Resource Allocation Optimization in a Data Center with Energy Storage Devices

Shuang Chen, Yanzhi Wang, Massoud Pedram
Department of Electrical Engineering
University of Southern California
Los Angeles, USA
{shuangc, yanzhiwa, pedram}@usc.edu

***Abstract*—As cloud computing is becoming the new generation of computing paradigm because of its many attractive attributes, huge data centers are built and operated to host the cloud services. Since these data centers usually incur a high electricity bill, the problem of reducing the electricity cost and maximizing the profit for a data center operator arises naturally. Because of the trend of dynamic pricing policies in the energy market, in which the electricity price changes across different hours of a day, the use of energy storage devices, such as batteries and supercapacitors, in a data center can be extended in addition to judicious computing/memory/storage resource management policies as another way to cut down on the operational cost. In this paper, we formulate a generalized optimization problem to minimize the linear combination of the electricity cost and the average request response time in a data center with energy storage devices. Solutions based on convex optimization techniques are proposed and the experimental results are discussed to demonstrate the effectiveness of the proposed formulation and the solution methods.**

I. Introduction

Cloud computing has been envisioned as the new generation of computing paradigm for its major advantages in on-demand self-service, ubiquitous network access, location independent resource pooling, and transference of risk [1]. In a cloud computing system, the computation and storage demands are shifted from users at network edges to a *Cloud*, in which resource abundant data centers are hosted. By accessing the business applications in the form of cloud services through the network, an application user can reduce or even eliminate the costs related to "in-house" resource provisioning for these applications. An overview of the cloud computing architecture can be found in [2].

To accommodate the ever-growing demand for cloud services, massive data centers are built which consume a large amount of electricity power, thus leading to soaring utility bills. For instance, Microsofts data center in Quincy, Washington consumes 48 megawatts which is enough to power 40,000 homes [3]. And according to the data from [4], [5], utility bills contribute an average of 24% of the total amortized cost of a data center. Consequently, controlling the utility costs is an essential problem for a data center operator to maximize its profit. On the other hand, due to the service level agreement (SLA) formed between the service provider and the clients, the service provider often ends up over-provisioning the resources [6], which causes an increase in the total power consumption of the servers and adds to the electricity bill. Therefore, the well known problem of optimal resource allocation has been studied by a series of prior work [7]–[11]. For instance, [10] addresses the resource allocation problem in a multi-tier cloud computing system aimed at profit maximization.

Because of the trend of dynamic energy pricing [12], the change energy price has become another factor that can significantly affect the total energy cost in addition to the total power consumption of a data center. In the case where the time-of-use (TOU) pricing policies are adopted by the utility companies, i.e. the electricity price changes at different times of a day, another natural thought to cut down on the utility bills is to reshape the energy consumption profile by adjusting the utilization level of the servers in the data center during different hours of a day. With the same energy consumption budget provided, raising utilization levels when the electricity price is relatively lower and lowering the utilization levels otherwise will no doubt result in lower electricity bill cost compared to the policy in which the utilization levels of the servers remain the same at all time. Although promising at first sight, the efficacy of this approach is greatly dependent on the incoming pattern of the service requests. Since one major motivation of dynamic pricing is to balance the power demand in different hours of a day, the peak hours with high energy prices are likely to overlap those with more user activities and higher request incoming rate. When a large number of service requests are arriving at the data center, the utilization levels of the servers cannot be lowered much even if the energy price is high, since under-provisioning of resources will cause SLA violations, the penalty of which can overwhelm the benefits of energy saving.

As a solution to the problem stated above, the usage of energy storage devices (ESDs), such as the UPS unit [13], which is originally designed to deal with power outages, can be extended to reserve energy in off-peak hours for the use in the peak hours. The utilization of ESDs in data centers is also studied from various aspects by some other work. Authors of [14] proposed an control algorithm to minimize the electricity cost based on the i.i.d assumption of the workload profile. And the problem of "peak shaving" is considered in [15] in order to reduce the infrastructure cost of a data center. Moreover, in the presence of energy sell-back policies, e.g. Renewable Energy Standard Offer Program[1], data center operators can also make use of the enlarged capacity of the ESDs to make profit from the electricity price difference between the peak hours and the off-peak hours.

