SLA–Mechanisms for Electricity Trading under Volatile Supply and Varying Criticality of Demand*

Georgios Methenitis[1,2], Michael Kaisers[1], and Han La Poutré[1,2]

[1] Centrum Wiskunde & Informatica, Amsterdam, The Netherlands
[2] Delft University of Technology, Delft, The Netherlands

Abstract. The increasing adoption of renewable power generation makes volatile quantities of electricity available, the delivery of which cannot be guaranteed. However, if not sold, the electricity might need to be curtailed, thus foregoing potential profits. In this paper we adapt service level agreements (SLAs) for the future smart electricity grid, where generation will primarily depend on volatile and distributed renewable power sources, and where buyers' ability to cope with uncertainty may vary significantly. We propose a contracting framework through SLAs to allocate uncertain power generation to buyers with varying preferences. These SLAs comprise quantity, reliability and price. We define a characterization of the value degradation of tolerant and critical buyers with regards to the uncertainty of electricity delivery (generalizing the Value of Lost Load, VoLL). We consider two mechanisms (sequential second-price auction and VCG) that allocate SLAs based on buyer bids. We further establish the settings of the proposed mechanisms, and show that both mechanisms ensure that no buyer has an incentive to misreport its valuation. We experimentally compare their performance and demonstrate that VCG dominates alternative allocations, while vastly improving the efficiency of the proposed system when compared to a baseline allocation that uses only the VoLL. This article facilitates distributed energy trading under uncertainty, thereby contributing an essential component to the future smart grid.

Keywords: SLAs, Electricity trading, Uncertainty

1 Introduction

Energy systems are in transition towards more sustainable and distributed generation portfolios, where smaller scale producers and consumers will participate as autonomous agents in decentralized markets. The main focus of this envisioned system is to maintain balance between available supply and demand. Maintaining balance becomes more challenging in face of generation from renewable resources such as the sun and wind, which are subjects to stochastic

* The extended abstract of this paper appears in the proceedings of AAMAS 2017 [20].

availability, and non-dispatchable. Their output cannot be regulated to match the demand, which is necessary to keep the system in balance. Therefore, demand side management (DSM) is necessary. DSM is the change in the behavior of the demand-side, which can be enabled through, e.g., financial incentives [9]. Dynamic pricing alongside scheduling of non-preemptive consumption loads are considered the main methodologies for balancing demand with uncertain supply. However, the former may introduce disruptive and unfavorable market behavior and thus planning and ahead prices are required [4], while the latter can violate the autonomy of consumer agents.

Service level agreements (SLAs) can provide the contracting framework for balancing volatile supply with demand between buyers (e.g., customers participating in retail tariff schemes) and sellers of electricity (e.g., small-scale producers that base their generation portfolio on distributed energy resources). Originally, SLAs define agreements between service providers and service users, specifying the service and its characteristics, which can vary depending on the application. To our best knowledge, for the first time in the context of electricity markets, we interpret SLAs as a direct extension of conventional electricity tariffs, which ensure delivery (100% quality) and a fixed kWh price (0% risk). In contrast to the current straightforward contracts, SLAs can be extended to include more features (e.g., delivery time, reliability, penalty for no delivery).

The process of specifying and allocating SLAs to buyers participating in the electricity market can be structured as a mechanism [12]. The design of the mechanism depends on the structure of the SLAs and the features they consist of. We specifically study SLAs comprising the following features:

Quantity The quantity of electricity that is subject to be transferred from the service provider to the user.

Reliability The probability of successful delivery of the quantity of electricity that is specified in the SLA.

Price The price per unit of the transferred quantity.

The features described above provide a basic SLA contracting framework for electricity trading between buyers and sellers.

Example 1. A seller holds a prediction of its generation for the next day during 1pm-2pm. The generation is not certain; there is 90% probability that the seller generates 1 unit and 50% probability 2 units of electricity. There are two buyer agents and both have unit demand. Let us now assume that the one buyer is a hospital that needs to perform a task, while the other buyer is an electric vehicle (EV) that needs to charge its battery. Assuming that there is no other seller in the system we consider, the two agents can agree on SLAs of 90% or 50% reliability with the seller for the unit demand they require. Considering that the importance of the task the hospital needs to complete is higher than the EV's, it is socially optimal to assign the most certain unit of generation to the hospital.

As illustrated in Example 1, it is socially desired that reliable energy is allocated to critical demand, and the risk of load-shedding is assigned to less critical buyers that in turn perform this task at lower social cost. The widest adopted concept to measure criticality in the literature as well as in practice is the *Value of Lost Load* (VoLL)[14]. The VoLL is defined as the estimated amount that customers receiving electricity through contracts would be willing to pay to avoid a disruption in their electricity service. In Example 1, the VoLL is higher for the hospital than for the EV.

In this paper, we study the problem of efficiently allocating SLAs for energy trading. We assume uncertain energy generation with a known distribution, and buyers with different preferences with regards to the uncertainty of being served. The main contributions of this work can be summarized as follows:

- We define a contracting framework through SLAs that enables energy trading under uncertain supply (Section 4).
- We propose a family of exponential functions that characterizes the buyers' varying degrees of criticality, thus generalizing the Value of Lost Load with costs associated to the risk of failed delivery (Section 4.2).
- We apply two mechanisms to assign SLAs to agents of different types, and incentivize truthfulness for strategic buyer agents (Section 5.2).
- Results show that the efficiency of the proposed system vastly improves in face of buyers with varying abilities to cope with uncertainty (Section 6).

