

Balancing Trade-offs in Coordinated PHEV Charging with Continuous Market-based Control

Klaas De Craemer, *Member, IEEE*, and Geert Deconinck, *Senior Member, IEEE*

Abstract—As the use of Plug-in Hybrid Electrical Vehicles is expected to rise, their rather high energy consumption and long charging times has the potential to impact power grid stability. While coordinated charging is generally considered the answer to such issues, often overlooked aspects in coordination schemes are reliability, fault tolerance and ease of implementation. In our research we adapt an existing market-based multi-agent coordination scheme to a real-world environment. The step from time-slots to an asynchronous and continuously moving scheme brings it closer to how an effective implementation would work and offers advantages regarding agent communication load and response time.

Index Terms—smart grids, PHEVs, coordinated charging, multi-agent systems

I. INTRODUCTION

THE future of electricity grids is shifting towards the concept of intelligent or smart grids. In a smart grid, beside other aspects, consumer appliances communicate with each other and energy suppliers to avoid peak loads by shifting their consumption in time. In particular, (PH)EVs are very suited for this type of control due to their high energy requirements and long charging times. This coordinated charging is to be seen as a business by charging service companies; the latter aggregate many Plug-in Hybrid Electrical Vehicles (PHEVs) in a fleet and can therefore act as a single trader on the energy market(s). By intelligently shifting charging schedules and avoiding peaks, electricity can be bought at lower tariffs.

To control charging PHEVs in a smart grid, several methodologies have been proposed [1] [2] [3] [4] [5] [6], but an attribute applicable to most is the fact that they make use of timeslots, e.g. time is divided into discrete time-slots of 15 minutes. While this is similar to the way energy is traded on the markets, in the real world, devices (and especially charging vehicles) operate continuously. Control and coordination actions should commence immediately, especially for fast-charging applications, and allow for quickly altering the fleets behavior if the need arises. Additionally, in the proposed methodologies, little attention is paid to the effect such discretization has on implementation and communication requirements. Such issues are typically abstracted or sidestepped.

In this paper, the results of a coordination scheme based on market-based control (MBC) and focusing on faster response

and lowered communication load are discussed. Results show the continuous system can trade off accuracy or optimality of the control action for the number of control messages exchanged with and the response speed of the PHEVs. The main contributions of this paper are:

- 1) Description of the typical problems involved in translating existing charging coordination algorithms to systems that work in the real-world. While optimality is important, not all algorithms or strategies are equally suitable for deployment.
- 2) Description of a market-based control scheme based on an existing algorithm, that is able to operate fully continuously, as opposed to a timeslot-based system.
- 3) Evaluation of this continuous MBC system with respect to communication load and accuracy or optimality of the performed coordination actions.

II. RELATED WORK

In recent years, many papers have appeared that focus on optimal strategies for controlling a fleet of charging EVs, sometimes situated in the more general domain of demand response (DR).

In [1], a system is proposed based on congestion pricing and Quality of Service (QoS) concepts in TCP/IP networks. A user utility function is employed to determine individual consumer demand in the next time-slot. The system is not fully distributed however, as a central entity is implied for the pricing signals sent to end-users. Effects of packet loss and delay of the control signal are considered by adding an unknown amount of noise to the price signals. However, this cannot be considered equivalent of message delay or loss. The system inherently has no feedback loop that allows continuous monitoring and intervention.

In [2], coordination in the form of a highly dimensional allocation problem is tackled via approximate dynamic programming. Vehicles send schedule updates (arrival, departure and SOC requirements) to a central aggregator which hosts the scheduling algorithm. Individual charging or discharging rates are then sent to each vehicle. All events are aligned on time-slot boundaries.

In [3], optimal demand-side management (DSM) is achieved using a model derived from game-theory. Each consumer's scheduler is required to broadcast its consumption schedule to all other participants in the DSM programme. Schedule granularity is chosen to be one hour.

In [5], vehicles send their requirements in the form of intentions to an aggregator. The aggregator has estimations on household consumption and production from residential

This research is partially supported by the IBBT-ICON projects SPARC and SmartE.

K. De Craemer and G. Deconinck are with the Department of Electrical Engineering ESAT-ELECTA, KU Leuven, Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium (e-mail: klaas.decraemer@esat.kuleuven.be and geert.deconinck@esat.kuleuven.be)

photovoltaic installations and coordinates charging as to flatten the local grid's power profile (known as *valley filling*).

In [7], an optimization model is defined and solved through linear programming, allowing consumers to adapt to electricity prices on an hourly basis. Future prices and their bounds are estimated via an ARIMA-based model. A two-way communicating smart meter or an Energy Management System (EMS) is assumed to be available in participating residences, but a case with estimated day-ahead instead of hourly pricing information is also considered.

