Colocation Demand Response: Joint Online Mechanisms for Individual Utility and Social Welfare Maximization

Qihang Sun, Chuan Wu, Zongpeng Li, Shaolei Ren

***Abstract*—Data centers with high yet elastic energy demand are ideal candidates for participation in demand response programs. This work studies emergency demand response (EDR) at multi-tenant colocation data centers (colocations). While the colocation has no direct control over tenants' servers, we design online mechanisms to incentivize and coordinate tenants' energy reduction. Our mechanism is online in nature, aiming to maximize not only social welfare but also tenant utility. Our main proposal is a truthful incentive auction that provides tenants monetary remuneration for EDR energy reduction, minimizing social cost, which combines seamlessly with an online primal-dual framework for each tenant to schedule their delay-tolerant workloads. The online optimization at each tenant targets its utility maximization, concurrently reporting valuation functions for the tenant to participate in the auction. Our online algorithms achieve long-term performance guarantees in both tenants' utility and social welfare maximization, while fulfilling the EDR requirement with minimal diesel generation. We validate the efficiency of our algorithms through both theoretical analysis and real-world trace-driven simulations.**

***Index Terms*—Colocation Data Centers, Emergency Demand Response, Primal-dual Online Algorithms**

I. Introduction

With the aging infrastructure and rapid incorporation of renewables leading to an increasing intermittency in power supply, large yet flexible energy loads are highly desirable demand response resources for maintaining grid stability. As canonical examples of such resources, megawatt data centers have been playing a crucial role in emergency demand response (EDR), a type of standby service that takes up 87% of all demand response capabilities [1]) and forms a last line of defense for grid stability [2]. For example, on July 22, 2011, hundreds of data centers cut their energy demands for EDR and helped avoid a wide-area blackout throughout North America [3]. Some developing countries have a poor power reliability and may experience outages regularly every day. Even in developed economies, major grid outages have significantly increased over the recent years due to the aging infrastructure, growing demand, and/or extreme weather.

Although EDR helps ease the adoption of renewables and stabilize the grid, the participation of data centers is far from being green. Typically, data centers contribute to EDR by turning on backup diesel generators, which are both costly and eco-unfriendly. For example, the toxic and carbon emissions of a diesel generator can exceed over 50 times the emission level of typical power plants and constitute a major source of pollution in regions such as California [4]).

Fortunately, recent advancements in data center energy management have made it possible to modulate server energy usage, partly in lieu of diesel generation, for cost-effective and *green* EDR. Some notable techniques include dynamically re-sizing the server cluster [5], scaling CPU frequency [6], deferring non-critical workloads [7], and migrating loads to other sites [8]. These techniques have endured real-world testings and have proven themselves appealing EDR options without affecting normal grid operation [9].

However, existing energy management techniques are not directly applicable for EDR in multi-tenant colocation data centers ("colocations"). In a colocation, multiple tenants house their own physical servers, while the operator manages the non-IT facility (*e.g.*, power supply and cooling). This is in contrast to Google-type private data centers whose operators have full control over both servers and facilities. Colocation data centers are already very common today, serving almost all industry sectors. Even Google and Microsoft have recently leased large capacities in colocations, while Apple houses 25% of its servers in colocations [10].Today, the U.S. alone has over 1,400 large colocations, which consume nearly as five times energy as Google-type data centers all combined (37.3% versus 7.8%, in percentage relative to all data center energy usage, excluding tiny server closets). Further, colocations are mostly located in populated areas and hence better targets for EDR, compared to Google-type data centers that are often built in rural areas where EDR is less needed.

In this work, we study cost-effective and green EDR in colocations. This poses unique challenges, however, due to "split incentives": the colocation operator desires EDR for rewards from the grid but cannot force tenants to cut energy, while tenants may not have incentives to contribute to EDR unless incentivized. Furthermore, the grid operator cannot directly reward tenants for EDR, since the grid cannot even meter an individual tenant's energy usage inside a colocation. Even though tenants are willing to contribute to EDR, which tenants should reduce energy and by how much still need to be carefully decided so as to minimize the overall performance loss incurred by participating tenants.

To address the above challenges, a few recent studies have begun to propose incentive mechanisms (*e.g.*, auction design [11] and supply function bidding [12], [13]), through which the colocation operator offers financial rewards to coordinate tenants' energy reduction for EDR. These studies often assume a simplified model: tenants each has a certain performance cost (measured in monetary values) when cutting energy, and a mechanism is executed to select tenants for en-

Q. Sun and C. Wu are with the Department of Computer Science, The University of Hong Kong, Hong Kong (e-mail: {qhsun, cwu}@cs.hku.hk). Z. Li is with the Department of Computer Science, University of Calgary, and School of Computer, Wuhan University (e-mail: zongpeng@ucalgary.ca). S. Ren is with the Department of Electrical and Computer Engineering, University of California, Riverside, USA (e-mail: sren@ece.ucr.edu).

ergy shedding and minimizing the colocation-wide social cost (including tenants' performance loss and diesel usage). This model implicitly assumes delay-sensitive jobs that forbid job shifting/deferring in the temporal domain. In reality, however, tenants have large batch workloads (*e.g.*, MapReduce-based data processing) whose executions span multiple time slots. These workloads are often delay-tolerant and ideal for EDR participation. The payoffs of (batch) workloads also depend on when they are completed [14], which in turn couples tenants' performance costs across multiple time slots, invalidating the model assumed in recent literature [11], [12], [13].

We consider a more practical EDR model where tenants run a set of batch jobs that span multiple time slots. We employ a penalty function to represent the performance costs when jobs are not completed prior to the initial *soft* deadline due to tenants' energy shedding during EDR (*i.e.*, deferring the execution of certain workloads). Our notion of penalty function is general and subsumes the case when tenants' workloads have hard deadlines.

As tenants' jobs and EDR events both span multiple time slots in practice, the previous incentive mechanisms designed over single time slots [11], [12], [13] are no longer applicable. Instead, we must design a new *online* mechanism. Nonetheless, tenants' cost valuations for reducing the same amount of energy are both time-varying (due to random new job arrivals) and coupled with prior scheduling decisions (captured through penalty functions), making the EDR mechanism design particularly challenging.

Our contributions. We leverage state-of-the-art and novel primal-dual online optimization techniques to design: (i) an online, polynomial-time algorithm for job scheduling, maximizing tenant utility based on a general and practical job model; (ii) a novel scheme for tenants to report their valuations to contribute to EDR without utility loss; and (iii) a truthful incentive mechanism to apportion energy reduction among tenants and to minimize the total cost. More concretely, we make the following contributions.

First, we model the entire colocation system, including operator and multiple tenants, under EDR events. We formulate the operator's total cost minimization problem and the tenant's utility maximization problem under a general and practical job model. To design an efficient online algorithm, we convert tenant's original job scheduling problem, including job deadline and penalty function, into a novel formulation.

Second, we design an efficient primal-dual online algorithm for tenant's utility maximization, based on the new tenant problem formulation, achieving a good competitive ratio.

Third, we design a truthful mechanism for the operator to minimize social cost, which decides the EDR reduction from tenants and the remuneration for participating tenants. Combining the operator's incentive mechanism and the tenant's individual online job scheduling algorithm, we further show a good competitive ratio in terms of global social welfare.

II. Related Work

Data center demand response has recently received much attention. For example, several studies focus on optimizing resource management for demand response by, *e.g.*, scaling CPU frequency [15], partial execution [16], geographic load balancing [17] and battery charging/discharging [18]. The recent work [19] formulates a problem similar to ours, but focuses on operator-owner data centers (*i.e.*, with only one tenant). Further, in [20], a prediction-based pricing scheme is studied, assuming that each data center is subject to an unknown cost function of its reduction amount; the cost function does not vary over time, which is different from our model considering dynamic pricing and online job scheduling. Moreover, field tests were recently conducted to demonstrate that data center can shed 20% energy using resource management approaches without significantly affecting normal operation for demand response [9]. These studies are different from and complementary to our work, as they study orthogonal problems and/or focus on Google-type data centers.

For colocation EDR, recent studies proposed incentive mechanisms [11], [12], [13], through which the colocation operator offers financial rewards to coordinate tenants' energy reduction for EDR. These studies cannot handle multiple time slots with online decisions, guarantee a good performance for tenants, or be used for modeling tenants running batch workloads. A more detailed comparison is already provided in Section I, and hence not repeated here.

Our work is also relevant to the literature on online job scheduling problems in cloud data centers. In [21], the authors propose an efficient heuristic algorithm by allowing job preemption for scheduling jobs with specific deadlines, to maximize the aggregate value of jobs completed by their deadlines. In [22], the authors propose LP-based algorithms achieving constant-factor approximation in optimal resource usage, for scheduling jobs with different start and end times. Another study [23] investigates the job scheduling problem without job deadlines, but with unknown job durations. All the existing studies do not investigate value-varying jobs according to different completion times, and none has combined job scheduling with demand response in a colocation data center.

III. System Model and Problem Formulation

A. Colocation EDR and Tenants' Job Models

We consider a colocation data serving N tenants. The *colocation operator* signs up with the grid *a priori* (*e.g.*, three years ahead in PJM [24]) and receives rewards for committed energy reduction during emergencies. An EDR event may last multiple time frames, where each frame can be an hour. In an EDR frame, the colocation operator is required to cap its grid power usage under a level dictated by the grid's EDR signal (or equivalently, reduce energy by a certain amount compared to a baseline value) [24]. A baseline usage, such as the colocation's average energy demand at the same hour over the past few weeks, is typically used to calculate the energy reduction [20].

The colocation can achieve energy reduction in an EDR frame through a combination of (i) tenants' energy reduction and (ii) on-site diesel generation. Without direct control over tenants' servers, the colocation incentivizes tenants' energy reduction by monetary remuneration using an inverse auction. Then, on-site (diesel) generator covers the discrepancy

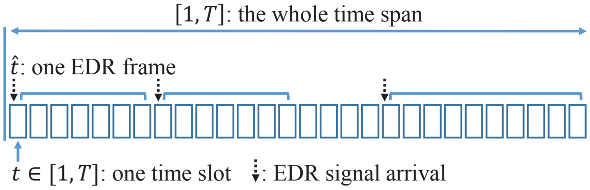

Fig. 1. An illustration of two timescales in our model.

between total EDR demand and total reduction from tenants, at a unit energy cost of α for fuel consumption.

We focus on delay-tolerant batch jobs that are interruptible during execution.[1] Each tenant manages its own servers in the shared data center, and receives batch jobs for execution from time to time. The job arrival typically occurs at a finer timescale than EDR time frames, *e.g.*, minutes versus hours. We refer to the time interval for job arrival and scheduling as "time slot". Let $[1, T]$ denote all the time slots in the system span; T can potentially be a very large number. An EDR signal arrives at a time slot, requesting energy shedding within the subsequent slots in an EDR time frame. Let $R^{(t)}$ denote the total energy reduction amount requested by the grid for each time slot t in the EDR frame. As a most general model, the time frames with EDR signals can be consecutive or non-consecutive, the EDR frame starting at time slot t lasts $\beta^{(t)}$ time slots, and the starting time of an EDR frame can be any time slot after the previous EDR frame ends. **Fig.** 1 illustrates the two time scales.

We assume that compared to the batch job execution time, each time slot is small enough to distinguish the arrival times of different jobs at each tenant. Tenant i receives new job $J_{it'}$ at time slot $t' \in [1, T]$, which requires $d_{it'}$ time slots for execution, with energy consumption $q_{it'}$ per slot when executed. $d_{it'}$ and $q_{it'}$ are zero if no job arrives at t'. Consider the simple example of a MapReduce job processing 10GB data, using 4 mapper servers first and 4 reducer servers subsequently. Suppose that a server consumes an average power of 250W, and that each time slot is 5 minutes. Then, $q_{it'}$ can be estimated as $4 \times 0.25 \times \frac{5}{60}$kWh. $d_{it'}$ can be estimated based on the input data size and historical processing statistics [25]. Suppose the job execution takes 2 hours in total, and then $d_{it'}$ is $2 \times \frac{60}{5}$ time slots.

