

Multi-Agent Coordination in the Electricity Grid, from Concept towards Market Introduction

Koen Kok

Energy research Center of the Netherlands (ECN)
Power Systems and Information Technology
P.O. Box 1, Petten, The Netherlands
j.kok@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 micro-economics. This concept has been demonstrated in a number of field tests of increasing scale. Currently, the focus is moving from proof of concept field tests proving the technology towards demonstrations developing commercial applications paving the way for large-scale application.

Categories and Subject Descriptors

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

General Terms

Algorithms, Experimentation, Performance

Keywords

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

1. INTRODUCTION

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. World-wide electricity use has been ever-growing. To ensure the infrastructure continues to run in the future, the increasing elec-

Cite as: Multi-Agent Coordination in the Electricity Grid, from Concept towards Market Introduction, Authors, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1681-1688
Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

tricity demand poses a serious threat. Additionally, increasing generation from renewables, and the surge of distributed generators (DG), (i.e. the generation capacity embedded in the (medium and low voltage) distribution networks) are forcing a change in infrastructure management. Intelligent distributed coordination will be essential to ensure the electricity infrastructure runs efficiently in the future. As compared to traditional grids, operated in a top-down manner, these novel grids require bottom-up coordination in a highly-distributed manner.

Both the rising share of renewable energy sources in the energy mix and the decentralization of electricity generation are changing the characteristics of power generation in three aspects:

- **Intermittency:** The energy sources for conventional power generation are continuously available and can be adjusted according to the electricity demand. Electricity from sustainable energy sources, such as wind and solar energy, can only be produced if the primary energy source is available. With the growing share of these intermittent energy sources it becomes more difficult to follow the fluctuating electricity demand.
- **Cardinality:** As a result of generation decentralization, the number of electricity production units is growing rapidly while individual capacities are decreasing.
- **Location:** The location of power generation relative to the load centers is changing. Due to decentralization, the distance between generation units in the grid relative to the location of electricity consumption is becoming smaller. However, central renewable generation is moving further away from the load centers as large-scale wind farms are being built off-shore and large-scale solar power plants in desert areas.

A widespread form of DG is *combined heat and power generation* (CHP), an efficient form of fossil-fueled electricity production combined with production of useful heat. CHP units are typically operated to follow heat demand. Consequently, these units are intermittent in nature as well.

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 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.

In this paper, we describe the PowerMatcher, a multi-agent systems solution for coordination in the emerging sustainable electricity system. In section 2, the concept and agent structure is described. This includes a novel analysis of agent strategies based on marginal cost for flexible electricity consuming or producing devices and installations in subsection 2.3. We show the existence of a bid strategy spectrum and determine the position of particular real-world devices and installations in this spectrum. In section 3, the results of two field tests are presented, while section 4 describes a field test, currently in roll-out, which combines the functionalities of the first two. At the end of the paper, in section 5, we briefly describe planned steps towards large-scale commercial application of the technology.

2. THE POWERMATCHER

PowerMatcher is a general purpose coordination mechanism for balancing demand and supply in 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 with the first two versions being tested in simulations and field experiments [3, 1, 6]. The third version is planned to be deployed in a number of field experiments [5] and real-life demonstrations with a positive business case.

2.1 Logical Structure and Agent Roles

The PowerMatcher implements *supply and demand matching* (SDM) using a multi-agent systems and market-based control approach. SDM is concerned with optimally using the possibilities of electricity producing and consuming devices to alter their operation in order to increase the over-all match between electricity production and consumption.

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

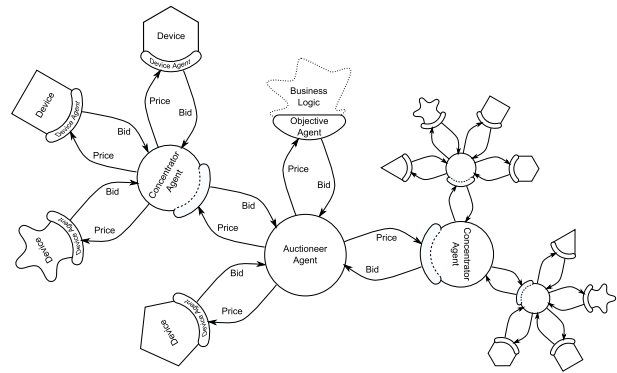


Figure 1: Example PowerMatcher agent cluster. See the text for a detailed description.

