

An Online Incentive Mechanism for Emergency Demand Response in Geo-Distributed Colocation Data Centers

Qihang Sun
The University of Hong Kong
qhsun@cs.hku.hk

Shaolei Ren
UC Riverside
sren@ece.ucr.edu

Chuan Wu
The University of Hong Kong
cwu@cs.hku.hk

Zongpeng Li
University of Calgary
zongpeng@ucalgary.ca

ABSTRACT

Deferring batch workload in data centers is promising for demand response to enhance the efficiency and reliability of a power grid. Yet operators of multi-tenant colocation data centers still resort to eco-unfriendly diesel generators for demand response, because tenants lack incentives to defer their workloads. This work proposes an online auction mechanism for emergency demand response (EDR) in geo-distributed colocation data centers, which incentivizes tenants to delay and shuffle their workload across multiple data centers by providing monetary rewards. The mechanism, called **BatchEDR**, decides the tenants' workload deferment/reduction and diesel usage in each data center upon receiving an EDR signal, for cost minimization throughout the entire EDR event, considering that only a limited amount of batch workloads can be deferred throughout EDR as well as across multiple data centers. Without future information, **BatchEDR** achieves a good competitive ratio compared to an omniscient offline optimal algorithm, while ensuring truthfulness and individual rationality over the auction process. Trace-driven experiments show that **BatchEDR** outperforms the existing mechanisms and achieves good social cost.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Design studies; Modeling techniques; I.1.2 [Algorithms]: Analysis of algorithms

Keywords

Colocation Data Centers; Emergency Demand Response; Primal-dual Online Algorithms

1. INTRODUCTION

The large yet flexible energy demand of data centers represents a valuable demand response resource for grid stability [8, 11, 32]. In particular, data centers are already im-

portant participants in emergency demand response (EDR). Taking up 87% of all demand response capacities in the U.S. [18], EDR protects the grid as the last line of defense and coordinates large energy consumers, including data centers, to shed consumption during emergency events (e.g., severe weather) [21]. EDR has been widely adopted throughout the world and can even be executed multiple times each day in some developing countries where power infrastructure is increasingly fragile. Further, as more renewables are incorporated into the grid and result in a higher volatility in power supply, we anticipate that EDR will be playing an even more crucial role.

The emergence of data center demand response, albeit highly desired, has also raised environmental concerns, because data centers typically participate by turning on back-up diesel generators that are both costly and environmentally unfriendly [6, 36]. On-site diesel generator produces over 50 times of NOx particles per unit energy generation as compared to a typical coal-fired power plant [25]. Several studies [5, 13, 30–32] have investigated server resource management techniques (e.g., switching off idle servers) to modulate energy demand, as a low-cost and green alternative for data center demand response [11].

While the existing efforts on server energy management for demand response are encouraging, they have been primarily focused on owner-operated data centers like Google, where the operator has full control over both servers and the data center facility. In practice, another type of data center—multi-tenant colocation data center (also called colocation or colo)—has become common, yet less studied. In a colocation data center, the operator manages the facility and provides support (including power, cooling, and space) to multiple tenants who control their own physical servers in a shared space. Colocations satisfy the data center needs of many industry sectors that do not want to completely rely on public clouds. For example, many companies house the entirety of their private clouds (or the private part of a hybrid cloud) in colocations. Even IT giants, including Microsoft and Amazon [6, 36], rely on colocations to complement their private data center infrastructure. The U.S. alone witnesses over 1,400 colocations, a number still in rapid rise, driven by the surging demand [7].

Compared to private data centers, colocations can be more appropriate candidates for demand response, given their typical metropolitan location where demand response is mostly needed. Nonetheless, colocation demand response presents a new “split incentive” challenge: colocation operator wants

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

e-Energy'16, June 21–24, 2016, Waterloo, ON, Canada

© 2016 ACM. ISBN 978-1-4503-4393-0/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2934328.2934331>

demand response for financial compensation from the grid, but relies on diesel generation to do so; tenants manage their own servers and have no incentive to shed consumption for demand response.¹

To overcome the split-incentive hurdle, recent studies [6, 24, 36] have proposed various incentive mechanisms through which the colocation operator passes down economic benefits to tenants for enabling cost-effective and green colocation demand response (without fully relying on diesel generation). These studies often have two limitations. *First*, they focus on a *single* occurrence of EDR and a single colocation, whereas EDR typically occurs over an extended period of time throughout a wide region, over which the colocation operator often manages multiple data centers. For example, EDR was triggered multiple times in January of 2014 throughout several eastern states in the U.S. due to extremely cold weather [19], and Equinix (a large colocation operator) has several data centers in these affected areas [9]. These colocation data centers are not *isolated*, because over 50% tenants house their servers in more than one colocations and can spatially route their workloads from one colocation to another [27]. *Second*, more importantly, prior studies on colocation demand response [6, 24, 36] implicitly assume that tenants reduce their IT energy by turning off servers which process user-facing production workloads (e.g., online shopping). Such interactive workloads are highly delay-sensitive and, in practice, tenants may have concerns with the associated performance degradation. Nonetheless, tenants are highly diverse in colocations and many workloads are delay-tolerant (e.g., lab testing, Hadoop), and are particularly suitable for demand response due to their scheduling flexibilities. In fact, LBNL’s field test of data center demand response used delay-tolerant batch workloads for experiments [11].

In this work, we study demand response by colocations, an important but less-explored type of data centers. Different from the existing research on colocation demand response [6, 24, 36], we focus on demand response in multiple colocations where tenants deploy their delay-tolerant batch workloads, which requires new mechanisms. Unlike delay-sensitive interactive workloads that must be processed promptly, batch workloads can often be flexibly re-scheduled, which also implies that tenants participating in demand response need to manage their batch workloads over a time span (e.g., a few hours). On the other hand, an EDR event often lasts a few hours or even a couple of days [19], over which participating tenants’ scheduling decisions of batch workloads must be carefully coordinated online without violating performance constraints (e.g., throughput for batch jobs over a period of time). Existing mechanisms for one-time EDR [6, 24, 36] that *greedily* schedule tenants’ energy reduction without accounting for batch workload scheduling over time are no longer applicable and can result in significant cost increases (see Section 6). Moreover, a tenant’s workloads located in multiple colocations affected by an EDR event need to be jointly coordinated.

We propose an online auction mechanism, called BatchEDR, to enable low-carbon and cost-effective EDR in colocations by incentivizing tenants to delay some batch workloads. BatchEDR collects bids from tenants online (in-

dicating their maximum workload deferment in each colocation) and, upon the arrival of each EDR signal, decides the participating tenants’ batch workload deferment subject to quality-of-service (QoS) requirements. BatchEDR does not require future information such as tenants’ future bids or EDR signals. Our main contributions are as follows.

- We formulate the problem of dynamically delaying tenants’ batch workloads in an online manner for EDR in geo-distributed colocations, with the goal of minimizing social cost (including both tenants’ performance cost and operator’s diesel cost).
- We design a new online mechanism, BatchEDR, which computes efficient online decisions to coordinate deferment of tenants’ batch workloads for EDR while achieving truthfulness. BatchEDR also meets each tenant’s requirements on the maximum one-round energy reduction across all data centers as well as maximum overall energy reduction throughout the entire EDR event. BatchEDR attains a good competitive ratio (around 1.05) under realistic setting in terms of overall social cost, even as compared with the offline optimum.
- We perform trace-based simulations to validate BatchEDR and our results show that BatchEDR achieves a close-to-minimum social cost and outperforms the best-known solution (e.g. Truth-DR [36]).

In the rest of the paper, we discuss the related work in Sec. 2, and present the problem model in Sec. 3. In Sec. 4 and Sec. 5, we propose the online algorithm framework and the one-round EDR mechanism. In Sec. 6 and Sec. 7, we show the experiment results and conclude this paper.

2. RELATED WORK

Data center demand response has quickly emerged as an important mechanism to transform data center’s huge energy demand from a negative to a valuable social asset. Ghatikar *et al.* [11] first verify the feasibility of data center demand response through field tests. Ghamkhari *et al.* [10] and Aikema *et al.* [1] consider ancillary services, i.e., voluntarily reducing workload, to optimize the resource management of data center to gain additional revenues from the utility. Several studies [5, 14, 17] achieve minimizing aggregate cost of data centers and utilities via dynamic pricing [14], tuning server power usage for regulation service [5], and minimizing social cost of data center demand response via pricing [17]. Other studies optimize resource management for data center demand response, e.g., partial execution [35], battery discharging [2], and geographic load balancing [22, 37]. Nevertheless, all these studies assume that data center operators have full control over the servers, and hence they are not applicable for colocations.

Several recent studies have investigated colocation demand response. Ren *et al.* [24] propose a *simple* mechanism for colocation, called iCODE, which is on a best-effort basis without satisfying EDR requirement and also is not truthful. Zhang *et al.* [36] propose a single-round randomized auction mechanism, Truth-DR, for EDR. [6] studies the same problem but uses parameterized supply function bidding. These studies [6, 24, 36] only focus on a single data center and, more importantly, assume that tenants manage delay-sensitive workloads for energy shedding for a *single* time slot, which may not be suitable for demand response in practice.

¹Tenants are not eligible for directly participating in demand response through the utility, because the utility cannot monitor individual tenants’ energy consumption.

Our study is also relevant to auction design, which has been applied in various contexts. For example, Zhou *et al.* [38] propose an online procurement auction mechanism in smart grid, achieving truthfulness, computationally efficiency, and a constant competitive-ratio. But, they assume that the bidders’ supply is always greater than demand. Babaioff *et al.* [3] design an online auction mechanism achieving constant approximation under a monotone hazard rate distribution. Goel *et al.* [12] propose a truthful auction achieving Pareto-optimal in an online manner, with total payment constraints. These studies, except Truth-DR [36], are not suitable for EDR, where both strategic tenants and the operator-controlled diesel generation can contribute to energy reduction; although Truth-DR is designed for EDR, it is intended for a single time slot and cannot achieve desirable properties, such as one-round reduction limitation at geo-distributed data centers, performance guarantee, and truthfulness, in the online EDR scenario.

