Truthful, Practical and Privacy-Aware Demand Response in the Smart Grid via a Distributed and Optimal Mechanism

Georgios Tsaousoglou, Konstantinos Steriotis, Nikolaos Efthymiopoulos, Prodromos Makris, *Member, IEEE*, and Emmanouel Varvarigos

***Abstract*—High penetration of Renewable Energy Sources in modern smart grids necessitates the development of Demand Response (DR) mechanisms as well as corresponding innovative services for the emerging flexibility markets. From a game-theoretic perspective, the key requirements for a DR mechanism are: efficiency in terms of social welfare, practical applicability, scalability, privacy and incentive compatibility, in the sense of making it a dominant strategy for each user to act truthfully according to his/her real preferences, leaving no room for cheating. Previous works typically address only a subgroup of these requirements. In this paper, we propose a DR architecture, including a mechanism based on Ausubel's clinching auction and a communication protocol, that provably guarantee both efficiency and truthful user participation. Practicality/easiness of participation is enhanced via simple queries, while scalability and user privacy are preserved via a distributed implementation. Simulation results confirm the desired properties, while also showing that the truthfulness property becomes even more important in markets where participants are not particularly flexible.**

***Index Terms*—Demand response, auction, flexibility, mechanism design, blockchain, incentive compatibility, game theory.**

NOMENCLATURE

Setting

$\mathcal{N}$	Set of participating users
n	Number of participating users
i	Index of user
T	Set of timeslots
m	Number of timeslots in the horizon
t	Index of timeslot.

Flexibility Service Provider

L	Aggregated consumption of all users
D	Reduction of aggregated consumption
a, b	Parameters of the FSP's reward function
λ	Per unit price for consumption reduction.

User

q	Consumption reduction
$\widetilde{q}$	Optimal consumption reduction
ω	Inelasticity parameter.

Modified Clinching Auction

k, ξ	Indexes of iteration
ε	Price step
ζ	Clinching (allocation) of consumption reduction for a user.

Appliance Models

p	Power consumption
p_{max}	Power consumption upper limit
T	Temperature
η, θ	Parameters of temperature dynamics
γ	EV arrival timeslot (plug-in time)
E	Energy required by an EV
δ	Minimum charging duration of an EV.

Manuscript received February 26, 2019; revised August 19, 2019, November 8, 2019, and December 29, 2019; accepted December 29, 2019. Date of publication January 9, 2020; date of current version June 19, 2020. This work was supported by the European Union's Horizon 2020 Research and Innovation Programme through FLEXGRID Project under Grant 863876. Paper no. TSG-00268-2019. *(Corresponding author: Georgios Tsaousoglou.)*

The authors are with the Institute of Communication and Computer Systems, National Technical University of Athens, 157 80 Athens, Greece (e-mail: geotsaousoglou@mail.ntua.gr).

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

Digital Object Identifier 10.1109/TSG.2020.2965221

I. INTRODUCTION

IN MODERN smart grids, the need to dynamically balance supply and demand, has brought a great deal of attention to the idea of Demand Response (DR). When there is a need for reducing energy consumption in real-time, an ad-hoc market is created where the operator offers to buy consumption reduction from the users. An intermediate entity is assumed, with the task to clear the ad-hoc flexibility market. We refer to this entity as the Flexibility Service Provider (FSP). Each user (consumer) is equipped with a smart meter that measures his/her consumption at all times. However, the users' local functions (related to their flexibility/comfort levels and consumption habits) are private to each user. This makes the task of the FSP quite challenging, especially when we consider users who act strategically and may misinterpret their local function if that makes them better-off.

In real-time DR, expecting the user to manually control his/her appliances and bids in real-time is not realistic. Rather,

an intelligent agent sites on each user's side that controls the user's DR actions and makes the user's optimal bidding decisions while respecting the user's energy needs. Several studies have leveraged methods from Artificial Intelligence and proposed learning algorithms for optimizing the decisions of such an energy management agent (e.g., [1]). However, strategic agent behavior can compromise the efficiency of the DR mechanism. In order to protect the system's efficiency, a mechanism needs to be not only optimal, but also incentive compatible/truthful. This requirement is widely overlooked in the DR literature. In the few cases where truthfulness is addressed, it comes with a sacrifice of practical implementation ability and user privacy.

In Mechanism Design terms, a mechanism is Dominant Strategy Incentive Compatible (DSIC) (or, equivalently, satisfies the truthfulness property) when it is at each user's best interest to truthfully implement his/her true preferences, regardless of what other users do (see [2, Sec.10.2.2] for a more detailed analysis). In the vast majority of the DR literature users are typically modeled to myopically best-respond at each iteration of the pricing mechanism, i.e., they decide upon their preferred consumption upon receiving a price signal. As analyzed in [3], such myopic "local rationality" does not necessarily imply "global rationality", i.e., given an iterative mechanism, it is not always to the user's best interest to repeatedly best-respond. Rather, a user might be better-off by submitting false bids through the process, and such strategic behavior may compromise the mechanism's efficiency [4]. In other words, mechanisms that are not incentive compatible, are no longer optimal when strategic players are involved.

In this paper we also address this requirement, defined as the capability of the mechanism to provoke strategic users to act truthfully in accordance to their preferences, which is overlooked in most of the DR literature. Moreover, we do so via an indirect and practical mechanism, which allows for distributed, scalable and privacy-preserving implementation, in contrast to the few studies that consider incentive guarantees.

II. Related Work

In the DR literature, the end user is typically modeled as a selfish player who participates in the mechanism with the purpose of maximizing his/her own payoff. The user's preferences are widely modeled as a convex function (e.g., [5], [6], [7]). In [8] the electricity bill is minimized while the user's satisfaction is maintained above a defined threshold. In [9], a similar framework was built for deciding the charging times of EVs under forecasted prices. In [10], a spread is applied to the real-time prices in order to penalize deviations from a predefined schedule. In these studies, the bill of a user depends only on his/her own actions and it is disengaged from the actions of others. Thus, the users' DR actions might not be well coordinated.

In [11], the authors assume that consumers voluntarily provide their consumption preferences to a central entity, which optimizes the social welfare. Similarly, in [12], users estimate their energy needs and report them to an aggregator. In [13], a set of users enter into a direct-load-control contract with a load serving entity, responsible to satisfy a DR event. However, in [11]–[13], users were assumed to honestly reveal their consumption preferences.

In contrast to the studies presented so far, [5] and [6], considered users that do not reveal their local preferences, and the FSP controls their consumption indirectly by iteratively updating prices and observing the aggregated consumption. The authors use a dual decomposition method to discover the optimal prices. A scalable approach is proposed in [14], where smoothing techniques facilitate fast convergence. In [15], the aggregator is modeled as a profit-maximizing entity and a simulated annealing algorithm was adopted for the price optimization problem. The authors in [16] configure the pricing scheme with a forecast component. In [17], the authors consider two simple billing rules and prove that best-response dynamics converges to Nash Equilibrium. In [18] and [19], pricing schemes are deployed with the objective of maximizing the fairness of the consumption allocation. In [20], the effect of the FSP's profit policy on the DR outcome was examined. However, in the studies of this paragraph, users are assumed to truthfully best-respond to each price query, and thus they don't compromise the algorithm's properties.

In mechanism design terms, the above mechanisms are not incentive compatible, because a strategic user can benefit by manipulating his/her responses. Note that the optimality guarantees of the above studies, would no longer hold in the case of strategic users. When considering strategic users the mechanism designer is confronted with a trade-off: the Vickrey-Clarke-Groves (VCG) mechanism is the unique welfare maximizing mechanism implemented in dominant (and not best-response) strategies [21], meaning that either a VCG approach is taken (like in [22], [23]) or welfare maximization is compromised (like in [24], [25], [26], [27]).

The main problem with the direct-revelation VCG approaches [22], [23] is that they require users to reveal their whole set of preferences to the FSP, while the latter makes all the calculations and decides the allocation and the rewards. This is clearly impractical, since real users cannot compactly express their preferences in closed-form mathematical functions and even when they can, they are not happy to compromise their privacy. These issues were also reported in [28] and [29], where the authors proposed that the available actions for each user were restricted to a predefined set in order to simplify the message space. However, the authors in [28] do not model the effect of the user's actions on the price (similarly to [8], [9], [10]) and the authors in [29] consider EV-charging users who are only interested in the overall energy consumption over the horizon (which is not suitable for other loads and neither for en-route charging EVs). Naturally, these kinds of approximations result in loss of efficiency.

In this paper, we opt for a VCG-like approach, so as to achieve social welfare maximization, but we omit the direct-revelation approach of the typical VCG mechanism. Instead, we design an iterative auction mechanism based on Ausubel's clinching auction, in which users are only required to make decisions regarding their consumption in the presence of price signals. The convergence of the proposed method can be dramatically accelerated, with a minimal loss of efficiency for which we also prove a theoretical upper bound.

TABLE I
CLASSIFICATION OF LITERATURE BASED ON REQUIREMENTS

	Optimal	Truthful	Scalable	Privacy-aware
[11] - [13]	✔	×	×	×
[18]- [20], [27]	×	×	✔	×
[5], [6], [14]- [17]	✔	×	✔	×
[22]	✔	✔	×	×
[26], [28], [29]	×	✔	✔	×
[30] - [40]	✔	×	✔	✔
This work	✔	✔	✔	✔

By adopting this approach, we are guaranteed the efficient and incentive-compatible VCG outcome but also allow for a scalable, distributed implementation and a privacy-preserving communication protocol.