Let $w_{it'}$ be the target *soft* deadline to complete job $J_{it'}$, such that its $d_{it'}$ time slots for job execution can be scheduled at any time slots from t' to $w_{it'}$ to yield a value of $b_{it'}$ for the tenant without utility losses. If the job is not completely by the deadline, a penalty function is applied to its value:

$$g_{it'}(\sigma_{it'}) = \begin{cases} g_{it'_+}(\sigma_{it'}) & \text{if } w_{it'} + \sigma_{it'} \in [w_{it'}, T] \\ +\infty & \text{otherwise} \end{cases} \quad (1)$$

where $\sigma_{it'}$ is the number of slots by which the deadline is violated, $g_{it'_+}(\sigma_{it'})$ is a non-decreasing function with $g_{it'_+}(0) = 0$, and $w_{it'} + \sigma_{it}$ represents the actual finishing time if the job execution goes beyond $w_{it'}$. An example penalty function is given in (19) in the performance evaluation section VI. Tenant

[1]Interactive workloads, *e.g.*, web search, cannot be delayed across multiple time slots due to latency requirement, and hence they are beyond the scope of our model for tenants' scheduling decisions.

TABLE I
NOTATION

Var	Definition
N	# of tenants
T	# of time slots
$J_{it'}$	arrival job of tenant i at time slot t'
$b_{it'}$	value of job $J_{it'}$
$d_{it'}$	# of required execution time slots of job $J_{it'}$
$q_{it'}$	per-slot energy consumption of job $J_{it'}$
$w_{it'}$	(soft) deadline of job $J_{it'}$
$\sigma_{it'}$	# of time slots job $J_{it'}$ completes beyond $w_{it'}$
$g_{it'}(\cdot)$	penalty function of job $J_{it'}$
$b_{it'l}$	value of job $J_{it'}$ under schedule l
$E_i^{(t)}$	energy capacity of tenant i in time slot t
$x_{it'}^{(t)}$	job $J_{it'}$ is executed in time slot t or not
$y_{it'}$	job $J_{it'}$ is accepted upon arrival in t' or not
$x_{it'l}$	job $J_{it'}$ is accepted by schedule l or not
$v_i^{(t)}(\cdot)$	cost valuation function of tenant i for time slot t
$c_i^{(t)}$	EDR energy reduction solicited from tenant i in t
$p_i^{(t)}(\cdot)$	energy cost function of tenant i for time slot t
$R^{(t)}$	EDR energy reduction demand from the grid for t
α	diesel cost per unit energy production
$\beta^{(t)}$	# of time slots the EDR frame starting from t contains
$z^{(t)}$	amount of diesel-generated energy in time slot t
$\eta_i^{(t)}$	total reserved energy usage at tenant i for time slot t
L_i	scaled minimum per-slot per-unit-energy-consumption job value
U_i	maximum per-unit-energy-consumption job value

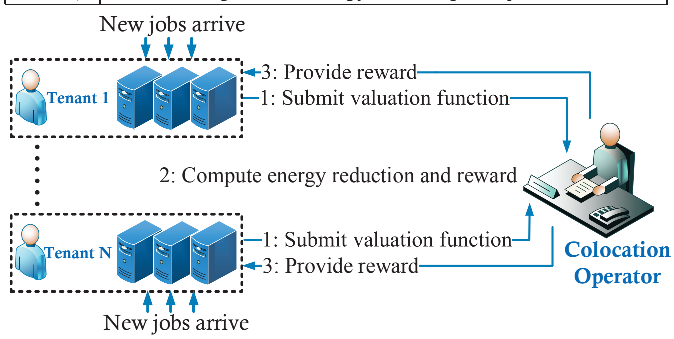

Fig. 2. An illustration of colocation EDR.

i's actual value for completing job $J_{it'}$ is $b_{it'} - g_{it'}(\sigma_{it'})$ in general. In summary, a job is characterized as:

$$J_{it'} = \{w_{it'}, d_{it'}, q_{it'}, b_{it'}, g_{it'}(\cdot)\}. \quad (2)$$

Let $E_i^{(t)}$ be tenant i's total energy capacity at t, based on a power contract between the tenant and the colocation upon server deployment (*i.e.*, power multiplies the duration of a time slot). Given the energy limitation, tenants carefully schedule their job execution, and may defer some jobs for energy shedding during EDR events, by turning idle servers into low power states. The job model is applicable for a large variety of practical workloads, such as those run on Amazon spot instances that may get interrupted at runtime and MapReduce data processing.

The colocation operator's demand for energy reduction by each participating tenant can be viewed as special jobs, which consume energy at the tenant for fixed time slots and can bring value to the tenant through monetary rewards. When an EDR signal arrives from the grid, the colocation invites bids from tenants. Each participating tenant decides and bids its cost valuation function (due to fewer jobs processed or deadline penalty incurred) in the upcoming EDR time frame $\hat{t}$, as denoted by $v_i^{(t)}(\cdot)$ for tenant i in $t \in \hat{t}$. The valuation function depends on the job runtime states that are also implicitly relevant to scheduling decisions over the previous time slots. After collecting bids from the tenants, the colocation carries out global cost minimization by optimally setting the amounts of energy to be reduced by each tenant ($c_i^{(t)}$, for time slot t in the EDR frame) and generated by diesel ($z^{(t)}$, for time slot t in the EDR frame), based on the cost valuations; it further computes the rewards to tenants ($r_i^{(t)}$, for each time slot t when tenant i sheds energy). The EDR work flow is illustrated in **Fig.** 2.

We next formulate the job scheduling and energy reduction allocation problems at each tenant and at the colocation operator, respectively.

B. Tenant Utility Maximization

Through an online optimization algorithm, each tenant schedules jobs and responds to EDR calls, in order to maximize its utility over the entire system span $[1, T]$. Let $x_{it'}^{(t)}$ indicate whether the job $J_{it'}$ is executed in time slot t ($x_{it'}^{(t)} = 1$) or not (0), and $y_{it'}$ indicate whether job $J_{it'}$ is accepted ($y_{it'} = 1$) or not (0). The offline utility maximization problem of tenant i can be formulated as follows. We do not explicitly include EDR energy reduction in the formulation, but will show later it can be seamlessly incorporated in our online algorithm.

Tenant i's utility maximization problem (without EDR):

$$\text{maximize} \quad \sum_{t' \in [1,T]} (b_{it'} - g_{it'}(\sigma_{it'}))y_{it'} \tag{3}$$

subject to:

$$\sum_{t' \in [1,t]} q_{it'} x_{it'}^{(t)} \le E_i^{(t)}, \ \forall t \in [1, T] \tag{3a}$$

$$t \cdot x_{it'}^{(t)} \le w_{it'} + \sigma_{it'}, \ \forall t' \in [1, T], \forall t \in [t', T] \tag{3b}$$

$$\sum_{t \in [t', T]} x_{it'}^{(t)} \ge d_{it'} y_{it'}, \ \forall t' \in [1, T] \tag{3c}$$

$$x_{it'}^{(t)} \in \{0, 1\}, \ \ y_{it'} \in \{0, 1\}, \ \ \sigma_{it'} \in [0, T], \ \forall t' \in [1, T], \forall t \in [1, T] \tag{3d}$$

Constraint (3a) models the energy capacity limitation of tenant i. (3b) ensures that each job $J_{it'}$ is executed in between its arrival time t' and $w_{it'} + \sigma_{it'}$, while (3c) ensures that at least $d_{it'}$ time slots are scheduled for an accepted job. The tenant utility maximization problem is an integer problem with mixed packing (3a) and covering (3c) constraints, as well as nonconventional constraints for deadline modeling (3b). Even in the offline setting, it is difficult to find an efficient algorithm to tackle the problem.

To design an efficient online algorithm for tenant job scheduling, we reformulate the problem into a *compact-exponential* integer problem with a nice packing structure, such that it can be treated by the primal-dual online optimization framework [26]. Consider a feasible schedule l of job $J_{it'}$ to be the vector $l = (\{x_{it'}^{(t)}\}_{\forall t \in [1,T]}, y_{it'}, \sigma_{it'})$, which satisfies constraints (3b), (3c) and (3d), and $\mathcal{L}_{it'}$ represents the set of feasible schedules of the job. Let $\mathcal{T}(l)$ be the set of time slots when job $J_{it'}$ is executed according to $l \in \mathcal{L}_{it'}$, *i.e.*, $\mathcal{T}(l) = \{t : t \in [t', T], x_{it'}^{(t)} = 1 \text{ in } l\}$. Let $x_{it'l}$ indicate whether the schedule $l \in \mathcal{L}_{it'}$ is selected (1) or not (0), and $b_{it'l}$ represent the valuation of job $J_{it'}$ following schedule l, *i.e.*, $b_{it'l} = b_{it'} - g_{it'}(\max_{t \in \mathcal{T}(l)} t - w_{it'})$. Problem (3) can be reformulated as follows:

Compact tenant utility maximization problem (without EDR):

$$\text{maximize} \quad \sum_{t' \in [1,T]} \sum_{l \in \mathcal{L}_{it'}} b_{it'l} x_{it'l} \tag{4}$$

subject to:

$$\sum_{l \in \mathcal{L}_{it'}} x_{it'l} \le 1, \ \forall t' \in [1, T] \tag{4a}$$

$$\sum_{t' \in [1,t]} \sum_{l \in \mathcal{L}_{it'} : t \in \mathcal{T}(l)} q_{it'} x_{it'l} \le E_i^{(t)}, \ \forall t \in [1, T] \tag{4b}$$

$$x_{it'l} \in \{0, 1\}, \ \forall t' \in [1, T], \forall l \in \mathcal{L}_{it'} \tag{4c}$$

Through the reformulation, the original constraints (3b), (3c) and (3d) are simplified into (4a) and (4c), and (3a) is equivalent to (4b). A feasible solution to (3) has a corresponding feasible solution in (4) and vice versa, and hence the optimal objective value of (4) is equal to that of (3). The price of the simplification is that the reformulated problem involves an exponential number of variables, as the number of possible feasible schedules for tenant i's job $J_{it'}$ is exponential in size. Nonetheless, we will show in Sec. IV that an efficient primal-dual online algorithm can update only a polynomial number of variables to solve (4) in polynomial time.

For primal-dual algorithm design, we formulate the dual of (4) by relaxing the integrality constraints and introducing dual variables $u_{it'}$ and $p_i^{(t)}$ to constraints (4a) and (4b), respectively:

$$\text{minimize} \quad \sum_{t' \in [1,T]} u_{it'} + \sum_{t \in [1,T]} E_i^{(t)} p_i^{(t)} \tag{5}$$

subject to:

$$u_{it'} \ge b_{it'l} - \sum_{t \in \mathcal{T}(l)} q_{it'} p_i^{(t)}, \ \forall t' \in [1, T], \forall l \in \mathcal{L}_{it'} \tag{5a}$$

$$u_{it'} \ge 0, \ \ p_i^{(t)} \ge 0, \ \forall t' \in [1, T], \forall t \in [1, T] \tag{5b}$$

C. Colocation Cost Minimization

The colocation operator aims to achieve EDR energy reduction with a minimal colocation-wide social cost in each EDR frame. At each time slot t, the total cost consists of *operator's cost*, its rewards to tenants plus the diesel cost ($\sum_{i \in [1,N]} r_i^{(t)} + \alpha z^{(t)}$), and *tenants' cost*, with each tenant's cost being its energy reduction cost valuation minus received reward ($v_i^{(t)}(c_i^{(t)}) - r_i^{(t)}$). Thus, the total cost equals $\sum_{i \in [1,N]} v_i^{(t)}(c_i^{(t)}) + \alpha z^{(t)}$, with the reward transfer cancelled. The cost minimization problem of the colocation operator in each time slot t of an EDR frame can be formulated as follows:

Operator's cost minimization problem:

$$\text{minimize} \quad \sum_{i \in [1,N]} v_i^{(t)}(c_i^{(t)}) + \alpha z^{(t)} \tag{6}$$

subject to:

$$\sum_{i\in[1,N]} c_i^{(t)} + z^{(t)} \geq R^{(t)}, \tag{6a}$$

$$z^{(t)} \geq 0,\; c_i^{(t)} \geq 0,\; \forall i \in [1,N] \tag{6b}$$

Constraint (6a) ensures that tenants' energy reduction and diesel generation cover the EDR demand. The operator's decisions, energy reduction $c_i^{(t)}$ by tenant i and the amount of diesel generation $z^{(t)}$, are decided based on tenants' self-reported valuation functions $v_i^{(t)}(\cdot)$. The decisions in t are indirectly coupled with decisions in other time slots through tenants' job scheduling.

