The HomeBots System and Field Test: A Multi-Commodity Market for Predictive Power Load Management

Fredrik Ygge † **Hans Akkermans** § **Arne Andersson**‡
Marko Krejic ¶ **Erik Boertjes** §

†EnerSearch AB
Chalmers Teknikpark
S-412 88 Gothenburg, Sweden
fredrik.ygge@enersearch.se
http://www.enersearch.se/~ygge

§Free University Amsterdam VUA
Department of Information Management & Software Engineering (IMSE)
De Boelelaan 1081a, NL-1081 HV Amsterdam, The Netherlands
{boertjes,hansa}@cs.vu.nl or HansAkkermans@compuserve.com
http://www.vu.cs.nl/~boertjes

‡Uppsala University
Computing Science Department
Box 325
S-751 05 Uppsala, Sweden
arnea@csd.uu.se
http://www.csd.uu.se/~arnea

†University of Karlskrona/Ronneby HK-R
Department of Computer Science (IDE)
S-372 25 Ronneby, Sweden
Marko.Krejic@ide.hk-r.se

Paper submission to PAAM'99.

Topics: Business agents for energy services; Negotiation; Price formation; Tools and techniques for process control.

Corresponding author: Erik Boertjes, Free University Amsterdam, Computer Science/IMSE Department, De Boelelaan 1081a, NL-1081 HV Amsterdam, The Netherlands. E-mail: boertjes@cs.vu.nl/hansa@cs.vu.nl; Fax: +31 20 444 7653; Tel: +31 20 444 7740/7700.

Abstract. *We present a system called* HOMEBOTS *for agent-based energy management services, realized by networked 'smart' industrial and household equipment communicating over the power line and other media. As a consequence of the deregulation of the electricity markets in many countries, energy utilities have started to pay high interest in offering value-added energy customer services rather than merely selling electricity (kWh). We discuss a number of important technical and business issues in launching such services, and describe some advanced solutions.*

First, we present a new computational market theory, implemented in the HOMEBOTS *system. It shows how large numbers of electrical loads can be automatically managed by autonomous agents, that communicate and negotiate in an electronic multi-commodity market leading to optimal use of electrical power. The advantages of this agent-based approach compared to traditional methods for power load management are described. Second, we demonstrate through simulated business scenarios that significant energy cost savings can thus be achieved. Third, our approach has been tested in a field experiment in an energy distribution area in the South-East of Sweden. The performed field tests show that the real-time requirements for agent communication over the power line in energy services are well met in realistic application settings.*

1 Agent-Based Energy Management

The electricity market in many countries has recently been deregulated or is currently under deregulation. At the same time we witness a rapid technological development, particularly in the area of access between global networks (such as the Internet) and individual households (e.g. through cable TV or even the electrical power grid), and in the area of incorporation of microcomputers in home appliances. Thus, an important question for power utilities is how this new ability to communicate with large numbers of electrical equipment (loads) can increase the competitiveness of the utility.

Value-added services. An interesting observation in this new setting is that utilities can offer new value-added services, such as providing a comfortable indoor temperature in a public building or providing ways to save costs by energy management. In the former case, the utility could have some contract with the customer which defines the economic terms for different levels of comfort (temperature). In the latter case, one could have contracts between utilities and customers with the aim to reduce the consumption at the customer side at certain (peak) hours. Namely, energy demand varies depending on the nature of the customer premises, and it shows big and costly fluctuations over time. One possibility to increase the efficiency of the energy system is therefore to try and manage customer equipment to reduce the temporal fluctuations in demand. The fundamental problem for the utility for these two examples is the same: When should the respective loads be at what consumption level (while considering the comfort constraints of the public building and the time constraints given by the contracts) such that costs are minimized?

Conventional power load management. This is called *power load management* [1, 2, 3]. The goal of load management is to move demand from expensive hours to cheaper hours. This reduces costs, curtails energy system over- or under-capacity, and enhances the utilization degree of investments in existing energy network assets. Energy load management is already an old idea used on a limited scale with large customers. When big peaks occur, utilities may be allowed to shut down certain customer processes in return for a certain financial compensation to the customer (direct load management). Or, such actions may be undertaken by customers themselves (indirect load management), often triggered by contractual penalties if the energy consumption exceeds specified ceilings. However, the current approaches to load management are restricted in scale and scope, as they essentially depend on one-way human decision making and control regarding a small number of energy-consuming devices and processes.

Due to the expanding capabilities of Information and Communication Technology, the current challenge is to automatically manage extremely large numbers of loads simultaneously. In particular, it should become possible to handle all electrical loads beyond the secondary substations on the 230V low voltage grid. In Europe, a typical 'cell' beyond such substations covers about 250 households, involving on the order of thousands of relevant loads. Similar numbers hold for office and plant sites. This is several orders of magnitude larger than in current forms of load management.

Given the different characteristics of different loads, different customer preferences etc., and the huge number of different loads, automatic load management can be viewed as a complex and highly distributed optimization problem.

Agent-based advances: HomeBots. Agent software technology, combined with recent advances in telecommunication, offers innovative ways to solve this problem. Electrical devices can nowadays possess communication and information-processing capabilities, by supplying them with networked microprocessors. This opens up new avenues for energy and telecom applications. In everyday language, it is now technologically possible that software-equipped communicating devices 'talk to', 'negotiate', 'make decisions' and 'cooperate with' one another, over the low-voltage grid and other media. So, we equip every load in the system with a software agent that acts as its representative. We call these software agents HOMEBOTS — a name inspired by smart equipment as featured in the tv series Star Trek and in Isaac Asimow's robot stories. The task of these agents is to take care of their electrical devices, such that they serve as a kind of personal assistant to help realize the customer's wishes.