[1] Toronto Hydro: http://www.torontohydro.com

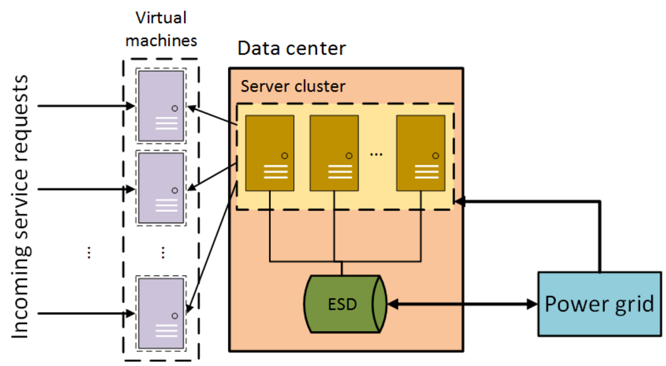

Fig. 1. System framework of a data center with an ESD

Despite all the work that focus on request dispatching/resource allocation or the usage of ESDs, none of them is able to identify the inter-dependency between the two problems and to formulate a combined problem that determine the management policies for both the server cluster and the ESDs. In fact, the workload of the data center, which is one of the major parameter in the problem formulation of [14], can be more accurately modeled using the information of the resource allocation scheme among the servers. At the same time, as is stated earlier, a data center operator can benefit from the extended usage of ESDs by exploiting resource allocation policies more aggressively. In this paper, we jointly consider the request dispatching, resource allocation, and the usage of ESDs to balance between the average response time of the requests and the total electricity bill. The response time of a service request is modeled using the generalized processor sharing (GPS) model [16], [17], the power consumption of the data center is calculated based on the utilization level of each server, and the charging/discharging of a unified ESD at the data center level is considered. A generalized optimization problem considering energy sell-back is formulated, and some solution methods based on convex optimization [18] techniques are proposed.

The rest of this paper is organized as follows: the system model is introduced in Section II. Section III shows the formulation of the optimization problem. Section IV presents the solution methods. The experimental results are presented in Section V. And the last section is the conclusion.

II. System Model

In this paper, we consider a data center comprised of a set of heterogeneous servers with different amount of computation resources and/or power consumptions, and an array of rechargeable batteries is installed as the ESD. Incoming service requests are dispatched among the servers, and the electricity power drawn from the power grid can be distributed between the server cluster and the battery array. The system framework is shown in Fig. 1.

A. Pricing model

The utility company is considered to adopt a TOU (time-of-use) pricing policy which has potentially different electricity prices in different time slots (from several minutes to one hour) of the day. Let L denote the number of time periods (slots) of a day. The duration of the l-th time period is denoted by T_l. During the l-th time period, the price of buying a unit amount of electricity energy is denoted by $Price_{B,l}$, while the price of selling a unit amount of electricity energy back to the power grid is denoted by $Price_{S,l}$. Generally speaking, $Price_{B,l}$ and $Price_{S,l}$ are higher during peak hours than off-peak hours.

B. Request arrival and processing model

In order to find an analytical form of the average response time of a service request, the arrival of the service requests is assumed to follow a Poisson process with an average arrival rate of λ_l during the l-th time period. For the k-th server in the data center which has K servers in total, the average processing rate is denoted by μ_k. For any given k and l, if the probability that a service request is dispatched to server k is denoted by $p_{k,l}$ during the l-th time period, then the arrival of service requests at the k-th server is known to follow a Poisson process with an average arrival rate of $p_{k,l} \cdot \lambda_l$ according to the properties of exponential distributions. According to the GPS model, if $\phi_{k,l}$ denotes the portion of resources allocated by the k-th server to process the incoming service requests during the l-th time period, the average response time of the requests processed by the k-th server during the l-th time period, denoted by $R_{k,l}$, can be calculated as

$$R_{k,l} = \begin{cases} \dfrac{1}{\phi_{k,l} \cdot \mu_k - p_{k,l} \cdot \lambda_l}, & p_{k,l} > 0 \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

The average response time of service requests over the whole day, denoted by $\bar{R}$, can be calculated as

$$\bar{R} = \frac{\sum_l \left(\lambda_l \cdot T_l \cdot \sum_k p_{k,l} R_{k,l}\right)}{\sum_l \left(\lambda_l \cdot T_l\right)} \tag{2}$$

C. ESD model

In this paper, we consider the case that a single array of rechargeable batteries is installed at the data center level with the maximum energy capacity of $E_{S,max}$. The amount of energy stored in the batteries at the beginning of the l-th time period is denoted by $E_{S,l}$. To retain the capability of handling power outages and to ensure safety, a minimum amount of energy, denoted by $E_{S,min}$, must be kept within the batteries during any time period. The amount of energy charged to and discharged from the batteries during the l-th time period are denoted by $E_{C,l}$ and $E_{D,l}$, respectively. There exists an upper bound of the amount of energy that can be charged to or discharged from the battery, denoted by $E_{C,max,l}$ and $E_{D,max,l}$, respectively, because of the limited power density of the batteries and power capacity of conversion circuitry. Moreover, a charging efficiency, denoted by η_C, and a discharging efficiency, denoted by η_D, are included in the proposed model to account for the energy conversion loss between the power grid, the batteries, and the power consuming components in the data center.Please note that although both $E_{C,l}$'s and $E_{D,l}$'s are introduced in order to simplify the problem formulation, it is intuitively true and can be proved straightforwardly that they cannot have non-zero values at the same time for any given time period l. In other words, the ESD will not be charged and discharged in the same time period.

