Colocation Data Center Demand Response Using Nash Bargaining Theory

Yuanxiong Guo, *Member, IEEE*, Hongning Li, *Member, IEEE*, and Miao Pan, *Member, IEEE*

***Abstract*—The huge yet flexible power consumption of data centers makes them promising resources for demand response, particularly for emergency demand response (EDR) which requires a certain amount of load curtailment during emergencies. However, current data centers often participate in EDR by starting up their backup diesel generators, resulting in both high costs and large carbon emissions. In this paper, we focus on cost-effective and eco-friendly demand response in colocation data centers by designing economic incentives for tenants to reduce their loads during emergency periods for EDR. In particular, we model and analyze the interaction among the data center operator and tenants by using Nash bargaining theory, and derive the optimal solutions for the load reduction and reimbursement for each tenant under two different bargaining protocols (i.e., sequential bargaining and concurrent bargaining). We prove that the derived solutions are Pareto-efficient and fair, and therefore self-enforcing and satisfactory for all entities. Numerical results based on trace-driven simulations show that the proposed bargaining approach is beneficial to both the data center operator and tenants, while also reducing the carbon emissions to the environment from data center demand response.**

***Index Terms*—Colocation data centers, emergency demand response, Nash bargaining, sequential and concurrent bargaining.**

I. INTRODUCTION

DATA centers are expanding in both numbers and scales to satisfy the exploding IT demand, which leads to significant power consumption. For instance, data centers in the U.S. alone consumed 91 billion kilowatt-hours of electricity in 2013, resulting in an estimated electricity bill of $9 billion. Moreover, the power consumption of U.S. data centers is expected to increase about 10% annually [1]. Although the huge power consumption of data centers is traditionally regarded as a burden to the power grid, data centers are being recognized as promising yet under-utilized demand response resources recently [2]. Data center demand response is made technically and economically feasible by leveraging the IT computing knobs such as geographical load balancing [3], [4], dynamic capacity provisioning [5], and workload shifting [6], as well as non-IT knobs including batteries and cooling systems [7], [8]. Through data center demand response, data centers receive financial benefit, and the power grid better balances the supply and demand in real time.

Emergency demand response (EDR) is the most widely-adopted demand response program in the U.S., accounting for 87% of demand reduction capabilities across all reliability regions [9]. Under emergency situations such as extreme weather or natural disasters, many large electricity consumers including data centers are coordinated under EDR to reduce their power consumption for preventing blackouts. Data centers, as large yet flexible loads, have been regarded as important resources for EDR by the U.S. EPA [10]. For instance, hundreds of data centers participated into EDR by reducing their power consumption on July 22, 2011 to prevent a cascading blackout that might lead to economic losses of billions of dollars [11].

Traditionally, data centers participate into EDR by starting up their backup diesel generators, which is both expensive and environmentally unfriendly. The high expenses incurred by running diesel generators might not even be covered by the financial benefit data centers receive from the power utility company as EDR is called more and more frequently with the increasing penetration of renewable energy sources. Moreover, the air pollutants emitted from running diesel generators for demand response have raised wide-spread concerns [12]. Therefore, data centers nowadays intend to participate into demand response by modulating server power consumption (see [13] for a survey on this topic), which is more cost-effective and eco-friendly. However, existing studies on data center demand response mostly focus on owner-operated data centers such as Google data centers whose operators have full control over both servers and facilities, and these data centers are actually not very suitable for EDR due to their geographical locations and workload nature as analyzed in [14].

Colocation data centers (simply called colos), which rent out spaces to multiple tenants to house their servers, are another important but under-explored type of data centers. In a colo, the colo operator is only responsible for facility support and has no control over the servers housed in the colo. In the U.S. alone, there are more than 1200 colos [15], and the market

Manuscript received May 18, 2016; revised September 12, 2016 and November 5, 2016; accepted November 14, 2016. Date of publication January 2, 2017; date of current version August 21, 2018. This work was supported in part by the U.S. National Science Foundation under Grant CNS-1343361, Grant CNS-1350230, and Grant CPS-1646607, in part by the National Natural Science Foundation of China under Grant 61602358 and Grant 61373170, and in part by the National Key Research and Development Program of China under Grant 2016YFB0800601. Paper no. TSG-00675-2016.

Y. Guo is with the School of Electrical and Computer Engineering, Oklahoma State University, Stillwater, OK 74078 USA (e-mail: richard.guo@okstate.edu).

H. Li is with the State Key Laboratory of Integrated Services Networks, Xidian University, Xi'an 710071, China (e-mail: hnli@xidian.edu.cn).

M. Pan is with the Department of Electrical and Computer Engineering, University of Houston, Houston, TX 77204 USA (e-mail: mpan2@uh.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSG.2016.2647246

of colo is around \$43 billion with annual growth rate to be 11% [1]. It has been reported that colos consume nearly 40% of the total data center energy in the U.S., while Google-type data centers collectively only account for less than 8%. Furthermore, colos are mostly located in urban areas such as New York and Silicon Valley which are residential intensive. Since EDR is more likely to be called in these areas, colos can provide a higher potential for peak power reduction than the owner-operated data centers which are usually located in rural areas.

Considering colo EDR, we need to address the split incentive challenge: while the colo operator desires satisfying EDR without relying on on-site backup generation, tenants who have control over their servers have little or even no incentive to behave in the interest of the colo operator since they are usually charged based on their subscribed peak or reserved power [16]. Moreover, the tenants cannot directly participate into EDR due to their small individual load reduction capabilities and invisibility from the power system operator. Therefore, proper incentives must be provided from the colo operator to tenants for reducing their server power consumption during EDR periods.

In this paper, we focus on cost-effective and eco-friendly EDR in a colocation data center. Specifically, we investigate how the colo operator should incentivize the tenants to reduce their power consumption during EDR periods, and how the economic surplus generated during EDR periods should be shared among the colo operator and tenants. Our proposed approach is based on Nash bargaining theory [17]: once the colo operator receives the EDR signal, the colo operator initiates bargaining with tenants to negotiate the amount of power reduction extracted from tenants and the corresponding reimbursement offered to them. The outcome, i.e., the load reduction and corresponding reimbursement, must be approved by both the colo operator and tenants. We quantify the benefit of the colo operator and tenants in collaborative load reduction during EDR events. The closed-form optimal solutions for the load reduction and reimbursement for each tenant are derived under two different bargaining protocols (i.e., sequential bargaining and concurrent bargaining). We prove that the optimal solutions satisfy Pareto efficiency and fairness. We finally use trace-driven simulations to verify that our proposed approach is beneficial to both the data center operator and tenants by decreasing their costs, while also reducing the carbon emissions from colo EDR.

The remainder of this paper is organized as follows. First, we present the related works in Section II. Then, we describe the system model in Section III. We further analyze the NBS under the sequential and concurrent bargaining in Section IV. Next, we show the performance evaluation in Section V. Finally, we conclude the paper in Section VI.

II. Related Work

Colocation data centers are receiving increasing attention recently in literature. One stream of research focuses on incentivizing the tenants in colos to reduce power consumption during EDR periods. Ren and Islam [16] first consider this problem and propose a heuristic solution. However, truthfulness of the tenants' bids is not guaranteed in [16]. Zhang *et al.* [18] investigate EDR in colo, and propose an auction based solution which guarantees approximate truthfulness. A parameterized supply function bidding based approach is proposed in [14] to study both EDR and economic DR in colo. Tran *et al.* [19] investigate incentive mechanism design in colo during the DR event using the Stackelberg game and the auction theory. Another stream of research investigates how to coordinate the tenants to reduce the operation cost of the colo. For example, Islam *et al.* [20] optimize the reward issued to the tenants to save the operation cost. A bidding based approach is proposed in [21] to minimize the carbon footprint in a colo. Decentralized energy management under the coordination of the colo operator in a colo is investigated in [22].