D. Social Welfare Maximization

Through designing efficient algorithms for tenant job scheduling and operator energy reduction allocation respectively, we seek to further achieve social welfare maximization. The social welfare in the colocation is the overall utility of tenants and the colocation operator. Tenant i's utility is the sum of valuations of its accepted jobs and the EDR rewards received, $\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l}+\sum_{t\in[1,T]} r_i^{(t)}$, when the EDR energy reduction and rewards are considered. The colocation operator's utility is its reward from the grid operator according to energy shedding amounts during EDR, fixed according to the pre-signed contract, minus its cost (rewards to tenants and diesel costs). Ignoring the constant reward that the operator receives from the grid,[2] the social welfare can be described by overall tenants' utility minus operator's overall cost, $(\sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} + \sum_{i\in[1,N]}\sum_{t\in[1,T]} r_i^{(t)}) - (\sum_{t\in[1,T]}\sum_{i\in[1,N]} r_i^{(t)} + \sum_{t\in[1,T]} \alpha z^{(t)}) = \sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} - \sum_{t\in[1,T]}\alpha z^{(t)}$. The offline social welfare maximization problem can be formulated as follows:

Global social welfare maximization problem:

$$\text{maximize} \quad \sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} - \sum_{t\in[1,T]} \alpha z^{(t)} \tag{7}$$

subject to:

$$\sum_{i\in[1,N]} c_i^{(t)} + z^{(t)} \geq R^{(t)},\; \forall t \in [1,T] \tag{7a}$$

$$\sum_{l\in\mathcal{L}_{it'}} x_{it'l} \leq 1,\; \forall i\in[1,N], t'\in[1,T] \tag{7b}$$

$$\sum_{t'\in[1,t]}\sum_{l\in\mathcal{L}_{it'}:t\in\mathcal{T}(l)} q_{it'}x_{it'l} + c_i^{(t)} \leq E_i^{(t)}, \quad \forall i\in[1,N], t\in[1,T] \tag{7c}$$

$$x_{it'l}\in\{0,1\},\; z^{(t)}\geq 0,\; c_i^{(t)}\geq 0, \quad \forall i\in[1,N], t\in[1,T], t'\in[1,T], l\in\mathcal{L}_{it'} \tag{7d}$$

The social welfare maximization problem includes tenant job scheduling decisions and operator's EDR energy reduction allocation decisions, with constraints (7b) and (7c) corresponding to (4a) and (4b) in tenant's utility maximization problem, and constraint (7a) equivalent to (6a) in operator's cost minimizing problem. In (7), energy reduction due to EDR is counted, such that the sum of energy consumption for running accepted jobs and EDR energy shedding does not exceed a tenant's contracted energy capacity. Note that no entity carries out algorithms to solve this social welfare maximization problem with complete knowledge of the entire system, which is formulated as a benchmark for performance analysis of the tenants' and operator's algorithms—we will show that aggregately, the online algorithms carried out by each tenant and the operator, towards their respective objectives, maximize social welfare as well.

[2]We will prove in Lemma 2 that the reward to tenants from the operator is always no more than the corresponding cost by using diesel. Hence, the reward from the operator to tenants is no more than the overall reward from the grid to the operator, and we can safely ignore the latter reward here.

We formulate the dual of (7) by relaxing integrality constraints, and introducing dual variable $\lambda^{(t)}$ for (7a) and the same dual variables as in dual of (4) for (7b) and (7c):

$$\text{minimize} \quad -\sum_{t\in[1,T]} R^{(t)}\lambda^{(t)} + \sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]} E_i^{(t)}p_i^{(t)} \tag{8}$$

subject to:

$$u_{it'} \geq b_{it'l} - \sum_{t\in\mathcal{T}(l)} q_{it'}p_i^{(t)},\; \forall i\in[1,N], t'\in[1,T], l\in\mathcal{L}_{it'} \tag{8a}$$

$$\lambda^{(t)} \leq \alpha,\; \forall t\in[1,T] \tag{8b}$$

$$p_i^{(t)} - \lambda^{(t)} \geq 0,\; \forall i\in[1,N], t\in[1,T] \tag{8c}$$

$$\lambda^{(t)}\geq 0,\; u_{it'}\geq 0,\; p_i^{(t)}\geq 0,\; \forall t,t'\in[1,T], i\in[1,N] \tag{8d}$$

IV. Online Algorithms for Tenant Utility Maximization and EDR Incentivization

We next design an online algorithm for each tenant's job scheduling and energy reduction bidding, and an auction mechanism for energy reduction incentivization at the colocation.

A. Online Algorithm for Tenants

We first design tenant's online algorithm considering job scheduling only, then show later how cost valuation and energy shedding during EDR fit in.

In the online setting, as jobs arrive at each tenant i, the variables and constraints of (4) emerge gradually. For example, upon arrival of job $J_{it'}$, there is a set of new primal variables, $x_{it'l}, \forall l\in\mathcal{L}_{it'}$, subject to $\sum_{l\in\mathcal{L}_{it'}} x_{it'l} \leq 1$, with $\mathcal{L}_{it'}$ identified based on the job's specification. The tenant must decide immediately whether to serve the job and if so, by which schedule $l\in\mathcal{L}_{it'}$, considering its energy capacity and existing jobs. If the tenant decides to execute this job ($J_{it'}$) following schedule l, then let $x_{it'l}=1$, and increase the amount of energy reservation by $q_{it'}^{(t)}$ in all time slots t scheduled to run the job according to l, *i.e.*, $\forall t\in\mathcal{T}(l)$. Otherwise, $x_{it'l}$ will be zero for all $l\in\mathcal{L}_{it'}$.

Job Scheduling. Towards optimally scheduling a new job over future time slots, we resort to the dual problem of (4) in (5), and apply the KKT conditions [27]. Corresponding to job $J_{it'}$, there is a dual variable $u_{it'}\geq 0$ subject to constraints (5a), *i.e.*, $u_{it'} \geq b_{it'l} - \sum_{t\in\mathcal{T}(l)} q_{it'}p_i^{(t)}$, $\forall l\in\mathcal{L}_{it'}$. The KKT conditions indicate that in the optimal solutions of (4) and (5), $x_{it'l}$ must be zero unless constraint (5a) is tight for schedule l. Thus, we let $u_{it'}$ be the maximum between 0 and the right hand side (RHS) of constraint (5a), $\forall l\in\mathcal{L}_{it'}$:

$$u_{it'} = \max\Big\{0, \max_{l\in\mathcal{L}_{it'}}\{b_{it'l} - \sum_{t\in\mathcal{T}(l)} q_{it'}p_i^{(t)}\}\Big\} \tag{9}$$

We decide whether to accept job $J_{it'}$ accordingly: if the RHS of (5a) is non-positive, then $u_{it'} = 0$ and the job is rejected, *i.e.*, $x_{it'l} = 0, \forall l \in \mathcal{L}_{it'}$; otherwise, $u_{it'} > 0$ and the job will be scheduled by the schedule l_i that maximizes the RHS of (5a) (breaking ties arbitrarily), *i.e.*, $x_{it'l_i} = 1$ and $x_{it'l} = 0, \forall l \neq l_i$.

The rationale is as follows. We interpret the dual variable $p_i^{(t)}$ as a marginal cost per unit energy consumption at tenant i in t, for deciding whether the new job has a high enough value to consume the required amounts of energy in its lifespan. The second term in the RHS of (5a) is the total cost of running the job by schedule l. The RHS of (5a) can be viewed as the net gain of job $J_{it'}$ if served by schedule l. The above method effectively serves a job by the schedule that maximizes its net gain, and jobs whose values are lower than incurred costs will be rejected.

Cost Function Design. Comparing the offline optimum and the online decisions, one main reason causing tenant utility loss is that low value jobs arriving earlier are accepted, while high value jobs cannot be accepted later due to insufficient remaining energy capacity. To prevent this, the marginal energy cost $p_i^{(t)}$ should be increased over time as the amount of reserved energy for t increases: when the energy surplus is abundant, the price is relative low for accepting all possible jobs; when remaining energy is scarce, the price should be significantly high so that only high value jobs are accepted.

Let $\eta_i^{(t)}$ denote energy usage in time slot t at tenant i, reserved for accepted jobs. $p_i^{(t)}$ should be a monotonically increasing function of $\eta_i^{(t)}$. Let

$$L_i = \frac{1}{2F} \min_{t' \in [1,T], l \in \mathcal{L}_{it'}} \{\frac{b_{it'l}}{q_{it'} d_{it'}}\} \tag{10}$$

which represents the minimum per-slot per-unit-energy-consumption job value among all jobs of tenant i over all possible schedules, scaled down by factor $\frac{1}{2F}$. Here $\frac{1}{F}$ is a lower bound of the ratio of a tenant's overall energy usage over its overall energy capacity, for all tenants, *i.e.*, $\frac{1}{F} \leq \frac{\sum_{t \in [1,T]} \eta_i^{(t)}}{\sum_{t \in [1,T]} E_i^{(t)}}, \forall i \in [1, N]$. Its value can be estimated based on historical energy utilization data at the tenants. Let

$$U_i = \max_{t' \in [1,T], l \in \mathcal{L}_{it'}} \{\frac{b_{it'l}}{q_{it'}}\}, \tag{11}$$

which denotes the maximum per-unit-energy-consumption job value among all jobs of tenant i over all possible schedules. We design the cost function as follows:

$$p_i^{(t)}(\eta_i^{(t)}) = L_i (\frac{U_i}{L_i})^{\frac{\eta_i^{(t)}}{E_i^{(t)}}} \tag{12}$$

When there is no energy usage in t yet, *i.e.*, $\eta_i^{(t)} = 0$, we have $p_i^{(t)}(0) = L_i$. The value of L_i is carefully designed to ensure that the initial marginal cost, $p_i^{(t)}(0)$, is low enough such that any job can be scheduled to run in this time slot. When the energy usage reaches the capacity of tenant i, *i.e.*, $\eta_i^{(t)} = E_i^{(t)}$, we have $p_i^{(t)}(E_i^{(t)}) = U_i$. The value of U_i ensures that any schedule using this time slot will lead to a negative net gain $u_{it'}$, hence no more job can be executed in t. $p_i^{(t)}(\eta_i^{(t)})$ grows exponentially with $\eta_i^{(t)}$ such that the scarcer the energy is, the higher the admission value becomes. The design of the cost function leads to tenant utility and social welfare maximization, which will be shown in Sec. V.

Energy Shedding Bid. When participating in the EDR auction, tenant i reports its energy cost valuation function $v_i^{(t)}(\cdot)$ for each time slot t in the upcoming EDR frame, as follows (assume truthful bidding):

$$v_i^{(t)}(c_i^{(t)}) = \begin{cases} \int_{\eta_i^{(t)}}^{\eta_i^{(t)}+c_i^{(t)}} L_i (\frac{U_i}{L_i})^{\frac{\eta}{E_i^{(t)}}} \mathrm{d}\eta, & 0 \leq c_i^{(t)} \leq E_i^{(t)} - \eta_i^{(t)} \\ +\infty & \text{otherwise} \end{cases} \tag{13}$$

As long as the energy shedding requested by the colocation operator for t ($c_i^{(t)}$) can be fulfilled by the tenant's remaining energy in t, the cost valuation for energy reduction (namely the minimum reward the operator should pay for energy reduction in this time slot) is computed as the aggregate cost of the energy reserved for EDR in t. Revealing the cost function in this form, the tenant ensures that any energy shedding demand with an overall reward no lower than the aggregate cost valuation, $\sum_{t \in \hat{t}} v_i^{(t)}(c_i^{(t)})$, is acceptable and will not harm the tenant's utility, as compared to not participating in EDR.

Tenant's online algorithm. Alg. 1 shows our online algorithm for job scheduling and EDR bidding at the tenants. Upon each job arrival, the tenant computes the best schedule and the corresponding largest net gain of the job (RHS of (5a)) (line 4), by calling a dual oracle in **Alg.** 2. Then $u_{it'}$ is decided according to (9) (line 5). If $u_{it'}$ is positive, the job is accepted, job scheduling decisions are set according to the best schedule (lines 6-9), and energy reservation and energy costs are updated in the time slots in which the job is executed (lines 10-12). Upon receiving bid invitation from the operator, the tenant bids cost valuation functions (lines 20-21), and then updates energy usage and energy costs in each time slot of the EDR frame, according to the solicited amount of energy shedding from the operator (lines 22-28).

The dual oracle in **Alg.** 2 computes in polynomial time the best schedule in $\mathcal{L}_{it'}$ achieving maximal net gain of job $J_{it'}$; since $\mathcal{L}_{it'}$ includes an exponential number of schedules, a naive exhaustive search is infeasible. Given a job completion time t_{end} within $[w_{it'}, T]$, we compute the best schedule l by which the job is executed in the $d_{it'}$ time slots no later than t_{end} with the lowest costs (lines 4-8). We go through different t_{end}'s and find the best schedule achieving the largest net gain (lines 3, 9-11). In this way, we do not need to find out all l in $\mathcal{L}_{it'}$, but can directly compute the best schedule for each t_{end}; the overall number of schedules we compare is $T - w_{it'} + 1$. Polynomial running time of the tenant's algorithm is proven in Theorem 2.

