Online Electricity Cost Saving Algorithms for Co-Location Data Centers

Linquan Zhang, Zongpeng Li, Chuan Wu, Shaolei Ren

***Abstract*—This work studies the online electricity cost minimization problem at a co-location data center, which serves multiple tenants who rent the physical infrastructure within the data center to run their respective cloud computing services. The co-location operator has no direct control over power consumption of its tenants, and an efficient mechanism is desired for eliciting desirable consumption patterns from the tenants. Electricity billing faced by a data center is nowadays based on both the total volume consumed, and the peak consumption rate. This leads to an interesting new combinatorial optimization structure on the electricity cost optimization problem, which also exhibits an online nature due to the definition of peak consumption. We model and solve the problem through two approaches: the *pricing* approach and the *auction* approach, and design online algorithms with small competitive ratios.**

I. Introduction

Co-location data centers (or *co-locations*) rent physical space and infrastructure support, *e.g.*, reliable power supply and cooling service, to multiple tenants for hosting their servers at a common site. They are rather different from private (owner-operated) data centers in which operators have full control of computing resources and site facilities. Co-locations offer a flexible data center solution to small and medium users who wish to run their own 'cloud' but are otherwise deterred by the daunting cost of constructing and maintaining their own data center. Even large users like Google and Akamai rely on co-locations as a cost-effective complement to their own data centers for achieving a global presence, particularly in regions of relatively low demand that do not justify a dedicated data center. The surging market of co-locations is expected to reach \$43 billion by 2018, with a compound annual growth rate of 11% [1]. The recent paradigm of modular datacenter design further facilitates the management and scalability of co-locations — accommodating a new tenant could be as simple as the plug-and-play of a few extra pods.

Electricity charges paid by a co-location is computed by an interesting formula that includes two components: i) the *peak charge*, determined by the peak demand within a billing cycle, *e.g.*, the maximum average power consumption measured over each 15-minute interval; ii) the *volume charge*, based on total energy consumption in the billing cycle [2], [3], [4]. Tab. I summarizes several electricity rate plans available in North America. The volume charge is relatively intuitive. The rationale behind the peak charge is that peak consumption shedding is critical to a power grid; even a small reduction in peak demand can provide significant cost savings and reduce greenhouse gas emission since a smaller number of peaking power plants, which usually burn natural gas or diesel oil and run only in peak hours with higher cost, are needed.

TABLE I
Summer Electricity Rates in North America

Utility Provider	Contract Type	Peak Charge (\$/kW)	Volume Charge (¢/kWh)
PG&E [5]	E20 Primary Firm	17.54	15.093
Duke Energy [6]	LGS-TOU-28	16.99	4.724
BC Hydro [2]	LGS	9.95	4.86
Georgia Power [7]	Power&Light PLH-8	16.90	0.5679
Mid American Energy [8]	Rate GD, Iowa	7.07	5.433

In practice, the peak charge component is seen to account for over 30% of the total electricity bill [3], [4]. For example, consider a co-location data center located in British Columbia, Canada, powered by BC Hydro with 24 MW peak demand and 15 MW average demand. The monthly peak charge is \$238,800, while the volume charge is \$524,880. In this case, the peak charge is 31% of the total monthly payment. The peak charge represents even higher portions of the bill in Georgia, where the monthly peak charge would be \$405,600 while the volume charge is \$61,333.2, under the same consumption pattern as in the British Columbia case. This suggests that *a well designed algorithm for shaping the power consumption profile and controlling the peak demand has a great potential in helping cut electricity cost at a co-location*.

However, different from the case of private data centers, the co-location operator has no direct control on which machines are on/off, since its role is to offer basic services such as stable power supply and cooling. The individual tenants at a co-location manage their own servers and control the corresponding power consumption. While the co-location has a strong incentive to cut peak consumption and therefore save cost, its tenants may or may not share that same interest, depending on the contract between the two sides.

Typical electricity pricing today between a co-location and its tenants is flat-rate based, and does not depend on the real consumption volume or pattern [9], [10]. The tenants have little incentive to reduce their electricity usage by shutting down under-utilized servers, or by modulating their consumption pattern via shifting computing jobs in the temporal domain to reduce peak consumption rate. Such actions desired by the co-location will not automatically happen without appropriate in-

Manuscript received on March 29, 2015; revised on July 15, 2015. Some preliminary results appear as a poster paper at ACM SIGMETRICS 2015.

L. Zhang and Z. Li are with the Department of Computer Science, University of Calgary. Email: {linqzhan, zongpeng}@ucalgary.ca. Z. Li is also with School of Computer, Wuhan University. C. Wu is with the Department of Computer Science, The University of Hong Kong. Email: cwu@cs.hku.hk. S. Ren is with the Department of Electrical and Computer Engineering, University of California, Riverside. Email: sren@ece.ucr.edu.

This work has been supported by AITF, NSERC (10006298), HK RGC (17204715 and HKU 718513), and U.S. NSF (CNS-1423137 and CNS-1551661).

centives. Co-locations are sometimes so desperate to cut peak consumption that they start their stand-by generators to cover part of its tenants' demand [4]. Such quick-start generation (*e.g.*, using diesel generators) is often not economical, nor is it environment friendly. A simple solution is to bill each tenant based on its peak charge and volume charge. Unfortunately such an approach may not help reduce the peak energy cost since many of the peaks do not coincide with the overall peak of the entire co-location. On the contrary, auction-based demand response mechanisms have the potential of efficiently providing incentive for tenants to cooperate, eliciting desired electricity consumption patterns with remuneration paid in return. A well designed auction may represent a win-win solution for both the co-location operator and its tenants.

The maximum power demand is dependent on the power consumption in all time intervals during a billing cycle. Decisions in different time slots are therefore coupled, leading to an inherent online nature of the problem of demand shaping. Even for the offline version of the problem (future prices and demands are perfectly known), computing the optimal solution efficiently is still highly non-trivial, since its underlying optimization problem is an integer program, which is NP-hard in general. The challenge escalates when one seeks to design an *online solution* for practical application, as knowledge on the demands, unit power prices as well as tenants' bids in the future are completely unknown.

Tenants could reduce energy consumption in various manners, *e.g.* shutting down under-utilized servers and shifting delay insensitive workloads in the temporal domain. The specific choice of a tenant's energy reduction is not our focus. We focus on designing mechanisms for eliciting desirable consumption patterns from the tenants. We model and solve the electricity cost saving problem in a co-location data center through two approaches: the *pricing* approach and the *auction* approach.

In the pricing approach, the co-location data center offers a price it is willing to pay for unit energy reduction by tenants, and the tenants decide and submit how much energy they are willing to save at that price. Finally the co-location data center determines which tenants' energy reductions are accepted. **In the auction approach**, the co-location data center invites the tenants to submit energy reduction bids including the amount of energy consumption to shed and the amount of remuneration asked. The co-location then conducts a reverse auction to determine winning bids along with their corresponding payments. The pricing approach is relatively simple and has been applied in real world demand response solutions, implemented in electric appliances on the market [11]. It is simple but requires the co-location to first come up with a good estimate on a unit reduction offer. The auction approach eliminates the need of such ad hoc guesses and resorts to the power of the market instead for automatic fair price revelation based on demand and supply. Yet, the auction design is more complex than the algorithm design in the pricing approach, and our solution to the former borrows techniques from the latter.

In each approach, we design efficient online algorithms with small competitive ratios. For the auction approach, we further apply a randomized auction design framework, which decomposes a fractional optimal solution to the second sub-problem into a convex combination of feasible integer solutions, to ensure truthfulness. Trace-driven simulation studies further verify the efficacy of the proposed algorithms, showing their close-to-optimum performance that is better than the worst-case bounds.

The rest of the paper is organized as follows. Sec. II reviews related work. The system model and the problem formulation are presented in Sec. III. We then study the pricing version and the auction version of the cost minimization problem in Sec. IV and Sec. V, respectively. Trace-driven simulation studies are presented in Sec. VI. Sec. VII concludes the paper.

II. Related Work

Energy efficiency in data centers attracted high attention from both academia and industry in recent years, since a small fraction of improvement in energy efficiency transfer into millions of dollars of cost savings. Several algorithms [12], [13] are proposed to ensure that power consumption is proportional to total workload, by dynamically turning on/off machines in data centers (dynamic capacity provisioning) or adjusting their speed (CPU speed scaling). Different from traditional data centers, co-location data centers are not able to directly control the servers' on/off status or speed.

A number of recent studies are devoted to demand response in data centers. Wang *et al.* [14] advocate that data center demand response can be an effective approach to improve power grid stability, to reduce the energy consumption as well as to increase the revenue of data centers. Liu *et al.* [15] propose predication-based pricing for demand response in data centers, where the data center operator has full control over all its servers. Again, co-locations have no direct control over the servers' power consumption, thus these methods are not directly applicable.

The peak-based electricity charge model is applied in the real-world for large business users as exemplified by data centers. Wang *et al.* [16] propose offline and online algorithms to minimize the electricity cost by delaying or dropping workloads. They prove a theoretical bound for the dropping only case. However, they assume that the dropping cost is proportional to the size of the workload, which ignores the fact that different jobs of the same size may be of different importance. Xu *et al.* [17] study the electricity cost minimization problem in a data center through partial execution while trying to satisfy the Service Level Agreement (SLA). Different from their work, we tackle the problem by eliciting voluntary energy reduction from the co-location tenants in an otherwise uncoordinated power consumption environment within a co-location data center. Bar-Noy *et al.* [18] advocate shaving the peak demand by using stored energy. Their solution focuses on optimizing the peak charge only without considering the volume charge and energy storage cost, while our work concentrates on the overall cost including the peak charge, volume charge as well as tenant costs.

The primal-dual approach [19] is a general tool for designing competitive online algorithms. A well designed primal-dual online algorithm sometimes achieves a small competitive ratio against the offline optimum. However, the method requires that the underlying problem is of a packing type or a covering type. The peak power demand charge model essentially makes our problem a mixed packing and covering problem. Directly applying the primal-dual approach is hence infeasible. Azar *et al.* [20] recently propose a technique to tackle mixed packing and covering problems. Yet it assumes that the packing constraints are given at the beginning while the covering constraints are revealed on the fly, and the problem only focuses on minimizing the maximum amount by which a packing constraint is violated. So their technique still has restrictions and is not directly applicable to our problem.

III. The Co-location Datacenter Model

We consider a co-location data center hosting a large number (thousands) of servers for its tenants. The co-location pays the utility company both a peak charge and a volume charge for electricity consumed by the data center. The system runs in a time-slotted fashion. The length of each time slot is τ, *e.g.*, 15 minutes [3]. The price of a unit amount of electricity fluctuates over time. Let f_t be the price at t. After a billing cycle, the peak demand among all time slots is identified, and used to compute the peak charge. Let f_{peak} be the peak demand price, known by the co-location at the beginning of the billing cycle. The co-location can save energy used towards server cooling when its tenants save energy by shutting down their servers. Let λ be the partial Power Usage Effectiveness (pPUE), which is the ratio between (a) the total energy consumption for IT and cooling and (b) the IT energy consumption. The pPUE varies as the ambient temperature varies, and a typical value ranges from 1.1 to 2.0, as shown in Tab. II.