This paper can be categorized into the first stream. The difference between this paper and existing works is that most existing works study the incentive mechanism design in a non-cooperative manner, and thus cannot model the cooperation among the colo operator and tenants. Our paper studies how the economic surplus generated when incentivizing the tenants to participate into EDR should be shared among the colo operator and tenants. Our analysis shows that the proposed approach can achieve Pareto efficiency and fairness. Moreover, the social welfare is proved to be maximized in the concurrent bargaining, which cannot be achieved using the Stackelberg game as proposed by [19]. Our paper outperforms the auction based works [14] and [18] in the following sense. The colo operator may not be suitable to perform as the auctioneer since it is self-interested, and the auctioneer selection can highly affect the auction result [23]. However, our proposed approach is decentralized and self-enforcing, i.e., both the tenants and the colo operator have the right to reject any outcome that impairs its profit so that a mutual beneficial outcome can be reached. Therefore, our paper is complementary to the existing works. To summarize, our paper captures the coordination between the colo operator and the tenants, and studies the fair economic surplus sharing among them. The proposed approach is proved to be Pareto-efficient and fair in a decentralized way.

III. System Model

We consider a colo operated by a colo operator with $\mathcal{N} = \{1, 2, \dots, N\}$ tenants. Each tenant $i \in \mathcal{N}$, who subscribes a peak power usage from the colo operator, has M_i homogeneous servers. The colo operator is responsible for the facility support such as cooling and security. In this paper, we focus on the EDR in a single period: under emergencies, the power utility company issues the EDR signal to the colo operator, requiring a certain amount of power reduction D from the colo. The EDR target D is mandatory and must be satisfied during the EDR period. Otherwise, extremely high penalty will be incurred. Traditionally, the colo operator starts up its on-site generator to satisfy the EDR target, which is both expensive and unfriendly to the environment.

A. Tenant

Power consumption model: Denote the idle power and peak power of each server of tenant i as P_i^{idle} and P_i^{peak}, respectively. Using the power consumption model [24], the average power consumption of tenant i when x servers are kept active is represented as

$$P_i(x) = x\Big[P_i^{\text{idle}} + u_i\Big(P_i^{\text{peak}} - P_i^{\text{idle}}\Big)\Big], \tag{1}$$

where u_i is the average CPU utilization level. Assume M/G/1/PS queueing model is adopted at each server [4]. Denote the mean arrival rate of workloads and the mean service rate of a server at tenant i as λ_i and μ_i, respectively. Then the average CPU utilization level is $u_i = \lambda_i/(\mu_i x)$.

We assume that all servers are kept active when tenant i does not participate into EDR programs. Then, its average power consumption is calculated as

$$P_i^{\text{ref}} = M_i\Big[P_i^{\text{idle}} + u_i^{\text{ref}}\Big(P_i^{\text{peak}} - P_i^{\text{idle}}\Big)\Big], \tag{2}$$

where $u_i^{\text{ref}} = \lambda_i/(\mu_i M_i)$. We further assume that the tenants participate into EDR by turning off the unused servers to reduce power consumption. Hence, when m_i servers are turned off at tenant i during the EDR period, its average power consumption P'_i is

$$P'_i = (M_i - m_i)\Big[P_i^{\text{idle}} + u'_i\Big(P_i^{\text{peak}} - P_i^{\text{idle}}\Big)\Big], \tag{3}$$

where $u'_i = \lambda_i/(\mu_i(M_i - m_i))$. From (2) and (3), the power reduction of tenant i when turning off m_i servers is calculated as

$$\Delta P_i = P_i^{\text{ref}} - P'_i = m_i P_i^{\text{idle}}. \tag{4}$$

Note that according to (4), the power reduction ΔP_i is discrete due to the fact that the number of turned off servers m_i is an integer. However, without loss of generality, in this paper, we relax the constraint that requires m_i to be integer due to the fact that data centers usually contain thousands of servers. Therefore, the power reduction ΔP_i can be regarded as continuous.

QoS constraint model: The maximum number of servers that can be turned off by each tenant is constrained by its quality of service (QoS) requirement. In this paper, we use the processing delay suffered by each tenant after turning off unused servers to quantify the QoS constraint. Denote $T_i^{\max}$ as the maximum average delay of the workloads that can be tolerated by tenant i. Then using queueing theory [25], the average response time is characterized as

$$T_i = \frac{1}{\mu_i - \frac{\lambda_i}{M_i - m_i}} \le T_i^{\max}. \tag{5}$$

Hence, it is straightforward to obtain that the number of the servers that can be turned off is bounded as follows:

$$0 \le m_i \le M_i - \frac{\lambda_i}{\mu_i - \frac{1}{T_i^{\max}}}. \tag{6}$$

Using (6), the DR capacity of tenant i is

$$\Delta P_i^{\max} = P_i^{\text{idle}}\left(M_i - \frac{\lambda_i}{\mu_i - \frac{1}{T_i^{\max}}}\right). \tag{7}$$

Therefore, the power reduction offered by each tenant i should satisfy the following constraint:

$$0 \le \Delta P_i \le \Delta P_i^{\max}. \tag{8}$$

Inconvenience cost model: Tenants that turn off servers will incur inconvenience cost due to performance degradation, and different tenants have different inconvenience costs. In this paper, we do not specify which type of inconvenience cost is imposed to the tenants. However, we make the following general assumption similar to [14].

Assumption 1: For each tenant $i \in \mathcal{N}$, the inconvenience cost function $C_i(\Delta P_i)$ is convex, strictly increasing and differentiable over the domain $0 \le \Delta P_i \le \Delta P_i^{\max}$, with $C_i(\Delta P_i) = 0$ when $\Delta P_i = 0$.

We denote the first-order derivative function of $C_i(\Delta P_i)$ as $C'_i(\Delta P_i)$, whose inverse function is denoted as $C_i'^{-1}(\cdot)$. Moreover, we assume that $\exists \Delta P_i : C_i(\Delta P_i) < \alpha\gamma\Delta P_i, \forall i$ to ensure the economic feasibility of tenant load reduction.

Tenant payoff: Each tenant i will be reimbursed when reducing its server power consumption. Denote the reimbursement received by tenant i as r_i. Then the payoff of tenant i can be represented as

$$U_i(r_i, \Delta P_i) = r_i - C_i(\Delta P_i). \tag{9}$$

In the following, we assume that each tenant i will interact with the colo operator by revealing its intended power reduction ΔP_i and reimbursement r_i.

B. Colo Operator

First, consider the scenario that no tenant is willing to reduce its power consumption during the EDR period. Therefore, to satisfy the power reduction target D issued by the utility company, the colo operator has to start up its on-site backup generator. In this case, we model the cost incurred by the colo operator as a linear function $G^0 = \alpha D$, where α is the unit on-site generation cost. Here we assume that the on-site generator has a sufficiently large capacity to satisfy the EDR target D, which is reasonable since most data centers are mission-critical and have enough backup generation capability in practice.