3. PRELIMINARIES AND PROBLEM FORMULATION

3.1 Incentivization for Colocation EDR

We consider a typical EDR event (e.g., PJM’s EDR [19]) that lasts for several hours over a large region. A colocation operator signs up for the EDR program in advance (e.g., three years ahead in PJM [19]) and receives rewards for energy reduction commitment: energy reduction is mandated during emergency that lasts up to a certain period of time and non-compliance incurs a heavy penalty.² It operates M geo-distributed data centers over the EDR region serving N tenants in total, each of which rents space and power to run their servers in all or a subset of the data centers. During an EDR event, the colocation operator’s data centers will contribute to energy reduction for a total of T time slots (e.g., hours) as per the contract with the grid. In each time slot, the grid sends an energy reduction signal to each data center j , specifying the amount of energy to reduce $R_j^{(t)}$ in this time slot.³ Such an EDR requirement can be fulfilled by energy usage reduction from tenants in the data center and/or energy production using on-site diesel generators. Let α_j denote the cost of producing one unit of energy using diesel in data center $j \in [M]$.

A tenant may run delay-sensitive interactive workloads (e.g., user-facing web service) and delay-tolerant back-end batch workloads (e.g., testing, back-end data processing) in these data centers. While interactive workloads must be processed with stringent delay requirements, batch workloads typically have large scheduling flexibilities, both over time and over locations [4, 16], and can be postponed to be executed after EDR ends. Thus, we consider that participating tenants reduce their servers’ energy for EDR by de-

²When signing up for EDR, the colocation operator may also share some of its received rewards with tenants to get them committed to energy reduction during EDR. Then, a *coordination* mechanism is still required to allocate energy reduction among self-interested participating tenants (with private workload information) during the actual EDR for social cost minimization. This is the focus of our study, and BatchEDR serves the purpose.

³Reduction is decided based on a predetermined reference value (e.g., the past average value or tenant’s power capacity subscription from the colocation) [19].

ferring batch workloads (which is also the approach adopted in LBNL’s field test of data center demand response [11]).

Each tenant’s batch workloads arrive online and can be quantified depending on the specific application, e.g., GB for sorting data (one of Hadoop’s default applications). While batch workloads are delay-tolerable, tenant still has a QoS constraint (e.g., throughput requirement for data processing) limiting the amount of workloads that can be deferred until EDR ends. Here, we consider two important constraints for both total (equivalent to “average” over the course of EDR) and per-slot amounts of deferrable workloads: first, let w_i represent the maximum amount of batch workloads that tenant $i \in [N]$ can defer during the entire T time slots of EDR across all data centers; second, let $v_i^{(t)}$ be the maximum amount of workloads that tenant i can delay in each time slot t over all of its data centers.⁴ **Note** that deferred batch workloads result in batch workload *reduction during EDR*, and hence we interchangeably use (batch) workload deferment and workload reduction without ambiguity. By deferring batch workloads, tenants can reduce energy consumption (e.g., via turning off unused servers [11]), and we use $f_{i,j}^{(t)}$ to convert a unit of workload reduction to the amount of energy reduction for tenant i in data center j during time t . The value of $f_{i,j}^{(t)}$ is reported to the colocation operator by tenant i , and as proved later, it is truthful.

Auction mechanism. BatchEDR incentives and coordinates tenants’ workload reduction based on an auction mechanism, where energy reduction is viewed as the “good” to be allocated to bidders. Specifically, to incentivize tenants’ contribution for EDR by deferring their workloads, the colocation operator hosts a procurement auction upon the arrival of each energy reduction signal $R_j^{(t)}$ and then *allocates* the energy reduction among participating tenants (as well as its diesel generator). The N tenants are bidders in the T auctions, each submitting a bid at the beginning of each time slot. The bid of tenant i includes a vector of three tuples: $(b_{i,j}^{(t)}, f_{i,j}^{(t)}, e_{i,j}^{(t)})$, $\forall j \in [M]$, where $e_{i,j}^{(t)}$ denotes the maximum amount of batch workload it can reduce in data center j in time slot t , $f_{i,j}^{(t)}$ is the factor converting the amount of deferred workload to energy reduction for data center j in t , and $b_{i,j}^{(t)}$ is its claimed cost (due to batch workload deferment) for per-unit energy reduction in data center j in t . Note that the bid’s values are all zero in data centers where the tenant does not deploy servers. As part of its bid, tenant i also submits $v_i^{(t)}$ to specify its maximum tolerable amount of workload reduction over all the data centers at time t . Further, in the first auction, tenant i also informs the colocation operator of its total workload reduction/deferment constraint w_i for the operator’s consideration over the T auctions.

Online decisions. In each time slot t , based on the collected bids and diesel cost in the data centers, the colocation operator makes the following decisions: (i) The amount of energy reduction (through batch workload reduction) by tenant i in data center j in this time slot. Here, we use $x_{i,j}^{(t)} \in [0, 1]$, $\forall i \in [N]$, $\forall j \in [M]$, each of which represents a percentage of tenant i ’s submitted maximum workload reduction in data center j in t , such that tenant i is asked to reduce energy at the amount of $f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}$. (ii) The re-

⁴Our analysis still applies when adding another constraint on deferrable workloads for each data center.

Table 1: Notation

Var	Definition
N	# of tenants
M	# of data centers
T	# of time slots
$b_{i,j}^{(t)}$	per-unit energy cost of tenant i in data center j at time t
$f_{i,j}^{(t)}$	workload-to-energy conversion factor of tenant i in data center j at time t
$e_{i,j}^{(t)}$	maximum energy reduction amount that tenant i can reduce in data center j at time t
w_i	maximum deferrable workload amount of tenant i across all data centers during all T time slots
$v_i^{(t)}$	maximum deferrable workload amount of tenant i across all data centers in time slot t
$x_{i,j}^{(t)}$	tenant i 's reduction % in data center j at t
$y_j^{(t)}$	amount of energy to produce by using diesel generators in data center j at time t
$R_j^{(t)}$	required energy reduction in data center j at t
α_j	per-energy diesel generator cost in data center j
$\phi_i^{(t)}$	tenant i 's reward at time t
$u_i^{(t)}$	tenant i 's utility at time t

spective financial awards, $\phi_i^{(t)}$, which is the total reward for tenant i in time t . (iii) The amount of energy to produce using diesel generator in each data center, $y_j^{(t)}, \forall j \in [M]$, in order to meet the EDR energy reduction requirement $R_j^{(t)}$. For ease of reference, we list important notations in Table 1.

Our goals of online auction mechanism design are three-fold: (1) *Truthfulness*: As a highly desired property of mechanism design to avoid cheating behaviors, truthfulness is formally defined in **Def. 1** for our auction model. (2) *Individual rationality*: Let $u_i^{(t)}$ represent the utility of tenant i in time slot t , equivalent to the reward it receives minus its true cost of energy reduction at the time. At any time, each tenant always obtains a non-negative utility, *i.e.* $u_i^{(t)} \geq 0, \forall i \in [N], \forall t \in [T]$. (3) *Social cost minimization*: Social cost is a standard measure to assess efficiency of a mechanism [29] and, in our system, it is the sum of the tenants' overall net cost and colocation operator's total cost. A tenant's net cost is its cost due to workload reduction minus the reward from the colocation operator, $\sum_{j \in [M]} \sum_{t \in [T]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} - \sum_{t \in [T]} \phi_i^{(t)}$. The colocation operator's cost is due to using diesel generator and offering rewards to tenants, $\sum_{j \in [M]} \sum_{t \in [T]} \alpha_j y_j^{(t)} + \sum_{i \in [N]} \sum_{t \in [T]} \phi_i^{(t)}$. As the rewards cancel each other, the social cost is equivalent to the sum of tenants' cost due to workload deferment and the operator's cost of using diesel, *i.e.* $\sum_{i \in [N]} \sum_{j \in [M]} \sum_{t \in [T]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \sum_{t \in [T]} \alpha_j y_j^{(t)}$, which represents the overall negative impact of EDR on the entire colocation system. Social cost has been a commonly metric in data center demand response studies [6, 17, 36], and minimizing it is also equivalent to maximizing social welfare.

DEFINITION 1. (Truthfulness) *The auction mechanism is truthful if for any tenant i in time slot t , reporting a bid $(b_{i,j}^{(t)}, f_{i,j}^{(t)}, e_{i,j}^{(t)}), \forall j \in [M]$ and $v_i^{(t)}$ that reveal its true information always maximizes its utility, regardless of the bids of other tenants.*

3.2 Social Cost Minimization Problem

We first formulate the offline social cost minimization problem which, assuming that the operator knows all the reduction signals and the tenants' true bids beforehand, provides the "perfect" solution and serves as a benchmark to compare BatchEDR with.

minimize:

$$\sum_{i \in [N]} \sum_{j \in [M]} \sum_{t \in [T]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \sum_{t \in [T]} \alpha_j y_j^{(t)} \quad (1)$$

subject to:

$$\sum_{j \in [M]} \sum_{t \in [T]} e_{i,j}^{(t)} x_{i,j}^{(t)} \leq w_i, \quad \forall i \in [N] \quad (1a)$$

$$\sum_{i \in [N]} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + y_j^{(t)} \geq R_j^{(t)}, \quad \forall j \in [M], \forall t \in [T] \quad (1b)$$

$$\sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)} \leq v_i^{(t)}, \quad \forall i \in [N], \forall t \in [T] \quad (1c)$$

$$x_{i,j}^{(t)} \leq 1, \quad \forall i \in [N], \forall j \in [M], \forall t \in [T] \quad (1d)$$

$$x_{i,j}^{(t)} \geq 0, \quad \forall i \in [N], \forall j \in [M], \forall t \in [T] \quad (1e)$$

$$y_j^{(t)} \geq 0, \quad \forall j \in [M], \forall t \in [T] \quad (1f)$$

Constraint (1a) specifies that for each tenant, the actual overall workload reduction at all the data centers throughout the entire EDR event cannot exceed the maximum workload reduction constraint w_i . We call this the time-coupling workload reduction budget constraint. (1b) means that the sum of tenants' reduced energy consumption⁵ and diesel-generated energy need to fulfill the EDR requirement set by the grid. (1c) limits the total workload reduction over all data centers for each tenant i at each time t within the limit $v_i^{(t)}$ that the tenant specifies. (1d) guarantees that the operator will not ask any tenant to reduce more workload than the maximum amount it specifies in its bid. Note that back-up diesel generator is designed to support the entire data center' operation for a couple of days and hence the colocation operator can always use diesel generator to fulfill any amount of energy reduction requirement for EDR [6]. In other words, there always exists a feasible solution to the social cost minimization problem.