In [6], a system rooted in online market mechanism design is presented and applied to the coordinated charging problem. Simulations use data from a real trial and results are compared to an offline benchmark that has prior knowledge on vehicle arrivals. Depending on the size of the time-steps and the vehicles' arrival interval, a realistic setting would be computationally demanding.

Summarized, while attaining good results regarding coordination, these approaches assume the existence of an ideal communication channel between the controlled vehicles and an aggregator in some form. Some require levels of local processing power or information from many other agents that cannot always be easily transferred over the available communication channels. Additionally, while the discretization in timeslots is convenient in market-linked systems, it is not representative of how real devices function and interact and limits the system's responsiveness.

III. BACKGROUND

As outlined in the previous sections, it is not a bad idea to take some practical issues into consideration when performing coordinated charging of EVs. Upon connection, a vehicle's charger or the station will identify cable rating and ensure some electrical safety before applying power. This is achieved through very basic signaling [8] (resistor coding or PWM) and at the moment there are no standards defining how and what additional information is exchanged (e.g. vehicle identification, electrical interlocking, charge profile ...). However, for coordinated charging to be feasible, more is needed;

- 1) A charger should at least be able to alter the charging rate of a vehicle (at continuous or discrete values). Some charging poles today can be remotely switched, but this feature is not intended for controlled charging.
- 2) An indication of the current state-of-charge (SOC) is required, together with some indication on the time the battery needs to be filled. This may or may not involve user interaction. Current charge poles are limited to metering of the vehicle's consumed power.
- 3) Most importantly, the chargers participating in coordinated charging must have means to communicate with each other or a central system, to send and receive information and coordination messages.

With charging poles "in the open", communication is typically achieved by cellular networks [9] [10]. This introduces complications regarding reliability, delay and data throughput. Furthermore, individual charging poles are resource constrained and therefore unsuitable for running computationally intensive software.

An additional problem for the coordination system is posed by the way energy is traded; on separate markets with different closure times. Good coordination should therefore not only focus on the optimality of the generated charging schedules, but also on minimizing data traffic while being fault tolerant in the case of communication failures.

IV. MARKET-BASED CONTROL

The agent-based approach used as basis for coordinated charging in this paper is based on the concept of market-based control (MBC) [11] [12] and Walrasian auctions. In MBC, supply and demand (bid) curves represent the willingness of consumers and producers of electricity to buy or produce at a certain price. In [11], devices are steered by device agents organized in a tree structure (see Figure 1). On the top of the tree is an auctioneer agent. Every quarter, agents at the bottom of the tree assemble bid functions representing their willingness to pay and consume, taking into account the specific constraints of the controlled device. Bids are sent upwards and the concentrators or auctioneer agent starts a matching process. Intermediate agents sum the bids of underlying agents and submit them as a single bid to a 'higher' agent. Practically, a device agent in this context could be seen as the software running on charge pole or inside a smart charger, while concentrator agents group multiple charge poles or parking areas. The number of intermediate levels is not fixed and can be varied as required by the application.

Eventually, the auctioneer agent determines the point at which consumers and (external) producers are in equilibrium. The mechanism is essentially the same as traditional demand and supply curves in economics. Afterwards, the equilibrium price is transferred downwards to all agents who will start consuming or producing according to their submitted bids. Note that the equilibrium price here is not necessarily represented in a currency. It can be a pure control signal and the term *priority* would be more suitable. The advantages of the MBC scheme are its Pareto-optimality and low (device) agent complexity.

Again, time is divided into timeslots during which bids apply (rounds). Decreasing the rounds interval would increase data traffic drastically, which is especially problematic for large clusters of agents. Also, updating the market equilibrium and sending it to all agents whenever a new agent arrives would trigger an avalanche of messages.

V. CONTINUOUS MARKET-BASED CONTROL

Our work focuses on improving the MBC scheme as implemented in [11] as to make it continuous and fault tolerant in the case of communication failures, while at the same time keeping communication and computation by the device nodes (the PHEV agents) down. The terms PHEV agents, intermediate agents and top agent are used to designate device agents, concentrators and the auctioneer agent respectively.

