Demand Response for Data Centers in Deregulated Markets: A Matching Game Approach

Shahab Bahrami[†], Vincent W.S. Wong[†], and Jianwei Huang[⋆]
[†]Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, Canada
[⋆]Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong
email: {bahramis, vincentw}@ece.ubc.ca, jwhuang@ie.cuhk.edu.hk

***Abstract*—With the fast development of deregulated electricity markets, a user can enter a contract with a utility company that offers the best rates among multiple competing utility companies. Meanwhile, a utility company is motivated to increase its market share by offering demand response programs with real-time pricing (RTP), which can help its customers to manage their energy usage and save money. In this paper, we focus on the demand response program in deregulated markets for data centers, which are often flexible in scheduling their workloads. We capture the stochastic workload process in a data center as a multiclass queuing system. We model the coupled decisions of utility company choices and workload scheduling of data centers as a *many-to-one matching game with externalities*. Analyzing such a game is challenging, as there does not exist a general algorithm that guarantees to find a stable outcome, where no player has an incentive to unilaterally change its strategy. We show that the data center matching game admits an *exact potential* function, whose local minima correspond to the stable outcomes of the game. We develop an algorithm that can guarantee to converge to a stable outcome. Compared with the scenario without utility company choices and demand response, simulations show that our proposed algorithm can reduce the cost of data centers by 15.4% and increase the revenue of those utility companies with lower tariffs by up to 82%. The peak-to-average ratio (PAR) of the customers' load demand is also reduced by 7.2%.**

I. Introduction

Recent advances in small-scale power plants and the integration of communication technologies into the power networks have enabled utility companies to enter the *deregulated electricity markets* [1]. In such a market, a customer is free to purchase electricity from one of several competing utility companies. Meanwhile, the utility companies can take advantage of such flexibility and choose their retail prices. This motivates the utility companies to deviate from today's common practice of flat-rate pricing and implement real-time pricing (RTP) [2], [3]. By implementing a demand response program with RTP, the utility companies can benefit from a smoother energy demand profile, and achieve a higher revenue by attracting more customers. The customers, on the other hand, can take advantage of the lower prices.

In this paper, we focus on the choices of utility company as well as demand response of a special type of customers – data centers. Data center owners often closely monitor and control the demand of their information technology (IT) equipment (e.g., servers, routers) and cooling facilities. Many typical workloads in data centers (e.g., high complexity scientific computations, and data analytics) are delay-tolerant, and hence may be rescheduled to off-peak hours [4]. This motivates a recent rich body of literature on the demand response algorithm design for data centers (e.g., [5]–[9]).

We can classify the related literature into two main threads. The first thread is concerned with the solution approaches for the workload management problem in data centers. Different techniques such as stochastic optimization [5] and convex optimization [6], [7] are used to tackle the workload management problem. In these works, the energy price of the utility company is fixed and the focus is to solve a cost minimizing problem for the data centers. The second thread is concerned with modeling the active pricing decisions of the utility companies for data centers. Wang *et al.* in [8] applied a two-stage optimization method to model the interactions of a utility company's pricing optimization and the data centers' energy demand optimization. The approach in [8] may not be directly applied to the case with multiple data centers. Tran *et al.* addressed a related problem in [9], where the utility companies need to obtain the closed-form solution to the data centers' cost minimization problem. This may not be always feasible in practice.

In this paper, we study the emerging deregulated markets, where multiple utility companies compete to supply electricity to the same group of geographically dispersed data centers. Each data center can choose which utility company to sign the contract and schedule its workloads to minimize its payment. If the utility companies adopt the RTP scheme, the data centers' payments will depend on the amount and the time of their electricity consumptions. Thus, the decisions of data centers are coupled among each other and with the pricing decisions of the utility companies. We capture the data centers' coupled decisions of utility company choices and workload scheduling as a *many-to-one matching game*. The underlying mechanism is a *matching with externalities* [10], as the payments of the data centers choosing the same utility company depend on the workload scheduling of each other. We characterize the stable outcome of the game, where no data center has an incentive to change its matched utility company and workload schedule *unilaterally*. Such characterization is quite challenging, as there does not exist a general algorithm that can guarantee to find a stable outcome in matching with externalities.

The contributions of this paper are as follows:

- *Data Center Workload Model*: We approximate the workloads' arrivals and executions in a data center over the contract period by a multiclass queuing system. Such a model enables us to schedule the number of operating servers, meanwhile satisfying the quality-of-service (QoS) requirements in executing different service requests.

- *Solution Method and Algorithm Design*: We characterize an exact potential function of the matching game, and show that the stable outcomes of the game correspond to the local minima of the potential function. We develop an algorithm that can be executed by the data centers and utility companies in a distributed fashion. We prove that the algorithm converges to a stable outcome of the game.
- *Performance Evaluation*: We perform simulations on a market with 50 data centers and 10 utility companies. The results show that the proposed algorithm reduces the cost of data centers and the peak-to-average ratio (PAR) in the aggregate demand of data centers connected to the same utility company by 15.4% and 7.2%, respectively. Furthermore, the utility companies that offer lower energy prices can increase their revenue by up to 82%, as they can attract more data centers as customers.

The rest of this paper is organized as follows. Section II introduces the system model. In Section III, we propose a matching game model for the data centers interaction. We also develop a distributed algorithm to obtain a stable outcome. In Section IV, we evaluate the performance of the proposed algorithm. Section V concludes the paper.

II. System Model

Consider a system with D data centers and U utility companies. Let $\mathcal{D} = \{1, \ldots, D\}$ and $\mathcal{U} = \{1, \ldots, U\}$ denote the set of data centers and the set of utility companies, respectively. Data center $d \in \mathcal{D}$ can purchase electricity from a utility company in a predetermined set $\mathcal{U}_d \subseteq \mathcal{U}$. Utility company $u \in \mathcal{U}$ is able to serve a predetermined subset of data centers denoted by $\mathcal{D}_u \subseteq \mathcal{D}$. Sets $\mathcal{U}_d$, $d \in \mathcal{D}$ and $\mathcal{D}_u$, $u \in \mathcal{U}$ are determined based on the the topology of the network and the geographic locations of the utility companies and data centers.

