

Data Center Demand Response in Deregulated Electricity Markets

Shahab Bahrami, *Student Member, IEEE*, Vincent W.S. Wong, *Fellow, IEEE*, and Jianwei Huang, *Fellow, IEEE*

Abstract—With the development of deregulated electricity markets, a customer can enter a contract with one of several competing utility companies. Meanwhile, a utility company is motivated to increase its market share by helping its customers manage their energy usage and save money through demand response programs. In this paper, we study the demand response program in deregulated electricity markets for data centers that often have significant flexibility in workload scheduling. We consider the real-time pricing (RTP) and model the data centers' coupled decisions of utility company choices and workload scheduling as a *many-to-one matching game with externalities*. To solve such a game, we show that it admits an *exact potential function*, whose local minima correspond to the stable outcomes of the game. We further develop a distributed algorithm that guarantees to converge to a stable outcome. Compared with the scenario without data centers' demand response, we show through simulation that the proposed algorithm can reduce the average contract payment of data centers by 18.7% and increase the revenue of the utility companies that offer lower electricity tariffs up to 80% by attracting more data centers as customers.

Keywords: deregulated electricity market, many-to-one matching game, stable outcome, distributed algorithm.

I. INTRODUCTION

The advances in small-scale power plants and the integration of communication technologies into the power grids have enabled *deregulated electricity markets* for many parts of the world [1], such as the states of Texas and Pennsylvania in the United States [2], Alberta in Canada [3], and Nordic countries in Europe [4]. In a deregulated 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 dynamically set their retail price to attract more customers and gain a larger market share [5]. This motivates the utility companies to deviate from today's common practice of flat-rate pricing and implement real-time pricing (RTP) [6]. When the customers respond to such a pricing scheme through demand response, the utility companies benefit from a smoother energy demand profile and sometimes a higher revenue. The customers, on the other hand, can take advantage of the lower prices.

Manuscript was received on Sept. 16, 2017, revised on Jan. 10, 2018, and accepted on Feb. 21, 2018. This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), 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-Chancellor's One-off Discretionary Fund of the Chinese University of Hong Kong (Project No. VCF2014016).

S. Bahrami and V.W.S. Wong are with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC, Canada, V6T 1Z4. J. Huang is with the Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong, email: {bahramis, vincentw}@ece.ubc.ca, jwhuang@ie.cuhk.edu.hk

In this paper, we focus on the data centers' utility company choices and price-based demand response. Data centers often monitor and control the demands of their information technology (IT) equipment (e.g., servers, routers). Many typical workloads (e.g., scientific computations, data analytics) in data centers are often delay-tolerant, hence may be rescheduled to off-peak hours [7], [8]. This motivates a recent rich body of literature on the design of the demand response algorithm for data centers, which can be divided into two main threads.

The first thread of literature is concerned with the data centers' workload management problem. Different techniques such as stochastic optimization [9], convex optimization [10], [11], and mixed-integer linear programming [12] have been used to tackle the workload management problem. These works often assumed fixed energy price of the utility company and addressed the cost minimizing problem from the data centers' point of view. The second thread of literature is concerned with the utility companies' pricing decisions for data centers. Wang *et al.* in [13] applied a two-stage optimization method to optimize the energy pricing rates for a utility company, in addition to the optimization of the energy demand profiles for data centers. Tran *et al.* in [14] studied the interactions among utility companies and data centers by a Stackelberg game model, where the utility companies obtain a closed-form solution to the cost minimization problems of data centers.

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 a utility company to sign the contract and schedule its workload to minimize its contract payment. We emphasize that signing contracts with multiple utility companies may not be practical for a data center in a deregulated electricity market. In particular, market deregulation promotes competition among the utility companies to offer better products, rates, and plans to retain and attract customers. Hence, in practice, a data center can choose the utility company with the best offer that fits its requirements. If a data center is not satisfied with its current utility company, it can switch to another utility company. A viable example for our proposed market mechanism is the deregulated electricity market in Houston in the United States, where there are 42 data centers [15] that can sign bilateral contracts with eight competing utility companies [16]. We consider the RTP from the utility companies, where the payments from the data centers depend on the amount and time of their electricity usage. With the RTP, the data centers' decisions are coupled together through their choices of utility companies. We capture the data centers' coupled decisions of utility company choices

and workload scheduling as a *many-to-one matching game*, where each utility company can supply electricity to multiple data centers, and a data center can match with one utility company. The underlying mechanism is a *matching with externalities* [17], due to the coupling decision of data centers. 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 challenging as there does not exist a general methodology for solving a matching with externalities. This paper is an extension of our previous work [18] by considering the energy demand uncertainties, as well as the preference of data centers and utility companies in their matching choice.

The contributions of this paper are as follows:

- **Data Center Workload Model:** We model the workloads' arrivals and executions in a data center by a time-dependent multiclass queuing system. Such a model provides a framework to compute the optimal number of active servers and the execution time of delay-tolerant services, subject to the constraints on the waiting time of the interactive real-time services over the contract period.
- **Risk-Aware Contract:** Data centers have uncertainty about their workload demand and local renewable generation. We introduce a risk measure called the conditional value-at-risk (CVaR), which enables the data centers to limit the risk of deviation in the energy demand from the contracted amount. To the best of our knowledge, this is the first work that incorporates risk management in the data center demand response problem.
- **Solution Method and Algorithm Design:** We characterize an exact potential function and show that the stable outcomes correspond to the local minima of the potential function. One can determine a local optimal solution to the potential function by solving a mixed-integer nonlinear optimization problem, which is NP hard. Instead, we develop an algorithm that can be executed by the data centers and utility companies in a distributed fashion and converges to a stable outcome of the game.
- **Performance Evaluation:** We perform simulations on a market with 50 data centers and 10 utility companies. When compared with the scenario without the demand response, our proposed algorithm reduces the cost of data centers and the peak-to-average ratio (PAR) of the aggregate demand of data centers connected to the same utility company by 18.7% and 8%, respectively. The proposed algorithm also enables the utility companies to attract more data centers as customers by setting lower energy tariffs, and as a result increase their revenue up to 80%. The computational complexity of the proposed algorithm is linear with the number of utility companies and independent of the number of data centers.

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 $\mathcal{D} = \{1, \dots, D\}$ and $\mathcal{U} = \{1, \dots, U\}$ 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 set $\mathcal{D}_u \subseteq \mathcal{D}$ of data centers. Sets $\mathcal{U}_d, d \in \mathcal{D}$, and $\mathcal{D}_u, u \in \mathcal{U}$, are determined based on the geographic locations of the utility companies and data centers as well as the topology of the power network. Fig. 1(a) shows a system with five data centers and three utility companies. Fig. 1(b) shows the corresponding bipartite graph representation. For example, utility company 1 can sell electricity to data centers in set $\mathcal{D}_1 = \{1, 2, 3\}$. Data center 4 can choose a utility company from set $\mathcal{U}_4 = \{2, 3\}$.

Each data center d possesses an energy management system (EMS) connected to the utility companies in set \mathcal{U}_d via a two-way communication network. It enables exchanging information such as the energy demands of the corresponding data center and the energy price for entering a bilateral contract.

A. Bilateral Contract and Contract Pricing Models

In deregulated markets, a data center can enter a bilateral contract with one utility company to purchase electricity. We can capture the contracts between data centers and utility companies as a many-to-one matching function [19].

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 (c) shows a feasible many-to-one matching, where a data center enters a contract with one utility company, while a utility company can enter contracts with multiple data centers. We assume that a data center can enter a short-term contract (e.g., one day) with a utility company [7], [20], [21]. Without loss of generality, we assume that the contract period of all data centers starts and finishes at the same time. We divide the intended contract period into a set $\mathcal{T} = \{1, \dots, T\}$ of T time slots with an equal length, e.g., 15 minutes per time slot.

In matching m , data center $d \in \mathcal{D}$ specifies its energy demand profile $e_d(t), t \in \mathcal{T}$, to be satisfied by its utility company choice. Each utility company purchases electricity from the wholesale market with a price $p(t), t \in \mathcal{T}$. The utility companies may offer RTP rates to the flexible large loads such as data centers. In an RTP scheme, the retail price of utility company u depends on both the volume and time of the energy consumption of all its customers. The (unit) retail price of utility company $u \in \mathcal{U}$, in time slot $t \in \mathcal{T}$, and matching m is an increasing convex function of the total energy demand $e_u(t) = e_u^{\text{other}}(t) + \sum_{d \in m(u)} e_d(t)$, where $e_u^{\text{other}}(t)$ denotes the demand in time slot t for the customers served by utility company u excluding the data centers. The retail price is greater than the wholesale price, in order to guarantee a positive profit for the utility company. We use the

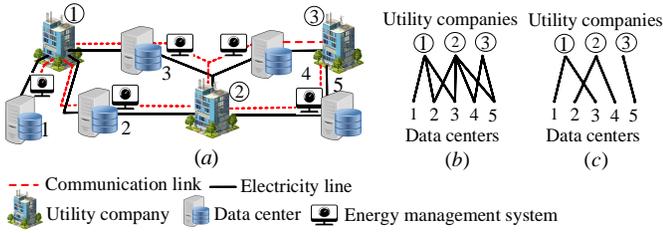


Fig. 1. (a) A system composed of five data centers equipped with EMS and three utility companies; (b) the corresponding bipartite graph representation; (c) a feasible many-to-one matching.

first-order Taylor approximation of the retail price near the wholesale price as follows:

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

where $\kappa_u(t)$, $u \in \mathcal{U}$, $t \in \mathcal{T}$, are nonnegative coefficients with the unit of $\$/(\text{MWh})^2$. The utility companies can determine $\kappa_u(t)$ according to the cost of supplying electricity to the data centers. The linear retail price model in (1) has been used in different problem settings such as the Cournot and Bertrand competitive markets [22] and retail power markets [23]. This linear price model can be extended to a piecewise linear model, which can be used to approximate nonlinear dynamic pricing schemes. The pricing scheme in (1) depends on the energy demand across all data centers. It motivates data center d towards scheduling its energy demand $e_d(t)$, $t \in \mathcal{T}$ to take advantage of the retail price fluctuations and lower its contract payment $\sum_{t \in \mathcal{T}} e_d(t) p_u^r(e_u(t), m)$ in matching m , where $m(d) = \{u\}$.