Complete future information (including energy reduction signals and tenants' true bids over the T time slots) is needed beforehand to solve the optimal offline problem (1). Nonetheless, the colocation operator needs to respond to EDR signals in an online manner without such future information. Thus, we seek to design an online auction mechanism to allocate energy reduction to participating tenants (and diesel generator), while guaranteeing truthful bidding, individual rationality, and competitiveness in terms of social cost. Towards this end, we employ a primal-dual algorithm design framework and formulate the dual of the primal social cost minimization problem (1), by associating dual variable vectors λ , z , σ , and δ with constraints (1a), (1b), (1c), and (1d), respectively. The correspondence between constraints and variables in the primal and dual problems is given in Table 2.

Dual Problem:

⁵Non-IT energy (e.g, cooling) decreases proportionally as tenants reduce server energy and hence, as in prior work [6, 36], is attributed to tenants.

Table 2: variable-constraint correspondence

Primal	(1a)	$x_{i,j}^{(t)}$	(1b)	$y_j^{(t)}$	(1c)	(1d)
Dual	λ_i	(2a)	$z_j^{(t)}$	(2b)	$\sigma_i^{(t)}$	$\delta_{i,j}^{(t)}$

maximize:

$$\begin{aligned}
& - \sum_{i \in [N]} \lambda_i w_i + \sum_{j \in [M]} \sum_{t \in [T]} z_j^{(t)} R_j^{(t)} \\
& - \sum_{i \in [N]} \sum_{t \in [T]} \sigma_i^{(t)} v_i^{(t)} - \sum_{i \in [N]} \sum_{j \in [M]} \sum_{t \in [T]} \delta_{i,j}^{(t)} \quad (2)
\end{aligned}$$

subject to:

$$z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} \leq b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \lambda_i e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)}, \quad \forall i \in [N], \forall j \in [M], \forall t \in [T] \quad (2a)$$

$$z_j^{(t)} \leq \alpha_j, \quad \forall j \in [M], \forall t \in [T] \quad (2b)$$

$$\lambda_i, z_j^{(t)}, \sigma_i^{(t)}, \delta_{i,j}^{(t)} \geq 0, \quad \forall i \in [N], \forall j \in [M], \forall t \in [T] \quad (2c)$$

Next, in Section 4, we first present our online algorithm framework that achieves a reasonably good competitive ratio (close to 1) as compared with the offline optimal social cost, by assuming that a truthful and optimal auction can be carried out in each time slot. Then, in Section 5, we discuss the one-round optimal mechanism, achieving truthfulness and individual rationality.

4. ONLINE ALGORITHM FRAMEWORK

In this section, we design an online algorithm framework \mathcal{A}_{online} as shown in Algorithm 1, which solves the offline optimization problem (1) and its dual (2), utilizing a one-round mechanism \mathcal{A}_{round} (shown in Algorithm 2) in each time slot. Below, we present the one-round social cost minimization problem and discuss the design rationale of our online algorithm.

4.1 One-Round Social Cost Minimization

In time slot t , assuming truthful bids (that will be proved later), we have the following *modified* social cost minimization problem in (3), subject to the same constraints as (1) except for the time-coupling workload reduction constraint (1a) that shall instead be handled by our online algorithm framework \mathcal{A}_{online} . In the modified objective function, instead of using $b_{i,j}^{(t)}$ as tenant i 's cost of per-unit energy reduction, we use $h_{i,j}^{(t)}$, which is larger than $b_{i,j}^{(t)}$ and decided based on tenant i 's amount of remaining workload reduction budget. The rationale will be detailed later. Thus, given tenants' bids and EDR requirement in time slot t , the one-round optimization gives the optimal solution $x_{i,j}^{(t)}, \forall i \in [N], j \in [M]$ and $y_j^{(t)}, \forall j \in [M]$, at t for problem (3) which uses a modified objective function.

minimize:

$$\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j y_j^{(t)} \quad (3)$$

subject to:

$$\sum_{i \in [N]} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + y_j^{(t)} \geq R_j^{(t)}, \quad \forall j \in [M] \quad (3a)$$

$$\sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)} \leq v_i^{(t)}, \quad \forall i \in [N] \quad (3b)$$

$$x_{i,j}^{(t)} \leq 1, \quad \forall i \in [N], \forall j \in [M] \quad (3c)$$

$$x_{i,j}^{(t)} \geq 0, \quad \forall i \in [N], \forall j \in [M] \quad (3d)$$

$$y_j^{(t)} \geq 0, \quad \forall j \in [M] \quad (3e)$$

Similar to the offline optimization problem (3), we introduce dual variable vectors z , σ , and δ , corresponding to constraints (3a), (3b), and (3c), and formulate the following dual problem. Then, according to the one-round dual problem (4), we adapt the solution to one-round dual (4) as a feasible solution to online dual (2).

maximize:

$$\sum_{j \in [M]} z_j^{(t)} R_j^{(t)} - \sum_{i \in [N]} \sigma_i^{(t)} v_i^{(t)} - \sum_{i \in [N]} \sum_{j \in [M]} \delta_{i,j}^{(t)} \quad (4)$$

subject to:

$$z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} \leq h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)}, \quad \forall i \in [N], \forall j \in [M] \quad (4a)$$

$$z_j^{(t)} \leq \alpha_j, \quad \forall j \in [M] \quad (4b)$$

$$z_j^{(t)}, \sigma_i^{(t)}, \delta_{i,j}^{(t)} \geq 0, \quad \forall i \in [N], \forall j \in [M] \quad (4c)$$

The one-round auction mechanism \mathcal{A}_{round} exactly solves the one-round optimization problem (3) to minimize the modified one-round social cost, as well as guarantees truthfulness and individual rationality. We delay the discussion of the one-round auction mechanism in Section 5, but first utilize the properties it achieves in analyzing the performance of our online framework.

4.2 The Online Algorithm

The solution to the one-round problem (3) should be as close as possible to the solution to the offline problem (1), in order to achieve a good competitive ratio, defined as the maximum ratio between the offline optimal social cost derived by solving (3) exactly and the social cost produced by solving the one-round problem (1) in each time slot. The difficulty arises from the time-coupling workload reduction budget constraint at each user: for each tenant, the maximum amount of batch workload reduction over all the T time slots is limited by w_i . Thus, when executing auctions in an online manner, we need to explicitly consider how current decisions affect the future ones, i.e., the overall social cost varies with how the total workload reduction budget w_i is split across T times. For example, if a tenant reduces too much batch workload in the early stage of the EDR event and exhausts its budget w_i prior to the end of EDR, the social cost can rise rapidly later on, since the colocation operator can no longer ask this tenant to reduce batch workload and instead may need to produce more energy using diesel generator for EDR.

The ideal scenario is that all tenants' workload reduction budgets can last for all T rounds of auctions, such that the colocation operator can explore the best energy reduction

strategy among all the tenants over the entire span for minimizing the social cost. Following this intuition, we should avoid exhausting tenants' workload reduction budgets too soon by not too greedily asking tenants to reduce energy at the beginning of the EDR event. Towards this end, we introduce an auxiliary variable λ_i for each tenant $i \in [N]$. Initially, λ_i equals 0, and its value increases with the decrease of tenant i 's workload reduction budget. Precisely, during the execution of our online algorithm, if λ_i has not reached $\Gamma^{[i]} = \min_{j \in [M], t \in [T]} b_{i,j}^{(t)} f_{i,j}^{(t)}$, the workload reduction budget of tenant i will not be exhausted. Instead of the actual cost $b_{i,j}^{(t)}$ for energy reduction, $h_{i,j}^{(t)} = b_{i,j}^{(t)} (1 + \frac{\Lambda \lambda_i^{(t-1)}}{b_{i,j}^{(t)} f_{i,j}^{(t)}})$ is used in the one-round social cost minimization problem (3), where $\Lambda = \max\{1, \max_{i \in [N], j \in [M], T \in [T]} \{\frac{f_{i,j}^{(t)}}{\Gamma^{[i]}} (\alpha_j - b_{i,j}^{(t)})\}\}$, to be solved in \mathcal{A}_{round} , if λ_i has not reached $\Gamma^{[i]}$. In this way, the cost of a tenant with a smaller remaining workload reduction budget will be evaluated higher at the colocation operator. As a consequence, when minimizing the one-round social cost online, the colocation operator will ask for less energy reduction from tenants with a smaller remaining workload reduction budget, and the tenants' total workload reduction budget w_i can last longer throughout the T time slots.