TABLE II
pPUE v.s. Ambient Temperature [21]

Outdoor Ambient (°F)	Return Air DB (°F)	Cooling Mode	pPUE
90	85	Compressor	1.31
70	85	Compressor	1.21
60	85	Mixed	1.17
50	85	Pump	1.10
25	85	Pump	1.05

Following empirical models from recent measurement studies [21], [22], we compute the pPUE using quadratic curve fitting based on data in Tab. II.

$$\lambda = 3.0825 \times 10^{-5}\theta^2 + 5.7154 \times 10^{-4}\theta + 1.0127 \quad (1)$$

where the formula is valid on the interval $\theta \in [25, 90]$, and θ is the ambient temperature. Since θ fluctuates over time, so does the pPUE. We use λ_t to denote the pPUE at time t.

Let D_t be the power demand at time t without energy reduction by tenants. At the end of slot $t-1$, D_t can be accurately predicted, and f_t is known. We assume that $\max_\tau D_\tau/D_t \leq \xi, \forall t$. Typically ξ ranges from $1.2 \sim 2$ since idle machines still consume around 50% of energy [23].

Diesel generators can also help shed the co-location's peak consumption from the power grid, yet they represent a less preferred option due to environment concerns [4] and are not considered in this work. As shown in Fig. 1, the system has two alternative ways to call for energy reduction from tenants for the next time slot, i) offering a p_t, determined by the co-location, as the unit price for energy reduction from tenants; ii) conducting a reverse auction by soliciting energy reduction bids from tenants. We assume that $p_t/\lambda_t \leq \kappa f_t$, *i.e.*, the co-location will not offer too high prices for the energy reduction. Let $\mathcal{M}_t$ be the set of tenants willing to join the process in i) at t. Each tenant $i \in \mathcal{M}_t$ submits $r_{i,t}$ to the co-location, indicating how many kilowatts it is willing to save at t. Let $\mathcal{N}_t$ be the set of tenants willing to participate in the power reduction reverse auction in ii) at t. Each tenant $j \in \mathcal{N}_t$ submits $(r_{j,t}, b_{j,t})$ to the co-location, where $r_{j,t}$ is its possible energy reduction at t, and $b_{j,t}$ is the remuneration asked for. We assume that $b_{j,t}/(\lambda_t r_{j,t}) \geq f_t, p_t/\lambda_t \geq f_t$, *i.e.*, the unit price for the energy reduction from tenants is higher than price of energy from the grid. The total energy reduction contributed by co-location tenants is a fraction of D_t and cannot exceed D_t, *i.e.*, $\lambda_t \sum_i r_{i,t} \leq \frac{1}{\rho} D_t, \forall t$. A typical ρ ranges from 2 to 5.

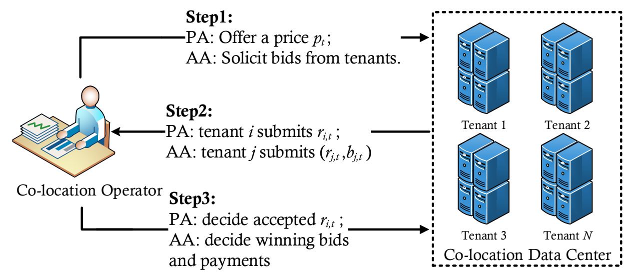

Fig. 1. An illustration of the two approaches for reducing energy consumption in a co-location data center. PA = the Pricing Approach, AA = the Auction Approach.

Let H_t be a continuous variable indicating the amount of energy purchased from the grid at time t. For each $r_{i,t}$ given by the tenant $i \in \mathcal{M}_t$, a binary variable $x_{i,t}$ indicates whether the co-location accepts $r_{i,t}$. For each bid $(r_{j,t}, b_{j,t})$ submitted by tenant $j \in \mathcal{N}_t$, we have a binary variable $x_{j,t}$ indicating whether this bid wins. A notation table is provided for ease of reference.

$\mathcal{M}_t$	the set of tenants participating in the pricing approach at t.
$\mathcal{N}_t$	the set of tenants participating in the auction at t.
$r_{i,t}$	the energy reduction at t for tenant i.
$b_{i,t}$	the bidding price at t for tenant i.
p_t	the unit price of energy reduction at t.
f_t	the unit cost of power consumption at t.
f_{peak}	the peak demand price.
D_t	the original power demand without energy reduction by tenants enabled at t.
H_t	power supplied by the grid at t.
$x_{j,t}$	binary variable: bid $(r_{j,t}, b_{j,t})$ wins or not.
λ_t	partial Power Usage Effectiveness (pPUE).
ξ	$\max_t(\max_\tau D_\tau/D_t)$.
κ	$\max_t\{p_t/(\lambda_t f_t)\}$.
ρ	$\min_t\{D_t/(\lambda_t \sum_i r_{i,t})\}$.

A. The Pricing Approach Model

In the pricing approach, we have the following constraint for meeting the power demand after considering energy reduction by tenants: $H_t \geq D_t - \lambda_t(\sum_i r_{i,t} x_{i,t}), \forall t$.

We focus on the total cost of the co-location: $\sum_t f_t H_t + f_{peak} \max_t H_t + \sum_{t,i} p_t r_{i,t} x_{i,t}$, where $\sum_t f_t H_t$ is the volume charge, $f_{peak} \max_t H_t$ is the peak charge, and $\sum_{t,i} p_t r_{i,t} x_{i,t}$ is the cost of paying energy reduction from the tenants. The cost minimization problem is then formulated as follows:

$$\text{minimize} \quad \sum_t f_t H_t + f_{peak} \max_t H_t + \sum_{t,i} p_t r_{i,t} x_{i,t} \tag{2}$$

$$\text{subject to:} \quad H_t + \lambda_t \sum_{i \in \mathcal{M}_t} r_{i,t} x_{i,t} \geq D_t \qquad \forall t \tag{2a}$$

$$x_{i,t} \in \{0,1\}, H_t \geq 0 \quad \forall i \in \mathcal{M}_t, t \tag{2b}$$

B. The Auction Approach Model

In the auction approach, the following constraint ensures that the power demand is met after considering the energy reduction by tenants: $H_t \geq D_t - \lambda_t (\sum_{j \in \mathcal{N}_t} r_{j,t} x_{j,t}), \forall t$. We focus on the overall social cost of the co-location and its tenants: $\sum_t f_t H_t + f_{peak} \max_t H_t + \sum_{t,j} b_{j,t} x_{j,t}$. Here $\sum_t f_t H_t$ is the volume charge, $f_{peak} \max_t H_t$ is the peak charge and $\sum_{t,i} b_{i,t} x_{i,t}$ is the cost of the tenants. Payments between the co-location and its tenants cancel themselves and are not reflected in the social cost. The social cost minimization problem can then be formulated as follows:

$$\text{minimize} \quad \sum_t f_t H_t + f_{peak} \max_t H_t + \sum_{t,j} b_{j,t} x_{j,t} \tag{3}$$

$$\text{subject to:} \quad H_t + \lambda_t \sum_{j \in \mathcal{N}_t} r_{j,t} x_{j,t} \geq D_t \qquad \forall t \tag{3a}$$

$$x_{j,t} \in \{0,1\}, H_t \geq 0 \quad \forall j \in \mathcal{N}, t \tag{3b}$$

IV. The Pricing Approach

Let H_{max} be the maximum of H_t over time, *i.e.*, $H_t \leq H_{max}, \forall t$. The optimization problem (2) then becomes a linear integer program (LIP):

$$\text{minimize} \quad \sum_t f_t H_t + f_{peak} H_{max} + \sum_{t,i} p_t r_{i,t} x_{i,t} \tag{4}$$

$$\text{subject to:} \quad H_t + \lambda_t \sum_{i \in \mathcal{M}_t} r_{i,t} x_{i,t} \geq D_t \qquad \forall t \tag{4a}$$

$$H_t \leq H_{max}, \qquad \forall t \tag{4b}$$

$$x_{i,t} \in \{0,1\}, H_t \geq 0 \quad \forall i \in \mathcal{M}_t, t \tag{4c}$$

LIPs are NP-hard in general. Even if complete future information on $\boldsymbol{D}, \boldsymbol{f}, \boldsymbol{p}, \boldsymbol{\lambda}$ and $\boldsymbol{r}$ are perfectly known, solving (4) optimally is still computationally challenging.

A. An Offline Approximation Algorithm

We first design an offline approximation algorithm, assuming all future information are known. We introduce a linear programming (LP) relaxation by relaxing the integer constraint $x_{i,t} \in \{0,1\}$ to $x_{i,t} \geq 0$. The main challenge in such offline algorithm design is to effectively modulate the power consumption over time, for controlling the peak charge during the billing cycle. We start analyzing such peak charge as follows:

$$\begin{aligned}
&\min_{\boldsymbol{x}, H_{max}, \boldsymbol{H}} \Big\{ \sum_t f_t H_t + f_{peak} H_{max} + \sum_{t,i} p_t r_{i,t} x_{i,t} \Big\} \\
=& \min_{H_{max}, \boldsymbol{H}} \Big\{ \sum_t \big(f_t H_t + p_t (D_t - H_t)/\lambda_t \big) + f_{peak} H_{max} \Big\} \\
=& \sum_t p_t D_t / \lambda_t + \min_{H_{max}, \boldsymbol{H}} \Big\{ \sum_t \big(f_t - p_t/\lambda_t \big) H_t + f_{peak} H_{max} \Big\}
\end{aligned}$$

Note that H_{max} is defined as the maximum $H_t, \forall t$. If $D_t \geq H_{max}$, then $H_t = H_{max}$. If $D_t < H_{max}$, then $H_t = D_t$.

$$\begin{aligned}
&\min_{H_{max}, \boldsymbol{H}} \Big\{ \sum_t \big(f_t - p_t/\lambda_t \big) H_t + f_{peak} H_{max} \Big\} \\
=& \min_{H_{max}} \Big\{ \sum_t \big(f_t - p_t/\lambda_t \big) (D_t - (D_t - H_{max}) \mathbf{1}_{D_t \geq H_{max}}) \\
&+ f_{peak} H_{max} \Big\} \\
=& \sum_t (f_t - p_t/\lambda_t) D_t + \min_{H_{max}} \Big\{ f_{peak} H_{max} \\
&+ \sum_t \big(p_t/\lambda_t - f_t \big) (D_t - H_{max}) \mathbf{1}_{D_t > H_{max}} \big) \Big\} \\
=& \sum_t (f_t - p_t/\lambda_t) D_t + \Phi
\end{aligned}$$

where $\Phi \triangleq \min_{H_{max}} \Big\{ f_{peak} H_{max} + \sum_t \big(p_t/\lambda_t - f_t \big) (D_t - H_{max}) \mathbf{1}_{D_t > H_{max}} \big) \Big\}$, and $\mathbf{1}_{D_t > H_{max}} = 1$ if $D_t \geq H_{max}$, and 0 otherwise. In order to minimize the total cost, we sort all D_ts in descending order. Let $D_{t_1} \geq D_{t_2} \cdots \geq D_{t_T}$. We analyze the total cost based on different H_{max} as follows.