A distinct family of studies has elaborated on how the consumption measurements of an individual user can be masked in order to protect the user's privacy (e.g., [31]). In [30] a distributed authentication method is proposed, while [32] exploits hash functions in order to serve secure data transmission. Furthermore, [33] evolves load hiding techniques and [34] proposes obfuscation technologies towards data privacy. In [35] and [36], the authors propose privacy-preserving data aggregation methods with minimum overhead, while [37] also accounts for the case of a malicious FSP. Finally, [38] exploits adaptive key evolution, while [39] and [40] focus on the consensus problem towards reliable communication in fully distributed systems.

However, the studies of this family do not contribute to the design of the pricing scheme per se and assume that prices only depend on the aggregated consumption. This class of pricing rules can result in an optimal allocation under assumptions but it is not incentive compatible. The iterative mechanism proposed in this paper, can be implemented in configuration with a self-organized architecture that ensures privacy while in the same time is able to exploit the aforementioned systems in order to further enhance its level of security (in contrast to the direct VCG mechanism).

Summarizing the above, our proposed DR architecture: i) is suitable for a distributed implementation (unlike [22], [23]), ii) achieves the VCG outcome and does not sacrifice efficiency (unlike [24]–[27]), and iii) is incentive compatible (unlike [5], [6], [8]–[20] and [30]–[40]). In Table I, we present the four relevant requirements for a DR scheme and compare the proposed scheme to the state of the art.

III. SYSTEM MODEL

We consider a flexibility market comprised of an FSP and a set $\mathcal{N} \triangleq \{1, 2, \ldots, n\}$ of n self-interested consumers, hereinafter referred to as users. We also consider a discrete representation of time, where continuous time is divided into timeslots $t \in \mathcal{T}$ of equal durations s, where set $\mathcal{T} \triangleq \{1, 2, \ldots, m\}$ represents the scheduling horizon. Each user possesses a number of controllable appliances, with each appliance bearing an energy demand. If the consumptions of different appliances are decoupled (independent of each other) the appliances can participate in the DR mechanism virtually as different users. Thus, we can consider one appliance per user for ease of presentation.

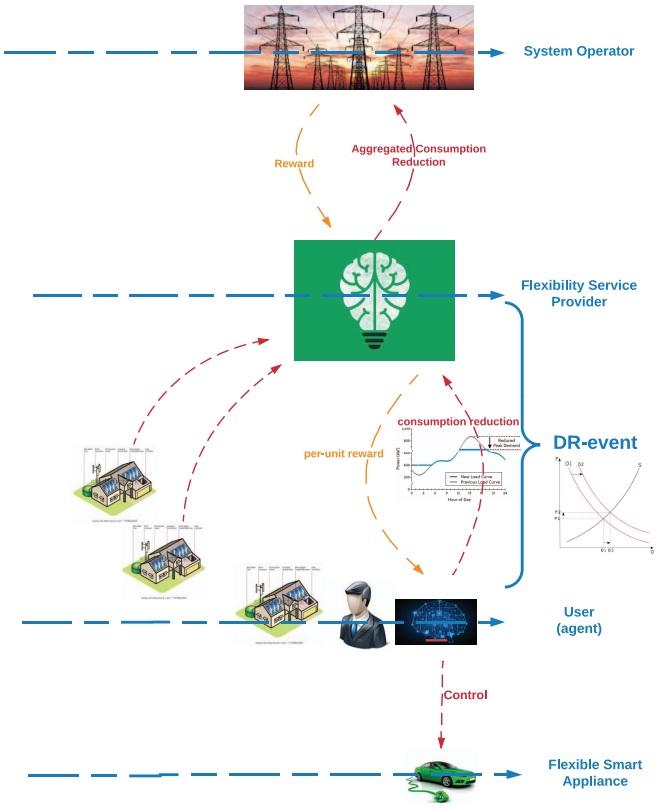

Fig. 1. System Architecture.

A. User's Consumption and Utility

The energy consumption of each user is measurable in real-time and can be shed upon request, in exchange for monetary compensation. Such a request for consumption modification is called a DR-event. In this paper we cope with real-time DR. Thus, the baseline consumption measurement is taken before the DR call, so the user cannot manipulate it since he/she does not know when a DR-event is going to occur.

The FSP takes on the task of providing the requested service by calling a DR-event among the users of its portfolio, which may include a micro-grid, or a local energy community [41], [42]. Upon a DR-event in timeslot t, the FSP offers a per-unit reward to the users for consumption reduction. User i can respond by reducing his/her consumption by a quantity q_i^t, assumed to be positive ($q_i^t \geq 0$), without loss of generality. The decisions for q can be taken by an intelligent agent (on behalf of the actual user and according to the user's preferences) in order to disengage the actual user from real-time participation. The system's architecture is depicted in Fig. 1. The consumption reduction q_i^t is characterized by its feasible set Q_i (defined by a set of constraints on q_i^t) and the discomfort function $d_i(q_i^t)$ of user i. The discomfort function is private to each user and expresses the minimum compensation in monetary units that a user requires, in order to reduce his/her consumption by the corresponding amount. We make the following assumptions on the form of function $d_i(q_i^t)$.

Assumption 1: Zero consumption reduction, brings zero discomfort to the user: $d_i(0) = 0$.

Assumption 2: The discomfort function is convex, so that additional increase of q_i^t brings increasing discomfort to the user:

$$q_i A^t \geq q_i B^t \Leftrightarrow d_i\big(q_i A^t + \varepsilon\big) - d_i\big(q_i B^t + \varepsilon\big) \geq d_i\big(q_i A^t\big) - d_i\big(q_i B^t\big), \forall \varepsilon, q_i A^t, q_i B^t > 0.$$

Detailed example appliance models (including operational constraints) are described in Section VII. The theoretical results to be presented in the following sections are valid for any user model that satisfies the assumptions above. A user's utility is defined as the difference between his/her discomfort and the reward $r_i(q_i^t)$ he/she receives:

$$U_i = \sum\nolimits_{t\in\mathcal{T}} \big(r_i\big(q_i^t\big) - d_i\big(q_i^t\big)\big) \tag{1}$$

In order to offer the rewards $r_i(q_i^t)$, the FSP draws on the reward offered by the operator requesting the reduction, as described in the following subsection.

B. DR-Event and the FSP

Let L^t denote the aggregated consumption of all users in $\mathcal{N}$, as seen by the operator, within a certain time interval t. The energy cost is modeled as a quadratic function of L^t, (like in [5]–[7], [14], [17], [22]):

$$C\big(L^t\big) = c_1 \cdot L^t + c_2 \cdot \big(L^t\big)^2$$

Upon a DR-event at t, let D^t denote the reduction from baseline consumption L_B^t to consumption $L_B^t - D^t$. The respective cost reduction is:

$$C\big(L_B^t\big) - C\big(L_B^t - D^t\big) = c_1 \cdot L_B^t + c_2 \cdot \big(L_B^t\big)^2 - c_1 \cdot \big(L_B^t - D^t\big) - c_2 \cdot \big(L_B^t - D^t\big)^2$$

which reads:

$$C\big(L_B^t\big) - C\big(L_B^t - D^t\big) = \big(c_1 + 2 \cdot c_2 \cdot L_B^t\big) \cdot D^t - c_2 \cdot \big(D^t\big)^2$$

We set $a = c_1 + 2 \cdot c_2 \cdot L_B^t$ and $b = c_2$. The cost benefit $C(L_B^t) - C(L_B^t - D^t)$ is denoted as a reward function $R(D^t)$:

$$R^t\big(D^t\big) = a \cdot D^t - b \cdot \big(D^t\big)^2, D^t \in \big[0, L_B^t\big], \tag{2}$$

where a, b are positive parameters with $a \geq 2bL_B^t$ so that it is an increasing function in the range of permitted values. The proposed DR architecture is open to any other choice of $R^t(D^t)$, provided it is an increasing and concave function. Thus, we assume that upon a DR-event, the operator offers a marginal per-unit reward for a reduction of D^t units:

$$\mu = \frac{d\big(R^t\big(D^t\big)\big)}{d(D^t)}. \tag{3}$$

IV. Problem Formulation

We would like to facilitate the allocation of consumption reduction among users so as to maximize social welfare. Social welfare is defined as the difference between the revenues $R^t(D^t)$ that the FSP receives from the operator, as defined in Eq. (2), and the sum of the discomforts caused to its users. This problem is formulated as:

$$\underset{q_i^t \in \mathcal{Q}_i, i\in\mathcal{N}}{\text{maximize}} \ \Big\{R^t\big(D^t\big) - \sum\nolimits_{i\in\mathcal{N}} d_i\big(q_i^t\big)\Big\} \tag{4}$$

$$\text{subject to} \ \ D^t = \sum\nolimits_{i\in\mathcal{N}} q_i^t. \tag{5}$$

Problem (4) is a convex optimization problem and could be solved efficiently if the local functions $d_i(q_i^t)$ were known (or truthfully disclosed). However, $d_i(q_i^t)$ of each user is not known and thus, problem (4) is typically solved via dual decomposition in the DR literature. In this approach, the FSP iteratively increases a per-unit reward λ asking the users their consumption reduction $q_i^t(\lambda)$ at each per-unit reward λ (auction query). At each iteration, each user i responds with his/her preferred $q_i^t(\lambda)$. A truthful (locally optimal) response by user i, denoted as $\widetilde{q_i^t}(\lambda)$, is the one that maximizes i's utility for reward λ. This is mathematically formulated as the solution to maximization problem (6):

$$\widetilde{q_i^t}(\lambda) = argmax_{q_i^t \in \mathcal{Q}_i, i\in\mathcal{N}} \big\{\lambda \cdot q_i^t - d_i\big(q_i^t\big)\big\} \tag{6}$$