Fig. 1 (a) shows a system with five data centers and three utility companies. Fig. 1 (b) shows the corresponding bipartite graph representation. Each data center d possesses an energy management system (EMS), which is connected to the utility companies in set $\mathcal{U}_d$ via a two-way communication network. The EMS enables exchanging information such as the energy consumption of the corresponding data center and the energy price for entering a bilateral contract. In deregulated markets, a data center can enter a bilteral contract with one utility company to purchase electricity. Meanwhile, a utility company can supply electricity to multiple data centers. We can capture the contracts between data centers and utility companies as a many-to-one matching [11], which is defined as follows.

Definition 1: A many-to-one matching among the data centers and utility companies is a function $m : \mathcal{D} \cup \mathcal{U} \rightarrow \mathcal{P}(\mathcal{D} \cup \mathcal{U})$, where $m(u) \subseteq \mathcal{D}_u$ represents the set of data centers served by utility company $u \in \mathcal{U}$, and $m(d) \subseteq \mathcal{U}_d$ with $|m(d)| = 1$ represents the utility company choice of data center $d \in \mathcal{D}$. Here, $|\cdot|$ denotes the cardinality and $\mathcal{P}$ is the power set of a set.

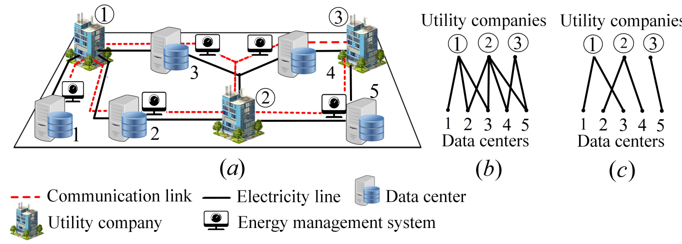

Fig. 1. (a) A system composed of five data centers equipped with EMS and three utility companies; (b) the corresponding bipartite graph representation; (c) a feasible many-to-one matching.

Fig. 1 (c) shows a feasible many-to-one matching. Although short-term contracts are not common for residential customers in today's deregulated markets, large loads such as data centers can enter a contract with utility companies for a period from a few hours to several days [4], [12]. We assume that a data center can enter a short-term contract (e.g., one day) with a utility company. Without loss of generality, we consider the same contract period for all data centers. We divide the intended contract period into a set $\mathcal{T} = \{1, \ldots, T\}$ of T time slots with an equal length, e.g., one time slot is 15 minutes.

In matching m, utility company u sets its retail price $p^{\rm r}_u(t)$, $t \in \mathcal{T}$, for the contracts with the data centers in set $m(u)$. Data center d specifies its demand profile $e_d(t)$, $t \in \mathcal{T}$, to be supplied by its utility company choice $m(d)$.

1) Contract Pricing Model: In general, a utility company purchases electricity from the wholesale market with a price $p(t)$, $t \in \mathcal{T}$, determined from the demand-supply balance in the wholesale market. The utility companies may offer dynamic electricity rates to the flexible large loads such as data centers. In the dynamic pricing scheme, the retail price of utility company u depends on the time of energy consumption, as well as the total energy demand from its customers. In particular, the retail price of utility company $u \in \mathcal{U}$ in time slot $t \in \mathcal{T}$ and matching m is an increasing function of the total energy demand $e_u(t) = e^{\rm other}_u(t) + \sum_{d \in m(u)} e_d(t)$, where $e^{\rm other}_u(t)$ denotes the demand in time slot t for the customers other than the data centers served by utility company u. The retail price is greater than the wholesale price, in order to guarantee a positive profit for the utility company. We consider the linear approximation of the retail price (around the wholesale price) as follows [13]:

$$p^{\rm r}_u(e_u(t), m) = p(t) + \kappa_u(t)\, e_u(t), \quad u \in \mathcal{U},\, t \in \mathcal{T}, \qquad (1)$$

where $\kappa_u(t)$, $u \in \mathcal{U}$, $t \in \mathcal{T}$, are nonnegative coefficients with the unit of \$/MWh^2. The utility companies can determine $\kappa_u(t)$ according to the cost of supplying electricity.

The dynamic pricing scheme in (1) motivates data center d towards scheduling its energy demand $e_d(t)$, $t \in \mathcal{T}$ to benefit from the retail price fluctuations. Next, we describe how a data center can manage its energy demand.

2) Data Center's Operation Model: A data center offers different service classes (e.g., video streaming, data analytics) to its customers. Let $\mathcal{C}_d = \{1, \ldots, C_d\}$ denote the set of service classes that are offered in data center $d \in \mathcal{D}$, where $C_d = |\mathcal{C}_d|$. To meet the QoS requirements, the delay in executing a workload is limited within a certain range. Let $\Delta_{c,d}$ denote the delay that the execution of a workload of service class $c \in \mathcal{C}_d$ can tolerate. A small $\Delta_{c,d}$ corresponds to the interactive services that are inflexible due to stringent delay requirements, such as web search, online gaming, and video streaming. A large $\Delta_{c,d}$ corresponds to the delay-tolerant services, such as scientific applications, data analytics, and file processing [5].