If $H_{max} = D_{t_1}$, then $\Phi = f_{peak} D_{t_1}$.

If $D_{t_1} > H_{max} \geq D_{t_2}$, then $\Phi = f_{peak} H_{max} + \sum_{k=1}^{1} (p_{t_k}/\lambda_{t_k} - f_{t_k})(D_{t_k} - H_{max}) = \sum_{k=1}^{1} (p_{t_k}/\lambda_{t_k} - f_{t_k}) D_{t_k} + (f_{peak} - \sum_{k=1}^{1} (p_{t_k}/\lambda_{t_k} - f_{t_k})) H_{max}$.

If $D_{t_\tau} > H_{max} \geq D_{t_{\tau+1}}$, then $\Phi = f_{peak} H_{max} + \sum_{k=1}^{\tau} (p_{t_k}/\lambda_{t_k} - f_{t_k})(D_{t_k} - H_{max}) = \sum_{k=1}^{\tau} (p_{t_k}/\lambda_{t_k} - f_{t_k}) D_{t_k} + (f_{peak} - \sum_{k=1}^{\tau} (p_{t_k}/\lambda_{t_k} - f_{t_k})) H_{max}$.

We illustrate the relation between H_{max} and Φ in Fig. 2, and observe that Φ reaches its minimum when $\exists \tilde{\tau}$, s.t. $f_{peak} = \sum_{k=1}^{\tilde{\tau}} (p_{t_k}/\lambda_{t_k} - f_{t_k})$, which implies that $H_{max} \in (D_{\tilde{\tau}}, D_{\tilde{\tau}+1}]$.

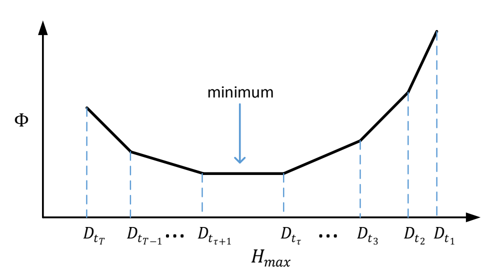

Fig. 2. The relation between H_{max} and Φ.

Energy reduction offered by co-location tenants may be limited, the optimal H^*_{max} is

$$H^*_{max} = \begin{cases} D_{\tilde{\tau}}, & \text{if } D_{\tilde{\tau}} > \tilde{H} \\ \tilde{H}, & otherwise \end{cases} \tag{5}$$

where $\tilde{H} = \max_t \{D_t - \lambda_t \sum_i r_{i,t}\}$.

Based on the analysis above, we design an approximation algorithm for LIP (4). At each t, the total amount of energy reduction by the tenants is determined once H^*_{max} is decided. Let H_t be the energy purchased from the grid at time t. We have a sub-problem for each time slot t:

$$\text{minimize} \quad p_t \sum_i r_{i,t} x_{i,t} \tag{6}$$

$$\text{subject to:} \quad \lambda_t \sum_{i \in \mathcal{M}_t} r_{i,t} x_{i,t} \geq D_t - H_t \tag{6a}$$

$$x_{i,t} \in \{0,1\}, \quad \forall i \in \mathcal{M}_t \tag{6b}$$

Problem (6) is a simplified version of (10), which we will discuss in Sec. V later. Alg. 5, to be presented later in Sec. V, is a polynomial approximation algorithm to the problem (6), verifying an integrality gap of 2, as proven in Theorem 5 later. By employing Alg. 5, we present our offline approximation algorithm as shown in Alg. 1.

Theorem 1. *Alg. 1 is a polynomial-time* 2*-approximation algorithm to (4).*

Proof: since the approximation algorithm uses the same H^*_{max} and H_t as the offline optimal solution does, we have

$$\begin{aligned}
&\sum_t f_t H_t + f_{peak} H_{max} + \sum_{i,t} p_t r_{i,t} x_{i,t} \\
\leq &\sum_t f_t H^*_t + f_{peak} H^*_{max} + 2 \sum_{i,t} p_t r_{i,t} x^*_{i,t} \\
= &2OPT_{LPR} \leq 2OPT_{LIP} \quad \square
\end{aligned}$$

Algorithm 1 An Offline Approximation Algorithm for the Pricing Approach

$x_{i,t} = 0, \forall i, t$; $H_t = 0, \forall t$;

Sort all D_ts in descending order;
Find the optimal H^*_{max} according to (5)
for all $t \in [1, T]$ **do**
  Determine H_t,

$$H_t = \begin{cases} D_t, & \text{if } H^*_{max} > D_t \\ H^*_{max}, & otherwise \end{cases}$$

  Solving the sub-problem (6) using Alg. 5 when energy reduction target is $D_t - H_t$;
end for

B. An Online Algorithm

An online algorithm works with hitherto information, and cannot access information that is available only in the future. New information comes on the fly and decisions for each time slot have to be made immediately without delay. Competitive analysis is employed to analyze the performance of an online algorithm compared with the offline optimum, to which full future information is available *a priori*. An online algorithm A for a minimization problem is γ-competitive if $A(I) \leq \gamma OPT(I), \forall I$, where I is an input and $OPT(I)$ is the offline optimum.

Drawing experiences from the offline case, we design an online algorithm in Alg. 2. The idea of Alg. 2 is that upon receiving a new request, it calculates and updates the optimal H'_{max} based on all information received so far. Then the amount of energy that the co-location data center needs to draw from the grid is determined as well.

Algorithm 2 An Online Algorithm for the Pricing Approach

$x_{i,t} = 0, \forall i, t$; $H_t = 0, \forall t$; $H_{max} = 0$;

for all $t \in [1, T]$ **do**
  Sort all D_ts received so far in descending order;
  Find $\tilde{\tau}$, s.t. $f_{peak} = \sum_{k=1}^{\tilde{\tau}} (p_{t_k}/\lambda_{t_k} - f_{t_k})$
  if such $\tilde{\tau}$ does not exist **then**
    $H_t = \max\{\min\{H_{max}, D_t\}, D_t - \lambda_t \sum_i r_{i,t}\}$;
  else
    $H_t = \max\{\min\{D_t, \max\{H_{max}, D_{\tilde{\tau}}\}\}, D_t - \lambda_t \sum_i r_{i,t}\}$;
  end if
  Solving the sub-problem (6) using Alg. 5 when energy reduction target is $D_t - H_t$;
  Update H_{max};
end for

Theorem 2. *Alg. 2 is* $(1 + 2(\kappa + 1)/\rho + 2)$*-competitive.*

Proof: We analyze the competitive ratio by comparing Alg. 2 with the offline optimal algorithm for the relaxed problem. For any time slot t, let H_t and $x_{i,t}$ be the decisions made by the offline optimal algorithm, while H'_t and $x'_{i,t}$ are those made by Alg. 2. As the requests are revealed one after another, the maximum amount of energy drawn from the grid that is allowed by Alg. 2 increases gradually, and finally approaches H^*_{max} defined in (5). Therefore, $H'_t \leq H_t, \forall t$. Let C_{off} and C_{on} be the total cost incurred by the offline optimal algorithm and by Alg. 2, respectively.

$$\begin{aligned}
C_{off} &= f_{peak} H^*_{max} + \sum_t f_t H_t + \sum_{i,t} p_t r_{i,t} x_{i,t} \\
&= f_{peak} H^*_{max} + \sum_t f_t H_t + \sum_t (p_t/\lambda_t)(D_t - H_t) \\
C_{on} &= f_{peak} H^*_{max} + \sum_t f_t H'_t + \sum_{i,t} p_t r_{i,t} x'_{i,t} \\
&\leq f_{peak} H^*_{max} + \sum_t f_t H'_t + 2 \sum_t (p_t/\lambda_t)(D_t - H'_t)
\end{aligned}$$

where the last inequality holds because Alg. 2 uses a 2-approximation algorithm to (6).

$$\begin{aligned}
\frac{C_{on}}{C_{off}} &\leq \frac{f_{peak} H^*_{max} + \sum_t f_t H'_t + 2 \sum_t p_t/\lambda_t (D_t - H'_t)}{f_{peak} H^*_{max} + \sum_t f_t H_t + \sum_t p_t/\lambda_t (D_t - H_t)} \\
&\leq 1 + \frac{\sum_t p_t/\lambda_t (D_t + H_t - 2H'_t)}{f_{peak} H^*_{max} + \sum_t f_t H_t + \sum_t p_t/\lambda_t (D_t - H_t)}
\end{aligned}$$

Note that $H^*_{max} = D_{t_{\tilde{\tau}}}$, where $f_{peak} = \sum_{k=1}^{\tilde{\tau}} (p_{t_k}/\lambda_{t_k} - f_{t_k})$.

$$\begin{aligned}
&f_{peak} H^*_{max} + \sum_t f_t H_t + \sum_t p_t/\lambda_t (D_t - H_t) \\
= &f_{peak} H^*_{max} + \sum_{k=\tilde{\tau}+1}^{T} f_{t_k} H_{t_k} + \sum_{k=1}^{\tilde{\tau}} f_{t_k} H_{t_k} \\
&+ \sum_{k=1}^{\tilde{\tau}} p_{t_k}/\lambda_{t_k} (D_{t_k} - H_{t_k})
\end{aligned}$$

$$=f_{peak}H^*_{max} + \sum_{k=\tilde{\tau}+1}^{T} f_{t_k}D_{t_k} + \sum_{k=1}^{\tilde{\tau}} f_{t_k}H^*_{max}$$
$$+ \sum_{k=1}^{\tilde{\tau}} p_{t_k}/\lambda_{t_k}(D_{t_k} - H^*_{max})$$
$$= \sum_{k=\tilde{\tau}+1}^{T} f_{t_k}D_{t_k} + \sum_{k=1}^{\tilde{\tau}} D_{t_k}p_{t_k}/\lambda_{t_k}$$

Similarly, we have

$$\sum_t p_t/\lambda_t(D_t + H_t - 2H'_t)$$
$$= \sum_{k=\tilde{\tau}+1}^{T} p_{t_k}/\lambda_{t_k}(D_{t_k} + H_{t_k} - 2H'_{t_k})$$
$$+ \sum_{k=1}^{\tilde{\tau}} p_{t_k}/\lambda_{t_k}(D_{t_k} + H_{t_k} - 2H'_{t_k})$$
$$= \sum_{k=\tilde{\tau}+1}^{T} p_{t_k}/\lambda_{t_k}(2D_{t_k} - 2H'_{t_k})$$
$$+ \sum_{k=1}^{\tilde{\tau}} p_{t_k}/\lambda_{t_k}(D_{t_k} + H^*_{max} - 2H'_{t_k})$$
$$\leq 2\sum_{k=\tilde{\tau}+1}^{T} (p_{t_k}/\lambda_{t_k})\lambda_{t_k}\sum_i r_{i,t_k} + 2\sum_{k=1}^{\tilde{\tau}} p_{t_k}/\lambda_{t_k}(\lambda_{t_k}\sum_i r_{i,t_k} + D_{t_k})$$

