Auctions for Congestion Management in Distribution Grids

Rens Philipsen*, Mathijs de Weerdt
Dept. Software Technology
Fac. Electrical Engineering, Mathematics, and Computer Science
Delft University of Technology
* r.m.philipsen@tudelft.nl

Laurens de Vries
Dept. Engineering Systems and Services
Fac. Technology, Policy, and Management
Delft University of Technology

***Abstract*—Large controllable loads, such as electric vehicles, are increasingly penetrating electricity distribution feeders. To avoid local congestion, their consumption behaviour must be steered, for which a real-time price propagated down from the transmission system does not suffice, as it does not reflect local grid conditions. To efficiently steer the charging of EVs by multiple self-interested parties, we propose an auction framework which accounts for local grid conditions, the limited flexibility of EVs, and the uncertainty inherent to small-scale networks. We formulate the EV charging problem as a job scheduling problem for self-interested aggregators, and auction network capacity for discrete time slots using sequentially-cleared auctions, which run in parallel. We simulate this auction on a local network using realistic data for EV driving behaviour and network capacity, showing this method leads to feasible allocations which are fairer in case one party is weaker than the other due to size or information asymmetry.**

I. Introduction

Changes in local electricity production (due to uptake in local renewable generation) or consumption (due to increasing use of electricity for transportation or heating) may lead to higher peak use on the distribution grid. Distribution grids have historically been overdimensioned to account for the inflexibility of consumption. Accounting for future high-consumption peaks in this manner is a costly approach, while local grid capacity problems are expected to be infrequent (but very expensive if they occur) [1]. Alternatively, flexible production and consumption could be be scheduled with the network constraints in mind, but this requires coordination between self-interested parties located on the same distribution network. Since increasing penetration of renewable electricity production will de-correlate the electricity price and network usage, an additional signal reflecting grid usage is required to ensure the grid is not overloaded. Aside from these technical concerns, Distribution System Operators (DSOs) may also have the more political objective of allowing equal entry to all parties.

The goal of this paper is therefore to design an efficient coordination mechanism for independent decision-makers controlling large flexible loads, with a focus on electric vehicles. Specifically, we propose an auction framework in which capacity is gradually made available over time, and the allocation and pricing functions promote early and honest reporting of valuations, while simultaneously allowing parties to account for uncertainty in both supply and demand of network capacity. We aim for a mechanism which encourages coordination for self-interested parties, regardless of their size or number.

Due to the differences between transmission systems and distribution systems, we cannot efficiently implement similar schemes from the transmission system. Most importantly, the radial network structure precludes a re-dispatch of generators as a means of relieving congestion; instead, flexibility must be offered by the consumers, who shift their demand in time. This flexibility is limited, however, by requirements from the consumers: cars simply must be charged. In allocating network capacity, we must account for two problems: the availability of capacity in future periods is uncertain, as some inflexible load exists, which is difficult to predict due to the small number of consumers; and the substitutability and complementarity of capacity, leading to complex valuations of (packages of) network capacity. We investigate a scenario in which there is high penetration of electric vehicle charging in a distribution grid. Electric vehicles are likely to play an important role in reaching the limits of the grid capacity, but at the same time their charging behaviour may be controlled to minimise cost and ensure the system limits are respected [2].

A. Related work

In economic literature, the allocation of multiple goods with interdependent valuations is often viewed through the lens of a combinatorial auction [3]. In such auctions, Vickrey-Clarke-Groves mechanisms (VCG) can induce truthful bidding, and lead to an equilibrium in which social welfare is maximised. Less is known, however, about their performance in the face of uncertainty in supply, and they impose heavy requirements on computation and communication. In order to make computation feasible even on short notice (15-minute intervals are commonly used), we aim to design an auction which is simple to understand, allows computation and communication to take place within a few seconds, and allows for any number of participants.

