

Research Report 15/97



Power Load Management as a Multi-Commodity Market

by

Fredrik Ygge

Department of
Computer Science and Business Administration
University of Karlskrona/Ronneby
S-372 25 Ronneby
Sweden

ISSN 1103-1581
ISRN HK/R-RES—97/15—SE

Power Load Management as a Multi-Commodity Market

by Fredrik Ygge

ISSN 1103-1581

ISRN HK/R-RES—97/15—SE

Copyright © 1997 by Fredrik Ygge

All rights reserved

Printed by Psilander Grafiska, Karlskrona 1997

Power Load Management as a Multi-Commodity Market

Fredrik Ygge

EnerSearch AB and University of Karlskrona/Ronneby
Department of Computer Science (IDE)

372 25 Ronneby, Sweden

Fredrik.Ygge@ide.hk-r.se

<http://www.rby.hk-r.se/~fredriky>

Abstract

The current report investigates the application of market-oriented programming to power load management. Power load management (or load management for short) is the management of loads at the customer side in order to manage energy systems more efficiently. As energy systems are very large and heavily distributed (typically including millions of loads in an area covering counties or countries), efficient and conceptually attractive methods are required for making load management successful. This report demonstrates how market-oriented programming can be utilized to meet the above demands.

1 Introduction

1.1 IT and Energy

During the development of more and more powerful and inexpensive computers, and faster, more reliable, and less expensive communication, energy utilities has increased the use of computers in the management of energy systems. For a few years it has been economically feasible to install more or less advanced energy control systems in industries and at other large customers' premises. As time passing these systems are getting more and more advanced and energy management is becoming an integrated part of the business processes for many companies.

In the development of Information Technology, IT, in recent years, 1994 – 1997, the most remarkable progress is maybe the connection of individual households to the Internet. The development of new communication means to households is extremely intensive, and new possibilities of using cable-TV networks or power-line communication for connections to global networks are investigated. In all, the costs for communicating to individual households, and the difficulties associated with it, can be expected to decrease dramatically the coming few years. This development generates new questions for energy utilities, e.g., "How can the energy utilities take most advantage of the ability to communicate with virtually every load for a low cost?" or "What kind of new services can be offered to customers?".

Many energy markets are facing deregulation, or have already been deregulated, and being able to offer services, that other utilities can not offer, is vital for survival. In the present work we focus on one such service, load management.

1.2 Load Management - Basic Principles

Load management is the concept of managing loads¹ at the customer side in order to run the energy system more efficiently. The very basic principle is to try to move load from expensive to less expensive time periods. Time periods can be expensive for many different reasons, stemming from either production or distribution.

When discussing load management, two different categories are normally used [Flood and Hyde, 1992, Ottosson *et al.*, 1989, Sorri and Salo, 1992, Wafa, 1996, Talukdar and Gellings, 1987]:

- direct (or utility based), and
- indirect (customer based, or time-of-use pricing based) .

¹A load is any device that consumes energy, e.g. a radiator, a light bulb or an electric motor.

Direct load management is the case where the provider of electricity controls the loads subject to certain constraints, as specified in some contract with the customer. With indirect load management, the customer faces different prices at different times and is itself responsible for making the decisions when to buy. In this report we focus on direct load management, though most of the presentation here is relevant for the indirect case as well.

Power load management can be used for many different purposes in different time-scales. One can distinguish between at least two different time-scales:

- the very fast, and
- the moderate time-scale.

An example of an application in the first time-scale is the damping of oscillations, see, e.g., [Samuelsson, 1997]. This application requires response times in the millisecond region. An example of the second time-scale is the application of avoiding, e.g., overheated transformers as a consequence of too high a load for too long times. Another example is using load management to run production units as optimally as possible. Here the relevant time frame is minutes up to one hour or a few hours. The respective problems related to the two approaches are fundamentally different.

In the first case the main problems are related to real-time constraints. Oscillation frequencies have to be identified on-line, a strategy for how to damp these oscillation with available loads must be computed, and finally this strategy has to be executed with the particular loads. All this has to be performed in a few milliseconds. On the other hand, the discomfort for a customer is typically negligible. If, e.g., a stove is disconnected a few seconds every hour, no one will take notice. It seems as only direct load management is relevant for this time scale.

The second case is more or less the other way around. The real-time demands are typically not hard, and a computation time for an allocation of a few seconds up to a few minutes is often acceptable. The impact of the control at the customer side is on the other hand significant. Here it is important to consider the consequences of the control and the problem is finding an satisfying strategy for what loads to connect or disconnect over a longer period of time. This report deals solely with this second case.

In the research so far we have only been dealing with active² power. Extending the approach presented here to manage reactive power as well seems straight forward, but is part of future research. We also only discuss the energy form electricity, but we believe the concepts here to be highly relevant for other forms as well, such as district heating.

²Active and reactive power is related to the angle, ϕ , between the voltage, U , and the current, I . The active and the reactive power are defined as $UI \sin \phi$ and $UI \cos \phi$ respectively. Generally one can say that active power is the "good" power, useful for heating etc., whereas reactive power is an undesired bi-product, caused by capacitors and inductors.

1.3 The Benefits of Load Management

In the past, energy utilities have mainly traded energy, kWh, for money. This is however a very low transaction level; the commodity energy is in itself uninteresting for most customers. What a customer mainly wants is to get certain services at certain costs. Such services could be heating of a house, light, heating of water etc. As an example, the most interesting service for managers of a public building might be to pay a certain annual amount of money for a comfortable indoor temperature, and get some compensation whenever the indoor temperature deviates from the agreed on temperature. The provider of the energy is then responsible for using the heating system as optimal as possible from an energy point of view, e.g. perform direct load management. This implies that a utility being able to perform highly efficient load management can offer this kind of high level services at a lower cost than its competitors. From the customer point of view it can mean that it can have a lower overall cost, since a utility may have knowledge in energy management that the customer has not.

The above example stands in bright contrast to the traditional view of direct load management where a contract between the customer and the utility typically allowed the utility to disconnect specific customer loads for specific amounts of time, e.g., one hour per day. For this, the customer would get some compensation, either as a rebate or through some additional services. Still, this approach can be very profitable for the customer. For example, if a customer has the alternatives of either buying energy for its warm water heater at different prices at each hour, or buying for a fixed price, but allowing the utility to disconnect the heater for, say 4 hours each day, it may be the most profitable choice to choose the second alternative. The customer might lack the interest and/or the knowledge to install an indirect load management system itself, or the utility might be able to install a corresponding system at a lower cost, and hereby offer a better deal. A natural protest to this argument is that there might be consultants doing the installation of the indirect load management system, and hereby an indirect load management system is at least as interesting. However, when two self-interested parties, the utility and the customer, trade with each other, the enforcement of the true revelation of expectations of the future is not an easily managed issue. As will be shown in the examples below, the true revelation of expected future demands is an important feature of our approach. This can make direct load management systems more efficient than indirect systems are. Hereby, higher profits can be obtained, resulting in better deals for the customer.

In summary, direct load management can generate significant profits for both the customers and the utility, see also, e.g., [Dag, 1996, Bergström, 1996].

1.4 Energy Market Deregulation

Many energy markets are facing deregulation, and others are already deregulated. An interesting question is of course whether or not direct load management has any future in a

deregulating setting. We argue that this is indeed the case. The basic principle of shifting loads from expensive to less expensive time periods will be just as important. The only difference being that whether or not a time period is expensive will be determined by the market as a whole and not by the energy system of an isolated monopolistic utility. This argument will be further supported by Example 4 below (Section 6.3). On the other hand it is clear that if a deregulated market leads to very small fluctuations in prices over time, load management will be uninteresting. Whether or not prices tend to smooth out in a deregulated market is a complicated issue beyond the scope of this report, but we do not find any really good arguments for that this will be the case.

Another interesting perspective of a deregulated market is that an entirely new type of direct load management company might appear. If the traditional energy utilities do not become experts on load management, others will. There is a clear economic potential in buying electricity on a deregulated market at competitive prices and sell it to customers at a relatively low fixed price while at the same time being allowed to disconnect the loads for certain time periods each day or week. Also the service discussed above – providing a comfortable indoor climate for a certain fee – is an interesting application for such a company. The company can be much more skilled on acting on an energy market and managing the heating of the building than the customer is. Hereby the customer can save money, while the company is making a profit. The key to success for such a company is, of course, to be an expert on direct load management.

In the final version some description of the Swedish energy market will be included in this section.

1.5 The Aim of this Report

In this report we aim at demonstrating how market-oriented programming can be utilized for direct load management. The goal has been to develop an integrated approach that can be used for both low and high level customer services (from, e.g., the disconnection of a load for one hour a day to high level energy services, such as climate control). At the same time, the concepts should be well suited for both monopolistic utilities as well as for utilities acting on competitive markets.

1.6 Placing the Current Work

The work presented in this report has been performed as a sub-project (sub-project number 8) of the Information/Society/Energy/System project – the ISES project – of EnerSeach. The project is a multi-disciplinary project with Ph.D. students and professors from, e.g., computer science, business administration and electrical engineering. Participating industrial partners of the project are ABB Network Partner, Electricité de France, IBM Utility,

2 Applying Market-Oriented Programming to Load Management

Market-oriented programming has been successfully applied to a number of applications [Wellman, 1993, Wellman, 1994, Wellman, 1996, Yamaki *et al.*, 1996, Ygge and Akkermans, 1997, Kurose and Simha, 1989, Walsh *et al.*, 1997]. In this section we describe how market-oriented programming can be applied to the application direct load management, and what the most important design steps are.

2.1 Basic Market Design

The first important issue is getting setting up the commodities properly. In our research we started with considering two commodities, money and the active power for the upcoming time-period, e.g. hour, see [Ygge *et al.*, 1996, Akkermans *et al.*, 1996, Ygge and Akkermans, 1996, Akkermans and Ygge, 1997]. The need at this upcoming period is then based on *expectations* of future prices. Even though we believe this approach to be very useful, an obvious weak point is the expectations. Agents at the local level (further described below) might have a very good view of the electricity needs for the near future, and there is a clear potential advantage in taking this into account. Furthermore, if the controlled loads constitute a significant part of the total load, it is hard or impossible to separate the issues of prediction and load management; an integrated view of the two would be advantageous.

