

Optimization of Charging Strategies for Electric Vehicles in PowerMatcher-Driven Smart Energy Grids

Pia Kempker
TNO, Performance of
Networks and Systems
The Netherlands
pia.kempker@tno.nl

Nico van Dijk
University of Twente
Dept. of Applied Mathematics
The Netherlands
n.m.vandijk@utwente.nl

Werner Scheinhardt
University of Twente
Dept. of Applied Mathematics
The Netherlands
w.r.w.scheinhardt@utwente.nl

Hans van den Berg
TNO, Performance of
Networks and Systems,
University of Twente
Dept. of Computer Science
j.l.vandenberg@tno.nl

Johann Hurink
University of Twente
Dept. of Applied Mathematics
The Netherlands
j.l.hurink@utwente.nl

ABSTRACT

A crucial challenge in future smart energy grids is the large-scale coordination of distributed energy demand and generation. The well-known PowerMatcher is a promising approach that integrates demand and supply flexibility in the operation of the electricity system through dynamic pricing and a hierarchical bidding coordination scheme. However, as the PowerMatcher focuses on short-term coordination of demand and supply, it cannot fully exploit the flexibility of e.g. electric vehicles over longer periods of time. In this paper, we propose an extension of the PowerMatcher comprising a planning module, which provides coordinated predictions of demand/price over longer times as input to the users for determining their short-term bids. The optimal short-term bidding strategy minimizing a user's costs is then formulated as a Stochastic Dynamic Programming (SDP) problem. We derive an analytic solution for this SDP problem leading to a simple short-term bidding strategy. Numerical results using real-world data show a substantial performance improvement compared to the standard PowerMatcher, without significant additional complexity.

CCS Concepts

Computing methodologies → Multi-agent planning; Mathematics of computing → Stochastic processes; Hardware → Smart grid;

Keywords

smart grids, stochastic dynamic programming, market-based coordination, electric vehicles, PowerMatcher

1. INTRODUCTION

The large scale coordination of demand response (see e.g. [9]), distributed generation and electricity storage will be crucial for power systems management in future smart energy grids. The PowerMatcher, a recognized technology, integrates demand and supply flexibility in the operation of the electricity system through the use of dynamic pricing and a hierarchical bidding coordination scheme, see e.g. [4, 7, 1, 11, 6]. This coordination scheme provides the PowerMatcher with attractive properties regarding scalability, stability and privacy [4, 3], which are prerequisites for practical usefulness. Recent field experiences and simulation studies (see [5, 14]) show the potential of the technology for network operations (e.g. congestion management and black-start support) and for market operations (e.g. virtual power plant operations).

As the PowerMatcher focuses on coordination of demand and supply on the short-term, it has limited means to fully exploit the flexibility of shiftable demands of electric power over longer periods of time (e.g. from electric vehicles) and to achieve the efficiency potentially attainable due to this flexibility. Therefore, as a next step, an extension of the PowerMatcher including forecasting and planning is investigated in this paper. The proposed extension comprises a planning module, which provides coordinated predictions of the demand/price over longer times e.g. a day ahead. This planning module uses a similar scheme as employed by the PowerMatcher for short-term coordination of demand and supply, thus preserving the important scalability, stability and privacy features. The output of the planning module then serves as additional input to the users for optimizing their short-term bids so as to meet their targets (e.g. charging their electric vehicle within a given time period) while keeping the charging costs the lowest possible.

Accordingly, the contribution of this paper is three-fold. In the first place, it describes the above mentioned 'two-time-scale' extension of the PowerMatcher in detail and discusses its main features.

Secondly, it addresses the question how the output of the planning module (an estimate of the long-term average

price) can be optimally exploited for establishing a short-term bidding strategy for a user with a certain amount of shiftable demand over the longer term (e.g. one night). Dealing with the lack of any additional information, we make the assumption that prices for the short-term intervals covered by the planning period are independent and identically distributed according to a normal distribution with mean value equal to the long-term average. We then show that finding the optimal short-term bidding strategy can be formulated as a Stochastic Dynamic Programming (SDP) problem. For this SDP problem, which can anyway be solved numerically, an explicit analytic solution is derived which in turn leads to a simple short-term bidding strategy. It is also explained how the SDP approach can be extended further for situations where more detailed information about (the fluctuations of) the prices during the planning period is available to the users for deciding about their short-term bids.

In the third place, a validation of the proposed PowerMatcher extension is provided. Numerical results show its performance compared to the ‘standard’ PowerMatcher and various other charging strategies. In particular, using simulations with detailed real-world data for wind production and household demands over a six months period, it is shown that the PowerMatcher extension with the simple bidding strategy works very well and in many cases provides a considerable improvement over the standard PowerMatcher.

Moreover, it appears that for the considered scenarios the outcomes of the PowerMatcher extension (even with the quite crude modeling assumptions mentioned above) are surprisingly close to the (theoretical) optimum achievable when all realized prices would have been known in advance. The latter result implies, at least within the present framework, that more precise estimations of the prices for successive short-term time intervals are not imperative since the algorithm is sufficiently robust w.r.t. the modeling assumptions.

The remainder of this paper is structured as follows. In Section 2, the PowerMatcher and its planning module extension (the two-time-scale PowerMatcher) are presented and discussed. Next, in Section 3, we provide the SDP formulation and its analytic solution leading to an explicit short-term bidding strategy for the extended PowerMatcher. Section 4 provides the numerical evaluation and validation, and Section 5 contains conclusions and topics for further research.

