Market-based Control in Decentralized Electrical Power Systems —Extended Version—

Koen Kok, Gerben Venekamp, Pamala Macdougall
Energy research Centre of the Netherlands (ECN)
Power Systems and Information Technology
Petten, The netherlands
{j.kok g.venekamp macdougall}@ecn.nl

ABSTRACT

Over the course of the 20th century, the electrical power systems of industrialized economies have become one of the most complex systems created by mankind. A number of ongoing trends will drastically change the way this critical infrastructure is operated. Demand for electricity keeps growing while the controllability of generation capacity is decreasing due to introduction of renewable energy sources. Further, there is an increase of distributed generators (DG), i.e. the generation capacity embedded in the (medium and low voltage) distribution networks. Intelligent distributed coordination will be essential to ensure the electricity infrastructure runs efficiently in the future. The PowerMatcher technology, a multi-agent coordination system, has been developed to provide this kind of coordination. The heart of the system is an electronic market on which local control agents negotiate using strategies based on short-term microeconomics. A proof-of-principle simulation study involving renewable power generation, demand response and distributed generation indicates the impact of the multi-agent coordination. The study focusses on a micro-grid setting where balancing is done by a diesel generator. Application of the PowerMatcher is shown to reduce the peak power delivered by the diesel generator by approx. 45% while the total diesel generated power decreased by approx 40%.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligencemultiagent systems, coherence and coordination

General Terms

Algorithms, Performance

Keywords

Multi-agent systems, market-based control, electronic markets, intelligent electricity infrastructures.

1. INTRODUCTION

The sustainable power systems of the future are characterized by a high penetration of both distributed generation (DG) and intermittent generation. DG are generation units embedded in the (medium and low voltage) distribution networks. The availability of intermittent generation units is dependent on the availability of a primary energy source, such as wind or solar radiation.

In the status quo, the balance between demand and supply is maintained by a relatively small number of big central power plants following load patterns that are, to a great extent, uncontrollable and partially unpredictable. As the supply side becomes more inflexible, a need emerges to utilize the flexibility potential of the demand side. With that, the nature of coordination within the electricity system is changing from a few centrally controlled power plants into coordination among a large number of generators and responsive loads. These generators and loads show time-varying levels of flexibility and a great variety in (production and consumption) capacity. Therefore, the standard paradigm of centralized control, which is used in the current electricity infrastructure, will no longer be sufficient. The number of system components actively involved in the coordination task will be huge. Centralized control of such a complex system will rapidly reach the limits of scalability, computational complexity, and communication overhead. An excellent domain for multi-agent systems coordination.

The PowerMatcher is a multi-agents systems based coordination mechanism designed for future sustainable power systems. An overview of the PowerMatcher, including a number of proof-of-concept field experiments, can be found in [5] and [3]. Since its incarnation in 2004, the PowerMatcher has been implemented in three major software versions. In a spiral approach, each software version was implemented from scratch with the first two versions being tested in simulations and field experiments. The third version is currently under development and is planned to be deployed in a number of field experiments and real-life demonstrations with a positive business case.

In this paper we describe the MAS principles behind the PowerMatcher, give an overview of the agent structure used and present the early results of a proof-of-principle simulation study conducted with the third version.

2. MARKET-BASED CONTROL

2.1 Multi-agent Systems

A multi-agent system (MAS) is a software system implemented as a collection of interacting autonomous agents [8]. A software agent is a self-contained software program that acts as a representative of something or someone (e.g., a device or a user). A software agent is goal-oriented: it carries out a task, and embodies knowledge for this purpose. For this task, it uses information from and performs actions in its local environment or context. Further, it is able to communicate with other entities (agents, systems, humans) for its tasks.

In multi-agent systems (MAS), a large number of such agents are able to interact. Local agents focus on the interests of local sub-systems and influence the whole system via negotiations with other software agents. MAS theory provides a paradigm for designing open, flexible, scalable, and extensible ICT systems aimed to operate in highly-complex environments [2]. While the complexity of individual agents remains low, the intelligence level of a well-designed MAS the global system is high. This phenomena is referred to as *emergence*. A MAS can be designed to exhibit specific emergent behavior. Hence, the design goal is to reach system-level intelligence through the interactions of high numbers fairly simple software agents.

