Online Mechanism Design for Vehicle-to-Grid Car Parks

Enrico H. Gerding
University of Southampton, UK
eg@ecs.soton.ac.uk

Sebastian Stein
University of Southampton, UK
ss2@ecs.soton.ac.uk

Sofia Ceppi
University of Edinburgh, UK
sceppi@int.ed.ac.uk

Valentin Robu
Heriot Watt University, UK
V.Robu@hw.ac.uk

Abstract

Vehicle-to-grid (V2G) is a promising approach whereby electric vehicles (EVs) are used to store excess electricity supply (e.g., from renewable sources), which is sold back to the grid in times of scarcity. In this paper we consider the setting of a smart car park, where EVs come and go, and can be used for V2G while parked. We develop novel allocation and payment mechanisms which truthfully elicit the EV owners' preferences and constraints, including arrival, departure, required charge, as well as the costs of discharging due to loss of efficiency of the battery. The car park will schedule the charging and discharging of each EV, ensuring the constraints of the EVs are met, and taking into consideration predictions about future electricity prices. Optimally solving the global problem is intractable, and we present three novel heuristic online scheduling algorithms. We show that, under certain conditions, two of these satisfy monotonicity and are therefore truthful. We furthermore evaluate the algorithms using simulations, and we show that some of our algorithms benefit significantly from V2G, achieving positive benefit for the car park even when agents do not pay for using it.

1 Introduction

Recent years have seen increasing interest in electric vehicles (EVs) as a key technology in the transition to a low-carbon transportation future. In AI and particularly multi-agent systems, recent work has discussed how the charging of EVs can be scheduled, e.g., to avoid bottlenecks in local distribution grids [Valogianni et al., 2015; Robu et al., 2013] or at EV charging stations [Gerding et al., 2013; Rigas et al., 2013]. However, EVs can also help ensure grid stability, especially as generation is shifting towards intermittent renewable resources. This is achieved through the so-called vehicle-to-grid (V2G) concept, whereby EVs are able to sell energy back to the grid during periods of scarcity and when energy prices are high.

In this paper, we use techniques from mechanism design and online scheduling to address the problem of charging and discharging multiple EVs. The setting we study is one of a public car park, where EVs arrive dynamically and stay for a period of time — during which they wish to be charged to a particular state of charge (SOC), but can also be used for V2G. Similar to existing work [Recalde Melo et al., 2014; Kahlen and Ketter, 2015], the reason to focus on this setting is that car parks have a larger size and "visibility" to the grid and hence are more likely to be offered an advantageous contract, or are able to trade in the wholesale market.

A crucial concern in addressing this problem, however, is that discharging has a cost for each EV, as it affects the remaining useful lifetime (RUL) of its battery. These costs are often EV-specific, as different vehicles have different types of batteries and age, and they are private to each EV driver. Moreover, when designing a charging/discharging schedule of each EV, a number of other constraints need to be taken into account, such as the driver's intended departure time and desired state of charge when leaving.

There has been considerable related work on V2G charging, typically focusing on how a demand response aggregator (such as a car park, or a local microgrid) can optimise its trading/balancing strategy using a number of available EVs (see also Section 2). However, many of the proposed solutions do not consider heterogeneous EVs with different costs for discharging, arrivals and SOC constraints. In addition, existing approaches assume the information about individual EVs to be known by the aggregator. In contrast, we consider EV owners to be self-interested agents with privately known preferences and constraints, who may misreport this information if this is to their benefit. Hence, we need to design incentives to truthfully elicit this information. Truthfulness (defined here using the strong notion of ex-post dominant-strategy incentive compatibility) is an important property, as it removes the strategic burden from EV owners. Moreover, since the information is used to optimise the balancing strategy, it results in more efficient outcomes.

Specifically, the contributions of our work are as follows:

- We formalise, for the first time, the V2G problem as an online mechanism design problem. Specifically, we consider a car park setting, where arriving EVs have differ-

ent discharging costs, required SOC and intended departure times, and this information is privately known. The aggregator (e.g., car park owner) decides on a charging and discharging scheduling algorithm and compensates agents for their incurred discharging costs to (1) incentivise truthful reporting, (2) minimise overall costs.

- We provide sufficient conditions for any scheduling and payment mechanism to be truthful in this setting.

- We propose two specific truthful mechanisms and a third mechanism, which is not necessarily truthful, but allows EVs to rely more on internal trading rather than with the grid. All three algorithms are scalable and can be used for large settings.

- Finally, we compare the performance of our three mechanisms experimentally using simulations, showing that they can achieve positive profits even when EVs do not pay for the charge they receive.

2 Related Work

The concept of V2G has received considerable attention in the area of power systems and energy economics. In this vein, the seminal work by [Kempton and Tomić, 2005] quantifies the revenues that EVs could achieve by participating in different electric markets: peak power shaving, spinning reserves and regulation services, but they do not design actual control algorithms for EV charging and discharging. Work which does consider such algorithms includes [Shafie-khah *et al.*, 2016; Kamboj *et al.*, 2011; Han *et al.*, 2010; Shi and Wong, 2011; Vasirani *et al.*, 2013]. [Shafie-khah *et al.*, 2016] use stochastic programming to provide a strategy for a plug-in EV aggregator agent to profitably participate in a demand response market. They consider several methods to reason about prices and to account for domain constraints (such as those coming from transformer loading). However, their analysis is done from the perspective of the aggregator, rather than individual EVs. Taking a multi-agent perspective, [Kamboj *et al.*, 2011] model a number of EVs providing anciliary services to the grid as a coalition, and discuss how coalitional game theory techniques could be used to divide the resulting rewards. [Han *et al.*, 2010] and [Shi and Wong, 2011] model the problem of EV charging as an MDP. Specifically, [Han *et al.*, 2010] use EVs for frequency regulation, and EVs are rewarded for the time they are available, irrespective of how much their battery capacity is discharged/recharged during their stay, unlike our approach, in which effort is measured proportionally to the amount discharged. Furthermore, none of the above approaches consider the problem of strategic EVs, requiring an appropriate mechanism to ensure participating EVs report their costs truthfully.