2. TWO-TIME-SCALE POWERMATCHER

The PowerMatcher ([4, 7, 5]) provides a very efficient and scalable market-based coordination scheme for *short-term* demand response in smart grids. Many devices, however, offer demand flexibility over longer time horizons, and the optimal utilization of this flexibility requires coordinated predictions over longer periods of time (e.g. day-ahead or week-ahead horizons): Electric vehicles (EVs) for example, can shift their demand within the time window between being plugged in and leaving again.

In the following, we introduce a coordinated planning extension to the PowerMatcher. The planning module uses information from the household devices along with forecasts of e.g. the wind power to predict the average price over the following long-term period, and then make this prediction available to all agents for use in their short-term strategies. This is achieved by combining two instances of the Power-

Matcher algorithm, using different time scales:

Long-term scale (planning module): Each agent sends a demand profile for the long-term horizon (total expected demand vs. average price). The planning agent then predicts the average price over this time horizon.

Short-term scale (matching module): Each agent sends its current demand capabilities for the short-term horizon. Shiftable devices (e.g. electric vehicles) can make their short-term demand flexibility dependent on the long-term average price. The matching agent then determines the market-clearing price for the next short-term interval.

The idea behind the two time scales is that many devices, e.g. electric vehicles, have a demand which is shiftable over the short term, but inflexible over the long term, e.g. because the vehicle must be fully charged by the end of the night. The bidding profiles for the matching horizon t are illustrated in Figure 2. The main difference in the bidding profiles for the planning horizon T , compared to t , is that shiftable devices (such as EVs) have a constant profile over T . The location of the threshold in Figure 2 depends on the outcome of the planning stage.




Figure 2: Short-term bidding profiles

For implementing the long-term scale, we apply a moving-horizon approach, recomputing the expected average price for the updated time window on a regular basis. In the later sections, where we focus on charging electric vehicles, we use a night-ahead time horizon with a fixed end time, since in the scenarios we consider, the EV has to be charged until the next morning. For other shiftable devices other long-term horizons might be preferable, e.g. a 24-hour-ahead moving horizon for battery storage.

As in [4], the agents are organized in a tree structure, with the planning/matching agent at the root. The combined planning and matching algorithm is as follows:

1. Each agent (illustrated as agent k in Figure 1) estimates its total energy demand profile for the next 24 hours, and sends this information to the next-higher agent in the tree.
2. Each non-leaf agent aggregates the received information and passes it on, and finally the aggregated long-term demand profile reaches the root.
3. The root finds the market-clearing price of the long-term matching problem: This price, which can be viewed as the expected long-term average price, is sent to all agents in the tree.

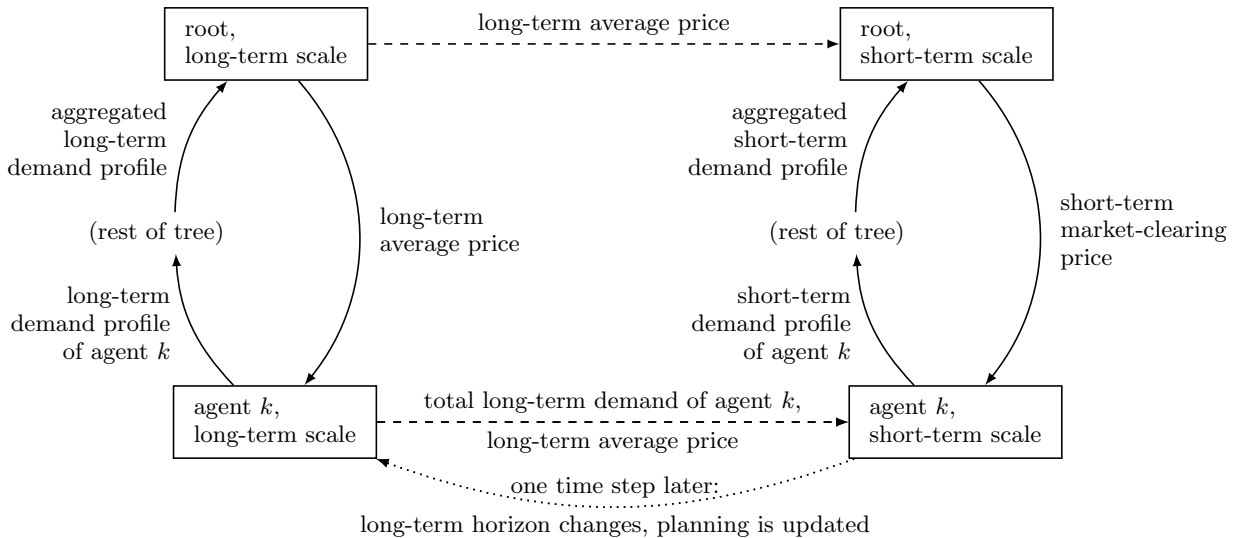


Figure 1: Combined long-term planning and near real-time matching algorithm

4. Knowing the long-term average price, each agent can now determine its long-term target demand.
5. Each agent can use the long-term average price and its long-term target demand as inputs for their bid on the short-term scale.
6. The short-term demand profiles of all agents are again aggregated and passed on, until the root receives the aggregated profile of all agents.
7. The root finds the short-term market-clearing price, and sends this values to all agents in the tree.
8. Each agent uses the market-clearing price to find its short-term target demand, and regulates its energy consumption/production accordingly.
9. One time step later, the time horizon shifts one time step ahead, and the entire process is repeated.