In [Ygge *et al.*, 1998] our model was extended to include future prices as well. The active power of different time slots represent the different commodities.⁴ The approach presented is not limited to any special division of time or number of commodities. For example, one could consider using 15 minute time slots the first hour, 30 minute time slots the next 2 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 finer grained time scale (larger number of commodities), the more accurate outcome, but also the heavier computational burden. It does seem very natural to add reactive power as additional commodities, but the feasibility of this has to be verified through further research.

For the prices and demands of the next few periods to be *exactly* correct a market with an

³For more information, refer to the web-site: <http://www.enersearch.se>.

⁴One can either use electric power or electric energy as the unit of each commodity, since one can always be obtained from the other at a given time interval. Here we consistently use electric power.

infinite number of commodities, covering time infinitely into the future, is required. This is of course practically impossible. Rather we set up a reasonable number of commodities, e.g., the 23 commodities described above, and let the agents base their decisions of these time periods on some rough estimations of future values, i.e. for the example here, rough estimates of the prices of the 49th hour and beyond. Two methods for such predictions that can be computed locally are:

1. assuming a periodic continuation, (i.e. with the commodities above let the prices of the 49th hour, the 97th hour and so on be the same as the prices during the first hour), and
2. assume the price of the last period covered by the market to be the price of every future time period.

Each agent should consider so many future time periods that the demands of the market does not change significantly by adding further future time periods. How to do this in practice is described in Section 3. One can also consider to send the agents globally available predictions of some future time periods not covered by the market.

Auctions (further 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 resource 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.

An example of a market outcome is given in Table 1. The market outcome provides the utility with a control strategy for the upcoming time period and useful predictions about future prices and demand in a compact and uniform manner.

2.2 Computational Agents Representing Loads – HOMEBOTS

Each load in the system is represented by a computational agent, a HOMEBOT. The responsibility of a HOMEBOT, illustrated in Figure 1, is to use electricity as efficient as possible given:

- the customer preferences,
- the load state,

Time Period:	1	2	3	4	5
Price:	1.23	1.43	1.65	1.12	1.01
Total production:	134	200	215	100	80
Total demand:	70	130	140	150	100
External demand:	64	70	75	-50	-20

Table 1: An example of a market outcome. The rows represent, from top to bottom, the time slots, the market price, the total production of the utility, the total demand of the utility’s customer, the external demand (i.e. what is traded on the external energy market). For time slots 1, 2, and 3, the utility sells resource to the external market and for time slots 3, 4, and 5, the utility buys resource from the external market.

- consumption predictions, and
- the load model.

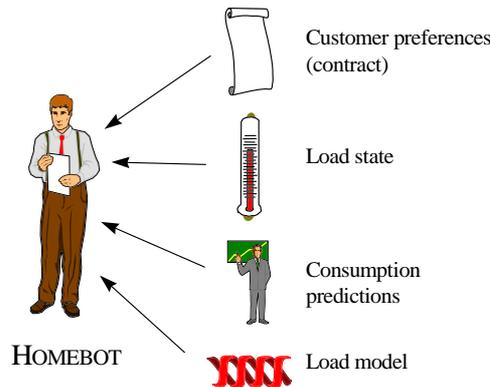


Figure 1: The important parameters of a HOMEBOT.

The *customer preferences* are represented by a customer contract. As mentioned above such a contract can establish, e.g., how long time the utility (i.e. the agent) might disconnect the load each day, or what indoor temperature the agent is responsible for maintaining, and what the cost for deviating from this temperature is.

For example, the indoor temperature of a building, or the amount of warm water in a warm water heater constitute the *load state*. Of course, a HOMEBOT responsible for an indoor temperature in a building has a much higher need for resource when the building is cold than when it holds a comfortable temperature.

The *consumption predictions* are preprogrammed or learned consumption patterns. If a HOMEBOT representing a warm water heater expects a customer to take a shower the coming hour, its action can be completely different from if it did not expect this.

Load characteristics, such as time constants, are contained in the *load model*. As an example, a HOMEBOT representing a load with a large time constant can avoid buying at high

prices if the high prices are merely temporary, since a temporary lack of resource will only affect the load state marginally.

A HOMEBOT is modeled as a *utility maximizing consumer* [Mas-Colell *et al.*, 1995, pp. 17–39, p. 314] holding a *utility function* [Mas-Colell *et al.*, 1995, Definition 1.B.2, p. 9], defined by

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

where $\mathbf{r} = [r_1, r_2, \dots, r_k]$ is the resource (active power) of the different time intervals, m (the "money"), corresponds to a real currency, such as SEK. f_i denotes the value (normalized by m) associated with a specific allocation, \mathbf{r} , i.e., captures the customer preferences, the load state, consumption predictions, and the load model for HOMEBOT α .

For every HOMEBOT there is a lower and upper limit for how much resource it can take. It can never take less than zero (i.e. it can not produce), and it can not take more than the power consumed when fully on. For certain HOMEBOTS there might be higher lower limits than zero, and lower upper limits than the load when fully on. See, e.g., the HOMEBOT representing the building in Section 3.

The task of a 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*) [Mas-Colell *et al.*, 1995, p. 20, p. 314], i.e. treat prices as exogenous rather than speculate about the effects of its own action on market prices [Sandholm and Ygge, 1997].

Also uncontrollable loads are represented by HOMEBOTS. In this case, however, the HOMEBOT is totally insensible to prices. Either a HOMEBOT can represent one or a number of loads. A HOMEBOT representing an uncontrolled load utilize a suitable method for prediction, see e.g. [Holst, 1977, Holst and Ekelund, 1987, Avilla, 1994, Kulkarni *et al.*, 1996].

2.3 Other Agent Types

Apart from the HOMEBOTS we use three other types of agents:

- production agents,
- distribution agents, and
- external demand agents.

A production agent represents one or more production units. Such an agent is modeled as *profit maximizing producer* [Mas-Colell *et al.*, 1995, pp. 135–137], i.e., agents solving the problem

$$\max_{\mathbf{r}} \mathbf{p} \cdot \mathbf{r} \Leftrightarrow c(\mathbf{r}), \quad (2)$$

where \mathbf{r} is a *produced* amount of resource, and $c(\mathbf{r})$ is the cost associated with producing \mathbf{r} .

The demand resulting from solving this problem then captures all constraints related to the production unit(s), e.g., current state and process model.

A distribution agent is very similar to a production agent, but here \mathbf{r} denotes the resource delivered to the HOMEBOOTS managed by the distribution agent.

External demand agents model the demand coming from consumers and producers outside of the utility's energy system. Thus, this agent type is used when the utility acts on a deregulated market. This agent is typically modeled directly by a demand, see further Section 6.3.

The three types will be demonstrated through examples below.

2.4 Agent Interaction and Structure

The agents interact through *auctions*. In an auction agents submit bids and the auctioneer computes a market equilibrium (i.e. a set of prices such that supply meets demand for each commodity). When market equilibrium has been obtained, the resource for the upcoming time interval is allocated.

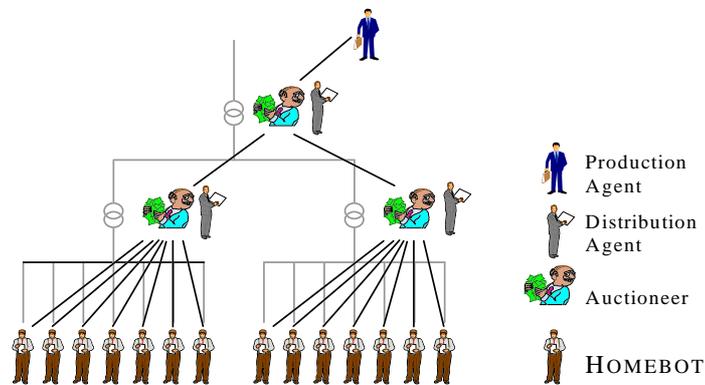


Figure 2: A typical structure of a load management system. The HOMEBOOTS are connected to auctioneers which, in turn, are connected to other auctioneers in a hierarchic manner.

The structure of the load management system is typically as in Figure 2. The power distribution system is inherently very hierarchic, and the structure of the load management system normally reflects this. Communication possibilities may vary significantly on different levels, and many different media and technologies can be utilized. For example, on the lowest level power line communication can be used. This communication is free of charge for power utilities, and hence it may be wise to perform auctions at this level more frequently than auctions on higher levels, *cf.* the discussions in [Andersson and Ygge, 1997].

An important remark here is that an agent is *not* equivalent to a computer. The structure shown in Figure 2 represents a communication structure between the agents, but it does not necessarily say that the first auctioneer should be placed in the secondary substation, nor that HOMEBOTS reside on hosts at the loads. The issue of where to place the agents is determined by the available hardware structure. For example, if the communication is very fast throughout the system and powerful micro-processors are available at every load, then it makes sense to place the HOMEBOTS at the loads for modularity reasons. If, on the other hand, the communication between the loads and the first auctioneer is rather slow and/or the computational power at the loads is very small, it makes sense to place all the HOMEBOTS of one auctioneer on the same host as the auctioneer and only communicate the required data for the HOMEBOT (see Figure 1) from the loads to the host. Another important issue when considering where the agents should be placed is the structure of the communication system. For example, the design of the computer system for energy management might be very different if cable-TV rather than power-line communication is used. However, the theory given here is essentially independent of how the agents are actually placed on different hosts. It should also be mentioned that it certainly is conceivable to have local production units, represented by production agents, among the HOMEBOTS at the lower levels.

2.5 Properties of the Market Outcome

From Theorem 2 in [Ygge *et al.*, 1998] we have that the market equilibrium in a market with the above agents is an optimal load management strategy in the sense that it is the solution to

$$\begin{aligned} \max \quad & \sum_{\alpha=1}^{n_h} f_{\alpha}(\mathbf{r}_{\alpha}) \Leftrightarrow \sum_{\alpha=1}^{n_{pd}} c(\mathbf{r}_{\alpha}) \\ \text{s.t.} \quad & \sum_{\alpha=1}^{n_h} \mathbf{r}_{\alpha} \Leftrightarrow \sum_{\alpha=1}^{n_{pd}} \mathbf{r}_{\alpha} = \mathbf{R}, \end{aligned} \quad (3)$$

where n_h is the number of HOMEBOTS representing controllable loads, n_{pd} is the number of production and distribution agents, and \mathbf{R} is the total amount consumed by the uncontrollable loads plus the amount sold to the external market (*or* minus the amount bought from the external market). How to determine what amount should be sold to the external market is described in Section 6.3.