In the first step, event based equilibrium price calculation is introduced; agents will only pass each other new bids and equilibrium prices when they consider that their situation has changed enough to warrant such actions. We refer to this as the caching of bids and equilibriums. This caching is

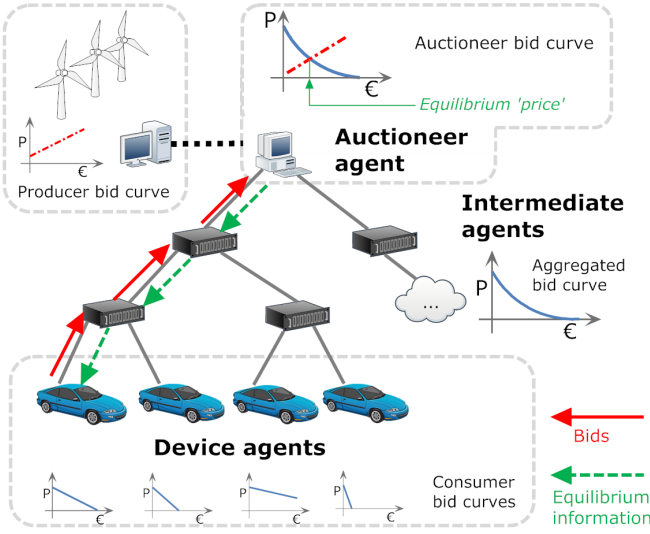


Fig. 1. MAS structure in the market-based control system. Bids from agents (consumers) are aggregated upwards and eventually the auctioneer agent matches these with a producer supply curve. The equilibrium price is then passed back downwards as a control signal.

achieved slightly differently at the hierarchical levels of the agent structure (being device agents, intermediate agents and the auctioneer agent).

When a device agent connects, the linked intermediate agent is contacted. In turn, the intermediate agent requests an initial bid from the device agent and stores this bid internally. Device agents are expected to only send new bid updates when the absolute distance to the previous bid has increased sufficiently (e.g. 150W). This threshold value is passed from intermediate agent to device agent upon first contact and can change per agent and during the charging process. As a safeguard against communication failures, intermediate agents will request new bids from the device agents at least once every hour. Intermediate agents employ a similar caching strategy; when their aggregated bid function has changed sufficiently, a new bid is sent to the auctioneer agent (or another layer of intermediate agents). The auctioneer agent is considered to have the largest resources (computation, memory, storage) available and capable of updating the equilibrium price almost as soon as it receives a new bid. For each intermediate agent connected to the top agent, the net effect of the change in equilibrium is determined. No reply is sent if the expected power change for each separate intermediate agent is too limited (depending on the application, e.g. $\pm 15\%$).

Moving back down the tree, intermediate agents receiving an equilibrium update will first determine total local impact and only update device agents that are most responsible for violating the local threshold (e.g. 1500W). Then, each agent is checked individually to avoid large equilibrium differences between agents. Additionally, when device agent power is very low (e.g. $\leq 200\text{W}$) no new updates are sent to that specific device agent, avoiding oversensitivity (which would translate into lots of messages for little absolute power changes).

To further reduce messages at the lowest level in the tree (typically the most “expensive” messages), a device agent

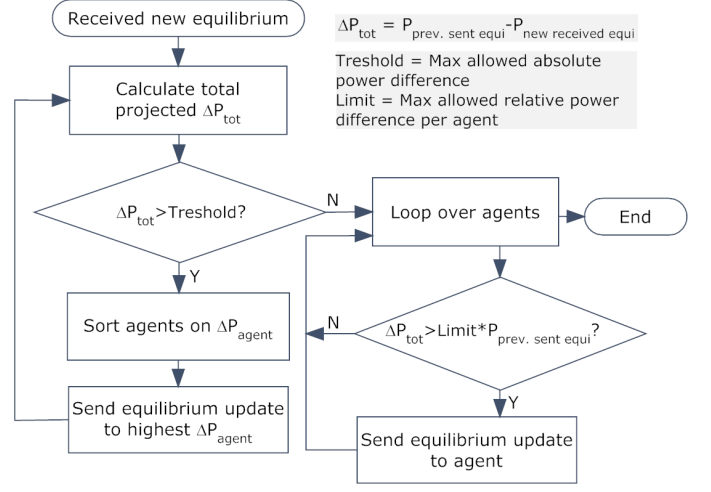


Fig. 2. Equilibrium caching decision flow chart for the intermediate agents

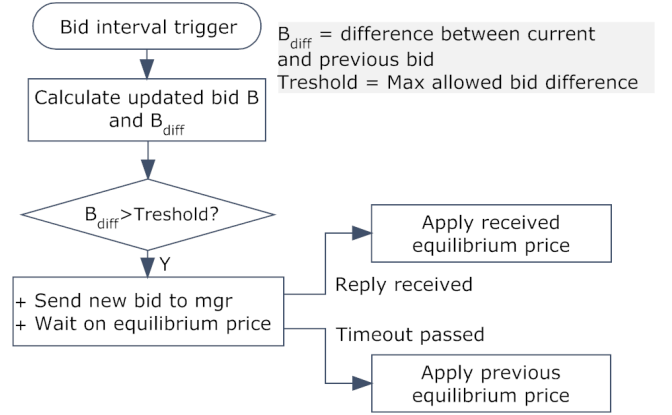


Fig. 3. Bid caching decision flow chart for the device agents

will only wait a limited amount of time for a reply from its governing intermediate agent before changing charge power. For example, after 120 seconds, the previously known equilibrium will be applied to the last submitted bid. Because this system is entirely event-driven, device agents are controlled asynchronously from each other. Furthermore, (drastic) changes in the equilibrium price propagate almost immediately (depending on the used intermediate agents parameters) to the device agents, allowing for a faster response to the top agents intervention.