D. Power consumption model

While modeling the total electricity power consumption of a data center, one should not only consider the power consumption of the servers, but also account for the power consumption overhead, e.g. the power consumption of the power delivery infrastructure, networking, and the cooling system, etc. If we use E_{usage} to denote the power usage effectiveness (PUE) factor [19] of the data center, which reflects the proportion of electricity power consumed by components other than the servers and ranges between 1.2 to more than 2.0, then, based on [20], the equivalent power consumption of the k-th server during the l-th time period, denoted by $P_{k,l}$, can be calculated as

$$\begin{aligned} P_{k,l} =& P_{idle,k} + (E_{usage} - 1) \cdot P_{peak,k} \\ &+ (P_{peak,k} - P_{idle,k}) \cdot \phi_{k,l} + \epsilon \end{aligned} \tag{3}$$

where $P_{idle,k}$ is the power consumption of the k-th server when it is idle, $P_{peak,k}$ is the peak power consumption of the k-th server when its utilization level reaches 100%, and ϵ is a constant. Since $P_{k,l}$ is a linear function of $\phi_{k,l}$, Eqn. (3) can be rewritten as

$$P_{k,l} = P_{con,k} + P_{lin,k} \cdot \phi_{k,l} \tag{4}$$

where $P_{con,k}$ and $P_{lin,k}$ are coefficients that can be calculated based on the values of $P_{idle,k}$, $P_{peak,k}$, E_{usage}, and ϵ. If the k-th server is turned off during the l-th time period, then the server does not consume any power and only a part of the power consumption overhead is incurred, which will be denoted by $P_{off,k}$. If we let $x_{k,l} = 1$ in the case that the k-th server is turned on during the l-th time period, and $x_{k,l} = 0$ otherwise, then the total power consumption related to the server cluster in the l-th time period, denoted by $E_{serv,l}$, can be calculated as

$$E_{serv,l} = T_l \cdot \sum_k \left[x_{k,l} P_{k,l} + (1 - x_{k,l}) P_{off,k} \right] \tag{5}$$

And the total amount of energy drawn from the power grid by the data center during the l-th time period, denoted by $E_{total,l}$, can be expressed as

$$E_{total,l} = E_{serv,l} + E_{C,l} - \eta_D \cdot E_{D,l} \tag{6}$$

$E_{total,l} > 0$ means that the energy flows from the power grid to the data center, while $E_{total,l} < 0$ means that the energy flows from the data center back to the power grid.

III. Problem formulation

Provided the request arrival rate, the energy price function, and specifications of the data center and the battery array, we formulate the problem of determining the resource allocation scheme for the servers and the charging/discharging scheme for the battery array as an optimization problem as shown below:

Find $x_{k,l}$'s, $p_{k,l}$'s, $\phi_{k,l}$'s, $E_{S,l}$'s, $E_{C,l}$'s, and $E_{D,l}$'s

Minimize

$$\sum_l C_{elec,l}(E_{total,l}) + k_{delay} \cdot \bar{R} \tag{7}$$

Subject to

$$\phi_{k,l} \cdot \mu_k - p_{k,l} \cdot \lambda_l \geqslant 0, \quad \forall k,l \tag{8}$$
$$\bar{R} \leqslant R_{max} \tag{9}$$
$$E_{S,l+1} = E_{S,l} + \eta_C \cdot E_{C,l} - E_{D,l}, \quad \forall l \tag{10}$$
$$E_{S,L+1} = E_{S,0} \tag{11}$$
$$\textstyle\sum_k p_{k,l} = 1, \quad \forall l \tag{12}$$
$$\phi_{k,l} \leqslant x_{k,l}, \quad \forall k,l \tag{13}$$
$$p_{k,l} \geqslant 0, \quad \forall k,l \tag{14}$$
$$\phi_{k,l} \geqslant 0, \quad \forall k,l \tag{15}$$
$$E_{S,l} \in [E_{S,min}, E_{S,max}], \quad \forall l \tag{16}$$
$$E_{C,l} \in [0, E_{C,max,l}], \quad \forall l \tag{17}$$
$$E_{D,l} \in [0, E_{D,max,l}], \quad \forall l \tag{18}$$
$$x_{k,l} \in \{0,1\}, \quad \forall k,l \tag{19}$$

where $C_{elec,l}(\cdot)$ is the energy cost function defined as

$$C_{elec,l}(E) = \begin{cases} Price_{B,l} \cdot E, & E \geqslant 0 \\ Price_{S,l} \cdot E, & E < 0 \end{cases} \tag{20}$$