Clearly, $\widetilde{q_i^t}(\lambda)$ is non-decreasing in λ, since $q_i^t \geq 0$. The auction terminates when λ reaches a value for which $\sum_{i\in\mathcal{N}} q_i^t(\lambda) = D^t(\lambda)$. The final price is called the market-clearing price and is denoted by λ_{mc}. The allocation at λ_{mc} is efficient if the users truthfully report their q_i^t at each query. However, truthful report may not be the best strategy for every user. To illustrate this, we present the following example:

Illustrative Example: Consider two users and a given timeslot t. User 1 operates a load of 10 kW while user 2 operates a 50 kW load. Their discomfort function is $d_i(q_i^t) = \omega_i \cdot (q_i^t)^2, i \in \{1, 2\}$, where their true flexibility parameters are $\omega_1 = \omega_2 = 0.1$. The reward function is $R^t(D^t) = 5 \cdot (D^t)$. Should they act according to their true discomfort function parameters, their utilities (given by Eq. (1)) at equilibrium would be $U_1 = U_2 = 4.875$. In case User 2 acts untruthfully according to $\omega_2^{fake} = 0.2$, his/her utility at equilibrium will be $U_2 = 7$. Therefore, the best strategy for User 2 is to be untruthful. ■

The previous *example* demonstrates how the market-clearing approach builds on the assumption that users behave myopically, by truthfully solving (6) at each iteration. The problem is that if we relax the truthfulness assumption and consider strategic users, market-clearing methods no longer result in efficient allocations. Thus, it is very important to design a mechanism that is not only efficient but also incentive compatible.

The Vickrey-Clarke-Groves (VCG) mechanism is the unique mechanism that is simultaneously truthful and efficient [21]. The VCG payment rule is the so called "Clarke pivot rule", which rewards each user i with an amount equal to the difference that i's presence makes in the welfare of other users. In the direct VCG mechanism, users are asked to declare their local functions $d_i(q_i^t)$ to the FSP (like in [22]). Because of the Clarke pivot rule, it is a dominant strategy for each user to make a truthful declaration [43]. In order to calculate the VCG rewards, problem 4 is solved $|\mathcal{N}| + 1$ times (one time with each user in $\mathcal{N}$ absent to calculate the payments, plus one time with all users present to calculate the allocation). The

major drawback of the direct VCG mechanism is the requirement that the users disclose their discomfort functions $d_i(q_i^t)$ to the FSP. This raises important issues such as privacy and difficulty of implementation. In the next section, we propose a modification of Ausubel's Clinching auction [44], allowing for a distributed implementation of VCG, designed to tackle these issues.

V. Ausubel's Clinching Auction for DR-Event Participation

The Clinching Auction (CA) is a well-known ascending price auction that halts when demand equals supply. However, in contrast to most auctions, allocation and rewards are not cleared exclusively at the final iteration. Rather, the goods (consumption reduction in our context) are progressively allocated as the auction proceeds and payments are also progressively built, while the auction design guarantees that the final allocation and payments coincide with the ones obtained through VCG. Thus, both allocation efficiency and incentive compatibility are achieved, while the aforementioned privacy and implementation drawbacks of the direct-VCG mechanism are effectively addressed.

In order for the CA to work in our setting, first we need to reverse the price trajectory. In the proposed Modified Clinching Auction (MCA), the FSP begins with a per-unit reward $\lambda = \lambda_{max}$ and in each iteration k the price λ^k is reduced by a small positive number ε. The size of ε adjusts the discretization level of MCA. By Eq. (3), reward λ_{max} is $\frac{d(R^t(0))}{d(\Delta L^t)} = a$, which, as analyzed in Section III, is the highest value possible given that R^t is concave. Users respond by bidding their preferred reduction $\widetilde{q_i^t}(\lambda)$ for each λ. We represent the user's response at λ as the solution to the user utility maximization problem (which is formally defined in Eq. (6) of the previous section).

The user's objective function is concave in q_i^t, since $\lambda \cdot q_i^t$ is linearly increasing and $d_i(q_i^t)$ is convex by Assumption 2. Also, the solution $\widetilde{q_i^t}$ is increasing in λ, which means that the user's response $\widetilde{q_i^t}$ gradually decreases as λ decreases. For the MCA, we relax constraint (5) to the inequality:

$$D^t \geq \sum\nolimits_{i\in\mathcal{N}} q_i^t \tag{7}$$

Consider an arbitrary iteration k of the MCA and let $D^t(\lambda^k)$ denote the operator's desired reduction for per-unit reward λ^k. The central idea of the MCA is the following: if there is a set $\mathcal{N}^j \subset \mathcal{N}$ for which we have

$$D^t\left(\lambda^k\right) - \sum\nolimits_{j\in\mathcal{N}^j} \widetilde{q_j^t}\left(\lambda^k\right) > 0 \tag{8}$$

then we allocate a reduction equal to $\zeta_i^k = D^t(\lambda^k) - \sum_{j\in\mathcal{N}^j} \widetilde{q_j^t}(\lambda^k)$ to each user $i \notin \mathcal{N}^j$ at a per-unit reward λ^k. We then say that user i "clinched" ζ_i^k units. The MCA auction terminates when set $\mathcal{N}^j$ that satisfies condition (8) and set $\mathcal{N}$, are equal, that is, constraint (7) is satisfied. After that, it allocates the remaining $D^t(\lambda^{(k-1)})$ proportionally to the users that bid in the second-to-last iteration. Algorithm 1 below, describes the proposed MCA.

Algorithm 1 MCA : Modified Clinching Auction

Initialize $\lambda^0 = \lambda_{max}$, $q_i^t(\lambda_{max})$, $D^t(\lambda_{max})$, $k = 0$
while $D^t(\lambda^k) < \sum_{i\in\mathcal{N}} \widetilde{q_i^t}(\lambda^k)$ **do**
  if there exists $\mathcal{N}^j : \sum_{j\in\mathcal{N}^j} \widetilde{q_j^t}(\lambda^k) < D^t(\lambda^k)$ **then**
    clinch units $\zeta_i^k = D^t(\lambda^k) - \sum_{j\in\mathcal{N}^j} \widetilde{q_j^t}(\lambda^k)$ for all $i \notin \mathcal{N}^j$ at per-unit reward λ^k
  else
    set $\lambda^{(k+1)} = \lambda^k - \varepsilon$ and $k = k + 1$
    ask each user a reduction query for λ^k and collect the responses $q_i^t(\lambda^k)$
    ask the operator for the desired total reduction $D^t(\lambda^k)$ at per-unit reward λ^k
  end if
end while
Clinch units $\zeta_i^k = \left(q_i^t(\lambda^{k-1}) - \sum_{\xi=0}^{k-1} \zeta_i^\xi\right) \cdot \frac{D^t(\lambda^{k-1})}{\sum_{i\in\mathcal{N}} q_i^t(\lambda^{k-1})}$ at per-unit reward λ^{k-1}, for each $i \in \mathcal{N}$

We are now in a position to prove the optimality of MCA in terms of social welfare performance:

Theorem 1: The social welfare loss at the final allocation of MCA is within $\frac{\varepsilon^2 + \lambda_{max}\cdot\varepsilon}{2b}$ of the maximum possible.

Proof: The Proof is given in Appendix A. ■

Since we cope with a real-time application, the trade-off between the mechanism's optimality and its computational time is of special importance. The latter directly relates to the price-step ε, which means that Theorem 1 gives a quantification of the trade-off described. In practice, for the relevant use cases of price-anticipating users (described in the introduction), the computational complexity of the MCA is small, which allows for a very small choice of ε. To emphasize this, it is useful to state the following corollary to Theorem 1.

Corollary 1: For $\varepsilon << 1$, the welfare loss grows linearly with ε.

Because the MCA includes a price-sensitive response also at the operator's side, we have to verify that the properties of efficiency and incentive compatibility still hold. This is proved in the following Propositions.

Proposition 1: Truthful bidding is a dominant strategy in MCA.

Proof: The proof is given in Appendix B. ■

Furthermore, the following properties of the VCG mechanism hold also for the MCA:

Proposition 2: MCA is individually rational, weakly budget-balanced, and achieves the maximum revenue for the FSP among all efficient and individually rational mechanisms.

