Coalitional Datacenter Energy Cost Optimization in Electricity Markets

Zhe Yu
Oklahoma State University
Stillwater, OK 74078
zhe.yu@okstate.edu

Yuanxiong Guo
Oklahoma State University
Stillwater, OK 74078
richard.guo@okstate.edu

Miao Pan
University of Houston
Houston, TX 77004
mpan2@uh.edu

ABSTRACT

In this paper, we study how datacenter energy cost can be effectively reduced in the wholesale electricity market via cooperative power procurement. Intuitively, by aggregating workloads across a group of datacenters, the overall power demand uncertainty of datacenters can be reduced, resulting in less chance of being penalized when participating in the wholesale electricity market. We use cooperative game theory to model the cooperative electricity procurement process of datacenters as a cooperative game, and show the cost saving benefits of aggregation. Then, a cost allocation scheme based on the marginal contribution of each datacenter to the total expected cost is proposed to fairly distribute the aggregation benefits among the participating datacenters. Finally, numerical experiments based on real-world traces are conducted to illustrate the benefits of aggregation compared to noncooperative power procurement.

CCS CONCEPTS

• **Hardware** → **Enterprise level and data centers power issues**; • **Theory of computation** → **Solution concepts in game theory**;

KEYWORDS

Cooperative game, datacenters, wholesale electricity market, cost allocation

ACM Reference format:
Zhe Yu, Yuanxiong Guo, and Miao Pan. 2017. Coalitional Datacenter Energy Cost Optimization in Electricity Markets. In *Proceedings of e-Energy '17, Shatin, Hong Kong, May 16-19, 2017,* 12 pages.
DOI: http://dx.doi.org/10.1145/3077839.3077860

1 INTRODUCTION

With the booming of Internet-based and cloud computing services in recent years, datacenters hosting these services have become ubiquitous in every sector of our economy, and their energy consumption has been skyrocketing. According to a report [1] by the Natural Resources Defense Council, datacenters in the U.S. consumed about 91 billion kWh of electricity in 2013, representing 2% of total U.S. electricity consumption and costing U.S. businesses $13 billion in annual electricity bills, and their total electricity consumption is estimated to be 139 billion kWh in 2020. Energy cost accounts for a significant fraction (about 42%) of the datacenter operating expense [2], and this fraction is growing at an alarming rate of 12% annually [3]. Consequentially, reducing energy cost has become a critical concern for datacenter operators.

In order to reduce the growing electricity bills of datacenters, from the demand side, substantial efforts have been made, ranging from hardware such as energy-efficient servers, storage devices, and network switches, to software such as virtualization and dynamic CPU speed scaling and capacity provisioning, which have led to dramatic improvements in the energy-efficiency of datacenters. On the other hand, it is also important for datacenters to manage their energy cost from the supply side. As large consumers, datacenters typically have multiple options to procure electricity to meet their power demand[1]. For instance, a datacenter may purchase power from a retailer such as a local utility company with a pre-specified rate by signing bilateral contracts beforehand [4]. It may also operate by leveraging on-site power generators and energy storage systems [5].

Given the significant power consumption and deregulation of electricity market, another promising opportunity to reduce datacenter energy cost is emerging: datacenters can directly participate in electricity market to meet their power demand. While it is typical for consumers to buy electricity from local utility companies, some independent system operators (ISOs), such as Electric Reliability Council of Texas (ERCOT) [6] and California ISO [7], have recently developed a market that allows consumers to purchase electricity directly from power suppliers by actively participating in the electricity market. Indeed, datacenter operators like Google have been granted the authority to trade in the wholesale electricity market for the purpose of managing their own energy cost [8]. *The key advantage for datacenters to procure electricity from the wholesale market instead of a local utility company is that they can avoid the insurance premiums, service charges, and mark-up included by utilities in retail rates* [9].

[1]Without ambiguity, we use energy and power interchangeably when considering electricity demand and cost.

e-Energy '17, Shatin, Hong Kong

DOI: http://dx.doi.org/10.1145/3077839.3077860

However, a major challenge for datacenters in procuring power directly from the wholesale market is the uncertainty of market prices and their power demand. In most regions of U.S., the market for electrical power is organized into a two-settlement structure: the day-ahead forward market and the real-time balancing market. The consumers need to make a commitment or bid about their scheduled energy usage to the day-ahead market at first, and then any deviations between the scheduled and actual usage are settled in the real-time balancing market and subject to financial penalties. Since the day-ahead market is often closed several hours (e.g., 14 to 38 hours in California ISO) ahead of the actual operating time, this leaves datacenters vulnerable to high deviation penalties due to their highly uncertain workloads and associated power demand when bidding in the day-ahead market. In addition, power demand and market prices are uncertain and hard to predict as well due to the dynamic nature of the market. Therefore, it is imperative for datacenters to mitigate risks associated with these sources of uncertainty in order to maximize the cost savings in procuring power from the wholesale market directly.

In this paper, we aim to address the above challenge and optimize datacenter participation strategies in the wholesale electricity market for minimizing energy cost and facilitating energy sustainability of datacenters. In particular, consider a scenario where multiple independent datacenters operated by different owners in the same region purchase power directly in the two-settlement electricity market. Although it is risky for datacenters to participate in the market individually due to the uncertainty of their workloads and associated power demand, this paper takes an aggregation-based approach that transforms these independent datacenters from isolated entities into coordinated ones in the market. Our essential idea is to exploit the statistical diversity of workloads across different datacenters and incentivize them to bid collectively in the day-ahead market. Intuitively, by aggregating workloads from different datacenters, the uncertainty in the mixture of workloads and associated power demand can be reduced, resulting in less chance of being penalized for deviations in the real-time balancing market and higher energy cost savings.

To incentivize aggregation and fairly distribute aggregation benefits among datacenters, we propose to use cooperative game theory. Specifically, the problem can be formulated into a cooperative game with transferrable utility. In this game, the set of players is the set of datacenters who seek to cooperate in reducing electricity cost. We first prove that coalitional formation can reduce energy cost compared to individual power procurement in the wholesale electricity market. Then our cooperative game is shown to be balanced and therefore has a nonempty core. Given that the two existing cost allocation methods, the Shapley value and nucleolus, are not applicable to our game, we design an efficient cost allocation scheme that can guarantee mutual benefits for all participating datacenters such that no one has the incentive to break up from the coalition and thus locate a cost allocation in the core.

The rest of the paper is organized as follows. A brief overview of cooperative game theory is given in Section 2. In Section 3, we describe the models for datacenter power consumption and two-settlement electricity market. In Section 4, we model the datacenter aggregation process as a cooperative game and quantify the benefits of aggregation. Then, the core of the formulated game is shown to be nonempty, and an efficient scheme is proposed to find a cost allocation belonging to the core in Section 5. Simulation results based on real-world traces are presented in Section 6. Related work is reviewed in Section 7. Finally, the conclusion is given in Section 8.

2 BACKGROUND: COOPERATIVE GAME THEORY

In this section, we will briefly introduce the fundamental concepts of cooperative game theory including the definition for a cooperative game with transferable utility, the solution concept (i.e., the core) of a cooperative game, two types of cooperative games with nonempty core (i.e., the convex games and balanced games), and widely-used cost allocation methods (i.e., the Shapley value and nucleolus).

2.1 Cooperative Game with Transferable Utility

In general, a cooperative game is defined by a pair $(\mathcal{N}, c)$. The first element is the set of players $\mathcal{N} := \{1, 2, \ldots, N\}$, indexed by $i \in \mathcal{N}$. Players may form different coalitions $S \subseteq \mathcal{N}$ to obtain a collective utility. The grand coalition $\mathcal{N}$ is the set of all players. Secondly, $c : 2^{\mathcal{N}} \to \mathbb{R}$ is the cost function that assigns a real cost (i.e., the negative of the utility) to each coalition $S \subseteq \mathcal{N}$. Transferable cost implies that the total cost represented by a real number can be divided in any manner among the coalitional members [10].

2.2 Imputations and the Core

The cost function of a cooperative game is said to be subadditive if it satisfies the following condition:

$$c(\mathcal{S}) + c(\mathcal{T}) \geq c(\mathcal{S} \cup \mathcal{T}),\ \forall \mathcal{S}, \mathcal{T} \subseteq \mathcal{N},\ S \cap T = \emptyset. \tag{1}$$

For such cooperative game, it is to the mutual benefit of the players to form the grand coalition $\mathcal{N}$, since by subadditivity the amount received, $c(\mathcal{N})$, is at least as small as the total amount received by any disjoint set of coalitions they could form. Next, we focus on how to fairly split this amount among participating players.