Since production and distribution can be considered as instant, and there is an ordering of the commodities (which is $m \succ r_i$) such that the input for every producer precedes the output, we know that the market outcome is implementable, see [Ygge, 1998].

3 Example 1: A Critical Section

We now exemplify the above concepts on a particular example. This example is a classical load management example, and it is traditionally treated as a scheduling problem

(e.g. [Wafa, 1996, Popovic, 1997]). Here we demonstrate that the problem can be successfully managed with a market approach, leading to very high quality solution. The problem setting analyzed is depicted in Figure 3.

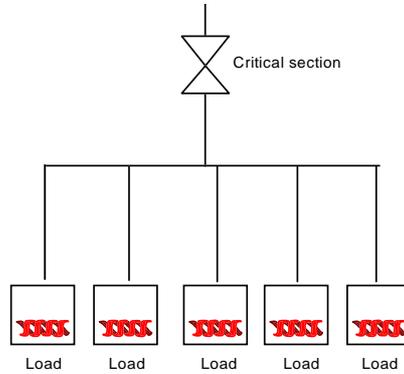


Figure 3: A critical section (bottle neck). A number of loads are delivered electricity through, e.g., a transformer, and at times the load is too high.

3.1 Load and Critical Section Characteristics

The loads of this example are warm water heaters. The investigated time-period can be thought of as morning hours where many people have taken morning showers and have gone to work. We assume that there is no consumption during the investigated time period. Thus, the heaters will at this point start to heat water until they are fully heated. The utility has a contract with each customer allowing the utility to switch off the loads for a certain amount of time during the 4 hour period. The interesting parameters for each load is how much it consumes when it is switched on, how much energy it is expected to need to be fully heated, and for how long it can be disconnected. The controllable loads of this example are described in Table 2.

Type	# Loads	Power (kW)	Need (kWh)	Time (h)
1	6	2	1	1
2	10	2	2	1
3	4	2	4	1
4	3	3	5	1/2
5	3	1	2	1

Table 2: The loads of example 1. The columns describe, from left to right, a type name, the number of loads of the current type, the power consumption when switched on, the current need, and the time that the load can be disconnected during the 4 hours.

In addition to the controllable loads there are some uncontrollable loads with the expected consumption 11, 10, 10 and 11 for each of the coming next 4 hours.

We assume the limitations of the bottleneck to be caused by too high temperatures. The cost of the bottleneck is assumed to be described by

$$cost(\mathbf{t}) = 10^{-4} \sum_{i=1}^k t_i^3, \quad (4)$$

(where $\mathbf{t} = [t_1, t_2, \dots, t_k]$, and t_i is the temperature, in $^{\circ}C$) i.e. for low and moderate temperatures the cost is very low, but for high temperatures the cost is significant as this can harm the bottleneck and possibly cause a breakdown. As numerical examples, if the temperature is $20^{\circ}C$ for every period the cost is 3.2SEK and if the temperature is $200^{\circ}C$, the cost is 3200SEK. The relation between the load and the temperature, at a sample interval of 1h, is assumed to be described by (cf. [Ygge and Akkermans, 1997])

$$\begin{aligned} t_i &= t_{i-1} \Leftrightarrow (t_i \Leftrightarrow 20) + 2 \cdot 10^{-2} r_i^2 \Leftrightarrow \\ t_i &= 0.5t_{i-1} + 10 + 10^{-2} r_i^2, \end{aligned} \quad (5)$$

where r_i is the total resource supplied through the bottleneck at time interval i . t_0 is assumed to be $90^{\circ}C$. The idea behind the relation between the resource and the temperature is that the heat generated in the bottleneck is $R \cdot I^2$, where R is the resistance of the bottleneck and I is the current through it, proportional to the power delivered.⁵

3.2 The HOMEBOTS and the Distribution Agent

We analyze a market of 4 commodities, representing 1 hour each. In this example each HOMEBOT will have a utility function as described by Eq. (1) with $f_i(\mathbf{r}) = 0$ when the consumption is such that the disconnection time, as defined in Table 2, is not violated, and $f_i(\mathbf{r}) = \Leftrightarrow C$, 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 very high cost. We assume every agent to have a perfect estimate of its future need. This is not required for the approach to work, but it simplifies the computations. If uncertainty about the need is present, the HOMEBOT should instead base its action on probability distributions of the need. Clearly, the more accurate estimates the HOMEBOT has, the better the outcome will be.

The agent will then act competitive on the computational market, with one small, but important, exception. The reason for the exception is best explained through examples. We observe a HOMEBOT having a load of type 2, and assume that $\mathbf{p} = [1, 1, 3, 1]$. Since the price of the third period is the highest, the ability to disconnect the load will not be used in this example and the demand, denoted $\mathbf{z}_\alpha(\mathbf{p})$, will be $\mathbf{z}_\alpha(\mathbf{p}) = [2, 2, 0, 0]$. Now, let $\mathbf{p} = [3, 1, 1, 1]$, then $\mathbf{z}_\alpha(\mathbf{p}) = [0, 2, 2, 0]$. Correspondingly, $\mathbf{p} = [1, 3, 1, 1]$ implies $\mathbf{z}_\alpha(\mathbf{p}) = [2, 0, 2, 0]$. But what about $\mathbf{p} = [2, 2, 1, 1]$? Now everything between

⁵This does not necessarily represent any real cable or transformer exactly. Rather the aim here is to give an example which has realistic characteristics – a rapidly rising cost for high temperatures, and high dependencies between the resource delivered at adjacent time periods – while still being quite simple.

$\mathbf{z}_\alpha = [2, 0, 2, 0]$ and $\mathbf{z}_\alpha = [0, 2, 2, 0]$ gives equally high utility and the agent is indifferent between the different allocations as long as it is disconnected for 1 hour during the first 2 hours. Furthermore, as soon as $p_1/p_2 \geq 1 + \delta$, where δ is an arbitrarily small positive constant, $z_1 = 0$. If many HOMEBOTS have this kind of shift it means that an equilibrium need not exist, neither for $p_1/p_2 \geq 1 + \delta$ nor for $p_1/p_2 \leq 1 - \delta$, regardless of how small δ is (not even for $\delta = 0$), even though the HOMEBOTS can be practically indifferent between the two prices. Therefore we let the shift from, e.g., $\mathbf{z}_\alpha = [2, 0, 2, 0]$ to $\mathbf{z}_\alpha = [0, 2, 2, 0]$ be linear for a small change in price. In our simulations we have used a $\delta = 10^{-6}$, let $z = [2, 0, 2, 0]$ for $p_1/p_2 \leq 1 - \delta$, $z = [0, 2, 2, 0]$ for $p_1/p_2 \geq 1 + \delta$, and let z_1 and z_2 decrease and increase linearly for $1 - \delta < p_1/p_2 < 1 + \delta$, as shown in Figure 4. Corresponding is done for all shifts, i.e. also for shifts related to p_3 . As will be shown below, this is a way to guarantee that equilibrium exists for many cases, without affecting the utility of the HOMEBOTS, or the total utility of the entire system, in any significant way.

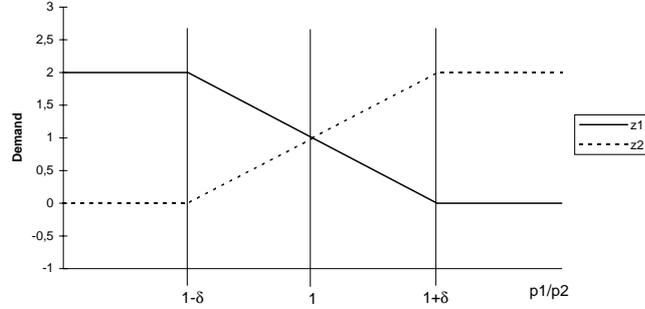


Figure 4: Illustration of the linear shift of demand at time 1 and 2 for a load of type 3. p_3 is here assumed to be (significantly) larger than p_1 and p_2 . In our implementations, $\delta = 10^{-6}$ has been used.

Hence, the utility function of the HOMEBOTS representing the controllable loads is

$$u_\alpha(\mathbf{r}, m) = f_\alpha(\mathbf{r}) + m, \quad (6)$$

and the agent will act competitively with the additional feature that there is a linear demand shift when prices are practically equal. As all the need is always satisfied within the 4 hours covered by the market, regardless of how the load contracts are utilized, future prices do not affect current demands. Thus, no predictions of price for time periods 5, 6, and so on, are required.

The uncontrollable loads are represented by a HOMEBOT with the demands 11, 10, 10 and 11 for the respective time intervals, independently of prices.

From Eq. (4) and Eq. (5), we see that the cost of the distribution can be expressed as a function of t_0 and \mathbf{r} . Hence, the *distribution agent* solves the optimization problem

$$\max \mathbf{p} \cdot \mathbf{r} \Leftrightarrow c(t_0, \mathbf{r}). \quad (7)$$

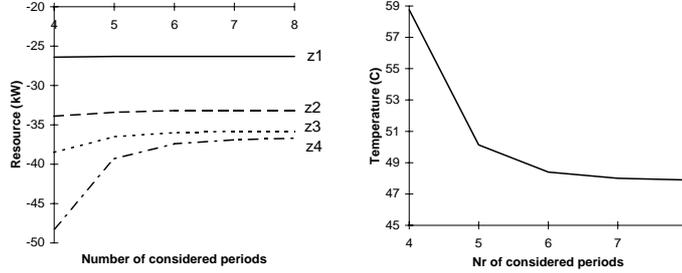


Figure 5: Leftmost, the effect on demand (of the four upcoming time periods) of different number of consider future periods at $p_1 = p_2 = p_3 = p_4 = 1$. The price of the future periods is here assumed to be 1. Rightmost, the temperature at the end of time period 4 is plotted.

As the demand of the time periods covered by the market are affected by future prices we will investigate the two approaches of future predictions that can be computed locally, as where described in Section 2.1. First we take a look at four demands as a function of the number of considered future periods. The result of this when assuming that p_4 will prevail as future price is shown in Figure 5. Together with the demands, the temperature at the end of the 4 hours is illustrated. As seen from the figure, considering 8 instead if 7 periods does not affect the demands or the temperatures significantly. Therefore, we are satisfied with using 8 periods. In the experiments we will consider both the case where p_4 is considered to be the prevailing price and the case where the price is assumed to be periodic, i.e. the case when $p_5 = p_6 = p_7 = p_8 = p_4$ and the case when $p_5 = p_1$, $p_6 = p_2$, $p_7 = p_3$, and $p_8 = p_4$.