The rest of the paper is structured as follows: Section 2 presents related work. Section 3 formulates the problem of purchasing electricity in the presence of uncertain supply. In Section 4, we introduce the structure of the proposed framework through SLAs, we further define the representation of the uncertain supply, and last, we propose a value function that determines different types and preferences of buyers. Section 5 studies and proposes different mechanisms to assign the SLAs, and it examines incentive compatibility issues arising in the domain. In Section 6, we evaluate our proposed setting, while Section 7 concludes this paper proposing possible extensions for future work.

2 Related Work

Several recent works have studied the problem of uncertainty in smart grids, where the increased volatility of such systems is the main reason for the potentially high costs associated with balancing supply and demand through conventional generation. For this reason, incentives for balancing can be forwarded to the demand-side [18, 27, 19]. Meanwhile, other works focus on the planning optimization of flexible consumers of electricity that can yield their own incentives (utility maximization) in the uncertain environment of the smart grid [2, 10].

In line with our assumptions, the problem of balancing volatile supply with flexible demand has been studied in recent works. In the presence of delay tolerant customers, service delays can be minimized via the Lyapunov optimization technique without the requirement of a-priory knowledge of the underlying statistics [21]. Furthermore, to deal with uncertain supply, potential scenarios of future renewable supply can be considered in Monte-Carlo planning. Based on the likelihood of scenarios, which is updated whenever new information about the supply is becoming known, an online mechanism that uses the concept of pre-commitment by the demand side has been proposed to allocate the available supply to flexible demand in order to maximize social welfare [24]. Similarly, the problem of matching uncertain supply with demand is considered as a multi-agent sequential decision making problem, where beliefs over states of the system are replaced with beliefs over future supply scenarios [28]. Each of the works above has a notion of cost (or criticality similar to VoLL), but none of those considers varying risk premiums for uncertainty.

SLAs have been considered as a tool for monitoring and coordination to ensure trustworthiness between different stakeholders, primarily with regards to the business processes, to ensure trustworthiness between different stakeholders [11], or as a negotiation protocol [1]. In contrast to our work, the discussion remains conceptual, and no quantitative implications on costs or efficiency are given. SLAs have also been used in resource allocation in computational grids to ensure the optimal allocation of computational resources and fair satisfaction of the participants through negotiations [23]. Here, the embedding is not in the energy domain, and the focus of the work is on strategic negotiation rather than the elicitation of truthful reports. Other related work has studied task allocation market mechanisms for multiple suppliers of finite or uncertain capacity [7].

In the closest state of the art work, the authors study the viability of selling uncertain quantities of wind generation with variable-reliability [3]. They further explore the connection between uncertainty in the generation and the costs for reserve capacity, and real-time markets. We follow a similar idea, but we focus our attention on the characterization of the demand with respect to its criticality, as well as the design of the mechanisms to assign demand through SLAs for electricity trading to strategic buyers with different preferences.

To the best of our knowledge we are providing the first work that adapts SLAs for energy trading under uncertain supply, providing both a discussion of the buyers' incentives and an empirical evaluation in illustrative settings. Here, SLAs are the resulting contracts that are allocated through auction-based mechanisms, to distribute uncertain supply to buyers of different types.

3 Uncertainty in Smart Grids

A *Smart Grid* is an electricity grid innovation that emphasises the transition from the traditional paradigm of passive distribution and consumption towards an energy network in which each node may take on an active role. This is exemplified by the increasing adoption of renewable generation (primarily from solar and wind) that makes households prosumers, serving some of their own load or even producing excess generation. The stochastic variation in generation introduces uncertainty in the supply. In addition, active control of loads such as heat pumps and batteries or charging controllers for electric vehicles introduce flexibilities – yet in the absence of clear incentives these flexible loads may be notoriously difficult to forecast.

In the proposed SLA allocation mechanism, participating buyer agents (i.e, electricity customers) purchase quantities of electricity through contracts of a specific quantity, reliability, and price, subject to uncertain supply and its availability. We consider only one seller and no outside option for the buyers. Let s denote the seller and $\mathcal{B}$ the set of buyer agents, there are n buyer agents in the set $\mathcal{B}$, such that $\mathcal{B} = \{1, \ldots, n\}$.

We make use of a simple two-step time model that serves as a fundamental model of the day-ahead auction process in current electricity markets. In the first step, *ahead* timestep, the mechanism holds a prediction (i.e., probability distribution function) of the available supply for the *realization* timestep. Let Q denote the random variable of the prediction of supply at the timestep ahead, $q \in \mathbb{R}^+$ is the observed realization of supply at the realization timestep. We further denote the cumulative density function of the random variable Q with $F(q) = \int_0^q f(x)\ dx$, where $f(q)$ is the probability density function. Similarly, each buyer agent i from the set $\mathcal{B}$ has a demand for electricity d_i, which we assume is fixed and known by the agent ahead of time.

Buyer agents, based on their preferences, can get assigned SLAs of certain quantity, reliability and price, ensuring that their demand will be satisfied with some probability in the realization timestep. The observed realization q of the electricity generation determines how much load can be served, and the mechanism determines the set of buyers that are indeed served such that the SLAs are satisfied in expectation with regards to their reliability.