2.2 Electronic Markets

Emergence of system-level intelligence can be achieved using *electronic markets*, which provide a framework for distributed decision making based on microeconomics. Using electronic markets, the interactions of individual agents can be made highly efficient. Microeconomics is a branch of economics that studies how economic agents (i.e., individuals, households, and firms) make decisions to allocate limited resources, typically in markets where goods or services are being bought and sold. One of the goals of microeconomics is to analyze market mechanisms that establish relative prices amongst goods and services and allocation of limited resources amongst many alternative uses. A distinctive feature of microeconomics is that it aims to model economic activity as an interaction of individual economic agents pursuing their private interests [7], [11].

Whereas, economists use microeconomic theory to model phenomena observed in the real world, computer scientists use the same theory to let distributed software systems behave in a desired way. Market-based computing is becoming a central paradigm in the design of distributed systems that need to act in complex environments. Market mechanisms provide a way to incentivize parties (in this case software agents), that are outside the sphere of direct control, to behave in a certain way [1, 10]. A microeconomic theory commonly used in MAS is that of general equilibrium. In general equilibrium markets, or exchange markets, all agents respond to the same price, that is determined by searching for the price that balances all demand and supply in the system. From a computational point of view, electronic equilibrium markets are distributed search algorithms aimed at finding the best trade-offs in a multidimensional search space defined by the preferences of all agents participating in the market [13, 15]. The market outcome is *Pareto* optimal, a social optimal outcome for which no other outcome exists that makes one agent better-off without making other agents worse-off.

In *Market-based Control*, agents in a MAS are competing for (one or more) resources on an equilibrium market whilst performing a local control task (e.g., classical feedback control of a physical process) that needs those resources as an input.

2.3 Market-based Control Example

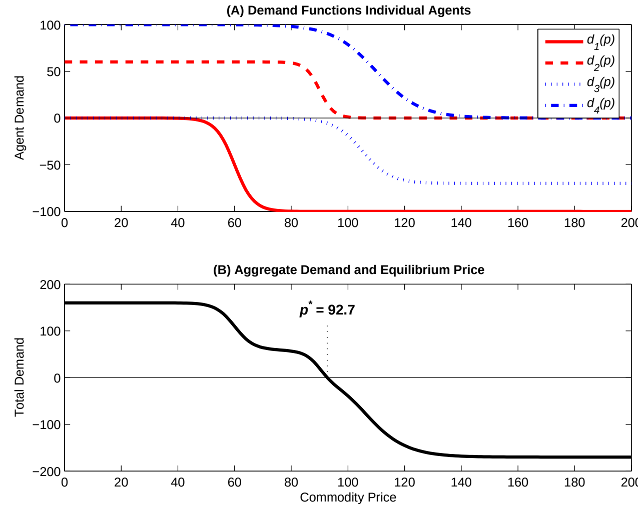

Figure 1: Example general equilibrium market outcome. (A) Demand functions of the four agents participating in the market. (B) The aggregate demand function. At price p^*, the market is in equilibrium: the sum of all supply and demand equals to zero.

Table 1: Agent demand levels for the two situations described in the text. Situation 1 corresponds to Figure 1, situation 2 to Figure 2.

	$d_1(p^*)$	$d_2(p^*)$	$d_3(p^*)$	$d_4(p^*)$	$\sum d_\alpha(p^*)$
S1	-99.9	15.1	-5.6	90.4	0.0
S2	0.0	0.0	-50.6	50.6	0.0

In a typical price-based market-based control problem, there are several producing and/or consuming agents and an auctioneer agent. Each market round the producers and consumers create their market bids and send these to the market agent. These bids are ordinary, or *Walrasian*, demand functions $d(p)$, stating the agent's demand d at a price of p. The demand function is negative in the case of production. After collecting all bids, the market agent searches for the equilibrium price, i.e. the price at which the market clears. This price is broadcast to all agents, who can determine their allocated production or consumption from this price and their own bid. Finally, all producing agents feed their allocated production into the flow network while all consuming agents extract their consumption from it.

