How Geo-Distributed Data Centers Do Demand Response: A Game-Theoretic Approach

Nguyen H. Tran, *Member, IEEE*, Dai H. Tran, Shaolei Ren, *Member, IEEE*, Zhu Han, *Fellow, IEEE*, Eui-Nam Huh, *Member, IEEE*, and Choong Seon Hong, *Senior Member, IEEE*

***Abstract*—We study the demand response (DR) of geo-distributed data centers (DCs) using smart grid's pricing signals set by local electric utilities. The geo-distributed DCs are suitable candidates for the DR programs due to their huge energy consumption and flexibility to distribute their energy demand across time and location, whereas the price signal is well-known for DR programs to reduce the peak-to-average load ratio. There are two dependencies that make the pricing design difficult: 1) dependency among utilities; and 2) dependency between DCs and their local utilities. Our proposed pricing scheme is constructed based on a two-stage Stackelberg game in which each utility sets a real-time price to maximize its own profit in Stage I and based on these prices, the DCs' service provider minimizes its cost via workload shifting and dynamic server allocation in Stage II. For the first dependency, we show that there exists a unique Nash equilibrium. For the second dependency, we propose an iterative and distributed algorithm that can converge to this equilibrium, where the "right prices" are set for the "right demands." We also verify our proposal by trace-based simulations, and results show that our pricing scheme significantly outperforms other baseline schemes in terms of flattening the power demand over time and space.**

***Index Terms*—Data centers (DCs), demand response (DR), Nash equilibrium, smart grids, Stackelberg games.**

I. Introduction

Data centers (DCs) are well-known as large-scale consumers of electricity (e.g., DCs consumed 1.5% of the worldwide electricity supply in 2011 and this fraction is expected to grow to 8% by 2020 [1]). A recent study shows that many DC operators paid more than \$10M [2] on their annual electricity bills, which continues to rise with the flourishing of cloud-computing services. Therefore, it is necessary for DC operators to both cut costs and increase performances. Recent works have shown that DC operators can save more than 5%–45% [3] operation cost by leveraging time and location diversities of electricity market prices to optimize geo-distributed DCs. However, most of the existing research is based on one important assumption: the electricity price applying to DCs does not change with demand. This assumption may not be true since an individual DC with enormous energy consumption (e.g., Facebook's DC in Crook County, Oregon can contributed up to 50% of the total load of its distribution grid [4]) will impact to the supply demand balance of its local utility, which in turn can alter the utility's price as shown in recent studies [5]–[7]. Furthermore, the power grid can be negatively affected due to this assumption. For example, blackouts might happen due to overloads in these areas where the DCs operator shifts all of its energy demand to a local utility with a low price and a high enough background load.

To make the power grid more reliable and robust, tremendous research and industry efforts have focused on building the next-generation power grids, known as smart grids. Due to its efficiency and potential, many studies consider how DC operators can run their geo-distributed DCs on smart grids that support two-way information exchange between utilities and customers [5], [8], [9]. An important feature of smart grids is demand response (DR). DR programs seek to provide incentives to induce dynamic demand management of customers' electricity load in response to power supply conditions. For example, just before the peak load hours, a utility can send the warning signal to customers' smart meters which will automatically schedule their demands to reduce the power consumption. Due to their huge and rapidly increasing energy consumption, DCs should be significantly encouraged to participate in the DR programs. Furthermore, with the recent trend in dynamic server capacity provision and flexibility of workload shifting, geo-distributed DCs have a great potential to easily adapt the DR programs. One of the DR programs is using real-time pricing schemes to reduce the peak-to-average (PAR) load ratio by encouraging customers to shift their energy demand away from peak hours. The challenge of an effective pricing scheme is how to charge the customers with a right price not only at the right time and right place but also on the right amount of customers' demand. A real-time pricing scheme is considered effective if it can mitigate the large fluctuation of energy consumption between peak and off-peak hours to increase power grid's reliability and robustness.

Manuscript received October 6, 2014; revised February 2, 2015; accepted March 30, 2015. This research was supported in part by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2014R1A2A2A01005900), and in part by the Ministry of Science, ICT, and Future Planning (MISP), Korea, under the Information Technology Research Center (ITRC) support program (IITP-2015-(H8501-15-1015) supervised by the IITP. Paper no. TSG-00990-2014. *(Corresponding author: Choong Seon Hong.)*

N. H. Tran, D. H. Tran, E.-N. Huh, and C. S. Hong are with the Department of Computer Engineering, Kyung Hee University, Yongin 446-701, Korea (e-mail: cshong@khu.ac.kr).

S. Ren is with the School of Computing and Information Sciences, Florida International University, Miami, FL 33199 USA.

Z. Han is with the Department of Electrical and Computer Engineering, University of Houston, Houston, TX 77004 USA.

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

Digital Object Identifier 10.1109/TSG.2015.2421286

In this paper, we consider the problem of using real-time pricing of utilities to enable the geo-distributed DCs' participation into the DR program. In this program, while geo-distributed DCs employ workload shifting and dynamic server provisioning in response to the price signal, the role of local utilities is how to set the real-time prices to flatten the customers' demand load. It can be observed that there is an interaction between geo-distributed DCs and their local utilities; and it is the first challenge of this DR problem that we call vertical dependency. Specifically, when participating in the DR program, a DCs' operator will distribute its energy demand geographically based on the electric prices adjusted intelligently by the local utilities. However, the utilities set their prices based on the total demand including the DCs' demand, which is only known when the price is available. We clearly see that this dependency makes it difficult for both DCs and utilities to make their decisions. The second important challenge, which is less obvious, is an interaction among local utilities feeding power to the geo-distributed DCs; and we call it horizontal dependency. Specifically, the DCs' decisions on workload shifting and server allocation depend on the electric prices set by local utilities; therefore, if any subset of the local utilities change their prices, it can lead to the DCs' decision changing. Since the utilities are noncooperative (i.e., no information exchange) in practice, how to design a pricing mechanism that can enable an equilibrium price setting profile is the bottleneck of this DR program.

To tackle the above discussed challenges, our contributions can be summarized as follows.

1) We transform the functional space of the geo-distributed DCs' DR program into a mathematical space of a formulated two-stage Stackelberg game. In this game, each utility will set a real-time price to maximize its own profit in Stage I; and given these prices, the DCs' operator will minimize its cost via workload shifting and dynamic server allocation in Stage II. We also utilize the backward induction method to find the Stackelberg equilibria of this two-stage game.
2) Based on the Stackelberg equilibria, our proposed scheme can deal with the inherent challenges of this DR as follows: first, the horizontal dependency between utilities are characterized as a strategic game in Stage I, and we show that there exists a Nash equilibrium in this game. Second, we propose an iterative and distributed algorithm to achieve the Stackelberg equilibrium. In this algorithm, the DCs and utilities exchange their information (i.e., DCs' demand and utilities' prices) iteratively until the algorithm converges. We also examine the algorithm's convergence where the "right prices" are set for the "right demands" as a solution for the vertical dependency issue.
3) Finally, we perform a real-world trace-based simulation to solidify the analysis. The results show that our proposed pricing scheme can flatten the workload not only over time but also over space to improve the power grid's reliability and robustness.

The rest of this paper is organized as follows. Section II is about related work. Section III presents the system model and the two-stage Stackelberg game. We analyze this game and propose a distributed algorithm in Section IV. Section V provides the trace-based simulation results. Finally, Section VI concludes this paper.

II. Related Work

DR is identified as one of high-prioritized areas for future smart grids [10]–[12] with its potential to reduce up to 20% of the total peak electricity demand of the U.S. [13]. Most DR proposals, which try to incentivize customers to manage their demand dynamically in response to the power supply conditions, mostly targeted to residential customers [14]–[17]. On the other hand, most of the existing research on DCs, which can be classified as medium or large industrial customers, mainly focus on their cost minimization that takes the electricity price for granted [3], [18], [19], which does not follow any DR programs. However, due to the important role of DCs in DR programs, DRs of DCs recently receive significant attention [4], [7]–[9], [20]–[22].

For those work considering DR of geo-distributed DCs, based on the interactions between DCs and utilities, we simply divide them into two categories.