B. Data Center's Operation Model

The energy demand of a data center includes the demand for the workloads execution [9], [14]. A data center offers different service classes (e.g., video streaming, data analytics) to its customers. Consider data center $d \in \mathcal{D}$. Let $\mathcal{C}_d = \{1, \dots, C_d\}$ denote the set of service classes, where $C_d = |\mathcal{C}_d|$. We assume that both the workloads' inter-arrival time and execution time follow the exponential distribution [13], [14]. We can model the workloads' arrivals and executions by a time-dependent multiclass $M/M/1$ queuing system [24].

To manage the energy demand for executing the workloads in a data center, we consider the possibility of deferring the execution of an incoming workload to future time slots. To meet the quality-of-service requirements, the delay in executing an incoming workload needs to be controlled within a certain range, which depends on the type of service request. Let $\Delta_{c,d}$ denote the maximum number of time slots that the execution of a workload of service class $c \in \mathcal{C}_d$, $d \in \mathcal{D}$, can be delayed. If $\Delta_{c,d} = 0$, then the service cannot be delayed due to its interactive nature. Examples of such interactive services include web search, online gaming, and video streaming. For the delay-tolerant flexible services, such as scientific applications, data analytics, and file processing, we have $\Delta_{c,d} \geq 1$ [25].

We propose a probabilistic workload scheduling framework. The EMS in data center d may defer the execution of a workload of service class c from time slot $t_1 \in \mathcal{T}$ to

$t_2 \in [t_1, t_1 + \Delta_{c,d}]$ with a "time-shift" probability $p_{c,d}(t_1, t_2)$. For data center d , let $\mathbf{p}_{c,d}(t_1) = (p_{c,d}(t_1, t_2), t_2 \in \mathcal{T})$ denote the time-shift probability mass function for the arriving workloads of service class c in time slot t_1 . The EMS of data center d decides on the probability mass functions $\mathbf{p}_{c,d} = (\mathbf{p}_{c,d}(t), t \in \mathcal{T})$ for service classes $c \in \mathcal{C}_d$. We have

$$p_{c,d}(t_1, t_2) = 0, \quad \text{if } t_2 \notin [t_1, t_1 + \Delta_{c,d}], \quad (2a)$$

$$p_{c,d}(t_1, t_2) \in [0, 1], \quad \text{if } t_2 \in [t_1, t_1 + \Delta_{c,d}], \quad (2b)$$

$$\sum_{t_2 \in \mathcal{T}} p_{c,d}(t_1, t_2) = 1, \quad t_1 \in \mathcal{T}. \quad (2c)$$

Let $\lambda_{c,d}(t)$ denote the average arrival rate of workloads of service class $c \in \mathcal{C}_d$ in data center d and time slot t . By workload scheduling, the workloads in time slot t include the newly arriving workloads, which will not be deferred to future time slots. The newly arriving workloads of service class c have an average arrival rate $\bar{\lambda}_{c,d}(t) = p_{c,d}(t, t) \lambda_{c,d}(t)$.

The workloads in time slot t also include the workloads that are initially in the system at the beginning of time slot t . These workloads include the deferred workloads from the previous time slots, and the workloads with incomplete job from time slot $t-1$. The average number of workloads of service class c in data center d that are deferred from the previous time slots to time slot t is

$$I_{c,d}(\mathbf{p}_{c,d}, t) = \sum_{t'=1}^{t-1} p_{c,d}(t', t) \lambda_{c,d}(t'). \quad (3)$$

We now compute the average number of workloads from time slot $t-1$ with incomplete job. Let $n_d(t)$ denote the number of operating (active) servers of data center d in time slot t . We consider homogeneous servers in data center d and use $c_{c,d}$ to denote the average time that it takes for a server in data center d to execute a workload of service class c . In the case with heterogeneous servers, we can use the average processing time over different server types. If there exists only one service class c in data center d , then all the operating servers execute the workloads of service class c in time slot t with the average execution rate $\mu_{c,d}(t) = n_d(t)/c_{c,d}$. Nevertheless, the operating servers execute multiple service classes simultaneously. In Appendix A, we show that the process of the workloads of service class c can be modeled by an equivalent single-class $M/M/1$ queuing system with the workloads' average arrival rate $\bar{\lambda}_{c,d}(t)$ and execution rate

$$\bar{\mu}_{c,d}(t) = \mu_{c,d}(t) \left(1 - \left(\rho_d(t) - \frac{\bar{\lambda}_{c,d}(t)}{\mu_{c,d}(t)} \right) \right), \quad (4)$$

where $\rho_d(t) = \sum_{c \in \mathcal{C}_d} \bar{\lambda}_{c,d}(t) / \mu_{c,d}(t)$ is the average server utilization in time slot t in data center d . We use the steady state approximation to compute the average number of backlog workloads at the end of time slot $t-1$ as follows [24]:

$$J_{c,d}(\mathbf{p}_{c,d}, t) = \frac{\bar{\lambda}_{c,d}(t-1)}{\bar{\mu}_{c,d}(t-1) - \bar{\lambda}_{c,d}(t-1)}. \quad (5)$$

By characterizing the workloads, we determine the bounds on the number of operating servers in data center d . The scheduled workloads in each time slot should be executed in a

short period of time. Let $\delta_{c,d}$ denote the maximum execution time (i.e., the waiting time in the queue plus the servers service time) for the workloads of service class c in data center d . Due to the large number of operating servers in a data center, the value of $\delta_{c,d}$ is generally much smaller than the length of one time slot, e.g., $\delta_{c,d}$ often corresponds to a few seconds. Our approach, however, is applicable to any value of $\delta_{c,d}$.

In Appendix A, we show that a workload of service class c in data center d experienced the maximum expected execution time either at the beginning or at the end of time slot t [26]. We use the steady state approximation of the expected execution time of an incoming workload of service class c at the end of time slot t . Thus, we have $1/(\bar{\mu}_{c,d}(t) - \bar{\lambda}_{c,d}(t)) \leq \delta_{c,d}$. By performing some algebraic manipulations, we obtain

$$\frac{\varsigma_{c,d}}{\delta_{c,d}} + \sum_{c' \in \mathcal{C}_d} \varsigma_{c',d} \bar{\lambda}_{c',d}(t) \leq n_d(t), \quad c \in \mathcal{C}_d, d \in \mathcal{D}, t \in \mathcal{T}. \quad (6)$$

Moreover, inequality $1/(\bar{\mu}_{c,d}(t-1) - \bar{\lambda}_{c,d}(t-1)) \leq \delta_{c,d}$ implies that $J_{c,d}(\mathbf{p}_{c,d}, t)$ in (5) is upper bounded by $\delta_{c,d} \bar{\lambda}_{c,d}(t-1)$. Hence, we can approximate the number of initial workloads of service class c at the beginning of time slot t as $I_{c,d}(\mathbf{p}_{c,d}, t) + \delta_{c,d} \bar{\lambda}_{c,d}(t-1)$. The expected execution time at the beginning of time slot t for a workload of service class c is $(1 + I_{c,d}(\mathbf{p}_{c,d}, t) + \delta_{c,d} \bar{\lambda}_{c,d}(t-1)) / \bar{\mu}_{c,d}(t)$. It should be less than or equal to $\delta_{c,d}$. By performing some algebraic manipulations, we obtain

$$\begin{aligned} & \frac{\varsigma_{c,d}}{\delta_{c,d}} + \sum_{c' \in \mathcal{C}_d} \varsigma_{c',d} \bar{\lambda}_{c',d}(t) + \varsigma_{c,d} (\bar{\lambda}_{c,d}(t-1) - \bar{\lambda}_{c,d}(t)) \\ & + \frac{\varsigma_{c,d} I_{c,d}(\mathbf{p}_{c,d}, t)}{\delta_{c,d}} \leq n_d(t), \quad c \in \mathcal{C}_d, d \in \mathcal{D}, t \in \mathcal{T}. \quad (7) \end{aligned}$$

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

$$n_d(t) \leq n_d^{\max}, \quad d \in \mathcal{D}, t \in \mathcal{T}. \quad (8)$$

Next, we approximate energy demand of data center d for workloads execution by the energy consumption of the operating servers, which have an average utilization $\rho_d(t)$, over the steady state period in time slot t . Let E_d^{idle} and E_d^{peak} denote the average idle energy consumption and the peak energy consumption per time slot of a single 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

$$e_d^w(t) = \eta_d(t) n_d(t) (E_d^{\text{idle}} + (E_d^{\text{peak}} - E_d^{\text{idle}}) \rho_d(t)), \quad (9)$$

where $\eta_d(t) > 1$ is the power usage effectiveness of data center d in time slot t . The typical value of $\eta_d(t)$ for most data centers is between 1.5 and 2 [27].

We assume that a data center possesses a small-scale renewable generator (e.g., photovoltaic (PV) panel) to *partially* supply its demand. We also assume that a data center possesses an energy storage system. The data center can charge and discharge the energy storage system to smooth out the fluctuations in the energy demand and renewable generation. Data center d schedules the energy storage's charging and discharging profile $e_d^b = (e_d^b(t), t \in \mathcal{T})$, where $e_d^b(t)$ denotes the amount of energy being charged ($e_d^b(t) > 0$) to or

discharged ($e_d^b(t) < 0$) from the battery energy storage in time slot t . The charging/discharging rate of the energy storage in data center d has limits $e_d^{\text{b,min}} < 0$ and $e_d^{\text{b,max}} > 0$. That is,

$$e_d^{\text{b,min}} \leq e_d^b(t) \leq e_d^{\text{b,max}}, \quad d \in \mathcal{D}, t \in \mathcal{T}. \quad (10)$$

Let $E_d^{\text{b,init}}$ denote the initial energy level of the energy storage in data center d . The stored energy in the storage of data center d until time $T' \leq T$ is nonnegative and upper bounded by the limit $E_d^{\text{b,max}}$. Thus, we have

$$0 \leq E_d^{\text{b,init}} + \sum_{t=1}^{T'} e_d^b(t) \leq E_d^{\text{b,max}}, \quad d \in \mathcal{D}, T' \leq T. \quad (11)$$

Finally, the total energy demand of data center d in time slot t is obtained by

$$e_d(t) = e_d^w(t) + e_d^b(t) - e_d^r(t), \quad d \in \mathcal{D}, t \in \mathcal{T}. \quad (12)$$

C. Risk-Aware Energy Demand Scheduling

The accurate prediction of renewable generation is a challenge. The predicted renewable generation $e_d^r(t)$ often does not exactly match with the actual generation level $\hat{e}_d^r(t)$ in time slot $t \in \mathcal{T}$. The predicted arrival rate $\lambda_{c,d}(t)$ of the workloads of service class c in data center d and time slot t often does not exactly match with the actual arrival rate $\hat{\lambda}_{c,d}(t)$ either.