We use this concept to achieve distributed load management in a novel fashion: by a cooperating 'society of intelligent devices'. It is the responsibility of each HOMEBOT agent to make the best use of electricity at the lowest possible cost. This may be realized by a proper timing of electricity use, by shifting the energy consumption of the represented device to cheaper periods. Of course, not many customers would want their tv set or microwave oven suddenly shut off for saving reasons, but there are many electric loads where time-of-use shifts are well possible. Examples are devices with recharge batteries, dish and laundry washers (for example in their use overnight), and especially space and water heating/cooling equipment, because these involve relatively slow thermodynamic processes with reasonable tolerances. Importantly, the latter are responsible for most (about 80% in many cases) of the energy consumption in residences and offices.

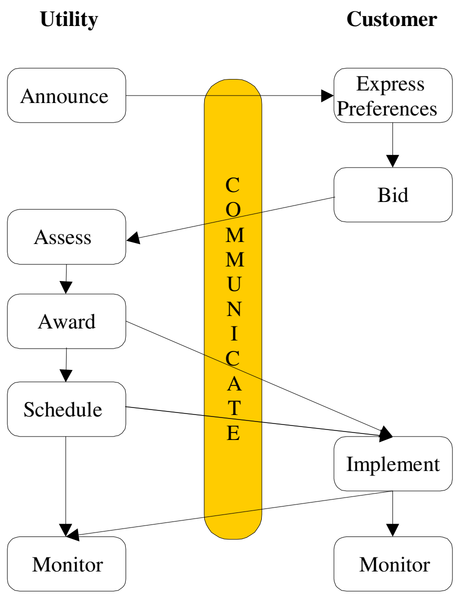

Figure 1: *A basic agent-based energy management scenario, depicting the agent tasks (boxes) and transactions (arrows) in an electronic market for energy load management.*

An important observation is that when each equipment agent would carry out time-of-use shifts to cheap hours in a completely individualistic fashion, a sub-optimal situation would result: all agents would move towards the same cheap spots, creating peak demands there instead. Hence, HomeBot equipment agents must communicate in order to arrive at the desired overall optimum load situation. This is achieved by their acting together on an electronic market for reallocating power by buying and selling power for different time slots. One possible, basic scenario for distributed load management based on agents acting on an electronic market is depicted in Figure 1. It shows the tasks carried out by software agents representing the utility and the customers, and the communication links (transactions) between the agent tasks. The bid-assess-award task cluster corresponds to the electronic market process.

In this paper, we elaborate this agent-based market for direct power load management. We introduce a novel multi-commodity computational market design for this problem in Section 2. A concrete example illustrating distributed load management is presented in Section 3. The simulations show that significant energy cost savings can be achieved, thus establishing the business case for new energy services such as agent-based load management. Our HOMEBOTS system has been implemented on top of a power-line telecommunication infrastructure in an electricity distribution area in the South-East of Sweden. Aim was to test the speed and reliability of equipment agent communication, with regard to the real-time requirements imposed on large-scale load management services. These tests have been performed in cooperation with a number of European utilities and system providers. In Section 4 we report on the setup and the results of the field experiments. The tests show that the real-time requirements can be successfully met with our agent-based approach. In Section 5 we summarize our conclusions, and point at uture research topics to be tackled in bringing new energy services on the market; many go beyond agent issues and deal with business process and customer interface aspects.

2 Multi-Commodity Electronic Market Design

This section describes our electronic market design for load management. The important design issues are the definition of the commodities, the agents, and the agent interaction protocol. The design deals with a multi-commodity market, an approach that represents a significant generalization over earlier two-commodity (electricity and money) approaches to load management (e.g. [4, 5]). It also takes some inspiration from similar systems for other application areas [6, 7, 8, 9, 10]. Furthermore, more realistic load models have been introduced compared to, e.g., [4].

2.1 Commodities

The power within different time slots represent the different commodities. A natural choice is to employ time slots of one hour, but it is well possible to employ more fine-grained and variable intervals. For example, one could consider using 15 minute time slots the first hour, 30 minute time slots the next two hours, 1 hour slots for the following 5 hours, and 4 hour time slots for the following 40 hours, i.e. a market of 23 commodities over a total time period of 48 hours. Of course, the more fine-grained the time scale (that is, a larger number of commodities), the more accurate the outcome, but also the heavier the computational burden.

For the prices and demands of the next few periods to be *exactly* correct a market with an infinite number of commodities, covering time infinitely into the future, would formally be required. Since this is impossible in practice, we set up a reasonable number of commodities, for example, the 23 commodities described above, and let the agents base their decisions for these time periods on some rough estimates of future values. For the example here, this comprises rough estimates of the prices of the 49th hour and beyond. In this paper we will let the estimate of those future prices be the price of the last time period covered by the market, although other alternatives are certainly conceivable [11].

Auctions (described below) can be performed regularly or when required. Typically, with the market setting defined above, a new auction would be performed every 15th minute. The reallocation of resources is performed as described by the bids of the upcoming time period. The future prices and allocations are established for two reasons: 1) the price and demand for the upcoming period can not be established without knowledge about future prices, and 2) the estimates of future prices and demands is useful information for the utility. The market outcome provides the utility with a control strategy for the upcoming time period and useful predictions about future prices and demands in a compact and uniform manner.

2.2 Agents

Every load in the system is represented by a computational agent, a HOMEBOT. The responsibility of a HOMEBOT is to use electricity as efficiently as possible, given the customer preferences, the load state, consumption predictions, and the load model. The *customer preferences* are represented by a customer contract. The indoor temperature of a building, or the amount of water in a warm water heater are examples of the *load state*. The *consumption predictions* are preprogrammed or learned consumption patterns. Physical load characteristics, such as time constants, are contained in the *load model*. This is illustrated in Figure 2.