1) One-Way Interaction: One of the most popular DR programs of DCs is coincident peak pricing (CPP), which is studied in [21]. CPP charges very high prices for power usage during the coincident peak hour at which the most electric demands is requested to the utility. By predicting the upcoming potential peak hours, the utilities send a warning signal (i.e., not a price) to help customers schedule their power consumption. However, current DCs do not respond actively to the warning signals due to the uncertainty of these warnings [21], which motivates researchers to devise more effective DR approaches. Wang *et al.* [7] used a "prediction-based" method where the customers (DCs) respond to the prices which are chosen based on a supply function. This supply function can be modeled using some data fitting methods based on history. Hence, in this paper only customers respond to a predicted price while there is no action from the power suppliers to set the prices corresponding to the demand.

2) Two-Way Interaction: There are three recent papers [5], [8], [9] in this category. The first two papers, which are highly related to this paper, consider dynamic pricing mechanisms with the coupling between utilities and DCs, whereas the last one proposed that DCs can participate in the spot market via a broker, which is a significant departure from our model. Moreover, the system model of [9] assumes that all utilities cooperate to solve a social optimization problem, which is not relevant to current practice since there is no information exchange between utilities in reality. On the other hand, the pricing scheme of [8] is based on a heuristic approach, which cannot maximize the utilities' profit as well as minimize their cost. This paper falls into this category of two-way interaction, yet is different from others in terms of its two-stage game-theoretic approach to tackle the vertical and horizontal coupling issues, which are not addressed in the literature, between geo-distributed DCs and local utilities.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

TABLE I
SUMMARY OF NOTATIONS

Notation	Description
I	Number of tenants
Λ	Total workload at front-end server
λ_i	Workload at DC i
d_i	Transmission delay from front-end server to DC i
γ	Weight of utility cost
e_i^d	Interactive-job energy
e_i^b	Batch-job energy
e_i	Total energy of DC i
ω	Weight factor of migration cost
p_i	Price at DC i
μ_i	Service rate of DC i
s_i	Number of active servers
B_i	Background energy demand of utility i
α_i, β_i	Parameters of background demand B_i model
D_i	Maximum average delay of DC i
C_i	Capacity of utility i

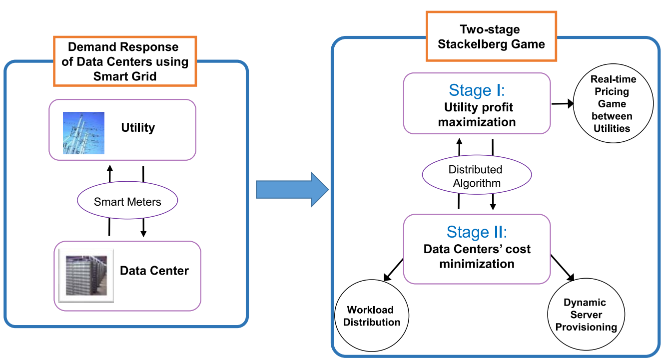

Fig. 1. Functional space of the geo-distributed DCs' DR on the left and its transformed mathematical space as a two-stage Stackelberg game on the right.

III. SYSTEM MODEL AND PROBLEM FORMULATION

We consider one-period DR as in [18] and [23], where its duration, which is controlled by a utility/load serving entity, matches an interval at which the DCs' decisions and utilities' real-time prices can be updated (such as 15 min or 1 h). Let $\mathcal{I} = \{1, \dots, I\}$ denote the set of sites with different electrical utilities' service regions where DCs are located. Such geo-distributed DCs are very common in practice, e.g., Google, Amazon, etc. Each DC i is powered by a local utility company and have S_i homogeneous servers. A DC with heterogeneous types of servers can be viewed as multiple virtual DCs each having homogeneous servers. For the ease of presentation, Table I lists key notations of this paper.

We incorporate the role of utility into the DR programs of DCs to regulate the power demand at each local site for load balancing the power grid. We illustrate a functional space and a mathematical space of this DR program in Fig. 1. In the functional space, we leverage the idea of using the advanced two-way communication of smart grid to facilitate the information exchange between utilities and DCs at each local site via smart meters. While utilities set prices to incentivize DCs to flatten the demand over time and locations to increase the power grid's reliability, as the price-takers the DCs will minimize their costs. In the mathematical space, we observe that there exists a special mutual interaction between DCs and utilities where utilities set prices based on the total demand, and DCs minimize their costs based on the prices. Therefore, we transform this DCs' DR program into a leader–follower game that can be studied using a two-stage Stackelberg game. Specifically, the utilities are the leaders that set the prices to maximize their profits in Stage I and DCs will make their decisions on workload shifting and dynamic server provisioning to minimize their costs in Stage II. We present this two-stage game formulation in the reverse sequence, starting with Stage-II optimization problem.

A. DCs' Cost Minimization in Stage II

We first describe the workload model of a typical DC. We then elaborate the DCs' cost focusing on the energy cost and delay cost model. Finally, we formulate the Stage-II DCs cost minimization.

1) Workload Model: Even though DCs can support a wide range of workloads, we generally classify them into two typical types of workload: interactive (noninterruptive) jobs and batch (interruptive) jobs. While the former is delay-sensitive (e.g., computing search, online game, etc.), the latter is delay-tolerant (e.g., backup tasks, MapReduce, etc.). We assume that each DC processes its batch jobs locally (i.e., batch jobs cannot be redirected to other DCs for load balancing) since without stringent delay constraint, they are flexible to be scheduled across a large time window at a local site, like [19]. For interactive jobs, we denote the total arrival rate to the DCs' front-end server [i.e., all DCs are managed by a DCs service provider (DCs provider)], by Λ and this front-end server is responsible for splitting the total incoming workload Λ into separate workloads of geo-dispersed DCs, denoted by $\{\lambda_i\}_{i \in \mathcal{I}}$. Even though we only consider workload shifting, the other control knobs for DR such as power load reduction (e.g., scaling down CPU frequencies and/or turning off unused servers) can also be integrated into our framework.

2) DCs Cost and SLA Model: We assume that the DCs provider tries not only to minimize its energy cost and migration cost but also to guarantee the service level agreement (SLA) requirements for the interactive jobs.

a) Energy cost: Since batch jobs are flexible to schedule in time domain, batch jobs processing is considered to consume an amount of energy e_i^b of each DC i with their dedicated servers. On the other hand, the energy consumption of interactive jobs at DC i is [2]

$$e_i^d = s_i\big(P_{\text{idle}} + \big(P_{\text{peak}} - P_{\text{idle}}\big)U_i + (\eta - 1)P_{\text{peak}}\big) \qquad (1)$$

where s_i is the number of active servers, μ_i is the service rate of a server, P_{peak} and P_{idle} are the server's peak and idle power, respectively, $U_i = \lambda_i / s_i \mu_i$ is the average server utilization, and η is the power usage effectiveness (PUE) measuring the energy efficiency of the DC. We can rewrite e_i^d as follows:

$$e_i^d = a_i \lambda_i + b_i s_i, \quad \forall i \in \mathcal{I} \qquad (2)$$

where $a_i = (P_{\text{peak}} - P_{\text{idle}})/\mu_i$ and $b_i = P_{\text{idle}} + (\eta - 1)P_{\text{peak}}$. Therefore, denoting the total energy by

$$e_i = e_i^d + e_i^b \tag{3}$$

and given a price p_i, the energy cost of DC i is $e_i p_i$.

b) Migration cost: Since migrating the workload from front-end server to geo-distributed DCs can be very costly [e.g., migrating virtual machines or video content requests over the Internet could be expensive due to reserving bandwidth from an Internet service provider (ISP)], we model the migration cost to DC i as $\omega d_i c_i(\lambda_i)$, where d_i is the transmission delay from the front-end server to DC i, ω is a weight factor and $c_i(\lambda_i)$ is a function which is assumed to be strictly increasing and convex. Since d_i is proportional to the distance, it is assumed to be a constant and we see that migrating more requests from the front-end server to a more distant DC is more costly. For analysis tractability, we choose a quadratic function $c_i(\lambda_i) = \lambda_i^2$ since it is widely used in many fields such as control, signal processing, communication networks, etc. to model a cost function [24].

c) SLA constraint: We assume that each delay-sensitive request imposes a maximum delay D_i that the DCs provider has to guarantee when shifting this request to DC i. Therefore, the SLA constraint in terms of delay guarantee can be modeled as follows:

$$\frac{1}{s_i\mu_i - \lambda_i} + d_i \le D_i, \quad \forall i \tag{4}$$

where $1/(s_i\mu_i - \lambda_i)$ is the average delay time of a request processed in DC i with arrival rate λ_i and service rate $s_i\mu_i$ by queueing theory, which has been widely used as an analytic vehicle to provide a reasonable approximation for the actual service process [19], [25].