A cost allocation for the coalition $S \subseteq \mathcal{N}$ is a vector $\pi \in \mathbb{R}^N$ whose entry π_i is the cost dispatched to each player i in the coalition S ($\pi_i = 0, i \notin S$). Further, a cost allocation π is said to be *efficient* if $\sum_{i \in \mathcal{N}} \pi_i = c(\mathcal{N})$, i.e., the total amount received by the players should be equal to $c(\mathcal{N})$. A cost allocation π is said to be *individually rational* if $\pi_i \leq c(\{i\})$, i.e., no player will be expected to receive more cost than acting individually. A cost allocation π for the grand coalition is said to be an *imputation* if it is both efficient and individually rational. In cooperative game theory [11, 12], the set of imputations for the game $(\mathcal{N}, c)$ is defined as

$$\mathcal{I} = \left\{ \pi \in \mathbb{R}^N : \sum_{i \in \mathcal{N}} \pi_i = c(\mathcal{N}),\ \pi_i \leq c(\{i\}),\ \forall i \in \mathcal{N} \right\}. \quad (2)$$

Next, we introduce the solution concept of a cooperative game. The *core* for the game $(\mathcal{N}, c)$ is defined as

$$\mathcal{C} = \left\{ \pi \in \mathbb{R}^N : \sum_{i \in \mathcal{N}} \pi_i = c(\mathcal{N}),\ \sum_{i \in S} \pi_i \leq c(S),\ \forall S \subseteq \mathcal{N} \right\}. \quad (3)$$

The core is a set of imputations such that no coalitions can obtain a cost which is less than the sum of cost assigned by forming the grand coalition. Obviously, if one can locate a cost allocation vector that lies in the core, then the grand coalition is optimal for the cooperative game.

2.3 Convex and Balanced Games

The core is always well-defined, but can be empty. However, the convex games and balanced games are two types of cooperative games which guarantee the existence of *nonempty* core [13, 14]. A cooperative game is said to be convex if the cost function satisfies the following condition:

$$c(\mathcal{S}) + c(\mathcal{T}) \geq c(\mathcal{S} \cup \mathcal{T}) + c(\mathcal{S} \cap \mathcal{T}),\ \forall \mathcal{S}, \mathcal{T} \subseteq \mathcal{N}. \quad (4)$$

This implies the cooperative game has a submodular cost function.

A map $\rho : 2^{\mathcal{N}} \to [0, 1]$ is said to be balanced if for all $i \in \mathcal{N}$,

$$\sum_{S \in 2^{\mathcal{N}}} \rho(S) \mathbf{1}\{i \in S\} = 1, \quad (5)$$

where $\mathbf{1}\{\cdot\}$ denotes the indicator function. Thus, the balanced map indicates that the sum of weights $\rho(S)$ assigned for each coalition including player i will be equal to 1. Then a cooperative game is said to be balanced if and only if for any balanced map ρ,

$$\sum_{S \in 2^{\mathcal{N}}} \rho(S) c(S) \geq c(\mathcal{N}). \quad (6)$$

2.4 Shapley Value

The Shapley value [15] as the cost allocation method is a unique mapping ψ that satisfies a series of characteristic axioms such as efficiency, symmetry, dummy and additivity. For a cooperative game $(\mathcal{N}, c)$ with transferable cost, the Shapley value $\psi_i(c)$ that distributes the cost for each player $i \in \mathcal{N}$ is defined as

$$\psi_i(c) = \sum_{S \subseteq \mathcal{N} \setminus \{i\}} \frac{|S|!(N - |S| - 1)!}{N!} \left[c(S \cup \{i\}) - c(S) \right]. \quad (7)$$

We observe that in (7), the marginal contribution of each player is represented as $c(S \cup \{i\}) - c(S)$ and the coefficient ahead of the marginal distribution is the probability that the player i randomly joins the coalition S. Thus, the Shapley value can be interpreted as the expected marginal contribution of player i in the grand coalition $\mathcal{N}$ when it joins the coalition S in a random order. It is guaranteed that the Shapley value lies in the core if the game is convex [13].

2.5 Nucleolus

The nucleolus [16] is another common cost allocation method. It uniquely exists in a cooperative game and satisfies the efficiency, individually rational, symmetry and dummy properties [10]. Different from axiomatically designing the cost allocation scheme to ensure fairness as in the Shapley value, the nucleolus aims at minimizing the dissatisfaction of the players. The dissatisfaction of a coalition S given an imputation π is measured by the excess. The definition of excess is given by

$$e(\pi, S) = \sum_{i \in S} \pi_i - c(S). \quad (8)$$

Since the core is defined as the set of imputations such that $\sum_{i \in S} \pi_i \leq c(S)$ for all coalitions $S \subseteq \mathcal{N}$, it follows that an imputation π is in the core if and only if all its excesses are negative or zero [17]. In order to find the nucleolus, we first need to locate an imputation that minimizes the maximum of the excesses $e(\pi, S)$ over all coalitions S by solving a linear program. After this is done, one may have to solve a second linear programming problem to minimize the next largest excess, and so on. Therefore, in the worst-case, $\mathcal{O}(2^N)$ linear programs need to be solved, which is computationally expensive [18].

3 SYSTEM MODEL

In this section, we start by introducing the datacenter power consumption model and characterize the uncertainty of power demand for each datacenter. Then, the two-settlement electricity market is described and the expected electricity cost for each datacenter when participating in the market individually is derived.

Consider a set $\mathcal{N} := \{1, 2, \ldots, N\}$ of independent datacenters participating in the same electricity market for power procurement. For instance, SoftLayer, Google, Microsoft and Amazon have all built large-scale datacenters in San Francisco Bay Area, and these datacenters are served by California ISO for wholesale power procurement. We explore the scenario in which individual datacenters form a coalition to collectively bid their aggregate power demand in the electricity market as a single entity for cost savings. Without loss of generality, we restrict our analysis to a single operating hour.

3.1 Datacenter Power Consumption Model

Assume each datacenter $i \in \mathcal{N}$ has M_i homogenous servers whose idle and peak power consumption are P_i^{idle} and P_i^{peak}, respectively[2]. Users submit their requests (e.g., search queries) to datacenters, and datacenters process these requests to satisfy the quality-of-service (QoS) requirement as

[2]Note that a datacenter with heterogenous servers can be also viewed as several datacenters, each having homogeneous servers. Therefore, we focus on the homogenous case in this paper.

indicated by the service-level agreement (SLA). When datacenter i keeps m_i active servers to process the arriving user requests, its IT power consumption can be estimated as [19]

$$P_i = m_i \left[P_i^{\text{idle}} + u_i (P_i^{\text{peak}} - P_i^{\text{idle}}) \right], \tag{9}$$

where u_i is the average CPU utilization level across all servers at datacenter i.

We adopt a M/GI/1 Processor Sharing (PS) queue to model the service process at each server [20]. The workload arrival rate at each datacenter i, measured in terms of the average number of arriving user requests per unit time, is assumed to be λ_i. Let μ_i denote the service rate at which user requests are processed by a server at datacenter i. Then the average CPU utilization level in datacenter i is calculated as $u_i = \lambda_i/(m_i\mu_i)$. Therefore, the power consumption model (9) can be rewritten as

$$P_i = m_i P_i^{\text{idle}} + \frac{\lambda_i}{\mu_i} \left(P_i^{\text{peak}} - P_i^{\text{idle}} \right). \tag{10}$$

Since each user request has a QoS requirement, datacenters need to turn on enough servers to meet that requirement. Here we use the average response time as the QoS metric. Based on the M/GI/1/PS queuing model, the average response time of user requests given m_i active servers in datacenter i is represented as

$$T_i = \frac{1}{\mu_i - \lambda_i/m_i}. \tag{11}$$

Let $T_i^{\max}$ denote the maximum average response time of user requests that can be tolerated at datacenter i. Then to ensure that $T_i \le T_i^{\max}$, we obtain the following feasible range for the number of active servers at datacenter i:

$$\frac{\lambda_i}{\mu_i - 1/T_i^{\max}} \le m_i \le M_i. \tag{12}$$

Here, we relax the constraint that requires m_i to be integer given the fact that datacenters usually contain thousands of servers. It is assumed that each datacenters turn on the minimal number of active servers without violating their QoS requirement using the dynamic capacity provisioning technique [21, 22]. Therefore the IT power consumption of each datacenter i is

$$P_i = \frac{\lambda_i}{\mu_i - 1/T_i^{\max}} P_i^{\text{idle}} + \frac{\lambda_i}{\mu_i} \left(P_i^{\text{peak}} - P_i^{\text{idle}} \right). \tag{13}$$