Second, if the tenants can be properly incentivized to reduce their power consumption, the cost (both economic and environmental) incurred by the colo operator will be reduced significantly. To capture the total power consumption of the colo including both the IT part and the non-IT part, we adopt power usage effectiveness (PUE) γ, which is defined as the ratio of the total power consumption to the IT power consumption. Then the cost of the colo operator with the participation of the tenants in EDR can be represented as

$$G(\mathbf{r}, \Delta\mathbf{P}) = \alpha\left(D - \gamma\sum_{i=1}^{N}\Delta P_i\right)^+ + \sum_{i=1}^{N} r_i, \tag{10}$$

where $\Delta\mathbf{P} = [\Delta P_1, \Delta P_2, \ldots, \Delta P_N]$ is the vector containing the power reduction amount of all tenants, $\mathbf{r} = [r_1, r_2, \ldots, r_N]$ is the vector containing the reimbursements issued to all tenants, and $(x)^+ = \max\{x, 0\}$. Note that (10) implies that the power reduction from the colo must satisfy the target D during the EDR period. When the total power reduction from the tenants in the colo level $\gamma \sum_{i=1}^{N} \Delta P_i$ cannot satisfy the EDR target D, the colo operator has to start up its on-site generator to satisfy the remaining $D - \gamma \sum_{i=1}^{N} \Delta P_i$ amount of power reduction.

Define the payoff of the colo operator as the cost saving with the tenants participating into EDR, which is stated as follows:

$$\begin{aligned} V(\mathbf{r}, \Delta\mathbf{P}) &= G^0 - G(\mathbf{r}, \Delta\mathbf{P}) \\ &= \alpha \min\left\{D, \gamma \sum_{i=1}^{N} \Delta P_i\right\} - \sum_{i=1}^{N} r_i. \end{aligned} \tag{11}$$

Two observations can be made from (11). First, if the reimbursement issued to each tenant can be properly designed, the colo operator can save significant costs when participating into EDR. Second, the colo operator cannot obtain extra payoff when the total power reduction from the tenants exceeds the EDR target D. Therefore, any proposal implying a power reduction that is larger than the EDR target will be not be compensated by the colo operator.

C. Social Welfare

The social welfare, which is defined as the sum of the payoffs of tenants and the colo operator, is usually adopted to evaluate the welfare of resource allocation at the aggregate level in economics. Denote the social welfare as $\Psi(\Delta\mathbf{P})$. We have

$$\begin{aligned} \Psi(\Delta\mathbf{P}) &= \sum_{i=1}^{N} U_i(r_i, \Delta P_i) + V(\mathbf{r}, \Delta\mathbf{P}) \\ &= \alpha \min\left\{D, \gamma \sum_{i=1}^{N} \Delta P_i\right\} - \sum_{i=1}^{N} C_i(\Delta P_i). \end{aligned} \tag{12}$$

In our case, to evaluate the performance of the power reduction profile $\Delta\mathbf{P}$ at the colo level, we are interested in maximizing the social welfare defined as the following optimization problem:

$$\max_{\Delta\mathbf{P}} \quad \alpha \min\left\{D, \gamma \sum_{i=1}^{N} \Delta P_i\right\} - \sum_{i=1}^{N} C_i(\Delta P_i) \tag{13}$$

$$\text{s.t.} \quad 0 \le \Delta P_i \le \Delta P_i^{\max}, \ \forall i. \tag{14}$$

IV. Nash Bargaining Approach

The challenge of cost-effective and eco-friendly colo EDR is that although the tenants can reduce their power consumption at a lower cost or carbon footprint than that by starting up the on-site generator, they have no incentive to reduce power consumption. To solve this challenge, we use the cooperative game theory, which is often used in situations where individuals have conflicting interests but have the means and incentives to coordinate and negotiate with each other to achieve a mutually beneficial outcome. In our problem, it is natural to assume that the colo operator and tenants can communicate and coordinate regarding power reduction and reimbursement decisions. Therefore, it is natural to study the colo EDR problem using the cooperative game theory. Specifically, in this paper, we model and analyze the colo EDR problem by using Nash bargaining theory [17]. As a specific branch of the cooperative game theory, Nash bargaining theory is well-suited to study the colo EDR problem and can yield a Pareto-efficient and fair outcome, hence self-enforcing and satisfactory for all entities.

In the remainder of this section, we first introduce the background on Nash bargaining theory and analyze the bargaining solution in the one-to-to bargaining case. Then we analyze the bargaining solutions under the sequential and concurrent bargaining, respectively.

A. Preliminary

Bargaining problems represent situations in which (1) there is a conflict of interest about agreements; (2) individuals have the possibility of concluding a mutually beneficial agreement; and (3) no agreement may be imposed on any individual without its approval. In [17] and [26], Nash established the following one-to-one bargaining framework. There is a set of two players $\mathcal{N} = \{1, 2\}$. Denote $\mathcal{X}$ as the set of possible agreements and D as the disagreement outcome. The two players either reach an agreement in $\mathcal{X}$ or fail to reach agreement where D is the outcome. Each player $i \in \mathcal{N}$ has preferences, represented by a utility function u_i over $\mathcal{X} \cup D$. Then the set of possible payoffs is defined as follows

$$\begin{aligned} &U = \{(v_1, v_2) \mid u_1(x) = v_1, u_2(x) = v_2 \text{ for some } x \in \mathcal{X}\} \\ &d = (u_1(D), u_2(D)) \end{aligned}$$

A bargaining problem is a pair (U, d), and a bargaining solution assigns every bargaining problem an outcome, which can be either an agreement or the disagreement outcome.

Rather than explicitly modeling the bargaining process as the strategic or noncooperative game model, Nash bargaining theory uses an axiomatic approach, which involves abstracting away the details of the process of bargaining and considers only the set of outcomes that satisfy four “reasonable” axioms, i.e., Pareto efficiency, symmetry, invariance to affine transformations, and independence of irrelevant alternatives. Nash proved that under mild technical conditions, i.e., U is a convex and compact set and there exists some $v \in U$ such that $v_i > d_i, \forall i$, there is a unique bargaining solution called Nash bargaining solution (NBS) that satisfies the above four axioms. Specifically, NBS is defined as follows.

Definition 1: A pair of payoffs (v_1^*, v_2^*) is a NBS if it solves the following optimization problem:

$$\max_{v_1, v_2} \quad (v_1 - d_1)(v_2 - d_2) \tag{15}$$

$$\text{s.t.} \quad (v_1, v_2) \in U, \tag{16}$$

$$(v_1, v_2) \ge (d_1, d_2). \tag{17}$$

B. One-to-One Bargaining

In this subsection, we consider the case that the tenant set $\mathcal{N} = \{i\}$, i.e., there is only one tenant i in the colo. To incentivize the tenant to participate into EDR, the colo operator bargains with the tenant to determine the power reduction and the corresponding reimbursement. In the remainder of this subsection, we study the one-to-one bargaining between the colo operator and the tenant.

We first consider that the bargaining ends at the disagreement point $(r_i^0, \Delta P_i^0) = (0, 0)$. The payoffs of the colo operator and the tenant at the disagreement point are denoted as $V^0 = 0$ and $U_i^0 = 0$, respectively. Next, we consider that the colo operator and the tenant reach an agreement. We note that the payoffs obtained by the colo operator and the tenant are calculated based on (9) and (11), respectively. The Nash bargaining solution (NBS) of the one-to-one bargaining between the colo operator and the tenant is obtained by solving the following optimization problem:

$$\max_{r_i, \Delta P_i} \quad \Big(U_i(r_i, \Delta P_i) - U_i^0\Big)\Big(V(r_i, \Delta P_i) - V^0\Big) \tag{18}$$

$$\text{s.t.} \quad U_i(r_i, \Delta P_i) - U_i^0 \geq 0, \tag{19}$$

$$V(r_i, \Delta P_i) - V^0 \geq 0, \tag{20}$$

$$r_i \geq 0, \tag{21}$$

$$0 \leq \Delta P_i \leq \Delta P_i^{\max}. \tag{22}$$

We present the NBS under the one-to-one bargaining case using the following lemma.

