systems*, vol. 25, no. 6, pp. 599–616, Jun. 2009.

[3] L. A. Barroso and U. Holzle, “The Case for Energy-Proportional Computing,” *Computer*, vol. 40, no. 12, pp. 33–37, 2007.

[4] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, “A view of cloud computing,” *Communications of the ACM*, vol. 53, no. 4, pp. 50–58, 2010.

[5] J. Baliga, R. W. A. Ayre, K. Hinton, and R. Tucker, “Green Cloud Computing: Balancing Energy in Processing, Storage, and Transport,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 149–167, Jan. 2011.

[6] R. H. Katz, “Tech Titans Building Boom,” *IEEE Spectrum*, vol. 46, no. 2, pp. 40–54, Feb. 2009.

[7] M. Ghamkhari and H. Mohsenian-Rad, “Energy and Performance Management of Green Data Centers: A Profit Maximization Approach,” *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 1017–1025, Jun. 2013.

[8] Hao Chen, A. K. Coskun, and M. C. Caramanis, “Real-time power control of data centers for providing Regulation Service,” presented at the the 52nd IEEE Annual Conference on Decision and Control (CDC), 2013, pp. 4314–4321.

[9] Yanzhi Wang, Shuang Chen, H. Goudarzi, and M. Pedram, “Resource allocation and consolidation in a multi-core server cluster using a Markov decision process model,” presented at the the 14th International Symposium on Quality Electronic Design (ISQED), 2013, pp. 635–642.

[10] A. Chandra, W. Gong, and P. Shenoy, “Dynamic Resource Allocation for Shared Data Centers Using Online Measurements,” in *Quality of Service — IWQoS 2003*, K. Jeffay, I. Stoica, and K. Wehrle, Eds. Springer Berlin Heidelberg, 2003, pp. 381–398.

[11] Q. Wu, P. Juang, M. Martonosi, L.-S. Peh, and D. W. Clark, “Formal Control Techniques for Power-Performance Management,” *IEEE Micro*, vol. 25, no. 5, pp. 52–62, Sep. 2005.

[12] M. Ghasemazar, E. Pakbaznia, and M. Pedram, “Minimizing energy consumption of a chip multiprocessor through simultaneous core consolidation and DVFS,” presented at the IEEE International Symposium on Circuits and Systems (ISCAS), 2010, pp. 49–52.

[13] D. Meisner, B. T. Gold, and T. F. Wenisch, “PowerNap: Eliminating Server Idle Power,” in *the 14th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, New York, NY, USA, 2009, pp. 205–216.

[14] H. Goudarzi and M. Pedram, “Energy-Efficient Virtual Machine Replication and Placement in a Cloud Computing System,” presented at the the 5th IEEE International Conference on Cloud Computing (CLOUD), 2012, pp. 750–757.

[15] H. Goudarzi and M. Pedram, “Geographical Load Balancing for Online Service Applications in Distributed Datacenters,” in *the 6th IEEE International Conference on Cloud Computing*, 2013, pp. 351–358.

[16] D. S. Palasamudram, R. K. Sitaraman, B. Urgaonkar, and R. Urgaonkar, “Using batteries to reduce the power costs of internet-scale distributed networks,” in *the 3rd ACM Symposium on Cloud Computing*, San Jose, California, 2012, pp. 1–14.

[17] J. Wood, “UltraBattery?? Cloud Energy Storage for the Grid: Positioning Data Center and Telecommunication Backup Resources as Smart Grid Assets That Support,” presented at the the 35th International Telecommunications Energy Conference “Smart Power and Efficiency” (INTELEC), 2013, pp. 1–5.

[18] P. Harsha and M. Dahleh, “Optimal sizing of energy storage for efficient integration of renewable energy,” presented at the 50th IEEE Conference on Decision and Control and Control Conference (CDC-ECC), 2011, pp. 5813–5819.

[19] A. L. Ott, “Experience with PJM market operation, system design, and implementation,” *IEEE Transactions on Power Systems*, vol. 18, no. 2, pp. 528–534, May 2003.

[20] “PJM Manual 11: Energy & Ancillary Services Market Operations.” PJM Interconnection, 2014.

[21] A. K. Parekh and R. G. Gallager, “A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-node Case,” *IEEE/ACM Trans. Netw.*, vol. 1, no. 3, pp. 344–357, Jun. 1993.

[22] J. Wilkes and C. Reiss, “Details of the ClusterData-2011-1 trace,” 10-May-2013. .

[23] B. Aksanli and T. Rosing, “providing regulation services and managing data center peak power budgets,” presented at the Design, Automation & Test in Europe (DATE), Germany, 2014.

[24] Hao Chen, M. C. Caramanis, and A. K. Coskun, " Dynamic Server Power Capping for Enabling Data Center Participation in Power Markets," Proceedings of the International Conference on Computer-Aided Design (ICCAD), 2013, pp. 105-112.

[25] Hao Chen, M. C. Caramanis, and A. K. Coskun, “The data center as a grid load stabilizer,” presented at the the 19th Asia and South Pacific Design Automation Conference (ASP-DAC), 2014, pp. 105–112.

[26] E. Pakbaznia and M. Pedram, “Minimizing Data Center Cooling and Server Power Costs,” in *Proceedings of the 14th ACM/IEEE International Symposium on Low Power Electronics and Design*, New York, NY, USA, 2009, pp. 145–150.

[27] L. Kleinrock, *Queueing systems.* New York: Wiley, 1974.

[28] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[29] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, Second Edition*, 2nd edition. McGraw-Hill, 2001.