The control variables are the set of servers to be turned on during each time period, the probabilities to dispatch a request to each server, the proportion of resources that each server allocates for the service requests, the amount of energy stored in the battery array in each time period, and the amount of the energy charged to/discharged from the battery array. The objective function is set as a linear combination of the total energy cost per day and the average response time per service request in order to achieve a desirable balance between the power consumption and the delay performance and to maximize the profit gained by the data center operator. $E_{total,l}$ is calculated as in Eqn. (6), $\bar{R}$ is calculated as in Eqn. (2), and k_{delay} is the coefficient that adjusts the weight of the average response time and may vary based on different SLAs. Constraint (8) ensures that the amount of resources allocated to an incoming request flow is enough to process it with finite waiting time. Constraint (9) sets the maximum tolerable average response time for the incoming service requests. Constraint (10) addresses the change of amount of stored energy in the battery array between different time periods. Constraint (11) makes sure that the amount of energy charged to and discharged from the battery array are balanced and the battery array will have the same amount of stored energy at the beginning of each day. Constraint (12) ensures that all the incoming service requests will be dispatched to a server. Constraint (13) prohibits any resource allocation from a server that is currently turned off. Constraints (14) - (19) set the domain of the control variables. It is worth noting that the case that energy sell-back is not allowed is also included in the formulation which correspond to the case that all energy sell-back prices, $Price_{S,l}$'s are set to zero. We omit the proof of equivalence since it is straightforward.

Generally speaking, the problem formulated as above is a mixed-integer non-linear programming (MINLP) problem, which does not have any efficient solution methods. Even if all the integer variables $x_{k,l}$'s are given in prior, the problem is still difficult to solve using conventional convex optimization techniques [18] because the objective function is neither convex nor concave.

Set initial values $p_{k,l}^{(0)}$'s, $\phi_{k,l}^{(0)}$'s, $E_{S,l}^{(0)}$'s, $E_{C,l}^{(0)}$'s, and $E_{D,l}^{(0)}$'s
$k \leftarrow 0$
repeat
  Find the optimal values for $p_{k,l}^{(k+1)}$'s, $E_{S,l}^{(k+1)}$'s, $E_{C,l}^{(k+1)}$'s, and $E_{D,l}^{(k+1)}$'s based on $\phi_{k,l}^{(k)}$'s
  Find the optimal values for $\phi_{k,l}^{(k+1)}$'s, $E_{S,l}^{(k+1)}$'s, $E_{C,l}^{(k+1)}$'s, and $E_{D,l}^{(k+1)}$'s based on $p_{k,l}^{(k+1)}$'s
  $k \leftarrow k+1$
until solution converges

Fig. 2. Pseudo code for the solution of resource allocation and battery management problem

IV. SOLUTION METHODS

Because of the non-convexity of the problem and the existence of the integer-valued variables, one can only obtain the optimal solution by using exhaustive search or some stochastic algorithm, e.g. simulated annealing [21] or genetic algorithm [22]. In this section, we propose two solution methods that use convex optimization techniques to derive near-optimal solutions.

First of all, we decompose the problem into two subproblems, namely, (i) the problem of resource allocation and battery management, which finds the optimal policy of request dispatching, resource allocation, and the amount charging/discharging energy, and (ii) the problem of server consolidation which determines the set of servers to be turned on. In this way, the continuous variables and the integer-valued variables are separated.

A. Solution to the problem of request dispatching, resource allocation, and battery management

While solving this problem, we assume that the values of all $x_{k,l}$'s are given. Based on the assumption that $Price_{S,l} \leqslant Price_{B,l}$ for any given l, i.e. the price to sell electricity back to the power grid cannot exceed the price to buy electricity from the power grid at any time (please note that this also includes the case that energy sell-back is not allowed), the energy cost function, $C_{elec,l}(\cdot)$, as defined in Eqn. (20), can be rewritten as

$$C_{elec,l}(E) = \max\left(Price_{B,l} \cdot E, Price_{S,l} \cdot E\right) \tag{21}$$

Since $E_{total,l}$ as calculated in Eqn. (6) is a linear function of $\phi_{k,l}$'s, $C_{elec,l}(\cdot)$ is a convex function that only depends on the values of $\phi_{k,l}$'s.