Lemma 1: The optimal power reduction of tenant i under the one-to-one bargaining maximizes the social welfare as defined in (12) with only one tenant and is given by

$$\Delta P_i^* = \min\Big\{D/\gamma, \Delta P_i^{\max}, C_i'^{-1}(\alpha\gamma)\Big\}. \tag{23}$$

The corresponding optimal reimbursement of tenant i under the one-to-one bargaining is

$$r_i^* = \frac{\alpha \min\big\{D, \gamma \Delta P_i^*\big\} + C_i(\Delta P_i^*)}{2}. \tag{24}$$

Proof: See Appendix A. ■

The above lemma implies that the NBS maximizes the social welfare. Intuitively, in our bargaining problem the total generated social welfare through bargaining can be freely transferred between players through the payment, and thus maximizing the product of their individual payoff gains when solving the NBS can only be achieved when maximizing the overall social welfare. Based on the one-to-one bargaining result, in the following, we derive the general one-to-many NBS under two bargaining protocols, i.e., the sequential and concurrent bargaining, and analyze the connection between the NBS and the social welfare maximization problem (13)-(14).

C. Sequential Bargaining

In this subsection, we focus on the NBS under the sequential bargaining. In this case, tenants do not coordinate with each other, and the colo operator bargains with the tenants sequentially based on a predefined order $\mathcal{N} = \{1, 2, \ldots, N\}$.[1] It is straightforward to see that the sequential bargaining is essentially N coupled one-to-one bargaining problems, where the colo operator's payoff under the disagreement at stage n is the accumulated payoff obtained via the NBS in the previous $n-1$ stages. Suppose the colo operator has finished the bargaining with all previous $n-1$ tenants, and is involved in the bargaining process with tenant n. Therefore, at the current stage, the bargaining outcomes for the previous $n-1$ stages $(r_i^*, \Delta P_i^*)$, $\forall i \in \{1, 2, \ldots, n-1\}$ have been obtained. Note that if the EDR target D has been satisfied at stage n, the colo operator will not agree on any possible agreement since it will not increase its payoff by accepting the agreement. In the following, we focus on the scenario that the EDR target has not been met at stage n.

At the current bargaining stage n, if the colo operator and tenant n fail to reach an agreement, tenant n will not reduce its power consumption and thus receive no payment from the colo operator, i.e., $\Delta P_n^0 = 0$ and $r_n^0 = 0$. Then the utility of tenant n under the disagreement U_n^0 is zero. For the colo operator, its payoff under the disagreement $V^0(\mathbf{r}_{n-1}^*, \Delta \mathbf{P}_{n-1}^*)$ is the accumulated payoff using the NBS obtained in the previous $n-1$ stages, where $\mathbf{r}_{n-1}^*$ and $\Delta \mathbf{P}_{n-1}^*$ are the vectors containing the reimbursements and power reductions of previous $n-1$ tenants, respectively. Specifically, the payoff of the colo operator under the disagreement $V^0(\mathbf{r}_{n-1}^*, \Delta \mathbf{P}_{n-1}^*)$ is represented as

$$V^0\big(\mathbf{r}_{n-1}^*, \Delta \mathbf{P}_{n-1}^*\big) = \alpha \min\left\{D, \gamma \sum_{i=1}^{n-1} \Delta P_i^*\right\} - \sum_{i=1}^{n-1} r_i^*. \tag{25}$$

If the colo operator and tenant n reach an agreement $(r_n, \Delta P_n)$, then the payoff of tenant n is

$$U_n(r_n, \Delta P_n) = r_n - C_n(\Delta P_n). \tag{26}$$

Note that (26) also denotes the payoff gain of tenant n if it participates into EDR since the payoff of tenant n under the disagreement is zero. The payoff of the colo operator at stage n can be represented as the accumulated payoff obtained through n stages

$$V(\mathbf{r}_n, \Delta \mathbf{P}_n) = \alpha \min\left\{D, \gamma\left(\sum_{i=1}^{n-1} \Delta P_i^* + \Delta P_n\right)\right\} - \sum_{i=1}^{n-1} r_i^* - r_n. \tag{27}$$

Then the payoff gain of the colo operator due to the power reduction of tenant n can be calculated as

$$V(\mathbf{r}_n, \Delta \mathbf{P}_n) - V^0\big(\mathbf{r}_{n-1}^*, \Delta \mathbf{P}_{n-1}^*\big) = \alpha \min\left\{D - \gamma \sum_{i=1}^{n-1} \Delta P_i^*, \gamma \Delta P_n\right\} - r_n. \tag{28}$$

According to the analysis above, the NBS at stage n can be obtained by solving the following optimization problem.

$$\max_{r_n, \Delta P_n} \quad U_n(r_n, \Delta P_n)\Big(V(\mathbf{r}_n, \Delta \mathbf{P}_n) - V^0(\mathbf{r}_{n-1}^*, \Delta \mathbf{P}_{n-1}^*)\Big) \tag{29}$$

$$\text{s.t.} \quad U_n(r_n, \Delta P_n) \geq 0, \tag{30}$$

$$V(\mathbf{r}_n, \Delta \mathbf{P}_n) - V^0(\mathbf{r}_{n-1}^*, \Delta \mathbf{P}_{n-1}^*) \geq 0, \tag{31}$$

$$r_n \geq 0, \tag{32}$$

$$0 \leq \Delta P_n \leq \Delta P_n^{\max}. \tag{33}$$

[1] We assume that the tenants in $\mathcal{N} = \{1, 2, \ldots, N\}$ have been sorted based on the order.

Constraints (30) and (31) guarantee the individual rationality of tenants and the colo operator, respectively. Constraints (32) and (33) define the bounds of the reimbursement and power reduction for each tenant, respectively.

The NBS under the sequential bargaining obtained by solving the optimization problem (29)-(33) is shown in the following lemma, which can be proved similar to Lemma 1.

Lemma 2: When $\gamma \sum_{i=1}^{n-1} \Delta P_i^* < D$, the NBS $(r_n^*, \Delta P_n^*)$ under the sequential bargaining is represented as follows:

- The power reduction of tenant n is

$$\Delta P_n^* = \min\left\{D/\gamma - \sum_{i=1}^{n-1} \Delta P_i^*, \Delta P_n^{\max}, C_n'^{-1}(\alpha\gamma)\right\}. \tag{34}$$

- The corresponding reimbursement issued to tenant n is

$$r_n^* = \frac{\alpha \min\left\{D - \gamma \sum_{i=1}^{n-1} \Delta P_i^*, \gamma \Delta P_n^*\right\} + C(\Delta P_n^*)}{2}. \tag{35}$$

Denote $\Psi(\Delta \mathbf{P}_{n-1}^*, \Delta P_n^*)$ and $\Psi(\Delta \mathbf{P}_{n-1}^*, 0)$ as the social welfare when the colo operator and tenant n reach the mutual beneficial agreement and disagreement, respectively. We can further characterize the payoffs of the colo operator and tenant n using the social welfare in the following lemma.

