

A FRAMEWORK FOR EVALUATION OF MULTI-AGENT SYSTEM APPROACHES TO LOGISTICS NETWORK MANAGEMENT

Paul Davidsson and Fredrik Wernstedt

*Department of Software Engineering and Computer Science, Blekinge Institute of
Technology, Soft Center, 372 25 Ronneby, Sweden*

Abstract: We study the applicability of multi-agent systems (MAS) to production and logistics network management. The goal is to create and evaluate sets of intelligent agents that can cooperatively support production and logistics network decisions, as well as to compare their performance to other more traditional methods. A short description of supply chains is given, as well as a formal characterization of the problem space under investigation. We outline a general simulator that allows for a systematic evaluation of different multi-agent approaches across the different parts of this problem space. This is illustrated by a case study on district heating systems. A major concern in this domain is how to cope with the uncertainty caused by the discrepancies between the estimated and the actual customer demand. Another concern is the temporal constraints imposed by the relatively long production and/or distribution times. In the case study we show how to lessen the impact of these problems by the usage of agent clusters and redistribution of resources.

1. INTRODUCTION

Despite the fact that more parties are becoming involved in the supply chain, the process is becoming more clearly separated. Each party in the chain has become more specialized and is only responsible for the performance of a small part of the total process. Furthermore, the chain is

getting more dynamic as a result of the constant changes and movements that is part of a today supply chain. These characteristics indicate the need for sophisticated software systems that connect the logistics flows of individual companies [1]. As the supply chain is getting more fragmented the need for communication and coordination grows. However, the concept of coordination has not been incorporated in system design until recently [2]. The common usage of enterprise resource planning (ERP) systems offers the promise of integration of supply chain activities. However, the flexibility of ERP systems has been less than expected or desired [3]. Advanced planning and scheduling (APS) is an upcoming development for controlling the logistics flow through the individual elements of a supply chain. Although it is not perfectly clear how to accomplish integrated control of supply chains, some commonly mentioned requirements are the possibility to monitor the state of all involved parties, to perform collaborative planning, to perform advanced scheduling, and being able to measure performance and cost.

In order to satisfy these and related requirements, a number of agent-based approaches has been suggested. Several authors propose agents for auction-oriented management of supply chains, e.g., Fan et al. [2] provide a theoretical design that could plan the operations in a supply chain, and Hinkkanen et al. [4] focus on optimisation of resource allocation within a manufacturing plant. A rule-based approach has been proposed by Fox et al. [5], which is concentrated on coordination problems at the tactical and operational level. Preliminary work towards collaborative inventory management has been performed by Fu et al. [6]. Furthermore, agent-based approaches concerning the simulation of the dynamics in supply chains have been considered by, e.g., Parunak et al. [7], and Swaminathan et al. [8].

We argue that time has come to try to develop methods and tools for systematic evaluation of these (and other) approaches in different types of supply chain management problems. The objective of this paper is to provide guidelines and a starting point for this type of research.

In the next section we give a short description of supply networks and management philosophy. This is followed by a formal characterization of the problem space under investigation. We then outline a general simulator for production and distribution and argue that this type of simulator allows for a systematic evaluation of different multi-agent approaches across the different parts of this problem space. Finally, a case study concerning district heating systems is described. A major concern in this domain (and many other supply chains) is how to cope with the uncertainty caused by the discrepancies between estimated and real demand. Another concern is the temporal constraints imposed by the relatively long production and/or distribution times. We show how to lessen the impact of these problems by the usage of agent clusters and redistribution of resources.

2. SUPPLY CHAIN NETWORKS

A supply chain is a network of autonomous or semiautonomous suppliers, factories, and distributors, through which raw materials are acquired, refined and delivered to customers. According to the simplified view that we will adopt here, and which is illustrated in Figure 1, supply networks can be outlined as having an hourglass shape [7].

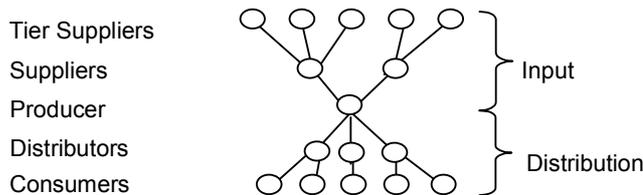


Figure 1. The hourglass shape of a simple supply network.

Of course, a supply chain is sometimes more complex with multiple end products that share components or facilities. The flow of material is not always along a tree structure, e.g., there could be various modes of transportation. Also, distribution, manufacturing and purchasing organizations along the supply chain often operate independently and have their own objectives, which may be in conflict, i.e., it can easily be argued that it is of most importance to coordinate the activities to attempt to achieve the desired global performance.