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.
- **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.

For a DER unit to be able to participate in a PowerMatcher cluster, its associated agent unit must communicate its momentary marginal cost to the Auctioneer Agent. This information is delivered in a *bid function* or *demand curve*: defining 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.

The logical agent structure follows the COTREE algorithm [7]. 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.

2.2 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 (subsection 3.1) and congestion management (subsection 3.2), this type of *near real-time coordination* suffices, as these processes take place on a similar timescale.

2.3 Agent strategies based on short-term economics

One of the key activities of a PowerMatcher cluster of agents is the delivery of near real-time balancing services. In order to operate such a near real-time coordination activity optimally, the agent society maintains a dynamic merit-order list of the (typically large number of) DER units participating. To make optimal decisions based on this list, the merit order needs to be based on the true marginal cost (or marginal benefit in case of demand response) of the individual DER units. The marginal electricity costs of most types of DER are highly dependent on local context and, hence, change over time. For example, the marginal electricity production cost for a CHP is highly dependent on the amount of heat demanded from the unit at a particular time. Thus, when the heat demand is high, the marginal cost for the electricity production is low and vice versa. The dynamic marginal cost levels of the units in the cluster cause the dynamic nature of the merit order list. As we will show later on, there exists a class of DER units for which, under circumstances, the marginal cost level cannot be determined unambiguously.

From a micro-economic viewpoint, the DER units are assumed to participate in a competitive market. This assumption holds when the number of DER units in the agent society is relatively high and their traded volumes are of the same order of magnitude. A competitive market leaves no room for speculation or gaming, and the best (*i.e.* the *dominant*) strategy for each participant is to optimize its own utility by truly bidding its marginal cost [4]. These locally-optimal strategies lead to a merit order list that results in an optimal allocation on the global level as well, as those

DER which are best fit to respond to a certain event are the first to be selected to do so.

The bidding strategy of a device agent is a mapping from its context history to a market bid. The context of a device agent includes:

- The process controlled by the agent, including the current state of the process and economical parameters such as marginal operating cost.
- The market environment in which this agent is situated, including the market mechanism and market prices.

In the extremes, there are two agent types that are forced to base their bid on either of the two context elements described above:

1. Those agents operating a DER unit that has clear and unambiguous levels of marginal costs. In a competitive market, the dominant strategy of these agents is to bid entirely according to their marginal operating costs.
2. Those agents operating a DER unit that does not have unambiguous marginal costs at all. In this case, the bidding strategy can only be based on market parameters, *i.e.* the market price (history).

As said, these cases are the extremes of a spectrum and hence, there is a group of agents whose bidding strategy is somewhere in between these extreme cases. In the next subsections we will give examples of these extreme and median cases.

2.3.1 Fully marginal-cost based

An example of a bidding strategy entirely based on the marginal cost level is that of a fuelled electricity generator set, for instance a gas generator set. The marginal cost for a given period of operation depends on the fuel price, the efficiency of the generator and the running-history dependent maintenance costs. Furthermore, each startup of such a generator causes additional costs for maintenance and fuel. The dominant strategy in this case is bidding a price equal to the marginal operation cost.

Thus, the optimal bidding function is given by:

$$d(p) = \begin{cases} 0 & \text{if } p < c_m \\ -P_g & \text{Otherwise} \end{cases} \quad (1)$$

where c_m is the marginal operation cost. Note that, by definition, $d(p)$ is negative in case of supply, hence the minus sign before the P_g term. It is clear that this bidding strategy depends entirely on the cost parameters of the generator. The market price history does not play a role in this strategy.

2.3.2 Fully price history based

At the other extreme is the bidding strategy of an electricity storage facility. Systems such as batteries, flywheels and pumped storage, charging from the electricity grid at one time and discharging to it at another. The aim of the agent is to buy electricity in periods of low prices, store it and resell in periods of high prices. Here, the notion of what defines a “high price” or a “low price” is crucial in the agent’s bidding strategy. Maximizing the agent’s utility comes down to determining the charge/discharge price that yields the best profit. This *optimal price set* is entirely dependent on

the dynamic price characteristics of the market environment plus the time needed for a full charge or discharge.