We now discuss how a data center can schedule the number of operating servers to meet the QoS requirements. We assume that both the workloads' inter-arrival time and execution time follow the exponential distribution [8], [9]. For data center d, a workload requesting service class $c \in \mathcal{C}_d$ arrives with an average rate of $\lambda_{c,d}(t)$, $t \in \mathcal{T}$, workloads per time slot. Let $\sigma_{c,d}$ denote the average time it takes for a server in data center d to execute a workload requesting service class c. Let $n_d(t)$ denote the average number of operating servers of data center d in time slot t. If all the servers in data center d execute the workloads of service class c, the corresponding average execution rate in time slot t is obtained as $\mu_{c,d}(t) = \frac{n_d(t)}{\sigma_{c,d}}$. However, the servers in data center d execute C_d service classes. Let $\rho_d(t) = \sum_{c\in\mathcal{C}_d} \frac{\lambda_{c,d}(t)}{\mu_{c,d}(t)}$ denote the server utilization of data center d in time slot t. The proportion of time that the servers are busy to execute the workloads of service classes other than c is $\rho_d(t) - \rho_{c,d}(t)$, where $\rho_{c,d}(t) = \frac{\lambda_{c,d}(t)}{\mu_{c,d}(t)}$. Hence, the proportion of time that the servers are busy to execute the workloads of service class c is $1 - (\rho_d(t) - \rho_{c,d}(t))$. Thus, we can model a data center by a multiclass $M/M/1$ queuing system, where the execution rate of the workloads of service class c is $\overline{\mu}_{c,d}(t) = (1 - (\rho_d(t) - \rho_{c,d}(t)))\mu_{c,d}(t)$.

We can show that a workload of service class c experiences the maximum expected waiting time either at the beginning or at the end of each time slot t [14]. The waiting time at the beginning of time slot t depends on the number of workloads whose jobs have not been completed yet. For simplicity, we will consider the steady state approximation of the average number of workloads at the end of time slot $t-1$. Hence, an incoming workload of service class c experiences the average waiting time of $\big(1 + \frac{\lambda_{c,d}(t-1)}{\overline{\mu}_{c,d}(t-1)-\lambda_{c,d}(t-1)}\big)/\overline{\mu}_{c,d}(t)$ at the beginning of time slot t. To satisfy the delay requirement, it should be less than or equal to $\Delta_{c,d}$. We can rewrite $\overline{\mu}_{c,d}(t-1)$ and $\overline{\mu}_{c,d}(t)$ in terms of $n_d(t-1)$ and $n_d(t)$, respectively. We can approximate $\frac{\rho_d(t-1)-\rho_{c,d}(t-1)}{\rho_d(t)-\rho_{c,d}(t)} \approx 1$ due to the small changes in the proportion of time that the servers are busy to execute the workloads of service classes other than c over two consecutive time slots. By performing some algebraic manipulations, for time slots $t-1$, $t \in \mathcal{T}$, we obtain

$$\frac{(\sigma_{c,d}/\Delta_{c,d})}{n_d(t)} + \frac{\sum_{c\in\mathcal{C}_d}\sigma_{c,d}\lambda_{c,d}(t-1)}{n_d(t-1)} \leq 1,\ c \in \mathcal{C}_d, d \in \mathcal{D}. \quad (2)$$

Inequality (2) implies that if the number of servers in time slot $t-1$ is small, then the number of workloads with incomplete jobs increases. Hence, the number of servers in time slot t should be sufficiently large to meet the maximum delay constraint at the beginning of time slot t.

We use the steady state condition to approximate the waiting time of an incoming workload of service class c at the end of time slot t. Thus, we need to satisfy $\frac{1}{\overline{\mu}_{c,d}(t)-\lambda_{c,d}(t)} \leq \Delta_{c,d}$. By performing some algebraic manipulations, we can rewrite the delay requirement as

$$\frac{\sigma_{c,d}}{\Delta_{c,d}} + \sum_{c\in\mathcal{C}_d} \sigma_{c,d}\lambda_{c,d}(t) \leq n_d(t),\ \ c \in \mathcal{C}_d,\ d \in \mathcal{D},\ t \in \mathcal{T}. \quad (3)$$

In data center d, the number of operating servers is upper bounded by $n_d^{\max}$. That is

$$n_d(t) \leq n_d^{\max}, \qquad d \in \mathcal{D},\ t \in \mathcal{T}. \quad (4)$$

Let E_d^{idle} and E_d^{peak} denote the average idle energy consumption and the peak energy consumption per time slot of a server in data center d, respectively. The average energy demand of data center $d \in \mathcal{D}$ in time slot $t \in \mathcal{T}$ can be obtained by

$$e_d(t) = \eta_d(t)\, n_d(t) \left(E_d^{\text{idle}} + (E_d^{\text{peak}} - E_d^{\text{idle}})\rho_d(t)\right), \quad (5)$$

where $\eta_d(t)$ is the power usage effectiveness (PUE) of data center d in time slot t. The typical value of $\eta_d(t)$ for a data center is between 1.5 and 2 [15].

III. Problem Formulation and Algorithm Design

Let $\boldsymbol{a}_d = (n_d(t),\ t \in \mathcal{T})$ denote the scheduling decision vector of data center d. Based on the pricing scheme in (1), the contract payment of data center d to utility company $u = m(d)$ depends on the matching m and the joint decision vector $\boldsymbol{a} = (\boldsymbol{a}_d,\ d \in \mathcal{D})$ of all data centers. Hence, we have

$$c_d(\boldsymbol{a}, m) = \sum_{t\in\mathcal{T}} e_d(t)\, p_u^{\text{r}}(e_u(t), m). \quad (6)$$

The decision making of data centers are interdependent. We capture the interactions among the data centers as a *many-to-one matching game*, which is defined as follows [11]:

Game 1 *Data Center Many-to-One Matching Game:*

- ***Players***: The set of all data centers $\mathcal{D}$.
- ***Strategies***: For data center d, the utility company choice $m(d) \in \mathcal{U}_d$ and scheduling decision $\boldsymbol{a}_d$ satisfy constraints (2)−(5). We denote the strategy of data center d by the tuple $\boldsymbol{s}_d = (\boldsymbol{a}_d, m(d))$. Let $\mathcal{S}_d$ denote the feasible strategy space for data center d defined by (2)−(5) and constraint $m(d) \in \mathcal{U}_d$. Let $\boldsymbol{s} = (\boldsymbol{s}_d,\ d \in \mathcal{D})$ denote the joint strategy profile of data centers. Let $\boldsymbol{s}_{-d}$ denote the strategy profile of all data centers except data center d.
- ***Costs***: Data center d incurs a cost $c_d(\boldsymbol{s}_d, \boldsymbol{s}_{-d})$ as in (6), which is a function of strategy profile $\boldsymbol{s}_d$ of data center d and the strategy $\boldsymbol{s}_{-d}$ of other data centers.