Figure 1 shows an example of price forming in a (single-commodity) general equilibrium market with four agents. The demand functions of the individual agents are depicted in graph (A). There are two consuming agents, whose demand decreases gradually to zero above a certain market price. Further, there are two producers whose supply, above a certain price, increases gradually to an individual maximum. Note that supply is treated as negative demand. In a control setting, the position of the inflexion point is typically determined by the current process state. The solid line in (B) shows the aggregate demand function. The equilibrium price p^* is determined by searching for the root of this function, i.e. the point where total demand equals total supply. The value of each agent's demand function at this prices is given in Table 1, Situation 1.

Suppose the commodity traded in this example is electrical power. Suppose further, the first agent is associated with a unit for combined heat and power generation (CHP), e.g. used to heat a swimming pool. While serving the local heat demand, the unit produces electricity at the same time. Its local control goal is to keep a large water-filled heat buffer between two temperature limits. This buffer serves heat demand coming from subsystems such as space heating and heating of pool water. In the situation depicted by Figure 1, the CHP unit runs at full capacity. Its produced electricity is consumed by the two consuming agents and its produced heat is heating up the buffer.

Suppose that some time later, the heat buffer temperature is approaching the upper temperature limit. Then, the agent's need to produce heat — and, thus, its willingness to deliver electricity to the other agents — will be much lower. Now, the agent wants to produce electricity only if it gets a really good price for it and updates its bid accordingly. Figure 2 and Table 1, Situation 2, show the new situation. Due to the change in demand function of the first agent, the equilibrium price rises to 109.8. This causes the consuming agents to lower their intake, for agent 2 virtually to zero. The resulting demand is met entirely by the production of agent 4.

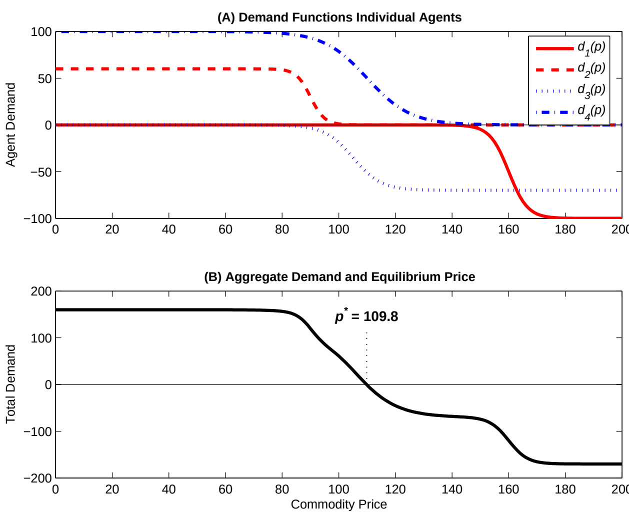

Figure 2: New market equilibrium after a change in the demand function of agent 1.

3. THE POWERMATCHER

This section describes *The PowerMatcher*, a general purpose coordination mechanism for near-real-time balancing of demand and supply in large clusters of *Distributed Energy Resources* (DER, distributed generation, demand response, and electricity storage connected to the distribution grid). These 'clusters' can be for example electricity networks with a high share of distributed generation or commercial trading portfolios with high levels of renewable electricity sources.

Since its incarnation in 2004, the PowerMatcher has been implemented in three major software versions. In a spiral approach, each software version was implemented from scratch

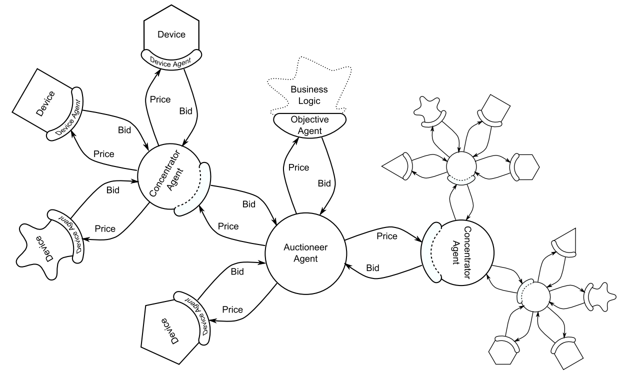

Figure 3: Example PowerMatcher agent cluster. See the text for a detailed description.

with the first two versions being tested in simulations and field experiments [6, 4, 12]. The third version is planned to be deployed in a number of field experiments [9] and real-life demonstrations with a positive business case.