In order to incorporate the non-IT (e.g. cooling, lighting) power consumption of datacenters, we denote the average power usage effectiveness (PUE) as γ_i, which is defined as the ratio of the total power consumption to the IT power consumption at datacenter i. It follows that the total power demand E_i of datacenter i is given by

$$E_i = \theta_i \lambda_i, \tag{14}$$

where θ_i is a constant defined as

$$\theta_i := \gamma_i \left(\frac{P_i^{\text{idle}}}{\mu_i - 1/T_i^{\max}} + \frac{P_i^{\text{peak}} - P_i^{\text{idle}}}{\mu_i} \right). \tag{15}$$

When datacenter i bids in the day-ahead market one day ahead, the user request arrivals for the next day are uncertain, and thus the average workload arrival rate λ_i can be modeled as a random variable whose probability distribution can be empirically estimated from historical data. It follows that the actual datacenter power demand $E_i(\lambda_i)$ as a linear function of the average workload arrival rate λ_i is also a random variable.

3.2 Two-Settlement Electricity Market

Consider a wholesale electricity market managed by an ISO with a two-settlement structure in the region through which the datacenters procure power. It consists of a day-ahead forward market and a real-time balancing market. In the day-ahead forward market, participants bid and schedule power transactions for each hour of the following day before the gate closure. After that, the ISO clears the market and calculates the day-ahead market clearing price for each hour as the intersection between the aggregate supply and demand curves. For instance, for California ISO, the day-ahead forward market closes for bids and schedules by 10 AM and clears by 1 PM on the day prior to the operating day. The schedules cleared in the day-ahead market are financially binding. Any deviations between the day-ahead committed schedule and actual power consumption/generation will be settled in the real-time balancing market during the operating day. If the actual consumption is more than or production is less than the committed schedule, the energy shortfall will be purchased in the balancing market at the negative imbalance price, which is usually higher than the day-ahead price. If the actual consumption is less than or production is more than the committed schedule, the energy surplus will be sold at the positive imbalance price, which is usually lower than the day-ahead price. Therefore, power deviations from day-ahead commitments normally result in penalties for participants.

Specifically, for the considered wholesale market, let $p^d \in \mathbb{R}^+$ be the market clearing price in the day-ahead forward market, $p^- \in \mathbb{R}^+$ be the negative imbalance price for energy shortfall, and $p^+ \in \mathbb{R}^+$ be the positive imbalance price for energy surplus. The datacenters are assumed to be price-taking because their energy consumption are often too small to influence the market. The market prices (p^d, p^-, p^+) are not known to the datacenters at the time of bidding in the day-ahead market and therefore modeled as random variables with known expected values denoted by μ_p^d, μ_p^-, and μ_p^+, respectively, which can be estimated empirically from historical market data. As explained before, without loss of generality, we assume $\mu_p^+ \le \mu_p^d \le \mu_p^-$. Moreover, the market prices (p^d, p^-, p^+) are assumed to be statistically independent of the workload arrival rates $(\lambda_i, \forall i)$.

Suppose that each datacenter $i \in \mathcal{N}$ bids a power procurement amount Q_i in the day-ahead market. With the above models and assumptions, it follows that the expected cost of datacenter i from participating in the market individually can be calculated as

$$\Phi_i = \mu_p^d Q_i + \mu_p^- \mathbb{E}[(E_i - Q_i)^+] - \mu_p^+ \mathbb{E}[(Q_i - E_i)^+], \quad (16)$$

where $(x)^+ := \max(x, 0)$, $\mu_p^d Q_i$ denotes the day-ahead trading cost, $\mu_p^- \mathbb{E}[(E_i - Q_i)^+]$ denotes the shortfall penalty, and $\mu_p^+ \mathbb{E}[(Q_i - E_i)^+]$ denotes the surplus profit.

4 COALITIONAL DATACENTER BIDDING

In this section, we start by introducing the datacenter aggregation model where multiple datacenters can form a coalition to bid in the day-ahead market collectively. Then, it can be verified that by bidding power demand aggregately in the day-ahead market, the total electricity bill can be effectively reduced based on the fact that datacenter aggregation can reduce the uncertainty of the total workload arrivals.

4.1 Datacenter Aggregation as a Cooperative Game

Datacenters can form a coalition and bid collectively in the day-ahead market. Any coalition $S \subseteq \mathcal{N}$ represents an agreement among the datacenters in S to act as a single entity in the market. The aggregated datacenter power demand of a coalition $S \subseteq \mathcal{N}$ is specified by

$$E_S = \sum_{i \in S} E_i. \quad (17)$$

Further, we denote the cumulative distribution function (CDF) of E_S as

$$F_S(e) = \Pr(E_S \le e). \quad (18)$$

The corresponding quantile function is given by

$$F_S^{-1}(\varepsilon) = \inf \{e \in [E_S^{\min}, E_S^{\max}] : \varepsilon \le F_S(e)\}, \quad (19)$$

where $E_S^{\min}$ and $E_S^{\max}$ are the lower and upper bounds of the aggregated power demand, which depends on the minimum and maximum workload arrival rates.

Next, we use cooperative game theory [23] to model this cooperation process as a cooperative game $(\mathcal{N}, c)$ with transferable cost since it is under a multi-agent scenario where each datacenter tends to minimize its own net cost. In our model, the set of datacenters $\mathcal{N}$ is the set of players in the cooperative game. Moreover, we assume each datacenter always seeks to minimize its own electricity cost, and then the cost function $c(S)$ associated with every coalition $S \subseteq \mathcal{N}$ is represented as its minimum expected energy cost calculated as

$$\Phi_S = \mu_p^d Q_S + \mu_p^- \mathbb{E}[(E_S - Q_S)^+] - \mu_p^+ \mathbb{E}[(Q_S - E_S)^+], \quad (20)$$

$$c(S) = \min_{Q_S \ge 0} \Phi_S, \quad (21)$$

where Q_S is the bid amount of any coalition S in the day-ahead market. We assume the market prices for the coalitional bid is the same as that of individual bids. This assumption is acceptable since the datacenters are assumed to be relatively small [24] compared to all other consumers participating in the electricity market so that their operations have little impact on the cleared prices of the day-head market or real-time market. Solving (21) as a news-vendor problem [25, 26], the optimal day-ahead bid and expected cost are given in the following theorem:

THEOREM 4.1. *The optimal day-ahead bid of any coalition S is given by*

$$Q_S^* = F_S^{-1}(\varepsilon^*), \quad \text{where} \quad \varepsilon^* = \frac{\mu_p^- - \mu_p^d}{\mu_p^- - \mu_p^+}. \quad (22)$$

The optimal expected cost is given by

$$c(S) = \mu_p^+ \int_0^{\varepsilon^*} F_S^{-1}(\theta)\, \mathrm{d}\theta + \mu_p^- \int_{\varepsilon^*}^1 F_S^{-1}(\theta)\, \mathrm{d}\theta. \quad (23)$$

PROOF. The proof is referred to Appendix A. □

4.2 The Benefits of Aggregation

Intuitively, no group of datacenters can do worse by joining a coalition than by acting noncooperatively since aggregation can reduce uncertainty. We will prove this by the following theorem:

THEOREM 4.2. *Given an arbitrary coalition $S \subseteq \mathcal{N}$, let $\{Q_1, Q_2, \ldots, Q_{|S|}\}$ be a set of $|S|$ individual day-ahead bids. For $Q_S = \sum_{i \in S} Q_i$ we have:*

$$\Phi_S(Q_S) \le \sum_{i \in S} \Phi_i(Q_i). \quad (24)$$

PROOF. The proof is referred to Appendix B. □

It is straightforward to see that the expected cost by participating in the market collectively is less than the sum of that by participating in the market individually. That is, the datacenters save the expected cost of $\sum_{i \in S} \Phi_i(Q_i) - \Phi_S(Q_S)$ collectively via aggregation. Further, we establish some properties of the cost function associated with every coalition.

