

UC Riverside

2016 Publications

Title

Extending Demand Response to Tenants in Cloud Data Centers via Non-intrusive Workload Flexibility Pricing

Permalink

<https://escholarship.org/uc/item/0d3054q9>

Authors

Zhan, Yong
Ghamkhari, Mahdi
Xu, Du
et al.

Publication Date

2016-11-15

DOI

10.1109/TSG.2016.2628886

Peer reviewed

Extending Demand Response to Tenants in Cloud Data Centers via Non-intrusive Workload Flexibility Pricing

Yong Zhan, *Student Member, IEEE*, Mahdi Ghamkhari, *Student Member, IEEE*, Du Xu, *Member, IEEE*, Shaolei Ren, *Member, IEEE*, and Hamed Mohsenian-Rad, *Senior Member, IEEE*

Abstract—Participating in demand response programs is a promising tool for reducing energy costs in data centers by modulating energy consumption. Towards this end, data centers can employ a rich set of resource management knobs, such as workload shifting and dynamic server provisioning. Nonetheless, these knobs may *not* be readily available in a cloud data center (CDC) that serves cloud tenants/users, because workloads in CDCs are managed by tenants themselves who are typically charged based on a usage-based or flat-rate pricing and often have no incentive to cooperate with the CDC operator for demand response and cost saving. Towards breaking such “split incentive” hurdle, a few recent studies have tried market-based mechanisms, such as dynamic pricing, inside CDCs. However, such mechanisms often rely on complex designs that are hard to implement and difficult to cope with by tenants. To address this limitation, we propose a novel incentive mechanism that is not dynamic, i.e., it keeps pricing for cloud resources unchanged for a long period. While it charges tenants based on a Usage-based Pricing (UP) as used by today’s major cloud operators, it rewards tenants proportionally based on the time length that tenants set as deadlines for completing their workloads. This new mechanism is called Usage-based Pricing with Monetary Reward (UPMR). We demonstrate the effectiveness of UPMR both analytically and empirically. We show that UPMR can reduce the CDC operator’s energy cost by 12.9% while increasing its profit by 4.9%, compared to the state-of-the-art approaches used by today’s CDC operators to charge their tenants.

Keywords—Demand response, monetary reward, split incentive, cloud data center, time-shiftable load, demand delaying.

I. INTRODUCTION

To support the emergence of numerous cloud computing services, power-hungry data centers have collectively consumed 38 GW electricity world-wide as of 2012 (an increase of 63% compared to 2011), placing a surging pressure on operators to optimize energy management [1]. Further, combined with the growing electricity prices, energy cost has taken up nearly 15% of data center operator’s total cost of ownership [2].

Consequently, as driven by the increasing popularity of real-time pricing [3] and peak pricing [4], *data center demand response*, broadly interpreted as reshaping the energy consumption of data centers, has been surfacing as a crucial approach for saving energy cost [5]. Towards this end, several resource management knobs, such as workload shifting/scheduling [6]–[11] have been extensively leveraged to shift workloads to time periods with lower electricity prices and/or shave peak power demand to avoid high demand charges. These knobs,

albeit appealing for energy cost saving, often result in some performance losses. For instance, although delay-tolerant batch workloads often do not have as stringent delay requirements as interactive workload, shifting them to time periods with low electricity prices for cost saving still results in longer completion time and performance degradation.

While (enterprise) data centers serving internal workloads may leverage various resource management knobs and trade performance degradation for energy cost saving or vice versa, public cloud data center (CDC) operators serving cloud tenants cannot readily do so, because workloads are managed by tenants while CDC operators must meet the rigid service level agreements (SLAs). Specifically, tenants in a CDC are often charged based on their usages of CDC resources [12], regardless of when they use the cloud resources. Thus, tenants would simply submit their workloads whenever available and complete them as soon as possible, even if their workload is delay-tolerant and can be shifted to later times without much inconvenience. On the other hand, due to the lack of control of tenants’ workloads, CDC operators can only *passively* process the submitted workloads in the order tenants submit them. As a consequence, there exists a “split-incentive” hurdle in CDCs: operators would like to defer some workloads to time periods with lower electricity prices for energy cost saving, but do not have control over tenants’ workloads to enforce their wish.

Towards breaking the “split incentive” hurdle, a few recent studies [4], [9], [13]–[15] have begun to investigate market-based mechanisms (e.g., dynamic pricing) that provide tenants with monetary incentives to cooperate with CDC operators for cost saving. These pricing mechanisms set prices for cloud resources dynamically, thereby encouraging tenants to use cloud resources at times with lower electricity prices such that the CDC operator can reduce its energy cost, too.

While many of the existing market-based mechanisms are promising for reducing the CDC’s energy cost and/or lowering tenants’ cloud costs, they often rely on complex designs, which may be difficult to implement in practice. For example, [4], [13] attribute each tenant’s cost based on its contribution to the overall energy bill, and this requires the CDC operator reveals to tenants sufficient information about its energy bill and each tenant’s contribution. In addition, time-dependent dynamic pricing, as often considered in these prior studies, result in cost uncertainties/concerns for tenants, who, in order to opportunistically launch their cloud resources for workload processing, have to frequently adjust their scheduling decisions based on real-time price information. This is particularly unappealing for small/medium tenants who may see their costs increased due to the lack of required expertise to properly respond to dynamic pricing or other market mechanisms.

Y. Zhan and D. Xu are with the Key Laboratory of Optical Fiber Sensing and Communications, University of Electronic Science and Technology of China, Chengdu, China. M. Ghamkhari, S. Ren, and H. Mohsenian-Rad are with the Department of Electrical Engineering, University of California, Riverside, CA, USA. This work was done when Y. Zhan was a Visiting Student at the Smart Grid Research Lab, University of California at Riverside.