In line 6 of **Alg. 1**, we carefully set the increment of $\lambda_i^{(t)}$, where $\frac{\sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}$ gives the percentage of current reduction over the overall workload reduction w_i , and $\frac{\sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}}{(\gamma-1)w_i}$ reflects the corresponding percentage of social cost increment, where $\gamma = (1 + W_{max}) \frac{1}{W_{max}}$ and $W_{max} = \max_{i \in [N], t \in [T]} \{\frac{\sum_{j \in [M]} e_{i,j}^{(t)}}{w_i}\}$. The way of setting $\lambda_i^{(t)}$ ensures that $\lambda_i^{(t)}$ is always increasing after tenant i reduces its workload, such that the social cost evaluated by the operator based on $h_{i,j}^{(t)}$ is also increasing. Furthermore, it also ensures that the workload reduction constraint (1a) is satisfied over the T rounds of online auctions and a good competitive ratio in social cost can be achieved, as shown in the proof of **Thm. 1**. In line 12, we set λ_i as $\Lambda \lambda_i^{(T)}$, where λ_i is the dual variable in the offline optimization problem (2), associated with the constraint (1a) throughout the online auction. Thus, the increment of $\lambda_i^{(t)}$ can be understood as the adjustment for dual solution, to approach the optimal solution of the offline dual problem (2).

The performance of the online algorithm framework **Alg. 1** is given in **Thm. 1**.

Theorem 1. *If there is a one-round auction mechanism \mathcal{A}_{round} which produces the optimal solutions of (3) and (4), \mathcal{A}_{online} achieves $(\frac{\gamma-1}{\gamma-1-\Lambda})(1 + \frac{(c_1-1)W_{max}}{1+c_0(c_2-1)})$ -competitive for optimization (1). Here, $W_{max} = \max_{i \in [N], t \in [T]} \{\frac{\sum_{j \in [M]} e_{i,j}^{(t)}}{w_i}\}$, $\gamma = (1+W_{max}) \frac{1}{W_{max}}$, $\bar{\alpha} = \max_{j \in [M]} \{\alpha_j\}$, $\underline{\alpha} = \min_{j \in [M]} \{\alpha_j\}$, $\underline{b} = \min_{i \in [N], j \in [M], t \in [T]} \{b_{i,j}^{(t)}\}$, $\bar{f} = \max_{i \in [N], j \in [M], t \in [T]} \{f_{i,j}^{(t)}\}$, $c_0 = \underline{\alpha}/\underline{b}$, $c_1 = \bar{\alpha}/\bar{b}$, and $c_2 = \frac{\sum_{j \in [M], t \in [T]} R_j^{(t)}}{\bar{f} \sum_{i \in [N]} w_i}$.*

PROOF. We prove the correctness and the competitiveness of \mathcal{A}_{online} by proving following three claims.

Claim 1. \mathcal{A}_{online} produces a feasible solution for dual (2).

Algorithm 1: The Online Algorithm Framework-
 \mathcal{A}_{online}

```

1  $\lambda_i^{(0)} = 0, \forall i \in [N]$ ;
2 for  $1 \leq t \leq T$  do
3    $h_{i,j}^{(t)} = \begin{cases} +\infty, & \text{if } \lambda_i^{(t-1)} \geq \Gamma^{[i]} \\ b_{i,j}^{(t)} (1 + \frac{\Lambda \lambda_i^{(t-1)}}{b_{i,j}^{(t)} f_{i,j}^{(t)}}), & \text{otherwise} \end{cases}$ ,
    $\forall i \in [N], \forall j \in [M]$ ;
4   Run the one-round mechanism  $\mathcal{A}_{round}$ , and let  $\mathcal{N}^{(t)}$ 
   denote the set of winning tenants;
5   foreach  $i \in \mathcal{N}^{(t)}$  do
6      $\lambda_i^{(t)} \leftarrow$ 
      $\lambda_i^{(t-1)} (1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}) + \frac{\sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}}{(\gamma-1)w_i}$ ;
7   end
8   foreach  $i \notin \mathcal{N}^{(t)}$  do
9      $\lambda_i^{(t)} \leftarrow \lambda_i^{(t-1)}$ ;
10  end
11 end
12  $\lambda_i \leftarrow \Lambda \lambda_i^{(T)}, \forall i \in [N]$ ;

```

The detailed proof of Claim 1 is given in Appendix A.

Claim 2. *Let $\Delta P^{(t)} = P^{(t)} - P^{(t-1)}$ and $\Delta D^{(t)} = D^{(t)} - D^{(t-1)}$, where $P^{(t)}$ and $D^{(t)}$ are the objective value returned by (1) and (2). In each time slot t , $\Delta D^{(t)} \geq (1 - \frac{\Lambda}{\gamma-1}) \Delta P^{(t)}$ during the process of \mathcal{A}_{online} .*

PROOF. We assume that \mathcal{A}_{round} provides an optimal solution, satisfying the strong duality (e.g. $p = d$), i.e., the value of primal objective function (3) equals the value of dual objective function (4). Let \mathcal{N} denote the set of winning tenants in \mathcal{A}_{round} . In each round, the increments of (3) (i.e., $\Delta P^{(t)}$) and (4) (i.e., $\Delta D^{(t)}$) are as follows:

$$\begin{aligned} \Delta P^{(t)} &= \sum_{i \in \mathcal{N}} \sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j y_j^{(t)} \\ \Delta D^{(t)} &= - \sum_{i \in \mathcal{N}} \Lambda (\lambda_i^{(t)} - \lambda_i^{(t-1)}) w_i + \sum_{j \in [M]} z_j^{(t)} R_j^{(t)} \\ &\quad - \sum_{i \in [N]} \sigma_i^{(t)} v_i^{(t)} - \sum_{i \in \mathcal{N}} \sum_{j \in [M]} \delta_{i,j}^{(t)} \\ \Delta D^{(t)} &= d - \sum_{i \in \mathcal{N}} \Lambda w_i (\lambda_i^{(t)} - \lambda_i^{(t-1)}) \end{aligned}$$

According to line 6 in **Alg. 1** and $d = p$, we further have

$$\begin{aligned} \Delta D^{(t)} &= \\ &= d - \sum_{i \in \mathcal{N}} \Lambda w_i \left(\frac{\lambda_i^{(t-1)} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i} + \frac{\sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}}{(\gamma-1)w_i} \right) \\ &= d - \Lambda \sum_{i \in \mathcal{N}} \left(\sum_{j \in [M]} \lambda_i^{(t-1)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \frac{\sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}}{(\gamma-1)} \right) \\ &= p - \Lambda \left(\sum_{i \in \mathcal{N}} \sum_{j \in [M]} \lambda_i^{(t-1)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \frac{\sum_{i \in \mathcal{N}} \sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}}{(\gamma-1)} \right) \\ &\geq p - \Lambda \sum_{i \in \mathcal{N}} \sum_{j \in [M]} \lambda_i^{(t-1)} e_{i,j}^{(t)} x_{i,j}^{(t)} - \frac{\Lambda}{(\gamma-1)} \Delta P^{(t)} \end{aligned}$$

According to (3) and line 3 in **Alg. 1**, we replace p with (b, f, e, x, α, y) and have

$$\begin{aligned} \Delta D^{(t)} &\geq \sum_{i \in \mathcal{N}} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j y_j^{(t)} \\ &\quad - \Lambda \sum_{i \in \mathcal{N}} \sum_{j \in [M]} \lambda_i^{(t-1)} e_{i,j}^{(t)} x_{i,j}^{(t)} - \frac{1}{(\gamma-1)} \Delta P^{(t)} \\ \Delta D^{(t)} &\geq \sum_{i \in \mathcal{N}} \sum_{j \in [M]} (b_{i,j}^{(t)} + \frac{\Lambda \lambda_i^{(t-1)}}{f_{i,j}^{(t)}}) f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j y_j^{(t)} \\ &\quad - \Lambda \sum_{i \in \mathcal{N}} \sum_{j \in [M]} \lambda_i^{(t-1)} e_{i,j}^{(t)} x_{i,j}^{(t)} - \frac{\Lambda}{(\gamma-1)} \Delta P^{(t)} \\ \Delta D^{(t)} &\geq \sum_{i \in \mathcal{N}} \sum_{j \in [M]} b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j y_j^{(t)} - \frac{\Lambda}{(\gamma-1)} \Delta P^{(t)} \\ \Delta D^{(t)} &\geq (1 - \frac{\Lambda}{\gamma-1}) \Delta P^{(t)} \end{aligned}$$

Hence, we could conclude that in each time slot t , $\Delta D^{(t)} \geq (1 - \frac{\Lambda}{\gamma-1}) \Delta P^{(t)}$ during the process of \mathcal{A}_{online} . \square

Claim 3. \mathcal{A}_{online} can produce an almost feasible solution for (1), the primal online problem. Specifically, its output satisfies constraints (1b), (1c), (1d), (1e), and (1f), and weakly satisfies constraint (1a): For all tenant $i \in [N]$,

$$\sum_{j \in [M]} \sum_{t \in [T]} e_{i,j}^{(t)} x_{i,j}^{(t)} \leq w_i (1 + W_{max}) \quad (5)$$

The proof of Claim 3 is given in Appendix B.

Let C_{noLast} represent the social cost where we ignore the last accepted bid of tenants, C_{primal} represent the social cost achieved by \mathcal{A}_{online} in problem (1), and C_{actual} represent the actual social cost where \mathcal{A}_{round} still minimizes problem (3), but does not reduce workload more than remaining. Apparently, $C_{primal} \leq C_{actual} \leq C_{noLast}$, if there is an upper bound with respect to $\frac{C_{noLast}}{C_{primal}}$, and then C_{actual} could be bounded by C_{primal} , that is, we obtain the competitive ratio to optimum. Intuitively, the gap from C_{noLast} to C_{primal} is caused by the tenants' last-time reduction.