4 Contracting Framework

4.1 Service–Level Agreements

As outlined in Section 1, an SLA is a triplet (d_i, γ_i, p_i), which comprises the quantity d_i, the reliability γ_i, and the price p_i per transferred unit of electricity for the buyer agent $i \in \mathcal{B}$. For the remainder of this paper we assume unit-demand buyers, $d_i = 1,\ \forall i \in \mathcal{B}$. We further assume that the delivery of the electricity of the assigned SLAs is either successful or not. Let $\hat{d}_i = \{0, d_i\}$ denote the transferred quantity to the buyer agent i. We define v_i as the marginal *value* that the successful delivery of electricity brings to the buyer agent, $v_i = \alpha_i \hat{d}_i,\ \alpha_i \in \mathbb{R}^+$. The value of α_i refers to the private value of the agent i when delivery is assured ($\gamma_i = 1$), or equivalently the VoLL. Considering the binary model for the value that the transferred quantity of electricity delivers to the buyer agent, the value can be either fully obtained or not ($v_i = \{0, \alpha_i d_i\}$). The expected value of the buyer, given the demand d_i and the reliability γ_i, is

$$v_i(\gamma_i) = \alpha_i d_i \gamma_i. \tag{1}$$

Let $S(q)$ denote the reliability function (also known as survival function) of the seller agent s. Figure 1 illustrates $S(q)$. Note that $S(q) = 1 - F(q)$, where $F(q)$ is the cumulative density function of Q. The reliability function $S(q)$ determines the probability that the generation exceeds a certain value q. The dotted area represents an SLA (no price p_i is determined here) between the seller agent and the buyer agent i. The demand of the buyer agent i is equal to d_i and the

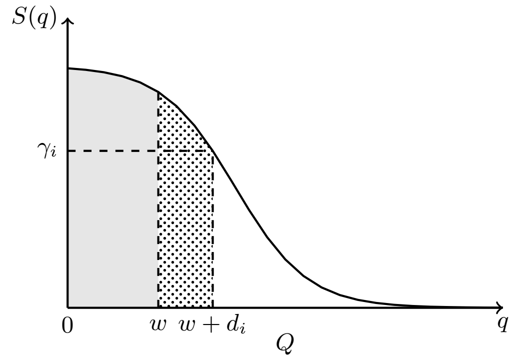

Fig. 1. The thick line illustrates the reliability function $S(q) = P(Q > q)$ of the random variable Q. The dotted area represents the portion of demand d_i of the buyer agent i with reliability $\gamma_i = S(w + d_i)$. The gray shaded area represents already assigned SLAs between the seller and the buyer agents.

reliability of the specific SLA is $\gamma_i = S(w + d_i)$, where w is the demand that is already deducted by previously allocated SLAs.

4.2 Critical & Tolerant Buyers

The expected value of a buyer in (1) is linearly dependent on the reliability γ_i of the SLA. Since the system gives raise to risk, we can distinguish between different attitudes of buyers towards risk, from critical to tolerant, as it is usual in economics and expected utility theory [13]. We define the *generalized* expected value function $\mathcal{V}_i(\gamma_i)$, where the reliability γ_i induces the risk in the form of uncertain delivery of the specified in the SLA quantity:

$$\mathcal{V}_i(\gamma_i) = \alpha_i d_i u_i(\gamma_i), \tag{2}$$

where $u_i(\gamma_i)$ encompasses the attitude of the buyer to the reliability γ_i. Note that for $u_i(\gamma_i) = \gamma_i$, $\mathcal{V}_i(\gamma_i)$ becomes equal to $v_i(\gamma_i)$ in (1). The generalized expected value function in (2) should embrace some common sense properties:

- Buyers have zero value for no reliability, i.e., $\mathcal{V}_i(0) = 0,\ \forall i \in \mathcal{B}$.
- Buyers have maximum value for no uncertainty, i.e., $\mathcal{V}_i(1) = \alpha_i d_i,\ \forall i \in \mathcal{B}$.
- Buyers have higher value for more certainty in the delivery, i.e., $\mathcal{V}_i(\gamma_i) \geq \mathcal{V}_i(\gamma_i - \varepsilon), \forall \varepsilon \in \mathbb{R}^+,\ \forall i \in \mathcal{B}$ (monotonicity).
- Buyers have positive value for any positive reliability, i.e., $\mathcal{V}_i(\gamma_i) > 0, \forall \gamma_i > 0, \forall i \in \mathcal{B}$ (buyers' willingness to participate).

We consider a variation of the exponential utility function [13]. In line with the aforementioned properties of $\mathcal{V}(\gamma_i)$, we define $u_i(\gamma_i)$ with regards to $\beta_i \in \mathbb{R}$.

$$u_i(\gamma_i) = \begin{cases} \dfrac{1 - e^{-\beta_i \gamma_i}}{1 - e^{-\beta_i}} & , \beta_i \neq 0 \\ \gamma_i & , \beta_i = 0 \end{cases} \tag{3}$$

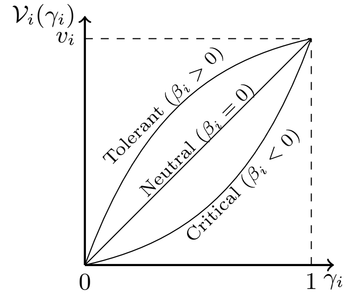

Fig. 2. The expected value function $\mathcal{V}_i$ of a buyer agent with regards to the reliability γ_i of an SLA. At full reliability ($\gamma_i = 1$) the expected value of the SLA by the buyer agent is equal to v_i which is determined by the private value α_i and the demand d_i. The value of β_i distinguishes different attitudes of buyer i towards the reliability γ_i. The values used for the illustration are $\beta_i = [-2.5, 0, 2.5]$.

This variation of the exponential utility function (3) can be substituted into Equation (2), yielding the expected value of the buyer agent given the reliability γ_i, where $\beta_i \in \mathbb{R}$ distinguishes the buyer type from critical ($\beta_i < 0$) to tolerant ($\beta_i > 0$). For $\beta_i = 0$ the expected value function $\mathcal{V}_i(\gamma_i)$ becomes equal to the expected value in (1). Figure 2 illustrates $\mathcal{V}_i(\gamma_i)$ for different values of β_i.