To address the aforementioned limitations of current pricing approaches [4], [9], [13]–[15], we propose a new and practical pricing policy, called Usage-based Pricing with Monetary Reward (UPMR), which *keeps pricing for cloud resources unchanged for a long period yet still being able to exploit tenants’ workload scheduling flexibilities via rewards*. Specifically, while UPMR charges tenants based on a Usage-based Pricing (UP) as used by today’s major cloud operators, it rewards tenants *proportionally* based on the time length that tenants set as deadlines for completing their workloads.

By using UPMR, tenants with no delay tolerance for their workloads are not affected, whereas tenants who are willing to defer their workloads for financial compensation are given an opportunity to reveal their workload scheduling flexibilities, i.e., by which time their workloads need to be completed. Thus, on one hand, UPMR is non-intrusive to tenants, and participating in UPMR is fully voluntary. On the other hand, the CDC operator can exploit the scheduling flexibilities provided by motivated tenants and properly schedule these workloads for energy cost saving subject to deadline constraints.

To derive UPMR, we will first present a detailed model for the interactions between tenants and the CDC operator, and then formulate the operator’s (non-convex) problem of energy cost minimization, which is decomposed into several convex sub-problems. Using real-world workload traces, we show that UPMR can effectively reduce the operator’s energy cost by 12.9% and increase its profit by 4.9%, compared to state-of-the-art approaches used by today’s CDC operators.

II. RELATED WORK

Optimizing data center energy efficiency has attracted great attention in the past years. For example, “power proportionality” has been investigated in various contexts, such as dynamically turning on/off servers for interactive workloads [7], [16], [17], tuning processor speed for online data-intensive services [18]. Other techniques include energy storage control in concert with IT resource management [19]–[21], thermal-aware load scheduling [22], geographic load balancing [23]–[25], renewable generation proportionality [26], [27], among others. More recently, there has been a growing interest in data center demand response [28]. These studies, however, all assume that data center operators have full control over the workloads, but this assumption does not hold in CDCs where workloads are managed by individual tenants.

To enable demand response within CDCs, various market-based pricing schemes have been recently proposed to align tenants’ interests with the operator’s, such that these two separate parties can cooperate towards a desired goal [4], [9], [13]–[15]. More specifically, [4], [15] considered fair energy cost attribution and split the cost among tenants based on their individual contributions to the overall energy bill. Similarly, in [13] the authors proposed to fairly share the energy bill among tenants by incentivizing them to disclose their workload flexibility to the CDC operator. In these studies, the CDC operator needs to provide to tenants information about electricity rates as well as the peak power demand, which may be difficult to implement. Further, tenants’ costs of using

cloud resources may become highly unpredictable under these pricing policies, causing business concerns/inconveniences for tenants’ budgeting. In [14], the authors proposed dynamic pricing for cloud resources, propagating the total energy bill onto tenants’ costs. In this case, tenants are incentivized to modulate their workloads based on the real-time price information provided by the CDC operator and hence need to periodically adjust their workload scheduling decisions. This represents a significant barrier, especially for small/medium tenants who do not have the required expertise to do so.

In contrast, UPMR is a practical mechanism that is easy to implement and uses monetary rewards to extract tenants’ workload scheduling flexibilities, allowing the CDC operator to manage these workloads accordingly for demand response and energy cost saving. A key advantage of UPMR, as compared to the existing solutions is that, tenants’ participation is fully voluntary: only tenants who are interested in trading their workload deferments for rewards need to provide scheduling flexibilities, while other uninterested tenants are not affected.

III. SYSTEM MODEL

In this section, we model both cloud tenants’ and the CDC operator’s decisions, based on a time-slotted system.

A. Cloud Tenants

Tenants with different workloads/jobs naturally have different sensitivities to postponing their workloads. For example, batch workloads such as MapReduce-based tasks can often tolerate a longer deferment without affecting tenants’ normal businesses, than user-interactive workloads such as web search. Here, we classify tenants into I types, each one having a different sensitivity with respect to workload latency. More specifically, we use a *revenue loss* factor $\kappa_i \in \mathbb{R}_0^+$ to quantify type- i tenants’ sensitivities to deferring their workloads: the more sensitive to deferment, the larger κ_i . For tenants running interactive applications with no delay tolerance, we set $\kappa_i \rightarrow \infty$ and refer to these tenants as *inelastic* tenants. All other types of tenants are referred to as *elastic* tenants for whom deferring workloads is acceptable to a certain extent.

Suppose that type- i elastic tenants’ (potential) revenue losses are nondecreasing and convex in the maximum number of time slots D_i by which their workloads can be deferred. In practice, tenants’ workloads are often processed prior to the maximum deferment of D_i time slots, and thus D_i can be viewed as the maximum delay tolerance for type- i elastic tenants. For example, we consider the following formulation to measure type- i elastic tenants’ revenue losses per Ψ requests¹:

$$L_i = \kappa_i D_i. \quad (1)$$

We also use the following reward function to be offered to type- i tenants per Ψ requests due to their workload flexibility:

$$\gamma_i = \rho \log(1 + D_i), \quad (2)$$

¹Our model also applies to any other revenue loss function, provided that it is nondecreasing and convex over deferment parameter D_i .

where the reward factor ρ is always non-negative and the deferment threshold specified by tenants should satisfy

$$0 \leq D_i \leq D_{max}, \quad (3)$$

in which $D_{max} \in \mathbb{Z}_0^+$ is a fixed positive integer number.