We introduce $C^{[L]} = \underline{b}(\bar{f} \sum_{i \in [N]} w_i) + \underline{\alpha}(\sum_{j \in [M]} \sum_{t \in [T]} R_j^{(t)} - (\bar{f} \sum_{i \in [N]} w_i))$ that represents the lower bound of social cost, and $\Delta C^{[U]} = (\bar{\alpha} - \underline{b}) W_{max} (\bar{f} \sum_{i \in [N]} w_i)$ that represents the upper bound of cost increment from C_{primal} to C_{noLast} . In these two terms, $(\bar{f} \sum_{i \in [N]} w_i)$ is the upper bound of total energy of all tenants, and $W_{max} (\bar{f} \sum_{i \in [N]} w_i)$ represents the upper bound of the overall reduced energy of all tenants through their last-time workload reduction.

$$\begin{aligned} \frac{C_{noLast}}{C_{primal}} &\leq \frac{C_{primal} + \Delta C^{[U]}}{C_{primal}} \\ &\leq 1 + \frac{\Delta C^{[U]}}{C^{[L]}} \\ &\leq 1 + \frac{(\bar{\alpha} - \underline{b}) W_{max} (\bar{f} \sum_{i \in [N]} w_i)}{\underline{b}(\bar{f} \sum_{i \in [N]} w_i) + \underline{\alpha}(\sum_{j \in [M]} \sum_{t \in [T]} R_j^{(t)} - (\bar{f} \sum_{i \in [N]} w_i))} \\ &\leq 1 + \frac{(c_1 - 1) \underline{b} W_{max} (\bar{f} \sum_{i \in [N]} w_i)}{\underline{b}(\bar{f} \sum_{i \in [N]} w_i) + c_0 \underline{b} (c_2 - 1) (\bar{f} \sum_{i \in [N]} w_i)} \\ &\leq 1 + \frac{(c_1 - 1) W_{max}}{1 + c_0 (c_2 - 1)} \end{aligned}$$

Hence, we could conclude that $C_{actual} \leq (1 + \frac{(c_1 - 1) W_{max}}{1 + c_0 (c_2 - 1)}) C_{primal}$, and then combining with Claim 2, the competitive-ratio of \mathcal{A}_{online} is $(\frac{\gamma-1}{\gamma-1-\Lambda})(1 + \frac{(c_1 - 1) W_{max}}{1 + c_0 (c_2 - 1)})$.

5. AUCTION MECHANISM

In this section, we introduce the one-round auction mechanism plugged in our online algorithm framework, which efficiently allocates required energy reduction to tenants according to their bids in each time-slot and guarantees individual rationality and truthfulness. We then complete our online auction design by showing the properties that it achieves.

5.1 One-round VCG Auction

The auction mechanism in each round optimally computes the amount of batch workload reduction that each tenant should reduce as well as diesel usage according to the one-round problem in (3), and decides the rewards for winning tenants. We apply the Vickrey-Clarke-Groves (VCG) mechanism for one-round auction, as it achieves the desired properties: concurrent truthfulness and social welfare maximization (equivalent to social cost minimization in our context), if the underlying resource allocation problem can be optimally solved [29]. The one-round social cost minimization problem in (3) belongs to linear program (LP), whose optimal solution can be computed in polynomial time using algorithms such as interior point methods [34].

In the one-round auction, the colocation operator solves (3) to decide the amounts of batch workload reduction for each participating tenant and the amount of diesel usage. Let $f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{*(t)}$ be the optimal energy reduction of tenant i in data center $j, \forall i \in [N], \forall j \in [M]$ in t , and $y_j^{*(t)}$ denote the optimal diesel energy production in data center j in $t, \forall j \in [M]$, computed by (3). Let $\widehat{C}_{-i}^{(t)}$ be the optimal social cost of (3) computed when tenant i is absent. The VCG reward to a winning tenant i is $\phi_i^{(t)} = \widehat{C}_{-i}^{(t)} - (\sum_{i' \neq i, i' \in [N]} \sum_{j \in [M]} h_{i',j}^{(t)} f_{i',j}^{(t)} e_{i',j}^{(t)} x_{i',j}^{(t)} + \sum_{j \in [M]} \alpha_j y_j^{(t)})$. Intuitively, the reward that a winning tenant i receives is the decrease in social cost attributed to i 's presence. The one-round auction mechanism is summarized in **Alg. 2**.

The utility function $u_i^{(t)}$ of tenant i in the VCG auction is typically defined as the difference between its reward received and its cost. In our online framework, a tenant's utility in each round should be related not only to its cost of energy reduction, but also to its remaining workload reduction budget in that round: intuitively, a smaller utility is obtained if a tenant wins in an one-round auction and reduces its workload when its remaining workload reduction budget is small, and larger otherwise. We characterize this property using a utility function: $u_i^{(t)} = \phi_i^{(t)} - \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)}$. Such form of utility function is consistent with the social cost calculation in the one-round problem (3) and also used in prior research [26] for different purposes. In this way, a tenant's workload reduction budget can last for a longer time; and thus the colocation operator can have a larger solution space to better schedule tenants' workload reduction for future EDR signals, efficiently reducing social cost over all T rounds of auctions. Next, we show in **Thm. 2** that truthful bidding is the best strategy for each tenant in the auction. Furthermore, each tenant gains a non-negative utility, based

on the VCG auction theory [29].

Algorithm 2: One-round Auction \mathcal{A}_{round} in time slot t

- 1 Solve LP (3) by interior-point method; let $x_{i,j}^{*(t)}$ and $y_j^{*(t)}, \forall i \in [N], \forall j \in [M]$ denote the solution.
 - 2 **foreach** $i \in [N]$ **do**
 - 3 $h'_{i',j} = \begin{cases} +\infty, & \text{if } i' = i \\ h_{i,j}^{(t)}, & \text{otherwise} \end{cases}, \forall i' \in [N], \forall j \in [M];$
 - 4 Solve LP of (3) where $h_{i,j}^{(t)}$ is replaced by $h'_{i,j}^{(t)}, \forall i \in [N], \forall j \in [M]$, by interior-point method; let $\widehat{C}_{-i}^{(t)}$ represent the optimal objective function value.
 - 5 $\phi_i^{(t)} = \widehat{C}_{-i}^{(t)} - (\sum_{i' \neq i} \sum_{j \in [M]} h_{i',j}^{(t)} f_{i',j}^{(t)} e_{i',j}^{(t)} x_{i',j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)});$
 - 6 **end**
-

Theorem 2. *The one-round VCG auction in Alg. 2 which produces percentages of tenants' workload reduction $x_{i,j}^{(t)}, \forall i \in [N], j \in [M]$, diesel energy production $y_j^{(t)}, \forall j \in [M]$, and rewards to winning tenants $\phi_i^{(t)}, \forall i \in [N]$, is truthful in $(b_{i,j}^{(t)}, f_{i,j}^{(t)}, e_{i,j}^{(t)}), \forall j \in [M]$ and $v_i^{(t)}$, and individually rational.*

PROOF. Suppose tenant i 's per unit energy bid is $b_{i,j}^{(t)}$, and then $h_{i,j}^{(t)}$ can be exactly calculated by definition. In our analysis, we omit the process of converting $b_{i,j}^{(t)}$ to $h_{i,j}^{(t)}$, and we assume that tenant i and other tenants $i' \neq i$ directly submit $h_{i,j}^{(t)}$ and $h_{i',j}^{(t)}$ as their per unit energy bids. Then, according to the rule for calculating reward $\phi_i^{(t)}$, when tenant i submits its truthful bid $h_{i,j}^{(t)}$, its utility is:

$$\begin{aligned} u_i^{(t)} &= \widehat{C}_{-i}^{(t)} - (\sum_{i' \neq i} \sum_{j \in [M]} h_{i',j}^{(t)} f_{i',j}^{(t)} e_{i',j}^{(t)} x_{i',j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}) \\ &\quad - \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} \\ u_i^{(t)} &= \widehat{C}_{-i}^{(t)} - (\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}) \end{aligned}$$

We know that this tenant's energy reduction amount $(f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)})$ and this diesel usage $y_j^{(t)}$ achieve the minimum of $\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}$, which represents the social cost in problem (3). Recall that $\widehat{C}_{-i}^{(t)}$ is the minimum social cost under the same EDR signal but one tenant less; and thus $\widehat{C}_{-i}^{(t)}$ is greater than or equal to $\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}$. We could conclude that each tenant's utility is always non-negative, $u_i^{(t)} \geq 0$.

Next, we compare the utilities under truthful bidding and untruthful bidding. Suppose tenant i reports an untruthful bid $(\vec{h}_{i,j}^{(t)}, \vec{f}_{i,j}^{(t)}, \vec{e}_{i,j}^{(t)}, \vec{v}_i^{(t)})$. Then tenants' energy reduction amount and diesel usage become $(\vec{f}_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)})_{\forall i \in [N], \forall j \in [M]}$ and $(\vec{y}_j^{(t)})_{\forall j \in [M]}$. Recall $v_i^{(t)}$ denotes the maximum amount of tenant i 's reduced workload in all data centers at time t , and it appears in constraint (3b), which is relevant to the

solution space of $x_{i,j}^{(t)}$. Intuitively, combining with $f_{i,j}^{(t)}$ and $e_{i,j}^{(t)}$, the solution space represents the capacity of tenant i 's reducible energy. The untruthful $\vec{v}_i^{(t)}$ might cause two possible changes in the solution space: (a) it shrinks the solution space, that is, it cannot increase the capacity of reducible energy; (b) it expands the solution space: however, the actual space which tenant i can achieve is still the same as before. Therefore, the false bidding cannot improve the maximum amount of its reducible energy. Tenant i 's utility under a false bid is:

$$\begin{aligned} \vec{u}_i^{(t)} &= \widehat{C}_{-i}^{(t)} - (\sum_{i' \neq i} \sum_{j \in [M]} h_{i',j}^{(t)} \vec{f}_{i',j}^{(t)} e_{i',j}^{(t)} x_{i',j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} \vec{y}_j^{(t)}) \\ &\quad - \sum_{j \in [M]} h_{i,j}^{(t)} \vec{f}_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} \end{aligned}$$

The difference between these two utilities is:

$$\begin{aligned} u_i^{(t)} - \vec{u}_i^{(t)} &= (\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}) \\ &\quad - (\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} \vec{f}_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}) \end{aligned}$$

Recall that $(f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)})$ and $y_j^{(t)}$ minimize the objective of (3), i.e., $\sum_{i \in [N]} \sum_{j \in [M]} h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} \alpha_j^{(t)} y_j^{(t)}$. Hence, in the right-hand side of the equation, the terms in the first pair of parentheses is no less than the terms in the second, and $u_i^{(t)} - \vec{u}_i^{(t)} \geq 0$. We can conclude that each tenant could obtain the maximum utility under truthfully reporting. \square

5.2 Online Auction Mechanism

By plugging the one-round VCG mechanism into the online algorithm framework \mathcal{A}_{online} (line 4) in **Alg. 1**, we obtain our online auction mechanism. As the VCG mechanism produces the optimal solution of problem (3) in each round, the competitive ratio of the online auction mechanism is the same as that of \mathcal{A}_{online} shown in **Thm. 1**. **Thm. 3** summarizes our result.