The purpose of a supply chain is to add value to its products as they pass through the supply chain (the input part) and to transport them into geographically dispersed markets in the correct quantities, at the correct time, and at a competitive cost (the distribution part). Supply chain management is concerned with integration of purchasing, manufacturing, transportation, and inventory activities. It also refers to integration of these activities across geographically dispersed facilities and markets. Finally, it also refers to temporal integration of these activities. The temporal horizon is usually separated into three levels, the strategic, the tactical, and the operational level. Long-term decisions are made on the *strategic* level, e.g., location and capacity of facilities in the supply chain. The *tactical* level is dependent on the strategic level and copes with medium term decisions, e.g., distribution planning and inventory level (buffer size). Finally, on the *operational* level, which is dependent on the strategic as well as the tactical level, the short-term decisions are made, e.g., scheduling of local transportation. Typically, the types of decisions made on the operational level are similar to those made on the tactical level but taken with a shorter time horizon in mind.

The decisions at the different levels are made in order to achieve one or more of the following goals: minimize production costs (e.g., by having an even production), minimize distribution and inventory costs, and maximize customer satisfaction. The overall goal often is to maximize profit, which almost always leads to a trade-off situation between these goals. Since this trade-off is application dependent, a specific balance between the goals is desired for a particular application.

A typical supply chain faces uncertainty in terms of both supply and demand. Thus, one of the most common problems faced by managers is to anticipate the future requirements of customers. Large errors in forecasts lead to large discrepancies between production and actual demand. This results in higher inventory costs, i.e., larger buffers are needed and/or the worth, or quality, of the products in the buffers decrease over time (deterioration). To deal with this problem, Just-In-Time (JIT) strategies have been developed. Monden [9] gives a brief definition of JIT as “producing the necessary items, in the necessary quantities at the necessary time”. Various benefits of JIT have been widely discussed in literature. However, most success stories take place in large manufacturers with stable demand, such as the automotive and electronic industries. One of the long-term aims of our work is to develop successful JIT strategies and JIT software tools for more dynamic situations.

3. A FORMAL CHARACTERIZATION OF THE PROBLEM DOMAIN

In this section, we formally define the problem space under investigation. We will restrict our attention to the distribution part of the supply network, see Figure 2.

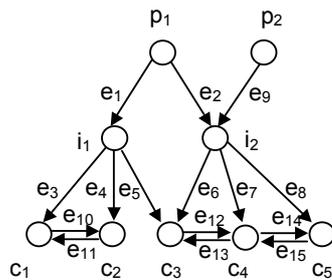


Figure 2. An example of a small production and distribution network where there are two producers ($p_{1,2}$), two distribution centers ($i_{1,2}$), and five customers (c_{1-5}).

Thus, we will not model complete supply chain networks, e.g., the details of the manufacturing process within the producer and interaction with sub-contractors are not modeled. However, we believe that the simplifications made do not change the applicability of the general approach we are suggesting.

We divide the description of the production and distribution network in three parts, *production*, *consumption* and *distribution*.

Production. Let $P = \{p_1, \dots, p_n\}$ be the set of all producers and $K = \{k_1, \dots, k_m\}$ be the set of all commodities. Then for each pair of producer, $p \in P$, and commodity, $k \in K$, we denote;

- the production time, α_{pk}
- the production cost, β_{pk}
- the production capacity, χ_{pk}
- the production at time t , δ_{pkt}

Consumption. Let $C = \{c_1, \dots, c_p\}$ be the set of all customers. Then for each pair of customer, $c \in C$, and commodity, $k \in K$, we denote;

- the consumption at time t , ϑ_{ckt}
- the demand at time t , σ_{ckt}

Distribution. Let the distribution network $D = (N, E)$ be a directed graph, where $N = P \cup C \cup I$ is the set of all nodes, $I = \{i_1, \dots, i_r\}$ is the set of all internal distribution nodes, C is the set of all customers, and $E = \{e_1, \dots, e_q\}$ is the set of all edges. Here an edge corresponds to a distribution channel between two nodes (there may be more than one edge between two nodes) and, N , is indexed as $(n_1, \dots, n_n, n_{n+1}, \dots, n_{n+p}, n_{n+p+1}, \dots, n_{n+p+r})$. Then for each pair of edge, $e \in E$, and commodity, $k \in K$, we denote;

- the distribution time, ε_{ek}
- the distribution cost, ϕ_{ek}
- the distribution capacity, γ_{ek}
- the distribution at time t , η_{ekt}

For each pair of node, $n \in N$, and commodity, $k \in K$, we denote:

- the buffer cost, μ_{nk}
- the buffer capacity, λ_{nk}
- the buffer usage at time t , ω_{nkt}

For each commodity, $k \in K$, we denote:

- the deterioration rate, θ_k

Although this model assumes that production and distribution costs etc are linear, we argue that it is possible to describe many interesting production and distribution problems using this formal model. Furthermore, there might be constraints and dependencies between different commodities concerning production, distribution, and buffer capacities. Note that the production, distribution, and buffer dynamics are part of the solution rather

than the problem, and that the possible amount of consumption is governed by these dynamics.