Notice that the cost of a data center d depends on the demand schedules of other data centers that are matched to the same utility company as d. Hence, our game is a *matching game with externalities* [10], [11]. The outcome of the game is a matching m and the joint scheduling decision profile $\boldsymbol{a}$ of the data centers. The outcome is stable when no data center will incur a lower cost from changing either its matched utility company or its action profile unilaterally [11].

Definition 2: A stable outcome of the matching game is the feasible strategy profile $\boldsymbol{s}^\star = (\boldsymbol{s}^\star_d,\ d \in \mathcal{D})$ such that for $d \in \mathcal{D}$

$$c_d(\boldsymbol{s}^\star_d, \boldsymbol{s}^\star_{-d}) \leq c_d(\boldsymbol{s}, \boldsymbol{s}^\star_{-d}), \qquad \boldsymbol{s} \in \mathcal{S}_d. \tag{7}$$

A data center's best response strategy is the choice that minimizes its own cost, assuming that the strategies of other data centers are fixed. That is

$$\boldsymbol{s}^{\text{best}}_d(\boldsymbol{s}_{-d}) \in \underset{\boldsymbol{s}_d \in \mathcal{S}_d}{\arg\min}\ c_d(\boldsymbol{s}_d, \boldsymbol{s}_{-d}), \quad d \in \mathcal{D}. \tag{8}$$

A stable outcome is a fixed point of the best responses of all data centers. That is, $\boldsymbol{s}^{\text{best}}_d(\boldsymbol{s}^\star_{-d}) = \boldsymbol{s}^\star_d$ for all $d \in \mathcal{D}$.

Problem (8) for data center d involves choosing a utility company, and it is a nonconvex optimization problem with discrete variables. However, under the given matching m, the objective function (6) and constraints (2)−(5) can be expressed as posynomials. Hence, problem (8) is a *geometric program* [16], which can be transformed into a convex optimization problem with variables $\boldsymbol{a}_d$. There are two steps involved in solving problem (8) for data center d under a given strategy profile $\boldsymbol{s}_{-d}$: (a) solving a convex optimization problem for a fixed matching m, and (b) comparing the objective value for all utility company choices for data center d.

In general, a stable outcome may not exist in a matching game with externalities [10]. We prove the existence of a stable outcome for Game 1 by constructing an *exact potential function* [17]. Such a function is defined as follows:

Definition 3: A function $P(\boldsymbol{s})$ is an exact potential for Game 1, if for any feasible strategy profiles $\boldsymbol{s} = (\boldsymbol{s}_d, \boldsymbol{s}_{-d})$ and $\widetilde{\boldsymbol{s}} = (\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d})$, we have

$$c_d(\boldsymbol{s}_d, \boldsymbol{s}_{-d}) - c_d(\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d}) = P(\boldsymbol{s}_d, \boldsymbol{s}_{-d}) - P(\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d}). \tag{9}$$

A potential function $P(\boldsymbol{s})$ tracks the changes in the data center's cost when its strategy changes. In the following theorem, we characterize an exact potential function for Game 1. There is no generic method of constructing a potential function, and it requires exploring the structure of the problem.

Theorem 1 *Game 1 admits an exact potential function*

$$P(\boldsymbol{s}) = \sum_{u \in \mathcal{U}} \sum_{t \in \mathcal{T}} \Big(\sum_{d \in m(u)} \Big(\big(p(t) + \kappa_u(t) e^{\text{other}}_u(t) \big) e_d(t) + \kappa_u(t)\, e^2_d(t) \Big) + \kappa_u(t) \sum_{d < d' \in m(u)} e_d(t)\, e_{d'}(t) \Big). \tag{10}$$

The proof can be found in Appendix A. Under a given matching m, the potential function (10) is a convex function of $\boldsymbol{a}$. Let $\boldsymbol{a}_m$ denote the global minimum of $P(\boldsymbol{s})$ under a given matching m. Let $\mathcal{M}$ denote the set of tuples $(\boldsymbol{a}_m, m)$ for all matchings m. In Theorem 2, we show that the stable outcomes of the matching game are in set $\mathcal{M}$.

Theorem 2 *Game 1 has at least one stable outcome. All stable outcomes are in set $\mathcal{M}$.*

The proof can be found in Appendix B. One can use the existing algorithms based on the best response update to determine a stable outcome [18]. These algorithms, however, often suffer from a low convergence rate, as only one single data center updates its strategy per iteration. We propose Algorithm 1 that can be executed by the data centers and utility companies in a distributed and parallel fashion to converge to a stable outcome. Let i denote the iteration index. The EMS of the data centers are responsible for the computations and message exchange. Fig. 2 shows the schematic of matching update in iteration i for five data centers and three utility companies in Fig. 1 (b). Our algorithm involves the initiation phase and matching phase.

Algorithm 1 The Data Center Matching Game Algorithm.

1: Set $i := 1$ and $\xi := 10^{-3}$.
2: Randomly assign each data center $d \in \mathcal{D}$ to a utility company $m^1(d) \in \mathcal{U}_d$, and initialize action profile $\boldsymbol{a}^1_d$.
3: Send parameters $\kappa_u(t), t \in \mathcal{T}$, to the data centers in set $\mathcal{D}_u$.
4: **Repeat**
5: Each data center d sends $e^i_d(t), t \in \mathcal{T}$ to utility company $m^i(d)$.
6: Each utility company u updates retail prices $p^{\text{r},i}_u(e^i_u(t), m^i)$ for $t \in \mathcal{T}$ using (1) and sends to the data centers in set $\mathcal{D}_u$.
7: Each data center d chooses a utility company in set $\mathcal{U}_d$ by computing its best response strategy in (8).
8: Each data center d sends termination request to its current utility company if it is different from the chosen one.
9: Each utility company u accepts at most one termination request.
10: Each data center d sends connection request to its chosen utility company if its termination request has been accepted.
11: Each utility company u accepts at most one connection request randomly.
12: Each data center d with an accepted connection request updates $m^{i+1}(d)$ with the chosen utility. Otherwise, $m^{i+1}(d) := m^i(d)$.
13: Each data center d, that changes its utility company, updates its action profile with its best response, i.e., $\boldsymbol{a}^{i+1}_d := \boldsymbol{a}^{\text{best},i}_d$.
14: Each utility company u communicates the retail price for the updated matching m^{i+1} to the data centers in $\mathcal{D}_u$.
15: Each data center d, that does not change its utility company, updates $\boldsymbol{a}^{i+1}_d$ according to (11).
16: $i := i + 1$.
17: **Until** No data center wants to change its strategy, i.e., $m^i = m^{i-1}$ and $||\boldsymbol{a}^i - \boldsymbol{a}^{i-1}|| < \xi$.