Lemma 3: The payoff of tenant n using the NBS $(r_n^*, \Delta P_n^*)$ at stage n satisfies

$$U_n\left(r_n^*, \Delta P_n^*\right) = \frac{\Psi\left(\Delta \mathbf{P}_{n-1}^*, \Delta P_n^*\right) - \Psi\left(\Delta \mathbf{P}_{n-1}^*, 0\right)}{2}. \tag{36}$$

Furthermore, the payoff of the colo operator at stage n is

$$V\left(\mathbf{r}_n^*, \Delta \mathbf{P}_n^*\right) = \frac{\Psi\left(\Delta \mathbf{P}_{n-1}^*, \Delta P_n^*\right)}{2} + \frac{\Psi\left(\Delta \mathbf{P}_{n-1}^*, 0\right)}{2} - \sum_{i=1}^{n-1} r_i^*. \tag{37}$$

Proof: We rewrite the objective function (29) as

$$U_n(r_n, \Delta P_n)\big(\Psi(\Delta \mathbf{P}_{n-1}^*, \Delta P_n) - \Psi(\Delta \mathbf{P}_{n-1}^*, 0) - U_n(r_n, \Delta P_n)\big),$$

which is obviously a quadratic function of $U_n(r_n, \Delta P_n)$. Using the optimality conditions with respect to the payoff of tenant n, we can see that (36) and (37) hold. ■

In the following, we briefly discuss the properties of the NBS under the sequential bargaining.

Fairness: Note that $\Psi(\Delta \mathbf{P}_{n-1}^*, \Delta P_n^*) - \Psi(\Delta \mathbf{P}_{n-1}^*, 0)$ represents the marginal social welfare at stage n, i.e., the surplus generated when tenant n agrees to reduce its power consumption. From Lemma 3, we can see that the marginal social welfare is equally shared by tenant n and the colo operator, which satisfies the fairness property. Note that here we imply max-min fairness since we assume that the colo operator and the tenant have the same bargaining power, which can be modeled by quantifying how myopic the colo operator and tenant are so that their payoffs are comparable. When the payoffs of the colo operator and tenant are not comparable, we imply proportional fairness. The intuition behind the fairness property is that any disproportional payoff gain will be denied by the party that receives the short end of the economic surplus.

Individual rationality: From Lemma 3, we can see that the payoff of tenant n is non-negative since the marginal social welfare $\Psi(\Delta \mathbf{P}_{n-1}^*, \Delta P_n^*) - \Psi(\Delta \mathbf{P}_{n-1}^*, 0)$ is non-negative. Hence, the individual rationality of tenant n is guaranteed. Moreover, based on the max-min fairness property, the individual rationality of tenant n implies that the individual rationality of the colo operator is also satisfied. Therefore, individual rationality is guaranteed under the sequential bargaining.

Marginal social welfare maximization: Using Lemma 2 and Lemma 3, it is straightforward to see that the optimization problem (29)–(33) maximizes the marginal social welfare $\Psi(\Delta \mathbf{P}_{n-1}^*, \Delta P_n) - \Psi(\Delta \mathbf{P}_{n-1}^*, 0)$. Therefore, the NBS $(r_n^*, \Delta P_n^*)$, $n \in \mathcal{N}$ under the sequential bargaining maximizes the marginal social welfare at every stage n.

Pareto efficiency: Based on Lemma 3, increasing the payoff of one party will lower that of other parties. Therefore, we can prove that the social welfare is shared by the colo operator and tenants in a Pareto-efficient way using the NBS.

D. Concurrent Bargaining

In this subsection, we focus on the NBS under the concurrent bargaining. In this case, all tenants coordinate and bargain with the colo operator concurrently. In essence, the concurrent bargaining is that N one-to-one bargainings happen simultaneously. In the following, we analyze the NBS under this protocol. We start with the analysis under disagreement. For a tenant $n \in \mathcal{N}$, if the colo operator and the tenant cannot reach an agreement, the tenant will not turn off any server for power reduction and thus receive no reimbursement from the colo operator, i.e., $r_n^0 = 0$, and $\Delta P_n^0 = 0$. Then the payoff of the tenant under the disagreement is zero, i.e., $U_n^0 = 0$. For the colo operator, its payoff under the disagreement in the worst-case scenario is also zero, i.e., no agreement is reached between the colo operator and any tenant, and the colo operator has to use its on-site generator to satisfy the EDR target D.

Next consider the case that a tenant $n \in \mathcal{N}$ and the colo operator reach an agreement $(r_n, \Delta P_n)$. Then the payoff of tenant n is represented as

$$U_n(r_n, \Delta P_n) = r_n - C_n(\Delta P_n). \tag{38}$$

The payoff of the colo operator can be obtained after finishing the bargaining with all tenants, which is calculated as follows:

$$V(\mathbf{r}, \Delta \mathbf{P}) = \alpha \min\left\{D, \gamma \sum_{n=1}^{N} \Delta P_n\right\} - \sum_{n=1}^{N} r_n. \tag{39}$$

Since the payoffs of the colo operator and tenants under the disagreement are zero, (38) and (39) also represent the payoff gains of the tenants and the colo operator, respectively. Therefore, we can obtain the NBS under the concurrent bargaining by solving the following optimization problem:

$$\max_{\mathbf{r}, \Delta \mathbf{P}} \quad V(\mathbf{r}, \Delta \mathbf{P}) \prod_{n=1}^{N} U_n(r_n, \Delta P_n) \tag{40}$$
$$\text{s.t.} \quad (30) - (33), \forall n.$$

We describe the NBS under the concurrent bargaining in the following lemma.

Lemma 4: Let $\mathcal{N}'$ be the set of tenants that participates into the EDR. The reimbursement issued to tenant $n \in \mathcal{N}'$ for reducing its power consumption is

$$r_n^* = C_n(\Delta P_n^*) + \frac{1}{N'+1}\left[\alpha \min\left\{D, \gamma \sum_{i=1}^{N} \Delta P_i^*\right\} - \sum_{i=1}^{N} C_i(\Delta P_i^*)\right]. \quad (41)$$

Proof: See Appendix B. ■

Regarding the power reduction profile $\Delta\mathbf{P}^*$ under the concurrent bargaining, we have the following lemma.

Lemma 5: When $D \geq \gamma \sum_{i=1}^{N} \Delta P_i^{\max}$, the power reduction profile $\Delta\mathbf{P}^* = [\Delta P_1^*, \Delta P_2^*, \ldots, \Delta P_N^*]$ under the concurrent bargaining is identical to that under the sequential bargaining.

Proof: Given that the concurrent bargaining problem is equivalent to the social welfare maximization problem which will be proved below, we have that both the sequential bargaining and concurrent bargaining maximize the social welfare when $D \geq \gamma \sum_{i=1}^{N} \Delta P_i^{\max}$. Moreover, based on Assumption 1, the social welfare maximization problem is strictly convex and has a unique optimal solution, which also solves the sequential bargaining problem and the concurrent bargaining problem. ■

The requirement $D \geq \gamma \sum_{i=1}^{N} \Delta P_i^{\max}$ is to ensure that the colo operator will be interested in initiating the bargaining process with all tenants under both bargaining protocols, especially the sequential bargaining. Otherwise, if D has been satisfied by the first few tenants under the sequential bargaining, the remaining tenants will automatically be rejected in the bargaining. Therefore, the condition $D \geq \gamma \sum_{i=1}^{N} \Delta P_i^{\max}$ guarantees that the power reduction profile under the sequential bargaining is independent of the bargaining order.

In the following, we briefly discuss the properties of the NBS under the concurrent bargaining.

Fairness: From (41), we observe that the max-min fairness is satisfied under the concurrent bargaining, i.e., the colo operator and the tenants that participate into EDR have equitable payoff gains. Specifically, each participant including the tenants and the colo operator receives $1/(N'+1)$ fraction of the social welfare.