Note that the algorithm needs U_i and L_i as input, whose exact values are not known before all jobs have arrived. We instead adopt estimated lower and upper bounds as input, *e.g.*, based on past experience. We will show the impact of inaccurate estimation in the simulations.

B. EDR Incentivization Mechanism at Colocation Operator

Allocation of Energy Reduction Amounts to Tenants. The operator solves the cost minimization problem in (6) to decide the amount of energy for each participating tenant to reduce

Algorithm 1 Tenant i's Online Algorithm, $\mathsf{A}_{\mathsf{tenant}}$

Input: $E_i^{(t)}, U_i, L_i, \forall t \in [1,T]$
Output: $x_{it'}^{(t)}, y_{it'}, \forall t' \in [1,T], t \in [1,T]$
Initialize:
$x_{it'}^{(t)} = 0, y_{it'} = 0, \eta_i^{(t)} = 0, p_i^{(t)} = p_i^{(t)}(0), \forall t' \in [1,T], \forall t \in [1,T]$;
Upon arrival of new job $J_{it'}$
  $(o_{it'}^*, l_{it'}^*) = \mathsf{Oracle}(J_{it'}, \{p_i^{(t)}\}_{\forall t \in [t',T]})$;
  $u_{it'} = \max\{0, o_{it'}^*\}$;
  if $u_{it'} > 0$ **then**
    // Update primal variables and reserved energy;
    $x_{it' l_{it'}^*} = 1$;
    $y_{it'} = 1, x_{it'}^{(t)} = 1, \forall t \in T(l_{it'}^*)$;
    $\eta_i^{(t)} = \eta_i^{(t)} + q_{it'}, \quad \forall t \in T(l_{it'}^*)$;
    // Update dual variables
    $p_i^{(t)} = p_i^{(t)}(\eta_i^{(t)}), \quad \forall t \in T(l_{it'}^*)$;
  end if
  if $y_{it'} = 1$ **then**
    Accept job $J_{it'}$ and execute it in time slots t specified in schedule $l_{it'}^*$ $(x_{it'}^{(t)} = 1)$;
  else
    Reject job $J_{it'}$;
  end if
End Upon
Upon invitation of energy shedding bids from colocation operator at beginning of EDR frame $\hat{t}$
  Bid cost valuation functions $v_i^{(t)}(\cdot), \forall t \in \hat{t}$ according to (13);
  for all solicited energy reduction from operator, $c_i^{(t)}, \forall t \in \hat{t}$ **do**
    if $c_i^{(t)} > 0$ **then**
      $\eta_i^{(t)} = \eta_i^{(t)} + c_i^{(t)}$;
      $p_i^{(t)} = p_i^{(t)}(\eta_i^{(t)})$;
    end if
  end for
  Carry out energy reduction in time slots t where $c_i^{(t)} > 0$;
End Upon

Algorithm 2 Dual Oracle for Tenant Job Scheduling, $\mathsf{Oracle}(\cdot)$

Input: $J_{it'}, \{p_i^{(t)}\}_{\forall t \in [t',T]}$
Output: $o_{it'}^*, l_{it'}^*$
$o_{max} = 0$;
$l_{max} = \emptyset$;
for all $t_{end} \in [w_{it'}, T]$ **do**
  $\Gamma = t_{end} - t' + 1$;
  // Sort time slots by $p_i^{(t)}$ in non-decreasing order.
  $(p_i^{\bar{t}_1}, p_i^{\bar{t}_2}, \ldots, p_i^{\bar{t}_\Gamma}) = \mathsf{sort}\left(p_i^{(t')}, p_i^{(t'+1)}, \ldots, p_i^{(t_{end})}\right)$;
  $l = (\bar{t}_1, \bar{t}_2, \ldots, \bar{t}_{d_{it'}})$;
  $o = b_{it'} - g_{it'}\left(\max_{j \in [1, d_{it'}]}\{\bar{t}_j\} - w_{it'}\right) - \sum_{j=1}^{d_{it'}} q_{it'} p_i^{\bar{t}_j}$;
  if $(o > o_{max})$ **then**
    $o_{max} = o$;
    $l_{max} = l$;
  end if
end for
$o_{it'}^* = o_{max}$;
$l_{it'}^* = l_{max}$;

$(c_i^{(t)})$ and the diesel usage $(z^{(t)})$ in each time slot t of the EDR frame, to fulfill the EDR energy shedding demand from the grid. As tenants' cost valuation functions for energy reduction, given in (13), are convex (proven in Lemma 1), (6) is a convex program and can be solved exactly in polynomial time, using an efficient algorithm such as an interior-point method.

Reward for Energy Reduction. For each $t \in \hat{t}$, let $M^{(t)}$ be the optimal objective value of $\sum_{\ddot{i} \in [1,N]} v_{\ddot{i}}^{(t)}(c_{\ddot{i}}^{(t)}) + \alpha z^{(t)}$, obtained by solving (6) over all tenants in $[1,N]$. Let $M_{-i}^{(t)}$ be the optimal objective value of $\sum_{\ddot{i} \in [1,N] \setminus i} v_{\ddot{i}}^{(t)}(c_{\ddot{i}}^{(t)}) + \alpha z^{(t)}$, computed by solving (6) assuming tenant i is absent from the auction. For each winning tenant i with $c_i^{(t)}$, the reward is computed by

$$r_i^{(t)} = (M_{-i}^{(t)} - M^{(t)}) + v_i^{(t)}(c_i^{(t)})$$

i.e., the sum of tenant i's marginal contribution to reducing cost in the system and its reported valuation for $c_i^{(t)}$. We always have $M_{-i}^{(t)} \geq M^{(t)}$ (more options leads to lower objective value in (6)), and hence the reward for reducing energy $c_i^{(t)}$ is always no smaller than the respective cost valuation, ensuring acceptance of the energy shedding amount at tenant i.

The EDR auction mechanism carried out by the operator is summarized in **Alg.** 3.

V. Theoretical Analysis

A. Analysis of Operator's EDR Auction Mechanism

Lemma 1. *The cost valuation function $v_i^{(t)}(\cdot)$ is convex for each tenant i and each time slot t.*

Proof. $v_i^{(t)}(c_i^{(t)})$ is the integral of $p_i^{(t)}(\eta_i^{(t)})$ (*i.e.*, (12)) if $0 \leq c_i^{(t)} \leq (E_i^{(t)} - \eta_i^{(t)})$. As $p_i^{(t)}(\eta_i^{(t)})$ is an exponential function on $\eta_i^{(t)}$ and its base, $(\frac{U_i}{L_i})^{\frac{1}{E_i^{(t)}}}$, is no smaller than 1, it is a convex function. The integral of a convex function is convex too [27]. Due to the convexity of extended-value extension [27], (*i.e.*, extending the domain of a convex function to $\mathbb{R}$ by defining its value to be $+\infty$ outside its domain), the function $v_i^{(t)}(c_i^{(t)})$ is still convex. □

Theorem 1. *The EDR auction mechanism* $\mathsf{A}_{\mathsf{operator}}$ *in **Alg.** 3, is truthful and individually rational, runs in polynomial time, and minimizes the total cost in each time slot of an EDR frame.*

Proof. (*Truthfulness*) An auction is truthful if bidding truth cost valuation functions as in (13) maximizes a tenant's net gain in the auction, *i.e.*, $\hat{u}_i^{(t)} = r_i^{(t)} - v_i^{(t)}(c_i^{*(t)})$. Without assuming truthful bidding, let $\hat{v}_i^{(t)}(\cdot)$ denote the reported cost valuation function from tenant i. We have

$$\begin{aligned} \hat{u}_i^{(t)} &= (M_{-i}^{(t)} - M^{(t)}) + \hat{v}_i^{(t)}(c_i^{*(t)}) - v_i^{(t)}(c_i^{*(t)}) \\ &= M_{-i}^{(t)} - \left(M^{(t)} - \hat{v}_i^{(t)}(c_i^{*(t)}) + v_i^{(t)}(c_i^{*(t)})\right) \end{aligned} \quad (14)$$

Algorithm 3 Colocation Operator's EDR Auction Mechanism, $\mathsf{A}_{\mathsf{operator}}$

```
Upon arrival of EDR signal R^(t̂) from the grid
    // Compute EDR Reductions
    Invite tenant bids v_i^(t)(·), ∀i ∈ [1, N], t ∈ t̂;
    for all t ∈ t̂ do
        Solve (6) by interior-point method; let c_i^{*(t)}, z^{*(t)}, ∀i ∈ [1, N], be the optimal solution;
        // Compute rewards
        M^(t) = Σ_{i∈[1,N]} v_i^(t)(c_i^{*(t)}) + α z^{*(t)};
        for all i ∈ [1, N] do
            r_i^(t) = 0;
            if c_i^{*(t)} > 0 then
                for all ï ∈ [1, N] do
                    v̇_ï^(t)(·) = { +∞, if ï = i ; v_ï^(t)(·), otherwise };
                end for
                Solve (6) using interior point method, where v_i^(t)(·) is replaced by v̇_i^(t)(·), ∀i ∈ [1, N], and let ċ_i^{*(t)}, ż^{*(t)}, ∀i ∈ [N], denote the optimal solution;
                M_{-i}^(t) = Σ_{ï∈[1,N]\i} v_ï^(t)(ċ_ï^{*(t)}) + α ż^{*(t)};
                r_i^(t) = (M_{-i}^(t) − M^(t)) + v_i^(t)(c_i^{*(t)});
            end if
        end for
    end for
    Send energy reduction amounts c_i^{*(t)}, ∀t ∈ t̂, to all tenants i ∈ [1, N];
    Produce energy at the amount of z^{*(t)} using the diesel generator in each time slot t in the EDR frame t̂;
End
```

As $M_{-i}^{(t)}$ is not related to tenant i's bid, we will show $M^{(t)} - \hat{v}_i^{(t)}(c_i^{*(t)}) + v_i^{(t)}(c_i^{*(t)})$ is minimized by truthful bidding. We have

$$
\begin{aligned}
&M^{(t)} - \hat{v}_i^{(t)}(c_i^{*(t)}) + v_i^{(t)}(c_i^{*(t)}) \\
&\quad = \Big(\sum_{j\in[1,N]:j\neq i} \hat{v}_j^{(t)}(c_j^{*(t)}) + \alpha z^{*(t)}\Big) + v_i^{(t)}(c_i^{*(t)})
\end{aligned}
\tag{15}
$$

Suppose all other tenants' reported valuation functions are fixed ($\hat{v}_j^{(t)}(\cdot), \forall j \in [1,N], j \neq i$). When i truthfully reports its valuation function (*i.e.*, $\hat{v}_i^{(t)}(\cdot) = v_i^{(t)}(\cdot)$), the objective function of (6) is identical with (15), and the optimal solution $(c_i^{*(t)}, \forall i \in [1,N]\}, z^{*(t)})$ of (6) minimizes (15) as well (constraints (6a) (6b) still apply). When i reports $\hat{v}_i^{(t)}(\cdot) \neq v_i^{(t)}(\cdot)$, (15) is different from the objective function of (6) which is now $\sum_{j\in[1,N]:j\neq i} \hat{v}_j^{(t)}(c_j^{(t)}) + \alpha z^{(t)} + \hat{v}_i^{(t)}(c_i^{(t)})$. Therefore, the optimal solution $\{\bar{c}_i^{*(t)}, \forall i \in [1,N]\}, \bar{z}^{*(t)}$ of (6) may not minimize (15), and the value of (15) computed at $\{\bar{c}_i^{*(t)}, \forall i \in [1,N]\}, \bar{z}^{*(t)}$ is no smaller than the value of (15) computed at $\{c_i^{*(t)}, \forall i \in [1,N]\}, z^{*(t)}$.

(*Individual rationality*) By definition, the EDR auction satisfies individual rationality if a winning tenant's reward is always no less than its cost valuation for energy reduction. We have $r_i^{(t)} - v_i^{(t)}(c_i^{*(t)}) = (M_{-i}^{(t)} - M^{(t)}) + v_i^{(t)}(c_i^{*(t)}) - v_i^{(t)}(c_i^{*(t)}) = M_{-i}^{(t)} - M^{(t)}$. As $M^{(t)}$ is computed over more options than $M_{-i}^{(t)}$ and hence no larger than $M_{-i}^{(t)}$, we have $M_{-i}^{(t)} - M^{(t)} \geq 0$.