LEMMA 4.3. *The optimal expected cost $c(S)$ of any coalition S has following properties:*

1. *Positive homogeneity: For any scalar $\beta \ge 0$, $c(\beta S) = \beta c(S)$.*
2. *Subadditivity: For any two disjoint coalitions S_1 and S_2, if coalition $S_1 \cup S_2$ forms, then $c(S_1 \cup S_2) \le c(S_1) + c(S_2)$.*

PROOF. The proof is referred to Appendix C. □

From positive homogeneity, we observe that when the aggregated power demand is scaled, the corresponding value of the optimal expected cost will also be scaled in the same proportion. From subadditivity, we observe that for rational datacenters who always try to minimize their cost, they will form a large-size coalition to benefit more from the aggregation. It is straightforward to see in our game that all the datacenters will form the grand coalition $\mathcal{N}$ in order to minimize their total expected cost.

5 COST ALLOCATION MECHANISM

In the section, we focus on how to find a cost allocation vector π as defined in Section 2.2 to split the total expected cost to each datacenter in the grand coalition. First, we show that the core of our cooperative game exists and is nonempty by proving it is a balanced game. Next, we verify that our game is nonconvex, and hence the Shapley value is not applicable to locate the core of our game. Last, we propose a cost allocation scheme based on the marginal contribution of each datacenter to the total cost in the grand coalition.

5.1 Existence of the Nonempty Core

As shown in Section 2, both the convexity and balancedness can guarantee the core of a cooperative game to be nonempty. Since the convexity of a cooperative game is a stronger condition compared to the balancedness, we prove the existence of the core in terms of balancedness by the following theorem:

THEOREM 5.1. *The cooperative game $(\mathcal{N}, c)$ for datacenter aggregation is balanced and has a nonempty core.*

PROOF. Given an arbitrary balanced map $\rho : 2^{\mathcal{N}} \to [0,1]$, by following the concept of the balanced game, we have

$$\begin{aligned}
\sum_{S \in 2^{\mathcal{N}}} \rho(S)c(S) &= \sum_{S \in 2^{\mathcal{N}}} c(\rho(S)S) && (25)\\
&\geq c\left(\sum_{S \in 2^{\mathcal{N}}} \rho(S)S\right) && (26)\\
&= c\left(\sum_{S \in 2^{\mathcal{N}}} \rho(S)\left(\bigcup_{i \in \mathcal{N}} \mathbf{1}\{i \in S\}i\right)\right)\\
&= c\left(\bigcup_{i \in \mathcal{N}}\left(\sum_{S \in 2^{\mathcal{N}}} \rho(S)\mathbf{1}\{i \in S\}\right)i\right) && (27)\\
&= c\left(\bigcup_{i \in \mathcal{N}} i\right) = c(\mathcal{N}),
\end{aligned}$$

where (25) is because of the positive homogeneity of $c(S)$, (26) is because of the subadditivity of $c(S)$, and (27) is derived by the definition of balanced map ρ. Therefore, the cooperative game $(\mathcal{N}, c)$ is balanced and has a nonempty core. □

5.2 Marginal Cost Allocation

Two prominent cost allocation schemes are described in Section 2. However, both of them are not applicable to solve our cooperative game. The Shapley value can be guaranteed to lie in the core if the cooperative game is convex. However, as shown through a counterexample in Appendix D, our game is not convex. Therefore, the Shapley value does not necessarily belong to the core and hence is not applicable to allocate cost in our game. The nucleolus uniquely exists and can be used as a cost allocation scheme in our game. However, as mentioned before, in the worst-case scenario, $\mathcal{O}(2^N)$ linear programs need to be solved in order to get the cost allocation vector, which is computationally expensive.

Here, we propose a cost allocation scheme based on the marginal contribution of each datacenter to the total expected cost when participating in the grand coalition and prove the resulting cost allocation vector is in the core. We define an aggregation level vector $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_N]^T$, where each element $0 \leq \alpha_i \leq 1$ represents the fraction of datacenter power demand E_i that participates in the aggregative power procurement. Thus, the weighted power demand of the aggregation with the aggregation level vector $\boldsymbol{\alpha}$ is denoted as

$$E_{\boldsymbol{\alpha},\mathcal{N}} = \sum_{i=1}^{N} \alpha_i E_i, \tag{28}$$

whose quantile function is represented by $F^{-1}_{\boldsymbol{\alpha},\mathcal{N}}(\varepsilon)$ and defined similar to (19). Then by applying Theorem 4.1, we can obtain the optimal expected cost of the weighted power demand as

$$c_{\boldsymbol{\alpha}}(\mathcal{N}) = \mu_p^+ \int_0^{\varepsilon^*} F^{-1}_{\boldsymbol{\alpha},\mathcal{N}}(\theta)\,\mathrm{d}\theta + \mu_p^- \int_{\varepsilon^*}^1 F^{-1}_{\boldsymbol{\alpha},\mathcal{N}}(\theta)\,\mathrm{d}\theta. \tag{29}$$

The positive homogeneity and subadditivity proved in Lemma 4.3 can be easily extended to the case where we consider the weighted optimal expected cost $c_{\boldsymbol{\alpha}}(\mathcal{N})$. Further, we show another property as follows:

LEMMA 5.2. *The weighted optimal expected cost $c_{\boldsymbol{\alpha}}(\mathcal{N})$ of any coalition S is nonincreasing over $\boldsymbol{\alpha}$, i.e., for any two aggregation level vectors, if $\boldsymbol{\alpha} \succeq \boldsymbol{\alpha}'$*[3]*, then $c_{\boldsymbol{\alpha}}(\mathcal{N}) \leq c_{\boldsymbol{\alpha}'}(\mathcal{N})$.*

PROOF. Given two aggregation level vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ where $\boldsymbol{\alpha} \succeq \boldsymbol{\alpha}'$, then for any element in the vector $\boldsymbol{\alpha} - \boldsymbol{\alpha}'$, we have $0 \leq \alpha_i - \alpha'_i \leq 1, \forall i \in \mathcal{N}$. Using the subadditivity property, we have

$$c_{\boldsymbol{\alpha}}(\mathcal{N}) \leq c_{\boldsymbol{\alpha}'}(\mathcal{N}) + c_{\boldsymbol{\alpha}-\boldsymbol{\alpha}'}(\mathcal{N}), \tag{30}$$

which indicates the nonincreasing property. □

According to Lemma 5.2, the optimal expected cost will be achieved when $\boldsymbol{\alpha} = \mathbf{1}$, where $\mathbf{1} \in \mathbb{R}^{N \times 1}$ is an all-one vector. Then it follows that $c_{\boldsymbol{\alpha}}(\mathcal{N})|_{\boldsymbol{\alpha}=\mathbf{1}} = c(\mathcal{N})$.

To distribute the total expected cost $c(\mathcal{N})$ among the datacenters in the grand coalition, we compute the expected cost for each datacenter i as

$$\pi_i = \frac{\partial c_{\boldsymbol{\alpha}}(\mathcal{N})}{\partial \alpha_i}\Big|_{\boldsymbol{\alpha}=\mathbf{1}}, \ \forall i \in \mathcal{N}. \tag{31}$$

Indeed, π_i can be decomposed as the multiplication of two terms:

$$\pi_i = \frac{\partial c_{\boldsymbol{\alpha}}(\mathcal{N})}{\partial E_{\boldsymbol{\alpha},\mathcal{N}}}\Big|_{\boldsymbol{\alpha}=\mathbf{1}} \times \frac{\partial E_{\boldsymbol{\alpha},\mathcal{N}}}{\partial \alpha_i}\Big|_{\boldsymbol{\alpha}=\mathbf{1}}, \ \forall i \in \mathcal{N}, \tag{32}$$

where the second term is exactly the power demand E_i of each datacenter. On the other hand, the first term is the partial derivative of the weighted optimal expected cost with respect to the weighted power demand and then evaluating

[3]The operator $\succeq$ represents component-wise vector comparison.

Table 1: Simulation Parameters

	M_i	μ_i (requests/s)	$T_i^{\max}$ (ms)
Datacenter 1	10000	200	100
Datacenter 2	12500	250	80
Datacenter 3	15000	300	60
Datacenter 4	17500	350	40

at the full aggregation level, i.e., $\boldsymbol{\alpha} = \mathbf{1}$, which can be considered as the marginal cost assigned to each datacenter. Therefore, the multiplication of the marginal cost and power demand gives the distributed cost to each datacenter. Further, we prove that the cost allocation vector $\pi = [\pi_1, \ldots, \pi_N]^T$ given in (31) lies in the core as shown in following theorem:

THEOREM 5.3. *The resulting cost allocation vector of the proposed cost allocation scheme is fair and lies in the core of our cooperative game.*

PROOF. The proof is referred to Appendix E. □

The most significant advantage of exploiting this method is its low computational complexity. Compared to using the nucleolus, we only need to calculate $\mathcal{O}(N)$ equations.