Charging and discharging a storage device is subject to *round-trip energy losses*. Note that, for the operation of a storage system to be profitable in the long run, the margin between the buy price and the resell price must exceed the costs for these losses. However, these costs do not influence the optimal price levels themselves.

Therefore, the agent requires some sort of function \mathcal{E} that yields estimates of the optimal charge and discharge prices given the current price history and the charging/discharging time:

$$\langle \bar{p}_c, \bar{p}_d \rangle = \mathcal{E}(H_p, T_s) \quad (2)$$

$$T_s = C_s / P_s \quad (3)$$

where P_s is the storage charging/discharging power, C_s is the storage capacity, T_s is the storage charging/discharging time, and H_p the price history vector. Based on these estimated price levels the bidding function can be defined by:

$$d(p) = \begin{cases} P_s & \text{if } p < \bar{p}_c \\ -P_s & \text{if } p > \bar{p}_d \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The long-run profit is highly dependent on the quality of the estimator \mathcal{E} , which must operate in dynamic market environments whose characteristics will be unknown at design time for most cases.

2.3.3 Median case: CHP/Gas heater combination

This case is based on configurations found in installations supplying heat to residential areas. A typical configuration combines a CHP, a more traditional gas heater and a heat storage buffer. An installation of this type was part of the field test cluster described in section 3.1.

The marginal cost levels depend on the following parameters.

η_{chp}^t	[]	Thermal efficiency of the CHP
η_{chp}^e	[]	Electrical efficiency of the CHP
η_{htr}^t	[]	Thermal efficiency of the heater
p_g	[ct/m ³]	Gas price
H_c	[kJ/m ³]	Gas combustion heat
T_{max}	[°C]	Upper limit inner temperature heat buffer
T_{min}	[°C]	Lower limit inner temperature heat buffer

Typically, the thermal efficiency of the heater will be higher than that of the CHP: $\eta_{chp}^t < \eta_{htr}^t$.

The heat demanded by the residential area is subtracted directly from the heat buffer. The local control goal of the CHP/heater combination is to keep the inner temperature of the buffer, T , between thermal limits T_{max} and T_{min} . Hence, the buffer level is defined as:

$$L_B = \frac{T - T_{min}}{T_{max} - T_{min}} \quad (5)$$

To prevent the buffer from over or under heating, three levels are defined at which special control actions are to be taken:

- L_H : High buffer level: just below the fill level of 100%. Above this level both the CHP and the heater must be

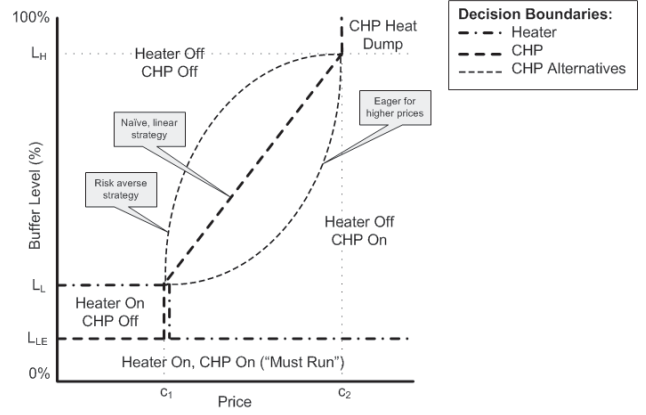


Figure 2: Bid strategy of a Heater/CHP combination as found in heat network systems delivering heat to residential areas. The strategy is well-defined below c_1 , the marginal cost for CHP-produced electricity when heat demand is high, and above c_2 , the CHP's marginal electricity cost when there is no heat demand at all.

switched off to prevent overheating. CHP operation is only possible in combination with heat dump, if that is technically possible (and ethically acceptable).

- L_L : Low buffer level: the level under which either the heater or the CHP must be switched on to prevent under heating.
- L_{LE} : Low emergency level: just above 0%. Below this level both heater and CHP must be switched on.

