IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 22 June 2022, accepted 30 June 2022, date of publication 18 July 2022, date of current version 1 August 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3192038

RESEARCH ARTICLE

Rebate Auction Mechanisms for Bidirectional Grid Balancing Using Cloud Workload Migrations

AHMED ABADA, MARC ST-HILAIRE, (Senior Member, IEEE), AND WEI SHI, (Member, IEEE)
Department of Systems and Computer Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada
Corresponding author: Ahmed Abada (ahmedabada@cmail.carleton.ca)

ABSTRACT Increasing the power-grid's flexibility is essential for expanding the integration of renewable energy sources in modern power grids. This work presents a new rebate auction framework that allows power grids to use cloud datacenters as managed loads to provide upward/downward grid flexibility. Since the energy consumption of datacenters is proportional to their computational workload, this paper presents a rebate auction framework that can induce cloud workload migrations between datacenters to correct energy imbalances. Unlike existing datacenter-based power grid balancing approaches that have only focused on providing downward flexibility and only considered owner-operated datacenters, the proposed framework provides bidirectional flexibility and accommodates both owner-operated and public cloud datacenters. Thus, providing a general framework for creating workload migrations between datacenters to balance the power grid. Because workloads on public cloud datacenters are managed by end-users, the proposed framework uses monetary incentives (auctioned rebates) to encourage large-scale end-users to migrate their cloud workloads between datacenters to correct energy imbalances. The use of monetary rewards as an incentive hides the complexity of grid-balancing from auction participants, who only participate in the auction to lower their cost, while grid-balancing happens as a result of workload migrations. This paper presents and compares two auction implementations under the proposed framework, a strategy-proof implementation that guarantees truthful bidding as a dominant strategy, but has NP-hard computational complexity, and an alternative implementation that does not guarantee truthful bidding, but has polynomial time complexity. Simulation results show that the proposed framework is effective in incentivizing cloud workload migrations to achieve the grid balancing goal and provides positive utility to all participants.

INDEX TERMS Datacenters, rebate auction, upward/downward grid balancing, renewable energy, Vickrey Clarke Groves (VCG) mechanism, VCG rebate auction (VRA), uniform rebate auction (URA).

I. INTRODUCTION

Grid balancing is essential for the proper operation of all power grids. Since they can not store energy in large quantities, power grids must maintain equal generation and consumption levels at all times to ensure system stability and reliable power delivery [1], [2]. To that end, power grids must allocate enough flexibility-resources capacity to perform upward/downward adjustments (increasing/decreasing the energy levels on the generation or consumption sides, respectively) as needed to ensure grid balance [3], [4]. However, the recent trend of greening the power supply by using more Renewable Energy Sources (RESs) for power generation has made the balancing task more challenging due to the inherent variability and unpredictability of such sources [5], [6]. Unpredicted imbalance events (due to sudden RESs' generation changes) can lead to high scarcity prices in events of energy shortage and negative prices in events of energy surplus, as well as possible system failures [7]–[9]. As such, the above balancing challenges are currently limiting the deployment of RESs in modern power grids to the levels manageable by the existing balancing capacity.

The associate editor coordinating the review of this manuscript and approving it for publication was Siqi Bu.

A. MOTIVATION AND BACKGROUND

1) EXISTING GRID BALANCING APPROACHES

Traditional power grid balancing approaches are generation-side oriented. They rely on adjusting the output of the generation-side of the power grid to follow the instantaneous demand variations of the load. Such approaches include Unit Commitment (UC), Economic Dispatch (ED), and Automatic Generation Control (AGC). UC optimizes the number of activated generation units over a future time horizon to minimize the operating cost [10], [11]. ED decides the operational set-points of the activated generation units selected by UC to minimize their operating cost and ensure that the balancing constraint is always met [12], [13]. AGC is a closed-loop feedback protocol that adjusts the output of the grid's operating reserves (flexible generation capacity) to balance against real-time demand variations [14], [15].

More recently, the advent of power-markets deregulation and smart-grid technologies has enabled new grid-balancing approaches that leverage demand-side flexibility. Power-markets deregulation has opened the energy markets to new players to foster competition. This has led to the creation of new energy markets, such as the ancillary services market, where different system reliability products are sold by generators and consumers [16]–[18]. On the other hand, smart-grid technologies have enhanced the power grids' communication, control and automation capabilities. The added capabilities have allowed them to introduce new Demand-Side Management (DSM) programs like Demand Response (DR), aiming to harness the flexibility of end-users' demand to maintain grid balance. DR programs were introduced by power-grids' Independent System Operators (ISOs) and local utilities to leverage the collective flexibility of commercial end-users for the purpose of grid balancing [19]–[21]. They encompass a variety of methods that use monetary incentives to control end-users' demands.

2) THE CHALLENGE OF BALANCING RENEWABLE ENERGY SOURCES

The above grid-balancing approaches are cost-efficient for conventional generation sources because their output can be precisely controlled. However, the addition of RESs into the generation mix introduces generation-side uncertainties that require a significant increase in provisioned flexibility resources to ensure reliable power delivery. The additional cost of balancing the output of RESs renders their energy more costly than conventional sources. Thus, limiting their integration in power grids to the levels at which their energy can be balanced reliably in a cost-efficient manner [22]. Therefore, it is crucial to increase the power-grid's flexibility (its ability to adjust generation and demand levels in response to imbalance events) in a cost-efficient manner (without having to rely on expensive operating reserves) to increase the integration of RESs into the power grid.

3) USING DATACENTERS FOR GRID BALANCING

Cloud datacenters are well suited to participate in DR programs because they are programmable, consume large amounts of energy [23], [24], and their energy consumption is proportional to their computational workload [25]–[27]. Therefore, they are well-suited to act as managed loads under DR to provide the needed DSM. However, public cloud datacenters (such as the ones operated by Infrastructure as a Service (IaaS) providers) cannot participate directly in DR programs because the workloads hosted on them are typically managed by the cloud end-users. Therefore, it is essential to elicit the participation of the end-users in adjusting the scheduling/placement of their computational workloads on such datacenters to provide the required DSM. The challenge with getting end-users' involvement in adjusting the scheduling/placement of their computational workloads on IaaS datacenters is that they are not direct customers of the power grids, and their energy consumption is not individually metered, so they cannot receive DR signals or rewards. Moreover, they are typically not exposed to the energy consumption aspect of their workloads on IaaS datacenters. As such, they cannot determine the amount of workload adjustments needed to correct an energy imbalance. Therefore, a new collaboration scheme is needed between the end-users, datacenters, and the power grid to enable cloud end-users' participation in adjusting the computational workloads of IaaS datacenters to correct energy imbalances.

B. THE PROPOSED APPROACH

This paper introduces a new rebate auction framework that auctions monetary rebates to cloud end-users to incentivize them to migrate their workloads between IaaS datacenters to correct energy imbalances. Rebate auctions are an opportunity for cloud end-users to reduce their operational costs by migrating their workloads in the direction required by the auction. While more simplistic fixed-price incentives could be offered to cloud end-users to incentivize the desired workload migrations, such a strategy would fail to capture current market conditions (as they may result in no migrations if the incentives are not enough or over-spending if the incentives are higher than what they should be). Therefore, auctions are used in this work to find the market value of the incentives that should be given to the end-users to perform the needed cloud workload migrations. Since workload migrations are a central component of the proposed rebate auction framework, finalizing the transactions cleared by the auction is made conditional on the completion of workload migrations. This paper further introduces two implementations under the proposed framework (using different types of auction mechanisms) and illustrates the properties and advantages of each.

C. CONTRIBUTIONS

The contributions of this paper are the following:

- A rebate auction framework that can incentivize cloud workload migrations between IaaS datacenters to achieve bidirectional power grid balancing.

- A Vickrey-Clarke-Groves (VCG) auction implementation based on the proposed framework. This implementation carries the strategy-proofness property of the VCG auction and guarantees that truthful bidding is a dominant strategy for all bidders wishing to maximize their utility. However, it has NP-hard computational complexity and is used as a benchmark for other implementations because it maximizes the total utility of all participants.
- A Uniform auction implementation based on the proposed framework. This implementation is more practical because its auction engine runs in polynomial time. However, it does not maximize the total utility, nor is strategy-proof.

The proposed framework introduces the ideas of selling monetary rebates as an incentive to perform a certain task, and making the finalization of the auction transaction conditional on completing the task. To the authors' best knowledge, this is the first use of both ideas in an auction setting. While traditional auctions sell items for monetary payments to generate revenue, the proposed auction framework sells rebates (money) in an auction as an incentive to perform a task. The idea of using the proposed auction framework is that the difference between the rebate received in the auction transaction and the payment required for it represents the incentive for performing the task in question, while the objective of the auction is to find the market value of this incentive. Since this is an auction sale of incentive, issuing the incentive is conditional on completing the task. While the framework in this paper focuses on providing an incentive for cloud workload migrations, the idea of selling rebates in an auction as an incentive to perform a task is extendable to cover other scenarios that require an auction-determined value for incentives.

D. PAPER ORGANIZATION

The rest of this paper is organized as follows: Section II reviews the related work in the area of grid balancing using cloud datacenters. Section III presents the system model and explains the interactions between its components. Section IV presents the proposed rebate auction framework and highlights the main steps of its operation. Section V illustrates the monetary effect of the proposed auction framework on all participants. Section VI presents two auction implementations under the proposed framework and explains their operation. Section VII presents a performance evaluation of the two provided implementations. Lastly, Section VIII concludes the paper.

II. RELATED WORK

Many research works (shown in Table 1) have already proposed using datacenters as managed loads by integrating them into DR programs to provide grid flexibility. Aside from Abada *et al.* [28], they have only focused on providing downward flexibility (to reduce energy demand during times of generation shortages) and only considered owner-operated datacenters where datacenter owners control the scheduling of computational workloads. This section outlines the previous work in the area of datacenters integration in DR to provide grid flexibility, and highlights the features of the current proposal.

TABLE 1. Summary of related literature and current work.

Research	Model parameters	Datacenter type	Flexibility direction
[29]	- Local energy storage - Variable energy prices	Owner operated	downward
[30]	- Price signals - On-site batteries	Owner operated	downward
[31]	- Variable energy prices - Future workload prediction	Owner operated	downward
[32]	- Reward for datacenter load reduction	Owner operated	downward
[33]	- Reward for datacenter load reduction - Ensures reliable load reduction	Owner operated	downward
[34]	- Dynamic energy pricing at different locations to control demand	Owner operated	downward
[35]	- On-site renewables and storage - Power procurement decisions	Owner operated	downward
[28]	- Sale of excess energy in an auction - Uses cloud workload migrations	Public IaaS	upward
Current work	- Sale of monetary rebates in an auction - Uses cloud workload migrations	Public IaaS	upward/ downward

Ghamkhari *et al.* [29] consider the possibility for datacenters with on-site energy storage to provide voluntary energy reduction. They develop a profit-maximization model that considers varying energy prices, datacenters revenue from providing their existing services, and compensations for providing the requested energy reductions. Nguyen *et al.* [30] design a game theory model for minimizing the maximum Peak-to-Average (PTA) load of a power grid during the day. They consider end-users with on-site batteries and different energy demand profiles and try to optimize the price signals during the day to minimize the maximum PTA energy demand while end-users try to minimize their cost. Liu *et al.* [31] consider the feasibility of using prediction-based-pricing market model to optimize the values of rewards given to datacenters for their participation in DR. Cao *et al.* [32] present a game theory model for bargaining between datacenters and a power grid to decide the datacenters' rewards for load reduction. Ma *et al.* [33] develop two systems that focus on reliable DR while considering end-users' performance on load reduction. The methods try to select the minimum number of users from a possible pool of participants to guarantee a certain level of reliability and minimize cost. Wang *et al.* [34] design a two-stage decision approach to model the interaction between the power grid and datacenters. In the first stage, the power grid sets the electricity price to balance the grid, while the datacenters decide their workload execution schedule in the second stage to minimize their cost. Kwon *et al.* [35] design a two-stage stochastic programming model to optimize power procurement and server provisioning decisions. Their model takes