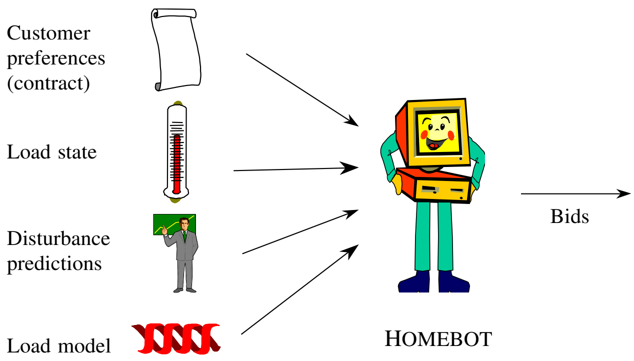

Figure 2: *The inputs and the output of a* HOMEBOT *agent.*

A HOMEBOT is modeled as a utility maximizing agent, whose preferences are given by a quasi-linear *utility function*[1] defined by

$$u_\alpha(\mathbf{r}, m) = f_\alpha(\mathbf{r}) + m, \tag{1}$$

where $\mathbf{r} = [r_1, r_2, \ldots, r_k]$ is the resource (active power) of the different time intervals $(1 \ldots k)$, $f_i(\mathbf{r})$ represents the benefit or cost (negative benefit), associated with $\mathbf{r}$. That is,

[1] A utility function is essentially a preference ordering: a high utility for some allocation expresses that this allocation is preferred over another allocation with a lower utility.

it captures the customer preferences, the load state, consumption predictions, and the load model. The money, m, corresponds to a monetary value in a real currency.

The task of each HOMEBOT is to maximize its utility through trade with other agents. In such trade the HOMEBOTS are programmed to act *competitively* (equivalent to acting as a *price-taker*) [12, p. 20, p. 314]. So, they treat prices as exogenous, rather than speculate about the effects of their own actions on market prices (*cf.* [13]).

Also uncontrollable loads are represented by HOMEBOTS. In this case, however, the HOMEBOT is totally insensitive to prices. Such a HOMEBOT can represent one or a number of loads. A HOMEBOT representing an uncontrolled load utilizes a suitable method for prediction, see for example [14, 15, 16].

Agents representing production units are modeled as profit maximizing producers, i.e. they solve the maximization problem:

$$\max_{\mathbf{r}} \mathbf{p} \cdot \mathbf{r} - cost_\alpha(\mathbf{r}), \tag{2}$$

where $\mathbf{r}$ is a *produced* amount of resource, and $cost_\alpha(\mathbf{r})$ is producer α's cost associated with producing $\mathbf{r}$. Also the production agents are programmed to act competitively.

2.3 Market Interaction

Agent system architecture. The structure of the load management system is typically as in Figure 3. The power distribution system is inherently very hierarchic, and the structure of the load management system normally reflects this.

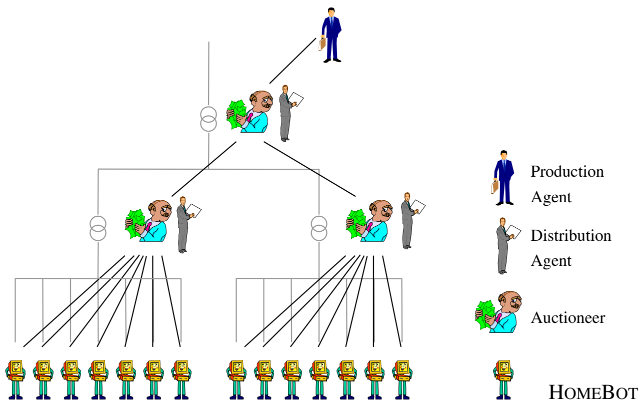

Figure 3: *A typical architecture of a load management system. The* HOMEBOTS *are connected to auctioneers which, in turn, are connected to other auctioneers in a hierarchic manner.*

Where the agents are physically placed, depends on communication and computation characteristics of the hardware. For example, if the communication is very fast throughout the system and if there are powerful micro-computers at each load (e.g. 'smart refrigerators') then it makes sense to place the HOMEBOTS at the loads for modularity reasons.

On the other hand, if the communication between the HOMEBOTS and the first auctioneer (for example placed in the secondary substation) is relatively slow and the processing power at the loads is very limited, then all agents belonging to the same secondary substation area might be placed on the same host for efficiency reasons.

The agents interact through *auctions*. In an auction agents submit competitive bids and an auctioneer computes a general equilibrium (i.e. a set of prices such that supply meets demand for each commodity). In a distributed setting typically a number of auctioneers are used in an hierarchic fashion, so that the computational burden is distributed and parts of the system can function even when failures are present, see further [17]. The market dynamics is depicted in Figure 4.

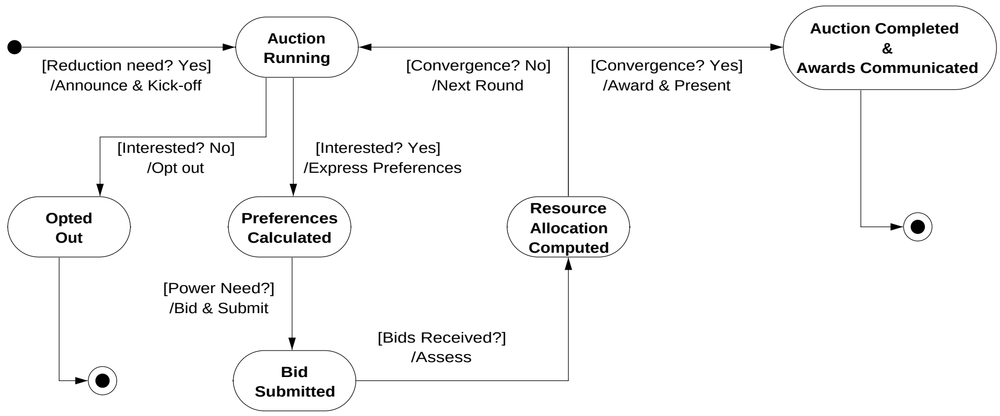

Figure 4: *The dynamics of the energy auction, in the form of a UML state diagram. Boxes indicate states, arrows stand for state transitions. The arrow labels tell under what conditions state transitions occur and what (trans)actions are then carried out by the HomeBots system. The solid ball (upper left) denotes the starting state of the auction, and the encircled balls are end states.*

More formally, with a price-oriented approach, each agent computes its *net demand* [18] (the change in allocation it desires at the going prices), denoted $z_{\alpha i}(\mathbf{p})$, for agent α and commodity i at the prices $\mathbf{p} = [p_1, p_2, \ldots, p_k]$, where p_i is the price for commodity i. An auctioneer for a set, A, of agents then computes a price vector such that $\forall i \sum_{\alpha \in A} z_{\alpha i}(\mathbf{p}) = 0$.