Planning demand-side response has received broad attention in research, but not always with grid capacity constraints in mind [4], [5], [6]. A second aspect that is often ignored is the uncertainty with regard to available grid capacity, electricity price, and driving schedules, which is amplified by the absence of large numbers of consumers; the increased variability is more problematic if it is uncertain. A third aspect is the effect of multiple parties being present in the market; work on EV charging often investigates centralised approaches [7], [8], [1]. Hu [9], by contrast, presents an iterative procedure to determine a grid capacity price in a setting with two aggregators. Our research differs by including uncertainty, and taking a different approach to the pricing mechanism.

II. Methodology and model description

We approach the problem of coordinating flexible consumption by designing an online auction framework in which network capacity is gradually made available over time. Our work assumes the following setting. On an electricity distribution feeder, electric vehicles (EVs) must be charged. We consider a single (possibly) congested point in the network (e.g. a substation), and auction capacity for this bottleneck. Time is discretised into periods, and we consider a limited time horizon. At any time we refer to the current period as τ, and some future period under consideration as t. There is a constraint $C_t^{\max}$ on the joint simultaneous charging by all EVs. This constraint is determined by the difference between the fixed available network capacity limit, and the varying inflexible consumption by households. The capacity available for EVs may therefore differ per period. Each period, the network operator makes a forecast of available capacity for the entire decision horizon, with increasing uncertainty for periods further into the future.

EVs are assigned to an aggregator, who make charging decisions on their behalf. The EVs arrive to the network, will depart at some deadline, and have a demand for electricity, which must be met before said deadline. We refer to such a charging task as a job. The arrival, demand and deadline (a, q, d) of an EV's next job are revealed to the aggregator upon arrival to the network. Before arrival, only a probability distribution is known for each of these variables. Between arrival and departure, EVs can charge any real number. We will refer to EVs and their associated jobs somewhat interchangeably with the index i. If a charging job is not completed, a penalty γ is

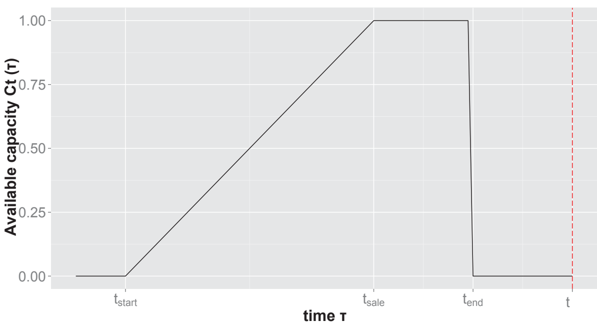

Fig. 1. Availability of network capacity over time

incurred. These aggregators are responsible for ensuring sufficient network capacity is available to meet their charging demands, and procuring the electricity itself. Their goal is to minimise their total cost, while meeting the charging demands of the EVs under their control. The procurement of electricity is associated with an exogenous procurement cost π. Like available capacity, this procurement price (set by the wholesale market) may vary per period.

A. Network capacity auction

We allocate capacity through the incremental sale of capacity in simultaneously-running, sequentially-cleared auctions. To model capacity availability, we utilise a simple piecewise linear curve (alternatively, any non-decreasing function can be used, e.g. a sigmoid) described by $t^{\text{start}}, t^{\text{sale}}$, and t^{end}. t^{start} is the period during which capacity for t is first made available, t^{sale} the moment all capacity is made available, and t^{end} the moment from which capacity is no longer sold.

$$C_t(\tau) = \begin{cases} C_t^{\max} & \text{if } t^{\text{sale}} \leq \tau < t^{\text{end}} \\ \frac{C_t^{\max}}{t^{\text{end}} - t^{\text{start}}} \cdot (\tau - t^{\text{start}}) & \text{if } t^{\text{start}} \leq \tau < t^{\text{sale}} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

The cumulative availability of capacity then follows the pattern as shown in Figure 1; A fully real-time market, of course, would occur if $t^{\text{start}} = t^{\text{sale}} = t^{\text{end}} = t$. Other known market designs can also be described in this manner (e.g. a forward market).