Individual rationality: From (41), we observe that the reimbursement (41) should compensate the inconvenience cost incurred by the tenant during the EDR period. Since non-negative social welfare will be generated during EDR, the tenants should receive non-negative payoff, implying that individual rationality is satisfied. Due to the max-min fairness property, the colo operator, who has the same utility gain as the participating tenants, receives non-negative utility gain, and thus the individual rationality of the colo operator is guaranteed. Therefore, the individual rationality is guaranteed under the concurrent bargaining.

Social welfare maximization: We now show the relationship between the NBS and the social welfare maximization problem in the concurrent bargaining. By substituting the optimal reimbursement (41) shown in Lemma 4 into the concurrent bargaining problem, it is straightforward to see that the optimization problem is equivalent to the social welfare maximization problem. Therefore, the NBS under the concurrent bargaining maximizes the social welfare, i.e., $\Delta\mathbf{P}^* = \mathrm{argmax}_{\Delta\mathbf{P}} \Psi(\Delta\mathbf{P})$.

Pareto efficiency: Based on the social welfare maximization property and the definition of Pareto efficiency, we can see that the social welfare is split in a Pareto-efficient manner since increasing the payoff of any player will decrease the payoffs of other players.

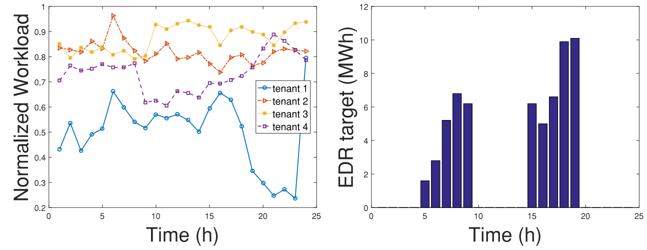

(a) (b)

Fig. 1. (a) Typical one-day workload traces of 4 tenants; (b) EDR reduction target.

V. Performance Evaluation

A. Simulation Setup

Colocation data center setup: We consider a colo with 4 tenants located at Ashburn, VA, which is a major data center market served by PJM interconnection. Each tenant i has $M_i = 2000$ homogenous servers. We set the idle power and peak power of each server as $P_i^{\text{idle}} = 150$W and $P_i^{\text{peak}} = 250$W, respectively. Furthermore, we set the PUE of the colo as $\gamma = 1.5$. Therefore, the corresponding power reduction of colo from tenant i is 1.5kW if the tenant reduces 1kW power consumption. The colo has a diesel generator, whose unit cost is $\alpha = \$0.5/\text{kWh}$.

Tenant workload description: We use the trace data of Google cluster as the workloads of the tenants [27]. All workloads have been normalized with respect to the maximum service capacity of each tenant. The normalized workloads for the tenants are depicted in Fig. 1a.

Tenant cost: The inconvenience cost of each tenant might include the delay cost and the wear-and-tear cost [5]. Without loss of generality, we focus on the wear-and-tear cost and model it as a linear function $C_i(\Delta P_i) = w_i \Delta P_i,\ \forall i \in \mathcal{N}$, where w_i denotes the unit wear-and-tear cost and is set to be uniformly distributed in $(0, 0.6]$\$/kW.

EDR setup: We use the EDR signals issued by PJM Interconnection on January 22, 2014 for our simulations [28]. The EDR target in each hour is depicted in Fig. 1b. As shown in Fig. 1b, there are 10 EDR events each lasting 1 hour. The EDR is called during 5–9 am and 3–7 pm.

B. Simulation Results

Social welfare maximization: We set the scenario that the colo operator only uses the on-site backup generator to satisfy the EDR targets without incentivizing any tenant as the benchmark. When tenants are incentivized to participate into EDR, according to our analysis, the social welfare increases comparing to the benchmark. In Fig. 2, we compare the social

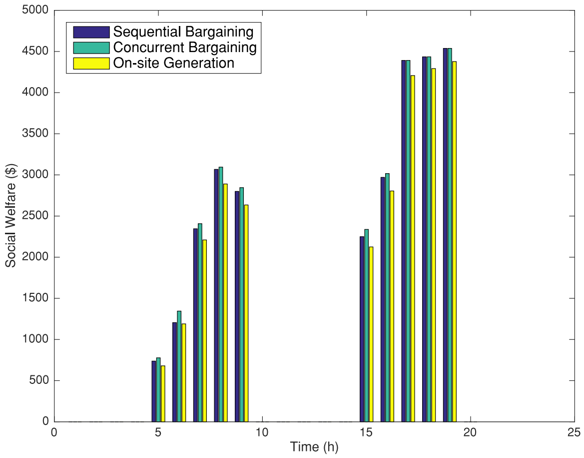

Fig. 2. Social welfare comparison of the sequential bargaining, concurrent bargaining and benchmark.

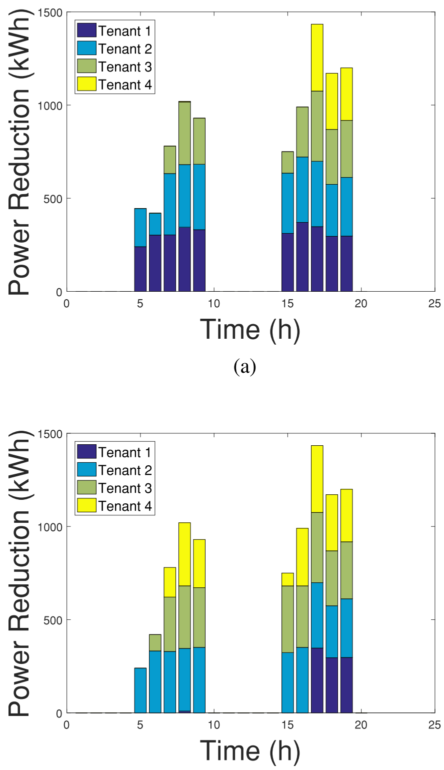

Fig. 3. Power reduction profiles under (a) the sequential bargaining and (b) the concurrent bargaining.

welfare under the sequential bargaining, concurrent bargaining, and benchmark. We observe that by incentivizing the tenants' participation, the social welfare increases significantly using our bargaining approach. Moreover, when the condition $D \geq \sum_{i=1}^{4} \Delta P_i^{\max}$ holds as is the case for the last three EDR periods, both the sequential and concurrent bargaining maximize the social welfare as shown in Fig. 2. As for the first seven EDR periods where $D < \sum_{i=1}^{4} \Delta P_i^{\max}$, the concurrent bargaining achieves higher social welfare than the sequential bargaining since it always maximizes the social welfare.

The NBS: In Fig. 3, we show the power reduction profile of all tenants at each hour. Specifically, we depict the power reduction profiles under the sequential and concurrent bargaining in Fig. 3a and Fig. 3b, respectively. Since the colo operator can start up its on-site generator when the power reductions of tenants are not sufficient, EDR targets will always be satisfied. Consider the first seven EDR periods. We observe that not all tenants participate into EDR since for these EDR periods, $D < \sum_{i=1}^{4} \Delta P_i^{\max}$. According to our analysis in Lemma 5, the bargaining order matters under the sequential bargaining. However, under the concurrent bargaining, we observe that tenant 2 who has the lowest unit inconvenience cost is preferred by the colo operator. Then we consider the last three EDR periods. We observe that all tenants are incentivized to participate into EDR since $D \geq \sum_{i=1}^{4} \Delta P_i^{\max}$. Note that the power reduction profiles under the sequential and concurrent bargaining are identical, which matches our analysis in Lemma 5. The corresponding reimbursements under the sequential and concurrent bargaining are depicted in Fig. 4. We observe that all the tenants can receive reimbursement from the colo operator once they participate into EDR.