A typical market algorithm for finding this price is a standard multi-variable Newton-Raphson scheme which reads[2]:

$$\mathbf{p}^{i+1} = \mathbf{p}^{i} - s \cdot \bigtriangledown \mathbf{z}^{-1}(\mathbf{p}^{i}) \cdot \mathbf{z}(\mathbf{p}^{i}), \tag{3}$$

where $i + 1$ and i denote iterations, s is a step size, and $\bigtriangledown \mathbf{z}(\mathbf{p})$ is the gradient matrix defined by $\bigtriangledown z_{ij}(\mathbf{p}) = \frac{\partial z_i(\mathbf{p})}{\partial p_j}$. The proper value of s can be determined at run-time by

[2]This is just an example of a *price-oriented algorithm*, but there exist many others. An alternative family of algorithms which is useful for finding the equilibrium are *resource-oriented algorithms* [19].

a backtracking algorithm [20, pp. 384 – 385]. As the computational task of the auctioneer includes summation of the n demand functions, solving the $k-1$ linear equation system the update scales with the number of agents and number of commodities as $\mathcal{O}(nk^{2.496})$ [21]. This is, under some smoothness conditions, a scheme with *quadratic convergence* [20]. Hence, the required number of iterations is $\mathcal{O}\left(\log(-\log\epsilon)\right)$, where ϵ is the error. Schemes used in many other computational markets [6, 9] have only linear convergence, scaling with $\mathcal{O}(-\log\epsilon)$.

The many simulations we have performed with our market schemes have shown that the HOMEBOTS approach is computationally very fast and efficient. A typical figure is that market convergence is achieved in just a few iterations [4, 5, 11]. The scalability also proves to be highly satisfactory: having even thousands of equipment agents simultaneously involved in an auction presents no problem whatsoever.

An important issue when using general equilibrium as a market mechanism is the existence and uniqueness of the equilibrium. A sufficient condition for the equilibrium to exist and be unique is that all demands are continuous, that the demand for a commodity always decreases with the price of the commodity, and that each agent has gross substitutability between the commodities, cf. [22]. Without going into detail here, we note that these conditions typically are fulfilled (with high accuracy) in an energy system as a whole [3, 11].

Market outcome. When market equilibrium has been obtained, the resources are allocated (the award task in Figure 1). It can be proven [11, Theorem 3.2] that this allocation is a globally *optimal* allocation, given the time division and available information. Thus, no other approach, agent/market-based or not, can do better, cf. [10]. Moreover, any *Pareto-efficient*[3] allocation with the described utility functions is a globally optimal allocation. Hence, any market mechanism which generates a Pareto-efficient allocation can be used for obtaining a global optimum. Our main reason for choosing the general equilibrium approach is the computational efficiency of such an approach.

Another salient feature of our market-based approach is that every agent can estimate the value of the contract it is holding, by comparing how much money is spent with and without the contract and compute the difference. This implies that it is possible to perform an on-line cost/benefit analysis for each and every contract in the system — a feature highly attractive for energy utilities.

[3] An allocation is *Pareto-efficient* if and only if there does not exist an allocation which is better for some agent without being worse for any other agent.

3 A Load Management Example

We will now demonstrate and simulate the market design for a specific example. The example includes one production unit, nine warm water heaters, one public building and a number of uncontrollable loads as depicted in Figure 5. While we mainly aim at demonstrating the general principles of the approach here, most numerical details have been omitted for readability and space reasons, but can be found elsewhere [11].

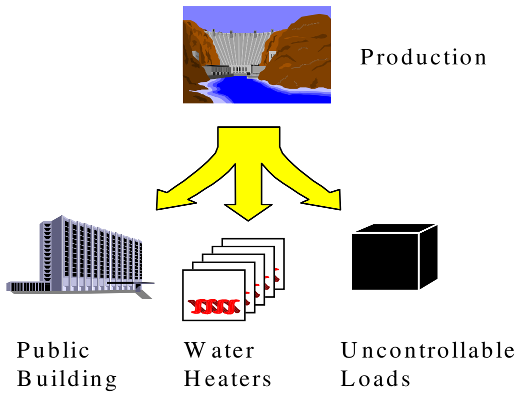

Figure 5: *A small part of an energy system — a producer and its customers.*

3.1 Load and Production Characteristics

We investigate a four-hour time period which can be thought of as morning hours where many people have taken morning showers and have gone to work. Thus, no further warm water is consumed the coming few hours. That is, the heaters will at this point start to heat water until they are fully heated. Let us suppose that the utility has a contract with each customer allowing the utility to switch off the loads for a certain amount of time during the four-hour period. The interesting parameters for each load are how much it consumes when it is switched on, how much energy it needs to be fully heated, and for how long it can be disconnected.

In addition to the water heaters, suppose there is a public building supplied by the producer. It is assumed that the utility has a contract with the customer making the utility responsible for maintaining a certain indoor temperature of the building for a certain annual fee. If the utility causes a deviation in temperature, the utility must pay the customer a compensation.