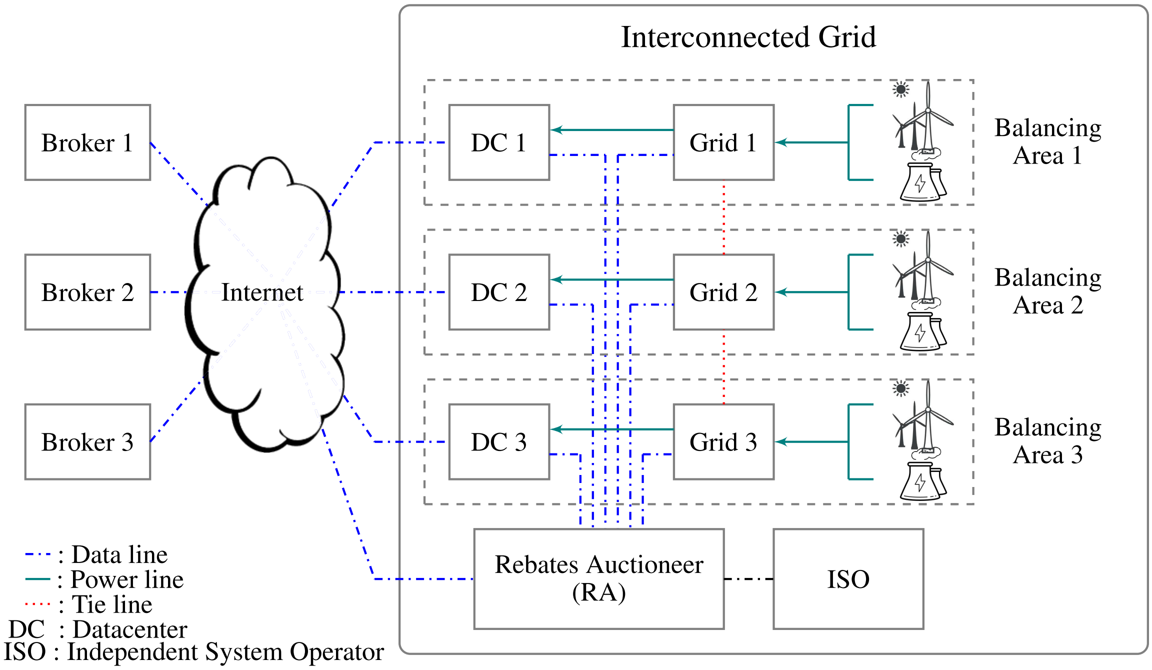

FIGURE 1. System model.

into account on-site renewables and energy storage. All of the above models target load reduction and assume total control of the datacenters on workload scheduling.

In [28], Abada *et al.* introduce the first proposal for using workload migrations of public cloud datacenters to provide upward flexibility and consume excess RESs' energy. The model there represented the excess RESs' energy as Energy Credits (ECs) and sold them to large-scale cloud end-users (cloud brokers) in an auction to incentivize them to migrate their workloads to the datacenter location where the ECs can be used to increase its computational workload and speed up the consumption of the excess energy. While it was successful in achieving the desired task of providing upward flexibility, it relied on restrictive assumptions that limited its practicality in current cloud settings. Mainly, it assumed that cloud brokers pay for the energy consumption of their workloads, and that they were provided with the energy consumption parameters of the different datacenters. This was required to allow them to compute the energy consumption of their workloads and its related cost on any datacenter. Lastly, since the model sold excess energy as ECs in an auction, it only provided upward flexibility, as it was not possible to sell "negative energy" in the case of energy shortage to provide downward flexibility.

To address the above limitations, this paper introduces the idea of using rebate auctions (auctioning of monetary rebates) instead of auctioning the energy itself. Therefore, it can handle both cases of upward and downward flexibility. This new approach makes the system more compatible with current cloud workflows, as it does not assume that cloud users pay for the energy consumption cost of their workloads nor need to have the energy consumption parameters of the different datacenters. However, it still applies the same principle of adjusting the energy consumption of datacenters through cloud workload migrations.

III. SYSTEM MODEL

This paper considers the interconnected power grid model (called an "Interconnection") of deregulated power markets [36], [37]. The interconnection model is comprised of multiple Balancing Areas (BAs) connected together by tie-lines, as shown in Figure 1. Each BA is managed by a local utility company (shown as the "Grid") that uses conventional as well as renewable energy sources and is responsible for maintaining local grid balance. BAs use existing UC, ED, AGC, and DR balancing approaches (as described in Section I) to maintain local balance and zero power flow on the tie-lines [15]. All BAs within the same interconnection are governed by a central ISO that has a global view of the system, runs the different energy markets within the interconnection, and must approve all system-wide operations to ensure reliable power delivery [38]. In addition, each BA is assumed to house one datacenter that operates as an IaaS cloud provider and agrees to participate in the proposed rebate auctions administered by the Rebate Auctioneer (RA). While the interconnection model of deregulated power markets may include other market players such as independent transmission and generation providers, this paper limits the scope of the model to the actors directly relevant to the proposed framework to simplify the presentation.

In normal operating conditions, BAs facing an unexpected energy imbalance would need to correct it either by using local flexibility resources or buying balancing capacity from real-time balancing markets administered by the ISO (in the case of deregulated power markets) [18], [39]. Neighboring BAs will also face the same choices if the imbalance is not corrected locally. Therefore, the purpose of the proposed balancing framework is to use IaaS datacenters' workload modulation to proactively correct energy imbalances locally and avoid having to use expensive balancing alternatives or disrupt the balance of neighboring BAs.

The model assumes that time is divided into equally sized intervals and that the proposed auction framework is executed at the beginning of each interval if needed. It also assumes that each cloud broker receives a random number of end-users' workloads in each time interval, and selects the lowest-cost datacnter to host the received workloads. When a new rebate auction is announced, cloud brokers can submit bids to migrate portions of their workloads in the direction required by the auction (always to a higher-cost datacenter) to take advantage of the rebate auction and further reduce their costs. The idea of a rebate auction is that it offsets all the costs associated with migrating workloads to a higher-cost datacenter and provides additional cost-savings to the winning bidders (determined by the auction). Thus, providing an incentive for cloud brokers to participate in the auction to lower their operational costs. The main actors considered in the system model are introduced in the following subsections.

1) CLOUD BROKERS

Cloud brokers aggregate and manage the cloud workflows of many end-users. They provide added-value services such as automated workload monitoring, management, and customized support to simplify the cloud deployment process for end-users. They do not maintain physical computing infrastructure; Therefore, they contract the physical resources needed to host their clients' workloads from IaaS cloud providers (the datacenters in the model). Brokers can use any pricing scheme they see fit to attract their customers' business. The details of their pricing schemes are beyond the scope of this paper since they do not affect the proposed framework. Brokers are responsible for maintaining the Quality of Service (QoS) performance guarantees in their Service Level Agreements (SLAs) offered to their end-users. As such, they can migrate end-users' workloads between datacenters to either maintain the SLA agreements or maximize their return. Cloud workloads received from end-users are described by the four attribute tuple $W = (w^c, w^m, w^s, w^d)$ that represents the required CPU, memory, storage, and time duration of each received workload.

2) IaaS DATACENTERS

The model assumes that datacenters have two internal cost components for the computational resources they offer, energy cost and operational cost. The energy cost component is the cost of energy used to run the allocated cloud resources. On the other hand, the operational cost component covers all non-energy-related costs such as rent, personnel, and profit margin. To simplify the cost calculation for their customers, datacenters announce per-unit cost prices for their computational resources. The announced prices take into account both of the above cost components.

The energy consumption of cloud datacenters is assumed to be linearly proportional to the amount of deployed computing resources. It can be calculated using the parameters $ConvD_k^c$, $ConvD_k^m$ and $ConvD_k^s$ that represent the amount of energy consumed per unit of deployed computational resources (CPU, memory, and storage respectively) per unit time. This is a widely used approximation for estimating the energy consumption of deployed computational resources in datacenters [40], [41]. Therefore, the energy consumed by an IaaS cloud datacenter j for hosting the cloud workload $W_{ij} = (w_{ij}^c, w_{ij}^m, w_{ij}^s, w_{ij}^d)$ of broker i is referred to as $KwhWD_{ij}$ and is computed as shown in Eq. 1.

$$\begin{aligned} KwhWD_{ij} = (&w_{ij}^c * ConvD_j^c + w_{ij}^m * ConvD_j^m \\ &+ w_{ij}^s * ConvD_j^s) * w_{ij}^d \end{aligned} \tag{1}$$

IaaS datacenters announce their computational resources prices $RateD_k^c$, $RateD_k^m$ and $RateD_k^s$ that represent the per-unit costs of their CPU, memory and storage resources respectively per time interval (as shown in table 2). Therefore, a cloud workload $W_{ij} = (w_{ij}^c, w_{ij}^m, w_{ij}^s, w_{ij}^d)$ submitted by a cloud broker i to an IaaS cloud datacenter j costs the broker a monetary value of $CostWD_{ij}$ (payable to the datacenter) as shown in Eq. 2.

$$\begin{aligned} CostWD_{ij} = (&w_{ij}^c * RateD_j^c + w_{ij}^m * RateD_j^m \\ &+ w_{ij}^s * RateD_j^s) * w_{ij}^d \end{aligned} \tag{2}$$

Figure 2 shows the energy and cost interactions between the system components for hosting a workload $W_{ij} = (w_{ij}^c, w_{ij}^m, w_{ij}^s, w_{ij}^d)$ of cloud broker i on datacenter j. According to Eq. 2, datacenter j receives a payment of $CostWD_{ij}$ from broker i for hosting its workload W_{ij}. On the other hand, it consumes an amount of energy equal to $KwhWD_{ij}$ as computed in Eq. 1 and it pays its local grid a moentary value equal to $(KwhWD_{ij} * RateD_j^e)$ for the consumed energy (where $RateD_j^e$ is the local energy price in BA j). Datacenters generate their profit by setting their computational resources costs $RateD_j^c$, $RateD_j^m$ and $RateD_j^s$ to cover all their energy and operational costs.

3) REBATE AUCTIONEER (RA)

The RA is a centralized entity that connects to the grid operators of the BAs, the datacenters, and the cloud brokers (as shown in Figure 1). It simplifies joining the proposed framework (a single point of contact). It is also connected to the ISO to get a global view of the interconnected grid to ensure that approved workload migrations do not cause new imbalance events. The job of the RA is to administer rebate auctions on behalf of BAs that want to correct an energy imbalance. They announce the availability of new auctions to the cloud brokers, receive their migration bids,

TABLE 2. Used variables.