Online mechanism design has previously been applied to EV charging, e.g., [Robu *et al.*, 2013; Stein *et al.*, 2012; Valogianni *et al.*, 2015], but these approaches do not consider V2G settings. Other work considers a number of EVs working together as a virtual power plant (VPP), either to balance a renewable resource [Vasirani *et al.*, 2013] or to sell electricity in the balancing market. In this vein, [Kahlen and Ketter, 2015] consider a car sharing service, in which drivers rent

and return vehicles at different locations. The system operators can decide which vehicles parked at each spot to use for rental, and which to use as part of the VPP to provide regulation services to the grid. Both of these consider the optimisation problem from the perspective of the VPP aggregator, and do not guarantee truthfulness of individual EVs.

3 Formal Model

We consider the problem of a single smart car park, henceforth called the *aggregator*, where electric vehicles, henceforth called *agents*, arrive and leave over time.

The Agents. Let I denote the set of agents. An agent $i \in I$ arrives at time a_i and has a privately-known type $\theta_i = \langle d_i, q_i, c_i \rangle$ representing the agent preferences and constraints, where d_i is the (actual) departure time, q_i is the required amount of electricity by the departure time, and c_i is a discharging cost parameter (more details below). Importantly, note that, in this model, the required charge is a hard constraint and, as a result, we do not need to include the value derived from a certain amount of charge.

In addition, the agent's battery has a (known) state of charge at time t, denoted by $SoC_{i,t}$ and the current charging rate is denoted by $r_{i,t} = dSoC_{i,t}/dt$, where $r_{i,t} > 0$ means the battery is charging and $r_{i,t} < 0$ means it is discharging. The entire charging and discharging schedule is denoted by r_i. Furthermore, limits are given by SoC_i^{max}, SoC_i^{min}, r_i^{max} and r_i^{min}. For simplicity we assume these limits are known by the aggregator, but they could also be elicited (i.e., be included as part of the private type). Given this, we can describe the total amount charged and discharged by time t as $q_{i,t}^+(r_i) = \int_{a_i}^{t} max(0, r_{i,t'})dt'$ and $q_{i,t}^-(r_i) = -\int_{a_i}^{t} min(0, r_{i,t'})dt'$ respectively, noting that both these values are positive.

Charging and discharging a battery incurs a monetary cost, since capacity generally decreases for each charge and discharge cycle. This loss depends on many factors, but the main ones are the state of charge and the rate of charge [Divya and Ostergaard, 2009]. This translates into monetary costs when the battery needs to be replaced, but also the inconvenience of having a battery with lower capacity and requiring more frequent charging. In general, we model the slope of the capacity loss by $w(SoC_{i,t}, r_{i,t})$, which is assumed to be common knowledge, and the total capacity loss incurred up until time t is $W_i(r_i) = \int_{a_i}^{d_i} w(SoC_{i,t}, r_{i,t})dt$. Using game-theoretic terminology, this loss is also referred to as the amount of *work* incurred [Archer and Tardos, 2001]. Importantly, we assume that the costs are linear w.r.t. to c_i, and therefore the total cost incurred by an agent is $c_i \cdot W_i(r_i)$. For some of the results and in the simulations, we furthermore assume a constant loss per unit of (dis)charge, i.e., $w(SoC_{i,t}, r_{i,t}) = a \cdot r_{i,t}$, for some constant a. This is a reasonable approximation if the state of charge never gets too low (i.e., if SoC_i^{min} is sufficiently high). By putting charging and discharging costs together (i.e., any discharging needs to be compensated by charging to reach the required charge q_i), the work function simplifies to $W_i(r_i) = a \cdot q_{i,d_i}^-(r_i)$.

The above results in the agent utility:

$$u_i(r_i, \tau_i) = -c_i \cdot W_i(r_i) + \tau_i \qquad (1)$$

if the following constraints are met:

- $\forall t \in [a_i, d_i] : SoC_i^{min} \leq SoC_{i,t} \leq SoC_i^{max}$ and $r_i^{min} \leq r_{i,t} \leq r_i^{max}$
- $SoC_{i,d_i} \geq q_i$

and is $u_i(r_i, \tau_i) = -\infty$ otherwise, where τ_i are monetary *transfers* from the aggregator, detailed in Section 4.

The Aggregator The aggregator determines the charging and discharging schedules and the transfers, and sells any excess elecricity to the grid and purchases any shortfall. We assume that electricity prices change over time and the buy and sell prices can also differ. Let $p_t^{buy} \geq p_t^{sell}$ denote the buy and sell prices at time t and furthermore let $r_{I,t}(r_i) = \sum_{i \in I} r_{i,t}$ denote the overall/net (dis)charge at t. Given this, we can compute the *aggregator utility*, u_{agg}, as follows:

$$u_{agg} = \int_{t \in T} -min(0, r_{I,t}) \cdot p_t^{sell} - max(0, r_{I,t}) \cdot p_t^{buy} dt$$
$$- \sum_{i \in I} \tau_i, \quad (2)$$

where T is a time interval, e.g., 24 hours.