The heating system consists of electric radiators. There is a lower constraint (of 10kW) for the minimal power the building can be assigned at the investigated time periods. (This lower constraint is often included in load management contracts in order to avoid unpleasant draft from windows etc.) In addition to the controllable loads, there are some uncontrollable loads with the expected consumption 10, 400, 100 and $80kW$ for each of the coming next four hours. Thus, there is a significant peak during hour two. The production cost is assumed to be $10^{-3}r_i^2$, where r_i is the amount of produced power at time i.

3.2 The HOMEBOTS and the Production Agent

We analyze a market of four commodities, representing one hour each. In this example the HOMEBOTS representing the warm water heaters will have a utility function as described by Eq. (1) with $f_i(\mathbf{r}) = 0$ when the consumption is such that the contracted disconnection time is not violated, and $f_i(\mathbf{r}) = -C$, where C is a large constant, otherwise. That is, as long as the contracted disconnection time is not violated, there is no cost for the load management, but if the contract is violated, there is a relatively high cost.

The HOMEBOT representing the public building has a utility function based on the characteristics of the building and the service contract with the customer. (For details, see [11, pp. 165 – 166].) The uncontrollable loads are represented by a HOMEBOT with the demands 10, 400, 100 and $80kW$ for the respective time intervals, independently of prices.

3.3 The Market Outcome

The above agents were implemented together with an auctioneer using a price-oriented Newton scheme, see Eq. (3). In Figure 6, the allocations in presence and absence of load management are compared, and the corresponding market prices that emerged from the auction are shown. We see that the result is a shift of load from time period two to the other periods, and that the market price of time period two has been reduced significantly. For this example, as little resource as possible is used for time period two by all involved controllable loads. However, it is still significantly higher than the prices of time period 1, 3, and 4, due to the large contribution of uncontrollable loads, and it could be profitable to arrange load management contracts with more customers in order to shift more load from time period two to periods 1, 3, and 4.

In addition to the fact that the peak hour cost has been reduced by approx. 33%, the total cost summed over time decreases by about 15%, as a consequence of our electronic market.[4] Thus, there is a significant business case for agent-based energy management services.

[4] The peak hour cost in the uncontrolled case is $1.03 \cdot 515$, whereas in the controlled case it is $0.83 \cdot 414.5$ from production plus 12.32 as compensation to the customer. The total cost during the four periods in the uncontrolled case is $0.456 \cdot 228 + 1.03 \cdot 515 + 0.4 \cdot 200 + 0.36 \cdot 180 = 779$. In the controlled case the corresponding figure is $0.48 \cdot 239.8 + 0.829 \cdot 414.5 + 0.46 \cdot 230 + 0.41 \cdot 205.2 = 648$ from production cost plus $2.80 + 1.09 + 12.32 + 0.53 = 16.74$ as compensation to the customer, that is 665 in total.

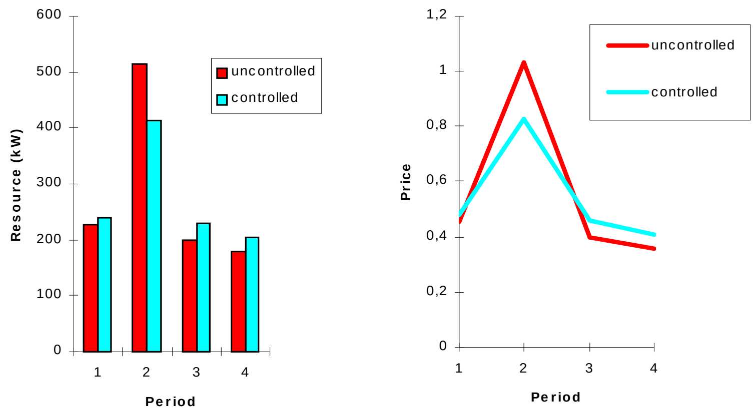

Figure 6: *To the left, the used power in different time periods with and without load control are shown. The result of the load control is that load is moved from period two to the other periods. During period two merely* $410kW$ *are used in the controlled case (*$400kW$ *from the uncontrollable loads and* $10kW$ *from the lower constraint of the resource allocated to the public building) compared to* $515kW$ *in the uncontrolled case. To the right, the market prices (the marginal production cost) of the different time periods are visualized. The market price of time period two has been reduced significantly. Total costs have been reduced from* 779 *to* 665*, i.e. by approx.* 15%.

4 Field Tests

4.1 Power Line Communication

Digital communication over the power line is becoming a promising channel for energy providers to implement new services [23]. Typical applications include remote meter reading, remote control tasks, load management, tariff-switching etc. Smart home technology is one of the areas where power-line communication is a key technology. Meter reading, tariff switching etc. are examples of applications with rather low demands on the communication system. Despite this fact there are several specific properties of the power line that have to be considered in the design of reliable systems involving power line communication. The main reason is that the power line was not designed for communication purposes and its properties as a communication channel are still not fully understood. The power line is a noisy medium, and the information bit rate it allows is still limited. The current state of the art thus poses constraints on the communication speed and bandwidth, a factor to be taken into account in designing practical applications such as load management, although technological progress is rapid. An important business consideration for the current strong interest in power line communication is, however, that *no new wiring* is required, a feature of prime importance in local telecom access technology [23].

4.2 Field Test Architecture

In the town of Ronneby (in Blekinge, a county in the South-East of Sweden), EnerSearch is conducting agent-based load management tests together with a number of European energy utilities and system suppliers, such as ABB, IBM, Iberdrola, PreussenElektra and Sydkraft. Equipment for communication over the power line has been installed in Villa Wega, a large villa used as an office building for the university, and in 70 households in Ronneby. The HOMEBOTS system has been connected to the various components of the heating installation of Villa Wega, consisting of 28 electrical radiators.