Let vectors $\hat{e}_d^r = (\hat{e}_d^r(t), t \in \mathcal{T})$ and $e_d^r = (e_d^r(t), t \in \mathcal{T})$ denote the actual and predicted generation profiles of the renewable generator in data center d , respectively. Let vectors $\hat{\lambda}_{c,d} = (\hat{\lambda}_{c,d}(t), t \in \mathcal{T})$ and $\lambda_{c,d} = (\lambda_{c,d}(t), t \in \mathcal{T})$ denote the profiles of actual and predicted arrival rate of the workloads of service class c in data center d , respectively. We define vectors $\hat{\lambda}_d = (\hat{\lambda}_{c,d}, c \in \mathcal{C}_d)$ and $\lambda_d = (\lambda_{c,d}, c \in \mathcal{C}_d)$ for data center d . The uncertainty in the renewable generation and workloads arrival causes the actual energy demand $\hat{e}_d(t)$ to deviate from the predicted energy demand $e_d(t)$. We use (9) and (12) to express $\hat{e}_d(t) - e_d(t)$ as

$$\begin{aligned} \hat{e}_d(t) - e_d(t) &= \sum_{c \in \mathcal{C}_d} \left(\varphi_{c,d}(t) (\hat{\lambda}_{c,d}(t) - \lambda_{c,d}(t)) \right. \\ & \quad \left. + \sum_{t'=1}^{t-1} \psi_{c,d}(t', t) (\hat{\lambda}_{c,d}(t') - \lambda_{c,d}(t')) \right) \\ & \quad - (\hat{e}_d^r(t) - e_d^r(t)), \quad (13) \end{aligned}$$

where

$$\varphi_{c,d}(t) = \varsigma_{c,d} \eta_d(t) p_{c,d}(t, t) (E_d^{\text{peak}} - E_d^{\text{idle}}), \quad t \in \mathcal{T},$$

and

$$\psi_{c,d}(t', t) = \varsigma_{c,d} \eta_d(t) p_{c,d}(t', t) E_d^{\text{peak}}, \quad t', t \in \mathcal{T}.$$

The excess energy demand of data center d in time slot t is equal to $[\hat{e}_d(t) - e_d(t)]^+$, where $[\cdot]^+ = \max\{\cdot, 0\}$. Utility company u can set penalties $p_u^+(t)$, $t \in \mathcal{T}$ with the unit of \$/MWh to prevent the data centers from under-estimating their demand. We define the *cost of risk* associated with the excess energy demand of data center d in matching m as

$$R_d(\hat{\lambda}_d, \lambda_d, \hat{e}_d^r, e_d^r) = \sum_{t \in \mathcal{T}} p_u^+(t) [\hat{e}_d(t) - e_d(t)]^+, \quad (14)$$

where $m(d) = \{u\}$. Utility company u can determine the penalties $p_u^+(t)$, $t \in \mathcal{T}$, by taking into account the real-time prices in the spot electricity market. In this paper, we consider the scenario that each data center purchases electricity from a utility company to meet its excess energy demand. In a scenario that data center d purchases electricity directly from the spot market to meet its excess energy demand, the cost model in (14) (without subscript u in $p_u^+(t)$, $t \in \mathcal{T}$) can be interpreted as the payment of data center d to the spot market.

We consider CVaR to determine appropriate vectors λ_d and e_d^r [28]. Optimizing the CVaR enables a data center to use the historical data record about its workloads and renewable energy generation to limit the risk of high penalty for the excess energy demand within a confidence level. The CVaR for data center d is defined for a confidence level $\beta_d \in (0, 1)$, and vectors λ_d and e_d^r as

$$\text{CVaR}_{d,\beta_d}(\lambda_d, e_d^r) = \mathbb{E}[R_d(\hat{\lambda}_d, \lambda_d, \hat{e}_d^r, e_d^r) \mid R_d(\hat{\lambda}_d, \lambda_d, \hat{e}_d^r, e_d^r) \geq \alpha_{\beta_d}], \quad (15)$$

where $\mathbb{E}[\cdot]$ is the expectation over the random variables $\hat{\lambda}_d$ and \hat{e}_d^r , and we have $\alpha_{\beta_d} = \min \{\alpha_d \mid \Pr(R_d(\cdot) \leq \alpha_d) \geq \beta_d\}$. Minimizing the CVaR in (15) results in the appropriate vectors λ_d and e_d^r that minimize the expected value of the penalty $R_d(\cdot)$ when it is higher than α_{β_d} .

In general, the explicit characterization of the probability distributions of the random variables $\hat{\lambda}_d$ and \hat{e}_d^r are not available. However, it is possible to estimate the CVaR in (15) by adopting the sample average approximation (SAA) technique [29]. We use the set $\mathcal{J} \triangleq \{1, \dots, J\}$ of J samples of random variables $\hat{\lambda}_d^j$ and $\hat{e}_d^{r,j}$ to estimate $\Pr(R_d(\lambda_d^j, e_d^{r,j}))$, the probability of the scenario with j th sample. The CVaR function in (15) can be approximated by [29]

$$\text{CVaR}_{d,\beta_d}(\lambda_d, e_d^r) \approx \min_{\alpha_d \in \mathbb{R}} \Gamma_{d,\beta_d}(\alpha_d, \lambda_d, e_d^r), \quad (16)$$

where

$$\Gamma_{d,\beta_d}(\alpha_d, \lambda_d, e_d^r) = \alpha_d + \sum_{j \in \mathcal{J}} \frac{\Pr(\lambda_d^j, e_d^{r,j})}{1 - \beta_d} [R_d(\lambda_d^j, \lambda_d, e_d^{r,j}, e_d^r) - \alpha_d]^+. \quad (17)$$

D. Preference of the Data Center and Utility Company

The total cost of data center d in matching m includes the bill payment and the CVaR function $\Gamma_{d,\beta_d}(\cdot)$. Let $\mathbf{a}_d = ((\mathbf{p}_{c,d}, c \in \mathcal{C}_d), \mathbf{e}_d^b, \boldsymbol{\mu}_d, \lambda_d, e_d^r, \alpha_d)$ denote the scheduling decision vector of data center d . The contract payment of data center d depends on its utility company choice $m(d) = \{u\}$ in matching m and the joint decision vector $\mathbf{a} = (\mathbf{a}_d, d \in \mathcal{D})$ of all data centers through the pricing scheme in (1). We have

$$c_d(\mathbf{a}, m) = (1 - \omega_d^{\text{cvar}}) \sum_{t \in \mathcal{T}} e_d(t) p_u^r(e_u(t), m) + \omega_d^{\text{cvar}} \Gamma_{d,\beta_d}(\alpha_d, \lambda_d, e_d^r), \quad (18)$$

where ω_d^{cvar} is a weight coefficient in the interval $[0, 1]$. For data center $d \in \mathcal{D}$, we define preference relation \succeq_d over the

pairs (\mathbf{a}, m) and (\mathbf{a}', m') as

$$(\mathbf{a}, m) \succeq_d (\mathbf{a}', m') \iff c_d(\mathbf{a}, m) \leq c_d(\mathbf{a}', m'). \quad (19)$$

Utility company u prefers a contract with a higher revenue. It also prefers a lower PAR of the aggregate energy demand to improve the performance of the energy network during peak hours. Reducing the PAR helps the utility company lower its retail price, hence can attract more customers. We consider the following objective function for utility company $u \in \mathcal{U}$:

$$f_u(\mathbf{a}, m) = (1 - \omega_u) f_u^{\text{rev}}(\mathbf{a}, m) - \omega_u f_u^{\text{PAR}}(\mathbf{a}, m), \quad (20)$$

where ω_u is a weight coefficient in the interval $[0, 1]$. The revenue of utility company $u \in \mathcal{U}$ in matching m is

$$f_u^{\text{rev}}(\mathbf{a}, m) = \sum_{d \in m(u)} \sum_{t \in \mathcal{T}} e_d(t) p_u^r(e_u(t), m),$$

and the PAR of the energy demand of the data centers in set $m(u)$ is

$$f_u^{\text{PAR}}(\mathbf{a}, m) = \frac{\max_{t \in \mathcal{T}} \{e_u^{\text{other}}(t) + \sum_{d \in m(u)} e_d(t)\}}{\frac{1}{T} \sum_{t \in \mathcal{T}} (e_u^{\text{other}}(t) + \sum_{d \in m(u)} e_d(t))}.$$

The total energy supply from utility company u is upper-bounded by e_u^{max} , due to the limited energy budget and transmission capacity in the network. We define the preference relation \succeq_u for utility company $u \in \mathcal{U}$ over the pairs (\mathbf{a}, m) and (\mathbf{a}', m') as

$$(\mathbf{a}, m) \succeq_u (\mathbf{a}', m') \iff \begin{cases} f_u(\mathbf{a}, m) \geq f_u(\mathbf{a}', m'), \\ e_u(t), e'_u(t) \leq e_u^{\text{max}}, t \in \mathcal{T}. \end{cases} \quad (21)$$

III. PROBLEM FORMULATION AND ALGORITHM DESIGN

Data center d aims to minimize its cost in (18) and achieve the most preferred scheduling decision and utility company choice based on the preference relation in (19). The decision making of data centers are interdependent, since the cost of a data center in (18) is a function of the joint decision vector \mathbf{a} of all data centers as well as the matching structure m . We capture the interactions among the data centers as a *many-to-one matching game* [19], which is defined as follows:

Game 1 (Data Center Many-to-One Matching Game):

- *Players:* The set of all data centers \mathcal{D} .
- *Strategies:* The strategy of data center d is the tuple $\mathbf{s}_d = (\mathbf{a}_d, m(d))$. Let \mathcal{S}_d denote the feasible strategy space for data center d defined by (2)–(13), (17), and constraint $m(d) \subseteq \mathcal{U}_d$. Let $\mathbf{s} = (\mathbf{s}_d, d \in \mathcal{D})$ denote the strategy profile of all data centers. Let \mathbf{s}_{-d} denote the strategy profile of all data centers except data center d .
- *Costs:* Data center d aims to minimize the cost $c_d(\mathbf{s}_d, \mathbf{s}_{-d})$ as in (18), which is a function of strategy profiles of data center d and other data centers.