Therefore

$$\frac{C_{on}}{C_{off}} \leq 1 + \frac{2\sum_{k=\tilde{\tau}+1}^{T}(p_{t_k}/\lambda_{t_k})\lambda_{t_k}\sum_i r_{i,t_k}}{\sum_{k=\tau\tilde{+}1}^{T} f_{t_k}D_{t_k} + \sum_{k=1}^{\tilde{\tau}} D_{t_k}p_{t_k}/\lambda_{t_k}}$$
$$+ \frac{2\sum_{k=1}^{\tilde{\tau}} p_{t_k}/\lambda_{t_k}(\lambda_{t_k}\sum_i r_{i,t_k} + D_{t_k})}{\sum_{k=\tau\tilde{+}1}^{T} f_{t_k}D_{t_k} + \sum_{k=1}^{\tilde{\tau}} D_{t_k}p_{t_k}/\lambda_{t_k}}$$
$$\leq 1 + 2(\max_t \frac{p_t}{\lambda_t f_t} + 1)/\rho + 2 \leq (1 + 2(\kappa+1)/\rho) + 2 \quad \square$$

Values of κ and ρ are related to the system setting. In our empirical studies in Sec. VI, $\kappa = 3$ and $\rho \approx 2.5$, which results in a competitive ratio guarantee of 6.2. We will further present in Sec. VI that the real competitive ratio observed ranges from 1.1 to 1.2, which is better than the theoretical worst-case bound of 6.2.

It is interesting to observe that, if full information is known, and then the online Alg. 2 degrades into the offline Alg. 1, whose approximation ratio is 2. We can view the competitive ratio $(1 + 2(\kappa + 1)/\rho) + 2$ as two separate components, where the $1 + 2(\kappa + 1)/\rho$ term is due to challenges from the online nature of the problem, while the term 2 results from computational challenges associated with linear integer optimization.

V. The Auction Approach

In the auction approach, the co-location conducts a reverse auction, *a.k.a.* a procurement auction, to solicit energy reduction bids from its tenants. Each tenant participating in the auction submits a biding price as well as the amount of potential energy reduction to the co-location. The co-location computes the winning bids and their corresponding payments. Different from the pricing approach, the bidding prices are determined by tenants themselves instead of the co-location, and are therefore heterogeneous, invalidating the analysis technique in Sec. IV.

A. An Online Algorithm

Similar to the pricing approach, we introduce H_{max}, so that (3) becomes an LIP:

$$\text{minimize} \quad \sum_t f_t H_t + f_{peak}H_{max} + \sum_{t,j} b_{j,t}x_{j,t} \tag{7}$$

$$\text{subject to:} \quad H_t + \lambda_t \sum_{j\in\mathcal{N}_t} r_{j,t}x_{j,t} \geq D_t \qquad \forall t \tag{7a}$$
$$H_t \leq H_{max} \qquad \forall t, \tag{7b}$$
$$x_{j,t} \in \{0,1\}, H_t \geq 0 \quad \forall j \in \mathcal{N}, t \tag{7c}$$

We first relax the above integer program into a linear program. A straightforward relaxation may lead to an unbounded integrality gap. Applying the technique of redundant LP constraints [24], [25], we introduce valid inequalities that are satisfied by all feasible mixed integer solutions of (3), to carefully bound the integrality gap. Such a bound is important for our auction design later. Let $S_t \subseteq \mathcal{N}_t$ be a subset of bids submitted at time slot t. Let $\delta(S_t) = D_t - \lambda_t \sum_{j\in S_t} r_{j,t}$ denote the remaining amount of energy when all bids in S_t are accepted to reduce the energy consumption. Let $r_{j,t}(S) = \min\{\lambda_t r_{j,t}, \delta(S)\}$ be the contribution of an additional bid j in making up the difference. We formulate the following *enhanced LP relaxation* to (7):

$$\text{minimize} \quad \sum_t f_t H_t + f_{peak}H_{max} + \sum_{j,t} b_{j,t}x_{i,t} \tag{8}$$

subject to:

$$H_t + \sum_{j\in\mathcal{N}_t\setminus S} r_{j,t}(S)x_{j,t} \geq \delta(S), \quad \forall t, S \subseteq \mathcal{N}_t : \delta(S) > 0 \tag{8a}$$
$$H_t \leq H_{max}, \qquad \forall t \tag{8b}$$
$$H_t, H_{max}, x_{j,t} \geq 0, \qquad \forall j, t \tag{8c}$$

We proceed to derive the dual LP of (8), by introducing dual variables $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ corresponding to primal constraints $(8a)$ and $(8b)$, respectively. The dual variables admit the following interpretation. α_S is unit energy price in set S, while β_t reflects how much power the co-location wants to draw from the grid. A higher β_t implies that the co-location intends to use power reduction from the tenants more than the grid. Usually high β_t is used when H_t is already too high and any increase in H_t may lead to a high peak charge.

$$\text{maximize} \quad \sum_{t,S\subseteq\mathcal{N}_t:\delta(S)>0} \delta(S)\alpha_S \tag{9}$$

$$\text{subject to:} \quad \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} \alpha_S - \beta_t \leq f_t, \quad \forall t \tag{9a}$$
$$\sum_t \beta_t \leq f_{peak} \tag{9b}$$
$$\sum_{S\subseteq\mathcal{N}_t:j\in\mathcal{N}_t\setminus S,\delta(S)>0} r_{j,t}(S)\alpha_S \leq b_{j,t}, \quad \forall j, t \tag{9c}$$

Based on the dual problem (9), we design a primal-dual online algorithm in Alg. 3. The high level idea behind Alg. 3

is the following. For each time slot t, the algorithm initializes an empty set S_t as a candidate tenant set, from which the co-location purchases energy reduction to reduce its total energy consumption. During each iteration, the dual variable α_{S_t} increases continuously. Once any constraint from $(9a)$ and $(9c)$ becomes tight, the co-location purchases energy from the corresponding tenant and the grid, respectively. Different allocation rules of β_t lead to different variations of the algorithm. We first design an algorithm using a rather straightforward way to allocate f_{peak} to all β_t, *i.e.*, equally allocating over all time slots (line 6).

Algorithm 3 A Primal-Dual Online Algorithm for the Auction Approach

```
1:  // Initialization
2:  x_{j,t} = 0, ∀j, t; H_t = 0, ∀t; α_S = 0, ∀t, S;
3:
4:  for all t do
5:      // when time slot t starts
6:      S_t = ∅; β_t = f_peak/T;
7:      // Iterative update of primal and dual variables:
8:      while δ(S_t) > 0 do
9:          increase α_{S_t} continuously until some constraint gets tight;
10:         if Σ_{S⊆N_t: j∈N_t\S, δ(S)>0} r_{j,t}(S) α_S = b_{j,t} then
11:             x_{j,t} = 1; S_t = S_t ∪ {j};
12:         end if
13:         if Σ_{S⊆N_t: δ(S)>0} α_S − β_t = f_t then
14:             H_t = δ(S_t); break;
15:         end if
16:     end while
17: end for
```

Lemma 1. *Alg. 3 computes feasible solutions to both primal LP (8) and dual LP (9).*

Proof: We first check whether the returned solution is feasible to the primal LP (7). Constraints $(7b)$ and $(7c)$ are always respected because of the settings in line 11 and line 14. If the `while` loop breaks due to $\delta(S_t) \leq 0$, then the demand constraint $(7a)$ is also satisfied. Otherwise it jumps out from the `while` loop because $\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S - \beta_t = f_t$. In this case, $H_t + \lambda_t \sum_j r_{j,t}x_{j,t} = D_t$, which implies the demand constraint is not violated, either. Therefore the solution is feasible to LP (7). We can verify that it is also feasible to LP (8).

We next examine the dual constraints. Once the constraint $(9a)$ becomes tight, the algorithm will break from the `while` loop, and therefore the constraint will not be violated. Constraint $(9b)$ is always respected due to the setting in line 6. Note that once tenant j is selected into the candidate set S_t, the corresponding item in $(9c)$, *i.e.*, $\sum_{S\subseteq\mathcal{N}_t:j\in\mathcal{N}_t\setminus S,\delta(S)>0} r_{j,t}(S)\alpha_S$ will not increase, and hence $(9a)$ will be respected as well. $\square$

We next investigate the component of the competitive ratio that is contributed by the peak demand, in the following lemma.

Lemma 2. *Alg. 3 produces result $(\boldsymbol{H}, \boldsymbol{x})$ such that: $f_{peak}H_{max} \leq \xi/(1-\frac{1}{\rho})\cdot OPT$, where OPT is the cost of the optimal offline solution.*

Proof: $f_{peak}H_{max} = f_{peak}\max_t \delta(S_t) = f_{peak}\delta_{S_{\tilde{\tau}}}$, where $\tilde{\tau} = \arg\max_t \delta(S_t)$.

Note that $\beta_t = f_{peak}/T$, and then $\sum_t \beta_t = f_{peak}$, which implies that:

$$\begin{aligned}
f_{peak}\delta(S_{\tilde{\tau}}) &= \delta(S_{\tilde{\tau}})\sum_t\Big(\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S - f_t\Big)\\
&\leq D_{\tilde{\tau}}\sum_t\Big(\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S - f_t\Big)\\
&\leq \sum_t\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S D_t\xi\\
&= \sum_t\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S D_t(1-\frac{1}{\rho})\xi/(1-\frac{1}{\rho})\\
&\leq \xi/(1-\frac{1}{\rho})\sum_t\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S\delta(S_t)\\
&\leq \xi/(1-\frac{1}{\rho})\sum_t\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S\delta(S)
\end{aligned}$$

The last inequality is due to $S \subseteq S_t, \forall S:\delta(S)>0$ and $\alpha(S)\neq 0 \Rightarrow \delta(S)\geq\delta(S_t), \forall S:\delta(S)>0$ and $\alpha(S)\neq 0$. Following weak duality, $\sum_{t,S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S\delta(S)\leq OPT$, which completes the proof. $\square$

The next lemma examines the competitive ratio incurred by the integrality constraint.

Lemma 3. *Alg. 3 produces a solution $(\boldsymbol{H}, \boldsymbol{x})$ that: $\sum_t f_t H_t + \sum_{j,t} b_{j,t}x_{j,t} \leq 2OPT$.*

Proof: For any $t\in[1,T]$, we analyze the costs by examining the following two cases.