4 The Mechanism Design Problem

As is common in the mechanism design literature, the main goal is to maximise the *social welfare*, SW, which equals the sum of utilities of all agents (Equation 1) and the aggregator's utility (Equation 2) [Nisan *et al.*, 2007, Ch.9].[1] Maximising this is also known as maximising *market efficiency*. Noting that transfers cancel out, this results in:

$$SW = \int_{t \in T} -min(0, r_{I,t}) \cdot p_t^{sell} - max(0, r_{I,t}) \cdot p_t^{buy}$$
$$- \sum_{i \in I} c_i w(SoC_{i,t}, r_{i,t}) dt, \quad (3)$$

subject to individual agent constraints.

Although obtaining optimal social welfare is generally not possible (due to uncertainty of prices and future arrivals, as well as tractability issues, as discussed in Section 5), in order to be effective, the aggregator first needs to elicit the private information from the agents. To this end, on arrival each agent i is asked to report their type, where the report is denoted by $\hat{\theta}_i = \langle \hat{d}_i, \hat{q}_i, \hat{c}_i \rangle$.[2] We assume that the aggregator can observe the arrival time, a_i, of the agents and their state

---

[1] An alternative goal is to maximise aggregator utility and, while here we focus on social welfare, we do measure the aggregator profit in the empirical evaluation (see Section 7).

[2] There is a subtle difference with traditional online mechanism design, where agents are also asked to report their arrival. In the parking scenario, the physical arrival and departure can be observed by the aggregator. However, to produce a (dis)charging schedule, the aggregator needs to know the intended departure, which is what $\hat{d}_i$ signifies here.

of charge. Furthermore, it can verify the departure time, d_i, *ex-post*, meaning once the agent has actually departed. Given this, the mechanism design problem is to determine the charging rate $r_{i,t}$ at each time $a_i \leq t \leq d_i$ and transfers (payments to the agents) such that the agents have an incentive to report truthfully, i.e., such that reporting $\hat{\theta}_i = \theta_i$ is a dominant strategy for each agent i. Such a mechanism is also said to be *truthful*. Next we provide sufficient conditions for a mechanism to be truthful in our setting.

Our results are an extension of [Archer and Tardos, 2001]. In [Archer and Tardos, 2001], necessary and sufficient conditions are given for a setting where the type is given by parameter c_i only. To account for the remaining parameters, let $\succ$ define an ordering over types such that $\theta_i' \succ \theta_i$ means that either $q_i' < q_i$, $d_i' > d_i$ or $c_i' < c_i$ (intuitively, θ_i' is a 'better' type for the mechanism in one of the dimensions). Furthermore, let $f_i(\hat{\theta}_i) = f_i(\hat{d}_i, \hat{q}_i, \hat{c}_i)$ determine agent i's allocation as a function of his reported type, i.e., $r_i = f_i(\hat{\theta}_i)$. Given this, we obtain the following result:

Theorem 1. *The mechanism is truthful if the following conditions are met:*

1. *(Monotonicity of Work) Whenever $\theta_i' \succ \theta_i$ then $W_i(f_i(\theta_i')) \geq W_i(f_i(\theta_i))$*

2. *The transfers are computed on (observed) departure and are given by:*

$$\tau_i = h_i + \hat{c}_i(W_i(f_i(d_i, SoC_{i,d_i}, \hat{c}_i)) - W_i^0)$$
$$+ \int_{\hat{c}_i}^{\infty} (W_i(f_i(d_i, SoC_{i,d_i}, x)) - W_i^0) dx, \quad (4)$$

where W_i^0 is the minimum required work to get to the state of charge on departure (i.e., without additional discharging), and h_i is a constant which does not depend on the report (but could, e.g., depend on the observed arrival and departure). Note that transfers are calculated based on actual departure, which we assume the mechanism can observe ex-post.

3. $SoC_{i,t} = q_i$ for $t \geq \hat{d}_i$

Proof. First, note that, due to the second term, the transfers always compensate the agent's costs for any *additional* work done (i.e., work done on top of getting to the required charge) and, in fact, typically, due to the third term in Equation 4, the agent is better off when doing more work (and is never worse off). As a result, the agent has no incentive to overreport q_i and report an earlier d_i since this will only lead to less work. If the agent truthfully reports d_i, then he also has no incentive to underreport q_i since, due to condition 3, his state of charge would be insufficient and his utility $-\infty$. An interesting case is reporting $\hat{q}_i < q_i$ and $\hat{d}_i > d_i$ when the agent knows the mechanism will temporarily charge above $\hat{q}_i$ with the intention to discharge later, and this could result in more work. However, since the transfer is based on the actual state of charge and on the actual departure, it is the same as if he had reported $\hat{q}_i = SoC_{i,d_i}$ and $\hat{d}_i = d_i$. Therefore, while the agent has done more work, he does not get compensated

for it and only incurs the costs. As a result of the above, the agent has no incentive to misreport q_i or d_i. Given this, we can use the result from [Archer and Tardos, 2001, Theorem 4.2] to show that the agent also has no incentive to misreport c_i. □

5 The Online Scheduling Problem

We now focus on the problem of setting the (dis)charging rate of the EVs in the car park in order to obtain a high social welfare (as defined by Equation 3). Now, optimising this problem is intractable for our setting. An additional challenge is that future arrivals as well as grid prices are generally unknown. To this end, we assume that a probability distribution of the prices is known, that the number of prices is finite and that charging decisions are made at discrete time steps. This way, the problem can be formulated as a Markov decision process (MDP), and solved by dynamic programming. While solving this optimally is tractable for a single EV, with multiple EVs we need to consider all combinations of states for each EV. This leads to an exponential increase in the number of states, which very quickly becomes intractable. Therefore, we propose three different heuristics:

Algorithm 1. This algorithm only bases decisions on current prices and does not consider future prices. Let ρ^{low} and ρ^{high}, $\rho^{low} \leq \rho^{high}$, denote two thresholds. Each EV i currently in the car park will be charged if $p_t^{buy} \leq \rho^{low}$ (i.e., if the price is below a threshold) at a constant rate as long as $SoC_{i,t} < q_i$. Furthermore, an EV will be discharged at a constant rate if $p_t^{sell} - c_i w(SoC_{i,t}, r_{i,t})/r_{i,t} \geq \rho^{high}$ (i.e., if the marginal profit is above a certain threshold) as long as there is sufficient time to recharge to get to the required SOC q_i by the departure time.