Symbol	Meaning
Sets and Indices	
I	Brokers
J	Datacenters
i	Broker index $0 < i \leq \|N\|$
j, k	Datacenter index $0 < j, k \leq \|M\|$
h	Time interval index $0 < h \leq T$, where T is the length of the scheduling horizon
Bids Variables	
Bid_{ijk}	$Bid_{ijk} = (rr_{ijk}, v_{ijk}, w^c_{ijk}, w^m_{ijk}, w^s_{ijk}, w^d_{ijk})$ that represents the bid of broker i to migrate its workload from datacenter j to datacenter k.
rr_{ijk}	Amount of rebates requested in the bid
v_{ijk}	Valuation offered for the requested rebates
$w^c_{ijk}, w^m_{ijk}, w^s_{ijk}$	The amounts of virtual resources (CPU, memory and storage respectively) to be migrated from j to k
w^d_{ijk}	The duration of migrated wokload
w^e_{ijk}	The energy consumption of migrated workload
p_{ijk}	The payment amount charged to a winning bid Bid_{ijk}
Datacenter Variables	
$RateD^{c,m,s}_k$	Cost per time interval for unit allocations of CPU, memory and storage resources at datacenter k
$RateD^e_k$	Energy cost at datacenter k
$ConvD^{c,m,s}_k$	Energy consumed by one unit of CPU, memory and storage resources allocated at datacenter k per time interval
$Capacity^{c,m,s}_k$	CPU, memory and storage capacities at datacenter k at time interval h
r_k	Amount of rebates available at datacenter k
Workload Variables	
W_{ij}	Total broker i workloads handled by datacenter j, $W_{ij} = (w^c_{ij}, w^m_{ij}, w^s_{ij}, w^d_{ij})$
$w^c_{ij}, w^m_{ij}, w^s_{ij}$	CPU, memory and storage components of W_{ij} respectively
w^d_{ij}	The duration of W_{ij}
Mig_{ijk}	The cost of migrating W_{ij} of broker i from datacenter j to datacenter k
$CostWD_{ij}$	Cost of hosting W_{ij} of broker i at datacenter j for one time interval
$KwhWD_{ij}$	Energy consumed by W_{ij} of broker i at datacenter j in one time interval
Decision Variables	
S_{ijkh}	Binary decision variable that represents the starting time interval during the scheduling time horizon T for the execution of the workload of Bid_{ijk} at datacenter k
A_{ijkh}	Binary decision variable that represents the allocated time intervals during the scheduling time horizon T for the execution of the workload of Bid_{ijk} at datacenter k

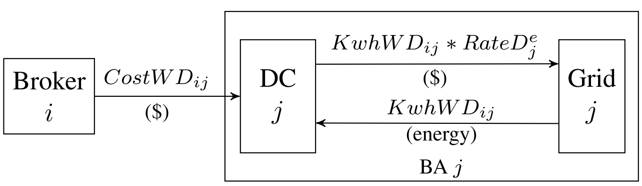

FIGURE 2. Energy and cost interactions.

calculate auction outcomes, verify workload migrations with destination datacenters and issue rebates to cloud brokers for completed migrations. The details of their operation are further illustrated in the following sections.

This paper makes the following assumptions regarding the system model:

- Time is divided into equal-sized intervals.
- A maximum of one rebate auction (one energy imbalance event) may happen at any given time.
- BAs know the total amount of energy imbalance that needs to be corrected. Forecast methods can be used for this purpose where the difference between forecasted generation and forecasted demand is the imbalance to be corrected [42], [43]. BAs also know the amount of rebate they need to offer in the auction. This paper assumes that imbalance penalties are the maximum rebate budgets that can be offered in an auction. This is because it would not make financial sense for BAs to spend more than that to correct an energy imbalance using the proposed framework if it is cheaper to pay the penalty. Both values (the amount of energy imbalance and the offered rebates) are provided to the proposed framework as input, as shown in the following section.
- It may not always be possible to correct energy imbalances in one auction run (in a single time interval) because of limited availability of datacenters' idle capacity or limited arrival of workload migration bids. Therefore, the proposed auctions may need to run for multiple iterations until energy balance is achieved.
- The proposed framework takes into account the workload migration cost in its formulation. However, its value is assumed to be zero to simplify the presentation.

IV. THE REBATE AUCTION FRAMEWORK

Since cloud brokers always select the lowest cost IaaS datacenter to host their end-users' workloads, the idea of the proposed framework is to use rebate auctions as a monetary incentive to make it cheaper for cloud brokers to migrate their workloads in the direction needed to correct energy imbalances. Using rebate auctions, a BA i that needs to balance a certain amount of energy e (either a surplus $e_{surplus}$ or a shortage $e_{shortage}$) would first need to decide the amount of rebate r that it is going to offer in the auction (this value is assumed to be known to the BA as per the assumptions in Section III). The value of r represents the available budget of rebates that can be used to incentivize cloud brokers to migrate their workloads. The BA sends this value (r) along with the amount of energy that needs to be balanced as a new auction tuple $auction_i = (e, r)$ to the RA to administer the auction. The RA announces the new auction to all cloud brokers and invites them to submit their bids for buying the rebates they need to migrate their workloads in the direction required by the auction. Auction announcements made by the RA only include the value r of available rebates (so brokers know the maximum rebate they can bid for) and the required workloads migration direction. For the case of energy surplus $e_{surplus}$, the migration direction would be towards the datacenter of BA i, while in the case of energy shortage $e_{shortage}$, the required migration direction would be away from the datacenter of BA i. It is important to note that auction announcements do not include the amount of energy e that needs to be balanced, as energy information is only needed internally by the RA

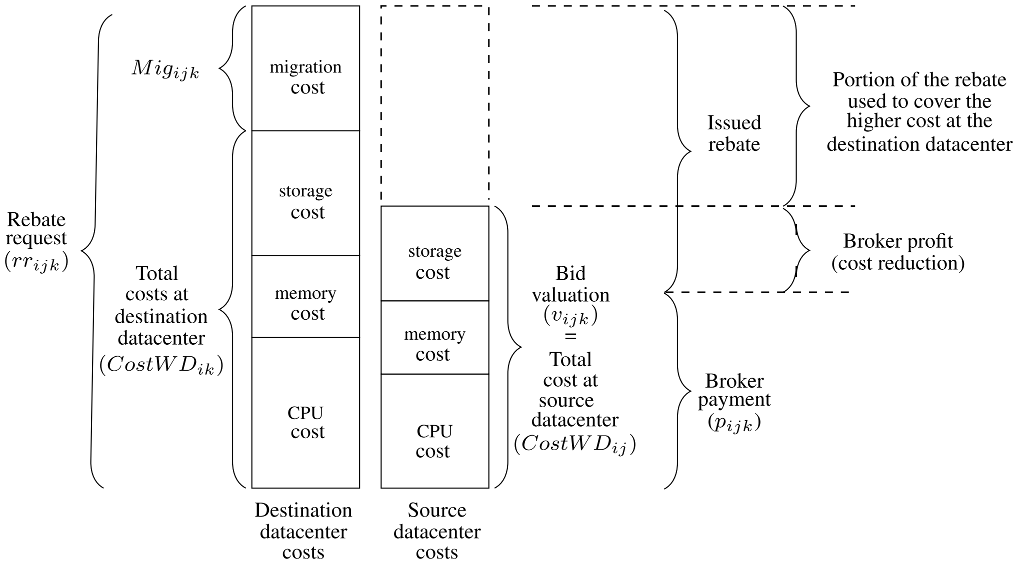

FIGURE 3. Bid calculation.

to ensure that it only allows the right amount of workload migrations needed to correct the imbalance. This simplifies the bidding process for the brokers, so they only participate in the auction to reduce their cost, while the RA takes care of the energy-balancing aspect of the auction.

Upon receiving a new auction announcement, interested cloud brokers prepare their bids for buying the rebates they need to migrate their workloads in the required direction and send them to the RA. The two possible scenarios for workload migrations are either towards a certain IaaS datacenter to increase its energy consumption (to consume excess energy) or away from a certain IaaS datacenter to reduce its energy consumption (to balance an energy shortage). In both cases, the migration always happens from a "Source Datacenter" j to a "Destination Datacenter" k. The bid of a broker i to migrate its cloud workload from datacenter j to datacenter k can be encoded as a bid tuple $Bid_{ijk} = (rr_{ijk}, v_{ijk}, cpu_{ijk}, mem_{ijk}, sto_{ijk}, w^d_{ijk})$. It specifies the amount of rebate requested to perform the migration (rr_{ijk}), the bid valuation (v_{ijk}), the computational workload to be migrated ($cpu_{ijk}, mem_{ijk}, sto_{ijk}$) and its duration ($w^d_{ijk}$) as shown in Table 2. Since workloads are always initially assigned to the datacenter with the lowest value of $CostWD_{ij}$ as per Eq. 2, any migration would always be towards a higher cost datacenter. To ensure that a broker i makes a profit when participating in a rebate auction, the paper assumes that it always requests a rebate amount rr_{ijk} equal to the cost of running its cloud workload at the destination datacenter k ($CostWD_{ik}$) plus the cost of migrating its workload from datacenter j to datacenter k (Mig_{ijk}) as shown in Figure 3 ($rr_{ijk} = CostWD_{ik} + Mig_{ijk}$). On the other hand, it offers a bid valuation v_{ijk} equal to its current cost at its source datacenter j ($CostWD_{ij}$). This guarantees that it would not lose money by migrating its workload since the required auction payment p_{ijk} on any accepted bid Bid_{ijk} is guaranteed not to exceed the bid's valuation v_{ijk}. While other bidding strategies could be used instead, such as requesting a larger amount of rebate rr_{ijk} for the same offered valuation v_{ijk}, such strategies reduce the bid's chance of winning (in the absence of information about competing bids). Therefore, the paper assumes that bidders always use the above bidding strategy (as shown in Figure 3) to maximize their chance of winning.

After bids are received by the RA, an auction mechanism is used to determine the winning bids and calculate their payments p_{ijk} (details are shown in the next section for two auction implementations). The RA then notifies the brokers and the destination datacenters of the auction outcome and identifies the workloads that should be migrated to complete the transactions of accepted bids. Upon completing the required workload migrations, destination datacenters confirm the completed migrations to the RA that then issues the remainder of the requested rebates rr_{ijk} to the cloud brokers (after deducting their payments p_{ijk} as shown in Figure 3). Since the auction transaction of a winning bid involves money exchanged in both directions between the RA and the winning cloud brokers (rr_{ijk} paid by the RA to the broker and p_{ijk} paid by the broker to the RA), the net rebate payment issued to a cloud broker after completing its workload migration is the difference between its requested rebate rr_{ijk} and its payment for the rebate p_{ijk}. This net rebate payment (issued rebate) can be thought of as the sum of two components, the cost of migrating to a higher cost datacenter and the broker's profit (cost reduction) for participating in the auction. The above description explains the steps taken during rebate auctions and how bids and net rebate payments are calculated. It applies to workload migrations for both cases of energy imbalances and can use any auction engine to find the winning bidders and determine their payments (as explained in the next section); As such, it represents the general framework of rebate auctions and is outlined by the message sequence diagram in Figure 4.

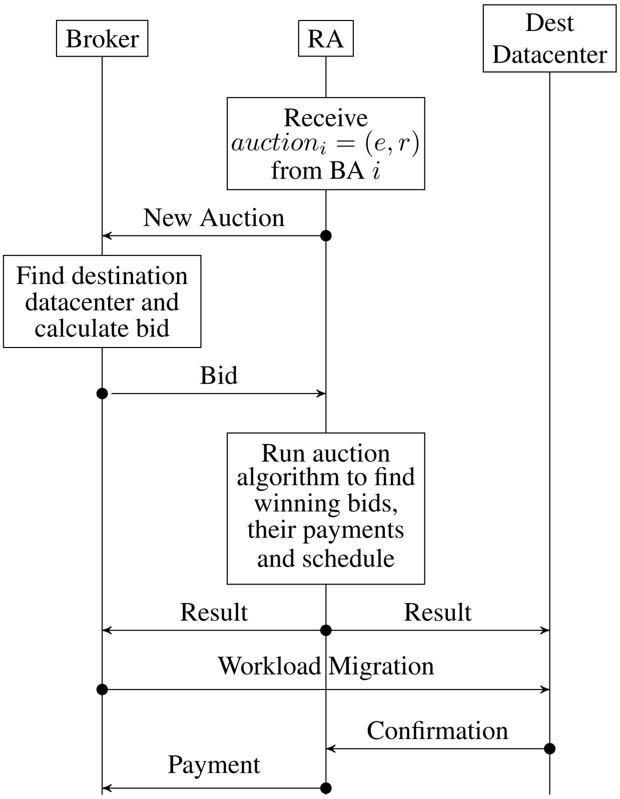

FIGURE 4. **Rebate auction framework.**

V. MONETARY EFFECT ON AUCTION PARTICIPANTS

This section explains the monetary transactions involved in the proposed auction framework and how they affect the auction participants. The calculations presented here apply irrespective of the auction algorithm used to implement the framework; As such, the auction engine used in the implementation is considered as a black box that selects winning bids and determines their payments.

A. EFFECT ON CLOUD BROKERS

Recall that when a bid Bid_{ijk} is accepted, the bidder must pay p_{ijk} to the RA and get rr_{ijk} in return. However, to simplify the transaction, they keep their p_{ijk} (since it is always less than rr_{ijk}) and they get paid the difference $(rr_{ijk} - p_{ijk})$ instead. Brokers can then use p_{ijk} and $(rr_{ijk} - p_{ijk})$ to finance the migration of their workload to the destination datacenter. Cloud brokers always benefit by joining the auction since their payments p_{ijk} are always less than their cost before migration $CostWD_{ij}$ (at the source datacenter) as shown in Figure 3.