Case 1, the `while` loop terminates due to $\delta(S_t)\leq 0$. We have $\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S - \beta_t \neq f_t$ before the termination. Therefore $H_t = 0$.

$$\begin{aligned}
\sum_j b_{j,t}x_{j,t} &= \sum_{j\in S_t} b_{j,t} = \sum_{j\in S_t}\sum_{S\subseteq\mathcal{N}_t:j\in\mathcal{N}_t\setminus S,\delta(S)>0} r_{j,t}(S)\alpha_S\\
&= \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\sum_{j\in S_t\setminus S} r_{j,t}(S)\alpha_S\\
&\leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\Big(\lambda_t\sum_{j\in S_t\setminus\{j_\omega\}} r_{j,t} - \lambda_t\sum_{j\in S} r_{j,t} + r_{j_\omega,t}(S)\Big)\alpha_S\\
&\leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\Big(D_t - \lambda_t\sum_{j\in S} r_{j,t} + r_{j_\omega,t}(S)\Big)\alpha_S\\
&= \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}(\delta(S) + r_{j_\omega,t}(S))\alpha_S\\
&\leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} 2\delta(S)\alpha_S
\end{aligned}$$

where j_ω is the last tenant added to the solution set S_t for time slot t. The first inequality is due to the definition of $r_{j,t}(S)$. The second inequality is because $\delta(S_t\setminus\{j_\omega\})>0 \Rightarrow D_t > \lambda_t\sum_{j\in S_t\setminus\{j_\omega\}} r_{j,t}$. The last inequality holds as a result of $\delta(S)\geq\min\{\delta(S),\lambda_t r_{j_\omega,t}\} = r_{j_\omega,t}(S)$.

Case 2, the `while` loop terminates because of the `break` in line 14, *i.e.*, $\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S - \beta_t = f_t$. We then have

$$\begin{aligned}
f_t H_t &= \delta(S_t)\Big(\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S - \beta_t\Big)\\
&\leq \delta(S_t)\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S \leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\alpha_S\delta(S)
\end{aligned}$$

where the last inequality is due to $S \subseteq S_t, \forall S : \delta(S) > 0$ and $\alpha(S) \neq 0 \Rightarrow \delta(S) \geq \delta(S_t), \forall S : \delta(S) > 0$ and $\alpha(S) \neq 0$. Similar to the analysis in Case 1, we have:

$$\begin{aligned}\sum_j b_{j,t}x_{j,t} &= \sum_{j\in S_t} b_{j,t} = \sum_{j\in S_t}\ \sum_{S\subseteq\mathcal{N}_t:j\in\mathcal{N}_t\setminus S,\delta(S)>0} r_{j,t}(S)\alpha_S \\ &= \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\ \sum_{j\in S_t\setminus S} r_{j,t}(S)\alpha_S \\ &\leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} (\lambda_t\sum_{j\in S_t} r_{j,t} - \lambda_t\sum_{j\in S} r_{j,t})\alpha_S \\ &\leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} (D_t - \lambda_t\sum_{j\in S} r_{j,t})\alpha_S \\ &= \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} \delta(S)\alpha_S\end{aligned}$$

where the last inequality is because $\delta(S_t) > 0 \Rightarrow D_t > \lambda_t\sum_{j\in S_t} r_{j,t}$. Therefore $f_tH_t + \sum_j b_{j,t}x_{j,t} \leq 2\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\delta(S)\alpha_S$ for the second case. Thus, for any t, we have $f_tH_t + \sum_j b_{j,t}x_{j,t} \leq 2\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\delta(S)\alpha_S$ in either case. In summary: $\sum_t f_tH_t + \sum_{j,t} b_{j,t}x_{j,t} \leq 2OPT$. □

Theorem 3. *Alg. 3 is* $(2 + \xi/(1-\frac{1}{\rho}))$*-competitive.*

Proof: Following Lemma 2 and Lemma 3, we have $\sum_t f_tH_t + f_{peak}H_{max} + \sum_{j,t} b_{j,t}x_{j,t} \leq (2 + \xi/(1-\frac{1}{\rho}))OPT$. Therefore the competitive ratio is $(2 + \xi/(1-\frac{1}{\rho}))$. □

Values of ξ and ρ depend on the specific system settings. In our trace-driven empirical studies in Sec. VI, $\xi \approx 1.2$ and $\rho \approx 2.5$, for which the competitive ratio is 4. The real competitive ratios observed range from 1.1 to 1.5, and are smaller than the worst case bound of 4.

B. A More Intelligent Online Algorithm

In Alg. 3, the dual variable $\boldsymbol{\beta}$ is handled in a somewhat simple way, leading to a competitive ratio of $(2 + \xi/(1-\frac{1}{\rho}))$. In particular, Alg. 3 does not track the current maximum H_t, which makes it less intelligent to the fluctuating power demand D_t as well as unknown bids $(b_{j,t}, r_{j,t})$. Our next goal is to manipulate $\boldsymbol{\beta}$ in a more sophisticated way, for a better performance guarantee. We introduce a new variable to record the maximum demand so far, as shown in Alg. 4.

Theorem 4. *Alg. 4 is* $(2+c)$*-competitive, where* $c = \sum_t \left\{\frac{\max_j b_{j,t}}{\min_j(\lambda_t r_{j,t})} - f_t\right\}/f_{peak}$.

Proof: We observe that i) the offline optimum algorithm will not accept any bids where $\frac{b_{j,t}}{\lambda_t r_{j,t}} \geq f_{peak} + f_t$, since the energy reduction will not reduce any cost in peak charge or volume charge. ii) H_{max} is updated only when the algorithm has to, therefore H_{max} in Alg. 4 is bounded by H^*_{max}. Thus $f_{peak}H_{max} \leq f_{peak}H^*_{max} \leq OPT$.

Since $\frac{b_{j,t}}{\lambda_t r_{j,t}} \geq f_t, \forall j, t$, $\sum_t f_tH_t \leq \sum_t f_tH^*_t + \sum_{j,t} b_{j,t}x^*_{j,t} \leq OPT$. Similar to the proof of the second case in Lemma 3, we have $\sum_{j,t} b_{j,t}x_{j,t} \leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0}\delta(S)\alpha_S$. However $\sum_t \beta_t$ may be increased to a value larger than f_{peak}, which violates constraint $(9b)$. Applying the classic technique of dual fitting, we estimate the upper bound of $\sum_t \beta_t$, and then scale down β_t by this upper bound. The dual variable $\boldsymbol{\alpha}$ is scaled down correspondingly. The scaled dual variables are feasible to the

Algorithm 4 A Smart Primal-Dual Online Algorithm for the Auction Approach

```
// Initialization
x_{j,t} = 0, ∀j,t; H_t = 0, ∀t; H_max = 0; α_S = 0, ∀t,S;

for all t do
    S_t = ∅;
    Eliminate bids where b_{j,t}/(λ_t r_{j,t}) ≥ f_peak + f_t;
    H_max = max{min{D_t, H_max}, D_t − λ_t Σ_j r_{j,t}};
    // Iterative update of primal and dual variables:
    while δ(S_t) > H_max do
        increase α_{S_t} and β_t continuously until some constraint
      in (9c) gets tight exactly before (9a);
        if Σ_{S⊆N_t: j∈N_t\S, δ(S)>0} r_{j,t}(S) α_S = b_{j,t} then
            x_{j,t} = 1; S_t = S_t ∪ {j};
        end if
    end while
    H_t = δ(S_t); Update H_max;
end for
```

dual problem (9), acting as a valid lower bound of the optimal solution.

Since α_{S_t} and β_t are increased continuously until a constraint in $(9c)$ becomes tight exactly before $(9a)$, therefore for each t:

$$\begin{aligned}&\min_j(\lambda_t r_{j,t}) \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} \alpha_S \\ \leq& \sum_{S\subseteq\mathcal{N}_t:j\in\mathcal{N}_t\setminus S,\delta(S)>0} r_{j,t}(S)\alpha_S \leq \max_j b_{j,t}\end{aligned}$$

where the first inequality holds because: i) the definition of $r_{j,t}(S)$ and ii) $\alpha_S = 0$ before tenant j is added to S_t.

Consequently, we set $\beta_t = \frac{\max_j b_{j,t}}{\min_j(\lambda_t r_{j,t})} - f_t$ so that some constraint in $(9c)$ becomse tight exactly before $(9a)$. Then the summation of all β_t is: $\sum_t \beta_t = \sum_t \left\{\frac{\max_j b_{j,t}}{\min_j(\lambda_t r_{j,t})} - f_t\right\}$.

Note that the constraint $(9b)$ requires $\sum_t \beta_t \leq f_{peak}$, thus the scaled down factor is: $c = \sum_t \left\{\frac{\max_j b_{j,t}}{\min_j(\lambda_t r_{j,t})} - f_t\right\}/f_{peak}$. We then have that: $(\boldsymbol{\alpha}/c, \boldsymbol{\beta}/c)$ is a feasible solution to $(9a)$.

$$\begin{aligned}\sum_{j,t} b_{j,t}x_{j,t} &\leq \sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} \delta(S)\alpha_S \\ &\leq c\sum_{S\subseteq\mathcal{N}_t:\delta(S)>0} \delta(S)\alpha_S/c \leq cOPT\end{aligned}$$

Therefore the total cost is

$$\sum_t f_tH_t + \sum_{j,t} b_{j,t}x_{j,t} + f_{peak}H_{max} \leq (2+c)OPT \quad \square$$

The value of c depends on system configuration. In our trace-driven empirical studies, $c = 1.49$, and the competitive ratio is 3.49, which is better than the competitive ratio of Alg. 3. The observed competitive ratios of Alg. 4 range from 1.1 to 1.2 in the empirical studies in Sec. VI.

C. A Truthful Auction Mechanism

We finally design an auction mechanism based on the online Alg. 4, to elicit truthful bids from co-location tenants for each time slot. While the celebrated Vickery-Clarke-Groves (VCG) mechanism is known to be truthful [26], it requires optimally solving social cost minimization multiple times, and is hence computationally infeasible. Our auction is based on the polynomial-time Alg. 4, and inherits a competitive ratio close to that of Alg. 4.