We emphasize that in the above-mentioned matching game, the utility companies are not players and the function forms of their prices are fixed based on (1), i.e., parameters $\kappa_u(t)$, $t \in \mathcal{T}$ are fixed for each utility company u . That is, Game 1 considers the perspective of the data centers. However, utility company u signs a contract with the most preferred data centers based on the preference relation in (21).

Here, the cost of data center d depends on the demand schedules of other data centers matched to the same utility company as d . Hence, our game is a *matching game with externalities* among the data centers [17], [19]. For our problem setting, the outcome of the game is a matching m and the joint scheduling decision profile \mathbf{a} of the data centers. The outcome is stable if there exists no data center that incurs a lower cost from changing either its matched utility company or its action profile unilaterally [17], [19].

Definition 2 (Stable Outcome): A stable outcome of the matching game is the feasible strategy profile $\mathbf{s}^* = (s_d^*, d \in \mathcal{D})$ such that for $d \in \mathcal{D}$

$$c_d(s_d^*, s_{-d}^*) \leq c_d(s, s_{-d}^*), \quad s \in \mathcal{S}_d. \quad (22)$$

In general, a stable outcome may not exist in a matching game with externalities [17]. Furthermore, there does not exist a general algorithm that guarantees to compute a stable outcome in a matching game with externalities. To study the stable outcomes in Game 1, we consider the concept of best response strategy of data center d , which is defined as

$$s_d^{\text{best}}(s_{-d}) \in \arg \min_{s_d \in \mathcal{S}_d} c_d(s_d, s_{-d}), \quad d \in \mathcal{D}. \quad (23)$$

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

Problem (23) for data center d involves choosing a utility company. Hence, it is a non-convex optimization problem with discrete variables. However, under the given strategy profile s_{-d} and matching m , problem (23) can be transformed into a convex optimization problem with quadratic objective function and linear constraints with variables \mathbf{a}_d . There are two steps involved in solving problem (23) for data center d under a given strategy profile 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 a data center.

We prove the existence of a stable outcome for Game 1 by constructing an *exact potential function* [30]. Such a function is defined as follows:

Definition 3 (Exact Potential Function): A function $P(\mathbf{s})$ is an exact potential for Game 1, if for any feasible strategy profiles $\mathbf{s} = (s_d, s_{-d})$ and $\tilde{\mathbf{s}} = (\tilde{s}_d, s_{-d})$, we have

$$c_d(s_d, s_{-d}) - c_d(\tilde{s}_d, s_{-d}) = P(\mathbf{s}) - P(\tilde{\mathbf{s}}). \quad (24)$$

A potential function $P(\mathbf{s})$ 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.

Theorem 1: Game 1 admits an exact potential function

$$\begin{aligned} P(\mathbf{s}) = & \sum_{u \in \mathcal{U}} \sum_{t \in \mathcal{T}} \left(\sum_{d \in m(u)} (1 - \omega_d^{\text{cvar}}) (\kappa_u(t) (e_d^2(t) \right. \\ & \left. + e_u^{\text{other}}(t) e_d(t) + \sum_{d' < d' \in m(u)} e_d(t) e_{d'}(t)) + p(t) e_d(t) \right) \\ & + \sum_{d \in \mathcal{D}} \omega_d^{\text{cvar}} \Gamma_{d, \beta_d}(\alpha_d, \boldsymbol{\lambda}_d, \mathbf{e}_d^r). \end{aligned} \quad (25)$$

The proof can be found in Appendix B. Under a given matching m , the potential function (25) is a convex function of \mathbf{a} . Let \mathbf{a}_m denote the global minimum of $P(\mathbf{s})$ under a

given matching m . Let \mathcal{M} denote the set of tuples (\mathbf{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, which is in set \mathcal{M} .

The proof can be found in Appendix C. We now propose Algorithm 1 executed by the data centers and utility companies in a distributed fashion to converge to a stable outcome. The proposed algorithm is based on the gradient decent method and best response update of multiple data centers.

Let i denote the iteration index. In Algorithm 1, Lines 1 to 3 describe the initialization for the data centers and utility companies. Lines 5 and 6 describe the information exchange between the data centers and utility companies about the energy demands and retail prices. Lines 7 to 11 describe how a data center d chooses a utility company and how a utility company u responds to the requests of the data centers. To compute the best response strategy $s_d^{\text{best}, i}(s_{-d}^i)$, data center d solves problem (23) for different possible utility company choices $m(d) \in \mathcal{U}_d$, as it is a convex optimization problem under the given strategy profile s_{-d}^i and utility company choice $m(d)$. If its current utility company is different from the utility company choice in its best response strategy, data center d sends a termination request to its current utility company. A utility company u uses the preference relation \succeq_u in (21) to accept the most preferred termination request. Next, each data center d sends a connection request to the utility company in its best response strategy if its termination request has been accepted. A utility company u uses the preference relation \succeq_u in (21) to accept the most preferred connection request. Note that each utility company can accept *at most one* termination request and *at most one* connection request per iteration.

Lines 12 to 15 describe how a data center d updates its strategy s_d^i . If data center d changes its matching, then it updates its scheduling decision profile with its best response, i.e., $\mathbf{a}_d^{i+1} := \mathbf{a}_d^{\text{best}, i}$. By receiving the updated retail price for the new matching m^{i+1} , each data center d (that has not changed its utility company) computes its updated action profile \mathbf{a}_d^{i+1} based on the following gradient updating process:

$$\mathbf{a}_d^{i+1} = [\mathbf{a}_d^i - \gamma_d^i \nabla_{\mathbf{a}_d^i} c_d(\mathbf{a}_d^i, \mathbf{a}_{-d}^i, m^{i+1})]_{\mathcal{F}}, \quad (26)$$

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 $[\cdot]_{\mathcal{F}}$ is the projection onto the feasible space defined by (2)–(13), and (17). The effect of externalities among the decision making of the data centers in Game 1 can be observed in Lines 6, 7, 9, and 11 of Algorithm 1. In Line 6, utility company u updates its retail price, which depends on the scheduling decisions of all data centers connected to that utility company. In Line 7, data center d solves the optimization problem (23), the solution of which depends on the retail prices of the utility companies. In Lines 9 and 11, the preference relation (21) of utility company u depends on the joint scheduling decisions of all data centers connected to that utility company.

Remark 1: In Lines 9 and 11 of Algorithm 1, each utility company is allowed to accept at most one termination request and at most one connection request per iteration. As we show

Algorithm 1 The Data Center Matching Game Algorithm.

Initiation phase

- 1: Set $i := 1$ and $\xi := 10^{-3}$.
- 2: Randomly assign each data center $d \in \mathcal{D}$ to a utility company $m^1(d) \subseteq \mathcal{U}_d$, and initialize action profile \mathbf{a}_d^1 .
- 3: Send parameters $\kappa_u(t), t \in \mathcal{T}$ to the data centers in set \mathcal{D}_u .

Matching phase

4: **Repeat**

Information exchange

- 5: Each data center d sends $e_d^i(t), t \in \mathcal{T}$ to utility company $m^i(d)$.
- 6: Each utility company u updates retail prices $p_u^{r,i}(e_u^i(t), m^i)$ for $t \in \mathcal{T}$ using (1) and sends to the data centers in set \mathcal{D}_u .

Utility company choice

- 7: Each data center d chooses a utility company in set \mathcal{U}_d by computing its best response strategy in (23).
- 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 the most preferred termination request based on the preference relation \succeq_u in (21).
- 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 the most preferred connection request based on the preference relation \succeq_u in (21).

Strategy update

- 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., $\mathbf{a}_d^{i+1} := \mathbf{a}_d^{\text{best},i}$.
 - 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 \mathbf{a}_d^{i+1} according to (26).
 - 16: $i := i + 1$.
 - 17: **Until** No data center wants to change its strategy, i.e., $m^i = m^{i-1}$ and $\|\mathbf{a}^i - \mathbf{a}^{i-1}\| < \xi$.
-

in Theorem 3, this assumption enables us to prove the convergence of Algorithm 1 to a stable outcome. The intuition behind this assumption is as follows. In Line 9, a utility company accepts the termination requests one by one per iteration, and thereby reducing its retail price gradually in order to retain those data centers who has submitted termination requests. In Line 11, a utility company accepts the connection requests one by one per iteration to avoid a sudden price increase and losing many data centers as customers.

In the following theorem, we show that Algorithm 1 converges to a stable outcome.

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

The proof can be found in Appendix D. In the following, we evaluate the per-iteration complexity of Algorithm 1 for each data center and the average number of iterations for Algorithm 1 to converge to a stable outcome.

Remark 2: In Line 7 of Algorithm 1, each data center d solves $|\mathcal{U}_d|$ convex optimization problems to determine its best response strategy. Furthermore, the data centers update their utility company choice and workload schedule in parallel. Hence, the per-iteration complexity of Algorithm 1 for data center d is independent of the number of data centers in the system and depends only on the number of utility companies

in set \mathcal{U}_d , i.e., $O(|\mathcal{U}_d|)$.

Remark 3: Depending on the number of termination requests and connection requests in Lines 8 and 10 of Algorithm 1, multiple data centers (between 0 and $|\mathcal{U}|$) update their strategies using their best responses. Hence, Algorithm 1 can be interpreted as a best response algorithm with multiple updates per iteration. In [31, Theorem 2], it is shown that in a potential game, the average number of iterations to converge for a best response algorithm (with one update per iteration) is linear with the number of players. The strategy update of multiple data centers in one iteration of Algorithm 1 can be roughly interpreted as multiple consecutive iterations with one update. Hence, we can use [31, Theorem 2] to conclude that under the given number of utility companies, the average number of iterations for Algorithm 1 to converge to a stable outcome is linear with the number of data centers, i.e., $O(D)$.

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, which are free to choose a utility company from a random predetermined subset of seven utility companies. Fig. 2(a) shows the wholesale market price on Oct. 10, 2016 of the Ontario's wholesale market [32]. Each data center is equipped with a PV plant. We use the historical generation data for Ontario, Canada power grid database from June 1, 2016 to Oct. 1, 2016 [32] to obtain the samples for PV generation. Fig. 2(b) shows the average output power of a PV unit over 24 hours obtained from the historical generation data. For each data center, we scale the available historical data such that the average generation level of the PV plants becomes 0.5 MW per time slot. Each data center is equipped with an energy storage system. The maximum charging/discharging rate and capacity of an energy storage is set to 0.05 MW and 0.5 MWh, respectively. The initial energy level of the storage system in each data center is set to 50% of the maximum capacity.