3) Problem Formulation: Our model focuses on two key controlling "knobs" of DCs' cost minimization: the workload shifting to DC λ_i and the number of active servers provisioned s_i at site i, $\forall i$. Then, the Stage-II DC cost minimization is given by

$$\begin{aligned} \textsf{DC} : \text{minimize} \quad & \sum_{i=1}^{I} e_i p_i + \omega d_i \lambda_i^2 && (5)\\ \text{subject to} \quad & \text{constraints (2)–(4)} \\ & \sum_{i=1}^{I} \lambda_i = \Lambda && (6)\\ & 0 \le s_i \le S_i, \ \forall i && (7)\\ & 0 \le \lambda_i \le s_i\mu_i \ \forall i && (8)\\ \text{variables} \quad & s_i, \lambda_i, \ \forall i. && (9)\end{aligned}$$

While constraints (2)–(4) are the definitions of the objective function and the SLA constraint, the remaining constraints are straight forward. In (6), all of the incoming workload must be served by some DCs. Moreover, (7) limits the number of active servers and (8) means that the total workload assigned to a DC must be less than its capacity. With thousands of servers in a DC, we can further relax the integer variables s_i as continuous variables so that this problem is tractable [18].

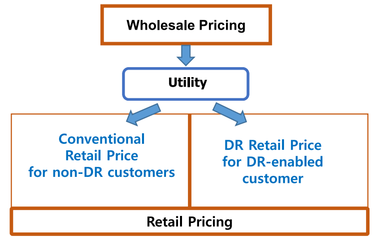

Fig. 2. Besides conventional wholesale and retail pricing, the utilities' DR real-time pricing is proposed for geo-distributed DCs and other DR-enabled customers.

B. Noncooperative Pricing Game in Stage I

In this stage, we first present the market structure. We next describe the utility's revenue and cost models and finally formulate the noncooperative pricing game.

1) DR Retail Price: Traditionally utilities involve many complex electricity markets. As buyers, utilities can participate in a wholesale market (day head, real-time balancing) to buy electricity from the generating companies with wholesale prices. As sellers, utilities make profit by selling retail to their customers with proper retail rates [5]. Since conventional customers (i.e., no DR) have inelastic demand with predictable patterns, utilities can predict and buy energy from wholesale market, then resell it at the conventional retail rates. However, DCs with workload shifting represent a new type of elastic-demand customers, which makes utilities difficult to predict their demands, impacting the grid's stability. Therefore, we propose a new DR retail pricing scheme for utilities to serve the unpredicted and elastic customers, e.g., not just load-shifting DCs but also for all DR-enabled customers. The basic idea of this scheme is that utilities and these DR-enabled customers can coordinate via smart-grid infrastructure to match supply with demand. Fig. 2 illustrates that utilities can apply the conventional and DR retail prices to their corresponding customers, which are complementary to each other so that the proposed scheme will not affect to the conventional scheme, similar to [9], [17], and [26]. Since conventional markets and customers are orthogonal to our model, henceforth, we only consider utility's profit model and the proposed real-time pricing scheme for DR-enable customers.

2) Utility's Revenue and Cost Model: We see that the optimal energy consumption of DCs that can be obtained from solving $\textsf{DC}$ depends on all utilities' prices. Denote the corresponding optimal power demand by $e_i(p)$, where $p := \{p_i\}_{i\in\mathcal{I}}$. We further assume that due to the grid regulations at each region, the lower and upper bound of the real-time price should be imposed and denoted by p_i^l and p_i^u, $\forall i, t$, respectively. Furthermore, besides the power demand of DCs, each utility has its own background load (e.g., residential/commercial/industrial demand). Since there are considerable works focusing on the residential DR programs, we assume that the background load of utility i, denoted by $B_i(p_i)$, also responds to the price and can be modeled by the following function:

$$B_i(p) = \begin{cases} B_i^l, & p_i \le p_i^l \\ \alpha_i - \beta_i p_i, & p_i^l \le p_i \le p_i^u \\ B_i^u, & p_i \ge p_i^u \end{cases} \tag{10}$$

where B_i^l and B_i^u are the minimum of maximum background demands of site i due to the physical constraints of consumers (i.e., minimum and maximum power of electric devices or vehicles). This function, which follows the linear demand model in [27], shows an inherent response of customers to the price: decrease the demand down to a lower-bound constraint when the price increases, and vice versa, where β_i is the decreasing slope and α_i models the physical upper-bound demand without price. Based on the history of customer's usage data, utilities can estimate α_i and β_i using some data fitting methods, similar to [7]. Based on the total power requested by DCs and background's demands, the revenue of utility i is given by

$$\mathcal{R}_i(p) = (e_i(p) + B_i(p_i))p_i. \tag{11}$$

On the other hand, every utility incurs a cost when it serves the customers' load. When the load increases, the utility's cost also increases since normally blackouts happen due to overload, which is a disaster to any utility. Hence, we can model the utility's cost based on a widely used electric load index (ELI) as follows:

$$\mathcal{C}_i(p) = \gamma \text{ELI} := \gamma r_i^2 C_i = \left(\frac{e_i(p) + B_i(p_i)}{C_i}\right)^2 C_i \tag{12}$$

where C_i is utility i capacity, γ reflects the weight of the cost, and r_i is a load ratio that measures the power load levels. A very high r_i can risk the utility's stability. ELI is motivated by the index measurement techniques used for load flattening in a power grid [9], [28], [29]. We see that ELI can weight different utilities' load ratio r_i by their capacities, providing feeder load-balancing capability. On the other hand, a utility with high γ shows that it is more concerned about the effect of ELI to the reliability, while a utility with low γ has more interest in making revenue and less concerned about the instability's threat.

3) Stage-I Pricing Game Formulation: In reality, the geo-distributed utilities usually have no communication exchange to optimize the social performance. Instead, each utility i has its own goal to maximize its profit, which is defined as the difference between revenue and cost as follows:

$$u_i(p_i, p_{-i}) = \mathcal{R}_i(p) - \mathcal{C}_i(p) \tag{13}$$

where p_{-i} denotes the price vector of other utilities except i. This notation comes from an observation that there is a game between utilities because the profit of each utility not only depends on its energy price but also on the others'. Hence, the Stage-I utility profit maximization game, denoted by $\mathsf{UP} = (\mathcal{I}, \{p_i\}_{i\in\mathcal{I}}, \{u_i\}_{i\in\mathcal{I}})$, is defined as follows.

1) *Players:* The utilities in the set $\mathcal{I}$.
2) *Strategy:* $p_i^l \le p_i \le p_i^u, \ \forall i \in \mathcal{I}$.
3) *Payoff function:* $u_i(p_i, p_{-i}), \ \forall i \in \mathcal{I}$.

IV. Two-Stage Stackelberg Game: Equilibria and Algorithm

In this section, we first apply the backward induction method to solve the Stackelberg game. Then, we propose an iterative algorithm to reach an equilibrium of this game.

A. Backward Induction Method

1) Optimal Solutions at Stage II: We realize that the Stage-II DCs' cost minimization can be decomposed into independent problems. Henceforth, we only consider a specific time period and drop the time dependence notation for ease of presentation. In this stage, DCs cooperate with each other to minimize the total cost by determining the workload allocation λ_i and the number of active servers s_i at each DC i. It is easy to see that the DCs' cost minimization is a convex optimization problem.

First, we observe that constraint (4) must be active because otherwise the DCs provider can decrease its energy cost by reducing s_i. Hence, we have (4) is equivalent to

$$s_i(\lambda_i) = \left[\frac{1}{\mu_i}\left(\lambda_i + \tilde{D}_i^{-1}\right)\right]_0^{S_i} \tag{14}$$

where $[.]_x^y$ is the projection onto the interval $[x, y]$ and $\tilde{D}_i := D_i - d_i$. In practice, most DCs can have a sufficient number of servers to serve all requests at the same time due to the illusion of infinite capacity of DCs [18]. Therefore, we adopt $s_i(\lambda_i) = 1/\mu_i(\lambda_i + \tilde{D}_i^{-1})$ in the sequel. By substituting this $s_i(\lambda_i)$ into the objective of DC, we have an equivalent problem DC' as follows:

$$\mathsf{DC}' : \min_{\lambda} \quad \sum_{i=1}^{I} f_i(\lambda_i) \tag{15}$$

$$\text{s.t.} \quad \sum_{i=1}^{I} \lambda_i = \Lambda \tag{16}$$

$$\lambda_i \ge 0, \ \forall i \tag{17}$$

where

$$f_i(\lambda_i) := \omega d_i \lambda_i^2 + p_i\left(a_i + \frac{b_i}{\mu_i}\right)\lambda_i + p_i\left(e_b + \frac{b_i\tilde{D}_i^{-1}}{\mu_i}\right).$$

It can be seen that DC' is a strictly convex problem, which has a unique solution. Since DCs provider likes to have $\lambda_i > 0, \forall i$, in order to utilize all DCs resources, we characterize the unique solution of DC' and a necessary condition to achieve this solution with the optimal $\lambda_i^* > 0, \forall i$, as the following result.