Although the problem is still non-convex because of the relationship between $R_{k,l}$, $p_{k,l}$, and $\phi_{k,l}$ as specified in Eqn. (1), it can be further decomposed into two set of convex problems: fixing the value of $p_{k,l}$'s, the problem is convex with respect to variables $\phi_{k,l}$'s, $E_{S,l}$'s, $E_{C,l}$'s, and $E_{D,l}$'s, while fixing the value of $\phi_{k,l}$'s, the problem is convex with respect to variables $p_{k,l}$'s, $E_{S,l}$'s, $E_{C,l}$'s, and $E_{D,l}$'s. Therefore, an iterative solution method consisting of a request dispatching phase and a resource allocation phase can be used to solve this problem. In the request dispatching phase, the values of $p_{k,l}$'s, $E_{S,l}$'s, $E_{C,l}$'s, and $E_{D,l}$'s are found subject to constraints (8) - (11), (13), (15) - (18). In the resource allocation phase, the values of $\phi_{k,l}$'s, $E_{S,l}$'s, $E_{C,l}$'s, and $E_{D,l}$'s are found subject to constraints (8) - (12), (14), (16) - (18). The pseudo code description of the algorithm is shown in Fig. 2.

In the special case that the purchase price and the sell-back price of electricity are the same at any time, i.e. $Price_l := Price_{B,l} = Price_{S,l}$, an alternative solution method based on a hierarchical solution framework can be used to achieve better result.

First, for any specific time period l, if the charging and discharging of the battery array is not considered and the maximum allowable latency is set to $R_{l,max}$, then the problem of finding the optimal request dispatching and resource allocation policy can be formulated as follows

Find $p_{k,l}$'s and $\phi_{k,l}$'s for a given l

Minimize

$$C_{serv,l}\left(R_{l,max}\right)$$

Subject to

$$\phi_{k,l} \cdot \mu_k - p_{k,l} \cdot \lambda_l \geqslant 0, \quad \forall k \tag{22}$$
$$\bar{R}_l \leqslant R_{l,max} \tag{23}$$
$$\textstyle\sum_k p_{k,l} = 1, \tag{24}$$
$$\phi_{k,l} \leqslant x_{k,l}, \quad \forall k \tag{25}$$
$$p_{k,l} \geqslant 0, \quad \forall k \tag{26}$$
$$\phi_{k,l} \geqslant 0, \quad \forall k \tag{27}$$

where $C_{serv,l}\left(R_{l,max}\right)$ is the minimum cost that can be achieved without using the ESD under the maximum average latency of $R_{l,max}$ and can be expressed as

$$C_{serv,l}\left(R_{l,max}\right) = Price_l \cdot E_{serv,l} + k_{delay} \cdot \frac{\lambda_l T_l \bar{R}_l}{\sum_l \lambda_l T_l} \tag{28}$$

$\bar{R}_l$ is the average response time in the l-th time period which can be calculated as

$$\bar{R} = \sum_k p_{k,l} R_{k,l} \tag{29}$$

with $R_{k,l}$'s specified as in Eqn. (1). Near-optimal solution can be obtained from this problem by iteratively finding the optimal values for $p_{k,l}$'s and $\phi_{k,l}$'s. The algorithm is shown in Fig. 3. By solving the problem with different values of $R_{l,max}$, one can build up a lookup table which stores the minimum overall cost with different average latency constraints. Please note that more generalized lookup tables that includes the values of $C_{serv,l}\left(R_{l,max}\right)$ under different electricity prices and/or request incoming rates can also be used to improve computational efficiency in the case that the user's behavior or the pricing policy changes frequently.

Provided that the solution for the problem of request dispatching and resource allocation can be found locally with any given $R_{max,l}$, from the global perspective, the original problem can be transformed into the problem of finding the optimal $R_{max,l}$'s as well as the charging/discharging policy for the battery array (i.e. $E_{S,l}$'s, $E_{C,l}$'s, and $E_{D,l}$'s), which is formulated as follows:

Find $R_{l,max}$'s, $E_{S,l}$, $E_{C,l}$'s, and $E_{D,l}$'s

Set initial values $p_{k,l}^{(0)}$'s and $\phi_{k,l}^{(0)}$'s
$k \leftarrow 0$
repeat
Find the optimal values for $p_{k,l}^{(k+1)}$'s based on $\phi_{k,l}^{(k)}$'s
Find the optimal values for $\phi_{k,l}^{(k+1)}$'s based on $p_{k,l}^{(k+1)}$'s
$k \leftarrow k+1$
until solution converges

Fig. 3. Pseudo code for solving the request dispatching and resource allocation problem for each time period

Minimize

$$\sum_l \left[C_{serv,l}\left(R_{l,max}\right) + Price_l \cdot \left(E_{C,l} - \eta_D \cdot E_{D,l}\right)\right]$$

Subject to

$$\frac{\sum_l \left(\lambda_l \cdot T_l \cdot R_{l,max}\right)}{\sum_l \left(\lambda_l \cdot T_l\right)} \leqslant R_{max} \tag{30}$$

as well as constraints specified in (10), (11), (16) - (18). Constraint (30) ensures that the maximum average latency in each time period is set in a way such that the global maximum average latency constraint can be satisfied.