To simulate the arrival rate of the workloads in a data center, we use the World Cup 98 web hits data [33], which consists of all the requests made to the 1998 World Cup website between Apr. 30, 1998 and Jul. 26, 1998 (i.e., 89 days). This dataset includes several service classes (e.g., HTML, JPEG) for different locations. For simulations, we randomly select 5 service classes for each data center and consider the average workload arrival rates of those service classes in the dataset. We also assume some additional specifications for the service classes. In data center d , the workloads of service classes $c = 1, \dots, 5$ can be delayed by at most $\Delta_{c,d} = 0, 4, 8, 16, 20$ time slots, respectively. We set $\delta_{c,d}$ at random from the interval of [0.1 sec, 3 sec] for service classes in each data center. The time of the requests are available in the World Cup 98 web hits dataset. We divide the collection period into 12 time intervals of seven consecutive days. Next, for each time interval, we determine the average number of workloads of service class c in time slot t . We obtain 12 samples of the

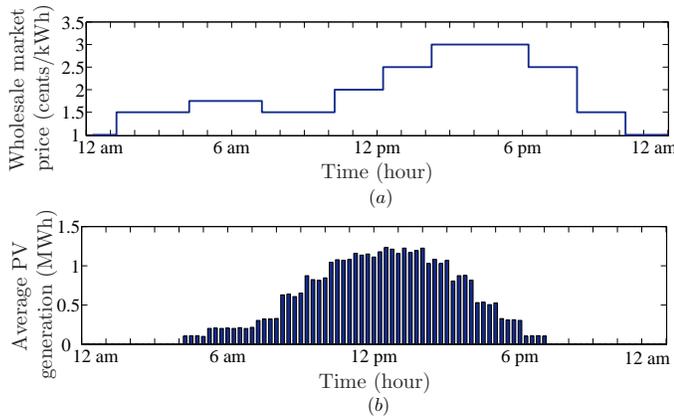


Fig. 2. (a) Wholesale day-ahead market prices over 24 hours; (b) Average PV generation over 24 hours.

average arrival rate for the workloads of service class c in each time slot. We consider $n_d^{\max} = 14,000$ homogeneous servers with power ratings $E_d^{\text{idle}} = 100 \text{ kW}$ and $E_d^{\text{peak}} = 200 \text{ kW}$ per time slot in each data center d . Parameters $s_{c,d}$, $c \in \mathcal{C}_d$ for the servers are chosen at random from interval $[0.1 \text{ sec}, 10 \text{ sec}]$. The coefficient ω_d^{cvar} is set to 0.5, and parameters $\eta_d(t)$, $t \in \mathcal{T}$ are set to 1.5 for all data centers. The confidence level β_d for data center d is chosen at random from interval $[0.75, 0.85]$.

Parameters $\kappa_u(t)$ for utility companies $u = 1, 2, \dots, 10$ are set to 0.224, 0.208, \dots , 0.08 $\$/(\text{MWh})^2$ for $t \in \mathcal{T}$, respectively. We set $p_u^+(t) = 20 \text{ \$/MWh}$, $t \in \mathcal{T}$, and $\omega_u = 0.8$ for all utility companies. We perform simulations using Matlab in a PC with processor Intel(R) Core(TM) i7-3770K CPU@3.5 GHz.

We first demonstrate how Algorithm 1 enables a data center to manage its energy demand. The step size in iteration i is set to $\gamma_d^i = 1/(50 + 0.07 \times i)$. For the sake of comparison, we consider the benchmark scenario, where each data center randomly chooses a utility company and does not schedule its workloads. Consider data center 1 as an example. Fig. 3(a) shows the predicted arrival rate of the delay-tolerant workloads in data center 1 requesting service class 5. With workload scheduling, the number of workloads during peak hours is reduced. Fig. 3(b) shows that with workload scheduling, the total number of operating servers in data center 1 decreases over the day, e.g., it is reduced from 14,000 to 12,000 around 5 pm. Moreover, Fig. 3(c) shows that the energy demand of data center 1 is reduced by 17.3% (from 11 MWh to 9.1 MWh on average) during the period from 12 pm to 6 pm, as a result of workload scheduling and reducing the number of servers.

We then study the changes in the cost of data centers when they implement Algorithm 1. Fig. 4 shows the daily cost of data centers 1 to 10 in three scenarios: (i) the benchmark scenario without demand response and utility company choice, (ii) the scenario without demand response and with utility company choice, and (iii) the scenario with both demand response and utility company choice. When comparing the second scenario with the first scenario, the total cost of the data centers decreases by 8.8% on average as a result of choosing their preferred utility company. When comparing the third scenario with the first scenario, the total cost of the data centers decreases by 18.7% as a result of both workload scheduling

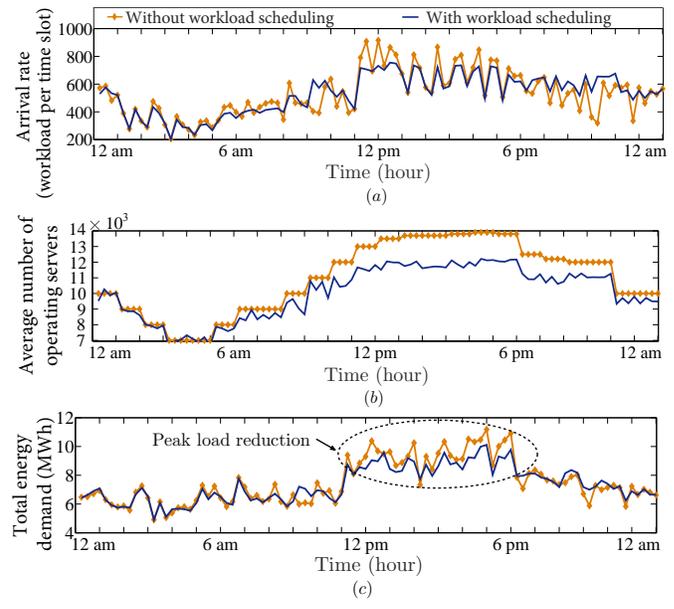


Fig. 3. (a) Arrival rates of the workloads of service class 5; (b) Total number of operating servers; (c) Total energy demand over 24 hours in data center 1 with and without workload scheduling.

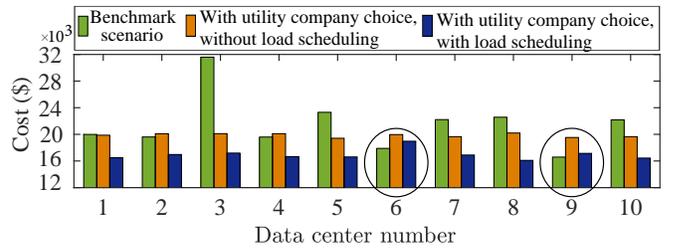


Fig. 4. Total daily cost of data center 1 by considering the opportunities of utility company choice or workload scheduling.

and choosing their preferred utility company. Note that in the second and third scenarios, there may exist some data centers (e.g., data centers 6 and 9) with a higher cost compared with the first scenario. This is due to the effect of externalities, i.e., the workload scheduling and utility company choices of other data centers on the price values. Nevertheless, in the stable outcome, no data center has an incentive to change its utility company choice and workload scheduling.

Next, we study the impact of the confidence level β_d on the energy demand prediction and risk of high excess energy demand for the data centers. It enables us to study the strategy of a data center in reducing the risk of penalty for excess energy demand in the contract. Fig. 5(a) shows the values of $\Gamma_{1,\beta_1}(\cdot)$ in (17) for data center 1 with β_1 in the range between 0.05 to 0.95. When β_1 increases, the data center becomes more risk averse and will try to decrease the risk of high excess energy demand by assigning a higher predicted workload energy demand and a lower predicted PV generation. Hence, the value of $\Gamma_{1,\beta_1}(\cdot)$ decreases when β_1 increases. Fig. 5(b) shows the contracted energy demand profiles of data center 1 with confidence levels $\beta_1 = 0.95$ and $\beta_1 = 0.8$. The contracted energy demand of the data center is larger in most time slots with the confidence level $\beta_1 = 0.95$. That is, the data center is risk-averse and prefers to sign a contract for a larger amount

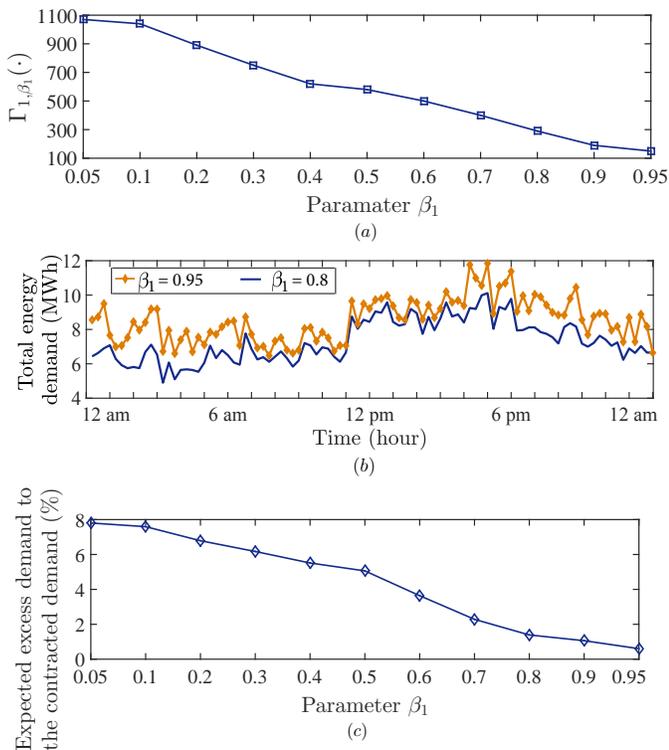


Fig. 5. (a) The value of $\Gamma_{d,\beta_d}(\alpha_d, \lambda_d, e_d^r)$ in (17) for data center 1 with β_1 within the range between 0.05 and 0.95; (b) Total energy demand profile of data center 1 with $\beta_1 = 0.95$ and $\beta_1 = 0.8$; (c) Total energy demand of data center 1 with β_1 within the range between 0.05 and 0.95.

of its energy demand, in order to reduce the risk of penalty payment for the excess energy demand. Fig. 5(c) shows the ratio of the *expected* excess energy demand to the contracted energy demand for data center 1 with the confidence level β_1 . When β_1 increases from 0.05 and 0.95, the expected excess energy demand (which should be supplied either by the utility company or directly from the spot market) decreases from 7.8% to 0.07% of the contracted energy demand. It also shows that in general this data center prefers to sign the bilateral contract with a utility company for most of its energy demand.