Lemma 1: Given a price vector p, we have the unique solutions of Stage-II DC problem

$$\lambda_i^* = \frac{\nu^* - p_i A_i}{2\omega d_i} > 0, \text{ and } s_i^* = \frac{1}{\mu_i}\left(\lambda_i^* + \tilde{D}_i^{-1}\right), \quad \forall i \tag{18}$$

only if

$$\omega > \omega_{\text{th}}^1 := \left(\hat{d} \max_i\{p_i A_i\} - \sum_{i=1}^{I} p_i A_i / d_i\right) \Big/ 2\Lambda \tag{19}$$

where $\hat{d} := \sum_{i=1}^{I} 1/d_i$, $A_i := a_i + b_i/\mu_i$, and $\nu^* = 1/\hat{d}(2\omega\Lambda + \sum_{i=1}^{I} p_i A_i/d_i)$.

Since all parameters to calculate ω_{th}^1 are available to DC i, we can consider condition (19) as a guideline for a DCs provider to choose an appropriate weight factor ω to ensure that all DCs have positive request rates.

2) Nash Equilibrium at Stage I: We continue to characterize the Nash equilibrium of the Stage-I game based on the Stage-II solutions. From (13), we have

$$u_i(p_i, p_{-i}) = \big(e_i^*(p) + B_i(p_i)\big)p_i - \gamma C_i\left(\frac{e_i^*(p) + B_i(p_i)}{C_i}\right)^2 \tag{20}$$

where $e_i^*(p) = (a_i\lambda_i^* + b_i s_i^*) + e_i^b$ (with λ_i^* and s_i^* obtained from Lemma 1) and can be presented as follows:

$$e_i^*(p_i, p_{-i}) = \frac{A_i^2 p_i}{2\omega d_i}\left(\frac{1}{\hat{d}d_i} - 1\right) + \frac{A_i}{2\omega\hat{d}d_i}\sum_{j\neq i}\frac{A_j p_j}{d_j} + \frac{A_i\Lambda}{\hat{d}d_i} + \frac{b_i}{\mu_i\tilde{D}_i} + e_i^b. \tag{21}$$

In the noncooperative game, one of the most important questions is whether there exists a unique Nash equilibrium. In the case of Stage-I game, we have the following definition of a Nash equilibrium.

Definition 1: A price vector $p^e := \{p_i^e\}_{i\in\mathcal{I}}$ is said to be a Nash equilibrium if no utility can improve its profit by unilaterally deviating its price from the Nash equilibrium

$$u_i\big(p_i^e, p_{-i}^e\big) \geq u_i\big(p_i, p_{-i}^e\big),\ p_i^l \leq p_i \leq p_i^u,\ \forall i. \tag{22}$$

Theorem 1: (Existence) There exist a Nash equilibrium of the Stage-I UP game.

In this Stage-I game, given all other utilities' strategies p_{-i}, a natural strategy of utility i is the best response strategy as follows:

$$\mathcal{BR}_i(p_{-i}) = \arg\max_{p_i\in\mathcal{P}_i} u_i(p_i, p_{-i}),\ \forall i \tag{23}$$

where $\mathcal{P}_i := \big[p_i^l, p_i^u\big]$. In order to find the best response, we set $\partial u_i(p)/\partial p_i = 0$. Then, the iterative best response updates can be obtained as follows:

$$p_i^{(k+1)} = \mathcal{BR}_i\Big(p_{-i}^{(k)}\Big) = \left[\frac{1/2 - \gamma N_i/C_i}{1 - \gamma N_i/C_i}\frac{h\Big(p_{-i}^{(k)}\Big)}{(-N_i)}\right]_{\mathcal{P}_i},\ \forall i \tag{24}$$

where $[.]_{\mathcal{P}_i}$ denotes the projection onto $\mathcal{P}_i$, k represents the iterations, $N_i := A_i^2/2\omega d_i(1/\hat{d}d_i - 1) - \beta_i$, and

$$h(p_{-i}) := \frac{A_i}{2\omega\hat{d}d_i}\sum_{j\neq i}\frac{A_j p_j}{d_j} + \frac{A_i\Lambda}{\hat{d}d_i} + \frac{b_i}{\mu_i\tilde{D}_i} + e_i^b + \alpha_i, \forall i. \tag{25}$$

When all utilities play best response strategies, a Nash equilibrium p^e is a profile that satisfies $p_i^e = \mathcal{BR}_i(p_{-i}^e)$, $\forall i$, i.e., every utility's strategy is its best response to others' strategies. However, there are two issues here.

1) There is no condition for general games such that the best responses converge to a Nash equilibrium.
2) Since multiple Nash equilibria can exist in the UP game, how the best response can converge to a unique Nash equilibria.

Hence, we next examine the convergence property of the best response (24) to a unique Nash equilibrium by using the concept contraction mapping.

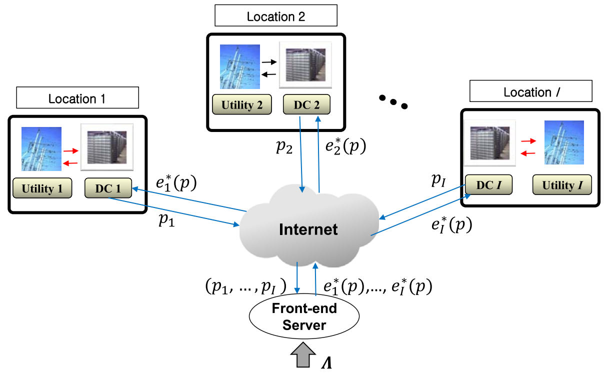

Fig. 3. Detailed operations of Algorithm 1, where red arrows represent steps 3 and 5 and blue arrows correspond to step 4.

We briefly introduce contraction mapping and its properties, all of which can be found in [30, Ch. 3]. Since many iterative algorithms have the form $x^{(k+1)} = \mathbf{T}(x^{(k)})$, $k = 0, 1, \ldots$, where $x^{(k)} \in \mathcal{X} \subset \mathcal{R}^n$, the mapping $\mathbf{T} : \mathcal{X} \mapsto \mathcal{X}$ is called a contraction if there is a scalar $0 \leq \sigma < 1$ such that

$$||\mathbf{T}(x) - \mathbf{T}(y)|| \leq \sigma||x - y||,\quad \forall x, y \in \mathcal{X} \tag{26}$$

where $||.||$ is some norm defined on $\mathcal{X}$. Furthermore, the mapping $\mathbf{T}$ is called a pseudo-contraction if $\mathbf{T}$ has a fixed point $x^* \in \mathcal{X}$ [i.e., $x^* = \mathbf{T}(x^*)$] and

$$\big|\big|\mathbf{T}(x) - x^*\big|\big| \leq \sigma\big|\big|x - x^*\big|\big|,\quad \forall x \in \mathcal{X}. \tag{27}$$

Both contraction and pseudo-contraction have the geometric convergence rate property: suppose the mapping $\mathbf{T}$ has a fixed-point, the sequence $\{x^{(k)}\}$ generated by $x^{(k+1)} = \mathbf{T}(x^{(k)})$ converges to a unique fixed point x^* geometrically satisfying

$$\Big\|x^{(k)} - x^*\Big\| \leq \sigma^k\Big\|x^{(0)} - x^*\Big\|,\quad \forall k \geq 0 \tag{28}$$

with any initial value $x^{(0)} \in \mathcal{X}$.

Based on the above properties of contraction mapping and Theorem 1, if we can show that the best response update (24) is a contraction mapping, then we can guarantee its convergence to a unique Nash equilibrium. Therefore, we establish the following sufficient condition.