Algorithm 2. The second algorithm considers future grid price distributions and optimises the *expected* utility for each vehicle i independently, given the price distributions, and assuming each agent buys and sells directly from the grid, i.e., assuming $\tau_i = \sum_{t=a_i}^{d_i}(-min(0, r_{i,t}) \cdot p_t^{sell} - max(0, r_{i,t}) \cdot p_t^{buy})$. The optimal solution is a *policy* which specifies a charging schedule for each realisation of the prices, and can be solved using dynamic programming.

Note that, in Algorithms 1 and 2, while the (dis)charging decisions are made independently for each agent, the demand and supply is first aggregated within the car park and electricity is only sold to (bought from) the grid if there is a net surplus (deficit). Therefore, whenever there is a situation where one EV charges and, at the same time, another discharges, instead of buying or selling from the grid, one EV will be used to charge another one, saving $p_t^{buy} - p_t^{sell} \geq 0$. Furthermore, note that Algorithm 2 is globally optimal when $\forall t : p_t^{buy} = p_t^{sell}$.

Algorithm 3. Algorithm 2 is optimal for an individual EV, but the optimal policy ignores potential benefits that can be achieved by discharging one vehicle in order to charge another one. To address this, Algorithm 3 takes the solution from Algorithm 2 as the basis, but then considers whether the social welfare can be improved further by adjusting the immediate

charging rate of the agents. In more detail, let $EU_{i,t}^*(SoC_{i,t})$ denote the expected utility of agent i from time t onwards given the current state of charge, assuming the agent buys and sells directly from the grid, and using the optimal policy from Algorithm 2. Then:

$$EU_{i,t}(r_{i,t}, SoC_{i,t}) = -min(0, r_{i,t}) \cdot p_t^{sell} - max(0, r_{i,t}) \cdot p_t^{buy}$$
$$-c_i w(SoC_{i,t}, r_{i,t}) + EU_{i,t+1}^*(SoC_{i,t} + r_{i,t})$$

is the expected utility when agent i (dis)charges at rate $r_{i,t}$ at time t and then proceeds with the individual agent's optimal policy from time $t + 1$ onwards (given the updated state of charge according to the current charging rate).

Note that, in Algorithm 2, the current charging rate is set to $r_{i,t}^* = \arg\max_{r_{i,t}} EU_{i,t}(r_i, SoC_{i,t})$. Instead, here we set the current charging rates of all agents in order to (myopically) maximise the social welfare as follows:

$$(r_1^*, r_2^*, \ldots) = \arg \max_{r_{i,t}, i \in I} \sum_{i \in I} EU_{i,t}(r_{i,t}, SoC_{i,t}) \quad (5)$$

Note that, unlike Algorithm 2, this algorithm may (dis)charge an EV if is not optimal for the individual EV, but when it contributes to the social welfare.

6 Algorithm Properties

In this section we prove that Algorithms 1 and 2 are truthful when the cost per unit of (dis)charge is constant (i.e., when $w(SoC_{i,t}, r_{i,t}) = a \cdot r_{i,t}$, see also Section 3). Recall that, according to Theorem 1, in order for these algorithms to be truthful, we need to show that ex-post monotonicity of work holds, for each of the type parameters. Note that we can consider each parameter independently, as long as monotonicity holds for a parameter independent of the remaining parameters. If monotonicity holds, the payments defined in Theorem 1 can be used to assure truthfulness.

Towards a proof, we first introduce bounds on the state of charge of a battery at any given time t for an agent i. In more detail, since we require that $SoC_{i,d_i} = q_i$ to guarantee truthfulness (Theorem 1), and due to the physical constraints on the (dis)charging rate and SOC of the agent, the algorithms above have an upper bound $ub_{i,t}$ and a lower bound $lb_{i,t}$ on the state of charge at time t given by:

$$ub_{i,t} = min\{SoC_i^{max}, q_i - r_i^{min} \cdot (d_i - t)\}$$

$$lb_{i,t} = max\{SoC_i^{min}, r_i^{max} \cdot (d_i - t) - q_i\}$$

We are now ready to show the monotonicity properties of the proposed algorithms w.r.t. each type dimension.

Theorem 2. *Algorithm 1 guarantees monotonicity of work w.r.t. cost parameter c_i, i.e., $c_i \geq c_i' \rightarrow W_i(f_i(\theta_i)) \leq W_i(f_i(\theta_i'))$, when the cost per unit charge is constant (i.e., independent of charging speed or state of charge).*

Proof. Let θ_i and θ_i' denote two identical types except for the costs c_i and c_i' respectively, and w.l.o.g. we assume that $c_i > c_i'$. In addition, let r_i and r_i' denote the respective allocations and $SoC_{i,t}$ and $SoC_{i,t}'$ the respective SOC. Note that $SoC_{i,a_i} = SoC_{i,a_i}'$ and $SoC_{i,d_i} = SoC_{i,d_i}' = q_i$. Furthermore, note that the bounds $ub_{i,t}$ and $lb_{i,t}$ are the same for

the two types θ_i, θ_i' because they do not depend on the cost (which is the only difference between them). To prove monotonicity we need to show that: $c_i \geq c_i' \rightarrow W_i(r_i) \leq W_i(r_i')$. However, since $W_i(r_i) = a \cdot q_{i,d_i}^-(r_i)$, it is sufficient to show that $q_{i,t}^-(r_i) \leq q_{i,t}^-(r_i')$ for any $a_i \leq t \leq d_i$. For this to hold, we need to show that, for any setting, one of two conditions always holds:

1. If $r_{i,t} > 0$ then $r_{i,t} \leq r_{i,t}'$, or, if $r_{i,t} < 0$ then $r_{i,t} \geq r_{i,t}'$, i.e., θ' does at least the same work as θ.