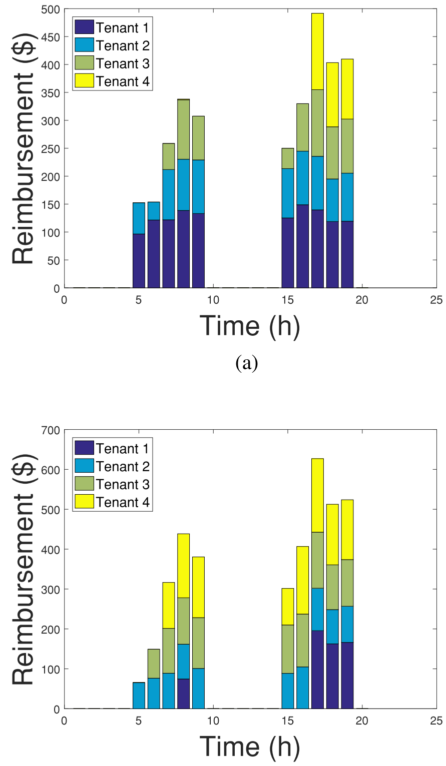

Fig. 4. Reimbursement profiles to the tenants under (a) the sequential bargaining and (b) the concurrent bargaining.

Individual rationality: We show the payoff received by each tenant in Fig. 5, where Fig. 5a and Fig. 5b depict the payoffs of the tenants under the sequential and concurrent bargaining, respectively. We observe that all tenants have non-negative payoffs once they reduce their power consumption during the EDR event. Therefore, the individual rationality of tenants is guaranteed. Then from Fig. 6, we can see that the payoff of the colo operator is non-negative under both the sequential and concurrent bargaining. Therefore, individual rationality of the system is guaranteed in both bargaining protocols.

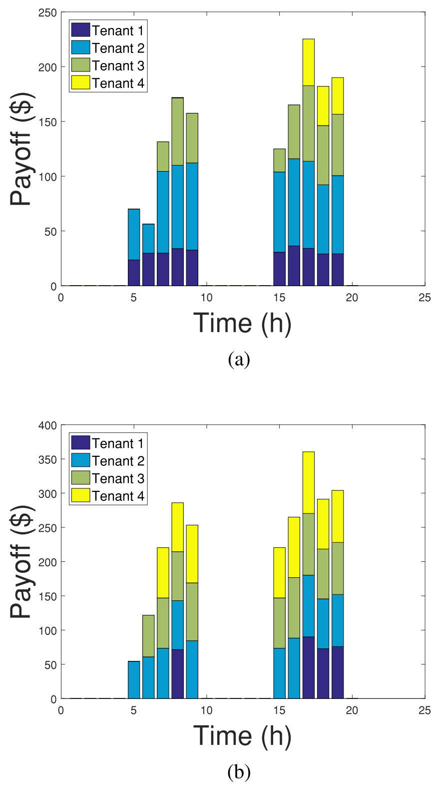

Fig. 5. Payoffs of the tenants under (a) the sequential bargaining and (b) the concurrent bargaining.

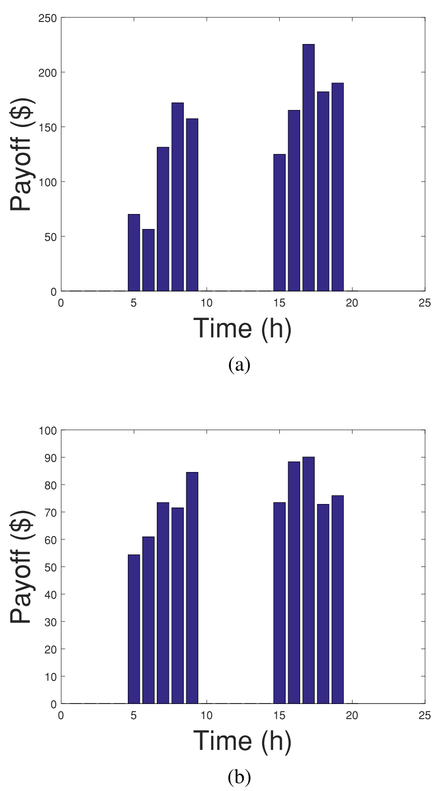

Fig. 6. Payoff of the colo operator under (a) the sequential bargaining and (b) the concurrent bargaining.

Fairness: We first consider the sequential bargaining scenario. By comparing Fig. 5 and Fig. 6a, we can observe that the payoff of the colo operator is equal to that of all the participating tenants. The reason is that the marginal social welfare is equally shared by the tenants and the colo operator. Then, we analyze the concurrent bargaining scenario. We can see that the payoffs of the colo operator and each participating tenant are the same, i.e., $1/(N'+1)$ fraction of the social welfare. Therefore, the max-min fairness is guaranteed using the proposed approach.

Payoff comparison between different protocols: From Fig. 5, we observe that the total payoff received by tenants is higher in the concurrent bargaining than that in the sequential bargaining for all EDR periods. This is because by simultaneously bargaining with the colo operator, the tenants have a larger negotiation power. However, it requires all tenants to coordinate beforehand in order to act simultaneously. On the other hand, the colo operator receives a lower payoff in the concurrent bargaining than that in the sequential bargaining for all EDR periods as shown in Fig. 6, which is reasonable as the NBS is Pareto-efficient as analyzed in Section IV.

VI. Conclusion

This paper has investigated how to incentivize the tenants to participate into colo EDR, and how the economic surplus generated during EDR should be shared by the colo operator and tenants. We have proposed an approach based on Nash bargaining theory to coordinate the tenants participating into EDR and derived solutions that are Pareto-efficient and fair under the sequential and concurrent bargaining protocol. Trace-driven simulations have been conducted to verify our theoretical analysis, which shows that the colo operator saves significant costs and the tenants receive reimbursements by participating into EDR.

Appendix A
Proof of Lemma 1

By taking the logarithm of the objective function of (18), we can see that the one-to-one bargaining problem has a unique solution due to its strict concavity. Denote π_i to be the tenant i's payoff as defined in (9) and $\Psi(\Delta P_i)$ to be the social welfare as defined in (12). We can rewrite the problem (18) as the following equivalent problem:

$$\max_{\Delta P_i, \pi_i} \quad (\Psi(\Delta P_i) - \pi_i)\pi_i$$
$$\text{s.t.} \quad \Delta P_i \in \left[0, \Delta P_i^{\max}\right], \Psi(\Delta P_i) - \pi_i \geq 0, \pi_i \geq 0.$$

Then we can solve the above problem by sequentially optimizing π_i and ΔP_i, and then deduce the optimal results ΔP_i^* and r_i^*.