Theorem 2: (Convergence and Uniqueness) If

$$\omega \geq \omega_{\text{th}}^2 := \max_i\left\{\frac{A_i\sum_{j\neq i}A_j/d_j - A_i^2\hat{d}\Big(1 - 1/\Big(d_i\hat{d}\Big)\Big)}{2\beta_i\hat{d}d_i}\right\} \tag{29}$$

then starting from any initial point, the best response updates (24) of the Stage-I UP game is a contraction mapping that converges to a unique Nash equilibrium p^e geometrically.

B. Distributed Algorithm

We first describe the detailed operations of the proposed algorithm. Next, we discuss practical implementation issues of the algorithm.

1) Proposed Algorithm's Operations and Convergence: We continue proposing a distributed algorithm, shown in Algorithm 1, which can achieve the Nash equilibrium. The detailed operations of Algorithm 1 are illustrated in Fig. 3.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Algorithm 1 DR of DC With Real-time Pricing

1: **initialize:** $k = 0$, ϵ is arbitrarily small, $p_i^{(0)} = p_i^u$, $\forall i$, and ω satisfies (29);
2: **repeat**
3: Utility i broadcasts its $p_i^{(k)}$ to all customers;
4: The front-end server collects $p^{(k)}$ from all DCs, updates $e_i^*(p)^{(k)}$ as (21) and sends it back to DC i, $\forall i$;
5: Each DC i reports its $e_i^*(p)^{(k)}$ to the local utility;
6: Utility i receives the demand responses from the local DC $e_i^*(p)^{(k)}$ and background users $B_i(p)^{(k)}$, then updates $p_i^{(k+1)} = \mathcal{BR}_i(p_{-i}^{(k)})$ as (24);
7: **until** $\left|p^{(k+1)} - p^{(k)}\right| < \epsilon$.

We assume that Algorithm 1 operates at the beginning of each pricing update period (i.e., 1 h) and the algorithm runs for many iterations (communication rounds with a parameter k) until it converges to a price setting equilibrium. Here, based on the total incoming workload, the front-end server of the DCs provider first collects all prices from its local DCs and calculates the optimal energy consumption as (21) (step 4). After that, the front-end server will feedback these energy consumption data to its local DCs, which then forwards its own information to the local utility (step 5). Each utility solves its own profit maximization problem (best response updates) to find an optimal price, then broadcasts this price to its local DCs and background customers (step 6). The process repeats until the game converges to the unique Nash equilibrium according to Theorem 2 (step 7). At this state the price setting is finalized and applied to the whole considered period.

Even though Algorithm 1 is presented in a scalable and synchronous fashion (i.e., all local utilities update and broadcast their prices at the same time), asynchronous distributed algorithm is preferred since in reality, the message-passing among front-end server, DCs and utilities usually incurs heterogeneous delays. Fortunately, with condition (29), Algorithm 1 can also work asynchronously since (29) is derived from establishing a contraction mapping with respect to a maximum norm $||.||_\infty$, which guarantees the asynchronous convergence of the mapping sequence [30, p. 431].

2) Practical Issues and Implementation Discussion: We discuss two issues here: the workload shifting assumption and the message-passing.

In terms of the former, we assume the DCs provider deploys a front-end server to distribute the incoming workload to DCs. This can be done by using various practical solutions such as incorporating the authoritative DNS servers (which is used by Akamai) or HTTP ingress proxies (which is used by Google and Yahoo) into the front-end servers. Furthermore, in reality there is only a sub-set of DCs to which a workload type can be routed to due to the availability resource constraint of each DC. This issue can be easily addressed by incorporating additional constraints into our model such as [31], and in practice we can implement it by classifying the workload types at the front-end server before routing.

In terms of the later, we assume that the two-way communication between a DC and its local utility can be enabled via communication networks of future smart grid. Regarding to the communications between DCs and its front-end server, a DC reports its utility's price by choosing one of the egress links of its ISP to send its packet through the Internet to the front-end server, and vice versa. Specifically, the total time of one iteration consists of the transmission time and computational time. While the transmission time from utilities to DCs (and vice versa) is from 1 to 10 ms over a broadband speed of 100 Mb/s, it is from 50 to 100 ms for a one-way communication between DCs and the front-end servers over a current ISPs path. The computational time depends on the processing power of the front-end server and smart meters on calculating the optimal energy (21) and maximizing the convex profit function (21), which are both low-complexity problems and can be in the time-scale of microsecond [24].

V. Trace-Based Simulations

In this section, we conduct trace-based simulations, implemented in the Python language with existing libraries including NumPy, SciPy, and Matplotlib, to validate our analysis and evaluate the performance of Algorithm 1.

A. Setups

We consider six geo-distributed DCs powered by their local utilities at the following ordered locations: 1) the Dalles, OR; 2) Council Bluffs, IA; 3) Mayes County, OK; 4) Lenoir, NC; 5) Berkeley County, SC; and 6) Douglas County, GA. These locations correspond to real Google's DCs [32]. All DCs' PUEs are set to 1.5 over time periods. The homogeneous servers have peak power of 200 W and idle power of 100 W, and the service rate of each server is chosen uniformly between 1.1 and 1.2. The migration weight ω is set to 1 unless otherwise stated. The delay SLA D_i are distributed uniformly between 100 and 300 ms and d_i is scaled by the vector [1.9, 1.0, 1.3, 2.5, 2.8, and 2.3] in which we assume that the front-end server is placed at Colorado.

We use realistic traces for the incoming workload Λ at the front-end server and the power demand of delay-tolerant batch jobs e_b at each DC. All of them are scaled with respective to service rates. We use an interactive workload trace collected from Microsoft Research (MSR) [33]. The workload can be predicted to a fairly reasonable accuracy using, e.g., regression techniques [3], [33]. Furthermore, we use Google trace for the power demand of delay-tolerant batch jobs e_b in recent study [34]. The batch job power demand and workload series spans over 30 days corresponding to a typical utility billing cycle and each point of series is a 1-h period.

Since lacking the public information of local utilities, we assume that all utilities have the capacities C_i uniformly in the range of 25 and 30 MW, which is a standard measure for a medium-size utility. While γ is set to 1 unless otherwise stated, α_i and β_i parameters are chosen uniformly in the range of $[25, 30]$ and $[0.25, 0.30]$, respectively.

We consider two baseline pricing schemes for comparison. The first baseline is based on the proposed dynamic pricing scheme of [8], which is briefly described as follows:

$$p_i(t+1) = \delta(\mathrm{PD}_i(t) - \mathrm{PS}_i(t)) + p_i(t) \tag{30}$$

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

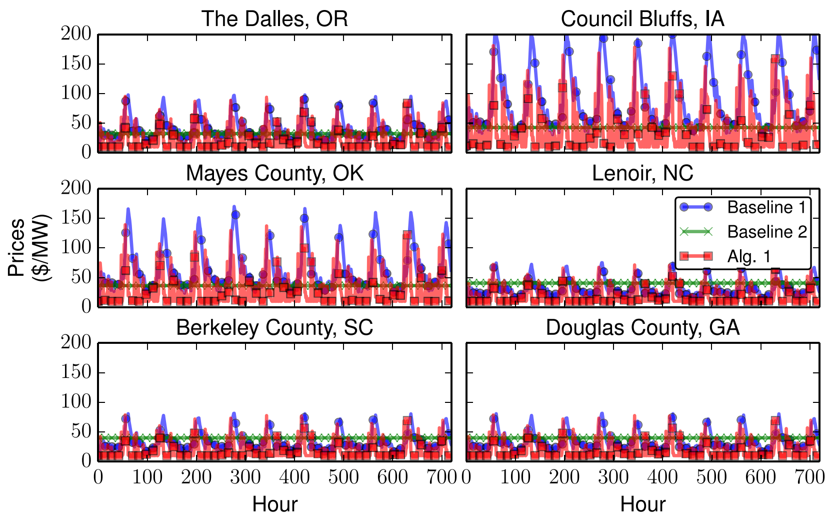

Fig. 4. Compared prices at six locations.