These levels define four different operational modes (see figure 2):

1. Below L_{LE} the high heat demand is the dominant factor in the operation of the installation. This is a must-run situation for both CHP and heater, regardless of the electricity price.
2. Between L_{LE} and L_L there is a heat demand that could be met by either the heater or the CHP. Hence, there is a choice of producing this heat using the heater or the CHP. In the latter case, the operating costs will be higher (as $\eta_{chp}^t < \eta_{htr}^t$) with additional electricity production in return. While the heat demand is covered by the CHP, the marginal cost of the additional electricity production is equal to:

$$c_1 = c_{chp}^t - c_{htr}^t \quad (6)$$

where c_{chp}^t is the marginal cost for heat produced by the CHP regardless the value of the co-produced electricity, and c_{htr}^t is the marginal cost for the heater-produced heat. With:

$$c_{chp}^t = \frac{p_g}{H_c} \eta_{chp}^t \quad (7)$$

$$c_{htr}^t = \frac{p_g}{H_c} \eta_{htr}^t \quad (8)$$

equation (6) can be expanded to:

$$c_1 = \frac{p_g}{H_c} (\eta_{chp}^t - \eta_{htr}^t) \quad (9)$$

The CHP is operated when the market price for electricity is higher than c_1 , otherwise the heater is operated.

- Above buffer level L_H , there is no heat demand. Hence, there is a choice to run the CHP and dump the produced heat. Even if the installation is not technically capable to discard CHP-produced heat, the marginal cost level of this option is of interest as it provides one of the strategy boundaries of the fourth operation mode, described below.

During CHP operation just for electricity production, the marginal cost for the electricity equals to:

$$c_2 = \frac{p_g}{H_c} \eta_{chp}^e \quad (10)$$

If the market price is above c_2 , it is profitable to run the CHP, even when the produced heat is discarded.

- In the region between L_L and L_H , there is a high level of freedom to let the CHP run dependent on the electricity price. At both boundaries of this region, the bidding strategy is well defined: at level L_L it is profitable to produce whenever $p > c_1$, while at level L_H it is profitable to produce whenever $p > c_2$. The ‘naive’ or ‘ignorant’ strategy would be to connect these two points linearly. However, dependent on both the dynamic price characteristics of the market *and* the used risk profile different trajectories are possible. In figure 2, two alternative strategies are shown. The risk-averse strategy tries to avoid must-run situations for both CHP and heater by taking the chance to fill the buffer whenever it is profitable to run the CHP. The other alternative strategy waits for higher prices to operate the CHP, with a higher risk of missing profit opportunities and ending in the must-run regions for heater and CHP.

2.3.4 Bid Strategy Spectrum

As becomes apparent, there exists a spectrum of DER bidding strategies. On one end of the spectrum, bidding strategies are based directly on true marginal cost or benefit. Along the spectrum, optimal bidding strategies become less dependent on marginal cost levels and more on the price dynamics in the (VPP) market context. As may be clear from the description of the CHP/Gas Heater combination, price-dynamics based strategies are not unambiguously defined but are dependent on a desired risk level.

In figure 3, the relative positions of a number of DER units are shown. Below, we discuss briefly the spectrum position of units not described previously.

- Generators of renewable power, such as wind turbines and photo-voltaic solar systems, typically have low marginal costs associated with them, as these consist mainly of maintenance costs. Fuel costs, the main marginal cost component for most other generation types, are essentially absent here. Therefore, the dominant strategy of renewables is to generate at any going electricity price. This positions them at the marginal-cost based extreme of the spectrum.
- CHP with heat buffer: In high-price situations, the bidding strategy of a solitaire CHP is similar to that




Figure 3: Bid Strategy Spectrum for Distributed Energy Resources based on momentary marginal cost levels.

of the CHP/Heater combination. The marginal cost for CHP produced electricity in the (theoretical) heat-dump case (c_2 in figure 2) is applicable here as well. However, the low-price behavior is dependent on the value attached (by the user) to a reliable heat supply and the risk level one allows for occasionally not being able to cover the heat demand entirely. Minimizing this risk is highly dependent on the prevailing price-dynamic characteristics. Hence, the position of CHPs on the right-hand side of the spectrum.