Since the combined planning and matching algorithm consists of a nested application of two instances of the PowerMatcher, its computation and communication requirements are twice that of the PowerMatcher. This means that the computational feasibility and high scalability of the PowerMatcher (see [4]) are retained. This sets the algorithm apart from iterative scheduling methods such as [8, 13], which may need several communication rounds to converge.

As is the case for any market-based coordination scheme, both the PowerMatcher algorithm and its two-time-scale extension are strategy-proof (i.e. robust to individual agents giving false information) if the network is so large that no individual agent can influence the price (i.e. all agents are price takers).

Since the individual demand profiles are aggregated to one combined profile before being passed on at each step in the communication process, a centralized collection of individual demand profiles is not possible. This ensures that the algorithm is in line with privacy concerns.

If predictions about the fluctuation of the demands within the time window are available, the same architecture can be used to estimate the standard deviation of the price over

T . This is the case in the numerical example in Section 4: Detailed wind power predictions can be derived from the weather forecast, and the aggregate fixed household demands follow a predictable daily pattern, from which a standard deviation is easily derived.

In the following sections we study the potential of the two-time-scale PowerMatcher, and discuss how to optimally use the extra information provided by the coordinated planning mechanism implemented in the two-time-scale PowerMatcher. We do this for the example of an electric vehicle in a network powered by wind turbines and a diesel generator (see Figure 3 below).

3. OPTIMAL CHARGING STRATEGIES

For the two-time-scale PowerMatcher with the long-term bidding strategy, we are confronted with the following question: Given the predictions of the long-term average price and its standard deviation, as provided by the planning module of the two-time-scale PowerMatcher, how can this knowledge be used in the creation of short-term demand profiles of shiftable devices?

Since different types of shiftable devices have different objectives and constraints, we focus on the particular case of an electric vehicle (EV) which needs to charge a certain amount overnight. Our objective is:

Objective

How to charge an electric vehicle overnight if the prices for the different time periods during the night are uncertain.

Note that we no longer speak of a ‘demand profile’, but rather about ‘charging’: Finding an appropriate demand profile for a specific time period means that we need to find an appropriate amount of energy to be charged during that period *for each possible value* of the price. Thus, even though we do not know the value of the price beforehand, we can treat it as given, and then concentrate on the objective above, where future prices are uncertain, but the current price is known.

This section will approach this objective in a number of

conceptual steps in order to determine an optimal charging strategy. In Section 3.1 the problem will be formulated as a general Stochastic Dynamic Programming (SDP) problem, using the modeling assumption that prices for different periods are independent. This assumption is needed to get a computationally attractive form.

Next in Section 3.2, we focus on i.i.d. prices, i.e., we now also assume that different periods have the same price distribution, as is in principle the case in the context of the two-time-scale PowerMatcher, where we only know the long run average price. For this special case the SDP is shown to have an analytic solution.

Finally, since the explicit solution is based on order statistics, which are in general not easy to compute, we present in Section 3.3 a simple heuristic which is more practical to implement within the two-time-scale PowerMatcher. Both the modeling assumptions and the heuristic are justified in simulations presented in Section 4.




Figure 3: Schematic representation of the scenario considered in this paper

For illustrative purposes, we consider the concrete setting of an EV which has to be charged with 8kWh of energy overnight (8pm – 8am), and which is connected to a network including wind and fossil fuel power (see Figure 3).

3.1 SDP problem formulation

If the energy prices were known for each period, then the charging problem could be regarded as a deterministic Knapsack Problem (see any introductory OR book for standard Knapsack formulations). However, from a single user’s point of view, these prices are not known in advance: They depend on external factors, such as weather conditions, and on other users’ demand profiles.

A stochastic modeling approach is therefore proposed to incorporate the complex bidding and price setting process. Accordingly, we introduce random variables P_t for the price per unit of energy in period t , where $t = 1, \dots, T$. Later on, in the numerical experiments of the two-time-scale PowerMatcher in Section 4, we will choose suitable distributions for these, with its expectation matching the long run expected average price from the planning module. For presentational convenience from now on we consider these prices to be given in discrete units (i.e., they will be represented by discrete random variables).

More precisely, our objective now is to minimize the total cost to charge an EV within T time intervals (e.g. for a one-night period, $T = 24$ periods of half an hour). At times $t = 1, 2, \dots, T$ a decision has to be made how much is to be charged during period t (i.e., within time interval $[t, t + 1)$), with a charging constraint. Let

P_t be the stochastic price variable for period t ,

L be the total amount of energy to be charged by the end of period T (i.e. during periods $1, 2, \dots, T$)

x_t be the amount still to be charged at time t , (i.e. during periods t, \dots, T),

u_t be the amount of energy to be charged during period t (decision variable), and

u_{\max} be the maximum amount of energy that can be charged during any period.

As visualized in Figure 4 for a simple example the charging problem inherently has the structure of a Decision Tree, or more precisely of a Markov Decision Problem (MDP) (see [10] for an extensive treatment of MDPs) with the following repetitive decision structure: Given the state at a particular time t , a decision is to be taken, due to which:

- There are immediate expected costs for period t
- The resulting state at time $t + 1$ is determined stochastically

Our goal is to determine optimal decisions u_t which will depend on the actual state at time t . This state description should contain sufficient information to determine the next state by a stochastic description. For our application, as opposed to standard knapsack formulations, this means that the state description should not only contain the amount x_t to be charged, but also the (‘known’) current price p_t . Hence, the state at time t is given by $(x_t; p_t)$.