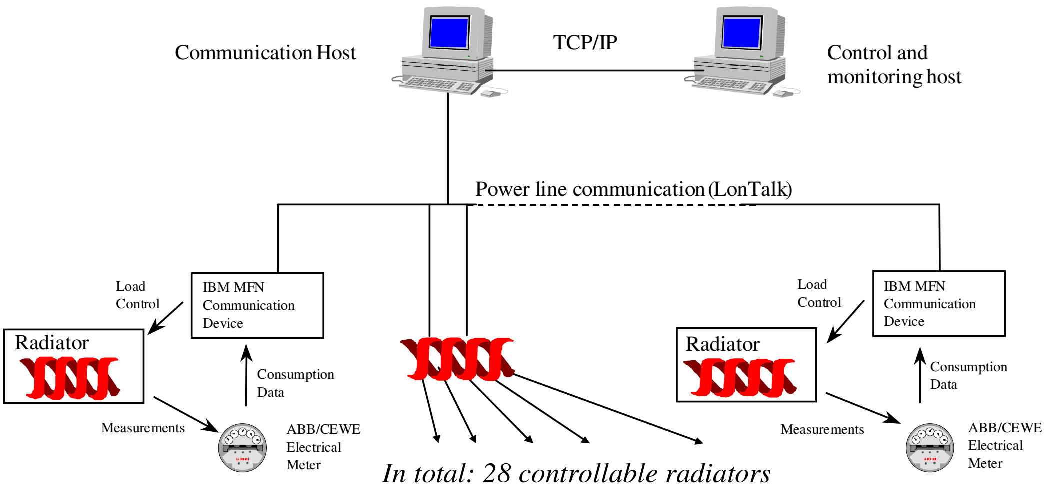

Figure 7: *The architecture of the Villa Wega field tests.*

Each individual radiator is separately metered and able to communicate with the HOMEBOTS system over the power line through the LonTalk protocol. This telecommunication includes the capabilities to read the current load state of any individual radiator, to change this state by giving on/off switching instructions over the power line, and to obtain real-time data concerning the accumulated energy consumption for each device. The installed hardware allows all agents to run on a secondary substation computer (essentially a standard IBM compatible PC). To this end, on each radiator an IBM 'Multi Functional Node' (MFN) is installed. Each MFN receives instructions from the HOMEBOTS system and is able to switch a radiator on and off. It is also capable of counting so called $s0$ pulses. These pulses are generated by an ABB energy meter that is also attached to each separate radiator. This meter generates a fixed number of pulses per kwH, thus indicating the amount of energy consumed by the radiator. The MFN returns these consumption values to the HOMEBOTS system. The hardware and telecommunication architecture of the field experiments in Villa Wega is presented in Figure 7.

4.3 Experimental Program

Performed experiments investigated the suitability of power-line communication for load management as described in this paper. All tests outlined below were carried out with the configuration in Villa Wega of Figure 7 and as have been done several times, during low, normal and high disturbance periods (night, working hours). Three sets of tests, each aimed at testing a different aspect of the system, have been carried out.

Tests on communication speed and reliability. A Read Load State message (i.e. a request for an indication of whether the radiator is switched on or off) was sent to every MFN, for approximately 1000 times in total. The following data were collected:
– Success/type of failure;
– Delay time;
– Load state value.

Similar tests have been performed for other types of messages: the Accumulated Consumption Message, indicating the energy consumption, and the Write Load state message, a request to switch on or off the radiator. Also sending combinations of messages was tested. This first test set has provided extensive information on functioning, speed and reliability of IBM's IDAM installation for power line communication.

Tests on the load scheduling methods. An important post-market process is the scheduling of loads, in order to implement load management in accordance with the outcomes of the market. This scheduling agent task (see Figure 1) does require power line communication with all individual loads. After the auction has been completed, its outcomes must be scheduled in accordance with the awarded power over some agreed period, say, the next hour. This is implemented through appropriate on/off switching of the involved loads over the power line. The awarded power must be scheduled over time, such that (i) the agreed amount of power is delivered averaged over time, while (ii) at the same time fluctuations and on/off switching costs are minimized. Computationally, this is a matter of satisfying a number of simultaneous constraints. Also this process is carried out automatically, and special algorithms have been developed for this scheduling of power delivery. The average on/off switching interval can thereby be chosen as a free parameter. Also this agent process has been tested in our field experiments.

As the initial allocation, a randomly generated one is taken. Tests aimed to study whether the scheduling algorithm worked properly in scheduling the various loads. The following data were collected:
– the allocated resource (per load and in total);
– per time slot, the current consumption for every load, plus the sum total. The total number of time slots is a few dozens, while a single time slot is on the order of one minute.

The same test has been performed but now with a number of (simulated) reallocations during the process, roughly every few minutes. For every reallocation, the allocated resources were collected. The same test has been performed again with different, more severe, hard constraints on the total resource.

Resource allocation plus scheduling full tests. These tests represent an integrated test of our whole load management scheme. Again the scheduling algorithm was tested, but now with the resource allocation computed from the utility functions and the associated market. The scenario is based on reduction (or increase) by a predefined amount of resources. The following data were collected:
– the allocated power resources (per load and in total, at every reallocation);
– per time slot, the current consumption for every load, plus the sum total;
– the temperature for every time slot;
– the utility function.
The same test has also been performed, based on a scenario where a different compensation price is offered for buying back power.

4.4 Results

Table 1 shows some typical results for the measured total communication times of the experiments described in the previous section. These communication times are defined as acknowledged messages to and from our Java application code to the processors at the loads. For a vast majority (over 90%) of the messages, the response times lie in the interval between 0.54 and 0.99 s, so that the median is significantly below 1 second. Average delay is 1.29 s and the standard deviation is 2.11 s. The clustering of communication times is related to re-send timer parameters. The success rate was over 98%.