• *Initiation phase:* Lines 1 to 3 describe the initialization for the data centers and utility companies.

• *Matching phase:* The loop from Lines 4 to 17 describes the matching phase. It includes the following parts:

a) Information exchange: Lines 5 and 6 describe the information exchange between the data centers and utility companies about the energy demands and retail prices. This step is shown in Fig. 2 (a).

b) Utility company choice: Lines 7 to 11 describe how data center d chooses a utility company and how utility company u responses to the requests of the data centers. For example, Fig. 2 (b) shows that data centers 2, 3, and 4 send termination request to utility company 2. Data center 5 sends termination request to utility company 3.

We allow a utility company to accept *at most one* termination request and *at most one* connection request in each iteration. A utility company accepts the requests at random, since it is indifferent between data centers. Fig. 2 (b) shows that utility company 2 accepts the termination request from data center 2. Utility company 3 accepts the termination request from data center 5. Fig. 2 (c) shows that data centers 2 and 5 send connection requests to utility companies 1 and 2, respectively. Fig. 2 (d) shows the updated matching structure.

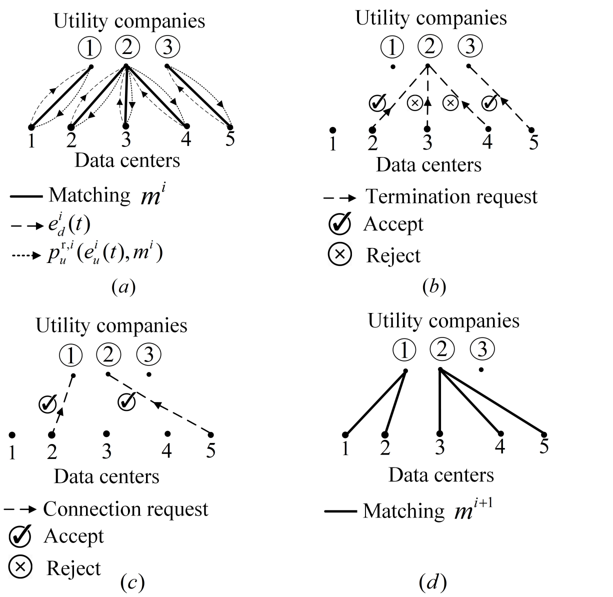

Fig. 2. Matching update procedure in Algorithm 1. (a) Information exchange among data centers and utility companies in matching m^i; (b) termination requests from data centers; (c) connection requests from data centers; (d) updated matching m^{i+1}.

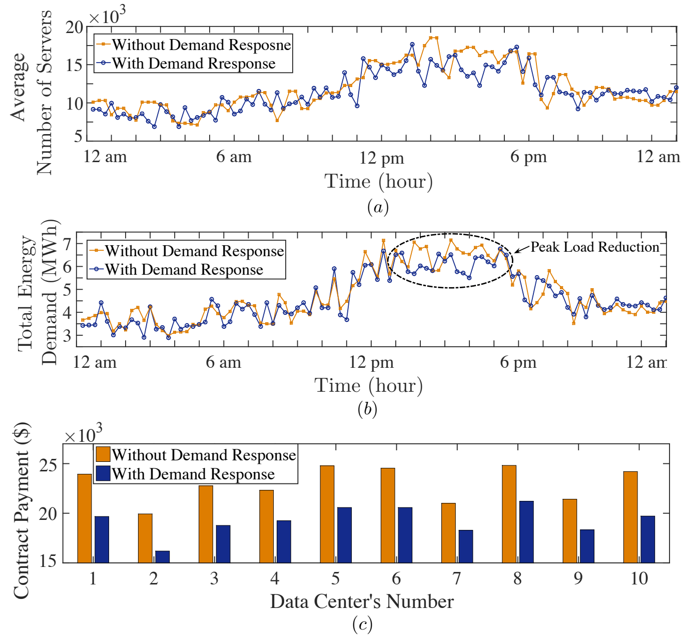

Fig. 3. (a) Total number of servers; (b) total energy demand of data center 1 over one day with and without demand response; (c) contract payment of data centers 1 to 10 with and without demand response.

c) Strategy update: Lines 12 to 15 describe how data center d updates its strategy $\boldsymbol{s}_d^i = (\boldsymbol{a}_d^i, m^i(d))$. If data center d changes its matching, then it updates its scheduling decision according to its best response. By receiving the updated price for the new matching, each data center d that has not changed its utility company will update its decision vector as follows.

$$\boldsymbol{a}_d^{i+1} = \left[\boldsymbol{a}_d^i - \gamma_d^i \nabla_{\boldsymbol{a}_d^i} c_d\left(\boldsymbol{a}_d^i, \boldsymbol{a}_{-d}^i, m^{i+1}\right)\right]_{\wp}, \qquad (11)$$

where $\gamma_d^i > 0$ is a diminishing step size with $\sum_{i=0}^{\infty} \gamma_d^i = \infty$ and $\sum_{i=0}^{\infty} (\gamma_d^i)^2 < \infty$, and $[\cdot]_{\wp}$ is the projection onto the feasible space defined by (2)−(5). In Algorithm 1, data centers use their best response strategies and update equation (11) for their utility company choice and workload scheduling. Each utility company accepts at most one termination request and at most one connection request in each iteration. We have

Theorem 3 *Algorithms 1 globally converges to a stable outcome of the data center matching game.*

The proof can be found in Appendix C.