2. If $r_{i,t} \geq r_{i,t}'$ and $r_{i,t} > 0$ then, after charging, the SOC of θ_i' is equal to or higher than the one of θ_i, meaning that $q_{i,t}^+(r_i) \leq q_{i,t}^+(r_i')$, or, if $r_{i,t} \leq r_{i,t}'$ and $r_{i,t} < 0$ then, after discharging, the SOC of θ_i' is equal to or lower than the one of θ_i, meaning that $q_{i,t}^-(r_i) \leq q_{i,t}^-(r_i')$.

Next, we divide all possible settings into four cases reflecting the fact that, at any time t, the system can either charge ($r_{i,t} > 0$) or discharge ($r_{i,t} < 0$) and that $SoC_{i,t}$ is either greater than, equal to, or lower than $SoC_{i,t}'$. Note that in Algorithm 1 the upper bound $ub_{i,t}$ is always equal to q_i because the algorithm never overcharges the battery, i.e., $SoC_{i,t} \leq q_i$ and $SoC_{i,t}' \leq q_i$ for any t.

Case 1: $SoC_{i,t} \geq SoC_{i,t}'$ and $r_{i,t} > 0$. The fact that the mechanism decides to charge type θ_i means that the price p_t^{buy} is advantageous for θ_i given her cost parameter c_i. Since $c_i > c_i'$, this necessarily implies that charging is the optimal choice also for θ_i'. Indeed, assume that the optimal decision for θ_i' is $r_{i,t}' < r_{i,t}$. This implies that p_t^{buy} is more advantageous for θ_i then for θ_i'. Since the price p_t^{buy} is the same for both types, the one with the highest benefit is the one with the lower cost, i.e., $c_i < c_i'$. However, this contradicts our assumption that $c_i > c_i'$. Thus, $r_{i,t} \leq r_{i,t}'$. Note that, if $ub_{i,t}$ constrains $r_{i,t}'$, then also $r_{i,t}$ is constrained by it because $SoC_{i,t} \geq SoC_{i,t}'$. Thus, Condition 1 is satisfied.

Case 2: $SoC_{i,t} \leq SoC_{i,t}'$ and $r_{i,t} < 0$. This case is symmetric to Case 1.

Case 3: $SoC_{i,t} \leq SoC_{i,t}'$ and $r_{i,t} > 0$. First note that, given the results of cases 1 and 2 and the fact that $ub_{i,t} = q_i$, we have that $SoC_{i,t} \leq SoC_{i,t}' \rightarrow q_{i,t}^+(r_i) \leq q_{i,t}^+(r_i')$. Now, as with case 1, if it is optimal to charge θ_i then charging is the optimal decision also for θ_i'. However, in this case $ub_{i,t}$ may impose a constraint on $r_{i,t}'$ but not on $r_{i,t}$, i.e., $r_{i,t} > r_{i,t}'$. Nevertheless, note that it is never possible that, after charging, the SOC of θ_i is higher than the one of θ_i' because $ub_{i,t}$ affects θ_i' only when her SOC becomes equal to q_i. Consequently, the decision of charging θ_i at t never leads to the case in which $q_{i,t^*}^+(r_i) > q_{i,t^*}^+(r_i')$ at a future time t^*. Thus, Condition 2 is satisfied.

Case 4: $SoC_{i,t} \geq SoC_{i,t}'$ and $r_{i,t} < 0$. This case is symmetric to Case 3. $\qquad\square$

Theorem 3. *Algorithm 2 guarantees monotonicity of work w.r.t. the cost parameter, i.e., $c_i \geq c_i' \rightarrow W_i(f_i(\theta_i)) \leq W_i(f_i(\theta_i'))$, when the cost per unit charge is constant (i.e., independent of charging speed or state of charge).*

Proof. The structure of this proof, the notation used, and the logical steps are the same as the ones in the proof of Theorem 2. The only difference is the analysis of the four cases that we now present.

Case 1: $SoC_{i,t} \geq SoC_{i,t}'$ and $r_{i,t} > 0$. First note that the upper bound $ub_{i,t}$ never constrains $r_{i,t}'$ to be lower than $r_{i,t}$ because $SoC_{i,t} \geq SoC_{i,t}'$.

The mechanism chooses to charge type θ_i as this option is optimal for her, i.e., given the current buying price p_t^{buy} and the expected future buying and selling prices, it is optimal for θ_i to buy electricity at time t.

Assume that $SoC_{i,t} \geq q_i$ and that the optimal option for θ_i' is $r_{i,t} > r_{i,t}'$, i.e, θ_i' is not charged as much as θ_i at time t. This naturally implies that θ_i' has a better buying opportunity in the future. If not, since $c_i > c_i'$, θ_i' would have charged the battery at t at least as much as θ because this would have guaranteed her at least the same benefit that θ_i gets. However, note that θ_i and θ_i' have the same market opportunities in terms of buying and selling prices because these prices are independent of the agents type. Thus, if there is a better buying option for θ_i' than the current one, the same option would have been available and optimal also for θ_i. But this contradicts the initial assumption. Thus, we proved that $r_{i,t} \leq r_{i,t}'$, i.e., Condition 1 holds.