The downside of this “caching” of bids and equilibriums is that, in contrast to the time-slot market-based control, agent(s) at the top of the tree have no exact information anymore on how much nodes at the lower levels are consuming or producing (in case of distributed generation), because their aggregated bids can be out of date or the equilibriums known to the device agents can differ slightly among them. This difference will have an influence on the equilibrium price and therefore on the optimality of the control action.

VI. AGENT BIDDING

One of the most crucial facets is the logic behind device agent bidding, as a bid represents the power flexibility or the

utility function of a PHEV. Different requirements determine the willingness-to-pay of a PHEV and thus its bid, including:

- 1) Maximum charge power P_{max} , limited by the vehicle's power electronics or the grid connection.
- 2) Time till departure Δt_{dep} .
- 3) Required energy ΔE needed by time t_{dep} .

A simple way of calculating such a bid is by using a corner price p_r to create a sloped curve, as depicted in the left of Figure 4. A higher p_r value will give rise to a flatter bid curve and expresses a vehicle agent's willingness to consume at higher prices. For convenience, p_r is kept within the range of [0,1].

A. Asymptotical bid function

By combining the stated requirements directly, the corner price p_r can be stated as shown in equation 1 (and as used in [13]).

$$p_r = \frac{\Delta E}{\Delta t_{dep} \cdot P_{max}} \quad (1)$$

After receiving the equilibrium price, the corresponding power point is set as the charging power. Vehicles having a p_r greater than the equilibrium will not charge, others will charge at a rate proportional to their individual requirements. Thus less flexible cars (with a less steep bid curve) will be charged first. However, due to the ever decreasing t_{dep} , p_r tends to increase asymptotically. This in turn leads to bids that keep edging towards higher charge power as departure time gets closer, as shown on the right of Figure 4.

B. Linear bid function

To avoid this effect, a linear bid function is also considered and shown in equation 2. The disadvantage of it is the need for a predetermined "offset" to scale the parameters. In this case, we used 12 hours as the maximum t_{dep} to consider, and the maximum allowed battery charge E_{max} to normalize the required energy ΔE .

$$p_r = 1 - \left(\frac{1}{2} \Delta t'_{dep} + \frac{1}{2} \Delta E' \right) \quad (2)$$

with

$$\Delta t'_{dep} = \begin{cases} \frac{\Delta t_{dep}}{12} & \text{if } \Delta t_{dep} < 12 \\ 1 & \text{if } \Delta t_{dep} \geq 12 \end{cases}$$

$$\Delta E' = \frac{\Delta E}{E_{max}}$$

VII. RESULTS

To assess the effect on coordinated charging of timeslot versus continuous MBC as explained in section V, a few benchmark scenarios have been constructed.




Fig. 4. Left: Basic PHEV device agent bid function as used in the current results. Right: Asymptotical bid function evolution as time to departure shrinks, giving rise to increasing charge power. By using a linear bid function this effect can be avoided.




Fig. 5. External producer bid function. Willingness to produce increases linearly with the offered priority or price.

A. Small-scale scenarios, comparing optimality

A first setting consists of a fleet of 96 PHEVs simulated during 3 days in a distribution grid. The task of the auctioneer agent comprises balancing a linear static producer bid function with the demands of the charging agents. This function is shown in Figure 5. While not entirely realistic, it makes it easier to analyze the different scenarios. A more realistic case could consist of a day-ahead profile that the auctioneer agent has to follow within certain boundaries. As a reference case, a timeslot MBC system with 15 minute intervals is used. In all cases, the hierarchical agent tree structure (also shown in Figure 1) consists of 1 top agent x 6 x 4 intermediate agents (two levels) x 4 PHEV agents. The maximum allowed power per vehicle is set at 3 kW (13 Ampere, corresponding to charging from regular outlets) and the driving behavior of the PHEVs is based on a statistical availability model [14]. Vehicles can be at home (plugged and available for charging), driving (discharging) or at work (not plugged).

The results of the simulated scenarios are evaluated by looking at two key values:

- 1) The number of messages exchanged (sent and received) by device agents. As explained in sections I and II, these messages are likely to be the most expensive in terms of bandwidth and monetary cost. Alternatively, it could also make sense to focus on the number of messages exchanged by intermediate agents.
- 2) RMS value of the power difference with the timeslot MBC (RMS deviation). This gives a first indication of how well or how bad a certain continuous MBC scenario performs. However, using the timeslot system as benchmark or reference case is not entirely justified because the use of timeslots introduces delay in the actions taken by device agents. Thus, another comparison

TABLE I
CONTINUOUS MBC SYSTEM SCENARIO PARAMETERS OVERVIEW.