Pricing for capacity at each t is determined ex-post, allowing the network operator more freedom in determining a price. While standard first- or second-price rules could of course be applied, we opt for an approach which ensures the aggregators do not pay excessively for abundant capacity; by letting them pay the n^{th} price when considering all auctions between τ and t, a pricing rule similar to the critical price [10]. Formally, we define a provisional price $\hat{p}$ as the price of the highest rejected bid (i.e. not in the set of accepted bids B^τ), and let p be the minimum over itself and later provisional prices.

Resultingly, the price paid by a winning bidder at time τ for capacity at time t, is

$$p_t^\tau = \min\{\hat{p}_t^n : \tau \leq n \leq t\} \quad (2)$$

$$\hat{p}_t^\tau = \max\{p_b^\tau : b \notin B_{\text{accepted}}^\tau\} \quad (3)$$

B. Aggregation of EV charging

The aggregators aim to minimise their cost, which consists of charging costs and penalties for missed deadlines. The charging cost in (4) is the product of charging decisions s (for each car i at each timestep t) and the combined cost for electricity (π^t) and network capacity (p^t). The penalty is the product of the binary variable u_j (1 if job j is completed, 0 otherwise) and penalty γ. Equation (5) defines c_i^t as the energy required for meeting the deadlines d_j for the jobs $j \in J_i$ of car i before time t; at any time t a car must charge at least this amount (6), but may never charge more than the maximum capacity of the battery $\text{SOC}^{\max}$ (7). Finally, there is a charging limit per car (8), and a network capacity limitation on all cars of aggregator a (9).

$$\min_s \sum_t \left(\sum_{i \in I_a} s_i^t \right) \cdot (\pi^t + p^t) + (1 - u_j) \cdot \gamma \quad (4)$$

$$\text{s.t. } c_i^t = \sum_{j \in J_i | d_j \leq t} (u_j) \cdot q_j \quad (5)$$

$$\sum_{n=0}^{t} s_i^n \geq c_i^t \ \forall i, t \quad (6)$$

$$\sum_{n=0}^{t} s_i^n - c_i^t \leq \text{SOC}^{\max} \ \forall i, t \quad (7)$$

$$0 \leq s_i^t \leq s_i^{\max} \ \forall i, t \quad (8)$$

$$\sum_{i \in I_a} s_i^t \leq C_a^t \ \forall t \quad (9)$$

$$u_j \in \{0, 1\} \ \forall j \quad (10)$$

In order to obtain network capacity, the aggregators must submit bids, which are defined as a (t, p, q) triple indicating the timeslot for which capacity is bought, the price they are willing to pay per unit, and the desired quantity. We define three different strategies for the aggregators to convert the result of the optimisation problem into a set of bids which they submit:

- Optimistic: agggregators delay their purchasing decisions as long as possible, submitting bids only if the capacity remaining after this auction round may not be enough to meet their current charging schedule.
- Conservative: agggregators aggressively buy capacity which they cannot shift, but behave optimistically otherwise.
- Best Response: aggregators place a set of best-response bids to its opponent's (expected) bids, not considering possible future reactions.

Of these strategies, the former two are rather simple, while the latter is an optimal response and assumes sufficient information is present. In our experiments, we examine the effects of these asymmetries. Uncertainty, finally, is dealt with by generating scenarios, for each of which the bids to be placed are computed in a deterministic fashion. The results per scenario are combined through a voting-like method: for each t, we compute the median of both p and q, and submit $(t, \bar{p}, \bar{q})$ as the final bid.

TABLE I
EV DATA. ARRIVAL-DEPARTURE IN 15-MINUTE INTERVALS.