Generally speaking, $C_{serv,l}\left(R_{l,max}\right)$ is a monotonically decreasing convex function with respect to $R_{l,max}$, i.e., the total cost related to the server will increase super-linearly if the maximum average latency is reduced. Therefore, the problem of finding the maximum average latency for each time slot as well as the charging/discharging policy is a convex optimization problem because the objective function and the inequality constraints are convex, and the equality constraints are affine. Standard solvers can be used to obtain the optimal solution with polynomial time complexity.

B. Solution to the problem of server consolidation

In the problem of server consolidation, we find the optimal set of servers to be turned on in each time period which is controlled by boolean variables $x_{k,l}$'s. When the size of the problem is small, one can run exhaustive searches or apply branch and bound algorithms to find the optimal consolidation policy. In the case that the problem size is relatively large and the computational efficiency is a major consideration, we can use a heuristic based on a greedy algorithm to find the near-optimal consolidation policy.

The algorithm uses a loop to find the values for $x_{k,l}$'s. The algorithms as proposed in Section IV-A are used as kernel algorithms in each iteration of the loop to solve the problem of optimal request dispatching, resource allocation, and ESD management and calculate the total cost. Initially, all $x_{k,l}$'s are set to 1, i.e. all servers are turned on. In each iteration, we search for the (k, l) pair which results in the lowest utilization level. The corresponding $x_{k,l}$ is set to 0 and the new total cost is calculated and compared with the previous one. The process continues until no more cost reduction can be achieved. The framework of the algorithm is shown in Fig. 4.

$x_{k,l} \leftarrow 1, \forall k, l$
Calculate the total cost, C_{total}
repeat
$C_{total,0} \leftarrow C_{total}$
Among all (k, l)'s such that $x_{k,l} = 1$, find (k_{min}, l_{min}) such that $\phi_{k_{min},l_{min}} = \min\{\phi_{k,l}\}$
Set $x_{k_{min},l_{min}} \leftarrow 0$
Calculate the total cost, C_{total}
until $C_{total} > C_{total,0}$

Fig. 4. Pseudo code for the algorithm to solve the problem of server consolidation

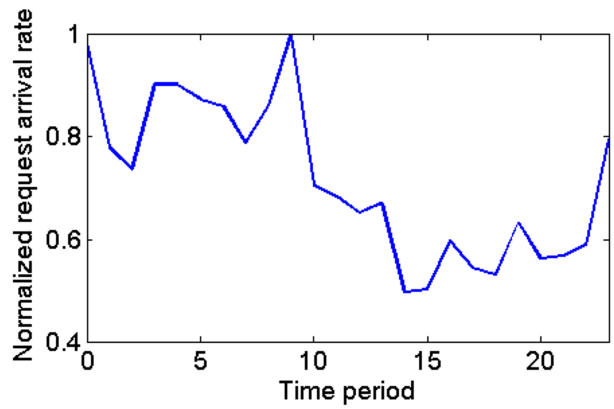

Fig. 5. Request arrival rate within one day

V. Experimental Results

In this section, we present the simulation results on some selected scenarios. Please note that we use normalized parameter values rather than real values for power and delay related parameters in the simulation.

In the first scenario, we consider the case that the electricity sell-back price is the same as the purchase price, and discuss how the ESD management, combined with proper request dispatching and resource allocation policies, can help reduce the total cost as defined in Eqn. (7). In this part, the server consolidation schemes are not considered, and the second algorithm proposed in Section IV-A which has a two-level hierarchical structure is applied. Each day is divided into 24 time periods, each lasting for one hour. A server cluster of 10 servers are considered. The request arrival rate in different time periods of a day is set according to the data extracted from the 29-day period Google cluster dataset[2] as shown in Fig. 5 with the maximum arrival rate among all time periods (i.e. $\max \lambda_l$) of 5. The processing rate of each server, defined as μ_k, follows a uniform distribution on $[1, 2]$. $P_{con,k}$ follows a uniform distribution on $[0.2, 0.3]$, $P_{lin,k}$ is set to two times $P_{con,k}$, and $P_{off,k}$ is set to 20% of $P_{con,k}$. The electricity price in each time period, defined as $Price_l$, follows a uniform distribution on $[0.1, 0.2]$. The total energy capacity of the ESD, defined as $E_{S,max}$, may vary from 0 to 50. The amount of reserved energy in the ESD, defined as $E_{S,min}$, is set to 10% of $E_{S,max}$, and the maximum amount of energy to be charged/discharged in each time period, defined as $E_{C,max}/E_{D,max}$, may vary from 10% to 90% of $E_{S,max}$ in all time periods. The efficiency of