IV. Performance Evaluation

In this section, we evaluate the performance of the stable outcome of the matching game. We set the contract period to be one day. We divide a day into $T = 96$ time slots, where each time slot is 15 minutes. We consider the electricity market with 10 utility companies serving 50 data centers, and each data center can choose a utility company from a random subset of seven utility companies. We use the wholesale market price on Oct. 10, 2016 of the Ontario's wholesale market [19]. The high price period is from 12 pm to 6 pm. Parameters $\kappa_u(t)$ for utility companies $u = 1, 2, \ldots, 10$ are set to $0.224, 0.208, \ldots, 0.08$ \$/(MWh)^2 for $t \in \mathcal{T}$, respectively.

To simulate the arrival rate of the workloads in a data center, we use the dataset from [20]. Each data center offers five service classes, and the workloads requesting service class $c = 1, \ldots, 5$ can be delayed by at most $\Delta_{c,d} = 0.05,\ 2,\ 12,\ 15,\ 20$ time slots, respectively. For service classes $c = 1, \ldots, 5$, we set $\sigma_{c,d}$ to 0.5, 10, 30, 50, 100 time slots, respectively. We consider $n_d^{\max} = 20{,}000$ homogeneous servers with power ratings $E_d^{\text{idle}} = 150\,\text{W}$ and $E_d^{\text{peak}} = 300\,\text{W}$ per time slot in each data center d. Parameters $\text{PUE}_d(t)$, $t \in \mathcal{T}$ are chosen at random from interval $[1.5, 2]$ for each data center. The step size in iteration i is set to be $\gamma_d^i = 1/(10 + 0.03 \times i)$.

We discuss how Algorithm 1 enables a data center to manage its energy demand. For the sake of comparison, we consider the scenario without demand response, where each data center randomly chooses a utility company and determines the number of servers based on constraints (2)−(4) without considering the price values. This is a nontrivial scenario, as a data center can delay the execution of a workload. Let us consider data center 1 as an example. Fig. 3 (a) shows that with demand response, the number of operating servers in data center 1 decreases during the peak hours, e.g., it is reduced from 17,000 to 14,000 around 4 pm. Fig. 3 (b) shows that the energy demand of data center 1 is reduced by 11.5% (from 7 MWh to 5.5 MWh during peak hours). Fig. 3 (c) shows that the contract payment of data centers is reduced by 15.4% on average as a result of server scheduling.

We discuss how Algorithm 1 affects the PAR of the aggregate demand and the revenue of the utility companies. We compare the PAR of the utility companies in the scenarios with and without data centers demand response. Fig. 4 (a) shows that, with data centers demand response, the PAR of the utility companies is reduced by 7.2% on average. A lower PAR improves the performance of the utility companies during peak hours. The revenue of utility companies depends on the matching structure. For the sake of comparison, we consider a

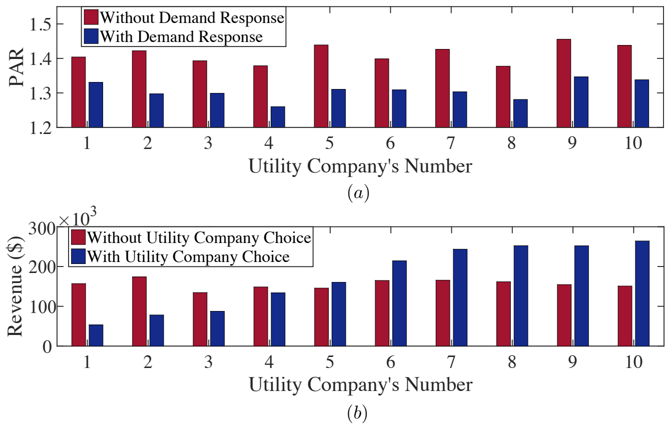

Fig. 4. (a) PAR in the energy demand; (b) revenue of the utility companies in the scenarios with and without utility company choice.

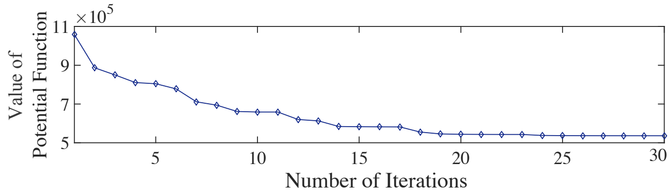

Fig. 5. The convergence of the potential function in Algorithm 1.

scenario where data centers randomly enter contracts with utility companies. In the stable outcome obtained from Algorithm 1, the number of data centers connected to utility companies 1 to 10 are 2, 2, 3, 4, 4, 4, 6, 6, 8, and 11, respectively. A utility company with a lower $\kappa_u(t)$ can attract more data centers as customers. When compared to the scenario without utility company choice, Fig. 4 (b) shows that the revenue of the utility companies with a higher $\kappa_u(t)$ decreases and the revenue of the utility companies with a lower $\kappa_u(t)$ increases (up to 82%). Thus, the results of Algorithm 1 is consistent with the utility companies' competition in deregulated markets.

Finally, we evaluate the convergence of Algorithm 1. Fig. 5 depicts the convergence of the potential function in one of our simulations with a random initial condition. The potential function decreases in each iteration and converges to a stable outcome in 30 iterations. The running time until convergence is 47 seconds. Regarding the computational complexity, in Line 7 of Algorithm 1, data center d solves $|\mathcal{U}_d|$ optimization problems to determine its best response strategy. Hence, the per-iteration complexity of Algorithm 1 for data center d is independent of the number of data centers and depends only on the number of utility companies in set $\mathcal{U}_d$, i.e., $O(|\mathcal{U}_d|)$.