3.3 The Market Outcome

The above agents were implemented together with an auctioneer using a price-oriented Newton scheme (as described in [Ygge, 1998]). The results are given below.

Figure 6 illustrates the resource usage in the different time periods, for the cases where no load control is used and load control with the two different predictions of the distribution agent. We see that load is shifted from time period 1 to the periods 2 and 3.

The market price of the different time periods is shown in Figure 7. As seen from the figure, the number of controllable is so large that the price of the first two periods is practically equal. Hence, there is no incentive for performing any shift of load between time periods one and two, cf. the discussion in Section 6.4.

The corresponding temperature of the bottleneck is plotted in Figure 8.

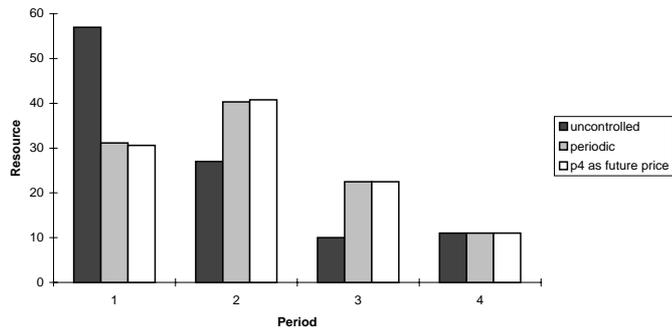


Figure 6: A comparison between the resource usage in the situations of a) no load control, b) load control where the distribution agent assumes the last price to prevail, and c) load control where the distribution agent assumes a periodic future price.

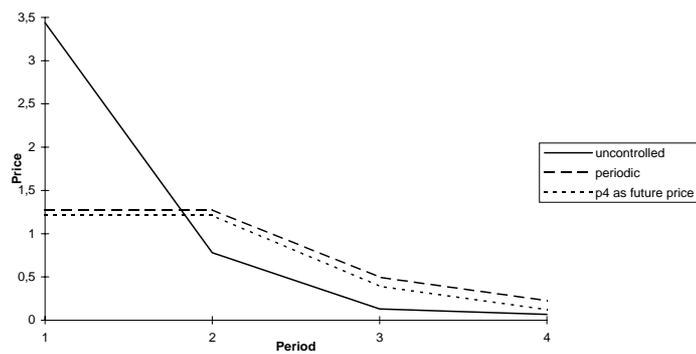


Figure 7: The market prices (the marginal cost of distribution) of the different time periods. In this example the ability to move demand from period one to period two is so great that no further shift is beneficial.

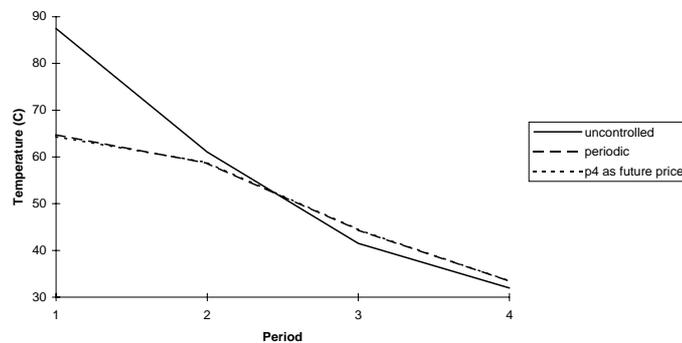


Figure 8: Illustration of the temperature of the bottleneck for the different settings.

The cost of the bottleneck (as given by Eq. (4) and Eq. (5)) without load control is 100.13. With load control and periodic prediction this cost is reduced to 59.71, and with p_4 as future price the cost is 59.53. We see that the improvement is significant.⁶ The existence of an equilibrium of this example was not a coincidence. As all demands, $z_\alpha(\mathbf{p})$, are continuous, $\partial z_{i,\alpha}(\mathbf{p})/\partial p_i \leq 0$, and $\partial z_{i,\alpha}(\mathbf{p})/\partial p_j \geq 0$ for $i \neq j$, the general equilibrium is guaranteed to exist (see, e.g., [Cheng and Wellman, 1997]).

The allocations of the controllable loads for the different time periods are given in Table 3. The results are from the simulation when p_4 is assumed to prevail as the future price.

Type	r_1 (kW)	r_2 (kW)	r_3 (kW)	r_4 (kW)
1	0.363	0.637	0	0
2	0.726	1.274	0	0
3	0.726	1.274	2	0
4	2.04	2.46	0.5	0
5	0.363	0.637	1	0

Table 3: *The allocations of example 1.*

4 Scheduling

In the previous section an allocation was computed (Table 3), and this allocation is known to be optimal in that it is a solution to Eq. (3)⁷, but nothing was said about when to actually switch on and off the loads. This issue is examined in this section.

4.1 The Problem of Example 1

All loads of type one were assigned the allocations 0.363 for the first time period. This implies that these loads should be switched off for 49.11 minutes during the first hour, but there is no guidance for which of the 60 minutes of the hour to use.

We use a uniform approach to when to switch on and off loads. When to do the switching is determined from to contradictory criteria:

⁶Even though the cost of the four periods is a highly relevant measure, it is not the entire truth. If the distribution agent only considers the four periods, without taking further periods into account, the cost would be lower. However, this would be sacrificing the future for the present. The presented result are optimal when taking the future into account as described by the predictions.

⁷Strictly speaking, the *true* market equilibrium is the optimal solution. As the utilized algorithm only computes an approximation of this equilibrium, we only obtain an approximate solution. However, the sum of the absolute value of the excess demand for every commodity was in the order of 10^{-7} in the simulation, so the approximation is indeed very accurate.

1. As the actual resource should at all times be as close as possible to the allocated resource, the switching should be made as often as possible. A too large deviation might, e.g., result in an unacceptable deviation of indoor temperature.
2. As there is a cost for switching the load on and off (wearing out relays etc.), the switching should be made as seldom as possible.

That is, too long a switching period results in a too large deviation from the allocated resource, and too short a switching period is too expensive. A suitable compromise must be established for each load. We refer to this compromise as the *optimal switching period*. For the loads of example 1 we assigned random optimal switching periods between 5 and 15 minutes. As an example, if a load of type one has an optimal switching period of 6 minutes, it will be on for 1.095 minutes and off for 4.905 minutes in a repetitive fashion. The result of running all loads in accordance with their optimal switching periods is illustrated in Figure 9.

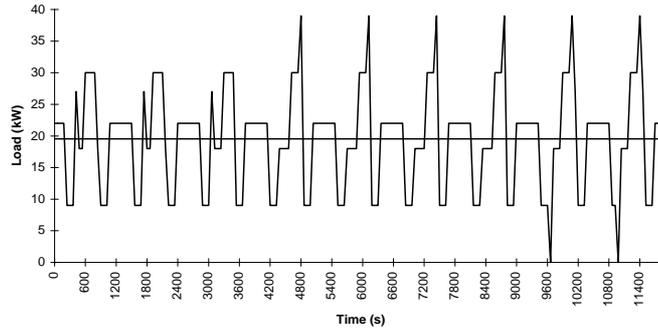


Figure 9: *The total load when the individual loads are unscheduled. The differences between the peaks and the lowest loads are significant.*

The result illustrated in Figure 9 is not very satisfying. Even though the total average is 19.55 over time, there are significant peaks. Since the heat of the bottleneck is described by Eq. (5), the generated heat can not be computed from the average load, but from an equivalent load defined by

$$r_{eq} = \sqrt{\frac{\sum_{i=1}^k r_i^2}{k}}. \quad (8)$$

That is, the heating effect resulting from the optimal switching periods, as described in Figure 9, is equivalent of a constant load of 21.15, which is notably higher than 19.55. Hence, even though we were able to compute the optimal allocation above, the outcome was not optimal because of the too high variation in load. For smoothing out the variations, the LOADSCHED algorithm⁸ is introduced.

⁸LOADSCHED has been developed by Dr. Arne Andersson at Lund University and Dr. Eric Schenk in cooperation with the author.

4.2 The LOADSCHED Algorithm

LOADSCHED is an on-line algorithm which considers discrete *scheduling intervals*, typically (but not necessarily) one minute. At the start of one interval, the expected consequences (evaluated for the end of the interval) of different actions are compared, and the least expensive one is chosen.

Two types of penalty functions, local and global, are used by LOADSCHED. On the local level the penalty of each load is a deviation penalty plus a possible switching cost. The deviation penalty is computed from

$$pd_{\alpha}^i(s^{i+1}) = (r_{ac,\alpha}^{i+1}(s^{i+1}) \Leftrightarrow r_{al,\alpha})^2, \quad (9)$$

where α denotes an agent, i (and $i+1$) denotes a time interval, pd_{α}^i is the deviation penalty, s^i is the load state (*on* or *off*), $r_{ac,\alpha}^i$ is the actual resource, and $r_{al,\alpha}$ is the allocated resource. The actual resource of the next time interval is

$$r_{ac,\alpha}^{i+1}(s^{i+1}) = \begin{cases} r_{ac,\alpha}^i + (l_{\alpha} \Leftrightarrow r_{al})t, & s^{i+1} = on \\ r_{ac,\alpha}^i \Leftrightarrow r_{al} \cdot t, & s^{i+1} = off \end{cases}. \quad (10)$$

where l_{α} is the power consumption of load α when switched on, and t is the scheduling interval.

If a load is currently off, the penalty, p_{α}^i , for keeping it off is computed from Eq. (9) and Eq. (10), and the penalty for switching it on is computed as the deviation penalty of Eq. (9) and Eq. (10) plus the *switching cost*. Similarly, if a load is currently on, the penalty for keeping it on is computed from Eq. (9) and Eq. (10), and the penalty for switching it off is computed as the deviation penalty of Eq. (9) and Eq. (10) plus the switching cost. (The switching cost, $c_{sw,\alpha}$ is computed from the optimal switching period and the deviation penalty of Eq. (9) so that the penalty over time is minimized by using the optimal switching period.) Hence, the penalty of a load can be expressed as

$$p_{\alpha}^i(s^i, s^{i+1}) = pd_{\alpha}(s^{i+1}) + (s^i \oplus s^{i+1})c_{sw,\alpha}, \quad (11)$$

where \oplus is an *exclusive or* operator, returning 1 if $s^i \neq s^{i+1}$, and 0 otherwise.