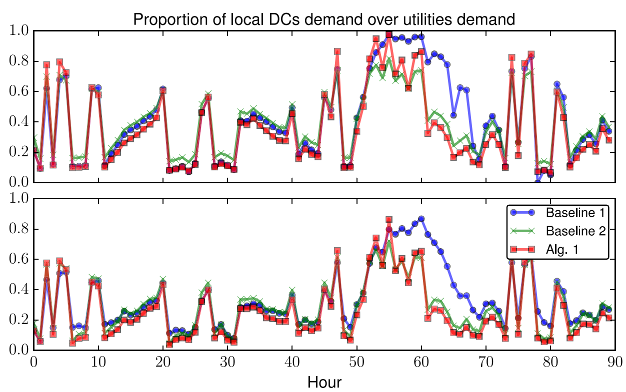

Fig. 5. Proportion of local DC demand over utility' total demand at Mayes County (top) and Lenoir (bottom).

where PD_i and PS_i are the power demand and supply of utility i. We set δ to 0.5 in all simulation scenarios. This baseline serves as a recent related benchmark.

The second baseline is based on the Google's contract with their local utilities. According to the empirical study in [32], there are six Google's DCs at six mentioned locations, where Google's DCs are inferred to have long-term contracts with their local utilities as the following fixed rates (i.e., energy charges) [32.57, 42.73, 36.41, 40.68, 44.44, and 39.97] \$/Wh, respectively. This baseline serves as an in-reality benchmark. We mainly use this baseline for the PAR comparisons since: 1) the Google long-term contract often negotiates a monthly electricity bill scheme that combines energy charges and demand charges that we do not know exactly, which can then influence the DCs' cost and utilities' profit and 2) it is not fair to compare a dynamic pricing scheme to a snapshot static pricing scheme in terms of cost and profit.

B. Results

We first provide the sample-path optimal prices of three schemes at six locations in Fig. 4. In all periods, we observe that Algorithm 1 can converge in less than ten iterations, where the stopping condition $\epsilon = 10^{-4}$. Since Baseline 1 and Algorithm 1 employ dynamic pricing mechanisms, we observe that the utilities' prices of these two schemes vary according to the workload pattern. We also observe the effect of migration cost to the optimal prices in this Fig. 4. Since the nearest DCs to the front-end server are sites 2 and 3, Fig. 4 shows that all dynamic pricing schemes set high prices at these two sites compared with the other sites. This can be explained as follows, due to the small migration cost at these sites which leads to high demand, the dynamic schemes set high prices to balance between energy cost and migration cost. Furthermore, we observe that Algorithm 1 can contribute less load to utilities than other schemes do most of the time; for example, this can be seen in Fig. 5 that shows the proportion of DCs' demand over utilities' total demand variations in three days at two sites.

TABLE II
AVERAGE OPTIMAL PRICES COMPARISONS WITH γ EFFECT

Sites	MSR			
	Baseline 1	Alg. 1 $\gamma = 1$	Alg. 1 $\gamma = 4$	Alg. 1 $\gamma = 8$
1	40.32	20.75	21.20	21.77
2	88.23	35.74	36.66	37.81
3	62.78	28.32	29.01	29.88
4	31.82	17.02	17.34	17.75
5	28.54	15.82	16.10	16.45
6	33.96	18.02	18.38	18.83

Furthermore, we also investigate the effect of γ to the pricing schemes. Table II shows that if we increase γ, then the Algorithm 1s optimal prices also increase since the higher the weight utilities' ELI cost factor is, the more conservative utilities are in terms of reliability by raising the prices. Finally, we can see that Baseline 1 always overprices Algorithm 1 in all scenarios since Baseline 1 is more aggressive than Algorithm 1 in terms of balancing the supply and demand. However, it could lead to high demand fluctuations (i.e., high PAR) as shown in the following results. We also observe that the average prices of Algorithm 1 are not affected by ω.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

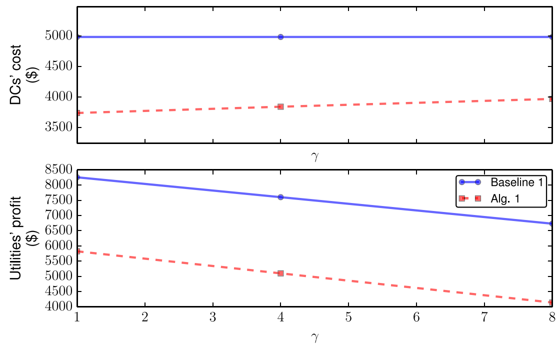

Fig. 6. Effect of γ to average DCs' cost and utilities' profit.

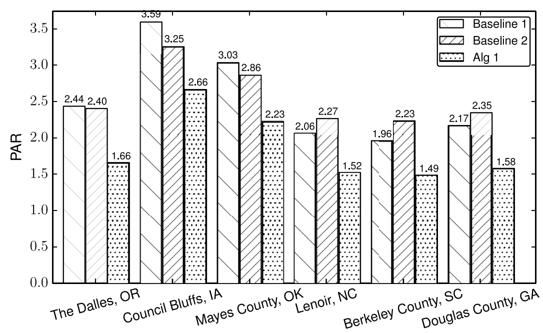

(a)

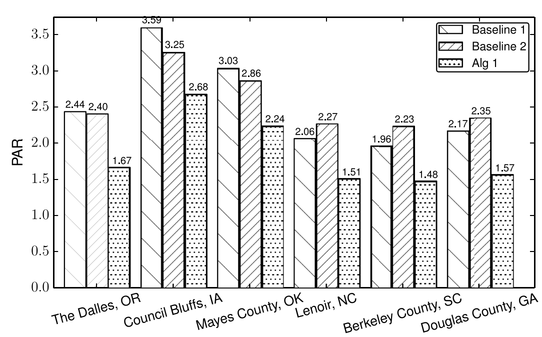

(b)

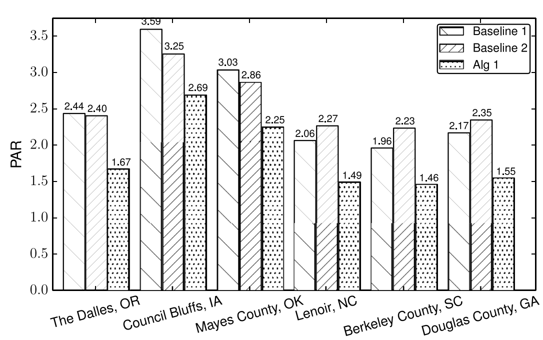

(c)

Fig. 7. PAR with respect to MSR trace at six locations with different γ. (a) $\gamma = 1$. (b) $\gamma = 4$. (c) $\gamma = 8$.

We also evaluate the effect of parameter γ to average DCs' cost and utilities' profit in Fig. 6. First, we can see that Baseline 1 with higher prices has higher DCs' cost and utilities' profit than those of Algorithm 1. In details, the share of DCs' energy cost of Algorithm 1 is 36.3%, 37.8%, and 38.7% when $\gamma = 1, 4$, and 8, respectively, whereas that of Baseline 1 (without γ impact) is 44.8%. Therefore, Algorithm 1 can give more incentives to encourage the DCs to join the DR program.

Second, we can see that when γ increases, the utilities' profit of both schemes decrease according to (20). Since the pricing scheme of Baseline 1 is independent with γ, we can see that γ has no effect to the DCs' cost of this baseline. However, we see that DCs' cost of Algorithm 1 increases when γ increases due to the corresponding increase of the optimal prices [see(20)]. With Algorithm 1, we see that small γ is favorable because it can provide low DCs' cost and high utilities profit. Furthermore, due to the background demand, we see that DCs' cost including the migration cost is lower than utilities' profit.

The final factor that we examine is the power demand PAR at each site, which is one of the most important metrics to measure the effectiveness of designs for smart grid since the fluctuation of energy consumption between peak and off-peak hours indicate power grid's reliability and robustness. PAR is calculated as $\max_t\{e_i^*(p(t)) + B_i(p_i(t))\}T / \sum_{t=1}^{T} e_i^*(p(t)) + B_i(p_i(t))$. Reducing PAR is the important goal of any DR program designs. Therefore, we extensively compare the PAR of three schemes with different γ in Fig. 7(a)–(c). The most important observation is that PARs performance of Algorithm 1 outperforms those of other schemes, either static or dynamic pricing, over time and space significantly. Specifically, considering the case $\gamma = 1$, Fig. 7(a) shows that for all sites 1–6, Algorithm 1 can achieve the lowest PAR value as expected, reducing the PAR to 32.3%, 27.0%, 28.1%, 28.0%, 25.8%, and 29.4% compared to Baseline 1, and 31.6%, 16.7%, 22.2%, 33.5%, 34.0%, and 34.0% compared to Baseline 2, respectively. We conclude that Algorithm 1 can spread out the demand not only over time but also over locations.