(*Polynomial time and total cost minimization*) By Lemma 1, (6) is a convex problem that can be solved in polynomial time by the interior point method, achieving minimal total cost. □

Lemma 2. *The optimal solution of the operator's cost minimization problem (6) satisfies $v_i^{(t)'}(c_i^{(t)}) \leq \alpha, \forall i \in [1,N], \forall t \in [1,T]$ (derivative of the cost valuation function is at most unit energy cost from diesel), and $v_i^{(t)}(c_i^{*(t)}) \leq r_i^{(t)} \leq \alpha c_i^{*(t)}, \forall i \in [1,N], \forall t \in [1,T]$ (the reward to a tenant is between the corresponding cost valuation and the corresponding diesel cost).*

Proof. We prove $v_i^{(t)'}(c_i^{(t)}) \leq \alpha$ by contradiction. Suppose there exists a pair of (i,t) such that $v_i^{(t)'}(c_i^{*(t)}) > \alpha$. We could then decrease $c_i^{*(t)}$ and increase $z^{*(t)}$ to obtain a smaller objective value, which contradicts the fact that $c_i^{*(t)}$ and $z^{*(t)}$ are optimal.

We have proven $r_i^{(t)} \geq v_i^{(t)}(c_i^{*(t)})$ in Theorem 1 (individual rationality). We prove $r_i^{(t)} \leq \alpha c_i^{*(t)}$ by contradiction. Suppose there exists a pair of $(\bar{i}, t)$ such that $r_{\bar{i}}^{(t)} > \alpha c_{\bar{i}}^{*(t)}$. Then we have $r_{\bar{i}}^{(t)} = M_{-\bar{i}}^{(t)} - (M^{(t)} - v_{\bar{i}}^{(t)}(c_{\bar{i}}^{*(t)})) > \alpha c_{\bar{i}}^{*(t)}$, where $M_{-\bar{i}}^{(t)}$ represents the minimum of (6) when tenant $\bar{i}$ is absent. We can construct another feasible solution $\{\ddot{c}_i^{(t)}, \forall i \in [T]\}, \ddot{z}^{(t)}$ of (6) with tenant $\bar{i}$ being absent as follows: $\ddot{c}_i^{(t)} = c_i^{*(t)}, \forall i \in [1,N], i \neq \bar{i}$ and $\ddot{z}^{(t)} = z^{*(t)} + c_{\bar{i}}^{*(t)}$, *i.e.*, using diesel to produce tenant i's energy $c_{\bar{i}}^{*(t)}$. We have

$$
\begin{aligned}
&M_{-\bar{i}}(t) - (M^{(t)} - v_{\bar{i}}^{(t)}(c_{\bar{i}}^{*(t)})) \leq \Big(\sum_{i\neq\bar{i}} v_i^{(t)}(\ddot{c}_i^{(t)}) + \alpha \ddot{z}^{(t)}\Big) - (M^{(t)} - v_{\bar{i}}^{(t)}(c_{\bar{i}}^{*(t)})) \\
&\leq \Big(\sum_{i\neq\bar{i}} v_i^{(t)}(c_i^{*(t)}) + \alpha(z^{*(t)} + c_{\bar{i}}^{*(t)})\Big) - \Big(\sum_{i\neq\bar{i}} v_i^{(t)}(c_i^{*(t)}) + \alpha z^{*(t)}\Big) \leq \alpha c_{\bar{i}}^{*(t)}
\end{aligned}
$$

which contradicts with our assumption. □

B. Analysis of Tenants' Online Algorithm

(1) Polynomial Running Time

Theorem 2. *Each tenant's online algorithm, $\mathsf{A}_{\mathsf{tenant}}$ in **Alg.** 1 and **Alg.** 2, runs in polynomial time upon new job arrivals.*

Proof. When executing the dual oracle, given t_{end}, finding the best schedule that completes the job no later than t_{end} takes $O(T \log T + T)$, due to sorting $\{p_i^{(t)}\}$ and computing penalty. Looping through all $t_{end} \in [w_{it'}, T]$, the complexity becomes $O(T \cdot (T \log T + T)) = O(T^2 \log T)$. When updating variables in lines 7-12 of **Alg.** 1, each instruction is executed at most T times. Hence the overall time complexity of A_{tenant} is $O(T^2 \log T)$. □

(2) Competitive Ratio We next analyze long-term utility maximization achieved by tenant's online algorithm in the following steps. We first analyze competitive ratio of the online job scheduling algorithm without considering EDR energy shedding, computed as the worst-case upper bound ratio of the tenant utility achieved by the offline solution of (3) to the overall tenant utility achieved by $\mathsf{A}_{\mathsf{tenant}}$ at the end of $[1,T]$.

We then show that participating in EDR energy shedding does not affect tenant's long-term utility guarantee.

Let P_0 and D_0 denote the initial values of the primal and dual problems in (4) and (5), respectively. Let ΔP and ΔD represent the changes in objective values of (4) and (5), before and after a scheduling event happens, *i.e.*, a new job is scheduled. We use OPT to represent the tenant utility achieved by the offline solution of (3).

Lemma 3. *If there exists $A \geq 1$ such that $\Delta P \geq \frac{1}{A}\Delta D$ for all scheduling events, and $C \geq 1$ such that $D_0 \leq \frac{1}{C}OPT$, then A_{tenant} in* **Alg.** *1 is $A \cdot \frac{C}{C-1}$-competitive.*

Proof. We have $P_0 = 0$ and $D_0 = \sum_{t\in[1,T]} E_i^{(t)} L_i$. The overall objective value of the primal problem, P_{final}, is the sum of ΔP's over all the scheduling events, and the overall objective value of the dual problem, D_{final}, is the sum of D_0 and all ΔD's. Since $\Delta P \geq \frac{1}{A}\Delta D$, we have $P_{\text{final}} \geq \frac{1}{A}(D_{\text{final}} - D_0)$. By the weak duality [27], we have $D_{\text{final}} \geq OPT$ and $(D_{\text{final}} - D_0) \geq (1 - \frac{1}{C})OPT$, hence we have $P_{\text{final}} \geq \frac{1}{A}(1-\frac{1}{C})OPT$. The competitive ratio is $\frac{1}{\frac{1}{A}(1-\frac{1}{C})} = A \cdot \frac{C}{C-1}$. □

We next define a relation between cost and energy usage before and after a scheduling event. Let $\eta_i^{(t)}$ and $\eta_i^{(t)+}$ be the reserved energy usage for time t for tenant i, before and after the event.

Definition 1. *The Usage-Cost Relationship with $A \geq 1$ is $p_i^{(t)}(\eta_i^{(t)})(\eta_i^{(t)+} - \eta_i^{(t)}) \geq \frac{1}{A}E_i^{(t)}(p_i^{(t)}(\eta_i^{(t)+}) - p_i^{(t)}(\eta_i^{(t)})), \forall i \in [1,N], \forall t \in [1,T]$.*

The *Usage-Cost Relationship* shows that the cost in time slot t due to the scheduling event is no smaller than $\frac{1}{A}$ of the increase of the term $E_i^{(t)}p_i^{(t)}$ in the dual objective of (5), due to the adjustment of the marginal cost.

Lemma 4. *If the Usage-Cost Relationship holds for a given $A \geq 1$, then we have $\Delta P \geq \frac{1}{A}\Delta D$, considering only job scheduling.*

Proof. If tenant i rejects a new job $B_{it'}$, we have $\Delta P = \Delta D = 0$, and the lemma holds. If the tenant accepts the job by schedule l and reserves corresponding energy usage for it, then $\Delta P = b_{it'l} = u_{it'} + p_i^{(t)}(\eta_i^{(t)})(\eta_i^{(t)+} - \eta_i^{(t)})$ and $\Delta D = u_{it'} + E_i^{(t)}(p_i^{(t)}(\eta_i^{(t)+}) - p_i^{(t)}(\eta_i^{(t)}))$. Since $u_{it'} \geq 0$ and $p_i^{(t)}(\eta_i^{(t)})(\eta_i^{(t)+} - \eta_i^{(t)}) \geq \frac{1}{A}E_i^{(t)}(p_i^{(t)}(\eta_i^{(t)+}) - p_i^{(t)}(\eta_i^{(t)}))$, we have $\Delta P \geq \frac{1}{A}\Delta D$. □

We next find a feasible A for which our cost functions $p_i^{(t)}(\cdot)$ satisfy the respective usage-cost relationship. Without loss of generality, we assume in the following discussions that the per-time-slot energy demand $q_{it'}$ of each job received at i is much smaller compared to the energy capacity of the tenant $E_i^{(t)}$ (considering typically a tenant is handling a large number of jobs), *i.e.*, $\eta_i^{(t)+} - \eta_i^{(t)}$ is small enough. We define a differential version of the usage-cost relationship:

Definition 2. *The Differential Usage-Cost Relationship with $A \geq 1$ is $p_i^{(t)}(\eta)\mathrm{d}\eta \geq \frac{E_i^{(t)}}{A}\mathrm{d}p_i^{(t)}, \forall i \in [1,N], \forall t \in [1,T]$.*

Lemma 5. *The cost function in (12) satisfies the Differential Usage-Cost Relationship for $A = \ln(\frac{U_i}{L_i})$, with L_i and U_i defined in (10) and (11), respectively.*

Proof. According to (12), we have $p_i^{(t)\prime}(\eta) = \frac{\ln(U_i/L_i)}{E_i^{(t)}} L_i(\frac{U_i}{L_i})^{\eta/E_i^{(t)}}$. Since $\mathrm{d}p_i^{(t)} = p_i^{(t)\prime}\mathrm{d}\eta$, the differential Usage-Cost Relationship becomes $L_i(\frac{U_i}{L_i})^{\eta/E_i^{(t)}} \geq \frac{E_i^{(t)}}{A}\frac{\ln(U_i/L_i)}{E_i^{(t)}} L_i(\frac{U_i}{L_i})^{\eta/E_i^{(t)}}$, which holds for $A \geq \ln(U_i/L_i)$. □

Theorem 3. *Without considering EDR energy reduction,* **Alg.** *1 is $2A$-competitive in tenant's utility, where $A = \ln(\frac{U_i}{L_i})$, with L_i and U_i defined in (10) and (11), respectively.*

Proof. The cost functions used in **Alg.** 1 satisfy the Differential Usage-Cost Relationship for $A = \ln(\frac{U_i}{L_i})$ (Lemma 5). Since the energy demand in a job $q_{it'}$ is much smaller compared to capacity $E_i^{(t)}$, we have $\mathrm{d}\eta_i^{(t)} = \eta_i^{(t)+} - \eta_i^{(t)} = q_{it'}$, and then

$$p_i^{(t)}(\eta_i^{(t)+}) - p_i^{(t)}(\eta_i^{(t)}) = p_i^{(t)\prime}(\eta_i^{(t)+} - \eta_i^{(t)}) = \mathrm{d}p_i^{(t)}.$$

So the Differential Usage-Cost Relationship (**Def.** 2) implies that the Usage-Cost Relationship (**Def.** 1) holds for A. According to **Alg.** 1, we have:

$$D_0 = \sum_{t\in[1,T]} E_i^{(t)} p_i^{(t)}(0) = \sum_{t\in[1,T]} E_i^{(t)} \frac{1}{2F} \min_{t'\in[1,T], l\in\mathcal{L}_{it'}} \{\frac{b_{it'l}}{q_{it'}d_{it'}}\}$$

Since $\frac{1}{F} \leq \frac{\sum_{t\in[1,T]}\eta_i^{(t)}}{\sum_{t\in[1,T]}E_i^{(t)}}$ (given when defining L_i in (10)), we know $\frac{\sum_{t\in[1,T]}E_i^{(t)}}{F}$ is the minimum amount of overall energy consumption of tenant i. In addition, $\min_{t'\in[1,T], l\in\mathcal{L}_{it'}}\{\frac{b_{it'l}}{q_{it'}d_{it'}}\}$ is the minimum per-timeslot per-unit-energy-consumption job value among all jobs of tenant i. We have $OPT \geq \frac{\sum_{t\in[1,T]}E_i^{(t)}}{F}\min_{t'\in[1,T], l\in\mathcal{L}_{it'}}\{\frac{b_{it'l}}{q_{it'}d_{it'}}\}$, and hence $D_0 \leq \frac{1}{2}OPT$. Then by Lemmas 3 ($C = 2$), 4 and 5, the theorem follows. □

When a tenant participates in EDR, its utility consists of job valuations and EDR rewards (*i.e.*, $\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} + \sum_{t\in[1,T]} r_i^{(t)}$), and its energy capacity bounds jobs' energy demand and EDR energy reduction (*i.e.*, $\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'l}} q_{it'}x_{it'l} + c_i^{(t)} \leq E_i^{(t)}, \forall t \in [1,T]$). Treating EDR energy shedding demand from the operator as a special job, before and after its scheduling, we have $\Delta P = \sum_{t\in\hat{t}} r_i^{(t)}$ and $\Delta D = \sum_{t\in\hat{t}} E_i^{(t)}(p_i^{(t)}(\eta_i^{(t)} + c_i^{(t)}) - p_i^{(t)}(\eta_i^{(t)}))$. by Lemma 2, we have $r_i^{(t)} \geq v_i^{(t)}(c_i^{(t)}), \forall i \in [1,N], \forall t \in [1,T]$. According to (13), we have:

$$\begin{aligned} r_i^{(t)} &\geq \int_{\eta_i^{(t)}}^{\eta_i^{(t)}+c_i^{(t)}} L_i(\frac{U_i}{L_i})^{\frac{\eta}{E_i^{(t)}}}\mathrm{d}\eta \\ &= \frac{E_i^{(t)}}{\ln(\frac{U_i}{L_i})}\cdot L_i(\frac{U_i}{L_i})^{\frac{\eta_i^{(t)}+c_i^{(t)}}{E_i^{(t)}}} - \frac{E_i^{(t)}}{\ln(\frac{U_i}{L_i})}\cdot L_i(\frac{U_i}{L_i})^{\frac{\eta_i^{(t)}}{E_i^{(t)}}} \\ &= \frac{1}{A}E_i^{(t)}(p_i^{(t)}(\eta_i^{(t)} + c_i^{(t)}) - p_i^{(t)}(\eta_i^{(t)})) \end{aligned}$$

Hence we have $\Delta P \geq \frac{1}{A}\Delta D$. By Lemma 3 and Theorem 3, $\mathsf{A}_{\text{tenant}}$ with energy shedding can obtain the same guarantee on long-term tenant utility as not participating in EDR.