Tenants' Decisions: We assume that tenants are rational, i.e., each tenant makes decision on its deferment threshold D_i for maximizing its net benefit (reward minus revenue loss). Thus, tenant's optimization problem can be formulated as

$$\underset{D_i}{\text{maximize}} \quad \gamma_i - L_i, \quad \text{s.t., constraint (3)}, \quad (4)$$

where the reward γ_i and revenue loss L_i are given by (2) and (1), respectively. Note that, although D_i should be set as an integer value due to time-slotted system model, we relax it to be continuous to simplify the problem (4). Thus, the optimization problem (4) is convex, whose solution is

$$D_i = \max \left(\min \left(\frac{\rho}{\kappa_i} - 1, D_{max} \right), 0 \right). \quad (5)$$

B. CDC Operator

The CDC operator earns revenue by serving tenants' workloads and incurs energy cost for workload processing. By using UPMR, the CDC operator also has a cost of rewards paid to incentivize workload deferment by elastic tenants.

Revenue: Without loss of generality, we assume that each request of user will consume the same amount of CDC resources, which includes CPU, memory, storage and bandwidth, etc. Suppose $\delta \in \mathbb{R}_0^+$ per Dollar denotes the price of CDC resources for processing Ψ requests. τ denotes the amount of time slots over a billing cycle. We can formulate the revenue that is earned by the CDC operator over a billing cycle as

$$\text{Revenue} = \delta \sum_{t=1}^{\tau} \sum_{i=1}^I \frac{\lambda_i[t]}{\Psi}, \quad (6)$$

where $\lambda_i[t] \in \mathbb{Z}_0^+$ denote the number of service requests that are generated by type- i tenants at time slot t .

Reward: We write the total rewards over a billing cycle as

$$\text{Reward} = \sum_{i=1}^I \gamma_i \sum_{t=1}^{\tau} \frac{\lambda_i[t]}{\Psi}. \quad (7)$$

We assume that the CDC operator knows or can estimate² tenants' revenue loss factors κ_i . Thus, the CDC operator can calculate tenants' responses to its offered reward rate according to (5). By replacing D_i in (2) with (5), we get

$$\gamma_i = \rho \log \left(1 + \max \left(\min \left(\frac{\rho}{\kappa_i} - 1, D_{max} \right), 0 \right) \right). \quad (8)$$

Wear-and-Tear Cost: The CDC operator incurs a wear-and-tear cost during operation, e.g., during server on/off and

battery charge/discharge cycles, which can be formulated as

$$\begin{aligned} \text{Wear} = & \zeta \sum_{t=1}^{\tau} \max\{-S[t], 0\} \\ & + w_{on} \sum_{t=1}^{\tau} m_{on}[t] + w_{of} \sum_{t=1}^{\tau} m_{of}[t], \end{aligned} \quad (9)$$

where $\zeta \in \mathbb{R}_0^+$ (\$/KWh) measures the estimated battery wear-and-tear cost, $S[t]$ (KWh) denotes the energy consumption for charging the storage system at time slot t , and $S[t] < 0$ means that the storage system discharges energy to power up the CDC as a supplement to grid power. As for parameters $m_{on}[t]$ and $m_{of}[t]$, they denote the number of physical machines that are turned on/off at time slot t , respectively. In addition, $w_{on} \in \mathbb{R}_0^+$ (\$) and $w_{of} \in \mathbb{R}_0^+$ (\$) measure the wear-and-tear cost of machine due to turning on/off machine, respectively.

Energy Cost: Typically, the CDC operator is charged for electricity by the utility based on two parts: *energy charge* and *demand charge*. Here, energy charge is calculated based on the amount of energy consumption, while demand charge is calculated according to the peak demand, e.g., measured over each 15 minutes interval within a billing cycle [30]. Thus, the CDC's energy cost over a billing cycle can be calculated as

$$\text{Bill} = \sum_{t=1}^{\tau} \alpha[t]P[t] + \sum_{j=1}^J \frac{\beta_j}{T} \max_{t \in A_j} P[t], \quad (10)$$

where T denotes the length of each time slot (e.g., hour), $P[t] \in \mathbb{R}_0^+$ (KWh) denotes the energy usage at time slot t , $\alpha[t] \in \mathbb{R}$ (\$ per KWh) denotes the energy price at time slot t , $\beta_j \in \mathbb{R}_0^+$ (\$ per KW) denotes the price of type- j demand charge,³ and A_j denotes the set of time slots falling into the type- j demand charge window. The considered energy billing model is quite general and includes Time-Dependent Pricing (TDP) as a special case, where the CDC is charged only for energy consumption based on time-dependent prices.

Energy consumption. The total energy usage of the CDC can be calculated at each time slot by taking into consideration the server energy consumption, non-IT energy (captured by power usage effectiveness, i.e., the ratio of total data center energy to IT energy), overheads due to turning machines on/off and the storage capacity of CDC. Thus, we can write

$$P[t] = \max(E_{pue}(P_m[t]T + P_o[t]) + S[t], 0), \quad \forall t, \quad (11)$$

where $E_{pue} \in [1, \infty)$ denotes power usage effectiveness, $P_m[t] \in \mathbb{R}_0^+$ (KW) denotes the power of *active* machines that are being used, and $P_o[t] \in \mathbb{R}_0^+$ (KWh) is the energy overhead for turning on/off machines at time slot t . As in [32], the power of active machines is expressed as

$$P_m[t] = m[t](P_{idle} + (P_{peak} - P_{idle})u[t]), \quad \forall t, \quad (12)$$

where $m[t] \in \mathbb{R}_0^+$ denotes the number of active machines, $u[t] \in [0, 1]$ denotes the average server utilization, $P_{idle} \in \mathbb{R}_0^+$ (KW) and $P_{peak} \in [P_{idle}, \infty)$ (KW) denote the energy usage

²In practice, the CDC operator can learn tenants' loss factors from experiments or historical data, using, e.g., *calibration period of pilot trials* [29].