Given the above definition of the generalized expected value function, we can categorize buyer agents with respect to their attitudes towards reliability:

Critical For $\beta_i < 0$, the expected value function is convex, representing a critical (risk-averse) buyer. There is a stiff degradation of the value with regards to the uncertainty of the delivery for the critical buyer, resulting from opportunity costs that arise in case of failed delivery.

Tolerant For $\beta_i > 0$, the expected value function is concave, the buyer is tolerant (risk-seeking). Lower reliability translates to a rather high expected value, resulting from opportunity value that arises in case of failed delivery.

Neutral For $\beta_i = 0$, the expected value is linearly dependent to the reliability as in (1), representing a neutral buyer.

The generalized expected value function outlines a realistic model for capturing buyers' preferences. The proposed function describes the graceful or stiff degradation of the VoLL with regards to the probability of successful delivery. The type of the buyer agent i is characterized by the tuple (α_i, β_i). In the context of electricity markets, the same quantity of electricity may have different value for different consumers, which is captured in our model by α_i. Furthermore, the incurred value of a lost load with regards to the probability of electricity delivery is determined in our model by β_i. The function $\mathcal{V}_i(\gamma_i)$ indicates the VoLL of a buyer agent given the reliability of the SLA as $VoLL = \mathcal{V}_i(1) - \mathcal{V}_i(\gamma_i)$. The generalized expected value function can be defined in both one-shot and repeated settings, where in the latter agents could vary their types with regards to the outcome of their earlier assignments.

The price p_i that is specified in an SLA, and hence the expected utility of the buyer agent i will be determined by the resultant allocation of the mechanism (see Section 5). Let $\mathcal{U}_i$ denote the expected utility of the buyer i,

$$\mathcal{U}_i = \mathcal{V}_i(\gamma_i) - d_i p_i \gamma_i, \tag{4}$$

where the expected payment that is transferred from the buyer to the seller upon delivery is subtracted from the expected value.

5 Auction-Based SLA Allocation

Auctions are widely used in competitive electricity markets that take place day-ahead [5], and they are known to yield efficient allocations even in cases there is uncertainty about buyers' valuations for items to be sold [15]. We consider auctions as the method to allocate SLAs among buyers with varying private types, with regards to the value of their demand and their flexibility under the presence of uncertainty in the energy market. Here, we assume no agency for the seller, who only serves as the mechanism to allocate SLAs to the buyers.

Let $A \in \mathcal{A}$ be an allocation from the set of all feasible allocations $\mathcal{A}$, as the triplet of vectors $(\mathbf{d}, \boldsymbol{\gamma}, \mathbf{p})$, where each entry $A^{(i)} = (d_i, \gamma_i, p_i)$ is an allocated SLA between the buyer agent i and the seller agent s. Considering unit-demand buyers, $d = d_i = 1$, $\forall i \in \mathcal{B}$, an allocation can be written as $A = \{o_1, o_2, \ldots, o_n\}$, where $o_i \in \mathbb{Z}_1^n$ denotes the ordering of an allocated SLA. Given the ordering o_i of the SLA (d_i, γ_i, p_i) the reliability γ_i is given by $\gamma_i = S(o_i d)$ (see Fig. 1). Following the definition of the reliability function, $\forall j \in \mathcal{B}$, $\gamma_i \geq \gamma_j$, where $o_i < o_j$, the reliability is monotonically decreasing with the ordering. The set of feasible allocations $\mathcal{A}$ includes all allocations A for which every element appears only once in the set. The value of an SLA by the buyer agent i is determined by (2), with regards to the reliability γ_i and consequently to the ordering o_i. We define the expected social value of the buyer agents as the sum of the expected values of the set of the buyer agents given the allocated SLAs as $\sum_{i \in \mathcal{B}} \mathcal{V}_i(\gamma_i)$, where $\gamma_i = S(o_i d)$ is determined according to A. The allocation A further determines the order that buyers get served, for each buyer $\forall i \in \mathcal{B}$, $served_i = (q \geq o_i d)$, which follows from $P(served_i) = \gamma_i$.

5.1 Sequential Second-Price Auction

We consider a sequential second-price auction (SSPA) [17, 25], as an SLA allocation mechanism for all supply that may become available at the realization timestep. Items, SLAs in this case, are auctioned off one at a time. Given the assumption of unit-demand buyers, the seller auctions off SLAs of quantity d. We consider that the seller starts auctioning SLAs of decreasing reliability, such that the first SLA has reliability of $S(d)$, the second $S(2d)$, and so on. The ordering of auctioned items in sequential auctions affects the auctioneer's revenue [8]. Given the monotonicity property of the generalized value function in (2), SSPA of decreasing reliability SLAs maximizes the revenue of the seller. However, in later sections of this paper we also evaluate the case where the seller auctions off SLAs of increasing reliability.

A second price auction (also known as Vickrey) [26] is held by the seller in every round k of the auction where an SLA of reliability $S(kd)$ is auctioned off. Let $\mathcal{V}'_i(S(kd))$ denote the reported value of the buyer agent i with regards to the reliability $S(kd)$ offered in the k-th round of the sequential auction. Each buyer i places a sealed bid z_i, which is equal to the reported value with regards to the reliability $S(kd)$. The winner agent $w \in \mathcal{B}$ is determined as the buyer who submits the highest bid, $w = \text{argmax}_i \, z_i$, while the winning agent pays to the seller the price p_w of the second highest bid, such that $p_w = \max_{i \neq w} z_i$. The winner agent w of every round k of the sequential auction is allocated an SLA of unit quantity d, reliability $S(kd)$, and price p_w, no further participating in the next rounds of the auction. In each round of the auction only the winner agent w knows the price of the assigned SLA (no *price-discovery*). In the second price auction a buyer cannot increase its probability of allocation by increasing its bid in case the second highest bid is lower. In the opposite case, the buyer could win the auction by increasing its bid. However, this would result in negative utility (no *over-bidding*).