In each time slot, Alg. 4 decides H_{max} that is the maximum amount of energy drawn from the grid. If $H_{max} = D_t$. There is no need to ask tenants to submit energy reduction bids. Otherwise, the energy reduction target is $D_t - H_{max}$. Then the optimization problem (7) becomes:

$$\text{minimize} \quad \sum_{j\in\mathcal{N}_t,t} b_{j,t}x_{j,t} \tag{10}$$

$$\text{subject to:} \quad \lambda_t \sum_{j\in\mathcal{N}_t} r_{j,t}x_{j,t} \geq D_t - H_{max}, \quad \forall t \tag{10a}$$

$$x_{j,t} \in \{0,1\}, \quad \forall j,t \tag{10b}$$

And its corresponding enhanced LP relaxation (8) becomes:

$$\text{minimize} \quad \sum_{j,t} b_{j,t}x_{j,t} \tag{11}$$

subject to:

$$\sum_{j\in\mathcal{N}_t\setminus S} r_{j,t}(S)x_{j,t} \geq \delta'(S), \quad \forall t, S\subseteq\mathcal{N}_t : \delta(S)>0 \tag{11a}$$

$$x_{j,t} \geq 0, \qquad \forall j,t \tag{11b}$$

where $\delta'(S) = D_t - H_{max} - \lambda_t \sum_{j\in S} r_{j,t}$. The corresponding dual problem is shown as follows:

$$\text{maximize} \quad \sum_{t,S\subseteq\mathcal{N}_t:\delta(S)>0} \delta(S)\alpha_S \tag{12}$$

subject to:

$$\sum_{S\subseteq\mathcal{N}_t:j\in\mathcal{N}_t\setminus S,\delta(S)>0} r_{j,t}(S)\alpha_S \leq b_{j,t}, \qquad \forall j,t \tag{12a}$$

$$\alpha_S \geq 0, \quad \forall S\subseteq\mathcal{N}_t \tag{12b}$$

We note that (11) and (12) are simplified versions of (8) and (9) respectively. Thus, a simplified version of Alg. 3 or Alg. 4, shown in Alg. 5 can still work on (11) and (12).

Algorithm 5 A Primal-Dual Online Algorithm for Sub Problem (10)

```
// Initialization
x_{j,t} = 0, ∀j,t; α_S = 0, ∀t,S; S_t = ∅;

while δ(S_t) > 0 do
    increase α_{S_t} continuously until some constraint gets tight;
    if ∑_{S⊆N_t: j∈N_t\S, δ(S)>0} r_{j,t}(S) α_S = b_{j,t} then
        x_{j,t} = 1; S_t = S_t ∪ {j};
    end if
end while
```

Theorem 5. *Alg. 5 is a polynomial-time* 2*-approximation algorithm to the linear integer program (10), and verifies an integrality gap of* 2 *as well.*

Proof: similar to the first case in Lemma 3. □

Now we convert Alg. 5 into a truthful auction by applying a randomized convex decomposition technique [26], [27], [28]. First we solve the LP relaxation in (11) optimally, and then decompose the fractional solution into a convex combination of a series of integer solutions by exploiting the underlying covering structure of the social cost minimization problem. We then pick an integer solution randomly with the corresponding convex combination weights viewed as selection probabilities. The payments to the winning tenants are calculated according to Step three, ensuring that the sufficient and necessary condition of truthfulness is satisfied. We describe the details of the randomized auction as follows.

Step 1), Computing the optimal fractional solution. The relaxed problem (11) can be solved efficiently using a standard LP solution method, such as the simplex algorithm or the interior-point algorithm. After solving (11), we obtain an optimal fractional solution $\boldsymbol{x}^*$.

Step 2), Decomposing fractional solution into integer solutions. We have designed a 2-approximation algorithm verifying an integral gap of 2 for the sub-problem. We use a convex decomposition technique, which employs the approximation algorithm as a plug-in module, to decompose the fractional solution into a set of integer solutions, *i.e.*, find the combination weight $\boldsymbol{\mu}$ where $\sum_{p\in\mathcal{J}} \mu_p = 1$, such that $\sum_{p\in\mathcal{J}} \mu_p \boldsymbol{x}^p = \min\{2\boldsymbol{x}^*, \mathbf{1}\}$ where $\mathcal{J}$ is the set of all feasible integer solutions to (10). The exact decomposition assures the sufficient and necessary condition of truthfulness stated in Theorem 6 is satisfied.

To find $\boldsymbol{\mu}$, we solve the following linear program:

$$\text{maximize} \quad \sum_{p\in\mathcal{J}} \mu_p \tag{13}$$

$$\text{subject to:} \quad \sum_{p\in\mathcal{J}} \mu_p \boldsymbol{x}^p = \min\{2\boldsymbol{x}^*, \mathbf{1}\} \tag{13a}$$

$$\sum_{p\in\mathcal{J}} \mu_p \leq 1 \tag{13b}$$

$$\mu_p \geq 0, \quad \forall p\in\mathcal{J} \tag{13c}$$

However $\mathcal{J}$ has an exponential number of elements, which make (13) have an exponential number of variables, and therefore directly solving (13) is difficult. We consider the dual of (13) instead, which has an exponential number of constraints. We derive the dual (14) by introducing dual variables $\boldsymbol{\nu}$ and η corresponding to (13a) and (13b), respectively.

$$\text{minimize} \quad \sum_{j\in\mathcal{N}_t} \min\{2x_j^*, 1\}\nu_j + \eta \tag{14}$$

$$\text{subject to:} \quad \sum_{j\in\mathcal{N}_t} x_j^p \nu_j + \eta \geq 1 \quad \forall p\in\mathcal{J} \tag{14a}$$

$$\eta \geq 0 \tag{14b}$$

The ellipsoid method [28] can solve the dual efficiently in polynomial time, even though it has exponentially many constraints. We employ Alg. 5 as a separation oracle, which can find violated constraints in the dual (14) and provide them as hyperplanes to the ellipsoid method for cutting the solution space. After receiving a polynomial number of hyperplanes, the ellipsoid method can find the optimal solution to the dual. During the process of the ellipsoid method, a feasible integer solution to (10) is generated when a hyperplane is found. Therefore, after optimally solving the dual (14), the decomposition (13) becomes a linear program with a polynomial number of non-zero primal variables, which can be efficiently solved using standard LP solution algorithms. The following lemma ensures that the convex decomposition method can correctly find a convex combination of integer solutions, *i.e.*, $\sum_{p\in\mathcal{J}}\nu_p = 1$.

Lemma 4. *[27] The decomposition technique can solve* (13) *optimally in polynomial time with optimal objective value* $\sum_{p\in\mathcal{J}}\mu_p = 1$, *meanwhile it finds a polynomial number of feasible integer solutions* $\boldsymbol{x}^p$ *to* (10) *as well as their corresponding convex combination weights* μ_p.

Briefly, if the optimal value of (14) is 1, then $\sum_{p\in\mathcal{J}}\nu_p = 1$ as well by the strong duality of linear program. Note that there is a feasible solution ($\boldsymbol{\nu} = \mathbf{0}, \mu = 1$) to (14). Thus the dual (14) is at most 1. By exploiting the plug-in approximation algorithm, we can find a contradiction if the dual (14) is strict smaller than 1. During the process, violated constraints which act as hyperplanes to the ellipsoid method are found in dual (14).

Step 3), Winner determination and payment calculation. After decomposing the fractional solution into a series of integer solutions, we randomly pick an integer solution $\boldsymbol{x}^p$ with its corresponding combination weight μ_p as the probability. Let $P_j(b_j)$ be the probability that tenant j with bidding price b_j wins. b_{-j} be the all bids except (b_j, r_j). We compute the payments according to the following sufficient and necessary condition of truthfulness.

Theorem 6. *[29], [30] A randomized auction with bids* $\boldsymbol{b}$ *and payment* $\boldsymbol{f}$ *is truthful in expectation if and only if: for any bidder j, i)* $P_j(b_j)$ *is monotonically non-increasing in* b_j*; ii)* $\int_0^\infty P_j(b)db < \infty$*; iii)* $E[f_j] = b_jP_j(b_j) + \int_{b_j}^\infty P_j(b)db$.

We examine the three conditions one by one as follows:

1) Since x_j is a binary variable, we have $P_j(b_j) = P_j(b_j)\times 1 + (1 - P_j(b_j)) \times 0 = E[x_j] = \min\{2x_j^*, 1\}$. For the relaxed problem (11), increasing b_j makes corresponding x_j^* non-increase. Hence, $\min\{2x_j^*, 1\}$ is non-increasing in b_j. So $P_j(b_j)$ is monotonically non-increasing in b_j;

2) Any bid with bidding price higher than $\lambda_t r_{j,t}(f_{peak}+f_t)$ will be removed from the candidate set at the very beginning. Therefore $\int_0^\infty P_j(b)db = \int_0^{\lambda_t r_{j,t}(f_{peak}+f_t)} P_j(b)db < \infty$.

3) For losing tenants, the payment is 0; for winning tenants, the payment

$$f_j = b_j + \frac{\int_{b_j}^\infty \min\{2x_j^*(b, b_{-j}), 1\}db}{\min\{2x_j^*(b_j, b_{-j}), 1\}db}$$

We can verify that the payment satisfies condition iii) in Theorem 6.

Theorem 7. *The auction in Sec. V, is truthful in expectation and achieves a* $(2+2c)$*-approximation ratio, where c is defined in Theorem 4.*

Proof: As the randomized auction satisfies the sufficient and necessary condition in Theorem 6, it is truthful in expectation. The expected social cost:

$$\begin{aligned} &E[\sum_t f_tH_t + \sum_{j,t} b_{j,t}x_{j,t} + f_{peak}H_{max}] \\ &\le 2OPT + \sum_t E[\sum_j b_{j,t}x_{j,t}] \\ &\le 2OPT + 2OPT_{D_t - H_{max}} \le (2+2c)OPT \end{aligned}$$

where $OPT_{D_t-H_{max}}$ is the optimal cost of (11) when the target is $D_t - H_{max}$. The last inequality is due to the dual fitting in Theorem 4.

VI. Performance Evaluation

Simulation Setup. We consider a practical scenario of a co-location data center located in Vancouver, Canada, with 15 participating tenants. The maximum power demand of the co-location data center is 21 MW. The data center is powered by BC Hydro with a peak charge of \$9.95/kW and a volume charge of ¢4.86/kWh. We then generate a series of time varying volume charge rate f_t by adding randomness to the volume charge of ¢4.86/kWh.

We next collect the hourly ambient temperature in Vancouver from July 1, 2014 to July 31, 2014 [31]. Based on the ambient temperature, we compute the pPUE according to the fitting formula in (1). The time-varying curve, from July 1 to July 7, 2014, is depicted in Fig. 3.

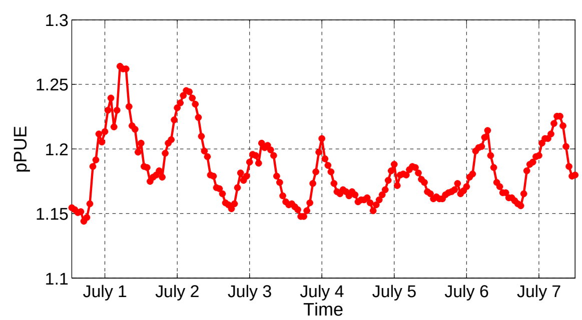

Fig. 3. pPUE values, generated based on the temperature in Vancouver from July 1, 2014 to July 7, 2014.