³Utilities may impose multiple demand charges for different time intervals, e.g., daytime/night or winter/summer demand charge [31].

of a machine in idle and fully utilized conditions, respectively. Next, we calculate the amount of active machines by

$$m[t] = \sum_{t'=1}^t (m_{on}[t'] - m_{of}[t']) + m[0], \quad \forall t, \quad (13)$$

where $m[0] \in \mathbb{Z}_0^+$ is the number of switched on machines at time 0. We formulate the average server utilization as

$$u[t] = \frac{\sum_{i=1}^I \hat{\lambda}_i[t]}{Nm[t]}, \quad \forall t, \quad (14)$$

where $\hat{\lambda}_i[t] \in \mathbb{R}_0^+$ is the service requests of type- i tenants scheduled at time slot t and $N \in \mathbb{Z}_0^+$ is the (average) number of service requests that can be hosted by a machine in a time slot. Here, the number of *scheduled* requests is modeled as

$$\hat{\lambda}_i[t] = \lambda_i[t] - \phi_i[t] + \eta_i[t], \quad \forall i, t, \quad (15)$$

where $\phi_i[t]$ and $\eta_i[t]$ denote the deferred requests of type- i tenants generated and scheduled at time slot t , respectively.

The power usage of turning machines on/off is modeled as

$$P_o[t] = o_{on}m_{on}[t] + o_{of}m_{of}[t], \quad \forall t, \quad (16)$$

where $o_{on} \in \mathbb{R}_0^+$ (KWh) and $o_{of} \in \mathbb{R}_0^+$ (KWh) denote the energy consumption of turning on/off a machine, respectively.

Operational Constraints: Next, we list the constraints that the CDC operator faces when it makes control decisions.

First, the reward factor and the number of machines that are turned on/off at each time slot must be non-negative:

$$\rho \geq 0, \quad (17)$$

$$m_{on}[t] \geq 0, m_{of}[t] \geq 0, \quad \forall t. \quad (18)$$

Second, the number of active machines should be large enough to process requests scheduled at each time slot:

$$\sum_{t'=1}^t (m_{on}[t'] - m_{of}[t']) + m[0] \geq \frac{\sum_{i=1}^I (\lambda_i[t] - \phi_i[t] + \eta_i[t])}{N}, \quad \forall t, \quad (19)$$

which ensures that the average utilization as shown in (14) is always no greater than 1.

Third, if the CDC is equipped with an on-site energy storage unit, then the energy storage unit's charge and discharge rates should be limited by the charger inverter's ratings:

$$-l_{out}T \leq S[t] \leq l_{in}T, \quad \forall t, \quad (20)$$

where $l_{out} \in \mathbb{R}_0^+$ (KW) and $l_{in} \in \mathbb{R}_0^+$ (KW) denote the energy storage unit's maximum discharge and charge rates, respectively. The amount of energy stored in the storage unit should be non-negative and not exceed its storage capacity:

$$0 \leq \sum_{t'=1}^t S[t'] + S[0] \leq C_s, \quad \forall t, \quad (21)$$

where $S[0] \in \mathbb{R}_0^+$ (KWh) is the stored energy at time 0 and $C_s \in \mathbb{R}_0^+$ (KWh) is the storage capacity.

Fourth, $\phi_i[t]$ and $\eta_i[t]$ satisfy

$$0 \leq \phi_i[t] \leq \lambda_i[t], \quad \forall i, t, \quad (22)$$

$$\eta_i[t] \geq 0, \quad \forall i, t. \quad (23)$$

That is, the amount of deferred requests generated at time slot t must be non-negative and no more than the amount of requests newly generated at time slot t . Also, the amount of deferred requests scheduled at time slot t must be non-negative.

Fifth, considering that the amount of requests generated from the beginning of a billing cycle up to any time slot must be no less than the amount of requests that are scheduled within these periods, we must also have

$$\sum_{t'=1}^t (\lambda_i[t'] - \phi_i[t'] + \eta_i[t']) \leq \sum_{t'=1}^t \lambda_i[t'], \quad \forall i, t, \quad (24)$$

which can be simplified to be as

$$\sum_{t'=1}^t (-\phi_i[t'] + \eta_i[t']) \leq 0, \quad \forall i, t. \quad (25)$$

Last but not least, to ensure that the requests of type- i tenants generated from the beginning up to the time slot t should be scheduled no later than time slot $t + \lfloor D_i \rfloor$, where $\lfloor \cdot \rfloor$ represents a floor function, we have

$$\sum_{t'=1}^{t+\lfloor D_i \rfloor} (\lambda_i[t'] - \phi_i[t'] + \eta_i[t']) \geq \sum_{t'=1}^t \lambda_i[t'], \quad \forall i, t, \quad (26)$$

which can be simplified to be as

$$\sum_{t'=t+1}^{t+\lfloor D_i \rfloor} \lambda_i[t'] + \sum_{t'=1}^{t+\lfloor D_i \rfloor} (-\phi_i[t'] + \eta_i[t']) \geq 0, \quad \forall i, t. \quad (27)$$

That is, the CDC cannot defer tenants' requests beyond their deferment thresholds. Note that, $\forall t > \tau$, $\lambda_i[t] = 0$, $\phi_i[t] = 0$ and $\eta_i[t] \geq 0$, which indicates that the CDC may schedule some of tenants' requests at time slot in the next billing cycle without violation of the thresholds specified by the tenants.