Theorem 1. *We assume that strategic buyers do not communicate their preferences to other participating agents, and there is no knowledge regarding the number and the distribution of the buyers participating in the auction. Furthermore, no buyer knows the reliability function of the seller, buyers only know that the reliability of the next SLA to be auctioned off is lower or equal to the reliability of the SLA that is being auctioned. The reliability function S is defined for all demand quantities. However, not all demand is guaranteed to be satisfied within an SSPA, since SLAs may be assigned with zero reliability. In SSPA, each round of the auction is an isolated Vickrey auction and therefore the mechanism is dominant-strategy incentive-compatible (DSIC).*

Proof. Each round of SSPA can be the last round or the round before with value arbitrarily close to zero and therefore can be treated as an isolated Vickrey auction [26]. Buyers' dominant strategy is to report their true value function $\mathcal{V}'_i(S(kd)) = \mathcal{V}_i(S(kd)), \forall i \in \mathcal{B}$. □

The proof of Theorem 1 exploits the property that no stochastic model can be built by the buyer regarding follow up rounds of SSPA. Given the assumptions made in the statement of Theorem 1, each round in which the buyer can wait without participating (bidding low or zero), does not add any information regarding the distribution of future bids of other agents. Consequently, there is no stochastic model which can compute an expectation of future utilities in case of waiting the next round to bid truthfully. To prove Theorem 1, we assume that the buyer is deterministic choosing to participate as this was the last round of the auction to maximize the likelihood of getting assigned an SLA of positive reliability. We showed that given the assumptions of the proposed setting there is no incentive for a strategic buyer to misreport its value function.

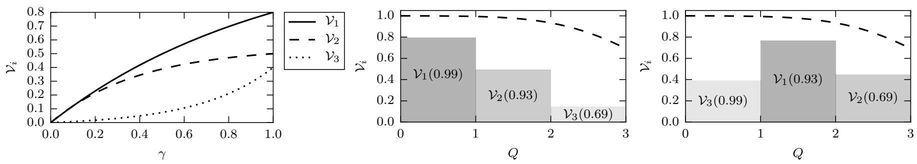

Fig. 3. (Left) The value functions $\mathcal{V}_i$ of the buyer agents with regards to the reliability γ of the possible SLAs. (Center) Resulting expected values of the buyer agents in the sequential second-price auction. (Right) Resulting expected values of the buyer agents in VCG mechanism. The dashed line illustrates the reliability over allocated SLAs.

5.2 Vickrey-Clarke-Groves Mechanism

Sequential second-price auctions can be suitable mechanisms to allocate SLAs to buyer agents, however the allocation of the SLAs depends more on the value α, and less on the criticality of the buyer, β.

Example 2. Consider two unit-demand buyers, agent 1 valuates 90% of successful delivery $\mathcal{V}_1(90\%) = a$ and $\mathcal{V}_1(50\%) = 3/4a$ and agent $2, \mathcal{V}_2(90\%) = a/2, \mathcal{V}_2(50\%) \approx 0$, using the sequential auction proposed in the previous section, agent 1 is assigned the SLA with 90% while agent 2 is assigned 50%. Considering zero payments, the resulting social value of the above assignment is a. However, the socially optimal solution would be agent 2 to be assigned 90% and agent 1 with 50% resulting in social value of $5/4a$.

Similarly to the above example, in Figure 3 (left) we illustrate the expected value functions of three buyers of diverse types. In Figure 3, center and right bar charts show two different allocations alongside the assigned reliability and the corresponding expected value. The dashed line presents the reliability of the allocated SLAs. In the greedy allocation (center), each slot is assigned to the buyer agent who has the highest bid (SSPA). On the contrary, the allocation that yields the optimal social value is illustrated on the right. SSPA by myopically allocating the highest bidder in each round (without any regards to the social value loss) results in a suboptimal social value.

Combinatorial auctions are means to derive socially optimal allocations [6]. In a combinatorial auction all buyers submit their bids for the whole bundle of items and the auctioneer computes the optimal allocation, which maximizes the social welfare. We consider Vickrey-Clarke-Groves (VCG) [22]. Given the vector $\boldsymbol{\gamma}$ of decreasing reliability of the available SLAs such that $\boldsymbol{\gamma} = (S(d), S(2d), \ldots, S(nd))$, each buyer submits a vector of bids $\mathbf{z}_i = (\mathcal{V}'_i(S(d)), \mathcal{V}'_i(S(2d)), \ldots, \mathcal{V}'_i(S(nd)))$, which includes the reported value of buyer i for each available reliability. Recall that $A \in \mathcal{A}$ is a feasible allocation $A = \{o_1, o_2, \ldots, o_n\}$, where o_i denotes the order over the decreasing reliability SLAs of the allocated agent. We further define $A^{-i} \in \mathcal{A}^{-i}$ as a feasible allocation of all buyer agents excluding agent i. Let A_{opt} denote the optimal allocation such that $A_{opt} = \text{argmax}_{A \in \mathcal{A}} \sum_{i \in \mathcal{B}} \mathcal{V}_i(A)$. The price p_i of each buyer agent is determined by its marginal contribution,

$$p_i = \sum_{j \in \mathcal{B} \setminus i} \mathcal{V}_j(A_{opt}^{-i}) - \sum_{j \in \mathcal{B} \setminus i} \mathcal{V}_j(A_{opt}), \tag{5}$$

where A_{opt}^{-i} is the optimal allocation without agent i present. Hence, each agent pays the loss incurred to the society by its presence. Under VCG mechanism, it is a dominant strategy for buyers to report their valuations truthfully [22].