Proof: The proof is given in Appendix C. ■

VI. Privacy-Preserving Distributed Implementation

In the MCA auction users are only required to respond to a specific sequence of queries, instead of communicating their discomfort function. Thus, each participant solves an optimization problem in parallel, while the mechanism still achieves the VCG outcome (and its nice properties). This allows the exploitation of blockchain services towards a DR

architecture which is not only efficient and truthful but also privacy-aware. In this section, we demonstrate how exactly the proposed optimal and incentive compatible mechanism can be configured with a scalable and privacy-preserving communication protocol that instantiates blockchain services. For this purpose, we exploit [45].

The proposed DR architecture exploits blockchain services [46], which are based on Distributed Hash Tables (DHT) [47] technologies, in order to execute MCA in a distributed fashion. In this context, users do not inform the FSP about their answers to the MCA's queries. Instead, the necessary aggregations are realized via a DHT, which is based on the scheme proposed by Kademlia [45]. Each user (node) is identified by a number (nodeID) that can be seen as a point in a specific virtual space. The nodeIDs do not serve only as identification, but they are also used for answering data base queries. Each piece of information is given as input to a hash function whose output belongs to the virtual space. Each node is responsible for a sub space of this virtual space according to its nodeID. Furthermore, participating nodes create and dynamically maintain routing tables in a bottom-up and self-organized way. Thus, they can collectively reach any point of this virtual space, by exploiting their routing tables, in order to store and get information. The distributed execution of MCA (DE-MCA) takes place through the following processes (see Algorithm 1 for the centralized version):

Process A - Data insertion: Each node i stores its bid $\widetilde{q_i^t}(\lambda^k)$ in another random node w through the use of the DHT system [47]. In more detail, i hashes its id and stores $\widetilde{q_i^t}(\lambda^k)$ and k in node w which is responsible for this id, based on the Kademlia architecture. It is highlighted that w is different for each i and k, as it is derived from the output of the hash function that Kademlia uses. This means that the set of nodes which are responsible for a specific data set is not determined from the data set owners. Thus, collusion of a relatively small number of malicious users to compromise privacy will fail.

Process B - Calculation of sums: The MCA algorithm requires the calculation of n+1 different sums at each iteration k. These are the n sums noted as $\sum_{j\in\mathcal{N}^j}\widetilde{q_j^t}(\lambda^k)$ (one for each user absent) and the sum $\sum_{i\in\mathcal{N}}\widetilde{q_i^t}(\lambda^k)$ of all user bids (see Algorithm 1). The proposed system exploits a tree structure and develops a distributed algorithm in order to calculate these sums. To do so, each node w that participates in the calculation requests from the subset of nodes in its routing tables, that dispose lower nodeID from it, to inform w on possible data values which they dispose in order to send them to w. The term "possible data values" refers to the aggregation of bids (up to w) for iteration k of MCA that is executed at that time instance. This process continues recursively until node *MAX*, which is the node with the highest id acquires the desirable data and calculates the sum $\sum_{i\in\mathcal{N}}\widetilde{q_i^t}(\lambda^k)$. At this point, this node also requests and receives $D^t(\lambda^k)$ and checks the termination condition of MCA. If the termination condition doesn't hold, *MAX* proceeds by broadcasting $\sum_{i\in\mathcal{N}}\widetilde{q_i^t}(\lambda^k)$ and $D^t(\lambda^k)$ to all nodes by using the aforementioned Kademlia tree.

Process C - Calculation of $\zeta_i^k(\lambda^k)$: Each node w calculates $\zeta_i^k(\lambda^k)$ by subtracting from the broadcasted sum $\sum_{i\in\mathcal{N}}\widetilde{q_i^t}(\lambda^k)$, the value that is stored in it. Note that this is not its own $\widetilde{q_i^t}(\lambda^k)$ value, and it doesn't know whose it is. If the result is negative then it sets $\zeta_i^k(\lambda^k)=0$.

Process D - Tuple update: It is highlighted here that in each iteration of MCA (e.g., the next iteration k+1) a different instance of Kademlia tree is created, so that $\zeta_i^{k+1}(\lambda^{k+1})$ is stored at a new node w^{k+1}, other than w^k. Thus, even in the case that a set of nodes are malicious, data privacy is not compromised. The tuple $A_i^k=\{\sum_{\xi=1}^k\zeta_i^\xi(\lambda^\xi),\lambda^\xi\cdot\sum_{\xi=1}^k\zeta_i^\xi(\lambda^\xi)\}$, containing the allocation and payments of user i until iteration k, is updated and passed from node w^k to w^{k+1}.

Process E - Final allocation and payments: At the final iteration, the updated tuples $A_i^{\varkappa}$ are communicated to the FSP. Note that the FSP receives only the final allocation and payments for each user, i.e., only the sum $\sum_{k=1}^{\varkappa}\zeta_i^k(\lambda^k)$ and not the intermediate values $\zeta_i^k(\lambda^k)$. In this way the aforementioned architecture ensures that the FSP and every other node that participates in the system do not have the data to estimate the local discomfort function $d_i(\cdot)$ of user i.

In Algorithm 2, the distributed execution of MCA is described. In case of a malicious FSP (i.e., with no hesitations to break the law), more strict privacy assumptions are needed, but this case is outside the scope of the present work. The interested reader can refer to the recent literature of privacy-preserving aggregation for the smart grid [30]–[38].

Since we cope with real-time DR it is important to note that, except for line 18 (distributed calculation of sum), all other operations of DE-MCA require constant time. Also, as is well known in DHTs, the latency of line 18 increases logarithmically with the number of users. This results in a very scalable implementation of the VCG mechanism, suitable for real-time DR. We also verify this property in the next section.

VII. User Models & Performance Demonstration

In this section, we present detailed appliance models taken from the literature and then use simulations to demonstrate the advantages of the MCA and verify its properties. We also compare MCA with the marginal cost pricing method [6] in terms of truthfulness and FSP's profits and with the direct-revelation VCG method [22] in terms of scalability. Simulations were run in MATLAB R2018b.

A. Detailed Appliance Models

The first model is taken from [6] and includes appliances that control the temperature of an environment, such as HVAC units. The user's most preferable temperature is denoted by parameter $T_i^{pref}(t)$ and was taken in our experiments to be uniformly distributed in the interval [75F, 79F]. The actual room temperature, denoted by $T_i^{in}(t)$, evolves according to

$$T_i^{in}(t)=T_i^{in}(t-1)+\eta\cdot\left[T^{out}(t)-T_i^{in}(t-1)\right]+\theta\cdot\left(p_{i,HVAC}^t-q_{i,HVAC}^t\right) \quad (9)$$

where $p_{i,HVAC}^t$ is the user's measurable power consumption before the DR-event occurrence and $q_{i,HVAC}^t$ is the curtailment resulting from the DR-event. Clearly, we have

$$p_{i,HVAC}^t-q_{i,HVAC}^t\geq 0, \quad (10)$$

Algorithm 2 DE-MCA : Distributed Execution of Modified Clinching Auction

FSP sets ε, initializes $\lambda^0 = \lambda_{max}$, $k = 0$, tuples $A_i^0 = \{0, 0\}$ and communicates them to all nodes
Each node i receives λ^0 and calculates $\widetilde{q}_i^t(\lambda^0)$
Instance 0 of Kademlia tree is created (its root is noted as MAX^0)
Data insertion (each node i puts $\widetilde{q}_i^t(\lambda^0)$ to instance 0 of Kademlia tree, see *Process A*)
Distributed calculation of $\sum_{i\in\mathcal{N}} \widetilde{q}_i^t(\lambda^k)$ which ends at node MAX^0 (see *Process B*)
Node MAX^0 requests $D^t(\lambda^k)$ from FSP and checks the termination condition
while $D^t(\lambda^k) < \sum_{i\in\mathcal{N}} \widetilde{q}_i^t(\lambda^k)$ **do**
  Node MAX^k broadcasts $\sum_{i\in\mathcal{N}} \widetilde{q}_i^t(\lambda^k)$ to all nodes
  Each node w^k in Kademlia tree, (where $w^k \neq i$) calculates $\zeta_i^k(\lambda^k)$ (see *Process C*)
  Node w^k updates the tuple A_i^k
  Node MAX^k sends to FSP signal to set $k = k + 1$
  FSP sets $k = k + 1$ and $\lambda^k = \lambda^{k-1} - \varepsilon$
  k^{th} Kademlia tree instance is created
  Each node w^{k-1} passes tuple A_i^{k-1} to node w^k (*Process D*)
  Node MAX^k communicates λ^k to all nodes
  User i receives λ^k and calculates $\widetilde{q}_i^t(\lambda^k)$
  Data insertion (each user i puts $\widetilde{q}_i^t(\lambda^k)$ to instance k of Kademlia tree, see *Process A*)
  Distributed calculation of $\sum_{i\in\mathcal{N}} \widetilde{q}_i^t(\lambda^k)$ at node MAX^k (see *Process B*)
  Node MAX^k requests $D^t(\lambda^k)$ from FSP and checks the termination condition
end while
MAX^k receives $D^t(\lambda^k)$ and broadcasts $\sum_{i\in\mathcal{N}} \widetilde{q}_i^t(\lambda^k)$ and $D^t(\lambda^k)$ to all nodes
Each node w^k calculates $\zeta_i^k(\lambda^k) = (\widetilde{q}_i^t(\lambda^k) - \sum_{\xi=0}^{k} \zeta_i^{\xi}) \cdot \frac{D^t(\lambda^k)}{\sum_{i\in\mathcal{N}} \widetilde{q}_i^t(\lambda^k)}$ and updates tuple A_i^k
All nodes communicate A_i^k to FSP (*Process E*)

and we also have the operational constraint

$$p_{i,HVAC}^t \leq p_{i,max}^t, \tag{11}$$

In the experiments $p_{i,max}^t$ was set to 5 kW. Outdoors temperature $T^{out}(t)$ was the same for all users and represented a typical summer day in Athens. Parameters η and θ where set to 0.9 and 3, respectively. The discomfort for such users was defined as the square difference between actual and desired temperatures:

$$d_i^{temp}(q_i^t) = \omega_{i,HVAC}^{temp} \cdot \left(T_i^{in}(t) - T_i^{pref}(t)\right)^2 \tag{12}$$

where parameter $\omega_{i,HVAC}^{temp}$ expresses the user's inelasticity in timeslot t and was randomly selected in the range [0.10, 0.50].