6 NUMERICAL EXPERIMENTS

In this section, we first introduce our simulation setup and then conduct trace-driven simulations to show the benefits of datacenter aggregation in purchasing power in the wholesale electricity market and the effectiveness of our proposed cost allocation scheme.

6.1 Simulation Setup

A set of four independent datacenters $\mathcal{N} = \{1, 2, 3, 4\}$ is considered in our simulations. The total number of servers for each datacenter is 10,000, 12,500, 15,000 and 17,500, respectively. Assume the idle power and peak power of each server is 150 W and 250 W, respectively. Besides, the average PUEs of all the datacenters are set to 1.5. The average service rate of a server in each datacenter is set to be 200, 250, 300 and 350 requests per second, respectively. The maximum average response time for each datacenter is set to be 100, 80, 60 and 40 ms, respectively. The above simulation parameters are summarized in Table 1.

The real-world dataset we use to simulate the workloads is from the Google cluster trace [27]. The selected dataset includes workload information over 29 days (i.e., 696 hours) during May 2011 for a cluster of 12,500 severs. We repeat the original data and extend it to 1008-hour workloads (i.e., 42 days). Then, we randomly choose 4 different 720-hour (i.e., 30 days) portions from the extended dataset as our datacenter workloads. Figure 1a shows the CDFs of the normalized datacenter workload arrival rate for four datacenters at hour 22. Then we can estimate the power demand of each datacenter according to (14). The CDFs of the power demand for four datacenters are depicted in Figure 1b.

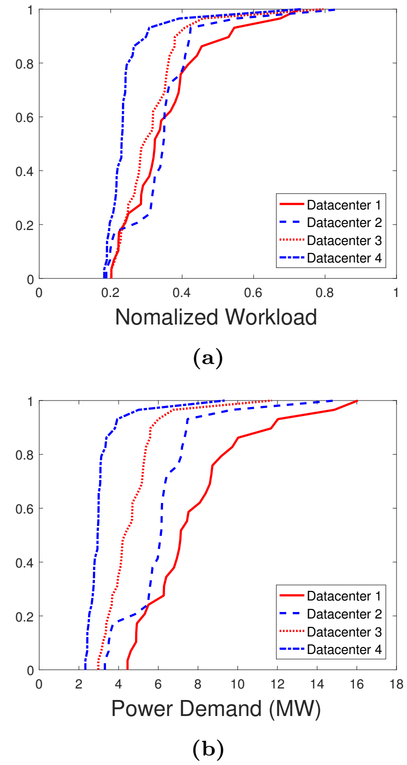

Figure 1: CDFs of the normalized datacenter workload arrival rate and power demand at hour 22.

In our simulations, datacenters can purchase power either individually or cooperatively by forming the grand coalition. Moreover, we assume datacenters bid their power demand in the day-ahead market for each hour in the following operating day. By default, the expected day-ahead price μ_p^d is set to be 50 cents per kWh, the expected negative imbalance price μ_p^- is set to be 2 dollars per kWh, and the expected positive imbalance price μ_p^+ is set to be 20 cents per kWh in the simulations.

Last, all our simulations are conducted on a desktop computer with an Intel Core i7-4790 3.60GHz CPU and 8GB RAM using MATLAB R2016a.

6.2 Experimental Results

In this section, we simulate and analyze how datacenters can benefit from forming the grand coalition to save their electricity cost when purchasing power in the wholesale electricity market. Here, we consider the case where each datacenter bids its power demand individually by minimizing its expected energy cost as the baseline scenario for comparison.

Benefits of Aggregation. We first observe the benefits of coalitional bidding in the wholesale market. Based on Theorem 4.1, we can calculate the optimal day-ahead bid Q_S^* of any coalition S. Figure 2 shows the resulting optimal day-ahead bidding level of our proposed method and the

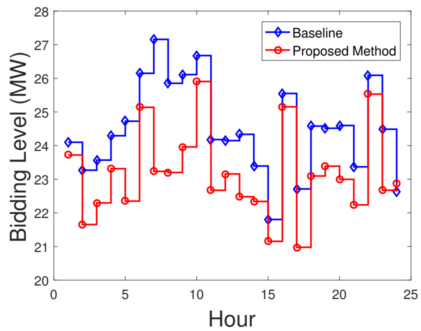

Figure 2: Day-ahead bidding level comparison over 24 hours.

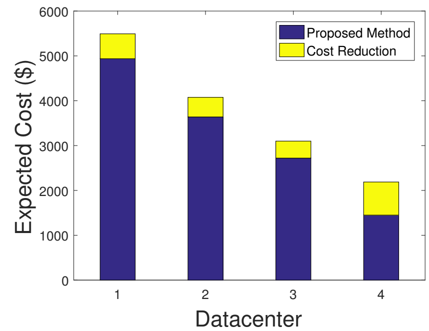

Figure 4: Cost allocation of each datacenter at hour 22 under the current setting.

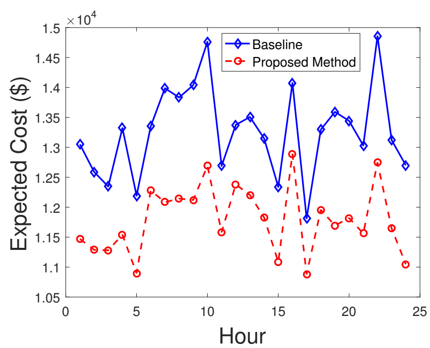

Figure 3: Total expected cost comparison over 24 hours.

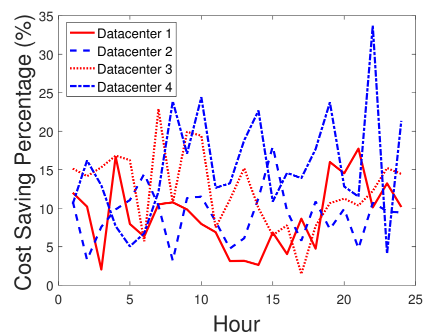

Figure 5: Individual cost saving percentage of each datacenter after coalitional day-ahead bidding over 24 hours.

sum of optimal individual bidding level in the baseline over 24 hours. Figure 3 shows the energy cost comparison of our proposed approach and the baseline. The result of the baseline scenario is obtained by adding up the optimal expected electricity cost of each datacenter when they bid in the day-ahead market individually, while the result of the proposed method is obtained by letting datacenters form the grand coalition to bid in the day-ahead market cooperatively. It is shown in Figure 3 that the total electricity cost is effectively reduced by cooperative day-ahead bidding, which validates the subadditivity property of our cooperative game given in Lemma 4.3. The average hourly cost saving is around 11.03% under the current setting.

Cost Allocation. Next we focus on how to fairly distribute the total energy cost after coalitional bidding among each participating datacenter using our proposed cost allocation method. We split the total expected cost based on the marginal contribution of each datacenter in the grand coalition by applying the proposed cost allocation scheme in Section 5.2. Figure 4 presents the cost allocation to each datacenter at hour 22. The height of each bar (yellow bar plus blue bar) denotes the expected cost of each datacenter when it bids individually in the day-ahead market at hour 22. The height of yellow bar shows the reduced cost of each datacenter after coalitional day-ahead bidding. The cost saving percentage of each datacenter over 24 hours in a day is given in Figure 5. It can be observed that our proposed allocation method can always ensure positive cost reductions for each datacenter and the cost saving amount of each datacenter is different, depending on its contribution to the aggregation benefits.

Table 2 presents the noncooperative and coalitional electricity cost of each coalition at hour 22. The last column gives the corresponding excesses $e(\pi, S)$ defined in (8). From row 1 to row 14, the calculated excesses are all negative which satisfies the condition of subgroup rationality, i.e., $\sum_{i \in S} \pi_i \leq c(S)$. The last row indicates that our cost allocation is efficient since $\sum_{i \in \mathcal{N}} \pi_i = c(\mathcal{N})$. It verifies that our proposed cost allocation lies in the core of the cooperative game since both subgroup rationality and efficiency conditions are satisfied.