In most problems, computing an optimal allocation lies in the class of NP-complete problems, and thereby computational infeasibility issues arise even for few participants.

Corollary 1. *The problem of SLA allocation among the unit-demand buyer agents can be solved optimally in polynomial time* $\mathcal{O}(n^3)$.

Proof. The unit-demand SLA allocation problem is equivalent to the *linear assignment problem*, where n agents have to be assigned n tasks while the number of tasks is equal to the number of agents. Each task stands for a slot in allocation $A \in \mathcal{A}$. It can be solved optimally in polynomial time, $\mathcal{O}(n^3)$, by the *Hungarian method* [16]. □

Given the polynomial complexity the proposed VCG mechanism could be used in a small scale electricity market (e.g., microgrid) under the presence of uncertain generation and varying buying preferences of the consumers.

6 Evaluation & Discussion

In this section we evaluate the performance of the studied mechanisms to allocate SLAs to buyers of different types. Specifically, we evaluate: the **VCG** mechanism (see Sec. 5.2), SSPA where the seller auctions off SLAs of decreasing reliability (**SPD**) as described in Section 5.1, and increasing reliability (**SPI**).

To study the efficiency of the proposed system when compared to a baseline allocation considering only the VoLL, we further compare against two mechanisms where only the value of buyers for certain delivery ($\mathcal{V}_i(1) = \alpha_i d_i$) is used for the allocation. In both mechanisms, a simultaneous second price auction is used for SLAs of certain delivery ($\gamma = 1$). In the first baseline mechanism (**POB**), the value function in (1) is used for the buyers, and thus we consider only neutral buyers (i.e., no added value or cost is generated as a result of the uncertainty). In the second baseline mechanism (**POC**), the generalized value function in (2) is used.

We evaluate and compare all the aforementioned mechanisms with regards to the *social value*, and the *social welfare*. For the remainder of this section, social value (SV) is defined as the average per member value. Following equation (2), SV $= \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \mathcal{V}_i(\gamma_i)$. Social welfare (SW) is the average of the expected utilities of the buyers, from (4), $SW = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \mathcal{U}_i$. Consequently, the expected seller's surplus can be written as $\mathcal{U}_s = \sum_{i \in \mathcal{B}} d_i \gamma_i p_i = |\mathcal{B}|(SV - SW)$, which is equal to the expected payments from the buyers to the mechanism (seller).

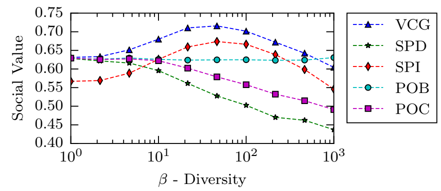

Fig. 4. Social value with regards to the diversity of criticality, for very low values of β-diversity all buyers approximate neutral attitude towards uncertainty in the delivery. For higher values of β-diversity buyers have increasingly varying criticality.

First, we analyze the influence of diversity in the criticality β of the buyers to the social value. We consider settings where buyers have similar private values α for electricity usage $\forall i \in \mathcal{B}, \alpha_i \sim \mathcal{U}(0.5, 1.0)$, which realistically captures buyers with similar needs and valuations for electricity (e.g., households). The random variable of the supply is normally distributed $Q \sim \mathcal{N}(\mu_Q = 20, \sigma_Q = 5)$, while the total demand exceeds by 20% the expected supply, $\sum_{i \in \mathcal{B}} d_i = 1.2\ \mu_Q = |\mathcal{B}|$ (24 buyers). We consider that $\forall i \in \mathcal{B}$, $\beta_i \sim \mathcal{U}(-D, D)$, where $D \in \mathbb{R}^+$ refers to the diversity of β values (β-Diversity). The higher the value of D, the larger the diversity in the criticality values β of the buyers.

Figure 4 illustrates the performance (social value) of each mechanism with regards to D (β-Diversity). For very low $D \approx 0$, buyers approximate the neutral behavior ($\beta \approx 0$). All mechanisms apart from SPI achieve similar performance. By auctioning off SLAs of increasing reliability to the buyers (starting from the lowest reliability) the likelihood to obtain higher values from buyers assigned SLAs early (and at low reliability) decreases. For $D \in [10^1, 10^2]$, there is a clear distinction between the performance different mechanisms achieve, where VCG mechanism obtains its maximum performance. VCG yields higher social value than all other mechanisms for almost the whole range of D. In the same range, SPI shows a significant increase in social value by prioritizing over tolerant buyers in the allocation[3]. The opposite is observed for SPD, which is the result of auctioning SLAs of decreasing reliability (See Sec. 5.2). For high $D \approx 10^3$ the performance of all mechanisms decreases below the performance of the baseline POB, which is not affected by the increasing diversity of β (neutral buyers). For very high diversity the probability of extremely critical buyers is increased. Consequently, the average social value is decreasing. In settings where buyers demonstrate extreme behavior with regards to the criticality ($\beta \ll 0$) the efficiency of the system is vastly affected. We showed how diversity in criticality β affects the social value achieved by the studied mechanisms, as well as the large improvement in the social value that VCG achieves.