[2]https://code.google.com/p/googleclusterdata/

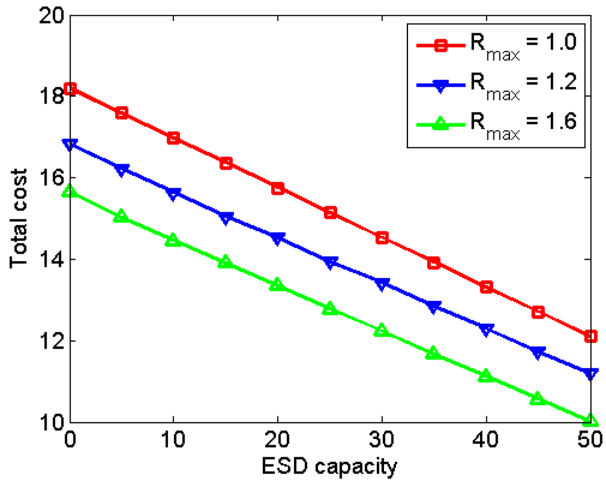

Fig. 6. Relation between the cost function and the ESD capacity

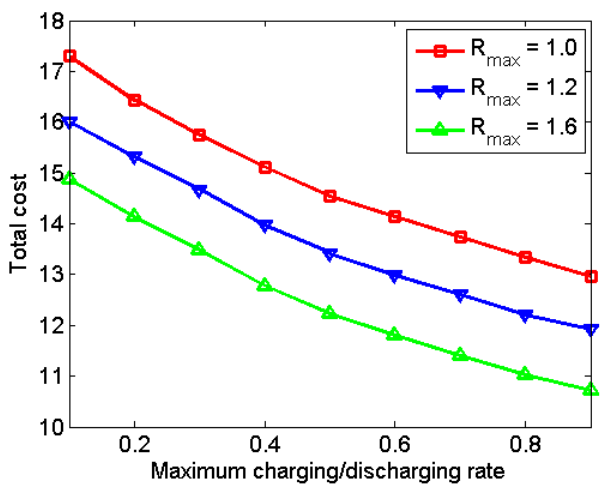

Fig. 7. Relation between the cost function and the maximum charging/discharging rate

charging and discharging are both set to 0.9. The coefficient k_{delay}, as specified in Eqn. (7), is set to 1. When setting $E_{C,max}$ and $E_{D,max}$ to 50% of $E_{S,max}$, and setting the maximum allowable average latency, defined as R_{max}, to different values, the relationship between the cost function and $E_{S,max}$ is shown in Fig. 6. As can be seen from the figure, the total cost is a monotonically decreasing function of $E_{S,max}$ with all R_{max}'s that are simulation. When $R_{max} = 1$, the reduction on the total cost is as large as 34% when having an ESD with a maximum capacity of 50 compared to the case in which no ESD is used. When setting $E_{S,max}$ to 30, and varying $E_{C,max}/E_{D,max}$ from 10% of $E_{S,max}$ to 90% of $E_{S,max}$, the total cost under different R_{max} conditions are shown in Fig. 7. As can be seen from the figure, increasing the maximum amount of charging/discharging energy in each time period will increase the amount of cost saving. Comparing to the case that $E_{C,max}/E_{D_max}$ is set to 10% of $E_{S,max}$, setting $E_{C,max}/E_{D_max}$ to 50% of $E_{S,max}$ further reduces the total cost by 16% when $R_{max} = 1$.