The second model represents temporally flexible loads (e.g., EVs) and is taken from [16]. The EV is plugged-in at timeslot γ_i (uniformly selected in the interval [3, 9], for one third of the users and in the interval [14, 20] for the remaining two thirds). Each EV charges at power $p_{i,EV}^t$ and has a total demand of $E_{i,EV}$ kWhs, where $E_{i,EV}$ was uniformly selected in the interval [6, 36]. The user wants the EV to be charged as soon as possible and any delay would bring discomfort. This model accurately represents en-route charging EVs. The desired charging duration, denoted as δ_i, was set to $\delta_i = 3$ timeslots for all users. The upper power limit $p_{i,max}^t$, was selected as $p_{i,max}^t = \frac{E_{i,EV}}{\delta_i}$. That is, if no DR-events occurred, each user would charge his/her EV in 3 consequent timeslots. An EV operational constraint is given as

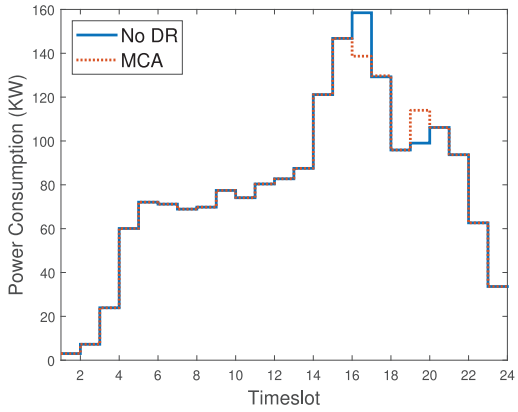

Fig. 2. Aggregated consumption as a function of time with and without the DR-event.

$$p_{i,EV}^t \leq p_{i,max} \tag{13}$$

The EV cannot be charged before arrival:

$$p_{i,EV}^t = 0, t < \gamma_i \tag{14}$$

and must be fully charged before leaving:

$$\sum\nolimits_{t\in\mathcal{T}} p_{i,EV}^t \geq E_{i,EV}. \tag{15}$$

During a DR-event a user may choose to curtail $q_{i,EV}^t$ units and shift charging to a later timeslot. This delayed charging (for timeslots after $\gamma_i + \delta_i - 1$), comes with a discomfort defined as

$$d_{i,EV}^{wait}(q_{iEV}^t) = \sum\nolimits_{t\in\{\mathcal{T}|t\geq\gamma_i+\delta_i\}} \left[\left(\omega_{i,EV}^{wait}\right)^{t-\gamma_i-\delta_i+1} \cdot p_{i,EV}^t\right] \tag{16}$$

where parameter $\omega_{i,EV}^{wait}$ expresses the user's inelasticity and was uniformly selected in [1, 1.5].

B. Simulation Results

Over a time horizon of 24 timeslots, with a duration of 15 minutes for each timeslot and for a setting of 50 users, we simulated a DR event in timeslot 17 (where there was a peak in the aggregated consumption). The parameters of the reward function were set to $a = 3$ and $b = 0.05$. We used step $\varepsilon = 10^{-3}$ in the MCA algorithm. Fig. 2 depicts the aggregated consumption along all 24 timeslots. As the figure shows, there is a consumption curtailment in timeslot 17 and a consequent shift of consumption to timeslot 20. Note that it could not be shifted to timeslots 18 or 19 because constraints (11) and (13) where already tight for these timeslots.

Next, we investigate the effect that cheating has on the FSP's profits, denoted by $\Pi^{truthful}$ for the case where users

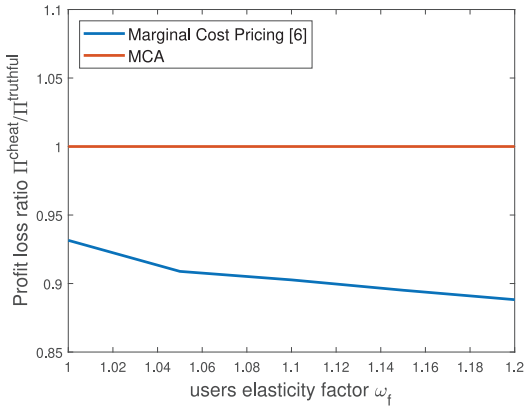

Fig. 3. Ratio $\Pi^{cheat}/\Pi^{truthful}$ as a function of ω_f.

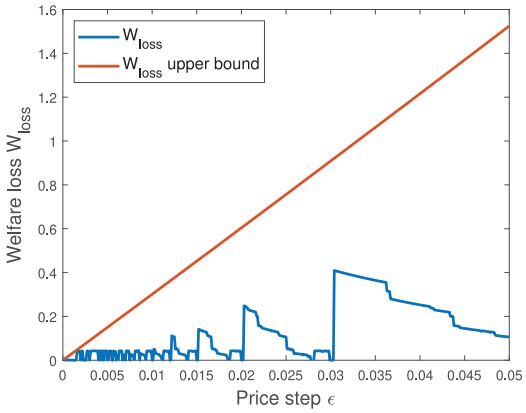

Fig. 4. Proportional welfare loss of MCA as a function of the price step ε.

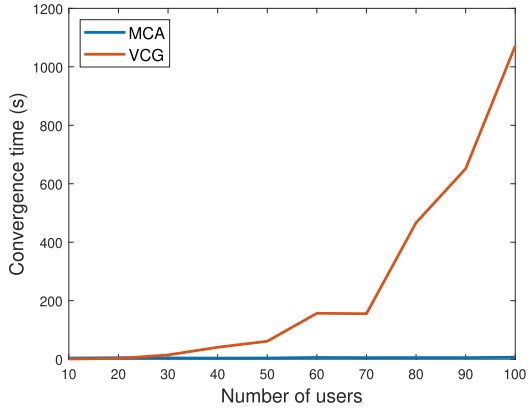

Fig. 5. Convergence time of MCA and VCG, as a function of the number of users.

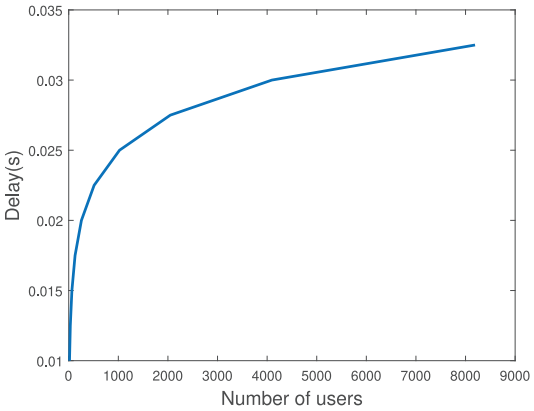

Fig. 6. Delay (latency) of privacy preserving protocol as a function of the number of participating users.

act truthfully and by Π^{cheat} for the case where they act according to what brings them the highest utility. We plot the ratio $\Pi^{cheat}/\Pi^{truthful}$ for different values of users' elasticities $\{\omega^{temp}_{i,HVAC}, \omega^{wait}_{i,EV}\}$. To do so, for each experiment we multiply the users' elasticity parameters by a positive factor ω_f. Higher values of ω_f indicate more inelastic users. Fig. 3 shows that the ratio $\Pi^{cheat}/\Pi^{truthful}$ is maximized and is equal to 1 for the MCA, verifying our theoretical results. For the marginal cost pricing method [6], the FSP's profit loss due to untruthfulness rises with ω_f (i.e., when users are less elastic), indicating that our scheme's truthfulness property becomes more important in markets where participants are not particularly flexible.

Next, we simulated the DR-event for different values of the step ε, measuring the proportional welfare loss

$$W_{loss} = \frac{W_{opt} - W_{MCA}}{W_{opt}}, \tag{17}$$

where W_{opt} is the optimal welfare and W_{MCA} is the welfare achieved by the MCA. The simulation results in Fig. 4 verify Corollary 1, which states that for small values of ε the upper bound on the welfare loss grows linearly with ε.