Next let $V_t(x; p)$ be the minimal expected cost during time intervals t, \dots, T , given the state at time t is $(x; p)$ (i.e., given that we need to charge x during $\{t, \dots, T\}$ and that the current price is $P_t = p$).

The optimal decisions for all states can be determined by iteratively solving the Stochastic Dynamic Programming equations: First, for $t = T + 1$ we have

$$V_{T+1}(x) = \begin{cases} 0 & \text{if } x = 0 \\ \infty & \text{otherwise} \end{cases} \quad (1)$$

Next, for $t = T, T - 1, \dots, 1$ and any $(x; p)$ we have

$$V_t(x; p) = \min_u \left[up + \sum_{p'} \mathbb{P}(P_{t+1} = p') V_{t+1}(x - u; p') \right]. \quad (2)$$

Here the sum is taken over all p' in the support of P_{t+1} , and the minimum over all u in $[0, u_{\max}]$. Finally, the charging problem with an amount of L to be charged over $[1, T + 1]$ and with price p at $t = 1$ is denoted by $V(L; p)$. The decision variable (or ‘action’, or ‘control’) u which minimizes (2) will be given by $u_t(x; p)$ and can be said to provide an optimal strategy. Thus, formally we define $u_t(x; p)$ as the amount we need to charge at time t to minimize the total expected cost, when we need to charge x during $\{t, \dots, T\}$ and $P_t = p$.

Returning to Figure 4, this figure also illustrates the example given earlier with $T = 24$, $L = 8$, $u_{\max} = 2$, and P_1, \dots, P_{24} i.i.d. with price distribution $\mathbb{P}(P_t = p) = 1/3$ for $p = 4, 5, 6$. Using (2) it is straightforward to compute the optimal value of $V_0(8) = 30.2$.

Remark 3.1. (Penalty costs) Instead of a strict charging requirement at the end of period T , we can also implement a penalty function for an rest amount not charged, e.g. by a fixed and a proportional penalty for $x_{T+1} = x_T - u_T$

$$M \mathbb{1}_{x_{T+1} > 0} + x_{T+1} b.$$

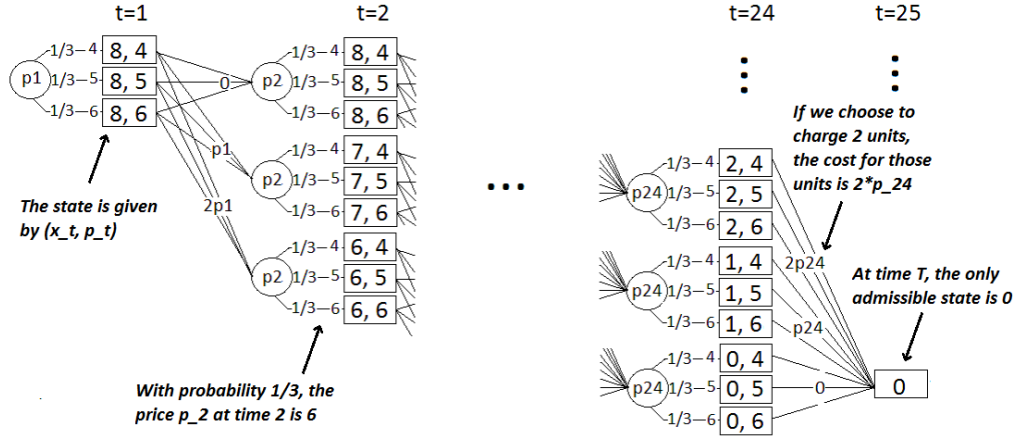


Figure 4: Decision tree for charging 8kWh over 24h, with a discrete price distribution $\mathbb{P}(P_t = p) = 1/3$ for $p = 4, 5, 6$, and possible charging decisions $u_t \in \{0, 1, 2\}$.

Remark 3.2. (History-dependent prices) The formulation of our SDP can easily be adapted for the case of history-dependent prices. In that case the state of the system should include, apart from x_t and p_t , also the price history $h_t = (p_1, \dots, p_{t-1})$. Furthermore, the probabilities in (2) need to be replaced by history-dependent conditional probabilities, so that (2) becomes

$$V_t(x_t; p_t, h_t) = \min_u [up_t + \sum_{p'} \mathbb{P}(P_{t+1} = p' | P_t = p_t, h_t) V_{t+1}(x_t - u; p', h_{t+1})].$$

However, the joint distribution of P_1, P_2, \dots, P_T will generally not be available (since this would require much more communication than provided in the two-time-scale PowerMatcher) and even if it is, the computation will be computationally prohibitive or will require a special approximate procedure.

By the independence assumption and the SDP approach, optimal strategies and corresponding values can now be computed directly and be numerically implemented in the two-time-scale PowerMatcher. Since we only have information available about the long-run average price, we assume that all prices will be i.i.d. with this expectation. In the next section we will see that in this case an explicit analytic optimal decision rule can be given, with its corresponding minimal cost. Note that real-world energy prices are in general not i.i.d., but follow daily and seasonal patterns. However, due to the robustness of the PowerMatcher algorithm, possible performance improvements due to more accurate models are limited (see Section 4).

3.2 Analytic solution for i.i.d. prices

In this section, we assume that the prices P_1, \dots, P_T are independent and identically distributed (i.i.d.) random variables. The particular choice of the underlying distribution is not of interest in this context, but for the two-time-scale PowerMatcher its expectation should match the long-run expected average cost resulting from the long-term bidding process.