V. Conclusion

In this paper, we studied the data centers' problem of choosing utility companies and scheduling workload in a deregulated electricity market. We modeled the interaction among data centers as a many-to-one matching game with externalities. We constructed an exact potential function, whose local minima correspond to the stable outcomes of the game. We developed an algorithm to determine a stable outcome. Simulation results showed that the data centers can decrease their cost by 15.4% with the proposed algorithm, as they can purchase electricity from their preferred utility companies and reduce their demands during peak hours. Meanwhile, the utility companies can achieve 7.2% reduction in the PAR. Those utility companies that offer lower tariffs can increase their revenue by up to 82%. For future work, we plan to extend the model by considering the competition among utility companies through price optimizations.

Acknowledgement

The work of S. Bahrami and V.W.S. Wong was supported by the Natural Sciences and Engineering Research Council of Canada. The work of J. Huang was supported by the Theme-based Research Scheme (Project No. T23-407/13-N) from the Research Grants Council of the Hong Kong Special Administrative Region, China, and a grant from the Vice-Chancellors One-off Discretionary Fund of The Chinese University of Hong Kong (Project No. VCF2014016).

Appendix

A. *Proof of Theorem 1*

To prove Theorem 1, we substitute (10) into the right-hand side of (9) and substitute (6) into the left-hand side of (9) for strategies $\boldsymbol{s}$ and $\widetilde{\boldsymbol{s}}$, and show that the results are the same. Data center d changes its utility company choice from u to $\widetilde{u}$, and its decision profile from $\boldsymbol{a}_d$ to $\widetilde{\boldsymbol{a}}_d$. Thus, the energy demand of data center d is changed from $e_d(t)$ to $\widetilde{e}_d(t)$ in time slot $t \in \mathcal{T}$. By substituting (10) into the right-hand side of (9) for $\boldsymbol{s} = (\boldsymbol{s}_d, \boldsymbol{s}_{-d})$ and $\widetilde{\boldsymbol{s}} = (\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d})$, we obtain

$$\begin{aligned}
&P(\boldsymbol{s}_d, \boldsymbol{s}_{-d}) - P(\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d}) = \\
&\textstyle\sum_{t\in\mathcal{T}} \Big(\big(p(t) + \kappa_u(t)\, e_u^{\text{other}}(t)\big)\, e_d(t) + \kappa_u(t)\, e_d^2(t) \\
&+ \kappa_u(t) \textstyle\sum_{d'\in m(u)\backslash d} e_d(t)\, e_{d'}(t) \\
&- \big(p(t) + \kappa_{\widetilde{u}}(t)\, e_{\widetilde{u}}^{\text{other}}(t)\big)\, \widetilde{e}_d(t) - \kappa_{\widetilde{u}}(t)\, \widetilde{e}_d^2(t) \\
&- \kappa_{\widetilde{u}}(t) \textstyle\sum_{d'\in \widetilde{m}(\widetilde{u})\backslash d} \widetilde{e}_d(t)\, e_{d'}(t) \Big).
\end{aligned} \tag{12}$$

In (12), the terms related to the data centers other than data center d cancel each other. Substituting (6) into the left-hand side of (9) for $\boldsymbol{s} = (\boldsymbol{s}_d, \boldsymbol{s}_{-d})$ and $\widetilde{\boldsymbol{s}} = (\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d})$, we have

$$\begin{aligned}
&c_d(\boldsymbol{s}_d, \boldsymbol{s}_{-d}) - c_d(\widetilde{\boldsymbol{s}}_d, \boldsymbol{s}_{-d}) = \\
&\textstyle\sum_{t\in\mathcal{T}} \big(e_d(t)\, p_u^{\text{r}}(e_u(t), m) - \widetilde{e}_d(t)\, p_{\widetilde{u}}^{\text{r}}(e_{\widetilde{u}}(t), \widetilde{m}) \big).
\end{aligned} \tag{13}$$

By substituting the retail price (1) into (13), the cost change for data center d will be equal to the potential function change in (12). This completes the proof. ■

B. *Proof of Theorem 2*

We first show that the global minimum of the potential function (10) is a stable outcome. Let $\boldsymbol{s}^\star = (\boldsymbol{a}_{m^\star}, m^\star)$ be the global minimum of $P(\boldsymbol{s})$. Thus, if data center d changes its action profile to $\boldsymbol{a}_d$ or its utility company to $m(d)$ unilaterally, then the value of the potential function increases. The change in the exact potential function is equal to the change in the cost of the deviating data center d. Hence, the cost of data center d increases as well. Consequently, no unilateral deviation from

$\boldsymbol{s}^{\star}$ can reduce the cost of any data center, and hence $\boldsymbol{s}^{\star}$ is a stable outcome of Game 1. As the exact potential function (10) has at least one global minimum, we know that the matching game has at least one stable outcome.

Next we show that an arbitrary stable outcome $(\boldsymbol{a}, m)$ is in set $\mathcal{M}$. We prove this by contradiction. Suppose that a stable outcome $(\boldsymbol{a}, m)$ is not in set $\mathcal{M}$. Hence, we have $\boldsymbol{a} \neq \boldsymbol{a}_m$. By definition, $\boldsymbol{a}_m$ is the global minimum of the potential function under matching m. We also know that $P(\boldsymbol{a}, m)$ is a convex function of $\boldsymbol{a}$. Thus, a unilateral change of $\boldsymbol{a}_d$ for any data center d in the opposite direction of the gradient $\nabla_{\boldsymbol{a}_d} P(\boldsymbol{a}, m)$ will reduce the potential function, and thus the cost of that data center. It contradicts the supposition that $(\boldsymbol{a}, m)$ is a stable outcome. Hence, $(\boldsymbol{a}, m)$ is in set $\mathcal{M}$. ■

C. *Proof of Theorem 3*

Since the potential function (10) is lower-bounded by zero, it is sufficient to show that the potential function decreases in each iteration of Algorithm 1. Line 17 of Algorithm 1 guarantees that if the algorithm converges, the result is a stable outcome. Next we provide the sketch of the proof.