IV. UPMR PRICING ALGORITHM: A PROFIT MAXIMIZATION APPROACH

A. Problem Formulation

We first formulate the CDC operator's profit over a billing cycle as its revenue minus its various expenses:

$$\text{Profit} = \text{Revenue} - \text{Reward} - \text{Wear} - \text{Bill}, \quad (28)$$

where Revenue, Reward, Wear and Bill are defined by (6)-(16). In this paper, we assume that the CDC operator optimizes its decisions at the beginning of each billing cycle with the assumption that tenants demands can be perfectly predicted. We leave online optimization under demand uncertainty as future work. Here, we seek to solve the following problem:

$$\begin{aligned} & \text{Maximize} && (28) \\ & \rho, m_{on}[t], m_{of}[t], \\ & S[t], \phi_i[t], \eta_i[t] \\ & \text{Subject to} && (17) - (23), (25), (27). \end{aligned} \quad (29)$$

Problem (29) is not convex caused by the definition of the reward rate (8) of objective function of (29) and constraint (27) and hence difficult to solve.

Algorithm 1 The Decomposition Method

Require:

- r (Loop times)
 - Solution(r) (Optimal solution of problem (35) at loop r)
 - 1: $r = 1$.
 - 2: $Lb = 0$.
 - 3: **while** 1 **do**
 - 4: Calculate upper bound of ρ via (30)-(32).
 - 5: Obtain a sub-domain of ρ as in (33).
 - 6: Build the corresponding sub-optimization problem (35).
 - 7: Solve problem (35).
 - 8: Save Solution(r).
 - 9: **if** $Ub = \infty$ **then**
 - 10: **return** $\max_r \{\text{Solution}(r)\}$.
 - 11: **end if**
 - 12: $Lb = Ub$.
 - 13: $r = r + 1$.
 - 14: **end while**
-

B. Decomposition Method

Next, we propose a decomposition method to tackle the difficulty in solving (29). Specifically, we decompose problem (29) into several sub-problems, each corresponding to a sub-domain of ρ . The details are presented in Algorithm 1.

We divide problem (29) iteratively as follow:

- **Step 1:** Initializing the lower bound of sub-domain of ρ , i.e., Lb , by letting $Lb = 0$ based on constraint (17).
- **Step 2:** From (5), calculate $\lfloor D_i \rfloor$ in constraint (27) corresponding to Lb , i.e., for each i , set

$$D_i^{Lb} = \left\lceil \max \left(\min \left(\frac{Lb}{\kappa_i} - 1, D_{max} \right), 0 \right) \right\rceil. \quad (30)$$

- **Step 3:** Setting the upper bound of sub-domain of ρ as

$$Ub = \begin{cases} \infty, & \text{If } \vartheta = \emptyset, \\ \min_{i \in \vartheta} \{(D_i^{Lb} + 2)\kappa_i\}, & \text{Otherwise.} \end{cases} \quad (31)$$

Here, the domain ϑ is defined as

$$\vartheta = \{i \in \{1, \dots, I\} | Lb < (D_{max} + 1)\kappa_i\}. \quad (32)$$

- **Step 4:** Obtaining a new sub-domain of ρ as defined by

$$Lb \leq \rho \leq Ub - \epsilon, \quad (33)$$

where ϵ denotes a small positive number approaches 0. Then, replacing $\lfloor D_i \rfloor$ in (27) by constant D_i^{Lb} that can be calculated by (30) and obtain a new affine constraint

$$\sum_{t'=t+1}^{t+D_i^{Lb}} \lambda_i[t'] + \sum_{t'=1}^{t+D_i^{Lb}} (-\phi_i[t'] + \eta_i[t']) \geq 0, \quad \forall i, t. \quad (34)$$

Next, build the corresponding sub-optimization problem

$$\begin{aligned} & \text{Maximize} && (28) \\ & \rho, m_{on}[t], m_{of}[t], && \\ & S[t], \phi_i[t], \eta_i[t] && (35) \\ & \text{Subject to} && (18) - (23), (25), (30) - (34). \end{aligned}$$




Fig. 1. Examples for data traces that we used in this paper: a) workload traces [33] b) energy price of Ameren [34].

TABLE I
SIMULATION CASES

Name	Pricing Policy	Demand Delaying	Storage Capability
UP	UP	No	No
UPS	UP	No	Yes
UPMR	UPMR	Yes	No
UPMRS	UPMR	Yes	Yes

- **Step 5:** Problem (35) is not convex. However, the optimum value of optimization variable ρ in problem (35) is $\rho = Lb$. It is straightforward to see that, substitution of $\rho = Lb$ in the definition of reward rate (8) of objective function of problem (35) gives a convex program that can be solved via convex tools such as CVX [33].
- **Step 6:** If $Ub < \infty$, update Lb by letting $Lb = Ub$ and go to Step 2 to traverse another sub-problem corresponding to another sub-domain of ρ . If $Ub = \infty$, we have traversed all sub-domain of ρ , i.e., all possible sub-optimization problems. Thus, we can find the global optimal solution of the original problem (29) among the optimal solutions of these sub-optimization problems.

Theorem 1: $\lfloor D_i \rfloor = D_i^{Lb}, \forall \rho \in [Lb, Ub - \epsilon], \forall i$.

The proof of Theorem 1 is given in Appendix A. From Theorem 1, in each loop in Algorithm 1, we can replace $\lfloor D_i \rfloor$ in constraint (27) by D_i^{Lb} and build an affine constraint (34) at step 4.

Theorem 2: The optimal value of optimization variable ρ in problem (35) equals Lb .

The proof of Theorem 2 can be found in Appendix B. From Theorem 2, we can substitute ρ in problem (35) by Lb and build a convex program at step 5 in Algorithm 1.

V. CASE STUDY

A. Setup

We consider a CDC serving three types of tenants: type-1 tenants are delay-sensitive (i.e., inelastic) while type-2 and type-3 tenants are elastic. The demands of delay-sensitive users are based on Youtube U.S. traffic from January 1, 2014 to January 31, 2014 [33]. The demands of type 2 and 3 users are constructed based on GMaps and Gmail U.S. traffic from January 1, 2014 to January 31, 2014, respectively [33]. Figure 1(a) shows a snapshot of these workload traces that are used in the simulation. Here, elastic workloads constitute approximately 50% of the overall workloads.