Simulation #	Node agents			Intermediate agents				Top agent			
	Bid interval	Bid timeout	Bid max diff	Bid interval	Bid timeout	Bid max diff	Total diff	Node diff	Low Threshold	Update interval	Total diff
Scenario 1	450s	120s	200W	45s	120s	200W	1000W	15%	200W	10s	15%
Scenario 2	240s	120s	10W	45s	120s	10W	50W	7.5%	50W	10s	5%
Scenario 3	900s	120s	400W	45s	120s	200W	1500W	22%	250W	10s	20%
Scenario 4	30s	15s	2W	30s	15s	2W	10W	2%	5W	5s	1%

TABLE II
SUMMARIZED SCENARIO RESULTS FOR A SETTING WITH 96 PHEVs,
ASYMPTOTICAL BID FUNCTION

Simulation	Node msgs	NRMSD Tslot	NRMSD Scen4
Timeslot	30000	0%	3.87%
Uncontrolled	0	33.88%	32.86%
Scenario 1	16825	6.01%	3.38%
Scenario 2	60415	4.27%	1.55%
Scenario 3	12182	13.17%	9.85%
Scenario 4	265595	3.68%	0%

TABLE III
SUMMARIZED SCENARIO RESULTS FOR A SETTING WITH 96 PHEVs,
LINEAR BID FUNCTION

Simulation	Node messages	NRMSD Tslot	NRMSD Scen4
Timeslot	26901	0%	12.55%
Uncontrolled	0	23.70%	22.30%
Scenario 1	4778	11.41%	5.02%
Scenario 2	15204	10.31%	3.78%
Scenario 3	3892	13.86%	7.69%
Scenario 4	52774	10.83%	0%

scenario is added, scenario 4, and parameterized to be as uncompromising as possible; agents bid and receive equilibrium price updates almost instantaneously. Such a scenario should therefore be very close to a theoretical optimum behavior of the original MBC concept. The RMS values are normalized to aid comparison.

Table I summarizes the different simulation parameters, while table II and III show the numerical results in the case of asymptotical and linear bid functions respectively. In the “uncontrolled” case, vehicles just start charging at full power as soon as they connect. Simulations 1 to 4 are continuous MBC variants with different levels of caching, with the last one, 4, serving as the uncompromising scenario.

More specifically, scenario 3 uses an aggressive level of caching (meaning that higher bid and equilibrium differences are allowed); scenario 2 aims for the opposite and is configured for accuracy, while scenario 1 fits in between at intermediate levels of caching. In the latter case, device agent message count was reduced by almost 70% when compared to the timeslot simulation.

1) *Asymptotical bid function*: Figure 6 shows the power consumption of the PHEVs during 3 days of simulation, for the asymptotic bidfunction. Time starts at 24 hours to cut out transient effects at the start of the simulation. The upper plot represents the power consumed by all connected PHEV device agents. In the uncontrolled case, power peaks during the night, when almost all of the vehicles are connected to a charger. In the controlled scenario, consumption is centered

around departure. Looking at the scenario-graphs, it can be seen that the more a scenario employs aggressive caching of bids and equilibriums, the more its power consumption profile will deviate towards the uncontrolled scenario. But the positive effect of the caching is the reduced need for communication between device agents, which can be seen in Table II. In the lower part of Figure 6, results are related to the equilibrium price as known to the auctioneer agent. Due to the caching of bids, the auctioneer agent underestimates the requested amount of power and equilibrium prices decrease. This is especially clear for scenario 3, having the most aggressive caching. In case of the timeslot MBC system, device agents have less opportunity to react to the equilibrium price, an effect which pushes the equilibrium price slightly up as device agents delay their charging to the point where even more of them need to be ready for departure. In the continuous system, this is avoided by having device agents send updated bids immediately whenever projected impact would be high enough.

2) *Linear bid function*: Figure 7 shows the same scenarios, but using the linear bid function from equation 2. Charging is now better spread out between arrival and departure times, as opposed to the asymptotic bid function where charging is postponed until close to departure. This has the effect of reducing total message count for the device agents as well, because charged vehicle agents “log out”. Again, as the caching becomes more aggressive, power profiles tend to shift towards the uncontrolled scenario and equilibrium prices become less pronounced. Because of the improved spread, power consumption and equilibrium price does not peak as much and the variation between the different scenarios is smaller. This bid function therefore seems better suited for application in our MBC system and illustrates the importance of consumer valuation functions and their impact on coordination effectiveness.