C. Competitive Ratio in Global Social Welfare

We next analyze the competitive ratio in social welfare, computed by the social welfare achieved by the offline solution of (7) divided by the overall social welfare achieved by online solutions of $\mathsf{A}_{\text{tenant}}$ and $\mathsf{A}_{\text{operator}}$ at the end of $[1,T]$, based on primal-dual analysis of the social welfare maximization problem in (7) and (8).

Let P^s_{final} denote the objective value of the primal problem in (7), as computed based on online solutions $x_{it'l}$ computed by $\mathsf{A}_{\text{tenant}}$ and $c_i^{(t)}$ and $z^{(t)}$ computed by $\mathsf{A}_{\text{operator}}$ over the entire system span $[1,T]$. Let D^s_{final} denote the objective value of the dual problem (8), computed by the dual solutions $u_{it'}$ and $p_i^{(t)}$ given by $\mathsf{A}_{\text{tenant}}$ and setting λ_i as the minimum among $\{p_i^{(t)}\}_{\forall t\in[1,T]}$ and α, *i.e.*, $\lambda_i = \min\{\alpha, \min_{t\in[1,T]}\{p_i^{(t)}\}\}$. We use OPT^s to represent the offline optimal objective value of (7), computed by solving (7) exactly based on complete knowledge of the system (job arrivals, EDR signals) over the entire span $[1,T]$.

Lemma 6. *If there exist $H \geq 0$ and G such that $P^s_{final} \geq \frac{1}{H}D^s_{final} - \frac{1}{G}OPT^s$, then $\mathsf{A}_{\text{tenant}}$ in **Alg.** 1 and $\mathsf{A}_{\text{operator}}$ in **Alg.** 3 together achieve $\frac{1}{\left(\frac{1}{H}-\frac{1}{G}\right)}$ competitiveness in social welfare.*

Proof. $u_{it'}$ and $p_i^{(t)}$ given by $\mathsf{A}_{\text{tenant}}$ and λ_i as the minimum among $\{p_i^{(t)}\}_{\forall t\in[1,T]}$ and α constitute a feasible solution to the dual problem in (8). Based on weak duality [27], we have $D^s_{\text{final}} \geq OPT^s$, and hence $P^s_{\text{final}} \geq \frac{1}{H}D^s_{\text{final}} - \frac{1}{G}OPT^s \geq \left(\frac{1}{H} - \frac{1}{G}\right)OPT^s$. □

We next prove the following lemmas 7, 8 and 9, which set up the relationship between online solutions of the primal problem in (7) and dual problem in (8), for proof of the ratio between P^s_{final} and D^s_{final}. Let $\bar{v}_i^{(t)}(\cdot)$ denote the valuation function, as given in (6), computed when EDR signal $R^{(t)}$ arrives. Let $\bar{\eta}_i^{(t)}$ denote the *final* energy consumption in time slot t at tenant i, *i.e.*, the actual amount used during the time slot due to executing jobs and/or EDR energy reduction.

Lemma 7. *The online primal and dual solutions derived by $\mathsf{A}_{\text{tenant}}$ and $\mathsf{A}_{\text{operator}}$ satisfy* $\sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} + \sum_{t\in[1,T]}\sum_{i\in[1,N]}\bar{v}_i^{(t)}(c_i^{(t)}) = \sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}\int_0^{\bar{\eta}_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta.$

Proof. Let $\eta_{i[B_{it'}]}^{(t)|-1}$ and $\eta_{i[B_{it'}]}^{(t)}$ denote the reserved energy consumption for time slot t at tenant i before and after job $B_{it'}$ is scheduled, respectively. According to line 5 of **Alg.** 1, $b_{it'l}$ with $x_{it'l} = 1$ satisfies $b_{it'l} = u_{it'} + \sum_{t\in\mathcal{T}(l)} q_{it'}p_i^{(t)}(\eta_{i[B_{it'}]}^{(t)|-1})$. Since the energy demand in each job is much smaller than the tenant's energy capacity, we further have

$$\begin{aligned} b_{it'l} &= u_{it'} + \sum_{t\in[1,T]} (\eta_{i[B_{it'}]}^{(t)} - \eta_{i[B_{it'}]}^{(t)|-1}) \cdot p_i^{(t)}(\eta_{i[B_{it'}]}^{(t)|-1}) \\ &= u_{it'} + \sum_{t\in[1,T]} \int_{\eta_{i[B_{it'}]}^{(t)|-1}}^{\eta_{i[B_{it'}]}^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta \end{aligned} \tag{16}$$

In the above, note that for $t \in T\setminus\mathcal{T}(l)$, no energy is reserved for t for job $B_{it'}$, hence we have $\eta_{i[B_{it'}]}^{(t)|-1} = \eta_{i[B_{it'}]}^{(t)}$ and the respective term in the summation is zero. Let $\eta_{i[c_i^{(t)}]}^{(t)|-1}$ and $\eta_{i[c_i^{(t)}]}^{(t)}$ denote the reserved energy consumption for time slot t at tenant i before and after EDR energy reduction $c_i^{(t)}$ is scheduled, respectively. According to (13), we have

$$\bar{v}_i^{(t)}(c_i^{(t)}) = \int_{\eta_{i[c_i^{(t)}]}^{(t)|-1}}^{\eta_{i[c_i^{(t)}]}^{(t)|-1}+c_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta = \int_{\eta_{i[c_i^{(t)}]}^{(t)|-1}}^{\eta_{i[c_i^{(t)}]}^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta \tag{17}$$

Combining (16) and (17), we have

$$\begin{aligned} &\sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} + \sum_{t\in[1,T]}\sum_{i\in[1,N]}\bar{v}_i^{(t)}(c_i^{(t)}) \\ &= \sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}:x_{it'l}=1} b_{it'l} + \sum_{t\in[1,T]}\sum_{i\in[1,N]}\bar{v}_i^{(t)}(c_i^{(t)}) \\ &= \sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} \\ &\quad + \sum_{i\in[1,N]}\sum_{t\in[1,T]}\Big(\sum_{t'\in[1,T]}\int_{\eta_{i[B_{it'}]}^{(t)|-1}}^{\eta_{i[B_{it'}]}^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta + \int_{\eta_{i[c_i^{(t)}]}^{(t)|-1}}^{\eta_{i[c_i^{(t)}]}^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta\Big) \\ &= \sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}\int_0^{\bar{\eta}_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta \end{aligned}$$

□

Lemma 8. *The online primal and dual solutions derived by $\mathsf{A}_{\text{tenant}}$ and $\mathsf{A}_{\text{operator}}$ satisfy* $\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}\int_0^{\bar{\eta}_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta \geq \frac{1}{B}\Big(\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}(E_i^{(t)}p_i^{(t)} - D_0)\Big),$ *where* $B = \max_{i\in[1,N]}\{\ln(U_i/L_i)\}$.

Proof. We have

$$\begin{aligned} \int_0^{\bar{\eta}_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta &= \int_0^{\bar{\eta}_i^{(t)}} L_i(\frac{U_i}{L_i})^{\frac{\eta}{E_i^{(t)}}}\mathrm{d}\eta \\ &= \frac{1}{\ln(U_i/L_i)}E_i^{(t)}L_i(\frac{U_i}{L_i})^{\frac{\bar{\eta}_i^{(t)}}{E_i^{(t)}}} - \frac{1}{\ln(U_i/L_i)}E_i^{(t)}L_i \\ &\geq \frac{1}{B}\left(E_i^{(t)}p_i^{(t)} - E_i^{(t)}p_i^{(t)}(0)\right) \end{aligned} \tag{18}$$

Based on $u_{it'} \geq 0$ and (18), we have

$$\begin{aligned} &\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}\int_0^{\bar{\eta}_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta \\ &\geq \frac{1}{B}\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \frac{1}{B}\sum_{i\in[1,N]}\sum_{t\in[1,T]}\left(E_i^{(t)}p_i^{(t)} - E_i^{(t)}p_i^{(t)}(0)\right) \\ &= \frac{1}{B}\Big(\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}(E_i^{(t)}p_i^{(t)} - D_0)\Big) \end{aligned}$$

□

Lemma 9. *The online dual solutions derived by $\mathsf{A}_{\text{tenant}}$ satisfy* $-\sum_{t\in[1,T]}\alpha R^{(t)} \geq -\frac{Q}{W}\Big(\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]}E_i^{(t)}p_i^{(t)}\Big),$ *where $Q = \max_{i\in[1,N]}\frac{\alpha}{L_i}$ and $W = \frac{\sum_{t\in[1,T]}\sum_{i\in[1,N]}E_i^{(t)}}{\sum_{t\in[1,T]}R^{(t)}}$.*

Proof. Since $\alpha = Q \cdot \min_{i\in[1,N]}\{L_i\}$ and $\sum_{t\in[1,T]} R^{(t)} = \frac{1}{W}\sum_{t\in[1,T]}\sum_{i\in[1,N]} E_i^{(t)}$, we have

$$\begin{aligned}\alpha \sum_{t\in[1,T]} R^{(t)} &= \frac{Q}{W}\sum_{t\in[1,T]}\sum_{i\in[1,N]} (E_i^{(t)} \cdot \min_{i\in[1,N]}\{L_i\})\\ &\leq \frac{Q}{W}\sum_{t\in[1,T]}\sum_{i\in[1,N]} E_i^{(t)} p_i^{(t)}\\ &\leq \frac{Q}{W}(\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{t\in[1,T]}\sum_{i\in[1,N]} E_i^{(t)} p_i^{(t)})\end{aligned}$$

The first inequality above is based on the definition of $p_i^{(t)}$ in (12), and the second inequality is due to $u_{it'} \geq 0$. □

Theorem 4. $\mathsf{A}_{\text{tenant}}$ *in* ***Alg.*** *1 and* $\mathsf{A}_{\text{operator}}$ *in* ***Alg.****3 together achieve* $\frac{1}{(\frac{1}{B}-\frac{Q}{W})-\frac{1}{B}(\frac{1}{2-\frac{Q}{W}})}$ *competitiveness in social welfare, where* $B = \max_{i\in[1,N]}\{\ln(U_i/L_i)\}$, $W = \frac{\sum_{t\in[1,T]}\sum_{i\in[1,N]} E_i^{(t)}}{\sum_{t\in[1,T]} R^{(t)}}$ *and* $Q = max_{i\in[1,N]}\{\frac{\alpha}{L_i}\}$.