Let the billing cycle be 30 days and $T = 1$ hour. Therefore, $\tau = 30 \times 24 = 720$. Let $\Psi = 10000$, $D_{max} = 24$, $\kappa_1 = \infty$,




Fig. 2. Comparison among UP, UPS, UPMR and UPMRS under peak pricing: a) Normalized profit gain and energy bill of CDC, b) Energy usage of UP and UPMR, c) Energy usage of UPS and UPMRS.




Fig. 3. Comparison among UP, UPS, UPMR and UPMRS under TDP: a) Normalized profit gain and energy bill of CDC, b) Energy usage of UP and UPMR, c) Energy usage of UPS and UPMRS.

$\kappa_2 = 0.1$ and $\kappa_3 = 0.11$. Let $N = 20$, $E_{pue} = 1.2$, $P_{idle} = 0.1$ KW, $P_{peak} = 0.2$ KW, $o_{on} = 0.02$ KWh and $o_{of} = 0.01$ KWh. Let $m[0] = 0$ and $S[0] = 0$ KWh. Let $w_{on} = \$0.003$ and $w_{of} = \$0.002$ as in [35].

In addition, we set the CDC's energy storage capacity C_s , such that it can power up CDC entirely at its peak power consumption up to 30 minutes. Both charge and discharge rates are set to be C_s KW. Let $\zeta = \$0.32$ per KWh, corresponding to a 1Ah lead-acid battery at 12V as mentioned in [36], which costs 2\$ and can stand 1300 recharge cycles at 40% of depth of discharge. In this paper, we evaluate UPMR under four cases, which are illustrated in Table I.

B. Under Peak Pricing

We now evaluate UPMR under a peak pricing scheme. In particular, the CDC's energy cost includes both energy charge and demand charge, where $\alpha[t] = \$0.05207$ per KWh and $\beta_j = \$15.59$ per KW based on electric rates of industrial power service offered by South Carolina Electric and Gas [37].

Figure 2(a) shows the CDC operator's normalized profit gain and energy cost with respect to the maximum profit gain and energy bill among the four cases, respectively. By utilizing energy storage, the profit gain can be enhanced by 1.1%, while the energy bill can be reduced by 2.5% when comparing with UP. Further, by employing UPMR without utilizing the energy storage system, the profit gain can be increased by 4.9% and the energy bill can be reduced by 12.9% when comparing with UP. Thus, UPMR can improve CDC's profit gain and reduce the energy bill more significantly when energy storage is not utilized than when it is. Additionally, UPMR can obtain nearly the same profit gain and energy reduction as UPMRS, indicating that CDCs employing UPMR do not need to utilize energy storage for power management.

Figure 2(b) and 2(c) show the detail of normalized energy usage of UP, UPS, UPMR and UPMRS, with respect to the maximum energy usage. We see that UPMR can effectively smooth the CDC's energy consumption. By comparing UPMR and UPMRS, we see that utilizing energy storage can also further but slightly smooth CDC's energy consumption.

C. Under Time-Dependent Pricing

We now turn to the evaluation of UPMR under TDP. Specifically, we consider that the CDC is charged through TDP based on energy consumption, where the prices are represented by the prices of Ameren real time price from January 1, 2014 to January 31, 2014 [34] as shown in Figure 1(b).

Figure 3 compares the results for UP, UPS, UPMR and UPMRS under profit maximization and TDP. We see that, under TDP, the CDC can obtain less profit gain improvement than the one under electricity rates of peak pricing via demand response. Specifically, CDC has less incentive to perform demand response under TDP. Utilizing the energy storage usually cannot greatly improve the CDC's profit or reduce its energy bill. The reason is that, in most cases, the wear-and-tear cost of battery is higher than the extra profit that can be gained from charging/discharging the energy storage under TDP. Also, under TDP, the CDC can still effectively increase CDC's profit while reducing its energy bill via UPMR.

VI. CONCLUSION

We proposed a simple yet effective pricing policy, called UPMR, which, on top of the commonly-applied usage-based pricing policy, includes a reward component offered to tenants to allow a CDC operator to defer their workloads. For maximizing the CDC operator's profit through UPMR, we

presented an effective and low-complexity decomposition-based algorithm to optimize the reward rate. With real-world workload traces, we evaluated UPMR in terms of maximizing the CDC's profit and reducing its energy cost, against UP (a commonly-used pricing policy in today's CDCs). Our results showed that UPMR can obtain 4.9% more profit gain and 12.9% more energy bill reduction without utilizing the energy storage system.

APPENDIX A PROOF OF THEOREM 1

Considering that $\lfloor D_i \rfloor$ is nondecreasing over ρ , Theorem 1 holds if and only if

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} = D_i^{Lb}, \quad \forall i. \quad (36)$$

We separately prove (36) when $Lb \geq \max_i \{(D_{max} + 1)\kappa_i\}$ and $Lb < \max_i \{(D_{max} + 1)\kappa_i\}$. Firstly, if $Lb \geq \max_i \{(D_{max} + 1)\kappa_i\}$, then from (31), we have

$$Ub = \infty. \quad (37)$$

Next, from (30), we get

$$D_i^{Lb} = D_{max}, \quad \forall i. \quad (38)$$

From (5) and (37), we get

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} = D_{max}, \quad \forall i. \quad (39)$$

Thus, from (38) and (39), (36) holds when $Lb \geq \max_i \{(D_{max} + 1)\kappa_i\}$.