Recall the definition of $V_t(x; p)$, and $u_t(x; p)$, and the SDP formulation in (2). To find explicit expressions for V_t and u_t for $t = 1, \dots, T$ we use the concept of order statistics, defined

as follows. Considering the prices P_t, \dots, P_T , we denote the k -th smallest value of these by $P_{(k)}^t$ for $k = 1, \dots, T - t + 1$, so that we have

$$P_{(1)}^t \leq P_{(2)}^t \leq \dots \leq P_{(T-t+1)}^t.$$

Thus, in particular $P_{(1)}^t$ and $P_{(T-t+1)}^t$ are respectively the minimum and maximum prices during $\{t, \dots, T\}$. For notational convenience we also define $P_{(0)}^t = 0$ and $P_{(T-t+2)}^t = \infty$. When some prices coincide, the ordering is not unique, but this is not important since we only need their values in the sequel. In fact we only need the expected values of the order statistics, but note that these depend on the whole price distribution.

An optimal control law and a corresponding value function are given in the following theorem.

THEOREM 1. Let $k = \lfloor \frac{x}{u_{\max}} \rfloor$. An optimal control law for (2), in the case of i.i.d. prices, is given by

$$u_t(x; p) = \begin{cases} 0 & \text{if } p > \mathbb{E}P_{(k+1)}^{t+1} \\ x - ku_{\max} & \text{if } \mathbb{E}P_{(k)}^{t+1} < p \leq \mathbb{E}P_{(k+1)}^{t+1} \\ \min(u_{\max}, x) & \text{if } p \leq \mathbb{E}P_{(k)}^{t+1} \end{cases} \quad (3)$$

and its corresponding value function (minimal cost) is

$$V_t(x; p) = \infty \quad \text{if } \frac{x}{u_{\max}} > T - t + 1, \text{ or else} \\ V_t(x; p) = u_{\max} \sum_{\ell=1}^k \mathbb{E}[P_{(\ell)}^t | P_t = p] + (x - ku_{\max}) \mathbb{E}[P_{(k+1)}^t | P_t = p]. \quad (4)$$

Intuitive explanation. Suppose we had perfect knowledge of all future prices within the time horizon under consideration: Then we would pick the $k = \lfloor \frac{x}{u_{\max}} \rfloor$ cheapest intervals out of the time horizon of $T - t + 1$ intervals, and charge u_{\max} during each of these intervals. If x is not a multiple of u_{\max} then we would charge the remainder $x - ku_{\max}$ in the next-cheapest interval available.

However, in the setting of this paper we do not have perfect knowledge of all future prices: At each time step, we only know the current price (or a set of possibilities for the bidding function) and (an estimate of) the probability distribution of the future prices. This means we have to work with expectations instead: At time t we decide to charge

u_{\max} if we expect the current given price to be among the k cheapest prices within the next $T - t + 1$ time intervals.

PROOF. W.l.o.g. we let $u_{\max} = 1$ by an appropriate choice of units. We will prove optimality of the control law given in (3) by showing that this control law and the corresponding expression (4) for V satisfy the Dynamic Programming equation in (2). We proceed by induction, as follows.

Base step ($t = T$). The constraint that a total of x needs to be charged by time $T + 1$ translates to

$$V_{T+1}(x; p) = \begin{cases} \infty & \text{if } x > 0 \\ 0 & \text{if } x = 0 \end{cases}$$

for all p . Using this, the minimizing argument from (2) shows that the control law for $t = T$ satisfies $u_T(x; p) = x$ for $x \leq 1$, which is in accordance with (3) since $\mathbb{E}P_{(1)}^{T+1} = \infty$ (minimum of an empty set). Also, for $x > 1$ we can choose any u as the minimizing argument in (2), so the control law given in (3) is indeed a valid optimal control law. Furthermore, from (2) we also find that

$$V_T(x; p) = \begin{cases} \infty & \text{if } x > 1 \\ up & \text{if } x \leq 1 \end{cases}$$

which satisfies statement (4) for $t = T$.

Induction step. Supposing that statements (3) and (4) hold for $t + 1$ and $u_{\max} = 1$, we need to prove that (3) and (4) also hold for t . We define the auxiliary function $v(u)$ as the function to be minimized in (2) and rewrite it as

$$\begin{aligned} v(u) &= up + \mathbb{E}_{P_{t+1}} \left[\sum_{\ell=1}^{\lfloor x-u \rfloor} \mathbb{E}[P_{(\ell)}^{t+1} | P_{t+1}] \right. \\ &\quad \left. + (x - u - \lfloor x - u \rfloor) \mathbb{E}[P_{(\lfloor x-u \rfloor + 1)}^{t+1} | P_{t+1}] \right] \quad (5) \\ &= up + \sum_{\ell=1}^{\lfloor x-u \rfloor} \mathbb{E}P_{(\ell)}^{t+1} \\ &\quad + (x - u - \lfloor x - u \rfloor) \mathbb{E}P_{(\lfloor x-u \rfloor + 1)}^{t+1} \end{aligned}$$

where the expectation $\mathbb{E}_{P_{t+1}}$ is w.r.t. P_{t+1} , the other expectations in the first line are w.r.t. P_{t+2}, \dots, P_T , and the expectations in the last line are w.r.t. $P_{t+1}, P_{t+2}, \dots, P_T$. To minimize $v(u)$, note that

$$\lfloor x - u \rfloor = \begin{cases} \lfloor x \rfloor, & \text{for } u \in [0, x - \lfloor x \rfloor] \\ \lfloor x \rfloor - 1, & \text{for } u \in (x - \lfloor x \rfloor, 1] \end{cases}, \quad (6)$$

and hence

$$\begin{aligned} \frac{\partial}{\partial u} v(u) &= p - \mathbb{E}P_{(\lfloor x-u \rfloor + 1)}^{t+1} \\ &= \begin{cases} p - \mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1} & \text{if } u \in [0, x - \lfloor x \rfloor] \\ p - \mathbb{E}P_{(\lfloor x \rfloor)}^{t+1} & \text{if } u \in (x - \lfloor x \rfloor, 1] \end{cases}. \end{aligned}$$