The *Hotmail* and *MSR* workload traces [12] and the *Wikipedia* workload trace [32], which are 24 hours long, are used to drive the simulation. Since the trace data is limited, we duplicate them with randomness of up to 20% to generate the 15 tenants' workloads for 30 days. All workloads are normalized with respect to each tenant's maximum service capacity. Fig. 4 illustrates the three traces for a 48-hour period.

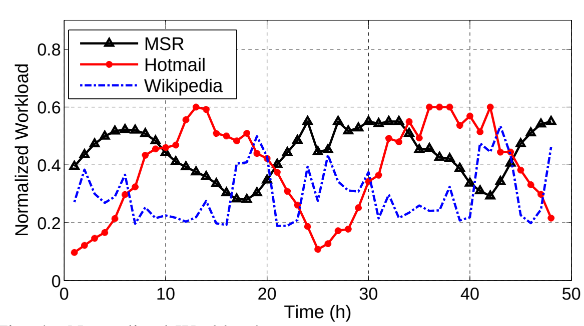

Fig. 4. Normalized Workload.

Based on the workload traces, we generate the power demand D_t. The tenants can achieve energy reduction by consolidating low workload machines and shutting down idle machines. We assume that the machines in idle status consume up to 50 % energy of their peak power, therefore the possible energy reduction is $(50\%+50\%\times W)C-100\%WC = \frac{1-W}{2}C$, where W is the normalized workload and C is the peak power demand when $W = 1$.

A tenant's cost is proportional to its total energy reduction. We set the unit price to be ¢$5 \sim$ ¢8.5 /kWh at random, which is more expensive than the volume charge.

Algorithms for the Pricing Approach. We set the price offered by the co-location $p_t = \kappa f_t$, where by default $\kappa = 3$. If $b_{i,t}/r_{i,t} \leq p_t$, then tenant i will participate in the energy reduction by submitting $r_{i,t}$ to the co-location, where $b_{i,t}$ is tenant i's true cost.

We first compare the approximation algorithm in Alg. 1 with the offline optimum, as shown in Fig. 5. We observe that Alg. 1 achieves almost the same H_{max} as the offline optimum does. Alg. 1 draws the same amount of energy from the grid in most time slots, which verifies the correctness of Eqn. 5.

We next investigate the performance of the online algorithm. Comparison between Alg. 1 and Alg. 2 is shown in Fig. 6. We have two key observations: i) Alg. 1 smooths the demand evenly by the help of tenants; ii) Due to lack of future information, Alg. 2 is unlikely to use high H_t; however the peak demand H_{max} finally approaches the optimal H^*_{max} as the number of requests that the algorithm receives increases. For example, H_{max} reaches H^*_{max} at around $t = 125h$.

Next we compare the overall cost for one month, *i.e.*, $T = 2880$, in Tab. III. The cost achieved by Alg. 1 is rather close to the offline optimum, and is much lower than suggested by the theoretically proven approximation of 2 in Theorem 1. We also notice that the overall cost achieved by Alg. 2 is close to the offline optimum as well, showing a competitive ratio of 1.09. That is also noticeably better than the theoretical bound in Theorem 2, which is $(1 + 2(\kappa + 1)/\rho) + 2 = 6.2$, where $\kappa = 3, \rho = 2.5$.

TABLE III
COST COMPARISON AMONG ALG. 1, ALG. 2, OFFLINE OPTIMUM AND A SIMPLE APPROACH WITH NO SCHEME, UNDER BC HYDRO'S PRICING SCHEME.

	Volume Charge ($)	Peak Charge ($)	Payment to tenants ($)	Total ($)
Algorithm 1	385,000	146,737	12,365	544,102
Algorithm 2	314,811	146,694	130,026	591,531
Offline Optimum	384,697	146,757	12,560	544,014
No Scheme	412,417	185,568	0	597,985

We compare the volume charge, peak charge and the payments to tenants among Alg. 1, Alg. 2 and offline optimum, under various pricing schemes used in Duke Energy, PG&E, Mid American Energy and Georgia Power, respectively. The results are illustrated in Fig. 7, Fig. 8, Fig. 9 and Fig. 10, respectively. Both Alg. 1 and Alg. 2 achieve close-to-optimal performance as compared with the offline optimum, under all pricing scheme. Alg. 2 incurs higher costs than Alg. 1 due to lack of future information. We also observe that the peak charge is the dominant cost under Georgia Power's pricing scheme, which is different from other three pricing scheme. This is because the volume charge of Georgia Power is an order of magnitude cheaper than other utility providers'. More importantly, Alg. 2 achieves a better performance under Georgia Power's pricing scheme as shown in Fig. 10, closer to the offline optimum, showing a competitive ratio nearly 1. This is due to that Alg. 2 is sensitive to the peak demand raise, and tries to avoid any unnecessary increase in peak demand. Therefore Alg. 2 works better under a pricing scheme where the peak charge is the dominant one.

We next compare the energy drawn from the grid by Alg. 1 under various κ, ranging from 1.4 to 4.4. Results are shown in Fig. 11. When κ is small, *i.e.*, the unit price p_t offered by the co-location is cheap, the system has to draw a large amount of energy from the grid due to tenants not willing to offer energy reduction. As κ increases, the amount of energy from the grid decreases. We also compare the energy drawn from the grid by Alg. 2 in Fig. 12, where similar trend is observed.

Algorithms for the Auction Approach. We compare the performance of the algorithms in Sec. V with the offline optimum.

In Fig. 13, Alg. 3 tries to save more energy than the offline optimum does to lower the peak demand, making itself spend more money in paying tenants' energy reduction at most time slots. However Alg. 3 does not memorize the maximum demand so far, so the amount of energy drawn from the grid is not stable over time. More importantly this makes Alg. 3 occasionally raise the peak demand to a high value, *e.g.*, $t = 125$, but fail to utilize the raised peak demand in the consecutive time slots. We also observe that the pattern of the amount of energy drawn from the grid by Alg. 3 exactly follows power demand $D(t)$, *i.e.*, it goes high as $D(t)$ goes high, and it drops as $D(t)$ drops, which reveals that Alg. 3 is not intelligent enough.

In Fig. 14, we compare the amount of energy drawn from the grid by Alg. 4 and that by the offline optimum. We observe that H_{max} in Alg. 4 approaches H^*_{max} as t increases. Compared with Alg. 3, Alg. 4 keeps tracking the maximum demand so far, and resorts to the energy reduction from tenants only when it would exceed the maximum demand so far. Results show that Alg. 4 does act more intelligently than Alg. 3.

We run the randomized auction in Sec. 5.3 for 20 times to obtain the average results. The amount of energy drawn from the grid by the randomized auction is depicted in Fig. 15, compared with Alg. 4. It follows the curve by Alg. 4 with a bit of fluctuation, which results from the randomized selection of integer solutions in the auction. We also observe that both Alg. 4 and the randomized auction jump up a step at $t = 125$. The reason is that the amount of energy drawn from the grid has to jump as the energy reduction from the tenants can not compensate the amount $D(t)$ jumps at $t = 125$.

We next compare the monthly cost in each component, *i.e.*, volume charge, peak charge and tenant cost, as shown in Fig. 16 and Tab. IV. Alg. 3 is the worst among them, but still achieves 1.104 times the offline optimum, which is much better than the theoretical worst-case competitive ratio in Theorem

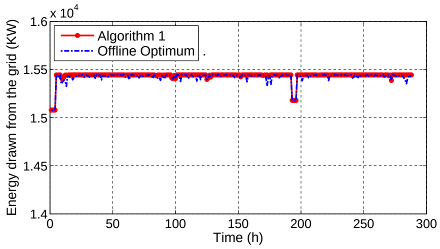

Fig. 5. Comparison between Alg. 1 and Offline Optimum.

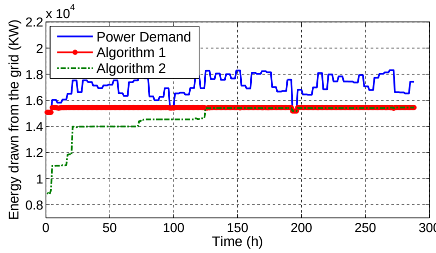

Fig. 6. Energy drawn from grid, July 1-3, 2014.

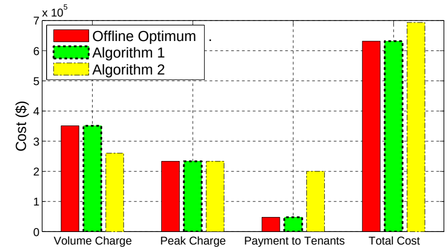

Fig. 7. Under Duke Energy's pricing scheme.

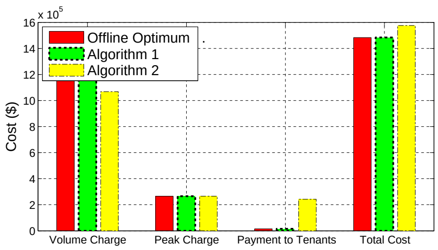

Fig. 8. Under PG&E's pricing scheme.

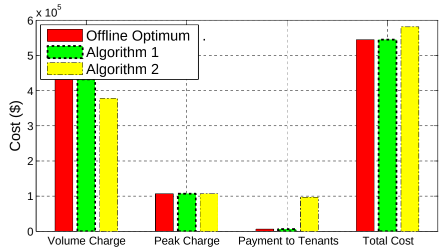

Fig. 9. Under Mid American Energy's pricing scheme.

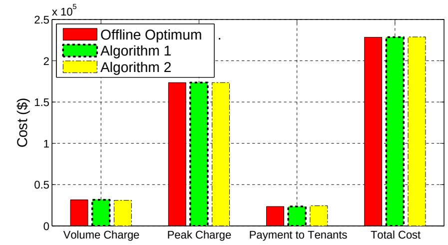

Fig. 10. Under Georgia Power's pricing scheme.

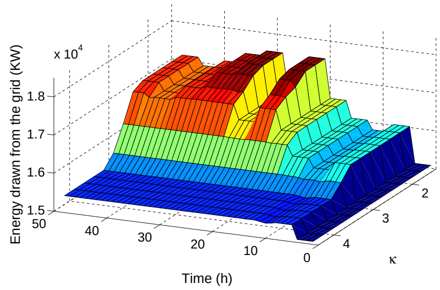

Fig. 11. Alg. 1 under various κ.

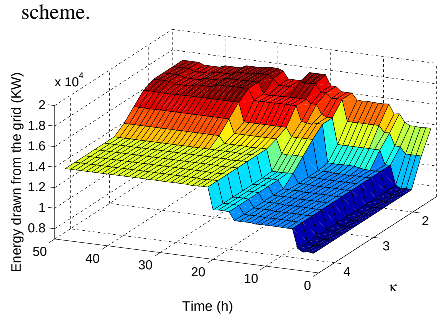

Fig. 12. Alg. 2 under various κ.