3.1 Logical Structure and Agent Roles

Within a PowerMatcher cluster, the agents are organized into a logical tree. The leaves of this tree are a number of *local device agents* and, optionally, a unique *objective agent*. The root of the tree is formed by the *auctioneer agent*, a unique agent that handles the price forming by searching for the equilibrium price. In order to obtain scalability, *concentrator agents* can be added to the structure as tree nodes. More detailed descriptions of the agent roles are as follows:

- **Local device agent:** Representative of a DER device. A control agent which tries to operate the process associated with the device in an economical optimal way. This agent coordinates its actions with all other agents in the cluster by buying or selling the electricity consumed or produced by the device on an electronic market. In order to do so, the agent communicates its latest bid (i.e., a demand function, see below) to the auctioneer and receives price updates from the auctioneer. It uses this received price, together with its latest bid, to determine the amount of power the agent is obliged to produce or consume.

- **Auctioneer agent:** Agent that performs the price-forming process. The auctioneer concentrates the bids of all agents directly connected to it into one single bid, searches for the equilibrium price and communicates a price update back whenever there is a significant price change.

- **Concentrator agent:** Representative of a sub-cluster of local device agents. It concentrates the market bids of the agents it represents into one bid and communicates this to the auctioneer. In the opposite direction, it passes price updates to the agents in its sub-cluster. This agent uses 'role playing'. On the auctioneer's side it mimics a device agent: sending bid updates to the auctioneer whenever necessary and receiving price updates from the auctioneer. Towards the sub-cluster agents directly connected to it, it mimics the auctioneer: receiving bid updates and providing price updates.

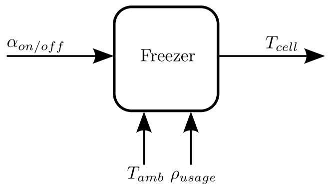

Figure 4: Freezer block model

- **Objective agent:** The objective agent gives a cluster its purpose. In absence of an objective agent, the goal of the cluster is to balance itself, i.e., it strives for an equal supply and demand within the cluster itself. Depending on the specific application, the goal of the cluster may be different. If the cluster has to operate as a *virtual power plant*, for example, it needs to follow a certain externally provided setpoint schedule. Such an externally imposed objective can be realized by implementing an objective agent. The objective agent interfaces the agent cluster to the *business logic* behind the specific application.

The logical agent structure follows the CoTree algorithm [14]. By aggregating the demand functions of the individual agents in a binary tree, the computational complexity of the market algorithm becomes $O(\lg a)$, where a is the number of device agents. In other words, when the number of device agents doubles it takes only one extra concentrator processing step to find the equilibrium price. Furthermore, this structure opens the possibility for running the optimization algorithm distributed over a series of computers in a network in a complimentary fashion to power systems architectures. We discuss the issue of scalability further in section .

3.2 Basic Device Agent Functionality

For a DER unit to be able to participate in a PowerMatcher cluster, its associated agent must communicate its momentary *bid curve* or *demand function* to the Auctioneer. As described before, this function defines the DER's electricity demand $d(p)$ for a given price p. An offer to produce a certain amount of electricity against a certain price is expressed by negative $d(p)$ values. As a convention, throughout this text we refer to these functions as a bid, even when (part of) the function expresses a production offer.

Lets's focus on an agent for an electricity-consuming device, say a freezer. A simple block model of the thermal process of a freezer cell and it's external influences is depicted in Figure 4. Input to the process model is the boolean control variable $\alpha_{on/off}$, switching the freezing element on or off. Further, the temperature in the freezing cell is influenced by two environment variables: the ambient temperature (T_{amb}) and a usage pattern (ρ_{usage}). The latter represents usage events like door opening & closing and goods being placed in or removed from the cell.

The control goal is to keep the inner cell temperature within the temperature band given by: T_{max} and T_{min}, the maximum inner cell temperature and the minimum inner cell temperature, respectively. In a conventional freezer, this is achieved by a standard on/off-controller with hysteresis. When participating in a PowerMatcher cluster, this conventional controller is replaced by a device agent. The goal of the agent is, again, to keep the cell temperature between the given limits, with an additional goal to consume in low-priced periods as much as possible.