B. EFFECT ON DATACENTERS

Datacenters are always paid at their regular prices $RateD_k^{c,m,s}$ regardless of whether brokers use rebate auctions or not. They also always pay the regular energy price for their energy usage (as shown in the cost model in Section III.2). Therefore, the only difference rebate auctions make to datacenters is that they attract workload migrations to datacenters within areas of energy surplus and decrease the workload of datacenters within areas of energy shortage.

C. EFFECT ON POWER GRIDS

The monetary effect of auction transactions on BAs depends on the type of energy imbalance. While BAs always receive brokers' payments p_{ijk} for accepted bids, in the case of energy surplus, the BA that auctioned the rebates also receives the regular energy cost of the consumed excess energy from the destination datacenter (its local datacenter). However, no such revenue is received in the case of energy shortage because no energy consumption happens after workloads are migrated away from the source datacenter. The following explains the monetary impact of rebate auction transactions for each type of energy imbalance.

1) THE CASE OF ENERGY SURPLUS

In the case of an energy surplus in BA k, rebates are auctioned to attract bids of workload migrations from datacenters in other BAs to the datacenter of BA k to consume its excess energy. Figure 5-a shows the energy and cost interactions for both of the source and destination datacenters when BA k has an energy imbalance. Under normal workflow conditions in BA k (no energy imbalance), its grid supplies the local datacenter with energy $KwhWD_{ik}$ and receives a payment of $(KwhWD_{ik} * RateD_k)$. On the other hand, when a rebate auction is used to incentivize the consumption of $KwhWD_{ik}$ (in the case of excess energy), the BA must still supply the energy $KwhWD_{ik}$ to the datacenter, but instead of only receiving a payment of $(KwhWD_{ik} * RateD_k)$ in return, it receives p_{ijk} from the bidding broker, $(KwhWD_{ik} * RateD_k)$ from the datacenter and it must pay the broker its requested rebate rr_{ijk}. Therefore, if the total revenue received by BA k $[pay_{ijk} + (KwhWD_{ijk} * RateD_k^e)]$ is greater than its total spending rr_{ijk} (as shown in Figure 5-b), the auction is said to be recovering a portion of the regular energy cost $(KwhWD_{ijk} * RateD_k^e)$. The Recovered Energy Cost (REC) can be calculated in this case as shown in Eq. 3.

$$REC_{ijk} = [pay_{ijk} + (KwhWD_{ijk} * RateD_k^e)] - rr_{ijk} \quad (3)$$

It is easy to see that in order for rebate auctions to recover any portion of the regular energy cost, both of the extra cost of migrating to a more expensive datacenter and the broker's profit must be covered by the energy revenue from the datacenter; otherwise, the rebate budget would have to be used to cover these cost components. Therefore, the less the cost to move to the destination datacenter and the less the broker's profit, the larger the portion of regular energy cost that can be recovered.

On the other hand, if the total revenue received by BA k $[pay_{ijk} + (KwhWD_{ijk} * RateD_k^e)]$ is less than its total spending rr_{ijk} (due to a higher cost difference between source and

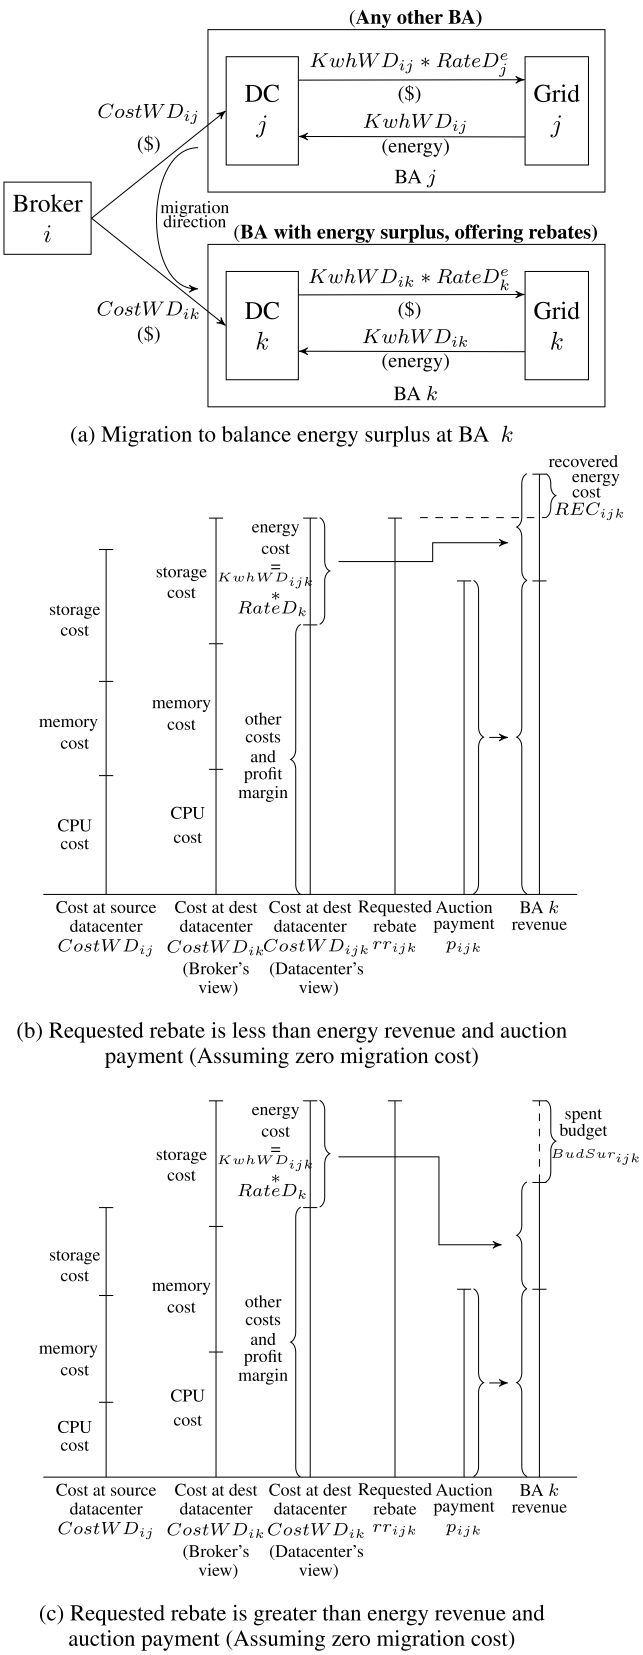

FIGURE 5. Monetary effect on power grids (Surplus).

destination datacenters or a larger broker's profit, as shown in Figure 5-c), the auction is said to be spending budget to consume the excess energy. The amount of budget spent in this case (*BudSur*) is calculated as shown in Eq. 4.

$$BudSur_{ijk} = rr_{ijk} - [pay_{ijk} + (KwhWD_{ijk} * RateD_k^e)] \quad (4)$$

Both of the above cases are illustrated in Figure 5, where zero migration cost (Mig_{ijk}) is assumed to simplify the presentation.

2) THE CASE OF ENERGY SHORTAGE

In this case, the BA that issued the rebates will no longer receive the energy cost from their local datacenters after migrating the workloads. Therefore, the BA would always need to spend from its available rebate budget *r* to incentivize brokers to migrate to a higher cost datacenter. The amount of budget spent in this case of energy shortage ($BudShr_{ijk}$) is shown in Figure 6-b and can be calculated as shown in Eq. 5.

$$BudShr_{ijk} = rr_{ijk} - p_{ijk} \quad (5)$$

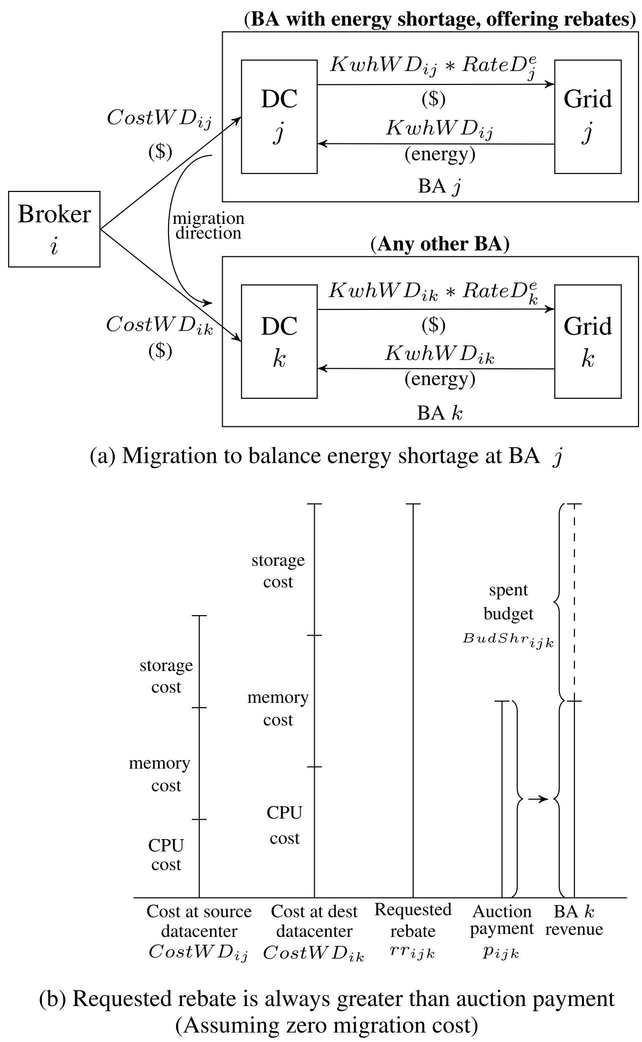

FIGURE 6. Monetary effect on power grids (Shortage).

VI. THE REBATE AUCTION MECHANISMS

This section introduces two auction implementations based on the proposed framework, the VCG Rebate Auction (VRA) and the Uniform Rebate Auction (URA), and outlines the advantages and limitations of each. Both auctions use the same idea of selling monetary rebates in an auction (conditional on performing cloud workload migrations) and accept the same bid format from auction participants. Both auctions must schedule the execution time of accepted workload migrations within a lookahead time horizon of T intervals, where $|T| > max(w^d_{ijk})\ \forall\, i, j, k$. The main tasks performed by the auction implementations are the following: **(1)** Find the energy consumption information of the received bids. This is used to ensure that only the right amount of workload migrations needed to correct an imbalance is allowed. Both implementations get this information by requesting it directly from the appropriate datacenters, **(2)** compute the auction outcome to determine winning bids, **(3)** Find the schedule of migrated workloads on destination datacenters, and **(4)** compute the payments for winning bids. Regardless of which auction implementation is used, the interactions between the bidders and the RA remain the same. The following subsections introduce the details of both implementations and how they perform the above tasks.

A. VCG REBATE AUCTION (VRA)

The VRA implementation uses the VCG auction mechanism to find the winning bids and compute their payments [44]. VCG gets its importance from being a Dominant Strategy Incentive Compatible (DSIC) mechanism, also called truthful or strategy-proof. It guarantees truthful bidding as the dominant strategy of auction participants. Therefore, eliminating the need for a bidding strategy as rational bidders will always choose to bid their true valuations to maximize their utility. VCG is socially efficient, meaning that it allocates the auctioned goods to the bidders with the highest valuation. It achieves social efficiency by maximizing the total utility of all participants, including the auctioneer. This is equivalent to selecting the bids that maximize the total valuation of auctioned goods. To illustrate, assume that the RA has accepted Bid_{ijk} of broker i, the utility of broker i can be computed as $UB_i = v_{ijk} - p_{ijk}$. On the other hand, the utility of BA j (the BA that issued the rebates) is $UBA_j = p_{ijk}$. The total utility can then be computed as $U_{total} = UB_i + UBA_j = v_{ijk}$. Therefore, maximizing the total valuation of accepted bids maximizes the total utility of all participants. However, the truthful property of the VCG auction only holds if its winner determination problem (allocation problem) above (maximizing the total valuation of accepted bids) is solved to optimality. Since this is a combinatorial auction problem where bidders are single-minded (as brokers are not interested in migrating half of their workload or receiving a portion of the requested rebate rr_{ijk}), it is NP-hard in computational complexity. Therefore, this implementation is economically efficient but not computationally efficient. It is introduced here because of its theoretical importance and to serve as a benchmark to compare other implementations against.