Delay (s)	Frequency
$0.54 \leq d < 0.99$	11929
$0.99 \leq d < 1.43$	473
$1.43 \leq d < 5.00$	3
$5.00 \leq d < 5.90$	395
$5.90 \leq d < 9.48$	9
$9.48 \leq d < 10.37$	134
$10.37 \leq d < 13.06$	1
$13.06 \leq d < 13.95$	257

Table 1: *Some typical experimental data from the* HOMEBOTS *field tests in an office building, for a total of* 13201 *messages.*

The messages have typical lengths of 10–160 bits, including additional data (like CRC info) required by the LonTalk protocol. Hence, although the maximum possible bitrate of power line telecommunication is rather limited with currently available commercial equipment, the test results show that it is adequate for power load management applications.

Thus, the various communication tests clearly show that the relatively non-intensive communication required between each agent and its respective load can be successfully performed over the electrical power line. Inter-agent communication in the secondary substation area is then efficiently managed by a single host, and interagent communication between larger groups of agents can be performed in a hierarchic manner on faster networks (cf. Figure 3 and [17]). In sum, the main conclusion of the performed field tests is that agent- and market-based load management is technologically feasible in many realistic customer settings: response times are sufficiently short and the reliability is acceptable.

5 Conclusions and Future Work

In this paper we introduced a novel multi-commodity market design for load management. This approach to load management has a number of advantages compared to existing methods (such as [2, 24, 3]):

1. It provides an integrated strategy for many different types of loads and contracts. They may vary from low level contracts allowing the utility to switch off loads for certain amounts of time to high level services, like indoor temperature control.

2. The outcome is of very high quality, typically a very close approximation of the theoretical optimum is obtained.

3. It enables natural decomposition, both from a software engineering perspective as well as from a computational perspective. All local characteristics are encapsulated by agents, communicating only through prices and demands while doing local optimization computations.

4. It can be efficiently implemented (for the two commodity case, see, e.g., [4, 17], and for the multi-commodity case, see, e.g., [25, 19, 11]).

5. The main abstractions used, price and demand, are probably the most natural abstractions to use for a utility.

6. The utility is provided with a compact and uniform estimate of the energy system characteristics (present and future) in terms of prices and demands.

7. A local estimate of the value of the load management contract is obtained. This enables the utility to do continuous on-line cost/benefit analysis of every load management contract in the entire system.

Clearly, it can be argued that there might be competing approaches (from, e.g., (distributed) mathematical optimization or resource allocation) that can be applied to load management fulfilling items (1), (2), (3), and (4). The problem does not necessarily have to be modeled as a market in order to set up utility functions and perform optimization in a distributed fashion. We do not disagree with such a position, but argue that the integrated view of all types of loads and contracts, the high quality outcome, the computational and conceptual decomposition of the global problem into small pieces of locally optimizing software, and the efficient algorithms, together form an attractive approach to load management. This holds regardless of if our approach is described with classical mathematical optimization and resource allocation concepts or if it is treated with market abstractions.

Our main arguments for a market view of the system are items (5), (6), and (7). Rather than the computational aspects, we believe the naturalness of the approach to be of vital importance. The successful integration of a load management system into the core business information management of a utility is heavily dependent on how the system is appreciated by the people in the organization. Staff responsible for energy trading is very familiar with concepts such as demand and prices. Therefore, the metaphor of local agents 'negotiating' over power and thereby obtaining market equilibrium is a very attractive one which simplifies the understanding of the principles of the system. The issue of naturalness is also important from the software engineering perspective. Along the same line of reasoning, the estimate of the energy system characteristics in terms of prices and demands is probably the most useful one for the utility. Furthermore, well known concepts from economics, such as price elasticity are directly applicable. Finally, the feature of local evaluation of load contracts for every load in the entire system is only inherent in a market approach, and we believe it to be very useful. Thus, the market design tightly fits the real situation conceptually.

In summary, our HOMEBOTS agent approach to power load management succeeds in reducing both peak loads and overall cost of power delivery and consumption. The performed field tests clearly demonstrate the technical feasibility of using the electrical power line as a communication medium for our approach.

There are a number of areas that we will further investigate in the near future. One issue is to further expand the range of application situations, business scenarios and customer interests that can be handled by electronic markets. The HOMEBOTS technology has the potential to support many more applications than could be pointed out here. We will develop detailed business scenarios that serve as a way to illustrate possible new interactive services. These scenarios will provide a basis for discussions with utilities and customers to explore needs and to elicit feedback. On this basis, a generic formal description of new interactive services will be derived.

Another important topic concerns information integration: load management will be but one of the new ICT-based services running on the same power-line information and communication infrastructure. A significant part of the information involved will be needed across different services. This raises the question how to connect and integrate relevant pieces of information. Here, reusable (service and system) models and architectures, and (ontology) mechanisms [26] for the sharing of information and meaning are to be researched. An issue in customer requirements engineering —not addressed in economic science or market theory, by the way— is how one actually determines a utility curve in a practical and individual case. In general, this will depend on various factors including personal characteristics of the customer, the underlying business models, as well as technical models and data concerning the functioning of devices (see Figure 2. Handling this aspect properly in preference- and style-sensitive user interfaces is a key element in any new ICT-based service, because it defines how the service is presented to and interacts with the individual customer. Now that the underlying software agent and market technology has successfully demonstrated its feasibility, the major next step is therefore to design and experiment with the full service chain from utility to customer.

Acknowledgements. We thank Richard Krejstrup for his work related to the field tests, and Rune Gustavsson for his support at HKR, Ronneby. This work has been partly sponsored by EnerSearch AB (`www.enersearch.se`), NUTEK (`www.nutek.se`) and the Netherlands Energy Research Foundation ECN (`www.ecn.nl`), in the context of the international university-industry research collaboration programme *IT in Energy*.