Impact of Price Penalty Ratio. Now we present how market prices affect the cost saving and the day-ahead bid of each datacenter when they form the grand coalition. According to Theorem 4.1, the optimal day-ahead bid depends on the quantile function where the percentile ε^* is decided by

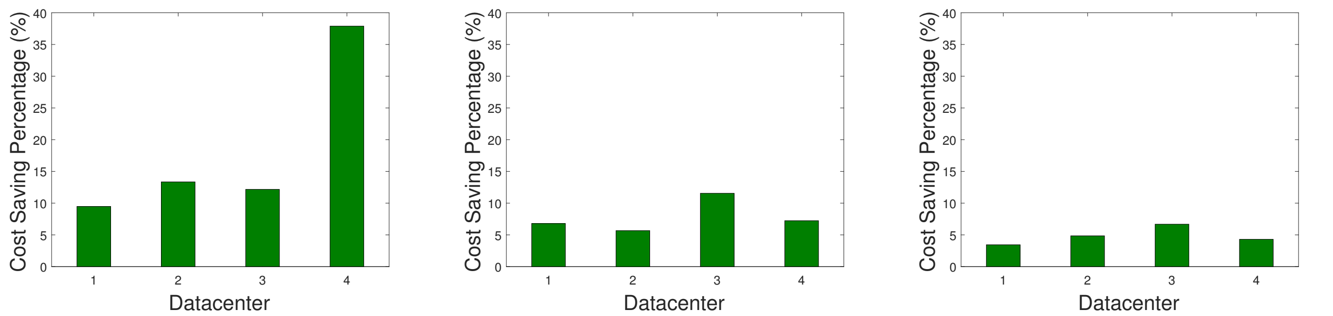

(a) $\omega = 0.25$ (b) $\omega = 0.50$ (c) $\omega = 0.75$

Figure 6: Cost saving percentage of each datacenter at hour 22 when the price penalty ratio ω is 0.25, 0.50 and 0.75, respectively.

Table 2: Cost comparison for all coalitions of four datacenters at hour 22

	S	$c(S)$	$\sum_{i\in S}\pi_i$	$\sum_{i\in S}\pi_i - c(S)$
1	$\{1\}$	5492.2	4937.7	-554.5
2	$\{2\}$	4075.6	3639.0	-436.6
3	$\{3\}$	3099.8	2720.8	-379.0
4	$\{4\}$	2187.8	1450.4	-737.4
5	$\{1,2\}$	8809.2	8576.7	-232.5
6	$\{1,3\}$	8132.5	7658.5	-474.0
7	$\{1,4\}$	7018.5	6388.1	-630.4
8	$\{2,3\}$	6748.7	6359.8	-388.9
9	$\{2,4\}$	5941.8	5089.4	-852.4
10	$\{3,4\}$	4966.8	4171.2	-795.6
11	$\{1,2,3\}$	11498.0	11297.5	-200.5
12	$\{1,2,4\}$	10300.0	10027.1	-272.9
13	$\{1,3,4\}$	9574.0	9108.9	-465.1
14	$\{2,3,4\}$	8454.0	7810.2	-643.8
15	$\{1,2,3,4\}$	12747.9	12747.9	0

Table 3: The percentage of the average cost saving of each datacenter under different price penalty ratios

	$\omega = 0.25$	$\omega = 0.50$	$\omega = 0.75$
Datacenter 1	10.11%	6.23%	3.08%
Datacenter 2	10.17%	7.05%	3.78%
Datacenter 3	14.67%	8.15%	4.98%
Datacenter 4	17.12%	8.23%	4.90%

expected electricity prices μ_p^d, μ_p^- and μ_p^+. For the simplicity of presentation, we set the expected positive imbalance price μ_p^+ to be 0 but it does not affect our analytical analysis beforehand. It follows that the percentile ε^* will reduce to $1-\mu_p^d/\mu_p^-$. Under this simplification, we introduce the price penalty ratio ω defined as μ_p^d/μ_p^- [25]. In order to obtain different price penalty ratios, we fix the expected day-ahead price μ_p^d as a constant and adjust the expected negative imbalance price μ_p^- to different values.

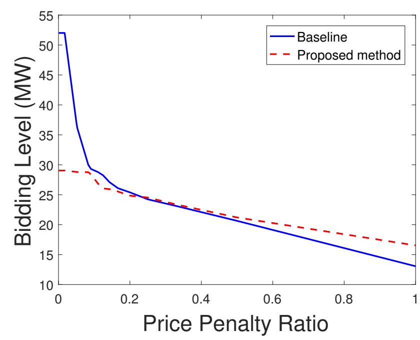

Figure 7: Day-ahead bidding level comparison under price penalty ratios from 0 to 1.

Figure 6 depicts the cost saving percentage of each datacenter at hour 22 when the price penalty ratio ω is 0.25, 0.50 and 0.75, respectively. Further, the percentage of the average cost saving of each datacenter over 24 hours is listed in Table 3. We can observe that the percentage of the average cost saving decreases when the price penalty ratio ω increases. This is intuitive since we have less chance to reduce cost through aggregation when the penalty price is lower. Indeed, when the expected negative penalty price is the same as the expected day-ahead electricity price, there is no need for aggregation since one could always buy any shortfall from the real-time market without penalty.

Figure 7 shows the changes of day-ahead bidding level of the baseline and the proposed method under different price penalty ratios. It can be observed that under both cases, the day-ahead bidding level decreases as the price penalty ratio increases. The reason is that when the price penalty ratio is near 0, datacenters behave more conservatively since the expected negative imbalance price is much higher than the expected day-ahead price. In order to avoid high penalty for energy shortfall, they tend to bid more power amount to lower the possible mismatch between committed power supply in the day-ahead market and realized power demand in the real-time market. On the other hand, when the price penalty ratio is approaching 1, datacenters can buy any shortfall in the real-time market without penalty and therefore tend to bid less. Moreover, the change rate of bidding level of our proposed method with respect to the price penalty ratio is smaller than that of the baseline. This is due to the fact that the proposed method has a smaller power demand uncertainty and therefore is less sensitive to the price penalty ratio.

7 RELATED WORK

In the past decade, multiple schemes have been proposed to reduce the electricity bill of datacenters. From the demand side, dynamic capacity provisioning [21, 28, 29] is developed to reduce energy cost by dynamically turning off servers. Dynamic CPU speed scaling [30–32] is shown to reduce the energy usage of datacenters by dynamically adapting the processing speed of a server to the current load. Geographical load balancing [20, 33–35] is developed to exploit the spatial diversity of electricity prices to minimize the energy cost of geographically distributed datacenters by dynamically routing the user requests to regions with lower energy prices. Different from their works, in this paper we consider datacenters to minimize the energy cost from the supply side by participating in the wholesale electricity market. From the supply side, datacenters can purchase electricity from the retail market [36, 37] with a fixed electricity price by signing bilateral contracts. They can also participate in the wholesale electricity market [4, 38–40] to exploit uncertainty of electricity prices and workloads to minimize their energy cost. Exploiting the temporal diversities of electricity prices to reduce energy cost using storage or delay-tolerant workload has also been investigated in [41–43]. These papers consider datacenters with the same owner participating in different wholesale electricity markets. However, in our paper we focus on the scenario where datacenters managed by different independent owners within the same region jointly participate in the wholesale electricity market. Therefore, we need to apply game-theoretic methods to model this multi-agent problem instead of optimization approaches.

8 CONCLUSION

In this paper, we have proposed a new approach to minimize the electricity cost for datacenters participating in the wholesale electricity market. The electricity cost can be effectively reduced by bidding in the day-ahead market collectively since aggregation can reduce the uncertainty of power demand. We model this aggregation process as a cooperative game and present a cost allocation mechanism based on the marginal contribution of each datacenter to the total expected cost to fairly distribute the optimal expected cost to each datacenter within the grand coalition. Finally, simulations based on real-world traces verify the effectiveness of our proposed cost saving method. In the future, we plan to investigate how to fairly distribute the actual realized cost instead of the expected cost to each participating datacenter.

ACKNOWLEDGMENTS

This work of M. Pan was partially supported by the U.S. National Science Foundation under grants CNS-1343361, CNS-1646607, and CNS-1350230 (CAREER).