Cluster #		1	2	3	4	5	6
Probability		15.6%	32.3%	13.5%	12.0%	14.2%	12.4%
Departure	μ	53,9	31,12	36,16	69,63	36,7	35,09
	σ	5,95	4,86	6,28	8,86	7,05	10,18
	min	41	0	0	51	0	0
	max	70	45	51	94	54	78
Arrival	μ	67,05	68,55	47,6	86,17	85,93	74,94
	σ	5,76	4,67	7,18	5,82	5,21	10,4
	min	53	54	2	72	74	36
	max	83	79	60	96	96	96
Demand (kWh)	μ	4,84	6,55	4,67	5,8	7,83	24,1
	σ	3,24	3,71	3,31	4,58	4,15	5,78
	min	1,56	1,54	1,51	1,53	1,58	15
	max	18	17,1	21	27	18,25	37,4

C. Data

For the EV mobility data, we clustered trips from a national mobility study [11] (grouping by departure, arrival, and trip distance) into 6 clusters using k-means clustering. We assigned each EV to one such group, and drew all values from their associated truncated normal distributions, adjusting arrival time for the previously-drawn departure time. We assumed 150Wh/km for demand. The distributions are presented in Table I. Arrival and departure times are described by a 15-minute interval (e.g. 32 indicates 07:45-08:00 a.m.). Inflexible demand patterns were derived from household consumption models, using a Markov chain to capture interdependencies between subsequent periods of high and low demand.

III. THEORETICAL PROPERTIES

Our method of incrementally making capacity available has advantages for both sides when it comes to uncertainty: the network operator can adjust the available capacity if uncertainty decreases over time, while aggregators can delay their purchasing decisions as the capacity cannot be purchased in its entirety by a competitor, ascertaining that at least some capacity will be available in the future. A result of this strategy is that at any moment τ, capacity is available for multiple future time slots, and these auctions run simultaneously. Since a valuation for one future time slot likely affects the

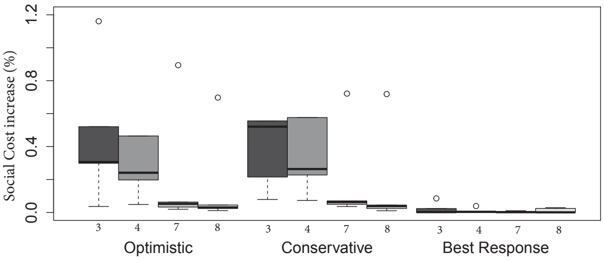

Fig. 2. Percentage of social cost increase over central optimum for different strategies and lead times.

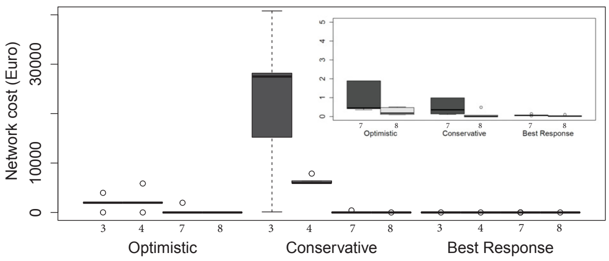

Fig. 3. Total price paid for capacity in deterministic setting for different strategies and lead times.

valuation for another (due to the substitutability of network capacity), we sequentially clear these markets, and allow aggregators to revise their bids in between. Both sides also benefit in determining the correct valuation of capacity: communication and computation are minimised for the grid operator (who would have to compute an optimal allocation given all possible preferences in a standard combinatorial auction), while aggregators can adapt their valuations as auction rounds finish, and their value for the remaining capacity changes. As a result, any allocation is final once made, but aggregators can take previous results into account when submitting their bids. The pricing rule, meanwhile, effectively means aggregators are never worse off by submitting a bid earlier; if they are able to obtain capacity for the same t at some $\tau' > \tau$ at a lower price, they still pay the price they would have paid at τ. This reduces the need for aggregators to reason about supply and demand in a single given auction, allowing them to submit bids which align with their true preferences, and ensures the price is 0 if there is excess capacity. This incentivises all bidders to bid as early as possible and at their true valuation. To illustrate this, consider an auction in which all capacity for a given t is made available instantaneously (which can be described using the given parameters); considered in isolation, this is an n^{th}-price auction for multiple units.