We now proceed to show that even in the case of large variations in the private value α, VCG mechanism achieves an advantage in social value over

[3] In SPI it is more likely that tolerant buyers are assigned low reliability SLAs, since their value for low reliability is higher than other types of buyers.

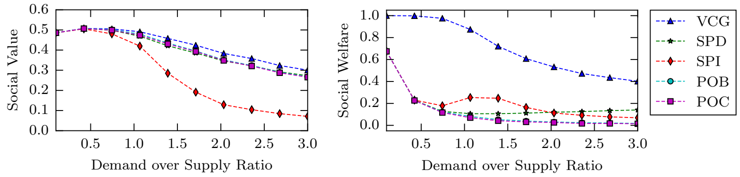

Fig. 5. Social value (left) and social welfare (right) achieved by different allocation mechanisms, namely, VCG, SSPA with decreasing reliability sequential auctions (SPD), increasing reliability sequential auctions (SPI), and two variants of a baseline method, POB and POC. The social value is computed as the average of all buyers' values with regards to the ratio of demand over the expected supply. Social welfare is normalized with respect to the social value obtained by the optimal VCG allocation.

the rest of the studied mechanisms. We use a diverse set of α, $\forall i \in \mathcal{B},\ \alpha_i \sim \mathcal{U}(0.1, 1.0)$, which captures the highly irregular private values of dissimilar buyers in electricity systems. In addition, we use small β-diversity, $\beta_i \sim \mathcal{U}(-5, +5)$, for buyers that do not exhibit extreme tolerant or critical behaviors. Figure 5 (left) illustrates the social value with regards to the ratio of the total demand over the expected supply. In the case where all the demand is served with high probability (ratio < 0.5), there is no significant difference between the different allocation mechanisms with respect to the social value. As the ratio increases it naturally follows that social value is decreasing for all allocation mechanisms. We can observe that the social value obtained by VCG is higher than any other allocation mechanism, and that there is a drop in the performance of SPI for demand to supply ratio higher than 1. The performance of the remaining methods does not vary significantly. VCG achieves the highest social value even in the case where α varies significantly among the buyers.

Finally, we show that the social welfare under VCG mechanism approximates the social value when the total demand is lower or approximately equal to the expected supply. Figure 5 (right) presents the social welfare achieved by all evaluated mechanisms, with respect to the ratio of the total demand over the expected supply. The social welfare is normalized with regards to the social value obtained by VCG mechanism (under the optimal allocation), (SW/SV_{VCG}). The normalized social welfare is equal to the ratio of social value remaining to the buyers, while the rest is transferred to the mechanism (seller) through payments. Up to ratio ≈ 1, the social welfare achieved by VCG mechanism is minimum at the 90% of the social value obtained using the optimal VCG allocation. The increased social welfare under SPI mechanism for ratio > 0.5 is a natural result, since more low reliability SLAs become available when the ratio increases[3]. On the contrary, SPD achieves around 15% of the optimal social value, however it exhibits a more stable (although lower) social welfare than SPI. The social welfare under the baseline mechanisms POB, POC approximates zero for high values of demand to expected supply ratio, and consequently most of the social value is transferred through payments to the mechanism. We presented the performance of the evaluated mechanisms in terms of the social welfare for different values of demand to expected supply ratio. We showed that VCG is a suitable mechanism to allocate SLAs in terms of the social welfare.

7 Conclusion

We proposed a contracting framework and mechanisms to allocate SLAs for electricity trading under uncertain supply and varying demand criticality of the buyers. We adapted SLAs as a direct extension of current conventional tariffs for use in electricity markets under uncertainty, and we defined the set of features that SLAs comprise, quantity, reliability, and price (Section 4.1). We further proposed a generalized value function for buyers with regards to the criticality of their demand in the face of uncertain delivery (Section 4.2). The proposed value function generalizes the concept of the Value of Lost Load with regards to the risk of unsuccessful delivery. The allocation of the SLAs to varying types of buyers as it results from two Vickrey based mechanisms (sequential second price auction, VCG) (Section 5). The two mechanisms ensure that no buyer has an incentive to misreport its value under certain conditions. Last, we evaluated the two mechanisms in an experimental study showing that VCG performance dominates all other allocations over a wide range of settings, and vastly improves the efficiency of the proposed system when compared to baseline allocation mechanisms considering only the VoLL (Section 6).

This work enables distributed energy trading under uncertainty, and may also serve as a broad basis for future extensions: (1) In this paper we have considered no agency for the seller, the study of the seller incentives to misreport its reliability function can therefore be a direct extension. (2) The exact characterization of the values functions of the buyers participating in the proposed electricity market under the presence of uncertain delivery. (3) The enrichment of the features included in the SLAs e.g., time, penalties for misreport, no-delivery penalty, (4) the multi-unit demand case, and (5) the presence of an outside option for the buyers, e.g., multiple sellers. In view of the attained properties and performance, we believe that using SLAs as we delineated here provides a promising avenue for addressing electricity trading in future smart grids. In particular, VCG allocation of SLAs is computable in $O(n^3)$, making it viable to assign the risk of demand curtailment to buyers that perform this task with low social cost. It can therefore be a tractable solution for peer-to-peer trading to balance local fluctuations in islanded grid scenarios or microgrids.

Acknowledgments

This work is part of the research programme Uncertainty Reduction in Smart Energy Systems (URSES) with project number 408-13-012, which is partly financed by the Netherlands Organisation for Scientific Research (NWO).