The RA achieves the **first auction task** of finding the energy consumption information of the received bids (w^e_{ijk}) by requesting it directly from the datacenters. In the case of an energy surplus (migration towards a destination datacenter), this information is requested from the destination datacenter, while in the case of an energy shortage (migration away from a source datacenter), it is requested from the source datacenter.

Since the truthfulness property of the VCG auction requires solving its allocation problem to optimality, the **second and third auction tasks** (computing the auction outcome and scheduling migrated workloads) are combined into one optimization problem to ensure that the accepted workload migrations have a valid schedule on destination datacenters. This means solving the objective function of the auction (that maximizes the valuation of accepted bids) under the capacity constraints of destination datacenters. Therefore, the RA queries destination datacenters for their available capacity during the lookahead horizon T to set the capacity constraints (Eqs. 7-9) accordingly. Other constraints are also added to the allocation problem to ensure that migrated workloads are allocated a number of time intervals equal to their duration w^d_{ijk}, and that they run uninterrupted for that duration. This combined optimization problem is formulated as a Mixed Integer Linear Program (MILP), as shown in Eqs. 6-16 and is executed by the RA at the beginning of time intervals when it has rebate to auction.

Objective function:

$$Max \sum_{i\in I}\sum_{j\in J}\sum_{k\in J}\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}} S_{ijkh} * v_{ijk} \tag{6}$$

Subject to :

Capacity Constraints :

$$\sum_{i\in I}\sum_{j\in J} w^c_{ijk} * A_{ijkh} < Capacity^c_{kh} \quad (h \in T) \tag{7}$$

$$\sum_{i\in I}\sum_{j\in J} w^m_{ijk} * A_{ijkh} < Capacity^m_{kh} \quad (h \in T) \tag{8}$$

$$\sum_{i\in I}\sum_{j\in J} w^s_{ijk} * A_{ijkh} < Capacity^s_{kh} \quad (h \in T) \tag{9}$$

$$\sum_{i\in I}\sum_{j\in J}\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}} S_{ijkh} * rr_{ijk} < r_k \tag{10}$$

Continuity Constraints :

$$\sum_{i\in I}\sum_{j\in J}\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}} \left(-w^d_{ijk} * S_{ijkh}\right) + \sum_{i\in I}\sum_{j\in J}\sum_{h\in T} A_{ijkh} = 0 \tag{11}$$

$$\sum_{i\in I}\sum_{j\in J}\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}} \sum_{d=0}^{w^d_{ijk}} A_{ijk(h+d)}$$

$$-\sum_{i\in I}\sum_{j\in J}\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}}\left(S_{ijkh} * w^d_{ijk}\right) \geq 0 \quad (12)$$

Single Allocation Constraint :

$$\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}} S_{ijkh} \leq 1, \forall i \in I, j \in J \quad (13)$$

Total Allocation Constraint :

$$\sum_{h\in T} A_{ijkh} \leq w^d_{ijk}, \forall i \in I, j \in J \quad (14)$$

Energy Consumption Constraint :

$$\sum_{i\in I}\sum_{j\in J}\sum_{\substack{h\in T\\ h<(|T|-w^d_{ijk}+1)}} S_{ijkh} * w^e_{ijk} < |e| \quad (15)$$

Integrality Constraint :

$$S_{ijkh}, A_{ijkh} \in \{0, 1\}, \forall i \in I, j \in J, k \in J, h \in T \quad (16)$$

The decision variables in the above optimization problem are the four-dimensional boolean arrays S_{ijkh} and A_{ijkh}. They represent the scheduling information of accepted migration bids. For both arrays, the indexes i, j, k and h are the broker identifier, source datacenter identifier, destination datacenter identifier, and the time interval within the lookahead time horizon T, respectively. The default value for both arrays is "0". S_{ijkh} and A_{ijkh} encode the time intervals at which accepted workload migrations are scheduled to start and their allocated time intervals, respectively. Therefore, an accepted Bid_{ijk} is going to have $S_{ijkh} = 1$ at the time interval h when its workload is scheduled to start executing at the destination datacenter and $A_{ijkh} = 1$ at all the time intervals h for which its workload is scheduled to execute. While either array is enough to deduce the schedule generated by the above optimization problem, both arrays are needed in the model to ensure that the allocated execution time of migrated workloads is not fragmented over the lookahead horizon (using the continuity constraints).

In the above MILP formulation, the objective function of Eq. 6 maximizes the total valuations v_{ijk} of accepted bids as per the VCG auction requirements. A bid B_{ijk} is accepted if it is scheduled to start at any interval during the lookahead horizon (if $S_{ijkh} = 1$ at any interval $h < T$). The capacity constraints in Eqs. 7-10 ensure the availability of enough capacity at destination datacenters to accommodate the migrated workloads and ensure that the sum of allocated rebates rr_{ijk} is less than the available rebates r_k. The continuity constraints (Eqs. 11, 12) ensure that the execution time intervals allocated to migrated workloads are offered in one contiguous block within the lookahead horizon T. The single allocation constraint (Eq. 13) ensures that accepted bids are scheduled only once during the lookahead horizon. The total allocation constraint (Eq. 14) ensures that accepted bids are allocated a number of time intervals equal to their duration w^d_{ijk}. The energy consumption constraint (Eq. 15) ensures that the total energy consumption of migrated workloads does not exceed the amount of imbalance that needs to be corrected. Lastly, the integrality constraint (Eq. 16) ensures no partial allocations on submitted bids.

The **final auction task** of finding the payments of winning bids is achieved in this implementation by using the VCG payment rule. VCG charges each winning bid a payment equal to its social cost. For a winning bid Bid_{ijk}, this value is equal to the loss in utility experienced by other participants due to its participation in the auction. It can be computed according to Eq. 17 as the difference between the total utility of other participants when Bid_{ijk} does not participate (where $v_{(ijk)_{-1}}$ indicates all valuations other than v_{ijk}) minus their total utility when it participates. This requires solving a new optimization problem under the same original constraints for each winning bid Bid_{ijk} (to find the total utility of other bidders when it does not participate). On the other hand, the total utility of other participants when Bid_{ijk} participates can be easily calculated as the optimal value of the objective function (Eq. 6) minus the v_{ijk} valuation of Bid_{ijk} as shown in the second term of Eq. 17.

$$p_{ijk} = max \sum_{i\in I}\sum_{j\in J} S_{ijkh} * v_{(ijk)_{-1}} * w^d_{ijk} - \left[\left(max \sum_{i\in I}\sum_{j\in J} S_{ijkh} * v_{ijk} * w^d_{ijk}\right) - v_{ijk}\right] \quad (17)$$

The operation sequence of VRA can be illustrated by the diagrams in Figures 7 and 8 (excluding the parts labeled for URA) for the cases of energy shortage and surplus imbalances, respectively. The auction always starts when a BA i sends a new rebate auction tuple $auction_i = (e, r)$ to the RA to correct an energy imbalance, as shown in Section IV. The RA announces the new auction to the brokers (announcements include the maximum rebate they can bid for and the required migration direction) and invites them to submit their bids for migrating their cloud workloads. Interested brokers then compute their bids as shown in Section IV and send them to the RA. For the energy shortage case (Figure 7), brokers need to find the lowest cost alternative datacenter to migrate their workloads to, and they use the cost at that datacenter to calculate their bids (as shown in Figure 3). In the case of energy surplus (Figure 8), the destination datacenter can only be the one within the BA i that issued the rebates (to consume its energy surplus), so its cost is used in the same way to calculate the migration bids. The RA achieves the first auction task of finding the energy consumption information of received bids by requesting it directly from the appropriate datacenter, as described above. To ensure that migrated workloads do not exceed the destination datacentrs' capacity, the RA queries destination datacenters for their available capacity over the lookahead time horizon T to set the capacity constraints of Eqs. 7-9 accordingly. The RA then achieves the second and third auction tasks by running the VCG auction (solving the combined auctioning/scheduling optimization problem in Eqs. 6 - 16) to determine the winning bids and their schedule (provided by the decision variables).

The last auction task is achieved by using Eq. 17 to calculate the required payments (p_{ijk}) of winning bids. The RA then notifies the brokers and the destination datacenters of the auction result. When brokers complete the required workload migrations, destination datacenters confirm the completed migrations to the RA, which then issues the remainder of the rebates to the brokers (as shown in Figure 3).

B. UNIFORM-PRICE REBATE AUCTION (URA)

The URA implementation uses the Uniform price auction mechanism to find the winning bidders and compute their payments [45], [46]. It addresses the main limitations of the VRA implementation (centralized scheduling by the RA and NP-hard complexity) by breaking the tight coupling between the auctioning and scheduling tasks performed jointly by the RA in VRA. This lets the RA to only be responsible for auctioning the available rebates (using the Uniform price auction), while workload scheduling is handled locally by the destination datacenters (using the polynomial-time First-Fit bin-packing algorithm [47], [48]). To achieve this division of labor between the RA and the destination datacenters, the RA first requests each possible destination datacenter to provide a tentative schedule of the workloads that it can accommodate within the scheduling horizon T; it then uses the provided schedules to determine the winning bids and their payments. This ensures that all accepted bids have a valid schedule on their destination datacenters. URA addresses the NP-hard computational complexity of the VRA implementation by using sub-optimal polynomial-time algorithms for the auctioning and scheduling tasks. However, it no longer retains the truthfulness and social welfare maximizing properties provided by the VCG auction.

The Uniform price auction sorts the received bids according to their bid densities (v_{ijk}/rr_{ijk}), it then accepts as many bids as possible (subject to auction constraints) in decreasing order of their bid densities, and it charges the accepted bids based on the valuation of the last accepted bid (the *critical bid*). Ordering the bids according to their bid density ensures that higher valued bids are considered first. Since the sorting step can happen in $O(log(n))$ time [49], the Uniform price auction has polynomial time complexity. On the other hand, the scheduling algorithm used by the destination datacenters (the First-Fit bin-packing algorithm) assigns each workload to the earliest time slot where it can be successfully scheduled within the scheduling horizon T under the computational resources' capacity constraints. Since a decision about the feasibility of scheduling each workload within the scheduling horizon T can be made in constant time (by checking if the workload is schedulable at any starting interval within T), the First-Fit scheduling algorithm also runs in polynomial time. Therefore, the whole URA implementation executes in polynomial time.