Now observe that, when $SoC_{i,t} < q_i$, the decision of charging θ_i is taken without considering that the electricity will be sold because $SoC_{i,d_i} < q_i$. Given this, the choice of charging is optimal as long as the state of charge is below q_i. However, this situation correspond to the one discussed in Case 3 of Theorem 2 where we show that the decision of charging θ_i at t never leads to the case in which $q_{i,t^*}^+(r_i) > q_{i,t^*}^+(r_i')$ at a future time t^*, hence Condition 2 holds.

Case 2: $SoC_{i,t} \leq SoC_{i,t}'$ and $r_{i,t} < 0$. The analysis of this case can be easily derived from the one of Case 1, as the two cases are symmetrical.

Case 3: $SoC_{i,t} \leq SoC_{i,t}'$ and $r_{i,t} > 0$. The only difference between this and Case 1 is that $SoC_{i,t} \leq SoC_{i,t}'$. The only effects of this difference is that $r_{i,t} \leq r_{i,t}'$ may not be possible. However, what was discussed in Case 3 of Theorem 2 can be applied here with the only difference that $ub_{i,t} \geq q_i$ (instead of $ub_{i,t} = q_i$).

Case 4: $SoC_{i,t} \geq SoC_{i,t}'$ and $r_{i,t} < 0$. The analysis of this case can be easily derived from the one of Case 3, as the two cases are symmetrical. $\qquad\square$

Theorem 4. *Algorithm 1 and Algorithm 2 guarantee monotonicity of work w.r.t. departure time d_i, i.e., $d_i \leq d_i' \rightarrow W_i(f_i(\theta_i)) \leq W_i(f_i(\theta_i'))$, and required amount of electricity q_i, i.e., $q_i \geq q_i' \rightarrow W_i(f_i(\theta_i)) \leq W_i(f_i(\theta_i'))$.*

Proof. Sketch Let θ_i and θ_i' denote two types which are identical except for either departure time, and in this case, w.l.o.g., we assume $d_i \leq d_i'$, or for required amount of electricity, in this case, w.l.o.g., we assume $q_i \geq q_i'$. The proofs by contradiction used in Theorem 2 and Theorem 3 hold also in case of departure time and required amount of electricity, i.e., if there is no constraint imposed by $ub_{i,t}$ and $lb_{i,t}$ then θ' is (dis)charged at least as θ. However, in this case $ub_{i,t}$ and $lb_{i,t}$

may be different from $ub'_{i,t}$ and $lb'_{i,t}$, i.e., the bounds may depend on the agent type. Nevertheless, the way in which an earlier departure and a higher requested amount of electricity affect these bounds is by making them more strict, i.e., $ub_{i,t} \leq ub'_{i,t}$ and $lb_{i,t} \geq lb'_{i,t}$. The proof follows from the fact that stricter bounds naturally impose stricter constraints on an agent's work. $\qquad\square$

7 Empirical Evaluation

In this section, we empirically evaluate our algorithms in a range of realistic settings, in order to determine their potential for implementing smart car parks. We are particularly interested in two key metrics: social welfare (Equation 3) and aggregator utility (Equation 2).

7.1 Experimental Setup

We simulate a large car park over a period of 24 hours, where most cars arrive in the morning, representing a typical workplace. Specifically, time is discretised into 1-hourly time slots, and the number of arrivals at every hour is drawn from a Poisson distribution with a mean that varies over time and peaks at 8am. Unless specified otherwise, the expected number of arrivals over the entire day is 100 cars, and each car stays in the car park for between 1 and 16 hours (drawn uniformly at random) or until midnight, whichever is earlier.

We discretise a car's battery into units representing 3 kWh each, which we assume is the electricity that can be charged or discharged in one hour. To determine the initial state of charge of agent i, SoC_{i,a_i}, and its maximum state of charge, SoC_i^{max}, we draw two integers from the interval 1 to 8 and assign the smaller to SoC_{i,a_i} and the larger to SoC_i^{max}. This means that vehicles have a battery capacity of up to 24 kWh, which is the typical capacity of current EVs. The required charge q_i is determined by randomly picking a feasible charge given the vehicle's duration of stay and maximum capacity. We assume constant capacity losses for charging and discharging a unit of electricity (a in Section 3) and draw these for each agent from a uniform distribution on [\$0.01, \$0.03]. Finally, the cost parameter c_i is drawn from the interval $[1, 5]$, discretised into steps of 0.1.[3]

Grid prices are simulated using a Markov chain that randomly starts in one of four states. Each state s is associated with a buying price $p^{buy,s}$ (respectively, these are \$0.03, \$0.15, \$1.20 and \$3.00 for states 1–4) and a selling price $p^{sell,s} = 0.75p^{buy,s}$. With probability 0.5, this Markov chain remains in the current state and otherwise transitions to one of its direct neighbours. These states represent off-peak and peak times, with some intermediate steps.

To test Algorithm 1, we consider a range of thresholds based on the prices that can occur in the various states. Specifically, we set $\rho^{low} = \rho^{high} = p^{buy,s}$ for each state s. In addition to Algorithms 2 and 3, we also consider a *NoV2G* algorithm, which does not allow discharging but charges to minimise the overall expected cost. This represents a base-

---

[3]This is to facilitate the computation of the transfers, as given in Theorem 1. In practice, the method described in [Babaioff *et al.*, 2010] could be used instead.

line benchmark against which we compare our algorithms, in order to quantify the benefit of using V2G for smart car parks.