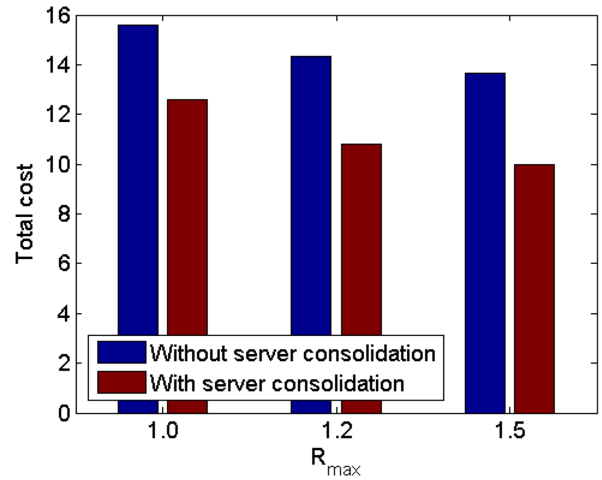

Fig. 8. The effect of server consolidation on the total cost

In the second scenario, we consider the case that energy sell-back is not allowed, and discuss the effect of server consolidation. In this case only the first algorithm proposed in Section IV-A as specified in Fig. 2 is applicable. $E_{S,max}$ is set to 50, $E_{C,max}$ and $E_{D,max}$ are set to 50% of E_{max}, and all other parameters are the same as in the first scenario. The total costs with and without server consolidation are shown in Fig. 8. As can be seen from the figure, the server consolidation results in 19% - 27% cost reduction.

VI. Conclusion

In this paper, a generalized problem is formulated to address the problem of request dispatching, resource allocation, and ESD management in a data center. The objective function of a linear combination of the total electricity cost per day and the average response time per request is considered. Solutions based on convex optimization techniques are proposed to solve the formulated problem efficiently. Experimental results show that both the ESD management algorithm and the server consolidation have significant effects on reducing the total cost.

References

[1] B. Hayes, "Cloud computing," *Commun. ACM*, vol. 51, no. 7, pp. 9–11, Jul. 2008.

[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, 2009.

[3] R. Katz, "Tech titans building boom," *Spectrum, IEEE*, vol. 46, no. 2, pp. 40–54, 2009.

[4] L. A. Barroso and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 4, no. 1, pp. 1–108, 2009.

[5] M. K. Patterson, D. Costello, P. Grimm, and M. Loeffler, "Data center tco; a comparison of high-density and low-density spaces," *Thermal Challenges in Next Generation Electronic Systems (THERMES 2007)*, 2007.

[6] L. A. Barroso and U. Hölzle, "The case for energy-proportional computing," *IEEE computer*, vol. 40, no. 12, pp. 33–37, 2007.

[7] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. De Rose, and R. Buyya, "Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," *Software: Practice and Experience*, vol. 41, no. 1, pp. 23–50, 2011.

[8] R. Buyya and M. Murshed, "Gridsim: A toolkit for the modeling and simulation of distributed resource management and scheduling for grid computing," *Concurrency and computation: practice and experience*, vol. 14, no. 13-15, pp. 1175–1220, 2002.

[9] K. Krauter, R. Buyya, and M. Maheswaran, "A taxonomy and survey of grid resource management systems for distributed computing," *Software: Practice and Experience*, vol. 32, no. 2, pp. 135–164, 2002.

[10] H. Goudarzi and M. Pedram, "Multi-dimensional sla-based resource allocation for multi-tier cloud computing systems," in *Cloud Computing (CLOUD), 2011 IEEE International Conference on*. IEEE, 2011, pp. 324–331.

[11] Y. Wang, S. Chen, H. Goudarzi, and M. Pedram, "Resource allocation and consolidation in a multi-core server cluster using a markov decision process model." in *ISQED*, 2013, pp. 635–642.

[12] A. Ipakchi and F. Albuyeh, "Grid of the future," *Power and Energy Magazine, IEEE*, vol. 7, no. 2, pp. 52–62, 2009.

[13] A. M. Jungreis, "Uninterruptible power supply," Feb. 6 2001, US Patent 6,184,593.

[14] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam, "Optimal power cost management using stored energy in data centers," in *Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '11. New York, NY, USA: ACM, 2011, pp. 221–232. [Online]. Available: http://doi.acm.org/10.1145/1993744.1993766

[15] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: What, where, and how much?" *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 1, pp. 187–198, Jun. 2012.

[16] 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 Transactions on Networking (TON)*, vol. 1, no. 3, pp. 344–357, 1993.

[17] Z.-L. Zhang, D. Towsley, and J. Kurose, "Statistical analysis of generalized processor sharing scheduling discipline," in *ACM SIGCOMM Computer Communication Review*, vol. 24, no. 4. ACM, 1994, pp. 68–77.

[18] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[19] R. Brown *et al.*, "Report to congress on server and data center energy efficiency: Public law 109-431," 2008.

[20] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," *ACM SIGARCH Computer Architecture News*, vol. 35, no. 2, pp. 13–23, 2007.

[21] P. J. Van Laarhoven and E. H. Aarts, *Simulated annealing*. Springer, 1987.

[22] D. Whitley, "A genetic algorithm tutorial," *Statistics and computing*, vol. 4, no. 2, pp. 65–85, 1994.