It must be noted however, that the use of the linear bid function reveals a consistent discrepancy between the timeslot system and even the ‘best’ continuous scenario (which is scen. 4). More specifically, the power profile and equilibrium values are slightly higher during periods with many departing vehicles, and slightly lower anywhere else. Because of this, NRMSD with the timeslot version is distorted upwards. The discrepancy is also visible on the simulations with the asymptotic bid functions, but more subdued. This seems to be an inherent effect of using a time-slot system.

Additionally, as briefly mentioned in section IV, the way device agents build their bids may exaggerate the underestimation of the consumed power by the auctioneer agent in

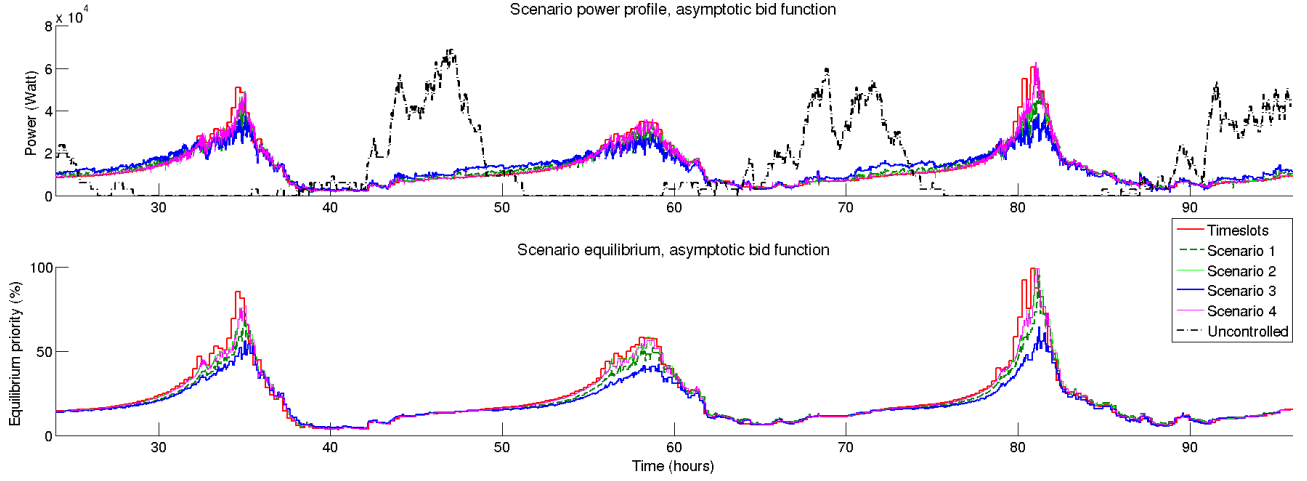


Fig. 6. Power consumption and equilibrium price during simulation of different scenarios with 96 PHEVs, 3 simulated days, using the asymptotic bid function.

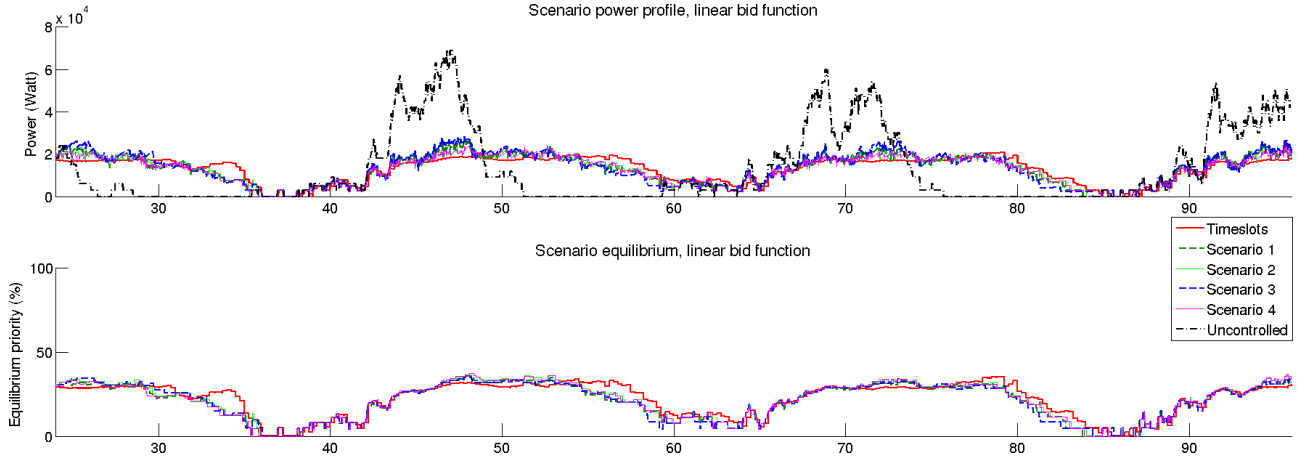


Fig. 7. Power consumption and equilibrium price during simulation of different scenarios with 96 PHEVs, 3 simulated days, using the linear bid function.