The global penalty, p_g^i , is computed as

$$p_g^i = (r_{ac,tot}^{i+1} \Leftrightarrow r_{al,tot})^2, \quad (12)$$

where $r_{ac,tot}^i$ is the total actual load at time i and $r_{al,tot}$ is the total allocated load.

The total penalty is computed from

$$p_{tot}^i(s^i, s^{i+1}) = \sum_{\alpha=1}^n p_{\alpha}^i(s_{\alpha}^i, s_{\alpha}^{i+1}) + C \cdot p_g^i. \quad (13)$$

where $\mathbf{s}^i = [s_1^i, s_2^i, \dots, s_n^i]$, n is the number of loads, and C is a constant.

It is also possible to set a total maximum, $r_{tot,max}$, that can never be exceeded.

The main steps of LOADSCHED are

1. Create a list, L , of all loads, sorted by increasing $p_\alpha^i(s_\alpha^i, on)$.
2. Set $s_\alpha^{i+1} = off, \forall \alpha$.
3. Let β be the first load in L .
4. Let \mathbf{v} be a vector defined by

$$v_\alpha = \begin{cases} s_\alpha^{i+1}, & \alpha \neq \beta \\ on, & \alpha = \beta \end{cases} . \quad (14)$$

5. If $(p_{tot}^i(\mathbf{s}^i, \mathbf{v}) < p_{tot}^i(\mathbf{s}^i, \mathbf{s}^{i+1})) \wedge (r_{tot}^{i+1} \leq r_{tot,max})$, then
 - (a) $\mathbf{s}^{i+1} = \mathbf{v}$.
 - (b) If β is not the last element of L , let β be the next element in L , and continue from 4.
6. Randomly select two loads, α and β , such that $s_\alpha^{i+1} = on$ and $s_\beta^{i+1} = off$ and set $s_\alpha^{i+1} = off$ and $s_\beta^{i+1} = on$. If p_{tot} is lower after the swap and $r_{tot,max}$ is not violated, keep the two new states, otherwise restore the old states. Repeat this m times.
7. Communicate the state change to all loads, α , for which $s_\alpha^{i+1} \neq s_\alpha^i$.

From empirical studies we have found that using a $C = 10$, an $m = 20$ works well for very general sets of configurations of loads. When selecting what loads to swap, we have found that the best result is obtained when the probability for selecting a load at a certain position in L is described by a normal distribution around the load in L which is the last load, α , for such that $s_\alpha^{i+1} = on$.

4.3 The Solution to Example 1

The result of applying LOADSCHED to Example 1 is plotted in Figure 10. The total allocated resource is 19.55. It is clear that the scheduler works extremely well for this example. The minimum load is 19 and the maximum load is 20. Since the smallest load is 1 we can not hope to obtain a smaller deviation.

In Figure 11, the equivalent resource, as computed from Eq. (8) is plotted. That is, the value at, e.g., 100 time periods is the value up to that point. In the case where the loads are

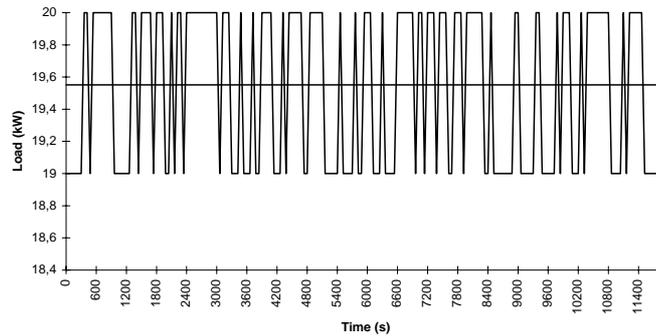


Figure 10: The total load when the individual loads are scheduled. As the smallest load is 1, this schedule is optimal in terms of the deviation from the allocated total resource.

scheduled, the expected load is practically equal to the allocated load. In the case where they are unscheduled, there is a noticeable difference. At the end of the 200 time periods, the equivalent load in the scheduled case is 19.53. This is of course only a temporary minimum. In the long run, the equivalent load must always be as large as the allocated load.

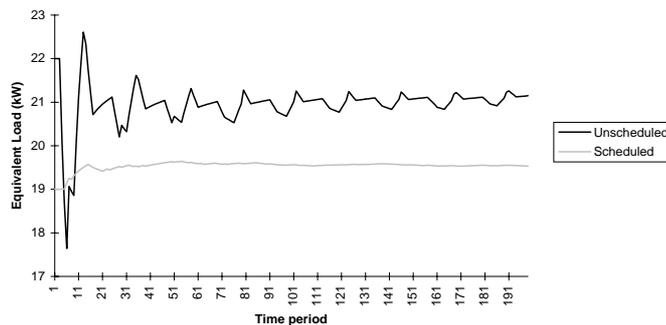


Figure 11: The equivalent load for the unscheduled and scheduled case. As the transients level out, the equivalent load in the unscheduled case ends up at approximately 21. In the scheduled case it ends up very close to 19.55, i.e. the allocated resource.

As demonstrated above, the scheduling of this example worked extremely well from the global perspective. The other interesting measure is how well it works at the local level. How much are the loads forced to deviate from the optimal scheduling period? One example of a schedule for an individual load is shown in Figure 12.

The result in Figure 12 is very typical; only a very small deviation from the optimal switching period is required in order to obtain a good global solution. In the long run, LOAD-SCHED guarantees that every load will have *exactly* what has been allocated.

One additional remark is probably appropriate for the scheduling of Example 1. When

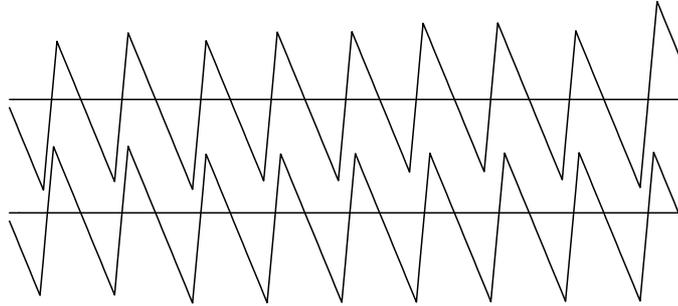


Figure 12: A typical effect of the scheduling of an individual load. The bottom curves are the resource resulting from the optimal switching period compared to the allocated resource. The top curves are the resource resulting from the scheduling compared to the allocated resource. The difference between the two is very marginal.

looking at the allocation of the loads of type 1, we had that the allocation of the first two time periods were 0.363 and 0.637 respectively. As the load of a heater when switched on is 2 for this type, it seems that the allocations above leads to that the loads must be switched off for 49.11 minutes during the first time period and 40.89 minutes during the second time period. This is more than the contracted one hour disconnection time, and it seems as if the HOMEBOT responsible for this load has violated the contract. Clearly, the load must be controlled 49.11 minutes during the first hour in order not to consume more than 0.363, but during the second hour the load will be fully charged with warm water once 0.637 has been allocated. The practical implication of this is that during the second hour, the load can only be forced to be switched off during 10.89 minutes, and after that it has to be considered as an unschedulable load until it has been fully charged. As the LOADSCHED algorithm is an on-line algorithm it will only be able to schedule the load as any other load as long as the load's 10.89 minutes are not consumed, and this might lead to a deviation from the theoretically optimal schedule. However, this problem can be made arbitrarily small by using a more fine grained time scale, i.e. use a larger number of commodities. This is a general truth; if the granularity of the market computations is such that the time interval is 1min, scheduling with a scheduling interval of 1min is even completely superfluous.

4.4 Discussion

Above we demonstrated LOADSCHED on a specific example. From many perspectives this example is quite simple. For example, the sizes of the loads and the optimal switching periods are of the same order of magnitude.⁹ Therefore, too many conclusions can not

⁹It can even be debated whether or not the unscheduled result is really *that* bad. Compared to uncertainties in estimates, measurements, etc., the deviation of the equivalent load from the allocated load (approx-

be drawn about the scheduling for more general examples. However, we have also tried LOADSCHED on a number of significantly harder problems, e.g. when one of the loads is substantially larger than the others. We have also tried examples where the maximum allowed total load ($r_{tot,max}$ above) has only been slightly larger than the allocated resource. So far, the results have indeed been positive.

An advantageous property is LOADSCHED is its efficiency. The total computation time required for the scheduling over the 200 time intervals demonstrated above is approximately 1s.¹⁰ This means that during the search for equilibrium it could be possible to consider the cost at the bottleneck based on the equivalent load rather than the allocated load.

Our conclusion at this point is that LOADSCHED works very well for our purposes and that for a wide class of loads, such as warm water heaters and radiators, the scheduling problem does not have to be considered when performing the market computations. Rather, it can be considered as an on-line problem separate from the market. See also the discussion in Section 7.2.

5 Example 2: Production Optimization

We now study a second example. In this example a new load type, representing a heating system of a customer having a high level contract, is introduced. Compared to Example 1, we this time study the use of a load management system from a production point of view rather than from a distribution point of view. The setting is illustrated in Figure 13.

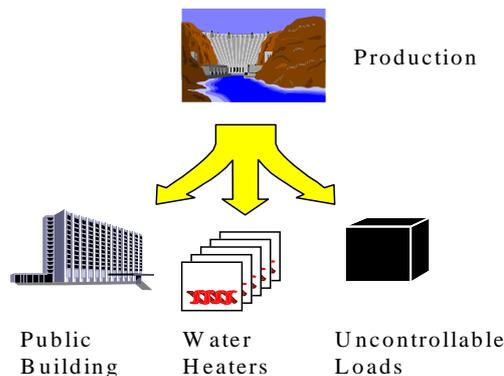


Figure 13: A small part of an energy system – a producer and its customers.

imetal 7%) might even be acceptable.

¹⁰LOADSCHED was implemented in Java and run on a standard PC (a 90MHz Pentium processor with 24MB of RAM).

5.1 Load and Production Characteristics

Three loads of each of the types 3, 4, and 5 above are used. A new load type is also introduced. It is assumed that the utility has a contract with a customer making the utility responsible for the indoor temperature of a public building which temperature is a function of the resource as described by (sample interval is 1h)

$$\begin{aligned} t_i &= t_{i-1} \Leftrightarrow 0.1(t_i \Leftrightarrow 10) + 10^{-2}r_i \Leftrightarrow \\ t_i &= \frac{1}{1.1}(t_{i-1} + 1 + 10^{-2}r_i). \end{aligned} \quad (15)$$

t_0 is assumed to be $19^\circ C$.