We now compare the PAR and revenue of the utility companies in three scenarios: (i) the benchmark scenario without data centers' demand response and utility company choice, and with the utility companies' randomly selection of a data center in responding to the connection or termination requests from data centers, (ii) the scenario with data centers' demand response and utility company choice, and with parameter $\omega_u = 1$ for the preference relation (21) for the utility companies, and (iii) the scenario with data centers' demand response and utility company choice, and with parameter $\omega_u = 1$ for the preference relation (21). Recall that the value of ω_u indicates the importance of reducing the PAR compared with increasing the revenue in responding to the connection or termination requests from data centers. Fig. 6(a) shows that, compared with the first scenario, the PAR is reduced by 4.6% and 8% on average in the second and third scenarios, respectively. The PAR is the lowest in the third scenario, as for $\omega_u = 1$, reducing the PAR is the most dominant factor for the utility companies in responding to the connection or termination requests.

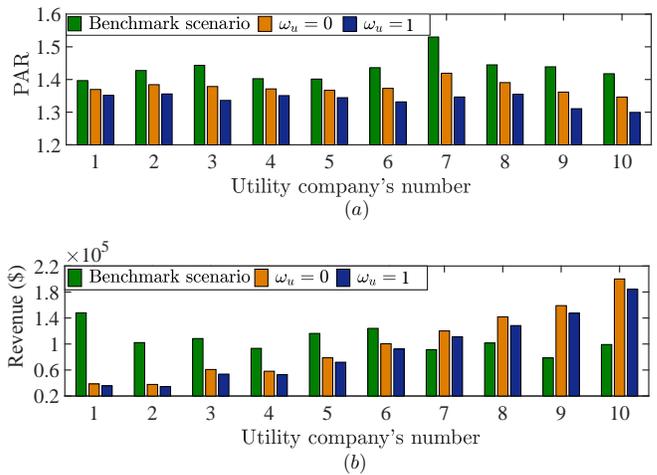


Fig. 6. (a) PAR in the generated power; (b) Revenue of the utility companies in the benchmark scenario, and the scenarios with $\omega_u = 0$ and $\omega_u = 1$.

Fig. 6(b) shows the revenue of the utility companies in the above-mentioned three scenarios. Utility company 1 has the highest parameter $\kappa_u(t)$ and utility company 10 has the lowest parameter $\kappa_u(t)$. Fig. 6(b) shows that the utility companies with a higher parameter $\kappa_u(t)$ have a higher average revenue in the benchmark scenario, since data centers randomly choose their matched utility company, and utility companies randomly accept the connection and termination requests from data centers. However, if data centers choose their matched utility company according to preference relation (19) and utility companies have the preference relation in (21) with $\omega_u = 0$, then the revenue of those utility companies with a higher $\kappa_u(t)$ decreases (up to 70%) and the revenue of those utility companies with a lower $\kappa_u(t)$ increases (up to 80%). In fact, in the stable outcome for the second scenario, the number of data centers connected to utility companies 1 to 10 are 2, 3, 3, 3, 4, 5, 5, 7, 8, and 10, respectively. It illustrates that parameter $\kappa_u(t)$ affects the market share, and hence the revenue for a utility company. When comparing the second scenario with the third scenario, the revenue of the utility companies is 9.8% higher on average, since increasing the revenue is the dominant criterion in the second scenario.

Finally, we evaluate the average number of iterations of Algorithm 1. Fig. 7(a) depicts the convergence of the potential function in one of our simulations for 50 data centers. The potential function decreases in each iteration and converges to a stable outcome in 40 iterations. Fig. 7(b) shows the average number of iterations versus the number of data centers for an error tolerance $\xi = 10^{-3}$. We perform simulations for 20 random initial conditions for the matching structure and data centers' energy demands. The number of utility companies is set to 10 in all scenarios. For each scenario, we increase the step size gradually and report the smallest number of iterations for convergence. The number of iterations increases in the number of data centers. However, our algorithm always converges in a reasonable number of iterations.

V. CONCLUSION

In this paper, we addressed the data centers' problem of choosing utility companies and scheduling workload in

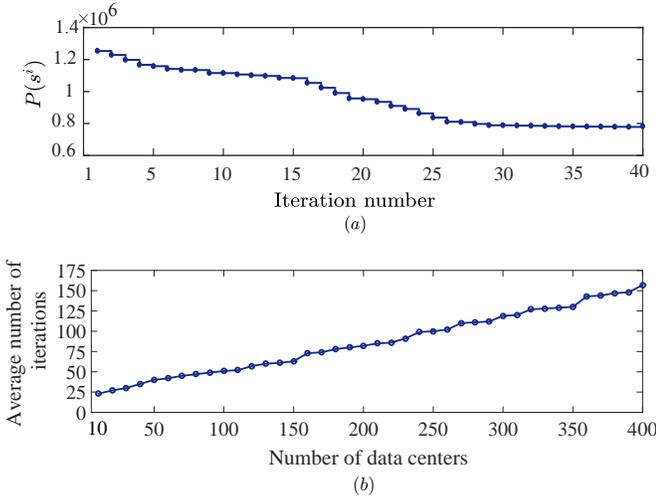


Fig. 7. (a) The value of potential function $P(s^i)$ per iteration for 50 data centers; (b) The required number of iterations for convergence versus the number of data centers.

a deregulated electricity market. We showed that with the RTP, the data centers' decisions of utility company choices and workload scheduling become coupled with each other and with the pricing decisions of the utility companies. We modeled the interaction among data centers as a many-to-one matching game with externalities. It was a challenge to prove the existence of the stable outcome of the underlying matching game. We addressed this challenge by constructing an exact potential function, whose local minima correspond to the stable outcomes of the game. Constructing an exact potential function also enabled us to develop a distributed algorithm to reach a stable outcome. Simulation results showed that the data centers can decrease their costs by 18.7% with the proposed algorithm, as they can purchase electricity from their preferred utility company and shift their demand to off-peak hours. Meanwhile, the utility companies can achieve an 8% reduction in the PAR. The utility companies with lower tariffs can increase their revenue by 80% through attracting more data centers as customers. For future work, we plan to extend the proposed data center's operation model by considering a time-dependent multiclass $G/G/1$ queuing system, and extend our proposed matching algorithm by taking into account the transmission lines congestion and power flow models.

APPENDIX

A. Multiclass $M/M/1$ Queuing System Model

In the first step, we determine an equivalent $M/M/1$ queuing model for the arrival and execution of the workloads of service class c in data center d and time slot t . In the second step, we use the transient behaviour of the underlying queuing system to show that a workload of service class c in data center d experienced the maximum expected execution time either at the beginning or at the end of each time slot t .

a) Consider data center $d \in \mathcal{D}$ in time slot $t \in \mathcal{T}$. There are multiple service classes in data center d . Hence, $\rho_d(t) - \bar{\lambda}_{c,d}(t)/\mu_{c,d}(t) = \sum_{c' \in \mathcal{C}_d, c' \neq c} \bar{\lambda}_{c',d}(t)/\mu_{c',d}(t)$ is the proportion of time slot t that the servers are busy with

executing the workloads of the service classes other than c . Therefore, $1 - (\rho_d(t) - \bar{\lambda}_{c,d}(t)/\mu_{c,d}(t))$ is the proportion of time slot t that the servers in data center d are not busy with executing the workloads of service classes other than c . Note that the parameter $\mu_{c,d}(t) = n_d(t)/\sigma_{c,d}$ is the average execution rate of the workloads in time slot t if all the operating servers *only* execute the workloads of service class c . Hence, in the underlying multiclass queuing system, by allocating $1 - (\rho_d(t) - \bar{\lambda}_{c,d}(t)/\mu_{c,d}(t))$ of time slot t to execute the workloads of service class c , the average workload execution rate becomes $\bar{\mu}_{c,d}(t) = \mu_{c,d}(t)(1 - (\rho_d(t) - \bar{\lambda}_{c,d}(t)/\mu_{c,d}(t)))$. Consequently, we can capture the process of the workloads of service class c in data center d and time slot t by an equivalent $M/M/1$ queuing system with an average workload arrival rate $\bar{\lambda}_{c,d}(t)$ and an execution rate $\bar{\mu}_{c,d}(t)$.

b) We use the transient behaviour of the $M/M/1$ queuing system corresponding to the workloads of service class c . We consider time slot t and use the continuous parameter τ to represent the time elapsed since the beginning of time slot t . Hence, τ varies from 0 to one time slot. We drop time index t from the parameters in the rest of the proof to avoid the potential of confusion. Consider the $M/M/1$ queuing system corresponding to the workloads of service class c in data center d and time slot t . The state of the queue represents the number of workloads in the system in time τ . Let $q_{c,d}(k_1, k_2, \tau)$ denote the probability of being in state k_2 in time τ when the initial state in time $\tau = 0$ is k_1 . For data center d , let $W_{c,d}(k_1, \tau)$ denote the expected execution time of a workload of service class c in time τ when the initial state is k_1 . We have

$$W_{c,d}(k_1, \tau) = \sum_{k=0}^{\infty} \left(\frac{k+1}{\bar{\mu}_{c,d}} \right) q_{c,d}(k_1, k, \tau). \quad (27)$$

We can rewrite (27) as

$$W_{c,d}(k_1, \tau) = \sum_{k=0}^{\infty} \left(\frac{k}{\bar{\mu}_{c,d}} \right) q_{c,d}(k_1, k, \tau) + \sum_{k=0}^{\infty} \left(\frac{1}{\bar{\mu}_{c,d}} \right) q_{c,d}(k_1, k, \tau). \quad (28)$$

The value of $\sum_{k=0}^{\infty} k q_{c,d}(k_1, k, \tau)$ in the first summation of (28) is equal to the expected number of workloads of service class c in time τ when the queue's initial state is k_1 . We denote this summation by $Q_{c,d}(k_1, \tau)$. For the summation in the second term, we have $\sum_{k=0}^{\infty} q_{c,d}(k_1, k, \tau) = 1$. We can rewrite (28) as

$$W_{c,d}(k_1, \tau) = \frac{Q_{c,d}(k_1, \tau) + 1}{\bar{\mu}_{c,d}}. \quad (29)$$

We now show that a workload of service class c in data center d experienced the maximum expected execution time either at the beginning or at the end of each time slot t . It is sufficient to show that function $W_{c,d}(k_1, \tau)$ in (29) is maximized when either $\tau = 0$ or τ is equal to one time slot.