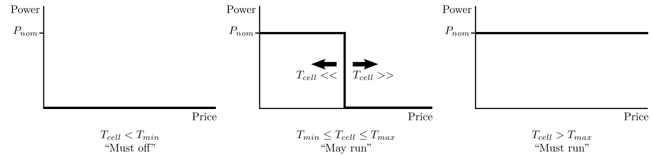

Figure 5: Three basic demand functions of a freezer.

Figure 5 gives the three basic bid shapes for the freezer. When the cell temperature is below its minimum (*left*), the freezing element must be switched off. Accordingly, the device agent sends a *Must Off* bid. Similarly, when the cell temperature is above its maximum (*right*), the agent sends a *Must On* bid. The agent is forced to accept any price in order to get the cell temperature back within its limits. When the cell temperature is within limits (*middle*), the agent has the flexibility to switch on or off the element dependent on the electronic market price. Since the freezer element can either be switched on or off the agent's bid is a step function: bidding either for the freezer's nominal power or for a power of zero. The position of the step flank reflects the agent's willingness to pay. When the cell temperature is still in the lower part of the temperature band, the agent is only willing to consume when the price is really low. However, when the temperature rises, the agent's willingness to pay increases with it. So, available flexibility is directly dependent on the device state (here the cell temperature), and the position of the step flank in the agent's bid directly reflects that. In order to optimize its strategy, the agent needs to have market-knowledge, as the notion of what defines a "high price" or a "low price" is crucial in the agent's bidding strategy. For an analysis of this aspect, we refer to [3].

3.3 Auctioneer and Concentrator Functionality

The core functionality of the auctioneer and the concentrators is to run the electronic market allocating the electrical power resource to the local device agents. The electronic market solves this allocation problem by finding the general equilibrium price p^* such that:

$$\sum_{a=1}^{N_a} d_a(p^*) = 0 \tag{1}$$

where N_a is the number of local device agents and $d_a(p)$ the demand function of agent a, stating the agent's demand or supply at a given price p.

The task of summoning all device agent's demand functions is divided over all concentrator agents and the auctioneer agent, here jointly referred to as *market agents*. Each market agent k summons the demand functions received from their attached agents. These functions originate from two different sources: (1) the device agents directly attached to k, and (2) the concentrator agents directly attached to k. The concentrated bid of k is calculated as:

$$a_k(p) = \sum_{j:x_j \in X_k} d_j(p) + \sum_{i:i \in Y_k} a_i(p) \tag{2}$$

where X_k is the set of local device agents directly connected to k and Y_k is the set of concentrator agents directly connected to k.

If k is a concentrator agent, it passes $a_k(p)$ on to the higher-level market agent it is attached to. If k is the auctioneer, it uses $a_k(p)$ to find the equilibrium price p^* such that the market is in equilibrium:

$$a_k(p^*) = 0 \tag{3}$$

Note that, in the latter case, a_k is the concentrated demand functions over all device agents:

$$a_k(p) = \sum_{a=1}^{N_a} d_a(p) \tag{4}$$

and that substitution of (4) in (3) yields the general market equation (1).

3.4 Communication Timing

The agents communicate in an event-based manner. Device agents update their bids whenever there is a change in the system state significant enough to justify a bid update. Typically, device agents update their bid once every few minutes or longer. Concentrators, in turn will not update their bid unless subsequent updated bids from lower agents result in a significant change in their concentrated bid. Likewise, the auctioneer will only communicate a new price after a considerable price change. In this way, coordination on a timescale of minutes is realized with low volumes of communicated data. For the two main application cases of the PowerMatcher, commercial portfolio balancing and congestion management , this type of *near real-time coordination* suffices, as these processes take place on a similar timescale. /todoPM chap: Add refs to applications.

4. CLUSTER-LEVEL BEHAVIOR

The self-interested behavior of local agents causes electricity consumption to shift towards moments of low electricity prices and production towards moments of high prices. As a result of this, the emergence of supply and demand matching can be seen on the global system level.

The aggregated, or concentrated, bid of all local control agents in the cluster — as is held by the auctioneer agent — can be regarded as a dynamic merit-order list of all DER participating in the cluster. On the basis of this list, the cluster as a whole is able to operate the (near-)real-time coordination activity optimally. In this section, we examine the behavior of a particular DER cluster in a number of simulated circumstances.