Finally, we compare MCA to the direct-revelation VCG method (proposed in [22]), in terms of scalability with respect to the number of users. Simulations are carried out on an Intel Core i7 4GHz, 64-bit, 16GB RAM, computer. The computational time of the method in [22] rises very quickly, which makes it impractical for real-time applications. In contrast, MCA scales remarkably well to any number of users, since the algorithm's convergence time does not depend on the size of set $\mathcal{N}$. In order to evaluate the latency of DE-MCA we assumed that the data network introduces a network delay (between any two nodes) that follows a uniform distribution between 5ms and 15ms. Fig. 6 depicts the latency introduced by DE-MCA. It is defined as the total time overhead that the proposed distributed implementation introduces due to data network delays between any two data network nodes. As it is known theoretically, in DHT technologies, this latency increases logarithmically with the number of users. This is verified in Fig. 6. In comparison to the timeslot duration (e.g., 15 minutes, which is a typical granularity for measurements and clearing of the balancing market), these results show that the proposed system is both scalable and efficient.

C. Incentive Compatibility and the Case of Inelastic Appliances

In this subsection we discuss the property of incentive compatibility. We verify the theoretical result of Proposition 1 and also experimentally study incentive compatibility in the case of inflexible appliances (where our Assumptions are not satisfied). We assume that one user misinterprets his/her discomfort by manipulating his/her $\omega^{wait}_{i,EV}$. The untruthful user is indexed by ch (for cheater). The cheater's utility U_{ch} is maximized for a certain choice of ω_{ch}. Fig. 7 shows U_{ch} as a function of ω_{ch}.

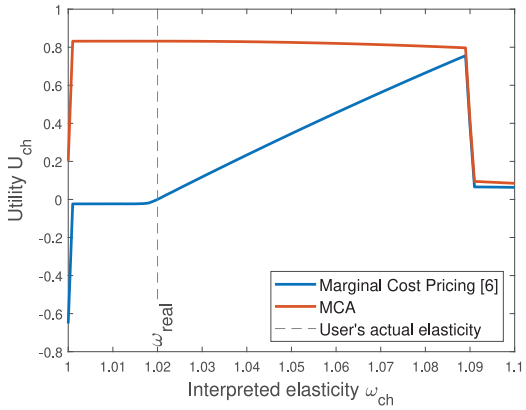

Fig. 7. Focal user's utility as a function of ω_{ch}.

The black vertical line represents the focal user's real $\omega_{i,EV}^{wait}$, denoted as ω_{real}. For the MCA, the user's optimal choice of ω (where U_{ch} is maximized) coincides with ω_{real}, which means that the user's best strategy is to act truthfully, in contrast to the marginal cost pricing method.

The result of Fig. 7 was expected, since it was already proven in Proposition 1. Although we cannot state a similarly strong theoretical guarantee for inflexible users, nevertheless our simulations show similar results. We study the case where an appliance is inflexible (Type 2 appliances of [14]), in the sense that it can only be turned on or off, but its consumption cannot take intermediate values:

$$p_{i,inel}^t \in \{0, p_{i,max}^t\}, \tag{18}$$

and thus,

$$q_{i,inel}^t \in \{0, p_{i,max}^t\}. \tag{19}$$

The user's discomfort for turning his/her load off, is denoted by $d_{i,inel}^{off}$. Thus, the user's discomfort function takes the form:

$$d_{i,inel}(q_{i,inel}^t) = \begin{cases} 0, & q_{i,inel}^t = 0 \\ d_{i,inel}^{off}, & q_{i,inel}^t > 0 \end{cases} \tag{20}$$

This kind of appliances violate Assumption 2. In fact, the form of the user's valuation exhibits complementarity (the user can either curtail all $p_{i,max}^t$ KWhs, but cannot make use of an allocation that is smaller than $p_{i,max}^t$). In the presence of such complementarities there is no tractable iterative auction that can achieve incentive compatibility [48]. We present an extension of MCA that accommodates inflexible users and evaluate it. Although we can no longer guarantee incentive compatibility, nevertheless simulation results show that, in practice, truthful bidding is still the best choice for each user.

Let $\mathcal{I}$ denote the set of inflexible users. Also, let $\varkappa$ denote the number of iterations until the auction halts (see Appendix A). The first step is to run the MCA algorithm. Then, we grant the MCA allocations $\zeta_i^k, \forall k$ only to elastic users $i \notin \mathcal{I}$. The remaining reduction $\sum_{k=1}^{\varkappa} \sum_{i \in \mathcal{I}} \zeta_i^k$, will be reallocated amongst the inflexible users, in a way that respects constraints (19). This is an instance of the knapsack problem. In order not to compromise the computational time guarantees of our real-time auction, we use a simple heuristic to solve it.

Algorithm 3 Extended Modified Clinching Auction

1: Run the MCA algorithm
2: set $q_{i,inel}^t = 0, \forall i \in \mathcal{I}$
3: sort users $i \in \mathcal{I}$, in increasing order of $d_{i,inel}^{off}/p_{i,max}$
4: set $q_{i,inel}^t = p_{i,max}$, for user $i \in \mathcal{I}$, in increasing order of the sorted list, until $\sum_{i \in \mathcal{I}} q_{i,inel}^t \geqslant \sum_{k=1}^{\varkappa} \sum_{i \in \mathcal{I}} \zeta_i^k$.

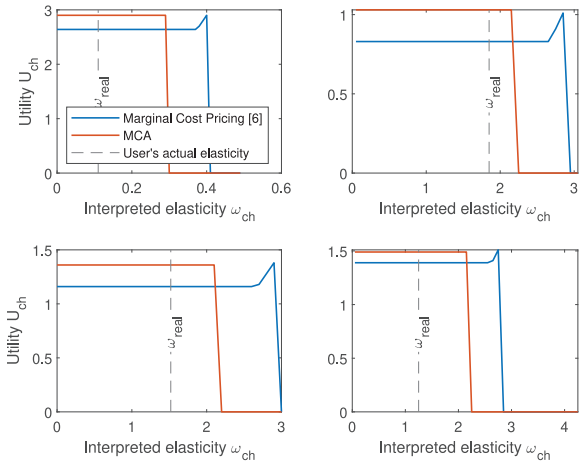

Fig. 8. Users' Utility as a function of user's interpreted elasticity ω.

Inflexible users are sorted in increasing order of their "bang-for-buck" i.e., their $d_{i,inel}^{off}/p_{i,max}^t$. We allocate $q_{i,inel}^t = p_{i,max}^t$ to user $i \in \mathcal{I}$, in increasing order of the sorted list, until $\sum_{i \in \mathcal{I}} q_{i,inel}^t \geq \sum_{k=1}^{\varkappa} \sum_{i \in \mathcal{I}} \zeta_i^k$. The procedure is depicted in Algorithm 3.

In Fig. 8, we present indicative results for various values of $p_{i,max}$ and $d_{i,inel}^{off}$, regarding truthfulness in the extended MCA. More specifically, we tested how well a user does (in terms of utility U_i, see Eq. (6)), by interpreting his/her elasticity with various (untruthful) values ω_{ch}. The user's actual discomfort for curtailing $p_{i,max}$ units is marked with a vertical dotted line. From the figure, it becomes clear that the user already achieves his/her maximum possible utility, by truthfully interpreting his/her discomfort and has nothing to gain by playing untruthfully. This is, again, in contrast to the marginal cost pricing approach [6].

VIII. Conclusion

In this paper we considered a setting of strategic, intelligent users and an FSP seeking to incentivize them in order to curtail part of their consumption in response to a real-time DR-event. We showcased the inefficiency of previous state-of-the-art approaches, which either do not consider user incentives, or adopt a direct-revelation approach, respectively leading to either lack of truthfulness and consequent inefficiency, or to lack of privacy and scalability. To overcome these shortcomings, we presented a novel iterative auction mechanism that implements the truthful and efficient VCG outcome but also allows for a distributed implementation and a privacy-preserving communication protocol. For this purpose, we identified a suitable mechanism from the field of mechanism design (Ausubel's Clinching auction) and modified

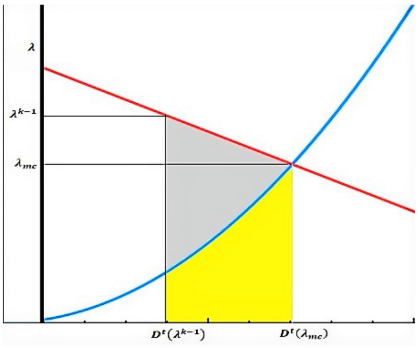

Fig. 9. $D^t(\lambda)$ and $\sum_{i\in\mathcal{N}} q_i^t(\lambda)$ as a function of λ.

it for a setting with continuous items (energy), where the number of items is not fixed but bears a cost function (reward function in our case). We extended the properties of the original mechanism to the modified mechanism that we propose. Furthermore, we proved a theoretical upper bound for the efficiency loss caused by an increase in the price step, which relates to a decrease in the mechanism's computational time.

Our theoretical and simulation results verified that the proposed scheme combines the desired properties with very good performance and small overhead. Moreover, we also implemented appliances that can only be turned on or off without the ability to adjust their power consumption. In our simulations, the proposed mechanism was shown to work very well in the DR setting, even with these types of loads, despite the fact that they do not satisfy the necessary assumptions.