The operation sequence of URA can be illustrated by the diagrams in Figures 7 and 8 (excluding the parts labeled for VRA) for the cases of energy shortage and surplus imbalances, respectively. URA rebate auctions start and progress

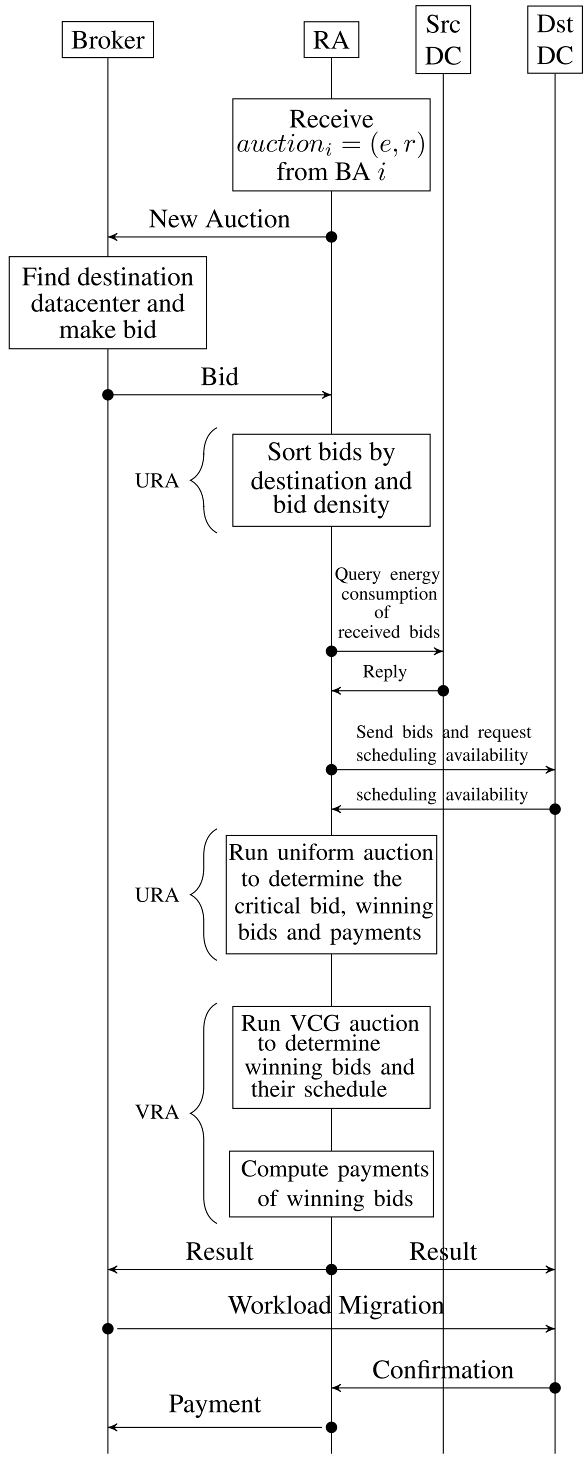

FIGURE 7. Energy shortage message sequences.

in the same way as their VRA counterparts up to the point when the RA receives the brokers' bids. For the energy shortage case shown in Figure 7, the RA achieves the **first auction task** of finding the energy consumption information of the received bids by requesting it directly from the source datacenter. The RA then groups the received bids in separate lists based on their sought destination datacenters and sorts

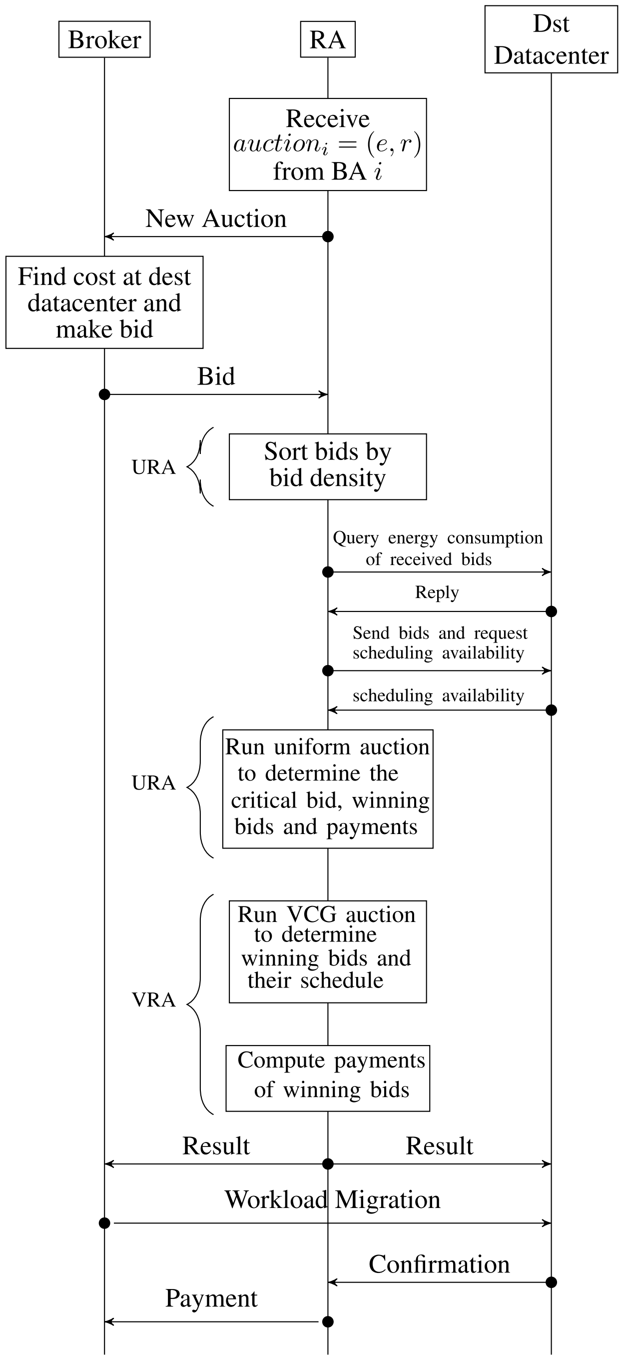

FIGURE 8. Energy surplus message sequences.

the bids of each list in decreasing order of their bid density v_{ijk}/rr_{ijk}. The RA then sends each list of ordered bids to its destination datacenter and requests each destination datacenter to identify the longest portion of its received list that it can schedule successfully (without changing the order of the bids) within the lookahead horizon T, and return the resulting schedule. After receiving the tentative lists of schedulable workloads from the destination datacenters, the RA approves their corresponding bids in decreasing order of their bid density either until all tentatively scheduled bids are approved or until the energy imbalance is corrected. The last approved bid by the RA is the critical bid. The second, third and fourth auction tasks (finding the winning bids, their schedule and payments) are all then achieved in one step; The winning bids are all the ones with bid densities higher than or equal to the critical bid, their schedule is as provided by their corresponding destination datacenters, and their payments are calculated according to the valuation given by the critical bid. The auction result (winning bids, their payments and schedule) is then sent to the brokers and the destination datacenters. When workload migrations are completed, destination datacenters notify the RA that the migrations have already taken place, so the RA can issue the remaining portion of the rebate to the winning brokers (after deducting their payments).

Similarly, for the case of energy surplus shown in Figure 8, the RA achieves its **first auction task** of finding the energy consumption information of the received bids by requesting it directly from the destination datacenter. The received bids are all grouped in one list (all have the same destination datacenter) and sorted in order of their bid density v_{ijk}/rr_{ijk}. The RA then sends the ordered list to the destination datacenter and requests it to identify the longest portion of the received list that it can schedule successfully (without changing the order of the bids) within the lookahead horizon T, and return the resulting schedule. After receiving the tentative lists of schedulable workloads from the destination datacenter, the RA approves their corresponding bids in decreasing order of their bid density either until all tentatively scheduled bids are approved or until the energy imbalance is corrected. The last approved bid by the RA is the critical bid. Similar to the energy shortage case, the **second, third and fourth auction tasks** (finding the winning bids, their schedule, and payments) are all then achieved in one step; The winning bids are all those with bid densities higher than or equal to the critical bid, their schedule is as provided by the destination datacenter and their payments are according to the valuation given by the critical bid. The auction result (winning bids, their payments and schedule) is then sent to the brokers and the destination datacenters. When brokers complete the required workload migrations, destination datacenters confirm the completed migrations to the RA, which then issues the remainder of the rebates to the brokers (as shown in Figure 3).

VII. PERFORMANCE EVALUATION

This section presents the simulation setup, used parameters, performance metrics, and discusses the simulation results. The simulation model was built using the Java programming language, and the auction optimization problem of the VRA implementation (Eqs. 6-16) was solved using the IBM ILOG CPLEX optimization solver V12.5.1 (via its Concert interface library for Java). Since existing datacenter-based grid balancing approaches use owner-operated datacenters (do not consider the case of decision-making end-users) and only provide uni-directional flexibility (as shown in Table 1), they are not directly comparable to the proposed systems. Therefore, the results in this section compare the performance of the proposed rebate auctions under the following performance metrics: 1) Grid balancing time, 2) Percentage of imbalance penalty avoided, 3) Brokers' savings, and 4) Time complexity. The considered metrics cover the grid-balancing aspect of the proposed auctions, their monetary effect on the auction participants, and their time complexity.

A. SIMULATION SETUP

The simulation model consists of 4 cloud datacenters and 6 cloud brokers, where datacenters are assigned randomly chosen cost parameters at the beginning of each simulation run. The time dimension is divided into equal-sized intervals of unit duration. Each cloud broker receives a random number of workloads (uniformly distributed ∈ U(1, 10)) at the beginning of each interval and automatically sends each received workload to the least-cost datacenter as per Eq. 2. Each simulation run starts with 10,000 kWh of energy imbalance at a randomly selected datacenter, and rebate auctions are used as described earlier to correct the imbalances. Simulation experiments were repeated 500 independent runs for each case of energy imbalance (surplus/shortage) with randomly generated inputs for each run. The presented results show the average of the independent runs with 95% confidence intervals.

Since the proposed rebate auctions require datacenters to compute the energy consumption of individual cloud workloads using the $ConvD_j^{c,m,s}$ parameters (as shown in Eq. 1), the simulations were conducted using synthetic input data because such parameters of workloads energy consumption are not made available by the commercial cloud datacenters. However, the presented results provide a proof of concept for the applicability of the proposed approaches, while the exact benefit for the auction participants would depend on their energy consumption and cost parameters, as shown in Section IV-A. The ranges for the parameters used in the simulation are listed in Table 3.

Since the VCG auction used in the VRA implementation guarantees truthful bidding as the dominant strategy for bidders wishing to maximize their utility, the simulation experiments assume that bidders always bid their true valuation (as shown in Figure 3) for both of the VRA and URA implementations. While bidding strategies are always considered in the context of Uniform price auctions [50], [51], no bidding strategies are considered in this performance evaluation since both auction implementations receive the same truthful bids. This is done to compare both implementations under the same conditions and evaluate the performance of the computationally efficient URA auction under truthful bidding.

Since the VCG auction charges each winning bidder its social cost (how much its presence negatively affects other bidders), it would not generate any revenue if the total demand is less than the available supply of auctioned goods. This is because in such a case, the presence of any bidder has no effect on the amount of goods allocated to other bidders. Therefore, it is important to ensure that the amount of rebates auctioned in each iteration is always less than the total amount

TABLE 3. Simulation parameters.

Parameter	Range
Datacenter Parameters	
$CostD^c$	∈ U(20 , 40) \$/GFLOPS
$CostD^m$	∈ U(20 , 40) \$/GB
$CostD^s$	∈ U(20 , 40) \$/GB
$CostD^e$	∈ U(60 , 100) \$/kWh
$ConvD^c$	∈ U(0.2 , 0.5)
$ConvD^m$	∈ U(0.1 , 0.2)
$ConvD^s$	∈ U(0.1 , 0.2)
$Capacity^c$	500 GFLOPS
$Capacity^m$	500 GB
$Capacity^s$	500 GB
$e_{surplus}$	10,000 kWh
$e_{shortage}$	10,000 kWh
Workload Parameters	
w^c	∈ U(2 , 5) GFLOPS
w^m	∈ U(2 , 5) GB
w^s	∈ U(2 , 5) GB
w^d	∈ U(1 , 3) Intervals

of requested rebates rr_{ijk} to generate payments from winning bidders. To do that, the conducted simulation experiments limit the amount of offered rebates in each iteration to a fraction of the total requested rebates rr_{ijk} in the received bids. The rebates availability fractions of 0.6, 0.7, 0.8, and 0.9 are used for this purpose, as shown in the results. This means that if the total amount of requested rebates rr_{ijk} in the received bids is equal to 100, the amount of auctioned rebates would be 60, 70, 80, and 90, respectively, according to the rebates availability fraction used.