When $Lb < \max_i \{(D_{max} + 1)\kappa_i\}$, we have

$$Ub = \min_{i \in \vartheta} \{(D_i^{Lb} + 2)\kappa_i\}. \quad (40)$$

In this case, we prove the correctness of (36) by four steps:

Step 1: For all $i \in \vartheta$, from (40), we have

$$Ub \leq (D_i^{Lb} + 2)\kappa_i, \quad \forall i \in \vartheta. \quad (41)$$

Bring (41) into (5), we have

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} \leq \lfloor \max(\min(D_i^{Lb} + 1 - \epsilon, D_{max}), 0) \rfloor. \quad (42)$$

From (30), we ensure that $D_i^{Lb} \geq 0$. In this case, $\min(D_i^{Lb} + 1 - \epsilon, D_{max}) \geq 0$. Combine it with (42), we get

$$\begin{aligned} \lfloor D_i \rfloor |_{\rho=Ub-\epsilon} &\leq \lfloor \min(D_i^{Lb} + 1 - \epsilon, D_{max}) \rfloor \\ &\leq \lfloor D_i^{Lb} + 1 - \epsilon \rfloor. \end{aligned} \quad (43)$$

Next, from (30), we ensure that D_i^{Lb} is integer. In this case, $\lfloor D_i^{Lb} + 1 - \epsilon \rfloor = D_i^{Lb}$. Combine it with (43), we ensure

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} \leq D_i^{Lb}, \quad \forall i \in \vartheta. \quad (44)$$

Step 2: For all $i \notin \vartheta$, from (30) and (32), we have $D_i^{Lb} = D_{max}$, $\forall i \notin \vartheta$. From (5), we have $\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} \leq D_{max}$, $\forall i$. Therefore,

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} \leq D_i^{Lb}, \quad \forall i \notin \vartheta. \quad (45)$$

Combine (44) and (45), we ensure

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} \leq D_i^{Lb}, \quad \forall i. \quad (46)$$

Step 3: Let

$$i' = \arg \min_{i \in \vartheta} \{(D_i^{Lb} + 2)\kappa_i\}. \quad (47)$$

From (40), we have

$$Ub = (D_{i'}^{Lb} + 2)\kappa_{i'}. \quad (48)$$

Since $i' \in \vartheta$, from (32), we have $Lb \in [0, (D_{max} + 1)\kappa_{i'}]$. If $Lb \in [0, 2\kappa_{i'}]$, from (30), $D_{i'}^{Lb} = 0$. Combine it with (48), we have

$$Ub = 2\kappa_{i'} > Lb. \quad (49)$$

If $Lb \in [2\kappa_{i'}, (D_{max} + 1)\kappa_{i'}]$, from (30), $D_{i'}^{Lb} = \lfloor \frac{Lb}{\kappa_{i'}} - 1 \rfloor$. Combine it with (48), we have

$$\begin{aligned} Ub &= \left(\left\lfloor \frac{Lb}{\kappa_{i'}} - 1 \right\rfloor + 2 \right) \kappa_{i'} \\ &> \left(\left(\frac{Lb}{\kappa_{i'}} - 1 \right) + 1 \right) \kappa_{i'} \\ &> Lb. \end{aligned} \quad (50)$$

Combining (49) and (50), we ensure that $Ub > Lb$. Meanwhile, since $\lfloor D_i \rfloor$ is nondecreasing over ρ , we have

$$\lfloor D_i \rfloor |_{\rho=Ub-\epsilon} \geq D_i^{Lb}, \quad \forall i. \quad (51)$$

Step 4: From (46) and (51), (36) always holds when $Lb < \max_i \{(D_{max} + 1)\kappa_i\}$, too.

APPENDIX B PROOF OF THEOREM 2

We notice that, the only part of the objective function of problem (35) that involves ρ is the definition of the reward rate (8). Meanwhile, (8) involves no variable other than ρ . Also, the only constraint of problem (35) that involves ρ is (33), which involves no other constraint except ρ . Therefore, the problem (35) can be decomposed in two separate optimization problems, among which the first one is to minimize the Reward term defined by (7) and (8), subject to the constraints (30)-(33). The optimal value of ρ , in this decomposed optimization problem is $\rho = Lb$, since the Reward term defined by (7) and (8) is a non-decreasing function of ρ .

REFERENCES

- [1] A. Venkatraman, "Global census shows datacentre power demand grew 63% in 2012," *Computer Weekly*, Oct 2012.
- [2] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: Research problems in data center networks," *SIGCOMM Comput. Commun. Rev.*, vol. 39, pp. 68–73, Jan 2009.
- [3] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Trans. on Smart Grid*, vol. 1, pp. 320–331, Dec 2010.
- [4] C. Wang, B. Urgaonkar, G. Kesidis, U. V. Shanbhag, and Q. Wang, "A case for virtualizing the electric utility in cloud data centers," in *Proc. of USENIX HotCloud*, Philadelphia, PA, Jun 2014.
- [5] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *Proc. of IEEE IGSC*, Dallas, TX, Nov 2014.
- [6] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," *Performance Evaluation*, vol. 70, pp. 770–791, Oct 2013.
- [7] M. Lin, A. Wierman, L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," in *Proc. of IEEE INFOCOM*, Shanghai, China, Apr 2011.