Theorem 3. *\mathcal{A}_{online} in Alg. 1 combining with \mathcal{A}_{round} in Alg. 2 constitutes a truthful, individually rational, and $(\frac{\gamma-1}{\gamma-1-\Lambda})(1 + \frac{(c_1-1)W_{max}}{1+c_0(c_2-1)})$ -competitive online auction.*

PROOF. The only difference between **Thm. 1** and **Thm. 3** is that we plug in a fractional VCG mechanism for ensuring truthfulness and the individual rationality. As the mechanism produces the optimum of problem (3), the competitive ratio of the online auction mechanism is the same as \mathcal{A}_{online} . Combining with **Thm. 2**, it can be seen that in each time-slot, if a tenant reports its true bid, it will obtain the maximum utility. According to **Def. 1**, our online auction mechanism for EDR is truthful. **Thm. 2** furthermore shows that tenant's utility is always non-negative. Hence, we conclude that our online mechanism achieves all desirable properties. \square

Theorem 4. *\mathcal{A}_{online} in Alg. 1 runs in polynomial-time in each round.*

The detailed proof of **Thm. 4** is given in Appendix C.

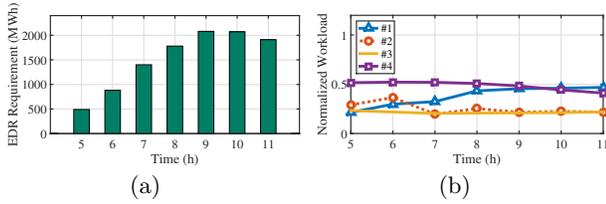


Figure 1: Trace data. (a) Total EDR energy reduction by PJM on January 7, 2014. (b) Normalized workload.

6. PERFORMANCE EVALUATION

Our theoretical analysis has shown that BatchEDR achieves a competitive ratio compared to the offline optimum in terms of social cost. Hence, the goal of this section is to investigate the performance of BatchEDR in a realistic scenario and validate our analysis via a trace-based simulation. Next, we first introduce the simulation setup and then show our results.

6.1 Data Sets and Simulation Setup

We evaluate our mechanism based on trace data, and simulate a real EDR event in a set of colocation data centers located in New Jersey and Virginia, which are primary data center markets served by PJM (a major regional transmission organization in U.S [20]). In our setting, the operator owns five colocation data centers, serving ten *participating* tenants. Each data center has a power usage effectiveness (PUE) factor of 1.5, which is common for colocation data centers [6, 36], i.e., one watt of IT power corresponds to 0.5 watt of non-IT power, adding to 1.5 watt for the whole data center. Each tenants houses 2,000 servers in one data center it rents. Each server has a normalized service capacity of 1.0 (corresponding to one unit of batch workload per time slot), an idle power of 120W and a busy power of 200-250W (accounting for different energy reductions by deferring a unit of batch workload and turning off one unused server). The diesel generator cost α_j is 200\$/MWh based on typical power generation efficiency and the current oil price as of 2015 [33]. In our simulation, the configuration of data centers is shown as follows:

	DC#1	DC#2	DC#3	DC#4	DC#5
# of Tenants	6	6	7	8	6
Location	Ashburn	Newark	Richmond	Trenton	Norfolk

6.1.1 EDR Requirement

Fig. 1(a) shows the total energy reduction requirement for EDR across PJM service areas on January 7, 2014 (severe weather on that day) [19]. For our evaluation, we scale down the data such that each data center is required to reduce energy equal to up to 25% of its peak IT energy (which is consistent with LBNL’s field test [11] that achieves around 25% energy reduction while still preserving data center operation). For example, a data center using 5000kW IT power will be required to reduce up to 1250kWh energy for EDR during one hour.

6.1.2 Workload and Energy

Fig. 1(b) illustrates four *samples* of the server utilization collected from real clusters like Microsoft Research and

Google [15, 23, 28], which we use as tenants’ batch workload arrival rate in each data center (normalized to their maximum service capacity). For instance, 0.3 workload is equal to the service capacity of 600 servers (each with a normalized capacity of 1.00) for a tenant with 2,000 servers. For our evaluation, we consider that tenants turn off unused servers (due to reduced/deferred batch workloads) to eliminate those servers’ power. Due to servers’ different idle powers and non-IT power that takes up 50% of IT power (based on PUE of 1.5), for a server with 200W to 250W power, the value of $f_{i,j}^{(t)}$ that maps batch workload reduction to energy reduction ranges from 300W to 375W per unit of batch workload.

6.1.3 Tenants’ Bids

We set the cost of each tenant $b_{i,j}^{(t)}$ to be 0.07~0.13\$/kWh, which is on a par with the electricity cost saving (had tenants operated their own data centers) and also consistent with prior research on colocation EDR [6, 15, 36]. In each data center, the maximum amount of batch workload that can be deferred/reduced by a tenant is its arrival workload, subject to total workload reduction constraints across all the data centers both over T time slots and over one round (as described below).

6.1.4 Overall Workload Reduction Constraint

For each tenant i , we set the maximum amount of overall workload reduction w_i to be 50% (varying in 40%~60% later) of its total batch workload arrival across all data centers over T time slots.

6.1.5 One-round Workload Reduction Constraint

For each tenant i at time t , we set the default one-round workload reduction limitation $v_i^{(t)}$ as 90% of its total batch workload arrival across all data centers at time t . This value will be varied in 80%~100% as we proceed with the simulation.

6.2 Results

Before presenting our results, we first introduce two benchmarks to compare BatchEDR with.

OPT: This is the optimal mechanism solving the offline problem and optimally deciding the workload reduction and the diesel usage in advance. Note that OPT is not implementable in practice due to lack of future information.

Truth-DR: Proposed in [36], Truth-DR is the best-known auction mechanism for colocation EDR, but it is designed for one-round auction in a single data center. We adapt Truth-DR to our setting by independently executing Truth-DR for each data center in each time, until the workload reduction budget is exhausted.

6.2.1 Close-to-Minimum Social Cost

We first compare the total social costs by time achieved by diesel only method, Truth-DR, BatchEDR, and offline optimal solution, as illustrated in **Fig. 2**. The result shows that BatchEDR is better than Truth-DR and also significantly outperforms the diesel-only method (without incentivizing tenants’ energy reduction). As Truth-DR is a one-round EDR mechanism, it may exhaust one tenant’s reducible workload at the early rounds and cannot reduce workload when cheaper bids emerge subsequently. Thus, BatchEDR outperforms Truth-DR by carefully accounting for tenants’ en-

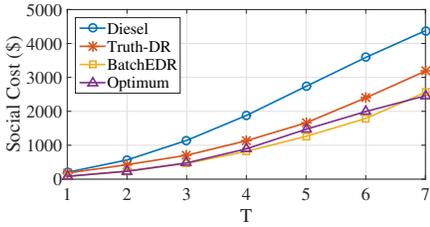


Figure 2: Comparison of social cost among different mechanisms.

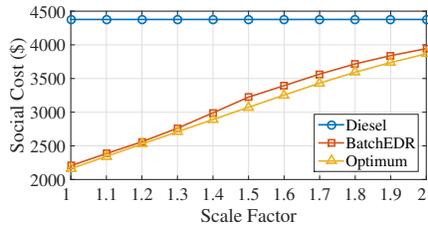


Figure 3: Comparison of social cost under bid scaling among different mechanisms.

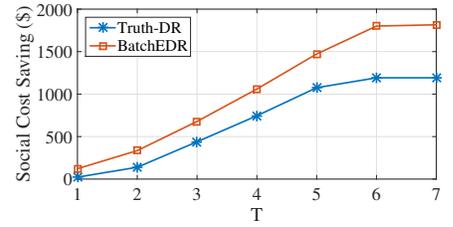


Figure 4: Comparison of social cost saving among different mechanisms.

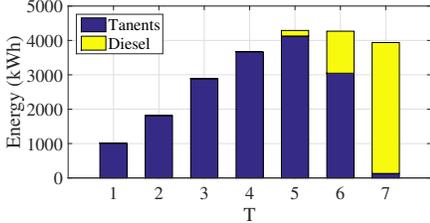


Figure 5: Comparison of reduced energy source between diesel and tenants' workload in BatchEDR.

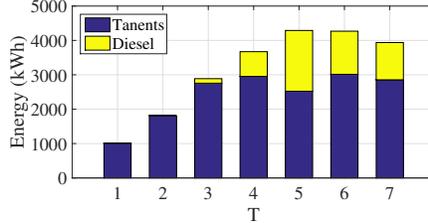


Figure 6: Comparison of reduced energy source between diesel and tenants' workload in OPT.