The heating system includes radiators of a total power of 300kW. There is a lower constraint for how little resource the building can be assign of 10, at the investigated time. (This lower constraint is often included in load management contract in order to avoid unpleasant draft from windows etc.) Thus, $10 \leq r_i \leq 300, \forall i$, for the building.

There is a contract between the utility and the customer stating that the utility is responsible for maintaining the indoor temperature $20^\circ C$ for a certain annual fee. The utility must pay the customer a compensation for deviation from the agreed on temperature of

$$c_{cust}(t_i) = 10 \cdot (t_i \Leftrightarrow 20)^2, \quad (16)$$

for every sample t_i at a 1 hour sample period.

In addition to the controllable loads there are some uncontrollable loads with the expected consumption 10, 400, 100 and 80 for each of the coming next 4 hours.

The production cost is assumed to be

$$c_{prod}(\mathbf{r}) = 10^{-3} \sum_{i=1}^k r_i^2 \quad (17)$$

5.2 The HOMEBOTS and the Production Agent

We analyze a market of 4 commodities, representing 1 hour each. The HOMEBOTS for the loads of type 3, 4, and 5 are exactly the same as the ones in Example 1.

The uncontrollable loads are represented by a HOMEBOT with the demands 10, 400, 100 and 80 for the respective time intervals, independently of prices.

Similarly to the distribution agent in Example 1, the building agent considers 8 time periods, and assumes p_4 to be the prevailing price after the time periods covered by the market. From Eq. (15) and Eq. (16), we see that the cost can be expressed as a function of t_0 and \mathbf{r} . Hence, the utility function of the HOMEBOT responsible for the building is

$$u_b(\mathbf{r}, m) = \Leftrightarrow \sum_{i=1}^k c_{cust}(t_0, \mathbf{r}) + m. \quad (18)$$

The maximization problem of the *production agent* is (cf. Eq. (17))

$$\max \mathbf{p} \cdot \mathbf{r} \Leftrightarrow c_{prod}(\mathbf{r}). \quad (19)$$

5.3 The Market Outcome

The above agents were implemented together with an auctioneer using a price-oriented Newton scheme.

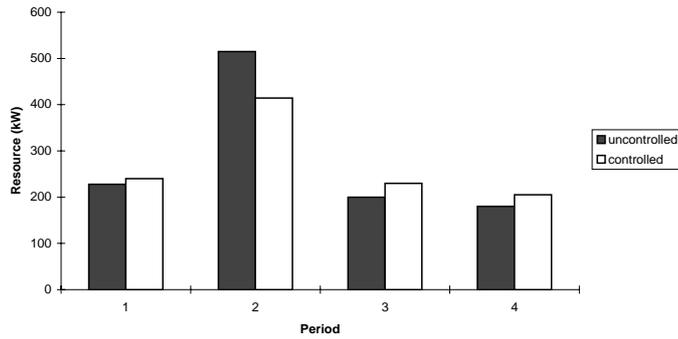


Figure 14: *The used resource in different time periods with and without load control. The result of the load control is that load is moved from period two to the other periods.*

In Figure 14 the allocations in presence and absence of load management are compared.¹¹ We see that the result is a shift of load from time period 2 to the other periods. The corresponding market prices are shown in Figure 15. The market price of time period 2 has been reduced significantly. However, it is still significantly higher than the prices of time period 1, 3, and 4, and it could be profitable to arrange load management contracts with more customers in order to shift more load from time period 2 to periods 1, 3, and 4, cf. Section 6.4.

The allocations of the different loads is given in Table 4.

As seen from Table 4, as little resource as possible is used for time period 2 by all involved controllable loads. (Note that the lower resource limit of the building is 10.)

The temperature of the public building is plotted in Figure 16. The most interesting thing of this figure is maybe that the HOMEBOT of the building has found that the optimal choice is to buy so much resource during time period 1 so that the temperature temporary is *above* the set point. This is really not surprising, and this control scheme has been used in previous approaches to load management, e.g., [Larsson *et al.*, 1988, pp. 33 – 35].

¹¹When load management is not present, the building is assumed to have a control system that maintains 20°C. The corresponding resources for doing this is (see Eq. (15)) 200, 100, 100, and 100.

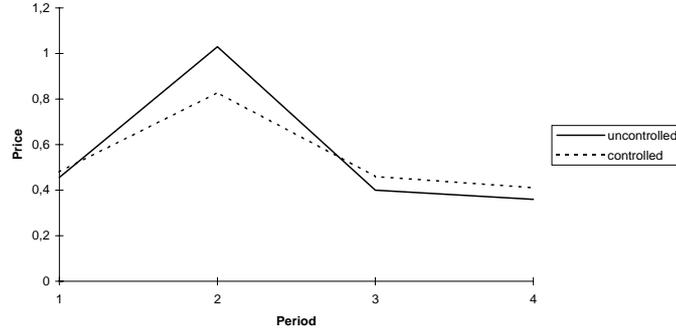


Figure 15: The market prices (the marginal production cost) of the different time periods. The market price of time period two has been reduced significantly.

Type	Uncontrolled				Controlled			
	r_1 (kW)	r_2 (kW)	r_3 (kW)	r_4 (kW)	r_1 (kW)	r_2 (kW)	r_3 (kW)	r_4 (kW)
1	2	2	0	0	2	0	2	0
2	3	3	0	0	3	1.5	1.5	0
3	1	1	0	0	1	0	1	0
Building	200	100	100	100	211.8	10	119.5	125.2
Uncontrollable	10	400	100	80	10	400	100	80
Total	228	515	200	180	239.8	414.5	229.99	205.2

Table 4: The allocations of Example 2. Note that the lower resource limit of the building is 10 and that loads of type 2 can only be controlled for 30min. When considering the total resource, keep in mind that there are three loads of each type 1, 2, and 3.

In [Larsson *et al.*, 1988], however, the scheme was hard-coded for certain hours each day, and in the present work it was the optimal choice in a certain situation; at another occasion it might be completely different.

As with Example 1, the general equilibrium for this example is known to exist. The total cost without load management is 389.61 and the cost with load management is 332.33 (of which 8.01 is compensation to the customer because of temperature deviations).

6 Utilizing the Load Management System

Above we gave two examples of the use of the load management system for two specific examples, related to distribution and production respectively. This section describes, in somewhat more general terms, how the load management system can be in different settings. Furthermore, the important issue of cost-benefit analysis is discussed.

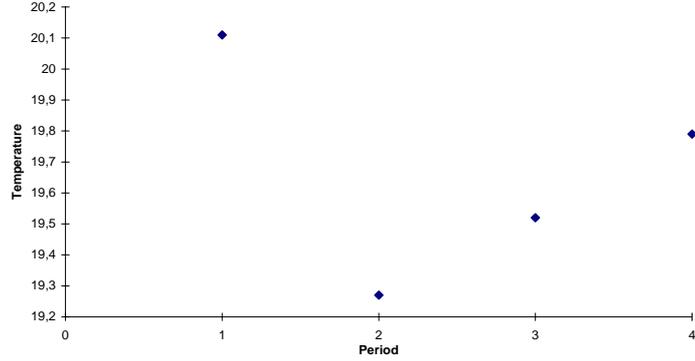


Figure 16: *The temperature of the public building. In order not to have a too big deviation from the agreed on temperature during period 2, when the resource is very scarce, the temperature is set above the set point in period 1.*

6.1 Multi-Level Distribution Agents

In Example 1 above, there was only one distribution agent. A somewhat more complicated distribution problem is shown in Figure 17.

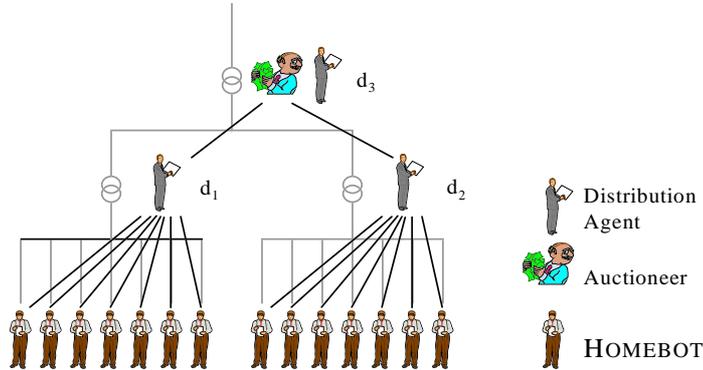


Figure 17: *Distribution agents on multiple levels.*

In this case the distribution agents, d_1 and d_2 , and the HOMEBOTS supported by the respective distribution agents, form non-separable units. Hence, a distribution agent has two different roles when trading on agents on higher levels compared to when it is trading on lower levels. When trading with agents on higher levels, the distribution agent will be seen as a utility maximizing consumer with the utility function

$$u_{d_i}(\mathbf{r}, m) = \epsilon_{d_i}(\mathbf{r}) + g(D_i, \mathbf{r}) + m, \quad (20)$$

where $c_{d_i}(\mathbf{r})$ is the cost of distributing the resource \mathbf{r} , D_i is the set of all agents supplied by d_i , and $g(D_i, \mathbf{r})$ is the summed utility of all agents in D_i , when \mathbf{r} has been

distributed among these agents so that $\sum_{\alpha \in D_i} u_{\alpha}(\mathbf{r})$ is maximized. From Theorem 2 in [Ygge *et al.*, 1998] we have that the optimal allocation of \mathbf{r} among the agents in D_i can be computed from the market equilibrium in a market consisting of the agents in D_i , plus an agent with the constant demand $\Leftarrow \mathbf{r}$.

For the rest, the market equilibrium is computed from the demand/price relations of d_3 , d_1 , and d_2 , just as it was computed in Example 1. In the case of more than two levels, the procedure is repeated recursively, i.e. $u_{d_3}(\mathbf{r})$ can be computed from $u_{d_1}(\mathbf{r})$, $u_{d_2}(\mathbf{r})$, and $c_{d_3}(\mathbf{r})$, just like $u_{d_1}(\mathbf{r})$ and $u_{d_2}(\mathbf{r})$ were computed from $c_{d_1}(\mathbf{r})$, $c_{d_2}(\mathbf{r})$ and the utility functions of the agents supplied by them.