Since $v(u)$ is continuous w.r.t. $u \in [0, 1]$ and $\mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1} \geq \mathbb{E}P_{(\lfloor x \rfloor)}^{t+1}$ there are three cases:

- If $p > \mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1}$ then $\frac{\partial v}{\partial u} > 0 \forall u \in [0, 1]$ and hence the minimum is attained at $u = 0$.
- If $\mathbb{E}P_{(\lfloor x \rfloor)}^{t+1} < p \leq \mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1}$ then $\frac{\partial v}{\partial u} \leq 0$ for $u \in [0, x - \lfloor x \rfloor]$ and $\frac{\partial v}{\partial u} > 0$ for $u \in (x - \lfloor x \rfloor, 1]$, and hence the minimum is attained at $u = x - \lfloor x \rfloor$.
- If $p \leq \mathbb{E}P_{(\lfloor x \rfloor)}^{t+1}$ then $\frac{\partial v}{\partial u} \leq 0 \forall u \in [0, 1]$, and the minimum is attained at $u = 1$ (or $u = x$ if $x < 1$).

This proves the optimality of (3). Now we can prove (4):

$$\begin{aligned} V_t(x; p) &= \min_u v(u) \\ &= \begin{cases} \sum_{\ell=1}^{\lfloor x \rfloor} \mathbb{E}P_{(\ell)}^{t+1} + (x - \lfloor x \rfloor) \mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1} & \text{if } p > \mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1} \\ (x - \lfloor x \rfloor)p + \sum_{\ell=1}^{\lfloor x \rfloor} \mathbb{E}P_{(\ell)}^{t+1} & \text{if } \mathbb{E}P_{(\lfloor x \rfloor)}^{t+1} < p \leq \mathbb{E}P_{(\lfloor x \rfloor + 1)}^{t+1} \\ p + \sum_{\ell=1}^{\lfloor x \rfloor - 1} \mathbb{E}P_{(\ell)}^{t+1} + (x - \lfloor x \rfloor) \mathbb{E}P_{(\lfloor x \rfloor)}^{t+1} & \text{if } p \leq \mathbb{E}P_{(\lfloor x \rfloor)}^{t+1} \end{cases} \\ &= \sum_{\ell=1}^{\lfloor x \rfloor} \mathbb{E}[P_{(\ell)}^t | P_t = p] + (x - \lfloor x \rfloor) \mathbb{E}[P_{(\lfloor x \rfloor + 1)}^t | P_t = p] \end{aligned}$$

To find the formulas for $u_{\max} \neq 1$ we replace x and u by $\frac{x}{u_{\max}}$ and $\frac{u}{u_{\max}}$ respectively, and multiply all prices with u_{\max} . \square

The control law given here is almost a bang-bang type strategy, in the sense that once x_t is a multiple of u_{\max} , the optimal law $u \in [0, u_{\max}]$ is either 0 or u_{\max} .

3.3 Heuristic

The implementation of the control law in (3) requires us to compute expectations of order statistics. This is a difficult task in terms of computational complexity. Therefore we propose an appealing heuristic which may replace the optimal rule, as follows. With $k = \lfloor \frac{x}{u_{\max}} \rfloor$, let the heuristic control law be given by

$$u_t(x; p) = \begin{cases} \min(u_{\max}, x) & \text{if } F(p) \leq \frac{k}{T-t+1} \\ x - ku_{\max} & \text{if } \frac{k}{T-t+1} < F(p) \leq \frac{k+1}{T-t+1} \\ 0 & \text{if } F(p) > \frac{k+1}{T-t+1}, \end{cases} \quad (7)$$

where F is the common price distribution function. At first sight this may seem equivalent to the optimal rule in (3): For strictly monotone F we have

$$p \leq \mathbb{E}P_{(k)}^{t+1} \Leftrightarrow F(p) \leq F(\mathbb{E}P_{(k)}^{t+1}).$$

$F(P_{(k)}^{t+1})$ is a random variable which is distributed as the (k) -th order statistic of $T - t$ i.i.d. uniform random variables on $[0, 1]$, so $\mathbb{E}F(P_{(k)}^{t+1}) = \frac{k}{T-t+1}$. Hence the heuristic control law (7) would be equivalent to the optimal control law (3) if $F(\mathbb{E}P_{(k)}^{t+1}) = \mathbb{E}F(P_{(k)}^{t+1})$. This is unfortunately not true in general.

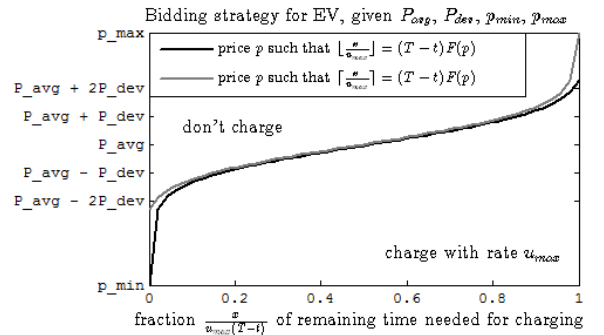


Figure 5: Bidding strategy of EV, with the threshold control law given in (7)

Even though the difference between the two rules may not be negligible, we still use it in the next section with good results, thus strengthening our belief that the two-time-scale PowerMatcher is a strong concept which seems to be robust w.r.t. the control law that is used.