4.1 Micro-grid Operation

Imagine a small island with a local electricity network with no connection to a greater network. The village of this island has 10 houses. Half of the houses are heated by heatpumps, the other half by micro-CHPs. Apart from the heatpumps, the energy consumption within the houses is inflexible and following standard household load profiles. Further, on the island there is a wind-diesel combination delivering that part of the momentary electricity demand not supplied by the CHPs. This combined unit is operated to balance the island system. When the local demand is higher than the CHPs and wind turbine are producing, the diesel generator is regulated to maintain the momentary system balance. On the other hand, when local demand is lower than the CHP and wind generated power, the wind turbine is curtailed and regulated to balance the network.

In a small-scale simulation, the impact of the PowerMatcher was analyzed for the hypothetical island system described above. The simulation has been carried out for two distinct cases:

1. **Reference Case.** This is the business as usual scenario. The heating systems are controlled by a standard thermostat on/off controller. The system is balanced entirely by the wind-diesel system.

2. **Coordinated Case.** In this case the micro-CHPs and the heat pumps (HPs) are coordinated by the PowerMatcher. The multi-agent system tries to match CHP production and HP consumption with the inflexible demand and supply of the households and wind turbine respectively. Any net surplus or shortage is still balanced by the wind-diesel combination.

Table 2 gives the characteristics of the units used. The wind turbine output followed the measured production profile of a real-world turbine (Figure 6). The heating systems, i.e. the micro-CHPs and the heat pumps, were used for space heating alone. At this stage, hot tap water demand was left out of the scope of the simulation. The heat demand was generated using a basic thermal model of a house. The main external variable of this model is the outside temperature, which was set to follow a standard reference pattern. The household electricity consumption followed a standard residential load profile. Goal of the simulation is to give a proof of principle of the coordination mechanism, illustrating the cluster-level behavior.

The simulation spans a period of two days. Figure 7 gives output power of the diesel generator in the two cases. Two important effects can be seen from the figure:

1. The total production of the diesel generator is lower in the coordinated case (approx. 40%).

2. The peak load served by the diesel generator is lower in the coordinated case (approx. 45%).

The first effect is an important result as the environmental footprint of the island's electricity system is improved. Apparently, the wind generated power is utilized better in the coordinated case. More wind power is consumed and the turbine has been curtailed less. The second effect is important from an investment point of view. If the peak load on the diesel system is lower, the unit's design capacity can be lower which leads to a lower investment.

Figures 8 and 9, show the temperatures in the rooms heated by the heat pumps. The local PowerMatcher device agents make use of the inherent energy buffer in the inner

Table 2: Electricity producing (P) and consuming (C) units in the island simulation. The flexible units can be coordinated by the PowerMatcher.

Type	P_{max}	Number	P/C	Flex?
Diesel generator	15 kW	1	P	yes
Wind Turbine	30 kW	1	P	no
Micro CHP	1 kW	5	P	yes
Heat pump	0.7 kW	5	C	yes
Household Load	1.1 kW	10	C	no

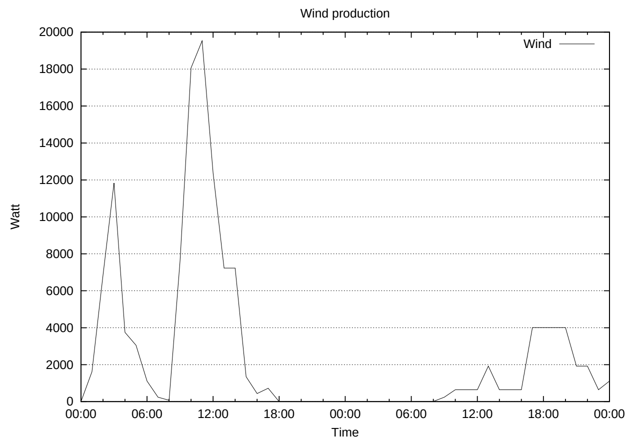

Figure 6: Power Output of the 30 kW wind turbine.