For a two-commodity implementation, we believe the COTREE algorithm (see [Andersson and Ygge, 1998]) to be very well suited for the aggregation of utility functions described here. For the general, multi-commodity, case, finding efficient algorithms is part of future research.

6.2 A Monopolistic Utility

In the case of a monopolistic utility, acting both as a distributor and producer, the situation is very similar to the situation of multi-level distributors, as described above. The only difference is that production agents are included at different levels. The production agents can be integrated into the system just as any HOMEBOT. This can, e.g., be as in Figure 18.

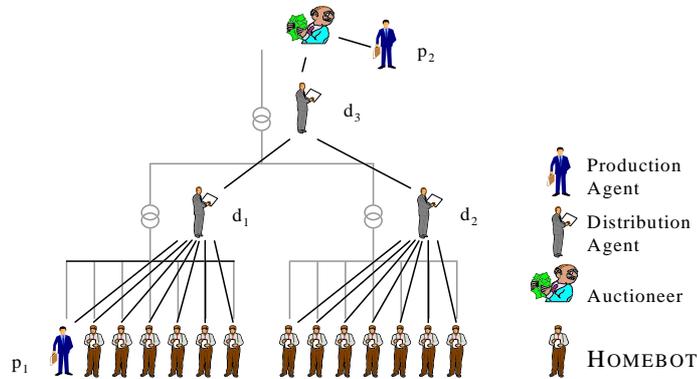


Figure 18: An example of a system configuration of a monopolistic utility.

In Figure 18 the utility function of the producer agent pa_1 is aggregated into $u_{d_1}(\mathbf{r})$ together with all the utility functions of the HOMEBOTS supplied by d_1 . $u_{d_3}(\mathbf{r})$ is computed as described in Section 6.1. Finally, the market equilibrium is computed from the competitive demand of d_3 and pa_2 .

6.3 A Utility on a Deregulated Market

The case when a producing utility is acting on a deregulated market is as easily managed as Example 2. In this section we give another example in order to describe the interaction with a market in some detail.

6.3.1 Load, Production and Market Characteristics

We use the same load and production characteristics as the ones used in Example 2, but add a market agent, representing the market at which the utility acts.

We assume the utility to have a perfect estimate of the market (though this is not required for the approach to work, *cf.* Section 3), and we assume the market demand to be

$$\mathbf{z}_{market}(\mathbf{p}) = \left[\frac{200}{p_1} \Leftrightarrow 10^3 \cdot p_1, \frac{2000}{p_2} \Leftrightarrow 10^3 \cdot p_2, \frac{1000}{p_3} \Leftrightarrow 10^3 \cdot p_3, \frac{200}{p_4} \Leftrightarrow 10^3 \cdot p_4 \right] \quad (21)$$

6.3.2 The HOMEBOTS, the Production Agent, and the Market Agent

The HOMEBOTS and the production agent are as described in Section 3. The market agent has the demand as described by Eq. (21).

6.3.3 The Market Outcome

The above agents were implemented together with an auctioneer using a price-oriented Newton scheme.

In Table 5 the most important data in presence and absence of load management are compared.

	Uncontrolled				Controlled			
Time Period:	1	2	3	4	1	2	3	4
Price:	0.4490	1.3404	0.8864	0.4300	0.4732	1.3072	0.8589	0.4645
Production:	224.5	670.2	443.2	215	237.15	653.6	429.45	232.35
Market demand:	-3.5657	151.692	241.759	25.1163	-52.626	222.79	305.38	-33.93
Building demand:	200	100	100	100	258.96	10	10	186.33

Table 5: *The most important data of Example 3.*

The total cost for the utility in the uncontrolled case is 311.10. That is, the cost has been reduced significantly compared to the case without the market (Example 2 in Section 5 above). It is even the case that the cost in the case *without* load management with the market is *lower* than the cost of the case *with* load management without the market. Therefore

it is tempting to believe that load management is not as important for a utility acting on a deregulated market as it is for a monopolistic utility. However, the cost in the case with load management *and* the market is 225.97, of which 17.03 is compensation to the customer of the public building. Thus, the gain of the load management system in this case is 85.13 compared to 57.20 in the case without the market. This is, of course, not a general truth. In other cases the presence of a market may reduce the need of a load management system. Generally, one can say that the less variation of prices on the market, the less need for load management.

Somewhat simplified we can say that the cause of the gain in this example was that resource could be bought from the market at inexpensive hours and sold at expensive hours. The building was "preloaded" with resource to a larger extent than in the case without the market. The resulting indoor temperature of the public building is plotted in Figure 19

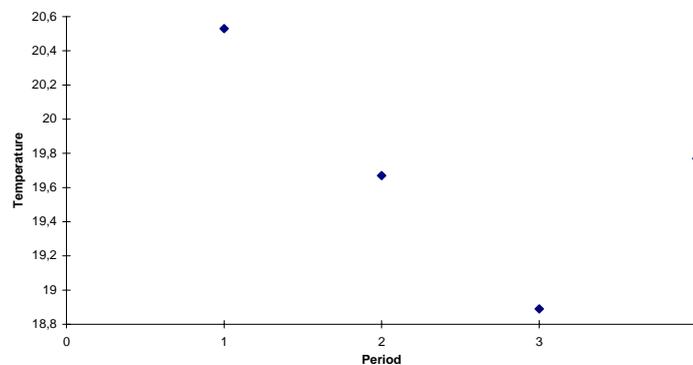


Figure 19: *The temperature of the public building. In order not to have a too big deviation from the agreed on temperature during period 2 and 3, when the resource is very scarce, the temperature is set above the set point in period 1.*

As in Example 2, the production and the loads are managed optimally at the given prices. However, there is a small potential gain for the utility in speculating about its own effect on prices. In most cases though, the potential gain is very small compared to the potential losses caused by uncertainty about the market, (*cf.* [Sandholm and Ygge, 1997]). Therefore, acting competitive is a very reasonable thing to do, and it is often the optimal behavior.

6.4 Cost-Benefit Analysis

Since every agent is facing the market price, it is able to compute an approximate value of its *profits* from the ability to control its load. If consequences on prices of the action of a single agent can be neglected, the profit is simply the cost for the consumed energy if the load was not controllable minus the cost when it is controlled.

As an example, consider a load of type 3 in Example 2 above. If the agent's effect on prices is negligible, the profit (resulting from moving 2kW from period two to period three) is $2 \cdot p_2 \Leftrightarrow 2 \cdot p_3 = 0.77$. The *actual* profit is the cost for the utility when the load is not controllable (i.e. is belonging to the uncontrollable loads) and the cost when the load is controllable. A simulation of the type in Example 2 was performed with one of the loads of type 3 considered as a uncontrollable load with consumption $[2, 2, 0, 0]$. That results in a total cost of 333.10 (of which 7.99 is compensation to the customer for temperature deviations of the public building). From Example 2 above we have that the cost when the load is controllable is 332.33. That is, the actual profit is $333.10 \Leftrightarrow 332.33 = 0.74$. Thus, even for this relatively small market the local estimation was quite accurate, and the less effect a single agent has on prices, the more accurate the estimate will be.

As another example, consider a load of type 1 in Example 1. Here the price of the first two time periods is practically equal. Hence, the profit of this contract is practically zero.

The local estimates provides the utility with extremely useful information. For each load in the system, an estimate of the value of the current load management contract is obtained. If a contract is generating less money than it costs, the utility has too many of these kind of contracts. In this situation the utility might gain money by lower the compensation for this type of contract, even if a number of customers are not willing to participate at the new lower compensation. Correspondingly, if the money generated by the contract is greater than the cost for it, the utility might profit from rising the price and hereby get more customers participating in the load management program.

7 More General Loads, Contracts, and Production and Distribution Units

In the above examples, loads, contracts, and production and distribution units were relatively simple. However, we believe the approach to work for very general settings. We will briefly discuss what the problems of more complicated settings are and how they can be approached for integration in a market-oriented system.

7.1 Non-Separable Loads

Theorem 2 in [Ygge *et al.*, 1998] is only valid for separable problems. That is, if a number of loads have interdependencies, like they often have in, e.g., industrial production, the market-oriented programming approach above is not directly applicable. Further research is required to determine whether or not a market-oriented programming approach to these kind of problems is feasible. However, a group of non-separable loads can form a separable unit. Such a unit can then successfully be integrated in the approach described

above. That is, a HOMEBOT can encapsulate a number of non-separable loads into a composed load that from the outside looks exactly like any other separable load. The approach is very similar to how the non-separability was managed in Section 6.1; the utility of the HOMEBOT for a certain amount of resource is the sum of the utility of the loads (managed by the HOMEBOT) when the resource is optimally distributed among them.

7.2 Complex Load Models

In industrial applications the resource energy is one of several inputs. Load management in such environments leads to more or less complicated production plans in which load management strategies is a part [Ericson, 1997, Björk, 1989]. In such a setting a HOMEBOT responsible for a production unit is far more complicated than the above, and the computational burden of determining the utility of different allocations of resource can be substantial, but the basic principles are the same. Hence, the approach here is not limited to any specific class of simple loads and contracts. The only constraint is that (a sufficiently accurate approximation of) the utility of different allocations must be computable in reasonable time. But this constraint is of course present for *any* approach – if it is impossible to evaluate different allocations in reasonable time, optimization is impossible.

It may also be the case that a group of loads, e.g., supplied by a critical section, are very hard to schedule. Then, this group might have to form a unit for which the scheduling aspect is included in the computation of the demand. See also Section 4.4.

In all, arbitrarily complex load models (and production and distribution models as well for that sake) can be captured by the market-oriented approach described here. The main difference between this approach and more centralized approaches, such as MIND [Nilsson and Söderström, 1992], is that with the market-oriented approach, local constraints are *encapsulated* by local agents, rather than being constraints in a global optimization computation. There is a uniform approach of encapsulating all units and determine a market equilibrium, independently of how complex local constraints are. This allows for a natural software engineering decomposition of the problem as well as for a natural decomposition of the computational problem, taking advantage of the inherent computational power of the distributed computer system.