Proof. Since $z^{(t)} = R^{(t)} - \sum_{i\in[1,N]} c_i^{(t)}, \forall t \in [1,T]$, we have P^s_{final} Since $z^{(t)} = R^{(t)} - \sum_{i\in[1,N]} c_i^{(t)}, \forall t \in [1,T]$, we have

$$\begin{aligned}&P^s_{final}\\ &= \sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} - \sum_{t\in[1,T]} \alpha(R^{(t)} - \sum_{i\in[1,N]} c_i^{(t)})\\ &\geq \sum_{i\in[1,N]}\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} + \sum_{t\in[1,T]}\sum_{i\in[1,N]} \bar{v}_i^{(t)}(c_i^{(t)}) - \sum_{t\in[1,T]} \alpha R^{(t)}\\ &= \sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]} \int_0^{\bar{\eta}_i^{(t)}} p_i^{(t)}(\eta)\mathrm{d}\eta - \sum_{t\in[1,T]} \alpha R^{(t)}\\ &\geq (\frac{1}{B} - \frac{Q}{W})\Big(\sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'} + \sum_{i\in[1,N]}\sum_{t\in[1,T]} E_i^{(t)} p_i^{(t)}\Big) - \frac{1}{B} D_0\\ &\geq (\frac{1}{B} - \frac{Q}{W})\Big(-\sum_{t\in[1,T]} R^{(t)}\lambda^{(t)} + \sum_{i\in[1,N]}\sum_{t'\in[1,T]} u_{it'}\\ &\qquad + \sum_{i\in[1,N]}\sum_{t\in[1,T]} E_i^{(t)} p_i^{(t)}\Big) - \frac{1}{B} D_0\\ &= \frac{W-BQ}{BW} D^s_{final} - \frac{1}{B} D_0\end{aligned}$$

The third and fourth rows are based on Lemma 2 and Lemma 7, respectively. The fifth row is based on Lemmas 8 and 9. The six row is due to $\lambda^{(t)} \geq 0$. Next, we have:

$$\begin{aligned}OPT^s &\geq \sum_{i\in[1,N]} \frac{\sum_{t\in[1,T]} E_i^{(t)}}{F} \min_{t'\in[1,T], l\in\mathcal{L}_{it'}}\{\frac{b_{it'l}}{q_{it'}d_{it'}}\} - \alpha \sum_{t\in[1,T]} R^{(t)}\\ &= \sum_{i\in[1,N]}\sum_{t\in[1,T]} (2E_i^{(t)} L_i) - \alpha \sum_{i\in[1,N]}\sum_{t\in[1,T]} \frac{E_i^{(t)}}{W}\\ &\geq \sum_{i\in[1,N]}\sum_{t\in[1,T]} \Big(E_i^{(t)}(2L_i - \frac{QL_i}{W})\Big)\\ &= \sum_{i\in[1,N]}\sum_{t\in[1,T]} \Big(E_i^{(t)} L_i(2 - \frac{Q}{W})\Big)\end{aligned}$$

The first inequality is due to $\sum_{t'\in[1,T]}\sum_{l\in\mathcal{L}_{it'}} b_{it'l}x_{it'l} \geq \frac{\sum_{t\in[1,T]} E_i^{(t)}}{F}\min_{t'\in[1,T], l\in\mathcal{L}_{it'}}\{\frac{b_{it'l}}{q_{it'}d_{it'}}\}$ and $z^{(t)} \leq R^{(t)}$. Then, $\frac{D_0}{OPT^s} \leq \frac{\sum_{i\in[1,N]}\sum_{t\in[1,T]} E_i^{(t)} L_i}{\sum_{i\in[1,N]}\sum_{t\in[1,T]}\big(E_i^{(t)} L_i(2-\frac{Q}{W})\big)} = \frac{1}{(2-\frac{Q}{W})}$, and hence $P^s_{final} \geq \frac{W-BQ}{BW} D^s_{final} - \frac{1}{B}\frac{1}{(2-\frac{Q}{W})} OPT^s$ Since $B \geq 1$ and $(1-B\frac{Q}{W}) \leq 1$, we have $\frac{BW}{W-BQ} = \frac{B}{1-B\frac{Q}{W}} \geq 0$. Based on Lemma 6 with $H = \frac{BW}{W-BQ}$, $G = B(2-\frac{Q}{W})$, the online algorithms together achieve a competitive ratio of $\frac{1}{(\frac{1}{B}-\frac{Q}{W})-\frac{1}{B}(\frac{1}{2-\frac{Q}{W}})}$ in social welfare. □

VI. Performance Evaluation

We now evaluate tenants' online scheduling algorithm in **Alg.** 1 and the operator's incentive mechanism in **Alg.** 3 by conducting simulation experiments modeling a large colocation data center based on real-world workload and EDR traces.

A. Simulation Setup

We consider a colocation data center with five tenants participating in EDR by default. Each tenant houses 2000 servers, and each server consumes a peak power of 250W. The energy capacity of each tenant ($E_i^{(t)}$) is decided as the total hourly energy consumption divided by the number of time slots per hour. We set the diesel generation cost α as 200\$/MWh, based on typical power generation efficiency and the average oil price in 2014 [28].

EDR event. We simulate a real EDR event that occurred on January 7, 2014, throughout the service region of PJM due to abnormally cold weather [29]. The EDR signals arrive hourly during a 7-hour span.[3] Each frame in our experiments equals one hour, the length of each time slot is 2 minutes, and the total system duration is $T = 210$ time slots. We scale the total PJM-wide energy reduction trace such that the EDR signals in our experiments indicate an overall energy shedding of at most 25% of the colocation's peak IT energy (consumed by all tenants). This is a reasonable setting, as LBNL's recent field test has demonstrated that data centers can reduce 25% of its server energy without significantly affecting the normal operation [9].

Tenant jobs. We simulate job arrivals at tenants following the workload trace collected from a Google cluster [30], a dataset widely used in emulating user workload in existing cloud literature [31], [32], which contains information of jobs submitted to the cluster, including normalized resource demand (CPU, RAM, Disk), arrival time and duration of each job. We convert the CPU demand of a job into a per-timeslot energy demand of the job ($q_{it'}$), by multiplying the normalized CPU demand by the energy capacity of the tenant. A job's execution duration ($d_{it'}$) is within $[15, 90]$ slots, according to the workload trace. We set $U_i/L_i = 30$ by default, and set the job valuation by multiplying overall energy demand in the bid (*i.e.*, $q_{it'}d_{it'}$) by a unit price randomly picked within different ranges according to the values of U_i and L_i. As the workload trace does not contain job execution deadlines, we randomly select a time, between the earliest possible finishing time and the ending time of the entire time horizon (*i.e.*, $[t'+d_{it'}, T]$), as the job's deadline ($w_{it'}$). We denote the time period $[t', w_{it'}]$ as the *no-penalty* execution window of job $B_{it'}$, *i.e.*, no penalty would be received if the job is completed within this time window. We use the following penalty function for each job:

[3] Note that our model and theoretical results cater to a much more flexible scenario, with EDR signals arriving anytime during a long time span. Here due to the limitation of available traces, we simulate 7 EDR signals arriving in 7 consecutive hours.

$$g_{it'_+}(\sigma_{it'}) = \begin{cases} \sigma_{it'} \cdot \frac{b_{it'}}{\theta(w_{it'}-t')}, \\ \quad \text{if } w_{it'} + \sigma_{it'} \in (w_{it'}, \min\{w_{it'} + \theta(w_{it'} - t'), T\}] \\ +\infty, \\ \quad \text{if } w_{it'} + \sigma_{it'} \in (\min\{w_{it'} + \theta(w_{it'} - t'), T\}, T] \end{cases} \tag{19}$$

Note that here we further split $g_{it'_+}(\cdot)$ in (1) into two pieces: if the job completion time $(w_{it'} + \sigma_{it'})$ falls within a *penalty* window beyond $w_{it'}$, but not beyond $w_{it'}$ plus θ fraction of the length of the no-penalty window $(w_{it'} - t')$, then the penalty increases linearly with the amount of time the job takes beyond $w_{it'}$ from 0 to $b_{it'}$ (the job's maximum value); otherwise, the penalty is infinity. θ is 0.5 by default. We will evaluate different lengths of the penalty window by varying θ, as well as different types of penalty functions (square and exponential functions).

Comparisons. We compare solutions given by our online algorithms with the offline optimum, computed by solving (7) exactly using CVX optimizer, assuming all tenants' job arrivals and all EDR signals are known in advance. We also compare with three other schemes. (1) *Diesel-only*: the colocation operator fully relies on diesel generation to meet the EDR energy shedding demand, and tenants schedule their jobs using $\mathsf{A}_{\mathsf{tenant}}$ without participating in EDR. (2) *Greedy*: each tenant greedily schedules its arrival jobs in time slots with available energy to meet the jobs' energy demands as early as possible, and uses the average valuation of scheduled jobs in a time slot as the cost valuation of energy reduction in this time slot to bid in the operator's auction. (3) *PropAllocation*: each tenant runs $\mathsf{A}_{\mathsf{tenant}}$ to schedule jobs, and the operator allocates energy reduction to tenants for EDR proportionally based on the percentage of their subscribed energy capacities ($E_i^{(t)}$) in the whole data center (in case tenants cannot possibly achieve the specified energy reduction for EDR, the operator will resort to diesel generation for the discrepancy). Note that we create the above themes as reasonable baselines, since no existing work in the literature has jointly studied tenants' scheduling decisions and the colocation operator's incentive mechanism.

B. Evaluation Results

1) Comparison with Offline Optimum: We first evaluate the competitive ratio in social welfare, calculated by dividing the offline optimal social welfare by social welfare obtained by our online solutions, which shows how well our online mechanism performs in terms of global social welfare without knowing future information. Due to time complexity of solving (7) exactly for offline optimum, the largest number of jobs is limited to 800. Job arrivals are extracted from the trace (scaled down to within the EDR duration), and an arrival job is randomly assigned to one of the five tenants.

In **Fig.** 3, we vary θ from 0.5 to 1.5, representing that the length of penalty window is 50%, 100%, or 150% of the non-penalty execution window, respectively. We also vary the number of tenants from three to seven, and the total number of jobs in the system changes accordingly. We observe that a smaller penalty window results in a smaller competitive ratio, *i.e.*, social welfare achieved by our algorithms is closer to the offline optimum. This is because a smaller penalty window leads to a smaller time window for job scheduling and hence restricts the decision space, which affects the offline optimal social welfare more as the offline optimum explores future information to perform better. We also observe that the competitive ratio is better with more tenants. In particular, we find that with global information, the offline optimum can achieve the perfect EDR reduction (*i.e.*, using tenant energy only but no diesel usage) when only few tenants participate in EDR; with more tenants, our mechanism relies less on diesel while the offline optimum does not change much in this aspect, leading to closer performance of our mechanism to the offline optimum. In all cases, the competitive ratio is within 2-3.

In **Fig.** 4, we evaluate the impact of the following penalty functions: a_1x, a_2x^2 and $a_3(e^x - 1)$. Let $\bar{w} = \theta(w_{it'} - t')$ denote the length of the penalty window. We use $a_1 = b_{it'}/\bar{w}$, $a_2 = b_{it'}/(\bar{w}^2)$, and $a_3 = b_{it'}/(e^{\bar{w}} - 1)$, such that the penalty always equals the job's value at the end of the penalty window no matter with which penalty function (a job brings no value if not completed before the end of the penalty window). Note a_1x is the linear function we define in (19). **Fig.** 4 shows that the competitive ratios are smallest with the linear penalty function, then with the square penalty function, and then with the exponential function. Under our penalty functions (zero at the start of the penalty window; $b_{it'}$ at the end of the window), the same delay from the soft deadline leads to the largest penalty with the linear penalty function (followed by the square function, then the exponential function), such that the job subject to the delay is more likely rejected by the tenant's online algorithm, leaving more energy for high-value future jobs. We also observe that the competitive ratio is worse with more servers: the offline optimum with future information can fully utilize the increased capacity, but our online mechanism cannot be as efficient; hence the welfare improvement of the offline optimum is larger.

In **Fig.** 5, we evaluate the impact of different EDR demands, where "50%, 100%, 200%" are the percentages used to scale the default EDR energy reduction demand in each frame. When the EDR demands go up, W in Theorem 4 is smaller and hence the ratio is larger, which is consistent with our theoretical result in the theorem. In **Fig.** 6, we further evaluate the impact of an estimation of U_i/L_i, with the estimated value varying between 80% and 120% of the actual U_i/L_i. We see that an underestimation leads to a better competitive ratio, and hence is more desirable than overestimation when running the online algorithm. We also observe that the more jobs, the better the competitive ratio is. This could be explained as follows: with more jobs, their energy demands approach the capacities of tenants, and there is less room for the offline optimum to achieve better social welfare by exploiting future information.