case of the continuous MBC. Device agents tend to increase charge power because of the decreasing slope of the bids but this increase is effectively cached as long as no thresholds are tripped in the intermediate agents. The results show that it is possible to balance coordination optimality (represented by the RMS difference with the timeslot system and scenario 4 of the continuous MBC) with communication load (the number of device agent messages) and responsiveness (determined in the parameter sets from Table 1). But scenario 2's profile is almost identical to that of scenario 4, showing that not much additional gain is achieved by further decreasing the caching thresholds beyond some point.

B. Randomized scenarios, trade-off behaviour

To visualize this trade-off behavior between coordination optimality and agent communication, sets of randomly generated scenarios were simulated and both evaluation parameters plotted against each other. Figure 8 shows the scaling behaviour when using the linear bid function. The interrupted line shows a hard “front”; no set exists beyond this border. Note that this is not the same as a full Pareto-front; examination reveals that sets close to the “corner” of the front have a

low device agent message count and good optimality, but do so by effectively sidestepping the intermediate agent level due to the use of very aggressive caching parameters there.

C. Scalability of the continuous MBC solution

To verify scalability of the continuous MBC scheme, the number of PHEVs is increased while maintaining the same hierarchy ratio for the agent network. The static producer bid was also scaled accordingly. Figure 9 shows the results for the different scenarios, using the asymptotical bid function. It can be seen that scaling performance is similar to a timeslot implementation which equals slightly better than linear. Using the linear bid function returned comparable numbers.

VIII. CONCLUSIONS AND FUTURE WORK

First of all, the results demonstrate the discrepancies that a time-slot based system implies in the case of coordinated charging, which are hard to avoid without switching to very small time steps. This is where a continuous or event-based system has an advantage; it has virtually infinitely small time steps. Unfortunately, such a system has to make a compromise

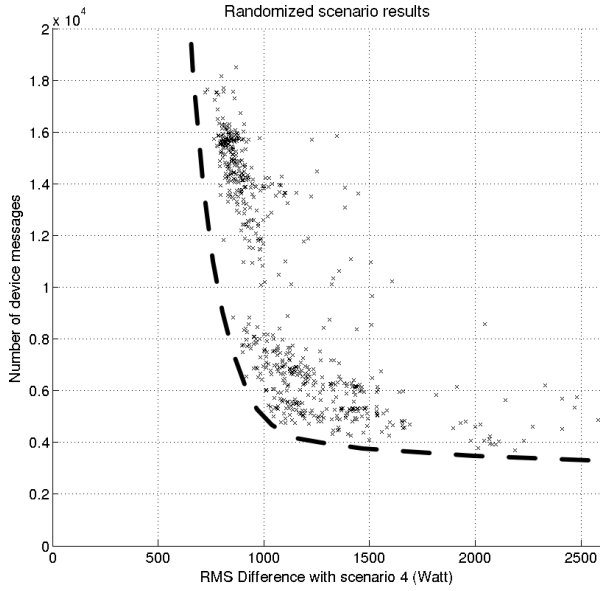


Fig. 8. Results from random sets of parameters for the continuous MBC system. The trade-off between the two key result parameters is demonstrated.

regarding optimality of the control actions when compared to its time-slot based counterpart.

Secondly, examining the proposed continuous system's results, one can see that by tuning the caching parameters a balance can be found between device agent messaging and optimality of the control action with the added advantage of fast response to device or auctioneer agent requests. In one scenario, device agent messages were reduced by more than 80% while the normalized RMS deviation of the solution was only 5.02% worse compared to an uncompromised continuous scenario. The other way round, another scenario used a comparable number of device agent messages but with an RMS deviation (again compared to the uncompromised continuous scenario) that was 70% smaller than the timeslot-based counterpart.

The presented approach to the problem of coordinated PHEV charging is a step towards a more robust, fault tolerant charging coordination system as an asynchronously moving scheme gives participating device and manager agents more room for recovery procedures. In future work, we will try to assess the performance of the presented scheme in a more market-oriented setting; meaning that we will replace the static producer bid function with a target such as balancing an energy retailer's portfolio. Preliminary results with such scenarios have shown one shortcoming of the MBC system lies in the instantaneous matching of demand and supply. It would be more efficient to incorporate information on future events (e.g. wind prediction or estimated vehicles arrivals) as this would remedy situations in which all the flexibility of the device agents is depleted and demand effectively becomes uncontrollable.

REFERENCES

- [1] Z. Fan, "A Distributed Demand Response Algorithm and Its Application to PHEV Charging in Smart Grids," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–11, 2012. [Online].