VI. Conclusion

We have investigated the DR of geo-distributed DCs with the help of emergence techniques of smart grid. We first characterize the challenged dependencies of this geo-distributed DCs' DR program where a utility' decisions not only depends on that of DCs, and vice versa, but also impacts on other utilities' decisions. We then formulate this DR program into a two-stage game to model these dependencies. In this game, the role of each utility is setting a price to maximize its profit, while the DCs minimize its cost by workload shifting and dynamic server allocation. We then characterize the existence and uniqueness of the Nash equilibrium of this game, and develop an iterative and distributed algorithm to reach this equilibrium. By using trace-based simulations, we validate and complement our proposal with the simulation results, which shows that our pricing schemes based on the two-stage game can flatten the energy demand of DCs over time and locations to increase the power grid's reliability and robustness.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Appendix A
Proof of Theorem 1

Since the strategy space of each utility i is a nonempty compact and convex subset of Euclidean space, it is sufficient for us to show that the continuous function $u_i(p_i, p_{-i})$ on this strategy space is a quasi-concave function, $\forall i$, such that there exists a Nash equilibrium for Stage-I game [35].

From (10) and (21), it can be seen that $e_i^*(p)$ and $B_i(p_i)$ are affine functions of p_i. Therefore, $(e_i^*(p) + B_i(p_i))^2$ is a convex function [24]. Furthermore, we have $\partial^2(B_i(p_i)p_i)/\partial p_i^2 = -\beta_i < 0$, $\forall i$ and $\partial^2(e_i^*(p)p_i)/\partial p_i^2 = A_i^2/2\omega d_i(1/\hat{d}d_i - 1) < 0$, $\forall i$, since $\hat{d}d_i > 1$, $\forall i$. Hence, both $e_i^*(p)p_i$ and $B_i(p_i)p_i$ are concave functions. Therefore, from (20) we see that $u_i(p_i, p_{-i})$ is the sum of two concave functions so that is also a concave (and hence quasi-concave as well) function.

Appendix B
Proof of Theorem 2

We first seek the condition such that the best response update (24) is a contraction mapping. Define a Cartesian product space $\mathcal{P} = \Pi_{i\in\mathcal{I}}\mathcal{P}_i$ and a vector $\boldsymbol{\mathcal{BR}}(p) := (\mathcal{BR}_i(p_{-i}))_{i\in\mathcal{I}}$. Since $\boldsymbol{\mathcal{BR}}(p)$ is continuous and differentiable on by $\mathcal{P}$, by the mean value theorem, we have

$$\|\boldsymbol{\mathcal{BR}}(p_1) - \boldsymbol{\mathcal{BR}}(p_2)\| = \left\|\frac{\partial \boldsymbol{\mathcal{BR}}(p)}{\partial p}\right\| \|p_1 - p_2\| \tag{31}$$

$\forall p_1, p_2 \in \mathcal{P}$, and p is on the segment connecting p_1 and p_2. Furthermore, the Jacobian $\partial\boldsymbol{\mathcal{BR}}(p)/\partial p$ is as follows:

$$\frac{\partial \mathcal{BR}_i(p_{-i})}{\partial p_j} = \begin{cases} 0, & \forall j = i \\ \frac{1/2-\gamma N_i/C_i}{(-N_i)(1-\gamma N_i/C_i)} \frac{A_i A_j}{2\omega\hat{d}d_i d_j}, & \forall j \neq i. \end{cases}$$

Then, by using the norm $||.||_\infty$ of the Jacobian, from (26) and (31), we see that (24) is a contraction mapping when

$$\left\|\frac{\partial \boldsymbol{\mathcal{BR}}(p)}{\partial p}\right\|_\infty = \max_i \left\{ \sum_{j\neq i} \left| \frac{1/2 - \gamma N_i/C_i}{(-N_i)(1-\gamma N_i/C_i)} \frac{A_i A_j}{2\omega\hat{d}d_i d_j} \right| \right\} < 1. \tag{32}$$

It is straightforward to see that the sufficient condition to satisfy (32) is $\max_i\{A_i/|2\omega\hat{d}d_i N_i| \sum_{j\neq i} A_j/d_j\} \leq 1$, which is equivalent to

$$\omega \geq \max_i \left\{ \frac{A_i \sum_{j\neq i} A_j/d_j - A_i^2\hat{d}\left(1 - 1/\left(d_i\hat{d}\right)\right)}{2\beta_i\hat{d}d_i} \right\}. \tag{33}$$

We have shown that with condition (33), the best response update is a contraction mapping. Furthermore, according to Theorem 1, we have the existence of a fixed-point of the mapping (24). Hence, based on the convergence property of contraction mapping, we complete the proof.

References

[1] J. Koomey, *Growth in Data Center Electricity Use 2005 to 2010*, vol. 1. Oakland, CA, USA: Anal. Press, Aug. 2011.

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

[3] L. Rao, X. Liu, L. Xie, and W. Liu, "Minimizing electricity cost: Optimization of distributed Internet data centers in a multi-electricity-market environment," in *Proc. IEEE INFOCOM*, San Diego, CA, USA, Mar. 2010, pp. 1–9.

[4] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 42, no. 1, pp. 111–123, 2014.

[5] J. Camacho, Y. Zhang, M. Chen, and D. M. Chiu, "Balance your bids before your bits: The economics of geographic load-balancing," in *Proc. 5th Int. Conf. Future Energy Syst. (e-Energy)*, Cambridge, U.K., Jun. 2014, pp. 75–85.

[6] A. H. Mahmud and S. Ren, "Online capacity provisioning for carbon-neutral data center with demand-responsive electricity prices," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 2, p. 26–37, Aug. 2013.

[7] P. Wang, L. Rao, X. Liu, and Y. Qi, "D-Pro: Dynamic data center operations with demand-responsive electricity prices in smart grid," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 1743–1754, Dec. 2012.

[8] Y. Li *et al.*, "Towards dynamic pricing-based collaborative optimizations for green data centers," in *Proc. IEEE 29th Int. Conf. Data Eng. Workshop (ICDEW)*, Brisbane, QLD, Australia, Apr. 2013, pp. 272–278.

[9] H. Wang, J. Huang, X. Lin, and H. Mohsenian-Rad, "Exploring smart grid and data center interactions for electric power load balancing," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 3, pp. 89–94, Jan. 2014.

[10] Office of the National Coordinator for Smart Grid Interoperability, "NIST framework and roadmap for smart grid interoperability standards," Nat. Inst. Stand. Technol., Gaithersburg, MD, USA, Tech. Rep. 1108, 2010.

[11] X. Fang, S. Misra, G. Xue, and D. Yang, "Smart Grid—The new and improved power grid: A survey," *IEEE Commun. Surveys Tuts.*, vol. 14, no. 4, pp. 944–980, Jan. 2012.

[12] F. Rahimi and A. Ipakchi, "Demand response as a market resource under the smart grid paradigm," *IEEE Trans. Smart Grid*, vol. 1, no. 1, pp. 82–88, Jun. 2010.

[13] A. Faruqui, "National assessment of demand response potential," Fed. Energy Reg. Comm., Washington, DC, USA, Tech. Rep., 2009. [Online]. Available: http://www.ferc.gov/legal/staff-reports/06-09-demand-response.pdf

[14] W. Saad, Z. Han, H. Poor, and T. Basar, "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications," *IEEE Signal Process. Mag.*, vol. 29, no. 5, pp. 86–105, Sep. 2012.

[15] A.-H. Mohsenian-Rad *et al.*, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 320–331, Dec. 2010.

[16] Z. Baharlouei, M. Hashemi, H. Narimani, and H. Mohsenian-Rad, "Achieving optimality and fairness in autonomous demand response: Benchmarks and billing mechanisms," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 968–975, Jun. 2013.

[17] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, San Diego, CA, USA, 2011, pp. 1–8.

[18] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Greening geographical load balancing," in *Proc. ACM SIGMETRICS Joint Int. Conf. Meas. Model. Comput. Syst.*, San Jose, CA, USA, 2011, pp. 233–244.

[19] S. Ren and Y. He, "COCA: Online distributed resource management for cost minimization and carbon neutrality in data centers," in *Proc. Int. Conf. High Perform. Comput. Network. Storage Anal. (SC)*, Denver, CO, USA, Nov. 2013, pp. 1–12.