From (29), it is sufficient to determine the time τ that maximizes function $Q_{c,d}(k_1, \tau)$. We can compute the derivative of function $Q_{c,d}(k_1, \tau)$ with respect to τ as [26]

$$Q'_{c,d}(k_1, \tau) = \bar{\lambda}_{c,d} - \bar{\mu}_{c,d} (1 - q_{c,d}(k_1, 0, \tau)). \quad (30)$$

We now consider the following three scenarios for k_1 , the initial number of workloads of service class c .

1) $k_1 = 0$: In this scenario, the queuing system corresponding to the workloads of service class c is initially empty. Thus, with probability of one, there is no workload in the queue in time $\tau = 0$, i.e., we have $q_{c,d}(k_1, 0, 0) = 1$. From (30), we have $Q'_{c,d}(k_1, 0) = \bar{\lambda}_{c,d}$, which is nonnegative.

When time τ elapses from 0 to one time slot, then the value of $q_{c,d}(k_1, 0, \tau)$ decreases from one to its steady state value. Meanwhile, from (30), the value of $Q'_{c,d}(k_1, \tau)$ decreases from $\bar{\lambda}_{c,d}$ and converges to zero gradually. Therefore, the value of $Q_{c,d}(k_1, \tau)$ increases gradually and converges to its steady state value, and hence functions $Q_{c,d}(k_1, \tau)$ and $W_{c,d}(k_1, \tau)$ are maximized when τ is equal to one time slot. In this scenario, the maximum value of $W_{c,d}(k_1, \tau)$ is $\frac{1}{\bar{\mu}_{c,d} - \bar{\lambda}_{c,d}}$.

2) $0 < k_1 \leq \frac{\bar{\lambda}_{c,d}}{\bar{\mu}_{c,d} - \bar{\lambda}_{c,d}}$: Initially there exist some workloads of service class c in the queue, but the number of initial workloads is at most $\frac{\bar{\lambda}_{c,d}}{\bar{\mu}_{c,d} - \bar{\lambda}_{c,d}}$. Thus, we have $q_{c,d}(k_1, 0, 0) = 0$. From (30), we have $Q'(k_1, 0) = \bar{\lambda}_{c,d} - \bar{\mu}_{c,d}$, which is negative.

When time τ elapses from 0 to one time slot, then the value of $q_{c,d}(k_1, 0, \tau)$ increases from zero to its steady state value. Meanwhile, $Q'_{c,d}(k_1, \tau)$ increases gradually from its initial value, i.e., $\bar{\lambda}_{c,d} - \bar{\mu}_{c,d}$, becomes positive, and then decreases to converges to zero. Therefore, the value of $Q_{c,d}(k_1, \tau)$ decreases at the beginning, and then increases gradually and converges to its steady state value. Hence, $Q_{c,d}(k_1, \tau)$ and $W_{c,d}(k_1, \tau)$ are maximized when τ is equal to one time slot.

3) $k_1 > \frac{\bar{\lambda}_{c,d}}{\bar{\mu}_{c,d} - \bar{\lambda}_{c,d}}$: We have $q_{c,d}(k_1, 0, 0) = 0$. From (30), we obtain $Q'(k_1, 0) = \bar{\lambda}_{c,d} - \bar{\mu}_{c,d}$, which is negative.

When time τ elapses from 0 to one time slot, then $q_{c,d}(k_1, 0, \tau)$ increases from zero to its steady state value, and hence the value of $Q'_{c,d}(k_1, \tau)$ increases from its initial negative value, i.e., $\bar{\lambda}_{c,d} - \bar{\mu}_{c,d}$ and converges to zero. Thus, $Q_{c,d}(k_1, \tau)$ decreases gradually and converges to its steady state value. In this scenario, $Q_{c,d}(k_1, \tau)$ and $W_{c,d}(k_1, \tau)$ are maximized when $\tau = 0$. The proof is completed ■

B. Proof of Theorem 1

To prove Theorem 1, we substitute (25) into the right-hand side of (24) and substitute (18) into the left-hand side of (24) for strategies \mathbf{s} and $\tilde{\mathbf{s}}$, and show that the results are the same. By changing the strategy of data center d from \mathbf{s}_d to $\tilde{\mathbf{s}}_d$, its utility company choice is changed from u to \tilde{u} , and its decision profile is changed from \mathbf{a}_d to $\tilde{\mathbf{a}}_d$. Hence, the matching structure is changed from m to \tilde{m} and the energy demand of data center d is changed from $e_d(t)$ to $\tilde{e}_d(t)$ in time slot $t \in \mathcal{T}$.

By substituting (25) into the right-hand side of (24) for $\mathbf{s} = (\mathbf{s}_d, \mathbf{s}_{-d})$ and $\tilde{\mathbf{s}} = (\tilde{\mathbf{s}}_d, \mathbf{s}_{-d})$, we obtain

$$\begin{aligned} P(\mathbf{s}_d, \mathbf{s}_{-d}) - P(\tilde{\mathbf{s}}_d, \mathbf{s}_{-d}) = & \\ \sum_{t \in \mathcal{T}} & \left((1 - \omega_d^{\text{cvar}}) \left(\kappa_u(t) \left(e_d^2(t) + e_u^{\text{other}}(t) e_d(t) \right) \right. \right. \\ & \left. \left. + \sum_{d' < d' \in m(u)} e_d(t) e_{d'}(t) \right) + p(t) e_d(t) \right) \end{aligned}$$

$$\begin{aligned} & - \kappa_{\tilde{u}}(t) \left(\tilde{e}_d^2(t) + e_u^{\text{other}}(t) \tilde{e}_d(t) \right) \\ & + \sum_{d' < d' \in \tilde{m}(\tilde{u})} \tilde{e}_d(t) e_{d'}(t) - p(t) \tilde{e}_d(t) \Big) \\ & + \omega_d^{\text{cvar}} \left(\Gamma_{d,\beta_d}(\alpha_d, \boldsymbol{\lambda}_d, \mathbf{e}_d^r) - \Gamma_{d,\beta_d}(\tilde{\alpha}_d, \tilde{\boldsymbol{\lambda}}_d, \tilde{\mathbf{e}}_d^r) \right). \end{aligned} \quad (31)$$

In (31), the terms related to the data centers other than data center d cancel each other, since the strategy of other data centers are unchanged in $\tilde{\mathbf{s}} = (\tilde{\mathbf{s}}_d, \mathbf{s}_{-d})$. By substituting (18) into the left-hand side of (24) for $\mathbf{s} = (\mathbf{s}_d, \mathbf{s}_{-d})$ and $\tilde{\mathbf{s}} = (\tilde{\mathbf{s}}_d, \mathbf{s}_{-d})$, we have

$$\begin{aligned} c_d(\mathbf{s}_d, \mathbf{s}_{-d}) - c_d(\tilde{\mathbf{s}}_d, \mathbf{s}_{-d}) = & \\ \sum_{t \in \mathcal{T}} & \left((1 - \omega_d^{\text{cvar}}) \left(e_d(t) p_u^r(e_u(t), m) - \tilde{e}_d(t) p_{\tilde{u}}^r(e_{\tilde{u}}(t), \tilde{m}) \right) \right) \\ & + \omega_d^{\text{cvar}} \left(\Gamma_{d,\beta_d}(\alpha_d, \boldsymbol{\lambda}_d, \mathbf{e}_d^r) - \Gamma_{d,\beta_d}(\tilde{\alpha}_d, \tilde{\boldsymbol{\lambda}}_d, \tilde{\mathbf{e}}_d^r) \right). \end{aligned} \quad (32)$$

By substituting the retail price (1) into (32), the cost change for data center d will be equal to the potential function change in (31). This completes the proof. ■

C. Proof of Theorem 2

We first show that the global minimum of the potential function (25) is a stable outcome. Let $\mathbf{s}^* = (\mathbf{a}_{m^*}, m^*)$ be the global minimum of $P(\mathbf{s})$. Thus, if data center d changes its action profile to \mathbf{a}_d or its utility company to $m(d)$ unilaterally, then 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. Thus, no unilateral deviation from \mathbf{s}^* can reduce the cost of any data center, and hence \mathbf{s}^* is a stable outcome of Game 1. As the exact potential function (25) has at least one global minimum, the matching game has at least one stable outcome.

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

D. Proof of Theorem 3

Since the potential function (25) is lower bounded (it is always positive), 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. Our proof involves two steps:

Step a) We show that the potential function decreases when the data centers update their utility company choices in Line 12 of Algorithm 1. Each data center updates its utility company choice to reduce its cost. Nevertheless, we cannot directly use (24) in Definition 3 to conclude that the potential

function decreases accordingly. In fact, equality (24) holds if a data center changes its utility company choice unilaterally. However, in Line 12 of Algorithm 1, several data centers may update their utility company choices simultaneously.

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.

Base case: Consider $k = 1$. If one data center updates its utility company choice, then it corresponds to a unilateral change in the strategy of the data center to reduce its cost. Hence, from (24) in Definition 3, the potential function $P(\mathbf{s}^i)$ decreases, when the cost of a data center decreases.