Appendix A
Proof of Theorem 1

The value of λ at which $D^t = \sum_{i\in\mathcal{N}} \widetilde{q_i^t}$ is denoted as λ_{mc}, which gives

$$D^t(\lambda_{mc}) = \sum\nolimits_{i\in\mathcal{N}} \widetilde{q_i^t}(\lambda_{mc}). \tag{21}$$

Let $\varkappa$ denote the number of iterations until the auction halts, that is,

$$\varkappa = \left\lceil \frac{\lambda_{max} - \lambda_{mc}}{\varepsilon} \right\rceil, \tag{22}$$

where $\lceil\cdot\rceil$, denotes the rounding to the nearest integer above. We have

$$\left\lceil \frac{\lambda_{max} - \lambda_{mc}}{\varepsilon} \right\rceil \geq \varkappa \geq 1 + \left\lceil \frac{\lambda_{max} - \lambda_{mc}}{\varepsilon} \right\rceil. \tag{23}$$

After the last clinchings (line 11 of Algorithm 1) we have efficiently allocated $D^t(\lambda^{\varkappa-1})$ reduction units to the users. The remaining $D^t(\lambda_{mc}) - D^t(\lambda^{\varkappa-1})$ are not allocated and this causes the loss of welfare W_{loss}, which is depicted as the grey area in Fig. 9, where the red line represents $D^t(\lambda)$ and the blue line represents $\sum_{i\in\mathcal{N}} \widetilde{q_i^t}(\lambda)$.

Since we remain agnostic of the closed form of $\sum_{i\in\mathcal{N}} \widetilde{q_i^t}(\lambda^k)$, we assume the worst case and calculate an upper bound on the sum of the gray plus the yellow area:

$$W_{loss} \geq \lambda_{mc} \cdot \left(D^t(\lambda_{mc}) - D^t\left(\lambda^{\varkappa-1}\right)\right) + \frac{1}{2}\left(\lambda^{\varkappa-1} - \lambda_{mc}\right) \cdot \left(D^t(\lambda_{mc}) - D^t\left(\lambda^{\varkappa-1}\right)\right). \tag{24}$$

By substituting $D^t(\lambda) = \frac{a-\lambda}{2b}$, from Eq. (3), we get

$$W_{loss} \geq \frac{\lambda_{mc} \cdot \left(\lambda^{\varkappa-1} - \lambda_{mc}\right)}{4b} + \frac{\lambda^{\varkappa-1} \cdot \left(\lambda^{\varkappa-1} - \lambda_{mc}\right)}{4b} \geq \frac{\left(\lambda^{\varkappa-1}\right)^2 - (\lambda_{mc})^2}{4b}. \tag{25}$$

By further substituting $\lambda^{\varkappa-1} = \lambda_{max} - \varepsilon \cdot (\varkappa - 1)$ and also substituting $\varkappa$, using the left inequality when $\varkappa$ appears with a minus sign and the right inequality when it appears with a plus sign, we finally obtain

$$W_{loss} \geq \frac{\varepsilon^2 + \lambda_{max} \cdot \varepsilon}{2b}. \tag{26}$$

Appendix B
Proof of Proposition 1

Fix an iteration k and assume that user i bids $q_{i,false}^t(\lambda^k) \neq \widetilde{q_i^t}(\lambda^k)$ in that iteration. From step 4 of MCA, we see that ζ_i^k does not depend on q_i^t but only on the other users' bids $q_j^t, j \neq i$. Thus, user i's bid can affect i's allocation only by changing the λ at which the termination condition holds. This means that a false bid $q_{i,false}^t(\lambda^k)$ will make a difference to i, only if k is the last iteration. However, by definition of $\widetilde{q_i^t}(\lambda^k)$ (see Eq. (6)), any bid $q_{i,false}^t(\lambda^k) \neq \widetilde{q_i^t}(\lambda^k)$ brings strictly lower utility to user i at any iteration k. Thus, truthful bidding brings the highest utility to user i.

Appendix C
Proof of Proposition 2

The MCA auction is welfare maximizing (by Theorem 1, for ε small enough) and truthful (by Proposition 1). Moreover, the class of VCG mechanisms is the unique class that simultaneously achieves these two properties [21]. Since MCA terminates with the VCG allocation and payments, it inherits the property of individual rationality.

Regarding the weak budget balance property, it suffices to show that our setting exhibits the no single-agent effect [2] 10.4.4. This is true if the aggregated utility of n–1 users does not improve by adding a n-th user to the system. This property holds in single-sided auctions with monotonous preferences, since dropping a user only reduces the competition for the remaining users, thus making them better-off. Moreover, by [49], the VCG mechanism maximizes the auctioneer's utility, which means that the FSP buys flexibility units from the users at the lowest possible price (among all efficient and individually rational mechanisms).

References

[1] F. Ruelens, B. J. Claessens, S. Vandael, B. De Schutter, R. Babuška, and R. Belmans, "Residential demand response of thermostatically controlled loads using batch reinforcement learning," *IEEE Trans. Smart Grid*, vol. 8, no. 5, pp. 2149–2159, Sep. 2017.

[2] Y. Shoham and K. Leyton-Brown, *Multiagent Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2009.

[3] N. Nisan, M. Schapira, G. Valiant, and A. Zohar, "Best-response mechanisms," in *Proc. ICS*, 2011, pp. 155–165.

[4] R. Johari, S. Mannor, and J. N. Tsitsiklis, "Efficiency loss in a network resource allocation game: The case of elastic supply," *IEEE Trans. Autom. Control*, vol. 50, no. 11, pp. 1712–1724, Nov. 2005.

[5] P. Samadi, A. H. Mohsenian-Rad, R. Schober, V. W. S. Wong, and J. Jatskevich, "Optimal real-time pricing algorithm based on utility maximization for smart grid," in *Proc. 1st IEEE Int. Conf. Smart Grid Commun.*, Gaithersburg, MD, USA, 2010, pp. 415–420.

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

[7] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis, "The role of aggregators in smart grid demand response markets," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 7, pp. 1247–1257, Jul. 2013.

[8] S. Althaher, P. Mancarella, and J. Mutale, "Automated demand response from home energy management system under dynamic pricing and power and comfort constraints," *IEEE Trans. Smart Grid*, vol. 6, no. 4, pp. 1874–1883, Jul. 2015.

[9] Z. Wang and R. Paranjape, "Optimal residential demand response for multiple heterogeneous homes with real-time price prediction in a multiagent framework," *IEEE Trans. Smart Grid*, vol. 8, no. 3, pp. 1173–1184, May 2017.

[10] G. Tsaousoglou, P. Makris, and E. Varvarigos, "Electricity market policies for penalizing volatility and scheduling strategies: The value of aggregation, flexibility, and correlation," *Sustain. Energy Grids Netw.*, vol. 12, pp. 57–68, Dec. 2017.

[11] M. Ahmadi, J. M. Rosenberger, W. J. Lee, and A. Kulvanitchaiyanunt, "Optimizing load control in a collaborative residential microgrid environment," *IEEE Trans. Smart Grid*, vol. 6, no. 3, pp. 1196–1207, May 2015.

[12] D. T. Nguyen and L. B. Le, "Joint optimization of electric vehicle and home energy scheduling considering user comfort preference," *IEEE Trans. Smart Grid*, vol. 5, no. 1, pp. 188–199, Jan. 2014.

[13] O. Erdinç, A. Taşcikaraoğlu, N. G. Paterakis, Y. Eren, and J. P. S. Catalão, "End-user comfort oriented day-ahead planning for responsive residential HVAC demand aggregation considering weather forecasts," *IEEE Trans. Smart Grid*, vol. 8, no. 1, pp. 362–372, Jan. 2017.

[14] S. Mhanna, A. C. Chapman, and G. Verbič, "A fast distributed algorithm for large-scale demand response aggregation," *IEEE Trans. Smart Grid*, vol. 7, no. 4, pp. 2094–2107, Jul. 2016.

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

[16] A. H. Mohsenian-Rad and A. Leon-Garcia, "Optimal residential load control with price prediction in real-time electricity pricing environments," *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 120–133, Sep. 2010.

[17] P. Jacquot, O. Beaude, S. Gaubert, and N. Oudjane, "Demand response in the smart grid: The impact of consumers temporal preferences," in *Proc. IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, Dresden, Germany, 2017, pp. 540–545.

[18] G. Tsaousoglou, N. Efthymiopoulos, P. Makris, and E. Varvarigos, "Personalized real time pricing for efficient and fair demand response in energy cooperatives and highly competitive flexibility markets," *J. Mod. Power Syst. Clean Energy*, vol. 7, no. 1, pp. 151–162, Jan. 2019.

[19] K. Steriotis, G. Tsaousoglou, N. Efthymiopoulos, P. Makris, and E. Varvarigos, "A novel behavioral real time pricing scheme for the active energy consumers' participation in emerging flexibility markets," *Sustain. Energy Grids Netw.*, vol. 16, pp. 14–27, Dec. 2018.

[20] K. Steriotis, G. Tsaousoglou, N. Efthymiopoulos, P. Makris, and E. Varvarigos, "Development of real time energy pricing schemes that incentivize behavioral changes," in *Proc. IEEE Int. Energy Conf. (ENERGYCON)*, Limassol, Cyprus, 2018, pp. 1–6.