[20] G. Ghatikar, V. Ganti, N. Matson, and M. A. Piette, "Demand response opportunities and enabling technologies for data centers: Findings from field studies," Lawrence Berkeley Nat. Lab., Berkeley, CA, USA, Tech. Rep. LBNL-5763E, 2012. [Online]. Available: http://drrc.lbl.gov/sites/all/files/LBNL-5763E.pdf

[21] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 1, p. 341–342, Jun. 2013.

[22] A. Wierman, Z. Liu, and H. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *Proc. IEEE Int. Green Comput. Conf. (IGCC)*, Dallas, TX, USA, 2014, pp. 1–10.

[23] M. Ghamkhari and H. Mohsenian-Rad, "Profit maximization and power management of green data centers supporting multiple SLAs," in *Proc. IEEE Int. Comput. Netw. Commun. (ICNC)*, San Diego, CA, USA, 2013, pp. 465–469.

[24] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, Mar. 2004.

[25] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Geographical load balancing with renewables," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 39, no. 3, pp. 62–66, Dec. 2011.

[26] L. P. Qian, Y. J. A. Zhang, J. Huang, and Y. Wu, "Demand response management via real-time electricity price control in smart grids," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 7, pp. 1268–1280, Jul. 2013.

[27] A. Mas-Colell, M. Whinston, and J. Green, *Microeconomic Theory*. New York, NY, USA: Oxford Univ. Press, 1995.

[28] W.-M. Lin and H.-C. Chin, "A current index based load balance technique for distribution systems," in *Proc. IEEE Int. Conf. Power Syst. Technol. (POWERCON)*, vol. 1. Beijing, China, 1998, pp. 223–227.

[29] M. Kashem and V. Ganapathy, "Three-phase load balancing in distribution systems using index measurement technique," *Int. J. Elect. Power Energy Syst.*, vol. 24, no. 1, pp. 31–40, Jan. 2002.

[30] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Englewood Cliffs, NJ, USA: Prentice-Hall, Jan. 1989.

[31] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's not easy being green," in *Proc. ACM SIGCOMM Conf. Appl. Technol. Archit. Protocols Comput. Commun.*, Helsinki, Finland, Aug. 2012, pp. 211–222.

[32] H. Xu and B. Li, "Reducing electricity demand charge for data centers with partial execution," in *Proc. 5th Int. Conf. Future Energy Syst. (e-Energy)*, Cambridge, U.K., Jun. 2014, pp. 51–61.

[33] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," in *Proc. IEEE INFOCOM*, Shanghai, China, Apr. 2011, pp. 1098–1106.

[34] C. Wang, B. Urgaonkar, and Q. Wang, "Data center cost optimization via workload modulation under real-world electricity pricing," in *Proc. 6th IEEE/ACM Int. Conf. Uitlity Cloud Comput. (UCC)*, Dec. 2013, pp. 1–14. [Online]. Available: http://csl.cse.psu.edu/publications/ucc13.pdf

[35] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave N-person games," *Econometrica*, vol. 33, no. 3, pp. 520–534, Jul. 1965.

Nguyen H. Tran (S'10–M'11) received the B.S. degree from the Vietnam University of Technology, Ho Chi Minh City, Vietnam, in 2005, and the Ph.D. degree from Kyung Hee University, Yongin, Korea, in 2011, both in electrical and computer engineering.

Since 2012, he has been an Assistant Professor with the Department of Computer Engineering, Kyung Hee University. His current research interests include queueing theory, optimization theory, control theory, and game theory to design, analyze, and optimize the cutting-edge applications in communication networks, such as cognitive radio, cloud-computing data center, smart grid, heterogeneous networks, and femto cell.

Dai H. Tran received the B.S. degree in information technology from the Fontys University of Applied Science, Eindhoven, The Netherlands, in 2011. He is currently pursuing the Master's degree with Kyung Hee University, Yongin, Korea.

Since 2013, he has been with the Department of Computer Engineering, Kyung Hee University. His current research interests include extending the capabilities of mobile system through mobile cloud computing.

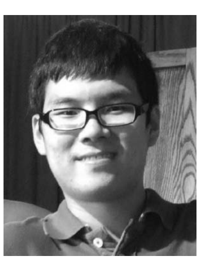

Shaolei Ren (M'13) received the B.E. degree from Tsinghua University, Beijing, China, in 2006; the M.Phil. degree from the Hong Kong University of Science and Technology, Hong Kong, in 2008; and the Ph.D. degree from the University of California–Los Angeles, Los Angeles, CA, USA, in 2012, all in electrical engineering.

Since 2012, he has been an Assistant Professor with the School of Computing and Information Sciences, Florida International University, Miami, FL, USA, and holds a joint appointment with the Department of Electrical and Computer Engineering. His current research interests include sustainable computing, data center resource management, and network economics.

Dr. Ren was a recipient of the Best Paper Award at the International Workshop on Feedback Computing (co-located with USENIX ICAC) in 2013 and the IEEE International Conference on Communications in 2009.

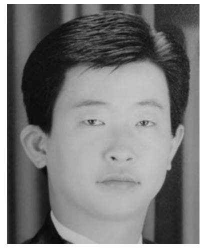

Zhu Han (S'01–M'04–SM'09–F'14) received the B.S. degree in electronic engineering from Tsinghua University, Beijing, China, in 1997, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, MD, USA, in 1999 and 2003, respectively.

From 2000 to 2002, he was a Research and Development Engineer with JDSU, Germantown, MD. From 2003 to 2006, he was a Research Associate with the University of Maryland. From 2006 to 2008, he was an Assistant Professor with Boise State University, Boise, ID, USA. He is currently an Associate Professor with the Department of Electrical and Computer Engineering, University of Houston, Houston, TX, USA. His current research interests include wireless resource allocation and management, wireless communications and networking, game theory, wireless multimedia, security, and smart grid communication.

Dr. Han was a recipient of the National Science Foundation CAREER Award in 2010 and the IEEE Fred W. Ellersick Prize Award in 2011. He has been an Associate Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS since 2010 and an IEEE Distinguished Lecturer since 2015.

Eui-Nam Huh (M'15) received the B.S. degree from Busan National University, Pusan, Korea; the M.S. degree in computer science from the University of Texas, Austin, TX, USA, in 1995; and the Ph.D. degree in computer engineering from Ohio University, Athens, OH, USA, in 2002.

From 2001 to 2002, he was the Director of the Computer Information Center and an Assistant Professor with Sahmyook University, Seoul, Korea. He was an Assistant Professor with Seoul Women's University, Seoul. He is currently a Professor with the Department of Computer Engineering, Kyung Hee University, Yongin, Korea. His current research interests include high-performance networks, sensor networks, distributed real-time systems, grids, cloud computing, and network security.

Prof. Huh was the Program Chair of the Workshop on Parallel and Distributed Real-Time Systems/International Parallel and Distributed Processing Symposium in 2003. Since 2002, he has been the Chair of the Korea Grid Standard Group. He was an Editor of the *Journal of Korean Society for Internet Information*.

Choong Seon Hong (S'95–M'97–SM'11) received the B.S. and M.S. degrees in electronic engineering from Kyung Hee University, Yongin, Korea, in 1983 and 1985, respectively, and the Ph.D. degree in computer engineering from Keio University, Minato, Japan, in 1997.

In 1988, he joined Korea Telecom (KT), Seoul, Korea, where he was a Technical Staff Member of Broadband Networks. In 1993, he joined Keio University. He was with Telecommunications Network Laboratory, KT, as a Technical Staff Senior Member and the Director of the Networking Research Team until 1999. Since 1999, he has been a Professor with the Department of Computer Engineering, Kyung Hee University. His current research interests include future Internet, ad hoc networks, network management, and network security.

Dr. Hong has served as a General Chair, the TPC Chair/Member, or an Organizing Committee Member for international conferences such as NOMS, IM, APNOMS, E2EMON, CCNC, ADSN, ICPP, DIM, WISA, BcN, TINA, SAINT, and ICOIN. He is an Associate Editor of the IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT, the *International Journal of Network Management*, and the *Journal of Communications and Networks*, and an Associate Technical Editor of the IEEE COMMUNICATIONS MAGAZINE. He is a member of ACM, IEICE, IPSJ, KIISE, KICS, KIPS, and OSIA.