2) Comparison with Heuristic Schemes: In **Fig.** 7 and **Fig.** 8, we compare the social welfare and the tenant's utility achieved by our online algorithms with the three heuristic schemes. **Fig.** 7 shows that our algorithms outperform all the other schemes, and **Fig.** 8 shows that the utility of individual tenants outperforms that achieved by the other schemes as well. All these reveal that our joint tenant and operator's online optimization algorithms, by efficiently incentivizing

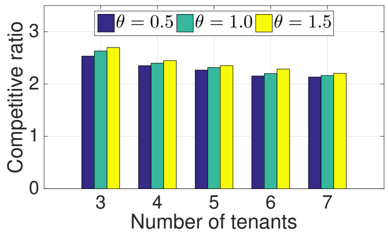

Fig. 3. Competitive ratio with different penalty window sizes

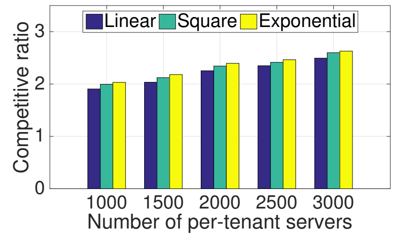

Fig. 4. Competitive ratio with different penalty functions

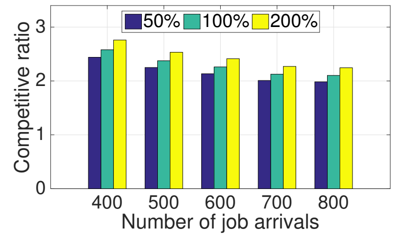

Fig. 5. Competitive ratio with different EDR demands

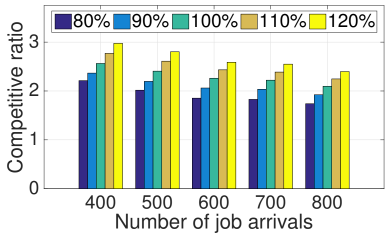

Fig. 6. Competitive ratio over different pct. of estimation over actual U_i/L_i

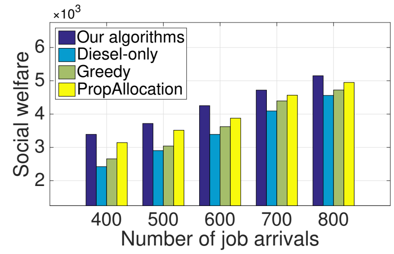

Fig. 7. Comparison of social welfare among different schemes

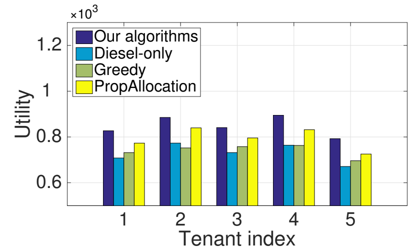

Fig. 8. Comparison of individual tenant's utility among different schemes

tenants to participate in EDR and optimally deciding tenant's job scheduling and cost valuation, improve the global social welfare and individual tenant's utility.

3) Breakdown of Energy Reduction: **Fig.** 9 shows the breakdown of energy reduction during each EDR frame, as provided by individual tenants and diesel generation. The sum of energy reduction amounts in each EDR frame is the overall energy shedding demand in the corresponding EDR signal. When the overall EDR energy shedding demand is higher, the diesel usage tends to be larger, since tenants may not afford more energy reduction due to job needs. **Fig.** 10 shows the diesel saving ratio in each EDR frame, calculated by dividing the overall tenants' energy reduction by the EDR demand from the grid in the frame. We observe that the diesel savings by our algorithms outperform the offline optimum in most times, with only 0 to 0.23 of the EDR demand covered by dirty energy.

4) With and Without Participating in EDR: In **Fig.** 11 and **Fig.** 12, we further illustrate the operator cost (*i.e.*, diesel cost plus payment to tenants) and tenant utility (*i.e.*, job valuation plus received reward) when the per-timeslot energy demand of each job at each tenant is scaled by different factors (such that the sever load is scaled). When the energy demand of a job is scaled, the job valuation is scaled accordingly. We also compare results by running our mechanism ("with EDR") with those obtained when no EDR auction is in place ("without EDR"): in the latter, the operator relies only on diesel generator and tenants schedule jobs according to **Alg.** 1 without participating in any EDR auction. We observe that both the operator and the tenants benefit from participating in our EDR mechanism. The operator cost remains stable while the tenant utility increases with the increase of energy demand and valuation of the jobs.

VII. Concluding Remarks

This work designs a set of online mechanisms, seamlessly jointing individual online optimization and centralized auction: where individual tenants execute online algorithms to schedule jobs, and meanwhile evaluate their energy valuation for participating EDR; and operator launches a truthful EDR mechanism to solicit energy from tenants and provides reward, targeting diesel usage saving and global social welfare maximization. We consider a general job model including deadline and penalty function, and handle it by novel primal-dual online optimization and compact-exponential technique. We creatively design a function to reveal the tenant's energy valuation during the online job scheduling and achieve a guarantee on global social welfare under an EDR event. To the best of our knowledge, compared with existing colocation EDR schemes, we are the first to handle the whole colocation system consisting tenant scheduling and operator EDR reducing, and we are also the first to reveal tenant's inexplicit energy valuation during a period. We combine a centralized mechanism and multiple decentralized online algorithms to achieve a global social welfare guarantee. Trace driven simulations validate our theoretical analysis and show good performance of our system (including online algorithm and EDR mechanism), as compared with offline optimum and different schemes.

VIII. Acknowledgements

This work was supported in part by Hong Kong RGC under grants HKU 718513, 17204715, 17225516 and C7036-15G(CRF), by NSERC and Wedge Networks, and by the NSF under grants CNS-1551661, CNS-1565474, and ECCS-1610471.

References

[1] K. Managan, "Demand response: A market overview," 2014, http://enaxisconsulting.com.

[2] 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.

[3] A. Misra, "Responding before electric emergencies," http://www.afcom.com/digital-library/pub-type/communique/responding-before-electric-emergencies/.

[4] Santa Barbara County, "Air pollution control district," http://www.ourair.org/do-you-really-need-a-diesel-generator/.

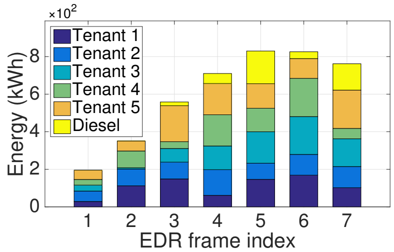

Fig. 9. EDR energy reduction by tenants and diesel generator

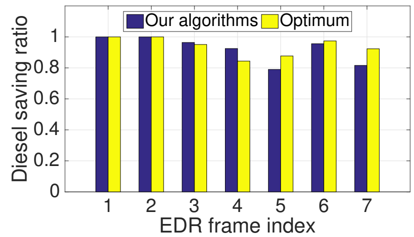

Fig. 10. Comparison of diesel saving ratio by our algorithms and offline optimum

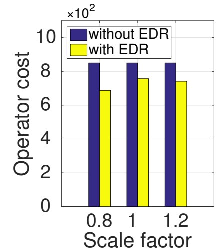

Fig. 11. Comparison of operator cost

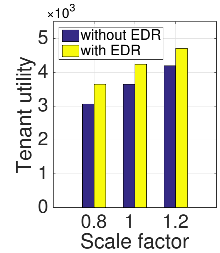

Fig. 12. Comparison of tenant utility

[5] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," in *INFOCOM*, 2011.

[6] J. R. Lorch and A. J. Smith, "PACE: A new approach to dynamic voltage scaling," *IEEE Trans. Computers*, vol. 53, July 2004.

[7] I. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Greenhadoop: Leveraging green energy in data-processing frameworks," in *Proc. ACM EuroSys*, 2012.

[8] D. Wang, S. Govindan, A. Sivasubramaniam, A. Kansal, J. Liu, and B. Khessib, "Underprovisioning backup power infrastructure for datacenters," in *Proc. ACM ASPLOS*, 2014.

[9] 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.

[10] Apple, "Environmental responsibility report 2016 - Apple," https://www.apple.com/environment/pdf/Apple_Environmental_Responsibility_Report_2016.pdf.

[11] Q. Sun, C. Wu, S. Ren, and Z. Li, "Fair rewarding in colocation data center: Truthful mechanism for emergency demand response," in *IEEE IWQoS*, 2015.

[12] N. Chen, X. Ren, S. Ren, and A. Wierman, "Greening multi-tenant data center demand response," in *IFIP Performance*, 2015.

[13] S. Ren and M. A. Islam, "Colocation demand response: Why do I turn off my servers?" in *ICAC*, 2014.

[14] M. A. Adnan and R. Gupta, "Utility-aware deferred load balancing in the cloud driven by dynamic pricing of electricity," in *DATE*, 2013.

[15] S. Li, M. Brocanelli, W. Zhang, and X. Wang, "Data center power control for frequency regulation," in *IEEE PES*, 2013.

[16] H. Xu and B. Li, "Reducing electricity demand charge for data centers with partial execution," in *ACM e-Energy*, 2014.

[17] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for internet-scale systems," in *SIGCOMM*, 2009.

[18] B. Aksanli and T. S. Rosing, "Providing regulation services and managing data center peak power budgets," in *DATE*, 2014.

[19] 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 *SIGMETRICS*, 2013.

[20] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," in *SIGMETRICS*, 2014.

[21] B. Lucier, I. Menache, J. S. Naor, and J. Yaniv, "Efficient online scheduling for deadline-sensitive jobs," in *SPAA*, 2013.

[22] S. Albagli-Kim, H. Shachnai, and T. Tamir, "Scheduling jobs with dwindling resource requirements in clouds," in *INFOCOM*, 2014.

[23] S. T. Maguluri and R. Srikant, "Scheduling jobs with unknown duration in clouds," in *INFOCOM*, 2013.

[24] PJM, "Retail electricity consumer opportunities for demand response in PJM's wholesale markets," http://www.pjm.com.

[25] I. Goiri, R. Bianchini, S. Nagarakatte, and T. D. Nguyen, "Approxhadoop: Bringing approximations to mapreduce frameworks," in *Proc. ACM ASPLOS*, 2015.

[26] N. Buchbinder and J. Naor, "The design of competitive online algorithms via a primal-dual approach," *Foundations and Trends in Theoretical Computer Science*, vol. 3, pp. 93–263, 2009.

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

[28] Wikipedia, http://en.wikipedia.org/wiki/Diesel_generator.

[29] PJM, "Demand response activity on January 7-8, 2014," http://www.pjm.com.

[30] C. Reiss, J. Wilkes, and J. L. Hellerstein, "Google cluster-usage traces: format + schema," Google Inc., Technical Report, Nov. 2011.

[31] H. Xu and W. C. Lau, "Task-cloning algorithms in a mapreduce cluster with competitive performance bounds," in *ICDCS*, 2015.

[32] W. Wang, B. Liang, and B. Li, "Multi-resource fair allocation in heterogeneous cloud computing systems," *IEEE TPDS*, vol. 26, 2015.

Qihang Sun received the B.S. degree from Wuhan University, China, in 2014, in Computer Science and Technology. He is currently working towards the Ph.D. degree in the Computer Science Department at The University of Hong Kong. His research interests include power-aware computing, online optimization, and cloud computing.

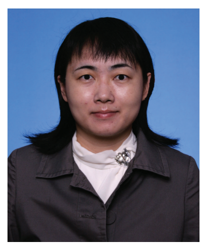

Chuan Wu received her B.E. and M.E. degrees in 2000 and 2002 from Department of Computer Science and Technology, Tsinghua University, China,and her Ph.D. degree in 2008 from the Department of Electrical and Computer Engineering, University of Toronto, Canada. She is currently an associate professor in the Department of Computer Science,The University of Hong Kong, China. Her research interests include cloud computing and online/mobile social network.

Zongpeng Li received his B.E. degree in Computer Science and Technology from Tsinghua University (Beijing) in 1999, his M.S. degree in Compute r Science from University of Toronto in 2001, and his Ph.D. degree in Electrical and Computer Engineering from University of Toronto in 2005. He is currently an associate professor in the Department of Computer Science in the University of Calgary. His research interests are in computer networks, particularly in network optimization, multicast algorithm design, network game theory and network coding.

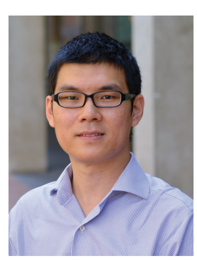

Shaolei Ren is an Assistant Professor of Electrical and Computer Engineering at University of California, Riverside. He received his B.E. from Tsinghua University in 2006, M.Phil. from Hong Kong University of Science and Technology in 2008, and Ph.D. from University of California, Los Angeles in 2012, all in electrical and computer engineering. His research interests include cloud computing, data centers, and network economics. He was a recipient of the U.S. NSF Faculty Early Career Development (CAREER) Award in 2015. He also received best paper awards from several conferences, including ACM e-Energy'16, IEEE ICC'16 and IEEE ICC'09.