[21] J. Green and J. J. Laffont, "Characterization of strongly individually incentive compatible mechanisms for the revelation of preferences for public goods," *Econometrica*, vol. 45, no. 2, pp. 427–438, 1977.

[22] P. Samadi, A. H. Mohsenian-Rad, R. Schober, and V. W. S. Wong, "Advanced demand side management for the future smart grid using mechanism design," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1170–1180, Sep. 2012.

[23] E. Nekouei, T. Alpcan, and D. Chattopadhyay, "Game-theoretic frameworks for demand response in electricity markets," *IEEE Trans. Smart Grid*, vol. 6, no. 2, pp. 748–758, Mar. 2015.

[24] N. Yaagoubi and H. T. Mouftah, "User-aware game theoretic approach for demand management," *IEEE Trans. Smart Grid*, vol. 6, no. 2, pp. 716–725, Mar. 2015.

[25] J. Ma, J. Deng, L. Song, and Z. Han, "Incentive mechanism for demand side management in smart grid using auction," *IEEE Trans. Smart Grid*, vol. 5, no. 3, pp. 1379–1388, May 2014.

[26] A. C. Chapman and G. Verbič, "An iterative on-line auction mechanism for aggregated demand-side participation," *IEEE Trans. Smart Grid*, vol. 8, no. 1, pp. 158–168, Jan. 2017.

[27] G. Tsaousoglou, K. Steriotis, N. Efthymiopoulos, K. Smpoukis, and E. Varvarigos, "Near-optimal demand side management for retail electricity markets with strategic users and coupling constraints," *Sustain. Energy Grids Netw.*, vol. 19, Sep. 2019, Art. no. 100236.

[28] S. Li, W. Zhang, J. Lian, and K. Kalsi, "Market-based coordination of thermostatically controlled loads—Part I: A mechanism design formulation," in *Proc. IEEE Power Energy Soc. Gen. Meeting (PESGM)*, Boston, MA, USA, 2016, pp. 1–11.

[29] E. Bitar and Y. Xu, "Deadline differentiated pricing of deferrable electric loads," *IEEE Trans. Smart Grid*, vol. 8, no. 1, pp. 13–25, Jan. 2017.

[30] H. J. Jo, I. S. Kim, and D. H. Lee, "Efficient and privacy-preserving metering protocols for smart grid systems," *IEEE Trans. Smart Grid*, vol. 7, no. 3, pp. 1732–1742, May 2016.

[31] F. Knirsch, G. Eibl, and D. Engel, "Error-resilient masking approaches for privacy preserving data aggregation," *IEEE Trans. Smart Grid*, vol. 9, no. 4, pp. 3351–3361, Jul. 2018.

[32] P. Gope and B. Sikdar, "Lightweight and privacy-friendly spatial data aggregation for secure power supply and demand management in smart grids," *IEEE Trans. Inf. Forensics Security*, vol. 14, no. 6, pp. 1554–1566, Jun. 2019.

[33] D. Egarter, C. Prokop, and W. Elmenreich, "Load hiding of household's power demand," in *Proc. IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, 2014, pp. 854–859.

[34] A. Awad, P. Bazan, and R. German, "Privacy aware demand response and smart metering," in *Proc. IEEE 81st Veh. Technol. Conf. (VTC Spring)*, 2015, pp. 1–5.

[35] P. Gope and B. Sikdar, "An efficient data aggregation scheme for privacy-friendly dynamic pricing-based billing and demand-response management in smart grids," *IEEE Internet Things J.*, vol. 5, no. 4, pp. 3126–3135, Aug. 2018.

[36] P. Gope and B. Sikdar, "An efficient privacy-friendly hop-by-hop data aggregation scheme for smart grids," *IEEE Syst. J.*, to be published, doi: 10.1109/JSYST.2019.2899986.

[37] A. Braeken, P. Kumar, and A. Martin, "Efficient and privacy-preserving data aggregation and dynamic billing in smart grid metering networks," *Energies*, vol. 11, no. 8, p. 2085, 2018.

[38] H. Li, X. Lin, H. Yang, X. Liang, R. Lu, and X. Shen, "EPPDR: An efficient privacy-preserving demand response scheme with adaptive key evolution in smart grid," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 8, pp. 2053–2064, Aug. 2014.

[39] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

[40] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.

[41] P. Makris *et al.*, "Digitization era for electric utilities: A novel business model through an inter-disciplinary S/W platform and open research challenges," *IEEE Access*, vol. 6, pp. 22452–22463, 2018.

[42] I. Mamounakis, N. Efthymiopoulos, G. Tsaousoglou, D. J. Vergados, P. Makris, and E. Varvarigos, "A novel pricing scheme for virtual communities towards energy Efficiency," in *Proc. IEEE Int. Energy Conf. (ENERGYCON)*, Limassol, Cyprus, 2018, pp. 1–6.

[43] V. Krishna, *Auction Theory*. New York, NY, USA: Academic, 2002.

[44] L. M. Ausubel, "An efficient ascending-bid auction for multiple objects," *Amer. Econ. Rev.*, vol. 94, no. 5, pp. 1452–1475, 2004.

[45] P. Maymounkov and D. Mazières, "Kademlia: A peer-to-peer information system based on the XOR metri," in *Proc. Peer-to-Peer Syst. IPTPS*, 2002, pp. 53–65.

[46] G. Zyskind, O. Nathan, and A. Pentland, "Decentralizing privacy: Using blockchain to protect personal data," in *Proc. IEEE Security Privacy Workshops*, San Jose, CA, USA, 2015, pp. 180–184.

[47] G. Xylomenos *et al.*, "A survey of information-centric networking research," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 2, pp. 1024–1049, 2nd Quart., 2014.

[48] F. Gul and E. Stacchetti, "Walrasian equilibrium with gross substitute," *J. Econ. Theory*, vol. 87, no. 1, pp. 95–124, 1999.

[49] V. Krishna and M. Perry, *Efficient Mechanism Design*, Game Theory Inf., Univ. Library Munich, Munich, Germany, 1998. [Online]. Available: https://EconPapers.repec.org/RePEc:wpa:wuwpga:9703010

Georgios Tsaousoglou received the Ph.D. degree from the National Technical University of Athens (NTUA) in 2019. He has been working in various H2020 EU projects in the area of smart grids as well as with the Greek Independent System Operator. He is currently a Senior Researcher with NTUA and a Marie Curie Fellow with the Technical University of Eindhoven (host) and Denmark Technical University (co-host). His research interests include multiagent systems, algorithmic game theory, optimization and artificial intelligence applied to the area of the smart grid and especially to electricity markets, demand response, and electric vehicles.

Konstantinos Steriotis received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 2013, where he is currently pursuing the Ph.D. degree. His research interests are in the area of smart grids and especially energy markets, demand side management, and energy storage systems.

Nikolaos Efthymiopoulos was born in 1980. He received the Ph.D. degree in computer science in 2010. He is currently a Senior Researcher with the National Technical University of Athens, Athens, Greece. Since 2004, he has been participated in the coordination of various projects (FP7-ICT-VITAL++, FP7-ICT-STEER, H2020-ICT-SOCIALENERGY, H2020-EE-FLEXGRID). His research activities include scalable optimization, theory of dynamical systems, pattern recognition, big data, distributed searching, market architectures, pricing, data networks, online social networks, and smart grids. He has more than 40 publications in the above areas.

Prodromos Makris (Member, IEEE) was born in 1985. He received the B.Sc., M.Sc., and Ph.D. degrees from the University of the Aegean, Mytilene, Greece, in 2007, 2009, and 2013, respectively. He is currently a Senior Researcher with ICCS/NTUA (National Technical University of Athens), Athens, Greece. During the last years, he has been actively participating in several national and EC-funded research and innovation projects. He has more than 40 publications in international conferences and journals. His research interests include context-aware mobile and wireless networking in the 5G and beyond era, Internet of Things and their applicability in smart energy networks, and resource management algorithms for energy efficiency. He has recently served as the Technical Manager for FP7 VIMSEN-GA-619547 and H2020-GA-731767 SOCIALENERGY projects. He is currently the Technical Manager of H2020-GA-863876 FLEXGRID project. He is a member of the Technical Chamber of Greece.

Emmanouel (Manos) Varvarigos received the Diploma degree in electrical and computer engineering from the National Technical University of Athens in 1988, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology in 1990 and 1992, respectively. He has held faculty positions with the University of California, Santa Barbara, CA, USA, (as an Assistant and later an Associate Professor), the Delft University of Technology (as an Associate Professor), the University of Patras (as a Professor), and Monash University (as a Professor and the Head of ECSE Department). He is currently a Professor with the ECE Department, National Technical University of Athens. He is a member of the board of the Computer Technology Institute—Diophantus, where he has been the Director and then the Scientific Director of the Greek School Network Division since 2003 the main public ISP in Greece connecting more than 20 000 schools and other educational units, which has a major role in the development of network technologies and telematic services in Greece. He has participated in more than 40 U.S.- and EU-funded research projects in the areas of networking, smart energy grids, and grid and cloud computing, and in many national research projects, and has been the Consortium Coordinator in eight of them. He has over 380 publications in refereed international journals and conferences.