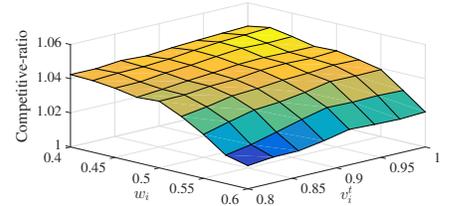


Figure 7: Competitive ratio in social cost : different w_i , different $v_i^{(t)}$.

ergy reduction budget over multiple rounds. OPT decides to reduce workload with future knowledge (*i.e.* knowing the future bids and EDR requirements), and thus it may not always try its best to reduce from tenants in early rounds, such that it can keep the workload reduction budget in the later rounds to reduce the total social cost.

In Fig. 3, we scale each tenant's per-energy cost $b_{i,j}^{(t)}$, showing the increasing trend of the total social cost when it is more costly to defer workloads (due to, e.g., higher business value of processing batch workloads). Nonetheless, BatchEDR still outperforms the diesel-only method and is fairly close to OPT, even though tenant's per-energy cost is two times of the default value.

6.2.2 Cost Saving Compared with Diesel-only

Fig. 4 shows the total social cost saving by BatchEDR and Truth-DR methods (compared to the diesel-only method). We see that BatchEDR yields a larger social cost saving compared with Truth-DR, *i.e.*, BatchEDR achieves up to 32% more social cost saving at the end of the EDR event, compared with Truth-DR.

6.2.3 Energy Reduction Source

Fig. 5 and Fig. 6 show the energy reduction source of BatchEDR and OPT in each time slot, respectively. We see that under BatchEDR, tenants contribute to EDR mostly during the first few time slots, while they cease their contributions during the last time slot due to workload reduction constraint. On the other hand, OPT can better plan tenants' energy reduction with the help of future information, but Fig. 2 has shown that such future information does not offer much gain in terms of social cost.

6.2.4 Competitive Ratio in Total Social Cost

Fig. 7 shows the competitive ratio achieved by BatchEDR in total social cost by varying $v_i^{(t)}$ and w_i (which capture performance constraints by limiting the percentage of a tenant's batch workloads that can be reduced during one round and over all T rounds across all of its data centers). We ob-

verse that the competitive ratio slightly increases with the decrease with w_i ratio and the increase with the v_i ratio. When w_i increases, more workload could be reduced, that is the solution space is expanded and more solution could attain a good social cost close to the optimum, and then BatchEDR could obtain a better solution. Moreover, when v_i is decreasing, the overall reduced workload might last for a longer time, and thus BatchEDR obtains a better social cost.

We have also conducted other experiments (e.g., competitive ratio under different numbers of tenants), but the results are similar and hence omitted for brevity.

7. CONCLUDING REMARKS

This paper investigates online incentive mechanisms for motivating tenants to voluntarily reduce their deferrable workload across multiple geo-distributed colocation data centers, in the EDR scenario. An online efficient and truthful auction mechanism, BatchEDR, is proposed for allocating energy reduction requirements among the tenants' online arrival workloads under one-round and overall limitations of reduced workload, and rewarding the tenants for their workload deferment/reduction. We prove that BatchEDR guarantees a performance bound in social cost, as compared with the offline optimum, and achieves truthfulness. Our trace-based simulation shows that BatchEDR performs well under various settings, compared with other alternative mechanisms. To best of our knowledge, our study is the first effort to design an online, truthful mechanism for EDR in multiple geo-distributed colocation data centers.

8. ACKNOWLEDGEMENTS

This work was supported in part by Hong Kong RGC under grants HKU 717812, 718513, 17204715, and C7036-15G(CRF), and by U.S. NSF under grants CNS-1551661 and CNS-1565474.

9. REFERENCES

- [1] D. Aikema, R. Simmonds, and H. Zareipour. Data Centres in the Ancillary Services Market. In *IGCC*, 2012.
- [2] B. Aksanli and T. S. Rosing. Providing Regulation Services and Managing Data Center Peak Power Budgets. In *DATE*, 2014.
- [3] M. Babaioff, L. Blumrosen, and A. Roth. Auctions with Online Supply. In *EC*, 2010.
- [4] N. Buchbinder, N. Jain, and I. Menache. Online Job Migration for Reducing the Electricity Bill in the Cloud. In *IFIP Networking*, 2011.
- [5] H. Chen, A. K. Coskun, and M. C. Caramanis. Real-Time Power Control of Data Centers for Providing Regulation Service. In *CDC*, 2013.
- [6] N. Chen, X. Ren, S. Ren, and A. Wierman. Greening Multi-Tenant Data Center Demand Response. In *IFIP Performance*, 2015.
- [7] DatacenterMap. Colocation USA, <http://www.datacentermap.com/usa/>.
- [8] EnerNOC. Ensuring U.S. Grid Security and Reliability: U.S. EPA's Proposed Emergency Backup Generator Rule, 2013, http://www.whitehouse.gov/sites/default/files/omb/assets/oira_2060/2060_12102012-2.pdf.
- [9] Equinix, www.equinix.com.
- [10] M. Ghamkhari and H. Mohsenian-Rad. Data Centers to Offer Ancillary Services. In *SmartGridComm*, 2012.
- [11] G. Ghatikar, V. Ganti, N. E. Matson, and M. A. Piette. Demand Response Opportunities and Enabling Technologies for Data Centers: Findings From Field Studies, 2012.
- [12] G. Goel, V. Mirrokni, and R. P. Leme. Clinching Auctions with Online Supply. In *SODA*, 2013.
- [13] S. Li, M. Brocanelli, W. Zhang, and X. Wang. Data Center Power Control for Frequency Regulation. In *PES*, 2013.
- [14] Y. Li, D. Chiu, C. Liu, L. T. 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 *ICDEW*, 2013.
- [15] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska. Dynamic Right-sizing for Power-proportional Data Centers. *TON*, 21(5):1378–1391, 2013.
- [16] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser. Renewable and Cooling Aware Workload Management for Sustainable Data Centers. In *SIGMETRICS*, 2012.
- [17] Z. Liu, I. Liu, S. Low, and A. Wierman. Pricing Data Center Demand Response. In *SIGMETRICS*, 2014.
- [18] K. Managan. Demand Response: A Market Overview, 2014, <http://enaxisconsulting.com>.
- [19] PJM. Demand Response Activity on January 7-8, 2014, <http://www.pjm.com>.
- [20] PJM. PJM Interconnection, <http://www.pjm.com/>.
- [21] PJM. Retail Electricity Consumer Opportunities for Demand Response in PJM's Wholesale Markets, <http://www.pjm.com>.
- [22] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the Electric Bill for Internet-Scale Systems. In *SIGCOMM*, 2009.
- [23] C. Reiss, A. Tumanov, G. R. Ganger, R. H. Katz, and M. A. Kozuch. Heterogeneity and Dynamism of Clouds at Scale: Google Trace Analysis. In *SoCC*, 2012.
- [24] S. Ren and M. A. Islam. Colocation Demand Response: Why Do I Turn Off My Servers? In *ICAC*, 2014.
- [25] Santa Barbara County. Air Pollution Control District, <http://www.ourair.org/do-you-really-need-a-diesel-generator/>.
- [26] W. Shi, L. Zhang, C. Wu, Z. Li, and F. C. Lau. An Online Auction Framework for Dynamic Resource Provisioning in Cloud Computing. In *SIGMETRICS*, 2014.
- [27] Uptime Institute. Data Center Industry Survey, 2014, <http://journal.uptimeinstitute.com/annual-data-center-industry-survey-2014/>.
- [28] 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 *SIGCOMM*, 2009.
- [29] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [30] 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 *MASCOTS*, 2014.
- [31] 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.*, 41(3):89–94, 2014.
- [32] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad. Opportunities and Challenges for Data Center Demand Response. In *IGCC*, 2014.
- [33] Wikipedia, http://en.wikipedia.org/wiki/Diesel_generator.
- [34] S. J. Wright. *Primal-dual Interior-point Methods*. SIAM, 1997.
- [35] H. Xu and B. Li. Reducing Electricity Demand Charge for Data Centers with Partial Execution. In *e-Energy*, 2014.
- [36] L. Zhang, S. Ren, C. Wu, and Z. Li. A Truthful Incentive Mechanism for Emergency Demand Response in Colocation Data Centers. In *INFOCOM*, 2015.
- [37] Z. Zhi, F. Liu, Y. Xu, R. Zou, H. Xu, J. C. S. Lui, and H. Jin. Carbon-aware Load Balancing for Geo-distributed Cloud Services. In *MASCOTS*, 2013.
- [38] R. Zhou, Z. Li, and C. Wu. An Online Procurement Auction for Power Demand Response in Storage-Assisted Smart Grids. In *INFOCOM*, 2015.

APPENDIX

A. PROOF OF CLAIM 1

PROOF. In **Alg.** 1 step 4, we obtain a feasible dual solution of (4), and hence it guarantees the constraint (4a); and based on the dual variable update rules (step 3, 6, 9, and 12

in **Alg. 1**), we have:

$$\begin{aligned}
z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} &\leq h_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)} \\
z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} &\leq \left(b_{i,j}^{(t)} + \frac{\Lambda \lambda_i^{(t-1)}}{f_{i,j}^{(t)}} \right) f_{i,j}^{(t)} e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)} \\
z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} &\leq b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \Lambda \lambda_i^{(t-1)} e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)} \\
z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} &\leq b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \Lambda \lambda_i^{(T)} e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)} \\
z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} &\leq b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \lambda_i e_{i,j}^{(t)} + \sigma_i^{(t)} e_{i,j}^{(t)} + \delta_{i,j}^{(t)},
\end{aligned}$$

which satisfies the constraint (2a).