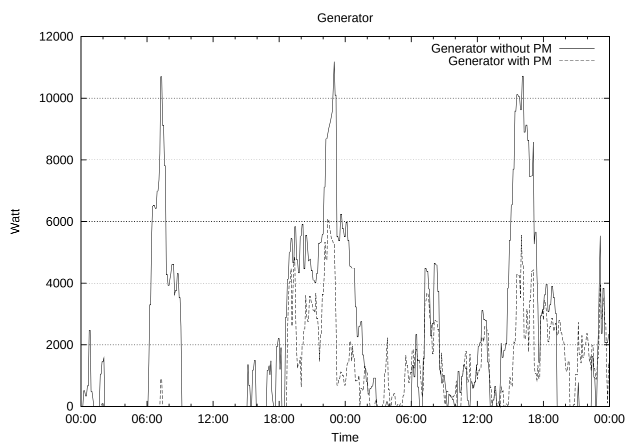

Figure 7: Diesel generator output power for the reference case (solid line) and the coordinated case (dashed line) over a period of two days.

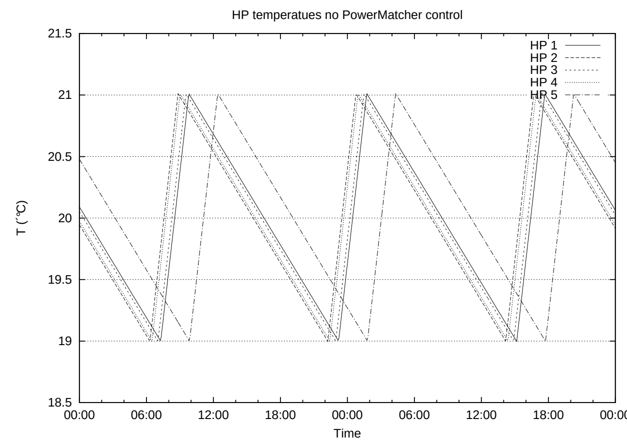

Figure 8: Room temperatures of the 5 heat pumps in the reference case. The basic On/Off controller behavior can clearly be seen.

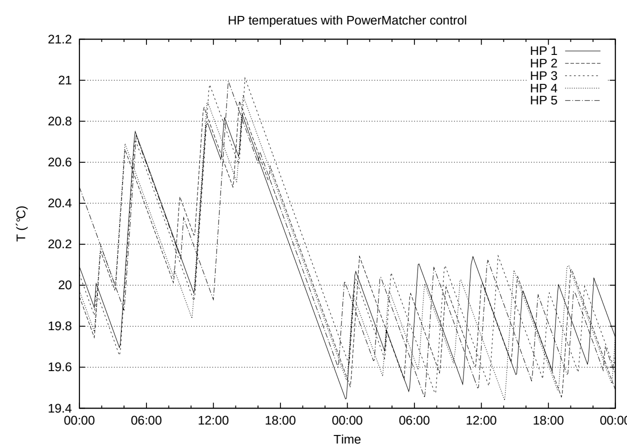

Figure 9: Room temperatures of the 5 heat pumps in the coordinated case. PowerMatcher control.

space of the houses to shift the heating operation. Note that at all times the comfort level is maintained. Figure 10 gives the price on the electronic market for the same time period. Note that the device agents in figure 9 try to heat the homes in the low-priced periods. The resulting price is influenced by a number of factors: (1) the momentary wind power availability, (2) the momentary household electricity demand, (3) the available operational flexibility of the micro-CHPs and the HPs. Note further that the diesel generator is only operated in the high-priced periods. Then, the cluster cannot provide the needed generation capacity, resulting in high prices and, in turn, utilization of the generator.

4.2 Weak Shore Connection: Congestion Management

Imagine, the island, as described in the previous subsection, does not have the diesel generator but a weak connection to a bigger network on-shore. In that case, the objective of the energy management system is not to balance the island system as much as possible and at all times. The objective is now to reduce the load on the connector at peak-load times.

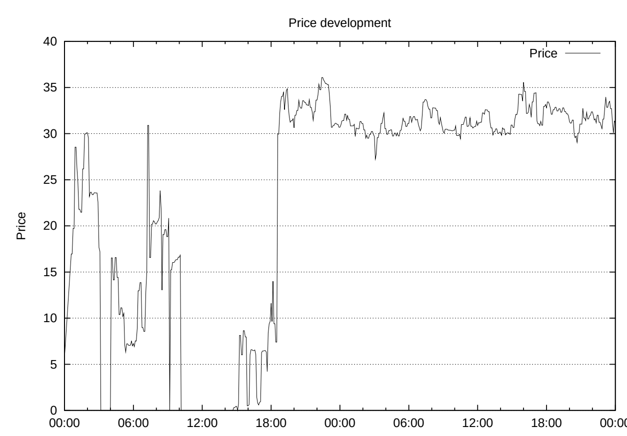

Figure 10: Price development of the PowerMatcher electronic market.