Appendix B
Proof of Lemma 4

Let $\mathcal{N}' \subseteq \mathcal{N}$ denote the set of N' tenants who participates into EDR. For any participating tenant $n \in \mathcal{N}'$, it is obvious that $U_n(r_n, \Delta P_n) > 0$. Otherwise, they will not be involved into the bargaining process. For the tenants $n \in \mathcal{N} \setminus \mathcal{N}'$, they cannot benefit from power reduction and thus have no incentive to participate in EDR. Therefore, we focus on the tenants in $\mathcal{N}'$. Let π_n denote the payoff of the tenant n as defined in (38) and $\Psi(\Delta\mathbf{P})$ denote the social welfare as defined in (12). By taking the log of the objective function (40), the NBS problem is transformed into the following:

$$\begin{aligned} \min \quad & \ln\left(\Psi(\Delta\mathbf{P}) - \sum_{n\in\mathcal{N}'} \pi_n\right) + \sum_{n\in\mathcal{N}'} \ln(\pi_n) \\ \text{s.t.} \quad & \pi_n > 0, 0 < \Delta P_n \le \Delta P_n^{\max}, \quad \forall n \in \mathcal{N}'. \end{aligned}$$

In the following, we will solve the above problem by sequentially optimizing $\Delta\mathbf{P}$ and $\pi_n, \forall n$. Given the optimal load reduction profile $\Delta\mathbf{P}^*$, we can use the first-order optimality condition to obtain π_i^* as

$$\pi_i^* = \Psi\big(\Delta\mathbf{P}^*\big) - \sum_{n\in\mathcal{N}'} \pi_n^*,$$

which is the same for all tenant $i \in \mathcal{N}'$. Therefore, we can obtain

$$\pi_i^* = \frac{1}{N'+1}\Psi(\Delta\mathbf{P}^*).$$

Substituting the definition of the payoff gain π_i and social welfare Ψ into the above equation, we can obtain the result (41).

References

[1] *Data Center Efficiency Assessment*, NRDC, New York, NY, USA, Aug. 2014. [Online]. Available: https://www.nrdc.org/energy/files/data-center-efficiency-assessment-IP.pdf

[2] R. E. Brown *et al.*, "Report to congress on server and data center energy efficiency: Public law 109-431," Lawrence Berkeley Nat. Lab., Berkeley, CA, USA, Tech. Rep. LBNL-363E, 2008.

[3] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for Internet-scale systems," in *Proc. ACM SIGCOMM*, Barcelona, Spain, 2009, pp. 123–134.

[4] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Greening geographical load balancing," in *Proc. ACM SIGMETRICS*, San Jose, CA, USA, 2011, pp. 233–244.

[5] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1378–1391, Oct. 2013.

[6] Z. Liu *et al.*, "Renewable and cooling aware workload management for sustainable data centers," in *Proc. ACM SIGMETRICS*, London, U.K., 2012, pp. 175–186.

[7] Y. Guo and Y. Fang, "Electricity cost saving strategy in data centers by using energy storage," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 6, pp. 1149–1160, Jun. 2013.

[8] Y. Guo, Y. Gong, Y. Fang, P. P. Khargonekar, and X. Geng, "Energy and network aware workload management for sustainable data centers with thermal storage," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 8, pp. 2030–2042, Aug. 2014.

[9] K. Managan. *Demand Response: A Market Overview*. Accessed on Feb. 2014. [Online]. Available: http://www.institutebe.com/

[10] *Ensuring U.S. Grid Security and Reliability: U.S. EPA's Proposed Emergency Backup Generator Rule*, EnerNOC, Boston, MA, USA, 2013. [Online]. Available: http://www.sustainableplant.com

[11] A. Misra. *Responding Before Electric Emergencies*. Accessed on Jan. 2016. [Online]. Available: http://www.afcom.com/digital-library/pub-type/communique/responding-before-electric-emergencies/

[12] G. Nelson. *Air Pollution: Dirty Diesel Generators Test EPA, Demand-Response Industry*. Accessed on Jul. 10, 2012. [Online]. Available: http://www.eenews.net/stories/1059967047

[13] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *Proc. IEEE IGCC*, Dallas, TX, USA, 2014, pp. 1–10.

[14] N. Chen, X. Ren, S. Ren, and A. Wierman, "Greening multi-tenant data center demand response," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 43, no. 2, pp. 36–38, 2015.

[15] *Data Center Map*. Accessed on Jan. 2016. [Online]. Available: http://www.datacentermap.com/

[16] S. Ren and M. A. Islam, "Colocation demand response: Why do I turn off my servers," in *Proc. USENIX ICAC*, Philadelphia, PA, USA, 2014, pp. 201–208.

[17] J. F. Nash, Jr., "The bargaining problem," *Econometrica J. Econ. Soc.*, vol. 18, no. 2, pp. 155–162, 1950.

[18] L. Zhang, S. Ren, C. Wu, and Z. Li, "A truthful incentive mechanism for emergency demand response in colocation data centers," in *Proc. IEEE INFOCOM*, Hong Kong, 2015, pp. 2632–2640.

[19] N. H. Tran *et al.*, "Demand response of data centers: A real-time pricing game between utilities in smart grid," in *Proc. USENIX Workshop Feedback Comput.*, Philadelphia, PA, USA, 2014.

[20] M. A. Islam, H. Mahmud, S. Ren, and X. Wang, "Paying to save: Reducing cost of colocation data center via rewards," in *Proc. IEEE HPCA*, Burlingame, CA, USA, 2015, pp. 235–245.

[21] M. A. Islam, S. Ren, and X. Wang, "GreenColo: A novel incentive mechanism for minimizing carbon footprint in colocation data center," in *Proc. IEEE IGCC*, Dallas, TX, USA, 2014, pp. 1–8.

[22] Y. Guo and M. Pan, "Coordinated energy management for colocation data centers in smart grids," in *Proc. IEEE SmartGridComm*, Miami, FL, USA, 2015, pp. 840–845.

[23] N. Lacetera, B. J. Larsen, D. G. Pope, and J. R. Sydnor, "Bid takers or market makers? The effect of auctioneers on auction outcome," *Amer. Econ. J. Microecon.*, vol. 8, no. 4, pp. 195–229, Nov. 2016.

[24] C. Lefurgy, X. Wang, and M. Ware, "Server-level power control," in *Proc. IEEE ICAC*, Jacksonville, FL, USA, 2007, p. 4.

[25] R. Nelson, *Probability, Stochastic Processes, and Queueing Theory: The Mathematics of Computer Performance Modeling*. New York, NY, USA: Springer, 2013.

[26] J. Nash, "Two-person cooperative games," *Econometrica J. Econ. Soc.*, vol. 21, no. 1, pp. 128–140, 1953.

[27] J. Wilkes. (Nov. 2011). *More Google Cluster Data, Google Research Blog*. [Online]. Available: http://googleresearch.blogspot.com/2011/11/more-google-cluster-data.html

[28] *PJM Interconnection*. Accessed on Mar. 2016. [Online]. Available: http://www.pjm.com/markets-and-operations/demand-response.aspx

Yuanxiong Guo (M'14) received the B.Eng. degree in electronics and information engineering from the Huazhong University of Science and Technology, Wuhan, China, in 2009, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Florida, Gainesville, FL, USA, in 2012 and 2014, respectively. Since 2014, he has been an Assistant Professor with the School of Electrical and Computer Engineering, Oklahoma State University, Stillwater, OK, USA. His current research interests include resource management for networked systems with particular applications to power networks, data centers, Internet of things, and wireless networks, and cybersecurity with emphasis on smart grid and big data.

Hongning Li (M'14) received the B.S. degree from the School of Science, Xidian University, Xi'an, China, in 2007, the M.S. degree from the School of Telecommunications Engineering, Xidian University in 2010, and the Ph.D. degree from the School of Computer Science and Technology, Xidian University in 2014. Her research interests are in the general area of security and privacy in cognitive radio networks and wireless networks and security.

Miao Pan (M'12) received the B.Sc. degree in electrical engineering from the Dalian University of Technology, China, in 2004, the M.A.Sc. degree in electrical and computer engineering from the Beijing University of Posts and Telecommunications, China, in 2007, and the Ph.D. degree in electrical and computer engineering from the University of Florida in 2012. He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, University of Houston. He was an Assistant Professor with the Computer Science, Texas Southern University from 2012 to 2015. His research interests include cognitive radio networks, cyber-physical systems, and cybersecurity.