The heuristic bidding strategy is illustrated in Figure 5. The distance between the thresholds $\lfloor \frac{x}{u_{\max}} \rfloor$ and $\lceil \frac{x}{u_{\max}} \rceil$ depends on x_0 and u_{\max} (in Figure 5, $x_0 = 8$ and $u_{\max} = \frac{1}{2}$).

4. NUMERICAL RESULTS

In this section, we compare the performance of different short-term charging strategies in simulations using real-world data for wind production and fixed household demands.

The data we use is available at [2], and consists of 15-minute averages for the network of Belgium over the first 180 days of 2015, of:

- the measured load, aggregated over all households and industries (we will consider this the fixed demand),
- the wind production, measured in some locations and upscaled to fit the network’s nominal capacity,
- long-term predictions of the aggregated load and of the wind production, updated once a day at 11am for all 15-minute intervals in the next 24 hours.

The Belgian network includes 4.8% wind production and 3.3% solar production: We exclude solar production here since the only shiftable devices considered in the simulation are electric vehicles, and they charge overnight when the solar production is zero. In order to compensate for this, and in order to incorporate the sharp increase in wind production over the recent and coming years, we scale the overall loads by a factor of $\frac{1}{5}$, while keeping the wind power as is.

For the purposes of this simulation, electric vehicles can charge from 8pm till 8am, and their demand is fixed at 8kWh/night. An extensive study of realistic driving behaviors can be found in [12].

In our example network, power is produced by wind turbines (generating cheap electricity whenever there is wind) and a diesel generator (which has virtually unlimited capacity but high unit prices). We add one EV to the network, in addition to the fixed loads provided in the data. Since the demand of the EV is very small compared to the fixed load, the EV is a *price taker*: while the current energy price influences the EV’s decisions, the EV’s actions have no significant influence on the prices. This would change if we added many EVs with similar demand patterns (e.g. one for every household), a possible topic for further research.

We compare the following strategies for charging the EV:

1. Charge as fast as possible, starting when plugged in,
2. Charge evenly over the whole night,
3. Strategy (7), using last night’s average and standard deviation as P_{avg} and P_{dev} (this corresponds to the standard PowerMatcher architecture),
4. Strategy (7), using coordinated night-ahead estimates of P_{avg} and P_{dev} , estimated once at the beginning of the night (this corresponds to the two-time-scale PowerMatcher approach with only one prediction round per night),

5. Strategy (7), using coordinated night-ahead estimates of P_{avg} and P_{dev} , updated every 15 minutes (this corresponds to the two-time-scale PowerMatcher approach)
6. For a lower bound on the cost, we pick the cheapest time slots in retrospect, using perfect knowledge or individual estimates of all prices.

With unit prices $p_{\text{wind}} = 1$ and $p_{\text{diesel}} = 10$, we find the average unit price at time t by calculating

$$p(t) = \frac{p_{\text{wind}} * \text{wind used} + p_{\text{diesel}} * \text{diesel used}}{\text{wind used} + \text{diesel used}}.$$

In Figure 6, the outcomes of the different strategies are illustrated for night 124 of 2015. All strategies lead to different charging times, as indicated by \star in the figure. The costs for charging 8kWh in night 124 of 2015 are:

strategy	1	2	3	4	5	6
cost	45.3	33.4	30.5	28.5	23.9	22.0
improvement	-	26%	33%	37%	47%	51%

Compared to charging when plugged in (strategy 1), the standard PowerMatcher (strategy 3) leads to 33% lower costs. The coordinated planning algorithm introduced in Section 2 (strategy 5) leads to another significant improvement (47% lower cost than strategy 1), and performs only 4% worse than the lower bound (strategy 6).

The total costs for charging one EV every night of the first 180 nights of 2015 are:

strategy	perfect foresight	%	day-ahead predictions	%
1	10861	-	10861	-
2	10672	1.7%	10672	1.7%
3	10626	2.2%	10626	2.2%
4	10111	6.9%	10190	6.2%
5	9841	9.3%	9971	8.2%
6	9683	10.8%	9909	8.8%

Some nights show no difference in costs (e.g. because there is no wind, and hence the unit prices are constant), while in other nights the difference is large (as in night 124). In total, over the first half of 2015 and using imperfect price predictions, the standard PowerMatcher (strategy 3) performs 2.2% better than strategy 1. Another improvement of 4% can be realized by running the two-time-scale PowerMatcher once at 8pm (strategy 4). Using the two-time-scale PowerMatcher every 15 minutes instead of once at 8pm (strategy 5) reduces the costs by another 2%, with only 0.6% difference compared to the lower bound of 9909.

Note that in Section 3.4 we assumed that the distribution $F(p)$ is known: In the simulations we modeled the prices as i.i.d. and normally distributed for the EV strategy. Even though the realizations of the prices resulting from the Belgian data are neither i.i.d. nor normally distributed, the algorithm still leads to large cost reductions, with little room for further improvement. This illustrates the robustness of the PowerMatcher algorithm w.r.t. modeling simplifications, and justifies the modeling assumption of i.i.d. prices.