Fig. 9. Number of device agent messages and RMS Deviation of the scenarios with different number of participating PHEVs.

Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6168819&isnumber=5446437>

- [2] V. W. Wong, "An approximate dynamic programming approach for coordinated charging control at vehicle-to-grid aggregator," in *2011 IEEE International Conference on Smart Grid Communications (Smart-GridComm)*. IEEE, Oct. 2011, pp. 279–284. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6102333>
- [3] A.-H. Mohsenian-Rad, V. W. S. 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*, vol. 1, no. 3, pp. 320–331, Dec. 2010. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5628271>
- [4] K. Clement, E. Haesen, and J. Driesen, "Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids," in *2009 IEEE/PES Power Systems Conference and Exposition*. IEEE, Mar. 2009, pp. 1–7. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4839973>
- [5] S. Vandael, K. De Craemer, T. Holvoet, N. Boucke, and G. Deconinck, "Decentralized coordination of plug-in hybrid vehicles for imbalance reduction in a Smart Grid," in *Power Engineering*, no. section 3, 2011, pp. 803–810. [Online]. Available: http://www.ifaamas.org/Proceedings/aamas2011/papers/B6_G65.pdf
- [6] V. Robu, S. Stein, E. H. Gerding, D. C. Parkes, A. Rogers, and N. R. Jennings, "An Online Mechanism for Multi-Speed Electric Vehicle Charging," in *Second International Conference on Auctions, Market Mechanisms and Their Applications (AMMA'11)*, New York, 2011. [Online]. Available: <http://eprints.soton.ac.uk/272424/>
- [7] A. J. Conejo, J. M. Morales, and L. Baringo, "Real-Time Demand Response Model," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 236–242, Dec. 2010. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5607339>
- [8] IEC, "IEC 62196-1 ed2.0: Plugs, socket-outlets, vehicle connectors and vehicle inlets - Conductive charging of electric vehicles - Part 1: General requirements," IEC, Tech. Rep., 2011.
- [9] Siemens, "Sitrafic Epos-I Charging point," 2011. [Online]. Available: <http://www.energy.siemens.com/fi/en/energy-topics/electromobility/portfolio/charging-pole/sitrafic-epos.htm>
- [10] Coulomb Technologies, "Coulomb Technologies Charging Stations," 2012. [Online]. Available: <http://www.coulombtech.com/products-chargepoint-stations.php>
- [11] K. Kok, K. Warner, and R. Kamphuis, "PowerMatcher: multiagent control in the electricity infrastructure," in *AAMAS 05: Proceedings of the 4th int. joint conf. on Autonomous Agents and Multiagent Systems, volume industry track*, New York, 2005, pp. 75–82. [Online]. Available: <http://www.powermatcher.net/fileadmin/>

powermatcher/user/documents/Articles/OverviewArticles/AAMAS_Article_PowerMatcher_DistributionVersion.pdf

- [12] K. Kok, Z. Derzsi, J. Gordijn, M. Hommelberg, C. Warmer, R. Kamphuis, and H. Akkermans, "Agent-Based Electricity Balancing with Distributed Energy Resources, A Multiperspective Case Study," in *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)*. IEEE, Jan. 2008, pp. 173–173. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4438877>
- [13] M. D. Galus, R. La Fauci, and G. Andersson, "Investigating PHEV wind balancing capabilities using heuristics and model predictive control," in *IEEE PES General Meeting*. IEEE, Jul. 2010, pp. 1–8. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5589267>
- [14] J. Van Roy, N. Leemput, S. De Breucker, F. Geth, T. Peter, and J. Driesen, "An Availability Analysis and Energy Consumption Model for a Flemish Fleet of Electric Vehicles," in *European Electric Vehicle Congress (EEVC)*, Brussels, 2011, p. 12.

Klaas De Craemer is a PhD Researcher at the department of Electrical Engineering (ESAT) in the division ELECTA. He received the M.Sc. degree in Electrical Engineering from the KU Leuven (Belgium) in 2009. His research interests include architectures and control systems for demand participation of smart homes and electrical vehicles, including their impact on grid stability.

Geert Deconinck is full professor at the Department of Electrical Engineering (ESAT) of the KU Leuven (Belgium). He performs research on the design of dependable system architectures for industrial control and smart grids and on lighting technology. He received his M.Sc. in Electrical Engineering and his Ph.D. in Engineering from the K.U.Leuven in 1991 and 1996, respectively. He is member of the IEEE SMC TC on Infrastructure Systems and Services and chairman of the TI society BIRA on industrial automation. He is a member of the Belgian Institute for Lighting, the Royal Flemish Engineering Society, the Institute of Engineering and Technology (IET), a senior member of IEEE.