IV. Experimental results

We experimentally evaluate our auction to determine how the gradual sale of capacity affects total costs, and the distribution of these costs among the aggregators. Revenue maximisation is not among the DSO's goals; instead, we may use the revenue alongside social welfare as a measure of how well the mechanism enforces coordination. First of all, we investigate how the DSO should make capacity available. We compute the total charging cost and network capacity cost, and compare our auction design with different settings for t^{open} in a setting where only the available network capacity is uncertain. We ran simulations for two days, time steps representing 15-minute intervals. We used two aggregators with 18 EVs each. EVs are recurrent; some time after departure, they may return again, at which point their battery's charge has been reduced by the previously demanded quantity. We assume that at any time, an EV requires charge for its next single trip only; we do not differentiate between different demand levels for which an EV has different utilities. We replicated each experiment five times, using different EV schedules in each replication. Presented values are the average of these replications.

We compare the experimental results to two benchmarks: first, to a central solution with perfect information, which provides a lower bound on the costs; and second, to a decentralised solution where aggregators repeatedly best-respond, which provides a lower bound on network capacity prices. We ignore the different possible settings for t^{close}, setting it equal to t for ease of analysis. Resultingly, network capacity is made available between t^{open} and t, with constant slope.

A. Total charging cost and network capacity cost

In a setting with EV schedules assumed to be known, we first evaluate the performance in terms of total cost for different lead times, i.e. $t^{\text{start}} - t^{\text{end}}$. Figure 2 shows the results. Note that with a short lead time, the outliers indicate jobs are missed by the simple strategies. As we increase lead time, however, even combinations of these strategies are able to complete all jobs due to the coordination enforced by the auction. In this setting, completing all jobs does not require much lead time. The decrease of social cost as lead time increases indicates coordination emerges from the auction mechanism, signalling capacity scarcity ahead of time and allowing aggregators to revise their plans accordingly. Next, we turn our attention to the different auction prices for different lead times, shown in Figure 3. Due to the pricing rule, high auction prices indicate either actual scarcity or absence of coordination. Observe that even in cases where all jobs can be completed, network capacity prices are very high. These high spikes are somewhat artificial as they are in part due to the arbitrary penalty on not completing a job, and we therefore offer a more qualitative interpretation; high average total network capacity cost is due to last-minute bidder "panick", as they were forced to pay a network capacity price close to their value of lost load. Furthermore, we also observe that a longer lead time avoids these spikes, enforcing cooperation through the market design itself. The inset shows the results for the

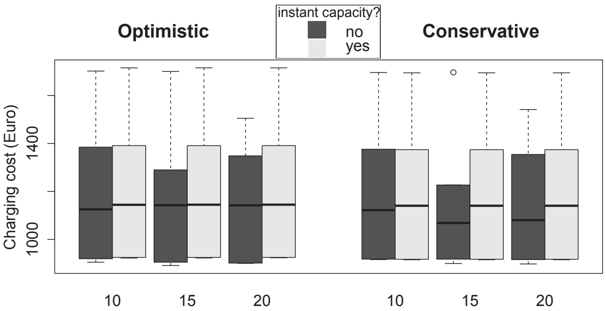

Fig. 4. Charging cost for weaker party without and with immediate availability, for different lead times

longer lead times. Again, even simple strategies rapidly approximate the benchmark result of best response. The case of best-responding aggregators is also of interest by itself. Irrespective of the lead time, network capacity costs are negligible. Due to the pricing rule, there is an incentive for demand reduction once outbidding the opponent is no longer profitable. By reducing demand until combined demand precisely matches supply, the resulting price drops to zero. Essentially, this is oligopolistic market sharing. Note that the likelihood of this result leans heavily on the assumption that aggregators have the correct information to determine their best response.

B. *Cost distribution*

Next, we examine the case where the aggregators are more asymmetric, one aggregator having a weaker position (e.g. due to information asymmetry or size). We compare our auction to a day-ahead benchmark, in which capacity is not released onto the market gradually, but immediately. Figure 4 shows the resulting cost for the weaker aggregator in different scenarios. We see that both the total charging cost increases for the weaker party if the network capacity is made available day-ahead. Gradually making capacity available can improve fairness by preventing the stronger aggregator from buying desirable capacity early on. It is also noteworthy to observe that a similar effect occurs when lead times are long. This suggests there is an optimal lead time, which depends on the model of the uncertainty involved (the EV driving schedule in this case).