In the proof, from the first line to the second line, we replace $h_{i,j}^{(t)}$ based on the equation in step 3. As $\lambda_i^{(t)}$ is non-decreasing (step 6 and 9) and $\lambda_i = \Lambda \lambda_i^{(T)}$ (step 12), the constraint (2a) is always feasible.

For each $h_{i,j}^{(t)}$ that equals $+\infty$ in (3) and (4), it implies $\lambda_i^{(T)} \geq \Gamma^{[i]}$. Also, according to (4b) and (4c), we have $\alpha_j \geq z_j^{(t)}$, $\sigma_i^{(t)} \geq 0$, and $\delta_{i,j}^{(t)} \geq 0$. Therefore, we have

$$\begin{aligned}
b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \lambda_i e_{i,j}^{(t)} &= b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \Lambda \lambda_i^{(T)} e_{i,j}^{(t)} \\
&\geq b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \Lambda \Gamma^{[i]} e_{i,j}^{(t)} \\
&\geq b_{i,j}^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)} + \left(\max_{i \in [N], j \in [M], T \in [T]} \{f_{i,j}^{(t)} (\alpha_j - b_{i,j}^{(t)})\} \right) e_{i,j}^{(t)} \\
&\geq \alpha_j f_{i,j}^{(t)} e_{i,j}^{(t)} \\
&\geq z_j^{(t)} f_{i,j}^{(t)} e_{i,j}^{(t)},
\end{aligned}$$

which satisfies the constraint (2a).

Therefore, we conclude that our online framework \mathcal{A}_{online} produces a feasible solution for dual (2) at the end. \square

B. PROOF OF CLAIM 3

We denote $W_{max} = \max_{i \in [N], t \in [T]} \left\{ \frac{\sum_{j \in [M]} e_{i,j}^{(t)}}{w_i} \right\}$, $\gamma = (1 + W_{max})^{\frac{1}{W_{max}}}$, and $\Gamma^{[i]} = \min_{j \in [M], t \in [T]} \{b_{i,j}^{(t)} f_{i,j}^{(t)}\}$.

PROOF. In step 5 of **Alg. 1**, at each round, we obtain a feasible solution for (3), which satisfies constraints (1b), (1c), (1d), (1e), and (1f). To show (5), we utilize the following equation, $\forall i \in [N], t' \in [T]$:

$$\lambda_i^{(t')} \geq \frac{\Gamma^{[i]}}{\gamma - 1} \left(\gamma^{\frac{\sum_{t=1}^{t'} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}} - 1 \right), 0 \leq t' \leq T \quad (6)$$

Our proof includes four parts.

(part 1): We prove **Eq. (6)** always holds in our online algorithm \mathcal{A}_{online} ; **(part 2):** We prove that if a tenant uses up its workload reduction budget, we will have that

$$\lambda_i^{(t')} \geq \frac{\Gamma^{[i]}}{\gamma - 1} \left(\gamma^{\frac{\sum_{t=1}^{t'} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}} - 1 \right) \geq \Gamma^{[i]}; \quad \textbf{(part 3):}$$

We prove that **Alg. 1** will not update $x_{i,j}^{(t)}$ in constraint (1a) when tenant i 's workload reduction budget is used up (*i.e.*, ensuring that the solution is feasible); **(part 4):** We analyze how much the reduction amount will be excessively allocated at most (*e.g.* (5)).

Proof of part 1: We will prove **Eq. (6)** by induction. At first, for $t' = 0$, the equation apparently holds; then, we suppose it holds for $t' - 1$; next, for t' .

$$\begin{aligned}
\lambda_i^{(t')} &= \lambda_i^{(t'-1)} \left(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \right) + \frac{\sum_{j \in [M]} b_{i,j}^{(t')} f_{i,j}^{(t')} e_{i,j}^{(t')} x_{i,j}^{(t')}}{(\gamma - 1)w_i} \\
&\geq \frac{\Gamma^{[i]}}{\gamma - 1} \left(\gamma^{\frac{\sum_{t=1}^{t'-1} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}} - 1 \right) \left(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \right) \\
&\quad + \frac{\sum_{j \in [M]} b_{i,j}^{(t')} f_{i,j}^{(t')} e_{i,j}^{(t')} x_{i,j}^{(t')}}{(\gamma - 1)w_i} \\
&\geq \frac{\Gamma^{[i]}}{\gamma - 1} \left(\gamma^{\frac{\sum_{t=1}^{t'-1} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}} \left(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \right) - 1 \right. \\
&\quad \left. - \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} + \frac{\sum_{j \in [M]} \frac{b_{i,j}^{(t')} f_{i,j}^{(t')} e_{i,j}^{(t')} x_{i,j}^{(t')}}{\Gamma^{[i]}} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \right) \\
&\geq \frac{\Gamma^{[i]}}{\gamma - 1} \left(\gamma^{\frac{\sum_{t=1}^{t'-1} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}} \left(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \right) - 1 \right)
\end{aligned}$$

The first inequality is based on the induction assumption (6).

Comparing with **Eq. (6)**, we next need to show

$$1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \geq \gamma \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \quad \text{to finish the proof.}$$

Based on the definitions of W_{max} and $x_{i,j}^{(t)}$, we have

$$0 \leq \frac{\sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i} \leq W_{max} \leq 1, \forall i \in [N], \forall t \in [T]. \quad \text{We utilize } \frac{\ln(1+x)}{x} \geq \frac{\ln(1+y)}{y}, \forall 0 \leq x \leq y \leq 1. \quad \text{Therefore, we have:}$$

$$\begin{aligned}
\frac{\ln(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i})}{\frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i}} &\geq \frac{\ln(1 + W_{max})}{W_{max}} \\
\ln(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i}) &\geq \ln(1 + W_{max}) \frac{1}{W_{max}} \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} \\
e^{\ln(1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i})} &\geq e^{\ln(1 + W_{max}) \frac{1}{W_{max}} \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i}} \\
1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} &\geq \left((1 + W_{max})^{\frac{1}{W_{max}}} \right)^{\frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i}} \\
1 + \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i} &\geq \gamma \frac{\sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')}}{w_i}
\end{aligned}$$

Hence, we prove **Eq. (6)** always holds in \mathcal{A}_{online} .

Proof of part 2: For a tenant i , at time t' , when its energy reduction $(\sum_{t=1}^{t'} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)})$ exceeds w_i , $\frac{\sum_{t=1}^{t'} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i} \geq 1$, since $\gamma \geq 1$, it is apparent that

$$\frac{\Gamma^{[i]}}{\gamma - 1} \left(\gamma^{\frac{\sum_{t=1}^{t'} \sum_{j \in [M]} e_{i,j}^{(t)} x_{i,j}^{(t)}}{w_i}} - 1 \right) \geq \Gamma^{[i]}, \quad \text{when tenant } i \text{ uses up its reduction budget. Combined with the proof of part 1, the inequality holds.}$$

Proof of part 3: If the workload has been used up, based on the proof of part 2, λ_i will be greater than or equal to $\Gamma^{[i]}$. According to line 3 in **Alg. 1**, when $\lambda_i^{(t-1)} \geq \Gamma^{[i]}$, the algorithm will not reduce tenant i 's workload anymore by setting corresponding $h_{i,j}^{(t)}$ as $+\infty$, which could ensure the solution feasible.

Proof of part 4: In our algorithm, $\sum_{j \in [M]} \sum_{t \in [T]} e_{i,j}^{(t)} x_{i,j}^{(t)}$ equals $\sum_{j \in [M]} \sum_{t=1}^{t'} e_{i,j}^{(t)} x_{i,j}^{(t)}$, where $t = t'$ is the first time tenant's total workload reduction amount exceeds its bud-

get. Hence,

$$\sum_{j \in [M]} \sum_{t=1}^{t'-1} e_{i,j}^{(t)} x_{i,j}^{(t)} < w_i, \text{ then:}$$

$$\begin{aligned} \sum_{j \in [M]} \sum_{t \in [T]} e_{i,j}^{(t)} x_{i,j}^{(t)} &= \sum_{j \in [M]} \sum_{t=1}^{t'-1} e_{i,j}^{(t)} x_{i,j}^{(t)} + \sum_{j \in [M]} e_{i,j}^{(t')} x_{i,j}^{(t')} \\ &\leq w_i + \max_{t \in [T]} \left\{ \sum_{j \in [M]} e_{i,j}^{(t)} \right\} \\ &\leq w_i (1 + W_{max}) \end{aligned}$$

Now, we have bounded how much the workload reduction exceeds the budget with the solution of the primal problem. Next, we analyze how the reduction excess influences the total social cost. \square

C. PROOF OF THEOREM 4

PROOF. Let $|N|$ and $|M|$ denote the number of tenants and the number of data centers respectively, and $|N||M|$ denote the multiplication of $|N|$ and $|M|$. In **Alg. 1**, for a specific t (*i.e.*, in a specific round), the variables $h_{i,j}^{(t)}$ are updated for at most $|N||M|$ times in line 3; the variables $\lambda_i^{(t)}$ are updated at most $|N|$ times in lines 6 and 9; and **Alg. 2** is executed once. In **Alg. 2**, the interior-point method is executed for at most $|N| + 1$ times for computing reduction and reward. As problem (3) is a linear program problem with a polynomial number of decision variables (*i.e.*, $|N||M| + |M|$) and constraints (*i.e.*, $|N| + 2|M| + 2|N||M|$), we can obtain the solution in polynomial time by interior-point method for each round of execution. The variables $h'_{i,j}^{(t)}$ and $\phi_i^{(t)}$ are updated for at most $|N||N||M|$ and $|N|$ times, which also can be finished in polynomial time. Thus, we conclude **Alg. 2** can be finished in polynomial time, and then **Alg. 1** runs in polynomial time in each round. \square