B. EFFECT OF REBATE AUCTIONS ON GRID BALANCING TIME

The main goal of rebate auctions is to correct energy imbalances. This performance metric compares the grid balancing capability of the proposed rebate auctions by showing the average time (the number of time intervals) needed by each implementation to correct the same amount of energy imbalance. Since the VCG auction requires the amount of auctioned rebates to be less than the total requested rebates rr_{ijk} to generate revenue, different rebate availability fractions are used (0.6, 0.7, 0.8, and 0.9) as discussed above. Figure 9 shows the results of the two auction implementations for correcting energy surplus imbalances compared to the time needed to consume the energy surplus if no auction is used. The results show that using rebate auctions speeds up the excess energy consumption as they concentrate the workloads at the datacenter that has the energy surplus, allowing it to consume the excess energy faster. It is also shown that the balancing time is reduced as more rebates are offered in the auction since more workloads can be migrated. Figure 9 also shows that the VRA implementation consumes surplus energy faster than the URA implementation because it is solved to optimality, thus, migrating more workloads. Figure 10 shows the results of the two auction implementations for correcting energy shortage imbalances where similar

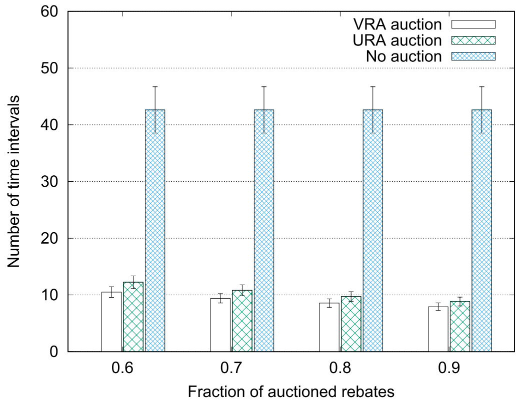

FIGURE 9. Balancing time for energy surplus.

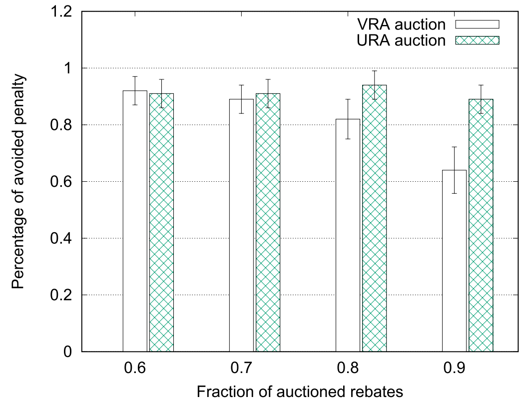

FIGURE 11. Imbalance penalty avoided (Surplus).

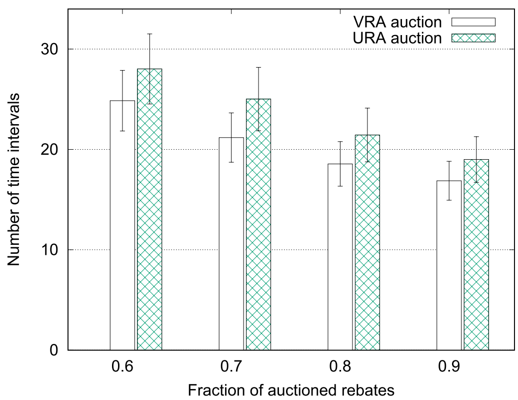

FIGURE 10. Balancing time for energy shortage.

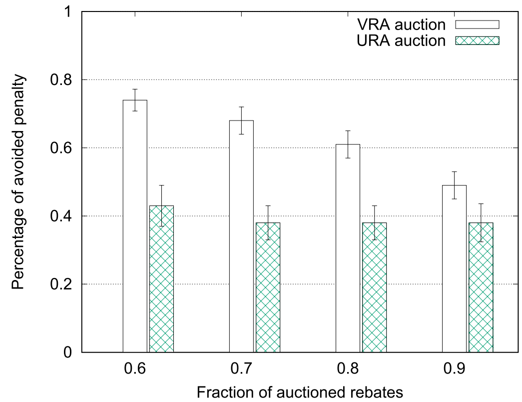

FIGURE 12. Imbalance penalty avoided (Shortage).

observations can be made in this case. No comparison is shown with the case of not using rebate auctions since balancing energy shortages requires a reduction of energy consumption, which can not happen without intervention. This is unlike the case of energy surplus imbalances, where excess energy is still consumed at its regular rate if no rebate auction is used.

C. EFFECT OF REBATE AUCTIONS ON THE PERCENTAGE OF IMBALANCE PENALTY AVOIDED

Regardless of the imbalance type and the required workloads migration direction (either towards or away from a certain datacenter), BAs that issue rebate auctions do so to avoid paying an imbalance penalty (such as the cost of turning on a backup generator). Therefore, imbalance penalties are considered the maximum budget of rebates that can be offered in a rebate auction because grid operators would not consider using rebates auctions to correct imbalances if it is cheaper to pay the imbalance penalty. This performance metric shows the savings that BAs can gain (shown as the percentage of imbalance penalty avoided) by using rebate auctions to correct energy imbalances instead of paying the penalty. The BAs' rationale for using the proposed rebate auction is that it should be cheaper to correct imbalances by using the auction than paying the penalty associated with the imbalance.

Figures 11 and 12 show the percentage of the remaining rebates budget (imbalance penalty avoided) after energy imbalances are corrected at different fraction levels of rebates availability (0.6, 0.7, 0.8, and 0.9). For the VRA implementation, Figures 11 and 12 show that more savings can be gained when smaller rebate availability fractions are used (fewer rebates are offered in the auction). This is because VRA can generate higher auction payments when fewer auctioned goods are available. On the other hand, the URA implementation generates the same level of savings at the different levels of rebate availability fractions because its uniform price rule charges all bidders using the same valuation (the valuation offered by the critical bid). The higher savings in the surplus case are due to the added revenue coming from the datacenters for consuming the excess energy. In contrast, no such revenue exists in the shortage case.

D. EFFECT OF REBATE AUCTIONS ON BROKER SAVINGS

This performance metric shows the effect of rebate auctions on brokers' savings. Broker savings are calculated as the percentage of reduction in brokers' costs as a result of taking part in the auctions and migrating their workloads. Figures 13 and 14 show the results of using the rebate auctions for both cases of energy surplus and shortage, respectively. In both cases,

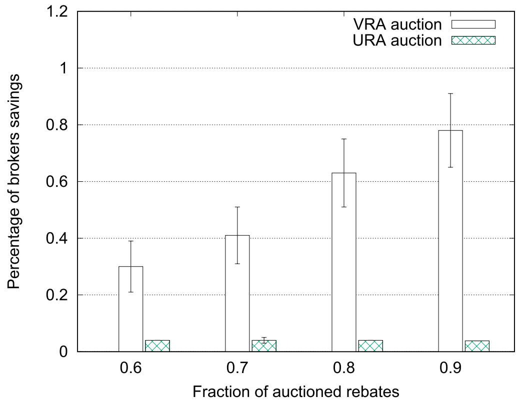

FIGURE 13. **Brokers savings (Surplus).**

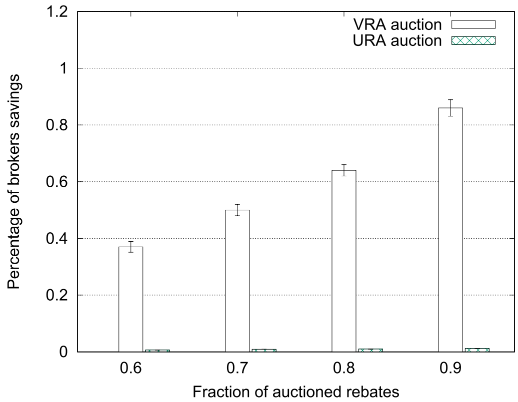

FIGURE 14. **Brokers savings (Shortage).**

brokers always save money by participating in the auction since their auction payments are always less than their cost at their "source datacenter" j. This is because their bids valuations v_{ijk} are always less than their requested rebates rr_{ijk}. Figures 13 and 14 show that the VRA implementation generates more savings for the brokers when more rebates are offered in the auction. This is because the more goods available in the auction, the lower the price paid. However, the URA generates fewer savings for the brokers because the accepted bids' valuations were close in value, and they were all charged the same price of the last accepted bid (the one that pays as bid, making no profit since no bidding strategy is used). As a result, they all see little benefit in total in this case. Therefore, this auction implementation highlights the need for participants to use a bidding strategy such as including a minimum amount of profit margin in their bids to maximize their return.

E. TIME COMPLEXITY

This performance metric shows the relative time complexity of the two rebate auction implementations at different problem input sizes. The VRA auction implementation was encoded as a MILP (Eqs. 6-16) and solved to optimality to maintain its truthfulness property. Such a setting is known to be NP-hard by a reduction from the Multidimensional Knapsack Problem (MKP) [52] given the capacity constraints along the different resource dimensions of the datacenters where workloads must be scheduled. On the other hand, the First-Fit algorithm used to schedule workloads in URA is a polynomial-time approximation algorithm. Figure 15 shows the running time of both auction implementations for different problem input sizes (the number of cloud brokers in the system). Results show that the VRA implementation does not scale well for large problem sizes and suggest the use of URA as the more viable option for large problem sizes.

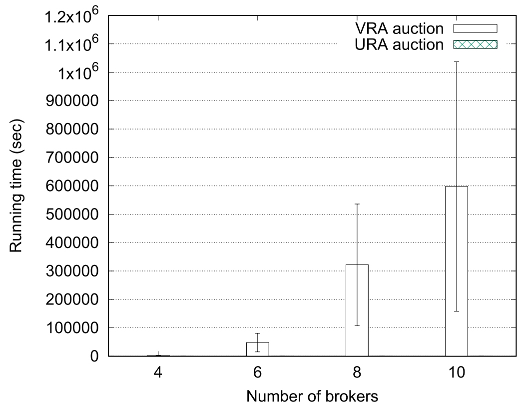

FIGURE 15. **Time complexity.**

VIII. CONCLUSION

Grid balancing is essential for the reliable operation of modern utility grids. Datacenters can play a major role in achieving effective grid balancing by providing demand-side flexibility using workload scheduling and migration. This flexibility is greatly needed to allow for increasing the integration of RESs into the generation mix of modern grids. This work presented a new rebate auction framework that uses cloud workload migrations between datacenters to provide bidirectional grid flexibility. Two alternative auction implementations were further provided under the proposed framework (based on the VCG and uniform price auctions), their properties were analyzed, and their performances were compared through simulation. Simulation results showed the effectiveness of the proposed rebate auctions in correcting energy imbalances and providing positive utility to all system participants.

REFERENCES

[1] M. Bestehorn and T. Borsche, "Balancing power consumption and production in smart grids," in *Proc. IEEE PES Innov. Smart Grid Technol., Eur.*, Oct. 2014, pp. 1–6.

[2] A. Zobaa, P. Ribeiro, S. A. Aleem, and S. Afifi, *Energy Storage at Different Voltage Levels: Technology, Integration, and Market Aspects*. London, U.K.: Institution Engineering Technology (IET), Sep 2018.

[3] O. M. Babatunde, J. L. Munda, and Y. Hamam, "Power system flexibility: A review," *Energy Rep.*, vol. 6, pp. 101–106, Feb. 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352484719309242

[4] M. Z. Degefa, I. B. Sperstad, and H. Sæle, "Comprehensive classifications and characterizations of power system flexibility resources," *Electric Power Syst. Res.*, vol. 194, May 2021, Art. no. 107022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S037877962100002X

[5] V. Mladenov, V. Chobanov, and A. Georgiev, "Impact of renewable energy sources on power system flexibility requirements," *Energies*, vol. 14, no. 10, p. 2813, May 2021. [Online]. Available: https://www.mdpi.com/1996-1073/14/10/2813

[6] R. Dorsey-Palmateer, "Effects of wind power intermittency on generation and emissions," *Electr. J.*, vol. 32, no. 3, pp. 25–30, Apr. 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1040619019300181

[7] K. De Vos, "Negative wholesale electricity prices in the German, French and Belgian day-ahead, intra-day and real-time markets," *Electr. J.*, vol. 28, no. 4, pp. 36–50, May 2015. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1040619015000652

[8] C. Brandstätt, G. Brunekreeft, and K. Jahnke, "How to deal with negative power price spikes?—Flexible voluntary curtailment agreements for large-scale integration of wind," *Energy Policy*, vol. 39, no. 6, pp. 3732–3740, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0301421511002795

[9] D. Quint and S. Dahlke, "The impact of wind generation on wholesale electricity market prices in the midcontinent independent system operator energy market: An empirical investigation," *Energy*, vol. 169, pp. 456–466, Feb. 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360544218323946

[10] M. F. Anjos and A. J. Conejo, "Unit commitment in electric energy systems," *Found. Trends Electr. Energy Syst.*, vol. 1, no. 4, pp. 220–310, 2017, doi: 10.1561/3100000014.