APPENDIX

A PROOF OF THEOREM 4.1

We first rewrite (21) as below:

$$c(S) = \min_{Q_S} \mu_p^d Q_S + \mu_p^- \int_{Q_S}^{E_S^{\max}} (u - Q_S) f_S(u)\,\mathrm{d}u - \mu_p^+ \int_{E_S^{\min}}^{Q_S} (Q_S - u) f_S(u)\,\mathrm{d}u, \quad (33)$$

where f_S is the corresponding probability density function (PDF) of the CDF as defined in (18). Then by applying the first order optimality condition associated with Leibniz integral rule, we have

$$\mu_p^d - \mu_p^- (1 - F_S(Q_S)) - \mu_p^+ F_S(Q_S) = 0, \quad (34)$$

$$Q_S^* = F_S^{-1}(\varepsilon^*), \quad \text{where} \quad \varepsilon^* = \frac{\mu_p^- - \mu_p^d}{\mu_p^- - \mu_p^+}. \quad (35)$$

Optimal expected cost is given by direct substitution of Q_S^* into (33):

$$\begin{aligned} c(S) &= \mu_p^d Q_S^* + \mu_p^- \int_{Q_S^*}^{E_S^{\max}} (u - Q_S^*) f_S(u)\,\mathrm{d}u \\ &\quad - \mu_p^+ \int_{E_S^{\min}}^{Q_S^*} (Q_S^* - u) f_S(u)\,\mathrm{d}u \\ &= \mu_p^d Q_S^* + \mu_p^- \int_{\varepsilon^*}^{1} (F_S^{-1}(\theta) - Q_S^*)\,\mathrm{d}\theta \\ &\quad - \mu_p^+ \int_0^{\varepsilon^*} (Q_S^* - F_S^{-1}(\theta))\,\mathrm{d}\theta \\ &= Q_S^* \underbrace{(\mu_p^d - \mu_p^- + \varepsilon^*(\mu_p^- - \mu_p^+))}_{=0} \\ &\quad + \mu_p^+ \int_0^{\varepsilon^*} F_S^{-1}(\theta)\,\mathrm{d}\theta + \mu_p^- \int_{\varepsilon^*}^{1} F_S^{-1}(\theta)\,\mathrm{d}\theta. \end{aligned} \quad (36)$$

B PROOF OF THEOREM 4.2

We introduce an ancillary random variable X_i and rewrite (20) in terms of X_i as follows:

$$X_i \coloneqq E_i - Q_i, \quad (37)$$

$$\Phi_S(Q_S) = \mu_p^d Q_S + \mu_p^- \mathbb{E}\big[\big(\sum_{i \in S} X_i\big)^+\big] - \mu_p^+ \mathbb{E}\big[\big(-\sum_{i \in S} X_i\big)^+\big], \quad (38)$$

$$\sum_{i \in S} \Phi_i(Q_i) = \mu_p^d \sum_{i \in S} Q_i + \mu_p^- \mathbb{E}\big[\sum_{i \in S} (X_i)^+\big] - \mu_p^+ \mathbb{E}\big[\sum_{i \in S} (-X_i)^+\big]. \quad (39)$$

By adopting the equivalent forms of $(x)^+$:

$$(x)^+ := \max(x,0) := \frac{x+|x|}{2}, \tag{40}$$

$$\begin{aligned}
&(38)-(39) = \\
&\quad \mu_p^- \mathbb{E}\left[\frac{\sum_{i\in S} X_i + \left|\sum_{i\in S} X_i\right|}{2} - \sum_{i\in S}\frac{X_i+|X_i|}{2}\right] \\
&\quad - \mu_p^+ \mathbb{E}\left[\frac{\left|\sum_{i\in S} X_i\right| - \sum_{i\in S} X_i}{2} - \sum_{i\in S}\frac{|X_i|-X_i}{2}\right] \\
&\quad = \left(\frac{\mu_p^- - \mu_p^+}{2}\right)\mathbb{E}\big[\big(\big|\sum_{i\in S} X_i\big| - \sum_{i\in S}|X_i|\big)\big] \le 0.
\end{aligned} \tag{41}$$

The above inequality holds according to the triangle inequality, i.e., $\left|\sum_{i\in S} X_i\right| \le \sum_{i\in S}|X_i|$ and also by assumption, we have $\mu_p^- \ge \mu_p^+$. Therefore, $\Phi_S(Q_S) \le \sum_{i\in S}\Phi_i(Q_i)$.

C PROOF OF LEMMA 4.3

First we prove the positive homogeneity. The CDF of the positively scaled E_S is denoted as

$$F_{\beta S}(u) = \Pr(\beta E_S \le u) = F_{\beta S}\left(\frac{u}{\beta}\right).$$

It follows that the quantile function of $F_{\beta S}(u)$ is given by

$$F_{\beta S}^{-1}(\varepsilon^*) = \beta F_S^{-1}(\varepsilon^*).$$

Using the results from Theorem 4.1, we can prove the positive homogeneity as

$$\begin{aligned}
c(\beta S) &= \mu_p^+ \int_0^{\varepsilon^*} F_{\beta S}^{-1}(\theta)\,\mathrm{d}\theta + \mu_p^- \int_{\varepsilon^*}^1 F_{\beta S}^{-1}(\theta)\,\mathrm{d}\theta \\
&= \beta\left(\mu_p^+ \int_0^{\varepsilon^*} F_S^{-1}(\theta)\,\mathrm{d}\theta + \mu_p^- \int_{\varepsilon^*}^1 F_S^{-1}(\theta)\,\mathrm{d}\theta\right) \\
&= \beta c(S).
\end{aligned} \tag{42}$$

Next we prove the subadditivity as

$$\begin{aligned}
c(S_1) + c(S_2) &= \min_{Q_{S_1}} \Phi_{S_1}(Q_{S_1}) + \min_{Q_{S_2}} \Phi_{S_2}(Q_{S_2}) \\
&= \Phi_{S_1}(Q_{S_1}^*) + \Phi_{S_2}(Q_{S_2}^*),
\end{aligned} \tag{43}$$

where $Q_{S_1}^*$ and $Q_{S_2}^*$ are the optimal day-ahead bids of their respective minimization problems. It follows from Theorem 4.2 that

$$\begin{aligned}
\Phi_{S_1}(Q_{S_1}^*) + \Phi_{S_2}(Q_{S_2}^*) &\ge \Phi_{S_1\cup S_2}(Q_{S_1}^* + Q_{S_2}^*) \\
&\ge \Phi_{S_1\cup S_2}(Q_{S_1\cup S_2}^*) \\
&= c(S_1 \cup S_2),
\end{aligned} \tag{44}$$

where $Q_{S_1\cup S_2}^*$ is the optimal solution of the expected cost minimization problem under coalition $S_1 \cup S_2$, while $Q_{S_1}^* + Q_{S_1}^*$ is a feasible solution of the minimization problem, then it follows that $c(S_1 \cup S_2) \le c(S_1) + c(S_2)$.

D PROOF OF NONCONVEX GAME

Consider a cooperative game involving three datacenters, indexed by $\mathcal{A}_1$, $\mathcal{A}_2$ and $\mathcal{A}_3$, respectively. We assume the marginal distribution of $\mathcal{A}_1$ and $\mathcal{A}_2$ are given by

$$\mathcal{A}_i = \begin{cases} 2, & \text{w.p. } 0.5 \\ 4, & \text{w.p. } 0.5 \end{cases} \quad \forall i = 1,2.$$

Further, assume $\mathcal{A}_3$ is perfectly positively correlated to $\mathcal{A}_2$, i.e., $\mathcal{A}_3 = \mathcal{A}_2$. Set the expected day-ahead, negative imbalance and positive imbalance prices as $\mu_p^d = 0.9$, $\mu_p^- = 1.4$ and $\mu_p^+ = 0.4$, respectively. Then based on Theorem 4.1, we have:

$$\begin{aligned}
\varepsilon^* &= \frac{1.4-0.9}{1.4-0.4} = 0.5, \\
c(\{1\}) &= c(\{2\}) = c(\{3\}) = 3.2, \\
c(\{1,2\}) &= c(\{1,3\}) = 5.9, \\
c(\{2,3\}) &= 6.4, \\
c(\{1,2,3\}) &= 9.1.
\end{aligned}$$

Here, we choose two coalitions as $\mathcal{S} = \{1,2\}$ and $\mathcal{T} = \{1,3\}$, and then from the above example, we have:

$$c(\{1,2\}) + c(\{1,3\}) = 10.8 \le c(\{1,2,3\}) + c(\{1\}) = 12.3,$$

which violates the definition of convex game given in (4). Therefore, our cooperative game is nonconvex.