References

1. Amato, A., Martino, B.D., Scialdone, M., Venticinque, S.: Towards a sla for collaborating smart solar-powered micro-grids. In: 2014 International Conference on Intelligent Networking and Collaborative Systems. pp. 592–597 (Sept 2014)
2. Angelidakis, A., Chalkiadakis, G.: Factored mdps for optimal prosumer decision-making. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems. pp. 503–511. International Foundation for Autonomous Agents and Multiagent Systems (2015)
3. Bitar, E., Poolla, K., Khargonekar, P., Rajagopal, R., Varaiya, P., Wu, F.: Selling random wind. In: 2012 45th Hawaii International Conference on System Sciences. pp. 1931–1937 (Jan 2012)
4. Braithwait, S., Hansen, D., O'Sheasy, M.: Retail electricity pricing and rate design in evolving markets. Edison Electric Institute pp. 1–57 (2007)
5. Contreras, J., Candiles, O., De La Fuente, J.I., Gomez, T.: Auction design in day-ahead electricity markets. IEEE Transactions on power Systems 16(3), 409–417 (2001)
6. Cramton, P., Shoham, Y., Steinberg, R.: Combinatorial auctions (2006)
7. Dash, R.K., Vytelingum, P., Rogers, A., David, E., Jennings, N.R.: Market-based task allocation mechanisms for limited-capacity suppliers. IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans 37(3), 391–405 (2007)
8. Elkind, E., Fatima, S.: Maximizing revenue in sequential auctions. In: International Workshop on Web and Internet Economics. pp. 491–502. Springer (2007)
9. Gellings, C.W., Chamberlin, J.: Demand-side management. Energy Efficiency and Renewable Energy Handbook, Second Edition edited by D. Yogi Goswami, Frank Kreith (Chapter 15) pp. 289–310 (1988)
10. Gerding, E., Stein, S., Ceppi, S., Robu, V.: Online mechanism design for vehicle-to-grid car parks. In: Proceedings of the 25th International Conference on Artificial Intelligence. AAAI Press (2016)
11. Gustavsson, R., Hussain, S., Nordström, L.: Engineering of trustworthy smart grids implementing service level agreements. In: Intelligent System Application to Power Systems (ISAP), 2011 16th International Conference on. pp. 1–6 (Sept 2011)
12. Hurwicz, L., Reiter, S.: Designing economic mechanisms. Cambridge University Press (2006)
13. Ingersoll, J.E.: Theory of financial decision making, vol. 3. Rowman & Littlefield (1987)
14. Kariuki, K.K., Allan, R.N.: Evaluation of reliability worth and value of lost load. IEE Proceedings - Generation, Transmission and Distribution 143(2), 171–180 (Mar 1996)
15. Krishna, V.: Auction theory. Academic press (2009)
16. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly 2(1-2), 83–97 (1955)
17. Leme, R.P., Syrgkanis, V., Tardos, E.: Sequential auctions and externalities. In: Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 869–886. SODA '12, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2012), http://dl.acm.org/citation.cfm?id=2095116.2095186
18. Ma, H., Robu, V., Li, N., Parkes, D.C.: Incentivizing reliability in demand-side response. In: the proceedings of The 25th International Joint Conference on Artificial Intelligence (IJCAI'16) (2016)

19. Methenitis, G., Kaisers, M., La Poutre, H.: Incentivizing intelligent customer behavior in smart-grids: A risk-sharing tariff & optimal strategies. In: Proceedings of the 25th International Conference on Artificial Intelligence. pp. 380–386. AAAI Press (2016)
20. Methenitis, G., Kaisers, M., La Poutré, H.: Sla–mechanisms for electricity trading under volatile supply and varying criticality of demand (extended abstract). In: Proceedings of the 2017 International Conference on Autonomous Agents & Multiagent Systems. AAMAS '17, International Foundation for Autonomous Agents and Multiagent Systems (2017)
21. Neely, M.J., Tehrani, A.S., Dimakis, A.G.: Efficient algorithms for renewable energy allocation to delay tolerant consumers. In: Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on. pp. 549–554. IEEE (2010)
22. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic game theory, vol. 1. Cambridge University Press Cambridge (2007)
23. Silaghi, G.C., erban, L.D., Litan, C.M.: A time-constrained {SLA} negotiation strategy in competitive computational grids. Future Generation Computer Systems 28(8), 1303 – 1315 (2012), http://www.sciencedirect.com/science/article/pii/S0167739X11002251, including Special sections SS: Trusting Software Behavior and SS: Economics of Computing Services
24. Ströhle, P., Gerding, E.H., de Weerdt, M.M., Stein, S., Robu, V.: Online mechanism design for scheduling non-preemptive jobs under uncertain supply and demand. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems. pp. 437–444. AAMAS '14, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2014), http://dl.acm.org/citation.cfm?id=2615731.2615803
25. Syrgkanis, V., Tardos, E.: Bayesian sequential auctions. In: Proceedings of the 13th ACM Conference on Electronic Commerce. pp. 929–944. EC '12, ACM, New York, NY, USA (2012), http://doi.acm.org/10.1145/2229012.2229082
26. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. The Journal of finance 16(1), 8–37 (1961)
27. Vinyals, M., Robu, V., Rogers, A., Jennings, N.R.: Prediction-of-use games: a cooperative game theoryapproach to sustainable energy tariffs. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems. pp. 829–836. International Foundation for Autonomous Agents and Multiagent Systems (2014)
28. Walraven, E., Spaan, M.T.: A scenario state representation for scheduling deferrable loads under wind uncertainty. In: MSDM 2015: AAMAS Workshop on Multiagent Sequential Decision Making Under Uncertainty, Istanbul, Turkey, May 2015 (2015)