Step a) Consider iteration i of Algorithm 1. We prove by induction that the potential function decreases when k data centers update their utility company choices simultaneously, where $k \geq 1$ is an arbitrary number. The base case (i.e., $k = 1$) corresponds to the unilateral change in the strategy of one data center. From (9), the potential function decreases, when the cost of a data center decreases. In the induction step, we consider $k = d$ and suppose that the potential function decreases when d data centers change their utility company choices. We prove that the potential function decreases when $k = d + 1$ data centers change their utility company choices. We divide the set of $d+1$ data centers into two sets with d data centers and one data center, respectively. We use the induction supposition for $k = d$ to show that when d data centers change their utility company choices, the potential function decreases.

Now, assume that data center $d+1$ decides to leave utility company $m^i(d+1)$ in iteration i to connect to utility company $m^{i+1}(d+1)$. In Algorithm 1, a utility company accepts at most one termination request from data centers per iteration. Thus, data center $d+1$ is the only one leaving utility company $m^i(d+1)$. Utility company $m^i(d+1)$ may accept a new connection request from other data centers, which increases its total demand. Thus, the payment of data center $d+1$ to utility company $m^i(d+1)$ will increase after updating the matching of other data centers. On the other hand, in Algorithm 1, a utility company accepts at most one connection request from data centers per iteration. Hence, data center $d+1$ is the only one that connects to utility company $m^{i+1}(d+1)$. Utility company $m^{i+1}(d+1)$ may accept a termination request from other data centers, which further decreases its total demand. Thus, the payment of data center d to utility company $m^{i+1}(d)$ will further decrease after updating the matching of other data centers. Consequently, the cost of data center $d+1$ decreases even other d data centers change their utility company choices. The exact potential function decreases when the cost of data center $d+1$ decreases. By the principle of induction, the potential function decreases, when multiple data centers change their utility company choices.

Step b) Under a given matching m^{i+1}, (9) implies that $\nabla_{\boldsymbol{a}_d^i} c_d\left(\boldsymbol{a}_d^i, \boldsymbol{a}_{-d}^i, m^{i+1}\right) = \nabla_{\boldsymbol{a}_d^i} P\left(\boldsymbol{a}_d^i, \boldsymbol{a}_{-d}^i, m^{i+1}\right)$. If data centers use (11) for the update, the potential function varies in the opposite direction of its gradient. Under a given matching, $P(\cdot)$ is a convex function of $\boldsymbol{a}^i$ and has a Lipschitz continuous derivative. Thus, for sufficiently small step sizes, the opposite gradient direction is a decreasing direction. ■

REFERENCES

[1] "2016 top markets report smart grid," Int.'l Trade Administration, U.S. Department of Commerce," Annual Report, Apr. 2016.

[2] Alberta Energy. [Online]. Available: http://www.energy.alberta.ca/Electricity/679.asp.

[3] Galvin Electricity Initiative. [Online]. Available: http://www.galvinpower.org/power-consumers/act/real-time-illinois.

[4] R. Basmadjian, J. F. Botero, G. Giuliani, X. Hesselbach, S. Klingert, and H. D. Meer, "Making data centres fit for demand response: Introducing GreenSDA and GreenSLA contracts," accepted for publication in *IEEE Trans. on Smart Grid*, 2017.

[5] Z. Liu, A. Wierman, Y. Chen, and B. Razon, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," in *Proc. of IEEE ACM Int.'l Conf. on Measurement and Modeling of Computer Systems*, New York, NY, Jun. 2013.

[6] Y. Guo and M. Pan, "Coordinated energy management for colocation data centers in smart grids," in *Proc. of IEEE SmartGridComm*, Miami, FL, Nov. 2015.

[7] T. Chen, Y. Zhang, X. Wang, and G. B. Giannakis, "Robust workload and energy management for sustainable data centers," *IEEE J. Sel. Areas in Commu.*, vol. 34, no. 3, pp. 651–664, Mar. 2016.

[8] H. Wang, J. Huang, X. Lin, and H. Mohsenian-Rad, "Proactive demand response for data centers: A win-win solution," *IEEE Trans. on Smart Grid*, vol. 7, no. 3, pp. 1584–1596, Dec. 2015.

[9] N. Tran, D. Tran, S. Ren, Z. Han, E. Huh, and C. Hong, "How geo-distributed data centers do demand response: A game-theoretic approach," *IEEE Trans. on Smart Grid*, vol. 7, no. 2, pp. 937–947, Mar. 2016.

[10] K. Bando, R. Kawasaki, and S. Muto, "Two-sided matching with externalities: A survey," *Journal of the Operations Research Society of Japan*, vol. 59, no. 1, pp. 35–71, Jan. 2016.

[11] A. E. Roth and M. Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1992.

[12] E. Hausman, R. Hornby, and A. Smith, "Bilateral contracting in deregulated electricity markets," The American Public Power Association, Synapse Energy Economics, Tech. Rep., Apr. 2008.

[13] N. Forouzandehmehr, M. Esmalifalak, H. Mohsenian-Rad, and Z. Han, "Autonomous demand response using stochastic differential games," *IEEE Trans. on Smart Grid*, vol. 6, no. 1, pp. 291–300, Jan. 2015.

[14] K. L. Rider, "A simple approximation to the average queue size in the time-dependent M/M/1 queue," *Journal of the ACM*, vol. 23, no. 2, pp. 631–367, 1976.

[15] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United States data center energy usage report," Ernest Orlando Lawrence Berkeley National Laboratory, CA, Tech. Rep. DE-AC02-05CH1131, 2016.

[16] S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A tutorial on geometric programming," *Optimization and Engineering*, vol. 8, no. 1, pp. 67–127, Apr. 2007.

[17] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 3124–143, May 1996.

[18] S. Durand and B. Gaujal, *Complexity and Optimality of the Best Response Algorithm in Random Potential Games*. Berlin, Heidelberg: Springer, 2016.

[19] Independent Electricity System Operator (IESO). [Online]. Available: http://www.ieso.ca.

[20] World Cup 98 web hits. [Online]. Available: http://ita.ee.lbl.gov/html/contrib/WorldCup.html.