Objective agent: measure connection load. When exceeding the maximum value: do a bid for the excess load.

4.3 Virtual Power Plant Operation

Strong shore connection. Use PM to let the island follow a traded profile.

5. CONCLUSIONS

Currently, two major trends are changing the characteristics of electricity generation in the power infrastructure: the increase in both distributed generation (DG) and intermittent power sources. Multi-agent technology and electronic markets form an appropriate technology to solve the resulting coordination problem. The PowerMatcher technology as described in this article is a market-based control concept for supply and demand matching (SDM) in electricity networks with a high share of distributed generation. The presented preliminary simulation results give a proof-of-principle of this approach. The concept is capable of utilizing flexibility in device operation via agent bids on an electronic power market. Via agent reactions on price fluctuations, the simultaneousness between production and consumption of electricity by intelligent devices is increased. The study focusses on a micro-grid setting where balancing is done by a wind-diesel combination. Application of the PowerMatcher is shown to reduce the peak power delivered by the diesel generator by approx. 45% while the total diesel generated power decreased by approx 40%.

6. ACKNOWLEDGMENTS

The authors wish to thank the complete PowerMatcher team at ECN. We especially thank Cor Warmer and Pamela Macdougall for their valuable input to the simulation case. This work has partially been funded by the EU in the framework of the SmartHouse/SmartGrid project (FP7-ICT-2007-224628).

7. REFERENCES

[1] R. K. Dash, D. C. Parkes, and N. R. Jennings. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, November/December 2003.

[2] N. R. Jennings. Building complex software systems: The case for an agent-based approach. *Communications of the ACM, Forthcoming*, 44:12–23, 2000.

[3] K. Kok. Multi-agent coordination in the electricity grid, from concept towards market introduction. In *AAMAS 2010: Proceedings of the 9th int. joint conf. on Autonomous Agents and Multiagent Systems*, volume industry track, 2010.

[4] 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 R. H. Sprague, editor, *Proceedings of the 41st Annual Hawaii International Conference on System Sciences*, page 173, Los Alamitos, CA, USA, 2008. IEEE Computer Society.

[5] K. Kok, M. Scheepers, and R. Kamphuis. Intelligence in electricity networks for embedding renewables and distributed generation. In R. Negenborn, Z. Lukszo, and J. Hellendoorn, editors, *Intelligent Infrastructures*. Springer, Intelligent Systems, Control and Automation: Science and Engineering Series, 2010.

[6] K. Kok, C. Warmer, 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, pages 75–82, New York, NY, USA, 2005. ACM Press.

[7] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[8] A. Newell. The knowledge level. *Artificial Intelligence*, (18):87–127, 1982.

[9] B. Roossien. Field-test upscaling of multi-agent coordination in the electricity grid. In *Proceedings of the 20th International Conference on Electricity Distribution CIRED*. IET-CIRED, 2009.

[10] T. W. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 201–258. The MIT Press, Cambridge, MA, USA, 1999.

[11] H. R. Varian. *Intermediate Microeconomics—A Modern Approach*. W. W. Norton and Company, New York, 1996. Fourth Edition.

[12] C. Warmer, M. Hommelberg, B. Roossien, K. Kok, and J. W. Turkstra. A field test using agents for coordination of residential micro-chp. In *Proceedings of the 14th Int. Conf. on Intelligent System Applications to Power Systems (ISAP)*. IEEE, 2007.

[13] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[14] F. Ygge. *Market-Oriented Programming and its Application to Power Load Management*. PhD thesis, Department of Computer Science, Lund University, Sweden, 1998. ISBN 91-628-3055-4.

[15] F. Ygge and H. Akkermans. Resource-oriented multicommodity market algorithms. *Autonomous Agents and Multi-Agent Systems*, 3(1):53–71, 2000. Special Issue Best Papers of ICMAS–98.