Induction step: Consider $k = d$. Suppose that the potential function $P(\mathbf{s}^i)$ decreases, when d data centers change their utility company choices in Line 12 of Algorithm 1. We prove that the potential function also decreases when $k = d + 1$ data centers change their utility company choices simultaneously. Without loss of generality, we denote the index set of data centers that change their utility company choices as $\{1, \dots, d+1\}$, which is partitioned into two sets: $\{1, \dots, d\}$ and $\{d+1\}$. The simultaneous changes in the utility company choices of data centers $\{1, \dots, d+1\}$ correspond to the simultaneous changes in the utility company choices of data centers $\{1, \dots, d\}$, and the change of utility company choice for data center $d + 1$. From our induction assumption, when data centers $\{1, \dots, d\}$ change their utility company choices, the potential function decreases. In the following, we show that the potential function further decreases when data center $d + 1$ changes its matched utility company. Suppose that data center $d + 1$ decides to leave utility company $m^i(d + 1) = \{u_{d+1}^i\}$ in iteration i and connect to utility company $m^{i+1}(d + 1) = \{u_{d+1}^{i+1}\}$. According to Line 7 of Algorithm 1, such a decision is the best response of data center $d + 1$. Hence, the cost of data center $d + 1$ should decrease from \tilde{c}_{d+1}^i to \tilde{c}_{d+1}^{i+1} , i.e., $\tilde{c}_{d+1}^{i+1} < \tilde{c}_{d+1}^i$ if other data centers do not change their utility company choices. However, data centers $\{1, \dots, d\}$ change their utility company choices at the same time. Therefore, the cost of data center $d + 1$ changes from c_{d+1}^i to c_{d+1}^{i+1} . We show that $c_{d+1}^i \geq \tilde{c}_{d+1}^i$ and $c_{d+1}^{i+1} \leq \tilde{c}_{d+1}^{i+1}$. Then, inequality $\tilde{c}_{d+1}^{i+1} < \tilde{c}_{d+1}^i$ will lead to inequality $c_{d+1}^{i+1} < c_{d+1}^i$.

We first show that $c_{d+1}^i \geq \tilde{c}_{d+1}^i$. A utility company accepts at most one termination request from data centers per iteration. Thus, data center $d + 1$ is the only one that leaves utility company u_{d+1}^i . However, utility company u_{d+1}^i may accept a new connection request from a data center in set $\{1, \dots, d\}$, which increases the total demand from this utility company. Therefore, the payment of data center $d + 1$ to utility company u_{d+1}^i will increase after updating the matching of data centers $\{1, \dots, d\}$. That is, we have $c_{d+1}^i \geq \tilde{c}_{d+1}^i$.

Next we show that $c_{d+1}^{i+1} \leq \tilde{c}_{d+1}^{i+1}$. 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 u_{d+1}^{i+1} . Nevertheless, utility company u_{d+1}^{i+1} may accept a termination request from a data center in set $\{1, \dots, d\}$, which decreases the total demand from this utility company. Hence, the payment of data center $d + 1$ to utility

company u_{d+1}^{i+1} will decrease after updating the matching of data centers $\{1, \dots, d\}$. That is, we have $c_{d+1}^{i+1} \leq \tilde{c}_{d+1}^{i+1}$.

We can conclude that after data centers in set $\{1, \dots, d\}$ change their utility company choices, we have $c_{d+1}^{i+1} < c_{d+1}^i$. From (24) in Definition 3, the exact potential function decreases when the cost of data center $d + 1$ decreases. We have shown that the potential function will decrease when $k = d + 1$ data centers change their utility company choices simultaneously. By the principle of induction, the potential function decreases when any number of data centers change their utility company choices simultaneously.

Step b) Under a given matching m^{i+1} , (24) implies that $\nabla_{\mathbf{a}_d^i} c_d(\mathbf{a}_d^i, \mathbf{a}_{-d}^i, m^{i+1}) = \nabla_{\mathbf{a}_d^i} P(\mathbf{a}_d^i, \mathbf{a}_{-d}^i, m^{i+1})$. If all data centers use (26) for update, the potential function varies in the opposite direction of its gradient, $\nabla_{\mathbf{a}_d^i} P(\mathbf{a}_d^i, \mathbf{a}_{-d}^i, m^{i+1})$. Under a given matching m^{i+1} , $P(\cdot)$ is a convex function of \mathbf{a}^i and has Lipschitz continuous derivative. Thus, for sufficiently small step size, the opposite gradient direction is a decreasing direction. We note that the step sizes γ_d^i are not required to be equal for all data centers $d \in \mathcal{D}$. Unequal but sufficiently small step sizes lead to updating the decision vector \mathbf{a}^i in the opposite subgradient direction of the potential function. ■

REFERENCES

- [1] "2017 top markets report smart grid," Int'l Trade Administration, U.S. Department of Commerce, Annual Report, Jan. 2017.
- [2] S. Borenstein and J. Bushnell, "The U.S. electricity industry after 20 years of restructuring," *Annual Review of Economics*, vol. 7, no. 1, pp. 437–463, Apr. 2015.
- [3] Alberta Electric System Operator (AESO) website. [Online]. Available: <http://www.aeso.ca>
- [4] Nordic Energy Regulators website. [Online]. Available: www.nordicenergyregulators.org
- [5] A. Mazer, *Electric Power Planning for Regulated and Deregulated Markets*. Hoboken, NJ: John Wiley & Sons, 2007.
- [6] Galvin Electricity Initiative. [Online]. Available: <http://www.galvinpower.org/power-consumers/act/real-time-illinois/enroll-now>.
- [7] 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.
- [8] J. Li, Z. Li, K. Ren, and X. Liu, "Towards optimal electric demand management for Internet data centers," *IEEE Trans. on Smart Grid*, vol. 3, no. 1, pp. 183–192, Mar. 2012.
- [9] 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.
- [10] M. Ghamkhari and H. Mohsenian-Rad, "Energy and performance management of green data centers: A profit maximization approach," *IEEE Trans. on Smart Grid*, vol. 4, no. 2, pp. 1017–1025, Jun. 2013.
- [11] T. Chen, Y. Zhang, X. Wang, and G. B. Giannakis, "Robust workload and energy management for sustainable data centers," *IEEE J. on Selected Areas in Commun.*, vol. 34, no. 3, pp. 651–664, Mar. 2016.
- [12] L. Rao, X. Liu, L. Xie, and W. Liu, "Coordinated energy cost management of distributed Internet data centers in smart grid," *IEEE Trans. on Smart Grid*, vol. 3, no. 1, pp. 50–58, Mar. 2012.
- [13] H. Wang, J. Huang, X. Lin, and A. 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.
- [14] 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.
- [15] Data Center Map website. [Online]. Available: www.datacentermap.com/texas/
- [16] Electric Choice website. [Online]. Available: www.electricchoice.com/tx/houston-electric-rates/

- [17] 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.
- [18] S. Bahrami, V.W.S. Wong, and J. Huang, "Demand response for data centers in deregulated markets: A matching game approach," in *Proc. of IEEE SmartGridComm*, Dresden, Germany, Oct. 2017.
- [19] A. E. Roth and M. Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1992.
- [20] 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.
- [21] J. Yang, J. Zhao, F. Luo, F. Wen, and Z. Y. Dong, "Decision-making for electricity retailers: A brief survey," accepted for publication in *IEEE Trans. on Smart Grid*, 2017.
- [22] J. Haraguchi and T. Matsumura, "Cournot–Bertrand comparison in a mixed oligopoly," *J. of Economics*, vol. 117, no. 2, Mar. 2016.
- [23] N. Forouzandehmehr, M. Esmalifalak, A. 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.
- [24] L. Kleinrock, *Queueing Systems, Volume 1: Theory*. Wiley-Interscience, 1976.
- [25] A. Wierman, Z. Liu, I. Liu, and A. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *Proc. of Int'l Green Computing Conference (IGCC)*, Dallas, TX, Nov. 2014.
- [26] 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. 631–667, 1976.
- [27] 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.
- [28] G. C. Pflug, "Some remarks on the value-at-risk and the conditional value-at-risk," in *Probabilistic Constrained Optimization*, P. U. Stanislav, Ed. Springer, US, 2000, pp. 272–281.
- [29] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *J. of Risk*, vol. 2, pp. 21–42, Apr. 2000.
- [30] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 3124–143, May 1996.
- [31] S. Durand and B. Gaujal, *Complexity and Optimality of the Best Response Algorithm in Random Potential Games*. Berlin, Heidelberg: Springer, Sept. 2016, pp. 40–51.
- [32] Independent Electricity System Operator (IESO). [Online]. Available: <http://www.ieso.ca>.
- [33] [Online]. Available: <http://ita.ee.lbl.gov/html/contrib/WorldCup.html>.



Vincent W.S. Wong (S'94, M'00, SM'07, F'16) received the B.Sc. degree from the University of Manitoba, Winnipeg, MB, Canada, in 1994, the M.A.Sc. degree from the University of Waterloo, Waterloo, ON, Canada, in 1996, and the Ph.D. degree from the University of British Columbia (UBC), Vancouver, BC, Canada, in 2000. From 2000 to 2001, he worked as a systems engineer at PMC-Sierra Inc. (now Microsemi). He joined the Department of Electrical and Computer Engineering at UBC in 2002 and is currently a Professor. His research areas include protocol design, optimization, and resource management of communication networks, with applications to wireless networks, smart grid, mobile cloud computing, and Internet of Things. Dr. Wong is an Editor of the *IEEE Transactions on Communications*. He has served as a Guest Editor of *IEEE Journal on Selected Areas in Communications* and *IEEE Wireless Communications*. He has also served on the editorial boards of *IEEE Transactions on Vehicular Technology* and *Journal of Communications and Networks*. He was a Technical Program Co-chair of *IEEE SmartGridComm'14*, as well as a Symposium Co-chair of *IEEE ICC'18*, *IEEE SmartGridComm ('13, '17)* and *IEEE Globecom'13*. He is the Chair of the IEEE Vancouver Joint Communications Chapter and has served as a Chair of the IEEE Communications Society Emerging Technical Sub-Committee on Smart Grid Communications. He received the 2014 UBC Killam Faculty Research Fellowship.



Jianwei Huang (F'16) is a Professor in the Department of Information Engineering at the Chinese University of Hong Kong. He is the co-author of 9 Best Paper Awards, including IEEE Marconi Prize Paper Award in Wireless Communications 2011. He has co-authored six books, including the textbook on "Wireless Network Pricing". He has served as the Chair of IEEE TCCN and MMTTC. He is an IEEE ComSoc Distinguished Lecturer and a Clarivate Analytics Highly Cited Researcher.



Shahab Bahrami (S'12) received the B.Sc. and M.A.Sc. degrees both from Sharif University of Technology, Tehran, Iran, in 2010 and 2012, respectively. He received the Ph.D. degree from the University of British Columbia (UBC), Vancouver, BC, Canada in 2017. Dr. Bahrami continued to work as a post-doctoral research fellow at UBC until Jan. 2018. His research interests include optimal power flow analysis, game theory, and demand side management, with applications to smart grid.