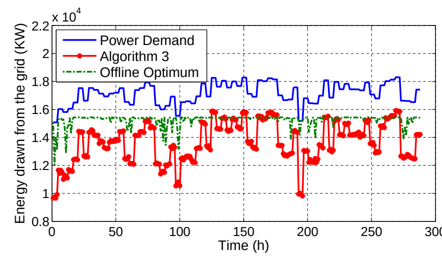

Fig. 13. Energy drawn from grid, under BC Hydro pricing.

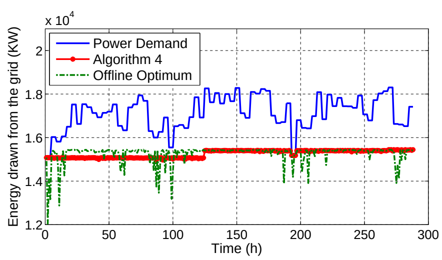

Fig. 14. Comparison among power demand $\boldsymbol{D}$, Alg. 4, Offline Optimum, under BC Hydro's pricing scheme.

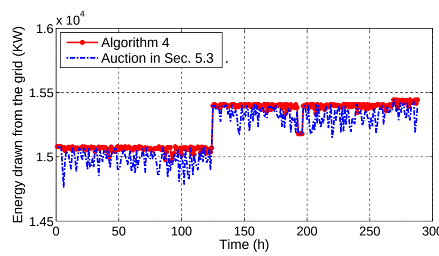

Fig. 15. Comparison between Alg. 4 and the Truthful Auction in Sec. 5.3.

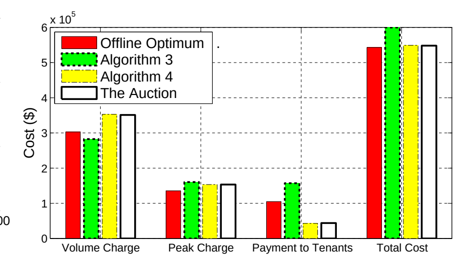

Fig. 16. Comparison between Alg. 4 and the Truthful Auction in Sec. 5.3.

$3 — 2+\xi/(1-\frac{1}{\rho}) = 4$, where $\xi = 1.2, \rho = 2.5$. Alg. 4 and the Auction both have a competitive ratio of 1.01, better than the theoretical worst-case competitive ratios in Theorem 4 and Theorem 7, respectively. In the trace, $c \approx 1.49$, then the ratios are 3.49 and 4.98, respectively.

TABLE IV
COST COMPARISON AMONG ALG. 3, ALG. 4, THE TRUTHFUL AUCTION AND OFFLINE OPTIMUM, UNDER BC HYDRO'S PRICING SCHEME.

	Volume Charge ($)	Peak Charge ($)	Payment to tenants ($)	Total ($)
Algorithm 3	282,561	160,195	157,084	599,840
Algorithm 4	353,078	153,323	42,728	549,129
Truthful Auction	351,119	153,323	43,695	548,137
Offline Optimum	302,776	135,452	104,894	543,122

VII. CONCLUSION

We studied the online electricity cost minimization problem at a co-location data center, considering that the electricity billing model applied to a data center is nowadays based on both the total volume consumed, and the peak consumption rate. We consider two approaches to provide incentive for tenants to shed energy consumption, for peak demand control. We designed online algorithms based on primal-dual techniques that exploit the salient feature of the data center electricity charge model, and proved guarantees on their competitive ratios. We further converted the online algorithm into an efficient auction mechanism that executes in an online fashion, runs in polynomial time, and guarantees truthful bidding and close-to-optimal social cost. Trace-driven simulation studies further verified the efficacy of the proposed algorithms, showing close-to-optimum performance in most cases that were studied.

REFERENCES

[1] "Colocation Market - Worldwide Market Forecast and Analysis (2013 - 2018)," http://www.marketsandmarkets.com/Market-Reports/colocation-market-1252.html.

[2] "BC Hydro," https://www.bchydro.com/.

[3] "Understanding Peak Demand Charges," http://energysmart.enernoc.com/bid/325336/Understanding-Peak-Demand-Charges.

[4] 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," in *Proc. of ACM SIGMETRICS*, 2013.

[5] "PG&E Electric Rates," http://www.pge.com/tariffs/electric.shtml.

[6] "Duke Energy Progress North Carolina Rates & Tariffs," http://www.duke-energy.com/pdfs/G10-NC-Schedule-LGS-TOU-dep.pdf.

[7] "Georgia Power: Electric Service Tariff," http://www.georgiapower.com/pricing/files/rates-and-schedules/large-business/5.10_PLH-9.pdf.

[8] "Mid American Energy Electric Tariffs," http://www.midamericanenergy.com/rates1.aspx.

[9] Enaxis Consulting, "Pricing Data Center Co-location Services," 2009, http://enaxisconsulting.com.

[10] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the Electric Bill for Internet-scale Systems," in *Proc. of ACM SIGCOMM*, 2009.

[11] "Demand Response," http://en.wikipedia.org/wiki/Demand_response.

[12] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic Right-sizing for Power-proportional Data Centers," in *Proc. of IEEE INFOCOM*, 2011.

[13] B. Guenter, N. Jain, and C. Williams, "Managing Cost, Performance, and Reliability Tradeoffs for Energy-aware Server Provisioning," in *Proc. of IEEE INFOCOM*, 2011.

[14] H. Wang, J. Huang, X. Lin, and H. Mohsenian-Rad, "Exploring Smart Grid and Data Center Interactions for Electric Power Load Balancing," *SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 3, pp. 89–94, Jan. 2014.

[15] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing Data Center Demand Response," in *Proc. of ACM SIGMETRICS*, 2014.

[16] C. Wang, B. Urgaonkar, Q. Wang, G. Kesidis, and A. Sivasubramaniam, "Data Center Cost Optimization Via Workload Modulation Under Real-World Electricity Pricing," 2013. [Online]. Available: http://arxiv.org/abs/1308.0585

[17] H. Xu and B. Li, "Reducing Electricity Demand Charge for Data Centers with Partial Execution," in *Proc. of ACM e-Energy*, 2014.

[18] A. Bar-Noy, M. Johnson, and O. Liu, "Peak Shaving through Resource Buffering," *Lecture Notes in Computer Science*, vol. 5426, pp. 147–159, 2009.

[19] N. Buchbinder and J. (Seffi) Naor, "The Design of Competitive Online Algorithms via a Primal: Dual Approach," *Found. Trends Theor. Comput. Sci.*, vol. 3, no. 2-3, pp. 93–263, Feb. 2009.

[20] Y. Azar, U. Bhaskar, L. Fleischer, and D. Panigrahi, "Online Mixed Packing and Covering," in *Proc. of ACM-SIAM SODA*, 2013.

[21] "Liebert DSE Precision Cooling System," http://www.emersonnetworkpower.com/en-US/Products/thermal-management/LargeRoomCooling/Documents/SL-18927.pdf.

[22] H. Xu, C. Feng, and B. Li, "Temperature Aware Workload Management in Geo-distributed Datacenters," in *Proc. of USENIX ICAC*, 2013.

[23] X. Fan, W.-D. Weber, and L. A. Barroso, "Power Provisioning for a Warehouse-sized Computer," in *Proc. of ACM ISCA*, 2007.

[24] R. D. Carr, L. K. Fleischer, V. J. Leung, and C. A. Phillips, "Strengthening Integrality Gaps for Capacitated Network Design and Covering Problems," in *Proc. of ACM-SIAM SODA*, 2000.

[25] T. Carnes and D. Shmoys, "Primal-Dual Schema for Capacitated Covering Problems," *Integer Programming and Combinatorial Optimization*, vol. 5035, pp. 288–302, 2008.

[26] R. Lavi and C. Swamy, "Truthful and Near-Optimal Mechanism Design via Linear Programming," in *Proc. of IEEE FOCS*, 2005.

[27] L. Zhang, S. Ren, C. Wu, and Z. Li, "A Truthful Incentive Mechanism for Emergency Demand Response in Colocation Data Centers," in *Proc. of IEEE INFOCOM*, 2015.

[28] R. Carr and S. Vempala, "Randomized Metarounding," *Random Struct. Algorithms*, vol. 20, no. 3, pp. 343–352, May 2002.

[29] A. Archer and E. Tardos, "Truthful Mechanisms for One-Parameter Agents," in *Proc. of IEEE FOCS*, 2001.

[30] R. B. Myerson, "Optimal Auction Design," *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58–73, 1981.

[31] "Climate, Government of Canada," http://climate.weather.gc.ca/.

[32] V. Vasudevan, A. Phanishayee, H. Shah, E. Krevat, D. G. Andersen, G. R. Ganger, G. A. Gibson, and B. Mueller, "Safe and Effective Fine-grained TCP Retransmissions for Datacenter Communication," in *Proc. of ACM SIGCOMM*, 2009.

Linquan Zhang received his B.E. degree in 2010 from the Department of Computer Science and Technology, Tsinghua University, China, and his M.Phil. degree in 2012 from the Department of Computer Science, the University of Hong Kong. Since September, 2012, he has been with the Department of Computer Science, at the University of Calgary, Calgary, Canada, where he is currently a Ph.D. candidate. His research interests are mainly in cloud computing, network function virtualization, and smart grids.

Zongpeng Li received his B.E. degree in Computer Science and Technology from Tsinghua University (Beijing) in 1999, his M.S. degree in Computer Science from University of Toronto in 2001, and his Ph.D. degree in Electrical and Computer Engineering from University of Toronto in 2005. Since August 2005, he has been with the Department of Computer Science in the University of Calgary. In 2011-2012, Zongpeng was a visitor at the Institute of Network Coding, Chinese University of Hong Kong. His research interests are in computer networks.

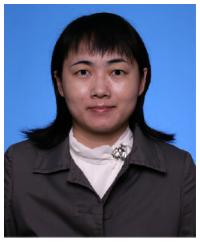

Chuan Wu received her B.Engr. and M.Engr. degrees in 2000 and 2002 from the Department of Computer Science and Technology, Tsinghua University, China, and her Ph.D. degree in 2008 from the Department of Electrical and Computer Engineering, University of Toronto, Canada. Since September 2008, Chuan Wu has been with the Department of Computer Science at the University of Hong Kong, where she is currently an Associate Professor. Her research is in the areas of cloud computing, online and mobile social networks. She was a co-recipient of the best paper award of HotPOST 2012.

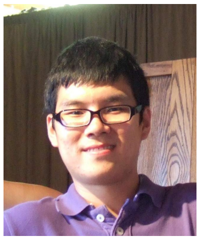

Shaolei Ren is an Assistant Professor of Electrical and Computer Engineering at University of California, Riverside. Previously, he was an Assistant Professor at Florida International University from 2012 to 2015. He received his B.E. from Tsinghua University in 2006, M.Phil. from Hong Kong University of Science and Technology in 2008, and Ph.D. from University of California, Los Angeles in 2012, all in electrical and computer engineering. His research interests include power-aware computing and systems, smart grid, and network economics. He was a recipient of the NSF CAREER Award in 2015.