5. CONCLUSION / FUTURE RESEARCH

In this paper we have developed a powerful extension of the well-known PowerMatcher for coordination of distributed demand and supply in smart energy grids. It achieves a better exploitation of the flexibility of shiftable demands from

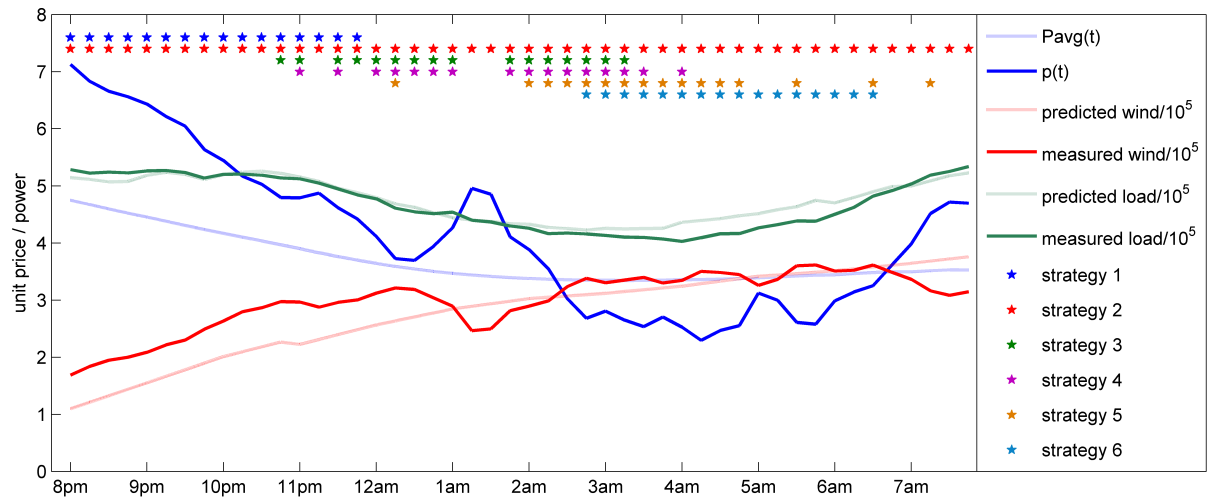


Figure 6: Comparison of different bidding strategies, prices and demand/production for one night

e.g. electric vehicles (EVs), and comprises a planning module together with a strategy for optimal utilization of the information provided by this planning module. The effectiveness of the novel combination of the planning/matching framework and the simple-to-use strategy for charging EVs was verified numerically, using real-world data.

An interesting next step is to run similar simulations with so many EVs that their combined decisions have a significant influence on the price. Another straightforward topic for further research is the extension of the methods used for finding the charging strategies in Section 3 to other shiftable devices (e.g. battery storage), or other types of controllable devices. A further interesting topic is the inherent trade-off between the amount and type of information provided by a planning module and the additional computation and communication required for this.

Acknowledgments

We gratefully acknowledge the contributions of Yvonne Prins, Pamela MacDougall, Koen Kok, and Leon Kester (all affiliated with TNO) to the research leading to this paper.

6. REFERENCES

- [1] G. Basso, N. Gaud, F. Gechter, V. Hilaire, and F. Lauri. A framework for qualifying and evaluating smart grids approaches: Focus on multi-agent technologies. *Smart Grid and Renewable Energy*, 4:333–347, 2013.
- [2] ELIA. Data download (2015). <http://www.elia.be/nl/grid-data/data-download>.
- [3] M. Hoefling, F. Heimgaertner, M. Menth, and H. Bontius. Traffic estimation of the PowerMatcher application for demand supply matching in smart grids. In *International Conference and Workshops on Networked Systems (NetSys)*, p. 1–6, 2015.
- [4] K. Kok. *The PowerMatcher: Smart coordination for the smart electricity grid*. PhD thesis, Vrije Universiteit, The Netherlands, 2013.
- [5] K. Kok, B. Roossien, P. MacDougall, O. van Pruissen, G. Venekamp, R. Kamphuis, J. Laarakkers, and C. Warmer. Dynamic pricing by scalable energy management systems – Field experiences and simulation results using PowerMatcher. In *IEEE Power and Energy Society General Meeting*, p. 1–8, 2012.
- [6] W. Lausenhammer, D. Engel, and R. Green. A game theoretic software framework for optimizing demand response. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT Europe)*, p. 1–5, 2015.
- [7] P. MacDougall, C. Warmer, and K. Kok. Mitigation of wind power fluctuations by intelligent response of demand and distributed generation. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT Europe)*, pages 1–6, 2011.
- [8] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, 2010.
- [9] A. Molderink, V. Bakker, J. L. Hurink, and G. J. M. Smit. Comparing demand side management approaches. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT Europe)*, p. 1–8, 2012.
- [10] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [11] S. Rafiei and A. Bakhshai. A review on energy efficiency optimization in smart grid. In *38th Annual Conference on IEEE Industrial Electronics Society (IECON)*, p. 5916–5919, 2012.
- [12] V. Silva et. al. Estimation of innovative operational processes and grid management for the integration of EV. Project deliverable D6.2, 2011.
- [13] T. van der Klauw, M. E. Gerards, G. J. Smit, and J. L. Hurink. Optimal scheduling of electrical vehicle charging under two types of steering signals. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT Europe)*, p. 122, 2014.
- [14] J. P. Wijbenga, P. MacDougall, R. Kamphuis, T. Sanberg, A. van den Noort, and E. Klaassen. Multi-goal optimization in PowerMatching city: A smart living lab. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT Europe)*, p. 1–5, 2014.