V. Conclusion & future work

Existing work on congestion management in the distribution grid does not account for the uncertainty that is inherent to a small-scale system due to the small number of loads. In this work, we have proposed an auction framework in which capacity is gradually made available over time. We have shown that the gradual auctioning scheme does not necessarily require sophisticated strategies by its users, but is able to signal impending scarcity through the design of the auction itself. For scenarios with both deterministic and unknown EV schedules, our auction design allowed for relatively simple strategies. In a deterministic scenario, too short lead times caused jobs to be missed, but these quickly disappeared as lead time increased. Moreover, increased lead time also led to the decrease of network capacity prices. Markets operating closer to real time require more sophisticated strategies, as careful planning is required if load must be shifted to earlier time steps. Furthermore, we have shown that the gradual sale of network capacity can protect weaker parties (here we used a simple strategy as an example, but this might also be due to size, information, or risk aversity) by ensuring late availability of network capacity. This prevents stronger parties from buying all available capacity early on. This lesson potentially extends to other capacity auctions, such as continental gas pipelines.

Avenues of research which remain open include filling the gap in the strategy space for sophisticated strategies under imperfect information, and assessing the performance of our auction for more intelligent agents. Furthermore, in our work we made a number of choices with respect to the pricing and allocation functions, which we aim to investigate more broadly.

References

[1] Remco A. Verzijlbergh, Laurens J. de Vries, and Zofia Lukszo. Renewable Energy Sources and Responsive Demand. Do We Need Congestion Management in the Distribution Grid? *IEEE Transactions on Power Systems*, 29(5):2119–2128, sep 2014.

[2] Kristien Clement-Nyns, Edwin Haesen, and Johan Driesen. The Impact of Charging Plug-In Hybrid Electric Vehicles on a Residential Distribution Grid. *IEEE Transactions on Power Systems*, 25(1):371–380, 2010.

[3] Peter Cramton, Yoav Shoham, and Richard Steinberg. Combinatorial auctions. 2006.

[4] Philipp Ströhle, Enrico H. Gerding, and Mathijs M. de Weerdt. Online mechanism design for scheduling non-preemptive jobs under uncertain supply and demand. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, pages 437–444, 2014.

[5] Woongsup Lee, Lin Xiang, Robert Schober, and Vincent W S Wong. Direct Electricity Trading in Smart Grid: A Coalitional Game Analysis. *IEEE Journal on Selected Areas in Communications*, 32(7):1398–1411, 2014.

[6] Pedram Samadi, Hamed Mohsenian-Rad, Robert Schober, and Vincent W S Wong. Advanced demand side management for the future smart grid using mechanism design. *IEEE Transactions on Smart Grid*, 3(3):1170–1180, 2012.

[7] Olle Sundstrom and Carl Binding. Flexible Charging Optimization for Electric Vehicles Considering Distribution Grid Constraints. *IEEE Transactions on Smart Grid*, 3(1):26–37, mar 2012.

[8] Francesco Baccino, Samuele Grillo, Stefano Massucco, and Federico Silvestro. A Two-Stage Margin-Based Algorithm for Optimal Plug-in Electric Vehicles Scheduling. *IEEE Transactions on Smart Grid*, 6(2):759–766, 2015.

[9] Junjie Hu, Shi You, Morten Lind, and Jacob Østergaard. Coordinated Charging of Electric Vehicles for Congestion Prevention in the Distribution Grid. *IEEE Transactions on Smart Grid*, 5(2):703–711, 2014.

[10] Noam Nisan, Tim Roughgarden, Éva Tardos, and VV Vazirani. *Algorithmic game theory*. 2007.

[11] Centraal Bureau voor Statistiek. Onderzoek Verplaatsingen in Nederland. Technical report, 2013.