- Direct Electrical Space Heating or Cooling: Modern building constructions show relatively high degrees of thermal inertness. This can give some degree of freedom in the operation of systems for space heating and cooling, but is dependent on the current temperature and the temperature desired by the user. As field experiences learn, it is possible to shift cooling or heating periods forward or backward in time without infringing user comfort [6]. Here, the agent strategy goal is to provide the desired comfort level against minimal electricity costs, shifting cooling/heating actions towards low-priced periods as much as possible. Comparable to the strategy for storage units, the notion of what ‘low prices’ actually are is crucial for a successful strategy. This locates this DER type directly in the price-history based end of the spectrum. However, as experiences with demand response programs aiming at influencing user behavior learn, most users are willing to offer some comfort in order to avoid periods of high tariffs. Due to this, we position Direct Electrical Space Heating or Cooling just left of the spectrum end.
- Freezer: The case of a freezer is similar to that of that of space heating/cooling described above, hence the position near the price-history based end of the spectrum. As a minor difference, for this instance, the cost of ‘lost service’ is known as this equals the total value of the stored food items.

3. FIELD TEST RESULTS

3.1 Commercial Portfolio Balancing

This subsection describes the first field experiment performed using the PowerMatcher. A more comprehensive and detailed description of the test, including an analysis of the business model, can be found in [1].

3.1.1 Value Driver: Balancing Responsibility

To prevent black-out situations, the instantaneous supply and demand balance in the electricity infrastructure needs

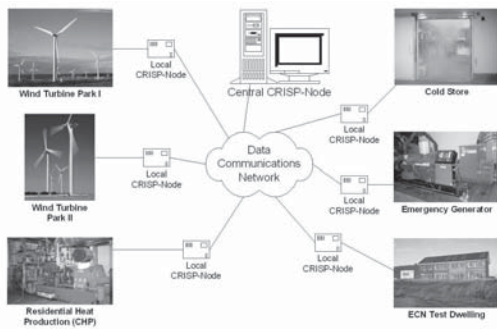


Figure 4: Configuration of the imbalance reduction field test.

to be maintained continuously. In regions where the electricity market is liberalized, the Transmission System Operator (TSO) maintains this balance. The TSO does so by monitoring the overall system balance on a timescale of seconds and adjusting contracted reserve power plants up or down to mitigate any occurring imbalance. A detailed description of this process is given in [2].

All parties active in the wholesale electricity trade are obliged to participate in the system of balancing responsibility. These parties are incentivized to balance their own commercial contract portfolio. This provides the TSO means to charge the costs made to maintain the real-time system balance to those wholesale market parties responsible for the imbalance. Central to this mechanism is the notion of *Balancing responsibility*, i.e., the obligation of wholesalers to plan their production and consumption and to make this plan available to the TSO. Parties having this responsibility are referred to as *balancing responsible parties* (BRPs). Although the TSO balances the over-all system on a seconds basis, settlement of imbalance caused by BRPs is done on a longer timescale, typically 15 or 30 minutes.

The system of balancing responsibility imposes imbalance risks to the market parties. Among BRPs, this risk will vary with the predictability and controllability of the total portfolio of the BRP. BRPs with low portfolio predictability are faced with higher imbalance risks. Typically, wind power suffers from low predictability. This gives higher imbalance costs resulting in a lower market value for electricity produced by wind turbines. In the last few years, day-ahead predictability of wind energy production has been improved substantially. However, a substantial error margin remains.

3.1.2 Field Test Set-up

For the purpose of the field test, five different installations were brought together in the portfolio of a virtual BRP. In reality, the installations represent a small part of the portfolios of two different BRPs, but for the sake of the experiment they were assumed to represent the full portfolio of one single BRP. Figure 4 illustrates the configuration of the field test. To all DER sites, hardware was added to run the local control agents. These agents interacted with the existing local measurement and control systems. Further, the local agents communicated with the auctioneer using a virtual private network running over a standard ADSL internet connection or (in one case) a UMTS wireless data connection.