References

[1] H. Ottosson, G. Lindbergh, B. C. T. Bengtsson, M. Håkansson, L. Larsson, B. Linder, T. Löfquist, and S. Örthagen. *Teknikstyrning av Elanvändningen, Rapport 1.* Sydkraft, Malmö, Sweden, 1989. Only available in Swedish.

[2] A. R. A. Wafa. Optimum load management securing minimum revenue loss in industry. In *Proceedings of Distribution Automation/Demand Side Management (DA/DSM) '96*. PenWell Conferences and Exhibitions, October 1996.

[3] S. Talukdar and C. W. Gellings. *Load Management*. IEEE Press, New York, 1987.

[4] F. Ygge and J. M. Akkermans. Power load management as a computational market. In M. Tokoro, editor, *Proceedings of the Second International Conference on Multi-Agent Systems ICMAS'96*, pages 393–400. AAAI Press, Menlo Park, CA, December 9–14 1996. (Available from http://www.enersearch.se/ygge).

[5] J. M. Akkermans and F. Ygge. Smart software as customer assistant in large-scale distributed load management. In *Proceedings of Distribution Automation/Demand Side Management (DA/DSM) '97*. PenWell Conferences and Exhibitions, October 1997. (Available from http://www.enersearch.se/ygge).

[6] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research (http://www.jair.org/)*, 1(1):1–23, 1993.

[7] M. P. Wellman. A computational market model for distributed configuration design. In *Proceedings of AAAI '94*, pages 401–407. Morgan Kaufmann Publishers, San Francisco, CA, 1994. (Available from http://ai.eecs.umich.edu/people/wellman/Publications.html).

[8] M. P. Wellman. Market-oriented programming: Some early lessons. In S. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter 4. World Scientific, 1996. (More info about the book at http://www.wspc.co.uk/wspc/Books/compai_hilites.html#mbc The paper is available from http://ai.eecs.umich.edu/people/wellman/Publications.html).

[9] H. Yamaki, M. P. Wellman, and T. Ishida. A market-based approach to allocating QoS for multimedia applications. In M. Tokoro, editor, *Proceedings of the Second International Conference on Multi-Agent Systems ICMAS'96*, pages 385–392. AAAI Press, Menlo Park, CA, December 9–14 1996. (Available from http://ai.eecs.umich.edu/people/wellman/Publications.html).

[10] F. Ygge and J. M. Akkermans. Making a case for multi-agent systems. In M. Boman and W. Van de Velde, editors, *Proceedings of MAAMAW '97*, pages 156–176. Springer Verlag, Berlin, May 13–16 1997. ISBN-3-540-63077-5, (Available from http://www.enersearch.se/ygge).

[11] F. Ygge. *Market-Oriented Programming and its Application To Power Load Management*. PhD thesis, Department of Computer Science, Lund University, 1998. ISBN 91-628-3055-4, CODEN LUNFD6/(NFCS-1012)/1-224/(1998) (Available from http://www.enersearch.se/ygge).

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

[13] T. W. Sandholm and F. Ygge. On the gains and losses of speculation in equilibrium markets. In *Proceeding of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97*, August 23–29, 1997. (Available from http://www.enersearch.se/ygge).

[14] J. Holst. *Adaptive Prediction and Recursive Estimation*. PhD thesis, Department of Automatic Control, Lund Institute of Technology, 1977. LUTFD2/(TFRT-1012)/1-206/(1977).

[15] F. Avilla. On-line short term forecasting system for distribution lines based on synthetical load profiles. In *Proceedings of Distribution Automation/Demand Side Management (DA/DSM) '94*. PenWell Conferences and Exhibitions, September 1994.

[16] A. D. Kulkarni, T. Ananthapadmanabha, and K. R. Rao. Knowledge based expert load shedding. In *Proceedings of Distribution Automation/Demand Side Management (DA/DSM) '96*. PenWell Conferences and Exhibitions, October 1996.

[17] A. Andersson and F. Ygge. Managing large scale computational markets. In H. El-Rewini, editor, *Proceedings of the Software Technology track of the 31th Hawaiian International Conference on System Sciences (HICSS31)*, volume VII, pages 4 – 14. IEEE Computer Society, Los Alamos, January 1998. ISBN 0-8186-8251-5, ISSN 1060-3425, IEEE Catalog Number 98TB100216. (Available from http://www.enersearch.se/ygge).

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

[19] F. Ygge and J. M. Akkermans. On resource-oriented multi-commodity market computations. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems ICMAS'98*, pages 365 – 371. IEEE Computer Society, July 4–7 1998. (An extended version is available from http://www.enersearch.se/ygge).

[20] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipies in C*. Cambridge University Press, 1994. Second Edition.

[21] V. Pan. How can we speed up matrix multiplication? *SIAM Review*, 26:393–416, 1984.

[22] J. Cheng and M. P. Wellman. The WALRAS algorithm – a convergent distributed implementation of general equilibrium outcomes. In *Computational Economics*, 1998. To appear. (Available from http://ai.eecs.umich.edu/people/wellman).

[23] J. M. Akkermans, D. Healey, and H. Ottosson. The transmission of data over the electricity power lines. Advisory report for the EU Commission. (Available from http://www.enersearch.se), 1998.

[24] Z. Popovic. Scheduling of direct load control using mixed integer programming. In *Proceedings of Distribution Automation/Demand Side Management (DA/DSM) '97*. PenWell Conferences and Exhibitions, October 1997.

[25] F. Ygge and J. M. Akkermans. Duality in multi-commodity equilibrium computations and any-time algorithms. In C. Zhang and D. Lukose, editors, *Proceedings of the Third Australian Workshop on DAI*, pages pp. 65 – 78, Perth, Australia, December 1997. (Available from http://www.enersearch.se/ygge).

[26] W. N. Borst, J. M. Akkermans, and J. L. Top. Engineering ontologies. *International Journal of Human-Computer Studies*, 46:365–406, 1997. ISSN 1071-5819.