7.2 Results

In the following, we vary a range of system parameters to test our algorithms in different settings. First, we are interested in how the number of vehicles in the car park affects the performance of the algorithms. To this end, we vary the expected number of arrivals over the course of the day from 10 to 300. Results are shown in Fig. 1a (social welfare) and 1b (aggregator utility).

Here, it is clear that Algorithm 1 performs poorly, even for the best threshold setting. This is mainly because the algorithm does not use a model of how prices evolve in the future and so takes poor decisions that lead to costly charging actions in the future. However, for the right choice of parameters, Algorithm 1 can perform better than the *NoV2G* baseline, in terms of both aggregator utility and social welfare. Algorithms 2 and 3 perform significantly better than any of the other benchmarks, achieving a positive aggregator utility in most cases. This means that the costs for buying electricity and paying vehicles for providing their battery for V2G are fully covered by the gains of selling electricity when prices are high.

While the utility is low in some settings, it is important to note that the aggregator here does not actually demand a payment for the electricity the agents receive or the parking service itself. In practice, this would be added (e.g., at a fixed rate per unit of required charge and per time parked), increasing the utility further. Furthermore, while the utility and welfare of Algorithms 2 and 3 are similar, their respective behaviours are not. Algorithm 2 tends to trade more with the grid, while Algorithm 3 relies more on charging and discharging internally.

Next, we consider how the availability of agents in the market affects the performance of the system. Specifically, Fig. 1c shows the effect of varying the maximum duration of an agent's stay in the car park. As welfare and aggregator utility are similar in all settings we tested (with utility being consistently higher), we now concentrate on the latter. Here, two opposing effects can be observed. Initially, as the duration increases, agents demand more charge as their feasible options increase, leading to a drop in welfare and aggregator utility (as noted above, this is partially because agents do not pay for the electricity they receive in this setup). However, as it increases further, the flexibility of agents increases, because they have more time beyond what is required to reach their desired charge. This leads to more opportunities for V2G, increasing the aggregator's utility.

We now turn to how features of the electricity market affect the performance of the V2G car park. To investigate this, we vary the probability of staying within a given price state, indicating how static the market is. This is shown in Fig. 1d, where it is clear that a more dynamic market (lower probability), leads to a higher utility for the aggregator. This is because there are more opportunities for buying at low prices and selling at high prices. A static market, in turn, means that the market may remain at a high price state for the entire day, leading to high electricity costs.
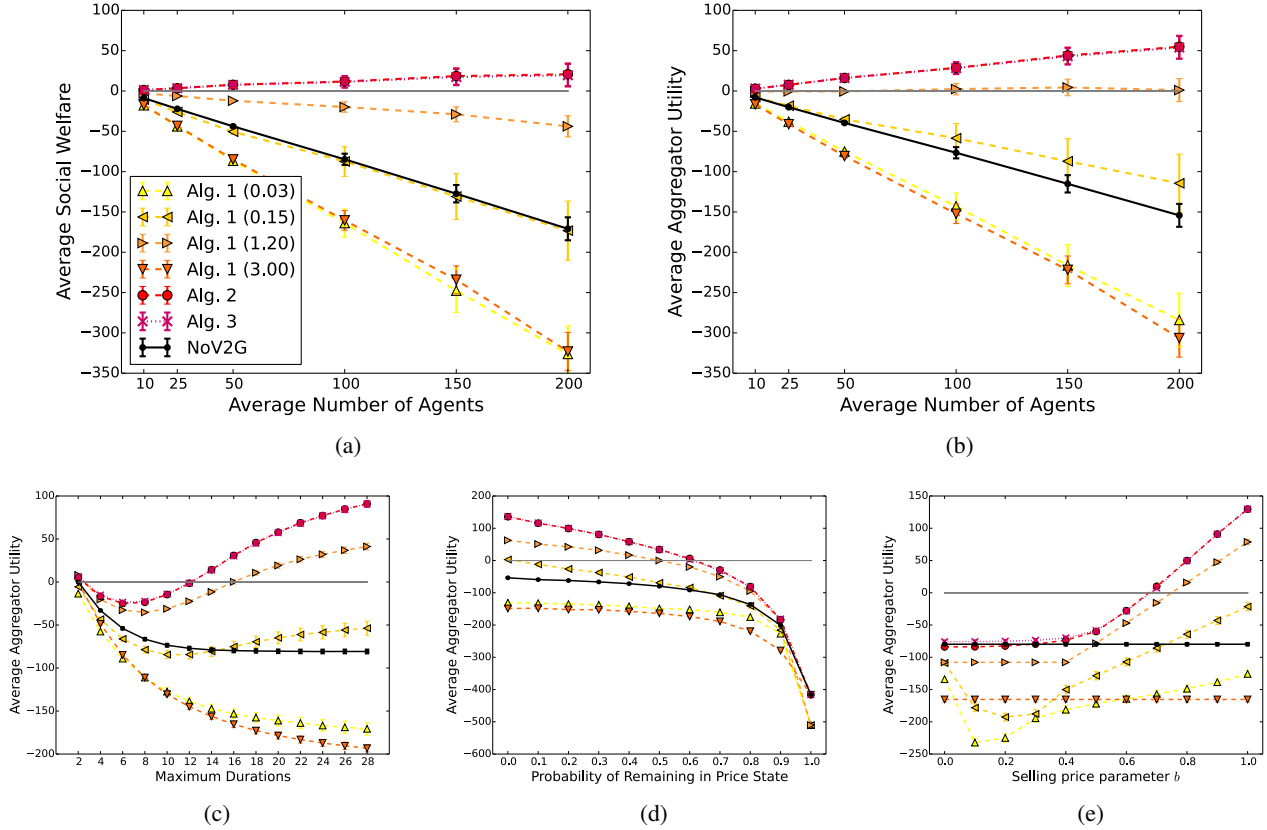
(a)   (b)

(c)   (d)   (e)

Figure 1: Results (with 95% confidence intervals, based on 1000 samples for each data point) in various settings. 1a and 1b show the social welfare and aggregator utility, respectively, when the number of agents is increased. 1c shows the effect of agents with increasing availability for charging, 1d varies the dynamism of prices and 1e varies the gap between buying and selling prices.