Table 1 gives an overview of the capacity of the individual installations included in the test. In order to give the smaller

sized installations a good influential balance compared to the larger ones, two of the sites were scaled up via an on-line simulation.

Table 1: Production (P) and Consumption (C) Capacities of the Field Test Installations

Site	P/C	Capacity	Simulated
Wind Turbine	P	2.5 MW	-
CHP	P	6 MW	-
Cold Store	C	15 kW	1.5 MW
Emergency Generator	P	200 kW	-
Heat Pump	C	0.8 kW	80 kW

P = Production; C = Consumption.

3.1.3 Imbalance Reduction Results

The field test ran for a number of months in the first half year of 2006. In the real-life DER portfolio, with a wind power dominated imbalance characteristic, the imbalance reductions varied between 40 and 43%. As seen from an electricity market perspective, these benefits are substantial. This makes the approach a good addition to the current options for handling wind power unpredictability, such as wind/diesel combinations, balancing by conventional power plants and large-scale electricity storage.

3.2 Congestion Management

This subsection describes the second field experiment performed using the PowerMatcher. A more comprehensive and detailed description of the test can be found in [6].

3.2.1 Value Driver: Deferral of Grid Reinforcements

In the Northwestern region of Europe, decentralized generation of heat and power by micro-CHP units in households is expected to penetrate the market at a high speed in the coming years. When the number of micro-CHP units in a region exceeds a certain limit, added value can be gained by clustered coordination via common ICT systems. In a field test, a cluster of five Stirling based micro-CHP units (1kW electric each) has been operated as a virtual power plant¹. The main goal of the field test was to demonstrate the ability of such a VPP to reduce the local peak load on the single low-voltage grid segment the micro-CHP units were connected to. In this way, the VPP supports the local distribution system operator (DSO) to defer reinforcements in the grid infrastructure (substations and cables) when local demand is rising. Although not all micro-CHP units included in the field test were connected to the same low-voltage cable, during the trial a connection to a common substation (i.e., low-voltage to mid-voltage transformer) was assumed.

3.2.2 Field Test Set-up

The field test focused on the network utilization factor of the local distribution grid in three different settings:

- **Baseline:** domestic load profile of 5 households.

¹In total 10 micro-CHPs were equipped to be part of the VPP. The results presented are realized with 5 of these 10 participating.




Figure 5: Typical measured day patterns for 5 micro-CHPs with PowerMatcher coordination: synchronisation of CHP output (dashed line) with domestic peak-demand (dotted) leading to peak load reduction at the transformer (solid line).

- **Fit-and-Forget:** load profile of 5 households plus micro-CHPs controlled by thermostat in the standard heat-demand driven manner.
- **VPP operation:** load profile of 5 households plus micro-CHPs controlled by PowerMatcher device agents coordinated to reduce peak-load, without any intrusion on comfort for consumers.

In the third setting, the micro-CHPs were controlled by local PowerMatcher device agents. These agents were clustered together with an objective agent monitoring the load on the shared transformer and demanding CHP electricity production when it exceeded a safety level.

The households participating in the field test were equipped with a Whispergen micro-CHP for heating of living space and tap water. For the latter, these systems were equipped with a tap water buffer of 120 liters. During the field test, the systems were extended with a virtual power plant node or VPP-node. The local agents ran on these VPP-nodes, communicating with the local infrastructure (micro-CHP, thermostat, and electricity meter) through power line communications and with the auctioneer agent through a TCP/IP connection. The end users communicated with the system by means of the thermostat.

The local agents aimed at producing CHP electricity in high-priced periods with a hard constraint of not infringing the user’s thermal comfort. When the transformer load exceeded the safety level, the objective agent issued a demand bid aiming at steering the load back to the safety level. This increase in demand caused a price rise on the electronic market, which, in turn, triggered those agents most fit to respond (i.e., the ones having the highest heat demand at that moment) to run their CHP. The micro-CHP units were only operated in case of local heat demand, either for space heating or for tap water heating. No heat was dumped. An additional simulation study was done to verify the findings in the field test and to investigate circumstances not engaged in the field experiment, such as winter conditions.