[11] A. Bhardwaj, V. K. Kamboj, V. K. Shukla, B. Singh, and P. Khurana, "Unit commitment in electrical power system—A literature review," in *Proc. IEEE Int. Power Eng. Optim. Conf.*, Melaka, Malaysia, Jun. 2012, pp. 275–280.

[12] F. N. Al Farsi, M. H. Albadi, N. Hosseinzadeh, and A. H. Al Badi, "Economic dispatch in power systems," in *Proc. IEEE 8th GCC Conf. Exhib.*, Feb. 2015, pp. 1–6.

[13] L. I. Dulau, M. Abrudean, and D. Bica, "Smart grid economic dispatch," *Proc. Technol.*, vol. 22, pp. 740–745, Jan. 2016, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2212017316000347

[14] J. Nanda, "Automatic generation control of an interconnected power system," *Proc. Inst. Elect. Eng.*, vol. 125, no. 5, pp. 385–390, May 1978. [Online]. Available: https://digital-library.theiet.org/content/journals/10.1049/piee.1978.0094

[15] K. Ullah, A. Basit, Z. Ullah, S. Aslam, and H. Herodotou, "Automatic generation control strategies in conventional and modern power systems: A comprehensive overview," *Energies*, vol. 14, no. 9, p. 2376, Apr. 2021. [Online]. Available: https://www.mdpi.com/1996-1073/14/9/2376

[16] G. Rancilio, A. Rossi, D. Falabretti, A. Galliani, and M. Merlo, "Ancillary services markets in Europe: Evolution and regulatory trade-offs," *Renew. Sustain. Energy Rev.*, vol. 154, Feb. 2022, Art. no. 111850. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032121011175

[17] P. Du, N. Mago, W. Li, S. Sharma, Q. Hu, and T. Ding, "New ancillary service market for ERCOT," *IEEE Access*, vol. 8, pp. 178391–178401, 2020.

[18] Q. Wang, C. Zhang, Y. Ding, G. Xydis, J. Wang, and J. Østergaard, "Review of real-time electricity markets for integrating distributed energy resources and demand response," *Appl. Energy*, vol. 138, pp. 695–706, Jan. 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261914010988

[19] M. H. Shoreh, P. Siano, M. Shafie-Khah, V. Loia, and J. P. S. Catalão, "A survey of industrial applications of demand response," *Electr. Power Syst. Res.*, vol. 141, pp. 31–49, Dec. 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378779616302632

[20] R. Deng, Z. Yang, M. Chow, and J. Chen, "A survey on demand response in smart grids: Mathematical models and approaches," *IEEE Trans. Ind. Informat.*, vol. 11, no. 3, pp. 570–582, Jun. 2015.

[21] X. Zhang, G. Hug, Z. Kolter, and I. Harjunkoski, "Industrial demand response by steel plants with spinning reserve provision," in *Proc. North Amer. Power Symp. (NAPS)*, Oct. 2015, pp. 1–6.

[22] H. Kondziella and T. Bruckner, "Flexibility requirements of renewable energy based electricity systems—A review of research results and methodologies," *Renew. Sustain. Energy Rev.*, vol. 53, pp. 10–22, Jan. 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032115008643

[23] C. Jennings, "As data center growth surges, so does the energy requirement," *Mil. Engineer*, vol. 109, no. 710, pp. 51–53, 2017. [Online]. Available: https://www.jstor.org/stable/26464625

[24] H. Rong, H. Zhang, S. Xiao, C. Li, and C. Hu, "Optimizing energy consumption for data centers," *Renew. Sustain. Energy Rev.*, vol. 58, pp. 674–691, May 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1364032115016664

[25] L. A. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, no. 12, pp. 33–37, Dec. 2007.

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

[27] A. Krioukov, P. Mohan, S. Alspaugh, L. Keys, D. Culler, and R. Katz, "NapSAC: Design and implementation of a power-proportional web cluster," in *Proc. 1st ACM SIGCOMM Workshop Green Netw.*, Aug. 2010, pp. 15–22.

[28] A. Abada, M. St-Hilaire, and W. Shi, "Auction-based scheduling of excess energy consumption to enhance grid upward flexibility," *IEEE Access*, vol. 10, pp. 5944–5956, 2022.

[29] M. Ghamkhari and H. Mohsenian-Rad, "Data centers to offer ancillary services," in *Proc. IEEE 3rd Int. Conf. Smart Grid Commun. (SmartGridComm)*, Nov. 2012, pp. 436–441.

[30] H. K. Nguyen, J. B. Song, and Z. Han, "Demand side management to reduce peak-to-average ratio using game theory in smart grid," in *Proc. IEEE INFOCOM Workshops*, Mar. 2012, pp. 91–96.

[31] 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, doi: 10.1145/2637364.2592004.

[32] X. Cao, J. Zhang, and H. V. Poor, "Data center demand response with on-site renewable generation: A bargaining approach," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2707–2720, Dec. 2018.

[33] H. Ma, V. Robu, N. Li, and D. Parkes, "Incentivizing reliability in demand-side response," in *Proc. 24th Int. Joint Conf. Artif. Intell.*, 2016, pp. 352–358. [Online]. Available: http://dl.acm.org/citation.cfm?id=3060621.3060671

[34] 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, doi: 10.1145/2567529.2567556.

[35] S. Kwon, L. Ntaimo, and N. Gautam, "Demand response in data centers: Integration of server provisioning and power procurement," *IEEE Trans. Smart Grid*, vol. 10, no. 5, pp. 4928–4938, Sep. 2019.

[36] Imdadullah, B. Alamri, M. A. Hossain, and M. S. J. Asghar, "Electric power network interconnection: A review on current status, future prospects and research direction," *Electronics*, vol. 10, no. 17, p. 2179, Sep. 2021. [Online]. Available: https://www.mdpi.com/2079-9292/10/17/2179

[37] T. Otsuki, A. B. M. Isa, and R. D. Samuelson, "Electric power grid interconnections in northeast Asia: A quantitative analysis of opportunities and challenges," *Energy Policy*, vol. 89, pp. 311–329, Feb. 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0301421515301993

[38] M. Greer, "Chapter 3—U.S. electric markets, structure, and regulations," in *Electricity Marginal Cost Pricing*, M. Greer, Ed. Boston, Massachusetts: Butterworth-Heinemann, 2012, pp. 39–100. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B978012385134500003X

[39] A. E. Brooks and B. C. Lesieutre, "A review of frequency regulation markets in three U.S. ISO/RTOs," *Electr. J.*, vol. 32, no. 10, Dec. 2019, Art. no. 106668. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1040619019302738

[40] A. Noureddine, R. Rouvoy, and L. Seinturier, "A review of energy measurement approaches," *ACM SIGOPS Operating Syst. Rev.*, vol. 47, no. 3, pp. 42–49, Nov. 2013, doi: 10.1145/2553070.2553077.

[41] A. Beloglazov and R. Buyya, "Adaptive threshold-based approach for energy-efficient consolidation of virtual machines in cloud data centers," in *Proc. 9th Int. Workshop Middleware Grids, Clouds E-Sci.*, New York, NY, USA, 2010, pp. 1–6, doi: 10.1145/1890799.1890803.

[42] L. Landberg, G. Giebel, H. A. Nielsen, T. Nielsen, and H. Madsen, ''Short-term prediction—An overview,'' *Wind Energy*, vol. 6, pp. 273–280, Jun. 2003.
[43] L. Hernandez, C. Baladron, J. M. Aguiar, B. Carro, A. J. Sanchez-Esguevillas, J. Lloret, and J. Massana, ''A survey on electric power demand forecasting: Future trends in smart grids, microgrids and smart buildings,'' *IEEE Commun. Surveys Tuts.*, vol. 16, no. 3, pp. 1460–1495, 3rd Quart., 2014.
[44] H. R. Varian and C. Harris, ''The VCG auction in theory and practice,'' *Amer. Econ. Rev.*, vol. 104, no. 5, pp. 442–445, May 2014. [Online]. Available: http://www.jstor.org/stable/42920977
[45] E. Markakis and O. Telelis, ''Uniform price auctions: Equilibria and efficiency,'' in *Algorithmic Game Theory*, M. Serna, Ed. Berlin, Germany: Springer, 2012, pp. 227–238.
[46] P. Milgrom, *Putting Auction Theory to Work* (Churchill Lectures in Economics). Cambridge, U.K.: Cambridge Univ. Press, 2004.
[47] B. Rieck, ''Basic analysis of bin-packing heuristics,'' 2021, *arXiv:2104.12235*.
[48] M. Hofri, *Bin Packing Heuristics*. New York, NY, USA: Springer, 1987, pp. 185–218, doi: 10.1007/978-1-4612-4800-2_5.
[49] C. A. Hoare, ''Quicksort,'' *Comput. J.*, vol. 5, no. 1, pp. 10–16, 1962.
[50] J. Burkett and K. Woodward, ''Uniform price auctions with a last accepted bid pricing rule,'' *J. Econ. Theory*, vol. 185, Jan. 2020, Art. no. 104954. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0022053119301048
[51] S. Bourjade, ''Uniform price auctions with asymmetric bidders,'' *BE J. Theor. Econ.*, vol. 19, no. 1, Jun. 2018, Art. no. 20160188, doi: 10.1515/bejte-2016-0188.
[52] J. Puchinger, G. R. Raidl, and U. Pferschy, ''The multidimensional knapsack problem: Structure and algorithms,'' *INFORMS J. Comput.*, vol. 22, no. 2, pp. 250–265, May 2010.

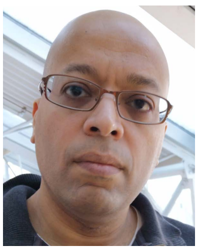

AHMED ABADA received the M.A.Sc. degree in electrical engineering from Carleton University, Ottawa, ON, Canada, where he is currently pursuing the Ph.D. degree. His research interests include renewable energy, power grid systems, cloud computing, computer networks, and optimization techniques.

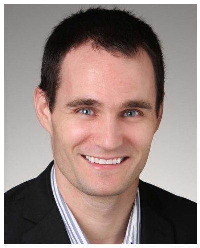

MARC ST-HILAIRE (Senior Member, IEEE) received the Ph.D. degree in computer engineering from Polytechnique Montréal, in 2006. He is currently a Professor with the School of Information Technology with a cross-appointment with the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON, Canada. He is conducting research on various aspects of wired and wireless communication systems. More precisely, he is interested in network planning and design, network architecture, mobile computing, and cloud computing. With more than 150 publications, his work has been published in several journals and international conferences. Over the years, he has received several awards, including the Carleton Faculty Graduate Mentoring Award, the Carleton Teaching Achievement Award, and several best paper awards. He is actively involved in the research community. In addition to serving as a member of technical program committees of various conferences, he is equally involved in the organization of several national and international conferences and workshops.

WEI SHI (Member, IEEE) received the Bachelor of Computer Engineering degree from the Harbin Institute of Technology (HIT) and the master's and Ph.D. degrees in computer science from Carleton University, Ottawa, ON, Canada. She is currently an Associate Professor with the School of Information Technology, cross appointed to the Department of Systems and Computer Engineering, Faculty of Engineering and Design, Carleton University. She is specialized in the design and analysis of fault tolerance algorithms addressing security issues in distributed environments, such as data-center networks, clouds, mobile agents and actuator systems, wireless sensor networks, as well as critical infrastructures, such as power grids and smart cities. She has also been conducting research in data privacy and big data analytics. She is also a Professional Engineer licensed in Canada.