7.3 Non-Divisible Resource

In Example 1 above we assumed that the contracted hour could be utilized arbitrarily during the four hours. We could, however, imagine that, for some reason, the contract states that the load can only be switched of for one entire hour without interruption. This causes no problem in terms of optimality – if the equilibrium exists, Theorem 2 in [Ygge *et al.*, 1998] gives that the solution is optimal – but as we can not use the linear shift as described in Figure 4, the equilibrium is no longer guaranteed to exist. A partial solution

to this is to do something very similar to what was done in the case when the resource was divisible – avoid that all loads shift at exactly the same time. For example, consider Example 1, but assume that all loads of type 1 can only be switched off during one entire hour. Now, instead of letting every load change from consuming at time period one to consuming at time period two at *exactly* $p_1/p_2 = 1$, let the loads shift at different relations, p_1/p_2 such that $1 \Leftrightarrow \delta \leq p_1/p_2 \leq 1 + \delta$. If possible, make sure that no loads shift at exactly the same relation. Then we know that the total excess demand can always be reduced to a number very close to the resource used by one single load (2 in this example). If the number of loads is large such a small total excess demand is certainly negligible, compared to, e.g., uncertainties about uncontrollable loads. If perfect feasibility is required some variant of the PROPORTION algorithm (see [Ygge and Akkermans, 1998a]) may be very useful. Furthermore, the larger the number of heterogeneous agents, the larger probability that a market equilibrium exists, even when demands are not continuous [Mas-Colell *et al.*, 1995, pp. 627 – 630]. This phenomena can turn out to be very useful in load management.

8 Conclusions

This report describes a new market-oriented approach to load management.

We claim that this approach to load management has a number of advantages compared to current methods (such as [Wafa, 1996, Popovic, 1997, Talukdar and Gellings, 1987, Larsson *et al.*, 1988, Karlsson, 1989]):

1. It provides an integrated strategy for many different types of loads and contracts. 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., [Ygge and Akkermans, 1996, Andersson and Ygge, 1998], and for the multi-commodity case, see, e.g., [Ygge and Akkermans, 1998a, Ygge and Akkermans, 1998b, Ygge, 1998]).
5. The main abstractions used, price and demand, or 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 maximize the sum of them in a distributed fashion. (For example, a non-market interpretation of Theorem 2 in [Ygge *et al.*, 1998] is given in [Andersson and Ygge, 1997].) Also the demonstrated scheduling approach is independent of whether or not the system is viewed as a market. We do not disagree with such a position, but argue that the two-level load management architecture (with longer term optimization and on-line scheduling), 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 a novel 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 described with market abstractions.

Our main arguments for a market view of the system are items (5), (6), and (7). We believe the naturalness of the approach to be of vital importance. The successful integration of a load management system in the core business information management of a utility is heavily dependent on how the system is appreciated by the people in the organization. The personnel responsible for the energy trade are very familiar with concepts such as demand and prices. Therefore, the metaphor of local agents "negotiating" over power and hereby 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 extremely useful.

In summary, our conclusion at this point in our research is that market-oriented programming is highly applicable to load management.

References

[Akkermans and Ygge, 1997] H. Akkermans and F. Ygge. Smart software as cus-

- tomers assistant in large-scale distributed load management. In *Proceedings of DA/DSM '97*. PenWell Conferences and Exhibitions, October 1997. (Available from <http://www.enersearch.se/~ygge>).
- [Akkermans *et al.*, 1996] H. Akkermans, F. Ygge, and R. Gustavsson. Homebots: Intelligent decentralized services for energy management. In *Proceedings of The Fourth International Symposium on the Management of Industrial and Corporate Knowledge, (ISMICK '96)*, October 1996.
- [Andersson and Ygge, 1997] A. Andersson and F. Ygge. Using methods for concave maximization problems to solve non-concave ones (working title). Technical report, Department of Computer Science and Business Administration, University of Karlskrona/Ronneby, 1997. In preparation. (To be made available from <http://www.enersearch.se/~ygge>).
- [Andersson and Ygge, 1998] A. Andersson and F. Ygge. Managing large scale computational markets. In *Proceedings of The 31th Hawaiian International Conference on System Sciences*, January 1998. In press. (To be made available from <http://www.enersearch.se/~ygge>).
- [Avilla, 1994] F. Avilla. On-line short term forecasting system for distribution lines based on synthetical load profiles. In *Proceedings of DA/DSM '94*. PenWell Conferences and Exhibitions, September 1994.
- [Bergström, 1996] U. Bergström. Simulation of a local energy system with focus on cost efficient dsm measures on a deregulated electricity market. In *Proceedings of DA/DSM '96*. PenWell Conferences and Exhibitions, October 1996.
- [Björk, 1989] C. Björk. *Industrial Load Management – Theory, Practice and Simulations*. Elsevier, Amsterdam, 1989.
- [Cheng and Wellman, 1997] J. Cheng and M. Wellman. The WALRAS algorithm – a convergent distributed implementation of general equilibrium outcomes. In *Computational Economics*, 1997. To appear. (Available from <http://ai.eecs.umich.edu/people/wellman>).
- [Dag, 1996] S. Dag. Modell för kostnadseffektis samverkan mellan Göteborg Energi och Volvo PV. Lic. thesis. ISBN: 91-7871-908-9, Division of Energy Systems, Department of Mechanical Engineering, Linköping Institute of Technology, 1996.
- [Ericson, 1997] L. Ericson. Dynamic load control – power peak saving applied to a foundry. Lic. thesis. ISBN: 91-88934-06-3, CODEN: LUTEDX/(TEIE-1015)/1-169/(1997), Department of Industrial Electric Engineering and Automation, Lund Institute of Technology, 1997.

- [Flood and Hyde, 1992] J. Flood and F.M. Hyde. An automated information and billing system and real-time pricing for large customer. In *Proceedings of DA/DSM '92*, January 1992.
- [Holst and Ekelund, 1987] J. Holst and J. Ekelund. Adaptiv prediktion av belastning på ett kraftnät. Technical report, Department of Mathematical Statistics, Lund Institute of Technology, 1987. LUTFD2/(TFMS-3043)/1-53/(1987), Only available in Swedish.
- [Holst, 1977] J. Holst. *Adaptive Prediction and Recursive Estimation*. PhD thesis, Department of Automated Control, Lund Institute of Technology, 1977. LUTFD2/(TFRT-1012)/1-206/(1977).
- [Karlsson, 1989] B. Karlsson. Laststyrningssystem för industrin. Lic. thesis. ISBN: 91-7870-462-1, Division of Energy Systems, Department of Mechanical Engineering, Linköping Institute of Technology, 1989.
- [Kulkarni *et al.*, 1996] A. D. Kulkarni, T. Ananthapadmanabha, and K. Raghavendra Rao. Knowledge based expert load shedding. In *Proceedings of DA/DSM '96*. PenWell Conferences and Exhibitions, October 1996.
- [Kurose and Simha, 1989] J. F. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5):705–717, 1989.
- [Larsson *et al.*, 1988] L. Larsson, T. Löfquist, H. Ottosson, A. Persson, I. Persson, I. Stolz, M. Edvinsson, and L. Möller. *Projekt Mintop*. Sydskraft, Malmö, Sweden, 1988. Only available in Swedish.
- [Mas-Colell *et al.*, 1995] Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [Nilsson and Söderström, 1992] K. Nilsson and M. Söderström. Optimizing the operating strategy of a pulp and paper mill using the MIND method. *Energy – The International Journal*, 17(10):945–953, 1992.
- [Ottosson *et al.*, 1989] H. Ottosson, G. Lindbergh, B.C.T. Bengtsson, M. Håkansson, L. Larsson, B. Linder, T. Löfquist, and S. Örtshagen. *Teknikstyrning av Elanvändningen, Rapport 1*. Sydskraft, Malmö, Sweden, 1989. Only available in Swedish.
- [Popovic, 1997] Z. Popovic. Scheduling of direct load control using mixed integer programming. In *Proceedings of DA/DSM '97*. PenWell Conferences and Exhibitions, October 1997.
- [Samuelsson, 1997] O. Samuelsson. *Power System Damping - Structural Aspects of Controlling Active Power*. IEA, LTH, Box 118, S-221 00 Lund, Sweden, 1997. ISBN 91-88934-05-5.

- [Sandholm and Ygge, 1997] T. 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>).
- [Sorri and Salo, 1992] V. Sorri and T. Salo. Influence of the cost control tariffs on the control of electricity consumption in Finland. In *Proceedings of DA/DSM '92*, January 1992.
- [Talukdar and Gellings, 1987] S. Talukdar and C.W. Gellings. *Load Management*. IEEE Press, New York, 1987.
- [Wafa, 1996] A. R. A Wafa. Optimum load management securing minimum revenue loss in industry. In *Proceedings of DA/DSM '96*. PenWell Conferences and Exhibitions, October 1996.
- [Walsh *et al.*, 1997] W. E. Walsh, M. P. Wellman, P. R. Wurman, and J. K. MacKie-Mason. Some economics of market-based distributed scheduling. Submitted. (Available from <http://ai.eecs.umich.edu/people/wellman/Publications.html>), 1997.
- [Wellman, 1993] 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.
- [Wellman, 1994] 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>).
- [Wellman, 1996] 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>).
- [Yamaki *et al.*, 1996] H. Yamaki, M. 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>).
- [Ygge and Akkermans, 1996] 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>).

- [Ygge and Akkermans, 1997] 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>).
- [Ygge and Akkermans, 1998a] F. Ygge and J.M. Akkermans. Market computations in time-critical environments. Submitted. (To be made available from <http://www.enersearch.se/~ygge>), 1998.
- [Ygge and Akkermans, 1998b] F. Ygge and J.M. Akkermans. On resource-oriented market computation. Submitted. (To be made available from <http://www.enersearch.se/~ygge>), 1998.
- [Ygge *et al.*, 1996] F. Ygge, R. Gustavsson, and H. Akkermans. Homebots: Intelligent agents for decentralized load management. In *Proceedings of DA/DSM '96*. PenWell Conferences and Exhibitions, October 1996.
- [Ygge *et al.*, 1998] F. Ygge, J.M. Akkermans, and A. Andersson. A multi-commodity market approach to power load management. Submitted. (To be made available from <http://www.enersearch.se/~ygge>), 1998.
- [Ygge, 1998] F. Ygge. *Market-Oriented Programming and its Application To Power Load Management*. PhD thesis, Department of Computer Science, Lund Institute of Technology, 1998. In preparation. (To be made available from <http://www.enersearch.se/~ygge>).