- [8] K. Le, J. Zhang, J. Meng, R. Bianchini, Y. Jaluria, and T. Nguyen, "Reducing electricity cost through virtual machine placement in high performance computing clouds," in *Proc. of IEEE SC*, Seattle, WA, Nov 2011.
- [9] C. Wang, B. Urgaonkar, Q. Wang, and G. Kesidis, "A hierarchical demand response framework for data center power cost optimization under real-world electricity pricing," in *Proc. of IEEE MASCOTS*, Paris, France, Sep 2014.
- [10] S. Ren, Y. He, and F. Xu, "Provably-efficient job scheduling for energy and fairness in geographically distributed data centers," in *Proc. of IEEE ICDCS*, Macau, China, Jun 2012.
- [11] H. Xu and B. Li, "Reducing electricity demand charge for data centers with partial execution," in *Proc. of ACM e-Energy*, Cambridge, UK, Jun 2014.
- [12] A. Li, X. Yang, S. Kandula, and M. Zhang, "Cloudcmp: comparing public cloud providers," in *Proc. of ACM IMC*, Melbourne, Australia, Nov 2010.
- [13] V. Ishakian, R. Sweha, A. Bestavros, and J. Appavoo, "Cloudpack* exploiting workload flexibility through rational pricing," in *Proc. of Springer Middleware*, Montreal, Canada, Dec 2012.
- [14] Y. Zhan, M. Ghamkhari, D. Xu, and H. Mohsenian-Rad, "Propagating electricity bill onto cloud tenants: Using a novel pricing mechanism," in *Proc. of IEEE GLOBECOM*, San Diego, CA, Dec 2015.
- [15] N. Nasiriani, C. Wang, G. Kesidis, B. Urgaonkar, L. Y. Chen, and R. Birke, "On fair attribution of costs under peak-based pricing to cloud tenants," in *Proc. of IEEE MASCOTS*, Atlanta, GA, Sep 2015.
- [16] B. Guenter, N. Jain, and C. Williams, "Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning," in *Proc. of IEEE INFOCOM*, Shanghai, China, Apr 2011.
- [17] A. Gandhi, M. Harchol-Balter, R. Raghunathan, and M. A. Kozuch, "Autoscale: Dynamic, robust capacity management for multi-tier data centers," *ACM Trans. on Computer Systems*, vol. 30, pp. 1–26, Nov 2012.
- [18] D. Meisner, C. M. Sadler, L. A. Barroso, W.-D. Weber, and T. F. Wenisch, "Power management of online data-intensive services," in *Proc. of ACM ISCA*, San Jose, CA, Jun 2011.
- [19] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters," in *Proc. of ACM ASPLOS XVII*, London, UK, Mar 2012.
- [20] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam, "Optimal power cost management using stored energy in data centers," in *Proc. of ACM SIGMETRICS*, San Jose, CA, Jun 2011.
- [21] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?," in *Proc. of ACM SIGMETRICS*, London, UK, Jun 2012.
- [22] Q. Tang, S. Gupta, and G. Varsamopoulos, "Thermal-aware task scheduling for data centers through minimizing heat recirculation," in *Proc. of IEEE CLUSTER*, Austin, TX, Sep 2007.
- [23] H. Mohsenian-Rad and A. Leon-Garcia, "Energy-information transmission tradeoff in green cloud computing," in *Proc. of IEEE GLOBECOM*, Miami, FL, Dec 2010.
- [24] Z. Liu, M. Lin, A. Wierman, S. Low, and L. Andrew, "Greening geographical load balancing," *IEEE/ACM Trans. on Networking*, vol. 23, pp. 657–671, Apr 2015.
- [25] L. Rao, X. Liu, L. Xie, and W. Liu, "Minimizing electricity cost: Optimization of distributed internet data centers in a multi-electricity-market environment," in *Proc. of IEEE INFOCOM*, San Diego, CA, Mar 2010.
- [26] M. Ghamkhari and H. Mohsenian-Rad, "Energy and performance management of green data centers: A profit maximization approach," *IEEE Trans. on Smart Grid*, vol. 4, pp. 1017–1025, Jun 2013.
- [27] M. Ghamkhari and H. Mohsenian-Rad, "Optimal integration of renewable energy resources in data centers with behind-the-meter renewable generators," in *Proc. of IEEE ICC*, Ottawa, Canada, June 2012.
- [28] Y. Li, D. Chiu, C. Liu, L. Phan, T. Gill, S. Aggarwal, Z. Zhang, B. T. Loo, D. Maier, and B. McManus, "Towards dynamic pricing-based collaborative optimizations for green data centers," in *Proc. of IEEE ICDEW*, Brisbane, QLD, Apr 2013.
- [29] S. Ha, S. Sen, C. Joe-Wong, Y. Im, and M. Chiang, "Tube: time-dependent pricing for mobile data," in *Proc. of ACM SIGCOMM*, Helsinki, Finland, Aug 2012.
- [30] A. Oudalov, R. Cherkaoui, and A. Beguin, "Sizing and optimal operation of battery energy storage system for peak shaving application," in *Proc. of IEEE Powertech*, Lausanne, Switzerland, Jul 2007.
- [31] <http://www.pge.com/>.
- [32] J. Li, Z. Bao, and Z. Li, "Modeling demand response capability by internet data centers processing batch computing jobs," *IEEE Trans. on Smart Grid*, vol. 6, pp. 737–747, Mar 2015.
- [33] "Browse real-time traffic to google products and services." <http://www.google.com/transparencyreport/traffic/explorer>.
- [34] "Real time prices - ameren." <https://www2.ameren.com/RetailEnergy/RealTimePrices>.
- [35] H. Qian and D. Medhi, "Server operational cost optimization for cloud computing service providers over a time horizon," in *Proc. of USENIX Hot-ICE*, Boston, MA, Mar 2011.
- [36] V. Kontorinis, J. Sampson, L. E. Zhang, B. Aksanli, H. Homayoun, T. S. Rosing, and D. M. Tullsen, "Battery provisioning and associated costs for data center power capping," tech. rep., UCSD, 2012.
- [37] "SCE&G electric rates." <https://www.sceg.com/for-my-business/manage-my-service/rates>.