E PROOF OF THEOREM 5.3

Our proof is similar to [44, 45] which focus on different aggregation problems. Here we only give a sketch of the proof process. The basic idea is that we could also use the noncooperative game theory to model the same problem by allowing power exchange within datacenters as well, and our proposed allocation method can find the Nash equilibrium of the formulated noncooperative game. Since the core of our cooperative game can be shown to be the same as the Nash equilibrium of the corresponding noncooperative game, our proposed cost allocation scheme is guaranteed to find the core of the cooperative game. Details about the proof process can be found in [44, 45].

REFERENCES

[1] National Resources Defense Council. Scaling up energy efficiency across the data center industry: evaluating key drivers and barriers. *Issue Paper*, 2014.
[2] James Hamilton. Cost of power in large-scale data centers. *Blog entry dated*, 11:28, 2008.
[3] Richard Brown et al. Report to congress on server and data center energy efficiency: Public law 109-431. Technical report, Lawrence Berkeley National Laboratory, 2008.
[4] Liang Yu, Tao Jiang, Yang Cao, and Qian Zhang. Risk-constrained operation for internet data centers in deregulated electricity markets. *IEEE Transactions on Parallel and Distributed Systems*, 25(5):1306–1316, 2014.
[5] Yuanxiong Guo, Yanmin Gong, Yuguang Fang, Pramod P Khargonekar, and Xiaojun Geng. Energy and network aware workload management for sustainable data centers with thermal storage. *Parallel and Distributed Systems, IEEE Transactions on*, 25(8):2030–2042, 2014.
[6] Electric reliability council of texas. http://www.ercot.com/.
[7] California iso. http://www.caiso.com/.

[8] Google gets go-ahead to buy, sell energy. https://www.cnet.com/news/google-gets-go-ahead-to-buy-sell-energy/, Feb. 2010.
[9] Mahdi Ghamkhari, Adam Wierman, and Hamed Mohsenian-Rad. Energy portfolio optimization of data centers. *IEEE Transactions on Smart Grid*, 2016.
[10] Walid Saad, Zhu Han, Mérouane Debbah, Are Hjorungnes, and Tamer Basar. Coalitional game theory for communication networks. *IEEE Signal Processing Magazine*, 26(5):77–97, 2009.
[11] H Peyton Young. *Cost allocation: methods, principles, applications.* North Holland Publishing Co., 1985.
[12] Stef H Tijs and Theo SH Driessen. Game theory and cost allocation problems. *Management Science*, 32(8):1015–1028, 1986.
[13] Lloyd S Shapley. Cores of convex games. *International journal of game theory*, 1(1):11–26, 1971.
[14] Lloyd S Shapley. On balanced sets and cores. *Naval research logistics quarterly*, 14(4):453–460, 1967.
[15] Roger B Myerson. Game theory: analysis of conflict. *Harvard University*, 1991.
[16] Guillermo Owen. *Game theory.* Emerald Group, 2013.
[17] Theo SH Driessen. *Cooperative games, solutions and applications*, volume 3. Springer Science & Business Media, 2013.
[18] Jayaram K Sankaran. On finding the nucleolus of an n-person cooperative game. *International Journal of Game Theory*, 19(4):329–338, 1991.
[19] Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. Renewable and cooling aware workload management for sustainable data centers. In *ACM SIGMETRICS Performance Evaluation Review*, volume 40, pages 175–186. ACM, 2012.
[20] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven H Low, and Lachlan LH Andrew. Greening geographical load balancing. In *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, pages 233–244. ACM, 2011.
[21] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking (TON)*, 21(5):1378–1391, 2013.
[22] Qiang Wu. Making Facebook's software infrastructure more energy efficient with Autoscale.
[23] Martin J Osborne and Ariel Rubinstein. *A course in game theory.* MIT press, 1994.
[24] Mahdi Kohansal and Hamed Mohsenian-Rad. Price-maker economic bidding in two-settlement pool-based markets: The case of time-shiftable loads. *IEEE Transactions on Power Systems*, 31(1):695–705, 2016.
[25] Eilyan Y Bitar, Ram Rajagopal, Pramod P Khargonekar, Kameshwar Poolla, and Pravin Varaiya. Bringing wind energy to market. *IEEE Transactions on Power Systems*, 27(3):1225–1235, 2012.
[26] Enrique Baeyens, Eilyan Y Bitar, Pramod P Khargonekar, and Kameshwar Poolla. Coalitional aggregation of wind power. *IEEE Transactions on Power Systems*, 28(4):3774–3784, 2013.
[27] Charles Reiss, John Wilkes, and Joseph L Hellerstein. Google cluster-usage traces: format+ schema. *Google Inc., White Paper*, pages 1–14, 2011.
[28] Brian Guenter, Navendu Jain, and Charles Williams. Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning. In *INFOCOM, 2011 Proceedings IEEE*, pages 1332–1340. IEEE, 2011.
[29] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. Online dynamic capacity provisioning in data centers. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*, pages 1159–1163. IEEE, 2011.
[30] Adam Wierman, Lachlan LH Andrew, and Ao Tang. Power-aware speed scaling in processor sharing systems. In *INFOCOM 2009, IEEE*, pages 2007–2015. IEEE, 2009.
[31] Rade Stanojevic and Robert Shorten. Distributed dynamic speed scaling. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–5. IEEE, 2010.
[32] Lijun Chen and Na Li. On the interaction between load balancing and speed scaling. *IEEE Journal on Selected Areas in Communications*, 33(12):2567–2578, 2015.
[33] Minghong Lin, Zhenhua Liu, Adam Wierman, and Lachlan L-H Andrew. Online algorithms for geographical load balancing. In *Green Computing Conference (IGCC), 2012 International*, pages 1–10. IEEE, 2012.
[34] Mohammad A Islam, Anshul Gandhi, and Shaolei Ren. Minimizing electricity cost for geo-distributed interactive services with tail latency constraint.
[35] Ying Zhang, Lei Deng, Minghua Chen, and Peijian Wang. Joint bidding and geographic load blancing for datacenters: Is uncertainty a blessing or a curse? In *INFOCOM, 2017 Proceedings IEEE*. IEEE, 2017.
[36] Cheng Wang, Bhuvan Urgaonkar, George Kesidis, Uday V Shanbhag, and Qian Wang. A case for virtualizing the electric utility in cloud data centers. In *HotCloud*, 2014.
[37] Zhenhua Liu, Iris Liu, Steven Low, and Adam Wierman. Pricing data center demand response. *ACM SIGMETRICS Performance Evaluation Review*, 42(1):111–123, 2014.
[38] Lei Rao, Xue Liu, Le Xie, and Wenyu Liu. Minimizing electricity cost: optimization of distributed internet data centers in a multi-electricity-market environment. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. IEEE, 2010.
[39] Lei Rao, Xue Liu, Le Xie, and Zhan Pang. Hedging against uncertainty: A tale of internet data center operations under smart grid environment. *IEEE Transactions on Smart Grid*, 2(3):555–563, 2011.
[40] Mahdi Ghamkhari, Hamed Mohsenian-Rad, and Adam Wierman. Optimal risk-aware power procurement for data centers in day-ahead and real-time electricity markets. In *Computer Communications Workshops (INFOCOM WKSHPS), 2014 IEEE Conference on*, pages 610–615. IEEE, 2014.
[41] Yuanxiong Guo, Zongrui Ding, Yuguang Fang, and Dapeng Wu. Cutting down electricity cost in internet data centers by using energy storage. In *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, pages 1–5. IEEE, 2011.
[42] Yuan Yao, Longbo Huang, Abhihshek Sharma, Leana Golubchik, and Michael Neely. Data centers power reduction: A two time scale approach for delay tolerant workloads. In *INFOCOM, 2012 Proceedings IEEE*, pages 1431–1439. IEEE, 2012.
[43] Yuanxiong Guo and Yuguang Fang. Electricity cost saving strategy in data centers by using energy storage. *IEEE Transactions on Parallel and Distributed Systems*, 24(6):1149–1160, 2013.
[44] Yue Zhao, Junjie Qin, Ram Rajagopal, Andrea Goldsmith, and H Vincent Poor. Wind aggregation via risky power markets. *IEEE Transactions on Power Systems*, 30(3):1571–1581, 2015.
[45] Dileep Kalathil, Chenye Wu, Kameshwar Poolla, and Pravin Varaiya. The sharing economy for the smart grid. *arXiv preprint arXiv:1608.06990*, 2016.