Finally, we investigate the relationship between buying and selling prices in Fig. 1e, by setting the selling prices to $p^{sell,s} = bp^{buy,s}$ and varying b from 0 to 1. Clearly, a higher selling price results in higher aggregator utility, and it is also evident that Algorithms 2 and 3 perform similarly to the *NoV2G* algorithm when prices are too low. Interestingly, Algorithm 3 here outperforms Algorithm 2 slightly, as it enables agents to trade electricity internally.

8 Conclusions and Future Work

To the best of out knowledge, this is the first work that formalises the V2G problem as an online mechanism design problem. We identify sufficient conditions for truthfulness, design three algorithms for this scenario, two of which are proved to be truthful, and empirically evaluate them.

There are several extensions we leave to future work. First, we aim to model more realistic battery types, where the cost increases with the number of charging-discharging cycles and costs are non-linear in the (dis)charged amount of electricity. Other issues relevant for practice include losses during discharging, and batteries with different round-trip efficiencies. Second, a range of other algorithms could be designed or adapted for this problem, such as those recently developed in the field of online stochastic combinatorial optimisation [Scott *et al.*, 2013]. Finally, we aim to evaluate our techniques using a real case study and data, such as modelling the arrival and departure of EVs in a shopping center.

Acknowledgements

References

[Archer and Tardos, 2001] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2001.

[Babaioff *et al.*, 2010] Moshe Babaioff, Robert D. Kleinberg, and Aleksandrs Slivkins. Truthful mechanisms with implicit payment computation. In *Proc. of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 43–52, New York, NY, USA, 2010. ACM.

[Divya and Ostergaard, 2009] K.C. Divya and Jacob Ostergaard. Battery energy storage technology for power systemsan overview. *Electric Power Systems Research*, 79(4):511 – 520, 2009.

[Gerding *et al.*, 2013] Enrico H. Gerding, Sebastian Stein, Valentin Robu, Dengji Zhao, and Nicholas R. Jennings. Two-sided online markets for electric vehicle charging. In *Proc. of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '13)*, pages 989–996, 2013.

[Han *et al.*, 2010] Sekyung Han, Soohee Han, and K. Sezaki. Development of an optimal vehicle-to-grid aggregator for frequency regulation. *IEEE Transactions on Smart Grid*, 1(1):65–72, 2010.

[Kahlen and Ketter, 2015] Micha Kahlen and Wolfgang Ketter. Aggregating electric cars to sustainable virtual power plants: The value of flexibility in future electricity markets. In *Proc. of the 29th AAAI Conference on Artificial Intelligence*, 2015.

[Kamboj *et al.*, 2011] Sachin Kamboj, Willett Kempton, and Keith S. Decker. Deploying power grid-integrated electric vehicles as a multi-agent system. In *Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, pages 13–20. IFAAMAS, 2011.

[Kempton and Tomić, 2005] Willett Kempton and Jasna Tomić. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1):268 – 279, 2005.

[Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge, 2007.

[Recalde Melo *et al.*, 2014] D.F. Recalde Melo, Gooi Hoay Beng, and T. Massier. Charging of electric vehicles and demand response management in a singaporean car park. In *Proc. of the 49th International Universities Power Engineering Conference (UPEC)*, pages 1–6, 2014.

[Rigas *et al.*, 2013] E.S. Rigas, S.D. Ramchurn, N. Bassiliades, and G. Koutitas. Congestion management for urban ev charging systems. In *Smart Grid Communications*, pages 121–126, 2013.

[Robu *et al.*, 2013] Valentin Robu, Enrico H. Gerding, Sebastian Stein, David C. Parkes, Alex Rogers, and Nicholas R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *Journal of Artificial Intelligence Research*, 48:175–230, 2013.

[Scott *et al.*, 2013] Paul Scott, Sylvie Thiébaux, Menkes Briel, and Pascal Hentenryck. Residential demand response under uncertainty. In *Proc. of the 19th International Conference on Principles and Practice of Constraint Programming (CP 2013)*, pages 645–660. Springer Berlin, 2013.

[Shafie-khah *et al.*, 2016] M. Shafie-khah, E. Heydarian-Forushani, M.E.H. Golshan, P. Siano, M.P. Moghaddam, M.K. Sheikh-El-Eslami, and J.P.S. Catalo. Optimal trading of plug-in electric vehicle aggregation agents in a market environment for sustainability. *Applied Energy*, 162:601 – 612, 2016.

[Shi and Wong, 2011] Wenbo Shi and Vincent W.S. Wong. Real-time vehicle-to-grid control algorithm under price uncertainty. In *Smart Grid Communications*, pages 261–266, 2011.

[Stein *et al.*, 2012] Sebastian Stein, Enrico H. Gerding, Valentin Robu, and Nicholas R. Jennings. A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, pages 669–676, 2012.

[Valogianni *et al.*, 2015] Konstantina Valogianni, Wolfgang Ketter, and John Collins. A multiagent approach to variable-rate electric vehicle charging coordination. In *Proc. of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '15)*, 2015.

[Vasirani *et al.*, 2013] M. Vasirani, R. Kota, R.L.G. Cavalcante, S. Ossowski, and N.R. Jennings. An agent-based approach to virtual power plants of wind power generators and electric vehicles. *IEEE Transactions on Smart Grid*, 4(3):1314–1322, 2013.