3.2.3 Congestion Management Results

The field test was conducted in May 2007, which was an exceptionally warm month for The Netherlands. Therefore

there was no space heating demand in the households, only demand for tap water heating. Figure 5 shows a typical day pattern during the field test when five micro-CHPs were participating in the VPP. The PowerMatcher shifts the micro-CHP production so that electricity is produced when there is a high demand for electricity. This lowers the peak load on the substation. The main findings of the field experiment and additional simulation studies were:

- The Fit-and-Forget policy did not provide benefits to the DSO in comparison to the baseline case. The average load on the transformer was lowered as compared to the baseline due to the local electricity generation from the micro-CHPs. However, the transformer peak load remained virtually unchanged.
- Adding VPP operation, based on PowerMatcher intelligent control, led to a peak-load reduction of 30% in summer (field test result) and 50% in winter (simulation outcome).

4. MULTI-STAKEHOLDER FIELD TEST

In the previous sections, we argued there is necessity to introduce distributed control in the electricity infrastructure in order to cope with the interrelated trends of increasing sustainable electricity sources and distributed generation. We have shown how a specific implementation of distributed control can be used for commercial portfolio balancing as well as for DSO congestion management. An important remaining question is: how to combine the two?

In real-life situations, large networks with many stakeholders involved and having multiple optimisation objectives for different scenarios should be expected. Each stakeholder has its own interest and these interests will conflict at certain periods in time. For example, a low price of electricity may stimulate consumption of electricity, but the immediate resulting increase in consumption may overload the grid locally. The key stakeholders are:

- **The Prosumer:** an end-customer that may be capable of generating electricity by means of devices such as a micro-CHP or a PV solar system. Such a Prosumer primarily wants to maximise the economic value of their investment in such devices as well as minimise the costs for their consumption of energy.
- **The Distribution System Operator (DSO),** who operates the grid, wants to limit load fluctuations as much as possible by optimizing the usage of their assets in this way.
- **The Energy Supplier,** trades energy on the wholesale market, delivers electricity to the Prosumers, and buys surplus electricity back from the Prosumers.

Figure 6 shows the dual optimization market that is adopted, and depicts the commercial and technical optimisation comprises on which these three stakeholders operate. The objective of this field test is to find an optimal control solution for the electricity consuming and producing devices, so that the interests of all stakeholders are respected as fairly as possible.

In the *PowerMatching City* field test, Multi-stakeholder optimization is being tested under real-life conditions. The



Figure 6: Orthogonal dual market-based architecture for agent-based multi-stakeholder optimization in future electricity systems.

field test consists of a cluster of about thirty real life households and two supplemented laboratory sites, resulting in a total of approximately one hundred DER devices. The devices range from micro-CHP, heat pumps, photovoltaic, (urban) wind, household appliances (laundry) as well as plug-in electric cars. Implementation of the coordination will be done by creating a Virtual Power Plant (VPP) based on the PowerMatcher concept. The field test is currently in the roll-out phase.

5. OUTLOOK

The described field tests are all mainly focussing on technology aspects. Now that the technology has been thoroughly proven, our aim will switch towards the development of large-scale commercial demonstration projects. These next step demonstrations need to involve all relevant stakeholders and technology providers. For market introduction we are:

- establishing a partnership with a software vendor and system integrator for product commercialization, and making these partners responsible for systems integration and roll-out to ensure smooth transfer to the market.
- developing off-the-shelf products based on the PowerMatcher technology. Two partnerships with home automation manufacturers exist.
- aiming at the development of a standard communication protocol that encompasses the PowerMatcher protocol.

6. ACKNOWLEDGMENTS

This overview was partially funded by the EU: Smart-House/SmartGrid (FP7-ICT-2007-224628). The work in subsection 2.3 was partially funded by the Dutch Government: EIT (EOSLT02008). The author thanks everyone who contributed in any way to the field tests described, and Pamela Macdougall comments that greatly improved the manuscript.

7. REFERENCES

- [1] 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.
- [2] 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.
- [3] 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.
- [4] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [5] B. Roossien. Field-test upscaling of multi-agent coordination in the electricity grid. In *Proceedings of the 20th International Conference on Electricity Distribution CIREN*. IET-CIREN, 2009.
- [6] 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.
- [7] 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.