4. A GENERAL SIMULATOR FOR PRODUCTION AND LOGISTICS NETWORK

We are currently developing a general simulator able to simulate all the relevant production and distribution problems that can be described by the formal model presented above. This includes problems at the strategic, tactical, as well as the operational level. Each part of the model corresponds to a set of explicit simulation parameters. Thus, by setting these parameters, it will be possible to simulate an arbitrary production and distribution problem.

The long-term goal is to systematically evaluate different agent-based approaches in the different parts of the problem space defined by the formal model (and the simulator). This ambition is in line with the more general ideas presented by Davidsson and Johansson [10]. We will, of course, also use the simulator to compare the performance of the agent-based approaches to more traditional approaches. Figure 3 illustrates the interaction between the control strategies and the general simulator. The control strategy manages production and distribution and the simulator controls consumption and simulates production, distribution and consumption.

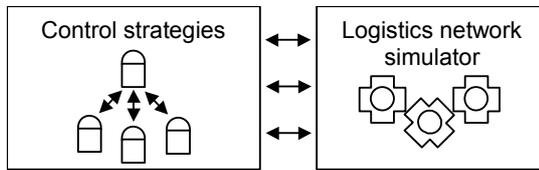


Figure 3. Schematic view of the relation between the simulator and the control strategy evaluated by the simulator.

However, we have thus far only investigated a small part of this problem space, namely that corresponding to district heating production and distribution. We will here focus on JIT production and distribution problems, which we define as situations where there are limited storage capacity or high deterioration rate, and a considerable delay from order to receipt of commodities (due to large production and/or distribution times). We have compared two different agent-based approaches in this domain, one semi-distributed and one centralized.

4.1 Case Study: Description of District Heating Systems Domain

This case is borrowed from ABSINTHE, a current collaboration project with Cetetherm AB, one of the world-leading producers of district heating substations [11]. The technological objective is to improve the monitoring and control of district heating networks through the use of agent technology. For more information on this project, see www.ipd.bth.se/absinthe.

The basic idea behind district heating is to use cheap local heat production plants to produce hot water. This water is then distributed by using one or more pumps at approximately 1-3 m/s through pipes to the customers where substations are used to exchange heat from the primary flow of the distribution pipes to the secondary flows of the building, see Figure 4. The secondary flows are used for heating both tap water and the building itself. In large cities district heating networks tend to be very complex, including tens of thousands of substations and hundreds of kilometers of distribution pipes with distribution times up to 24 hours.

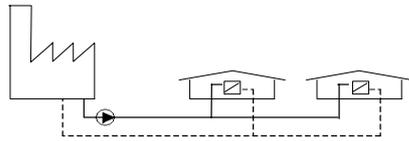


Figure 4. A simple district heating network containing one heat producer and two substations.

Let us describe a district heating system in terms of the general model. We here present a description corresponding to how parameters are set in the current version of the simulator, in which some simplifications have been made compared to actual district heating systems. For instance, we assume that there is only one production plant. However, we intend to remove such restrictions in future versions.)

Production. Since there is only one heat production plant, p_1 , and the only type of commodity, k_1 , is energy (hot water), we will here use a simplified notation that omits the p_1 and k_1 . We assume that raw materials are sufficient to support the production (up to the production capacity, χ). Moreover, we assume that:

- the production time, α , is negligible, i.e., 0 seconds,
- the production cost, β , is 1 cost unit for each kWh, and
- the production capacity, χ , is larger than the total demand (see below for definition of, σ).

Consumption. The set of customers, C , consists of 10 customers. Five customers, $c_1 - c_5$ are serving 40 households each and five customers, $c_6 - c_{10}$ are serving 60 households each. Moreover, we assume that:

- the demand of the customers, σ_{ct} , is composed by two parts: hot tap water demand and heating demand. The tap water demand is simulated using a model [12] based on empirical data where flow size and tapping durability is determined by the simulation of a random number, Y , from a certain distribution with cumulative distribution function F_Y , which can be performed using uniformly distributed numbers, and where the time between tapping is a non-homogenous Poisson process;

$$X = 1 - \exp\left(-\int_0^T \mu(u) du\right), \text{ where } \mu(u) \text{ is the time-varying opening}$$

intensity and T the time to derive. The opening intensities are derived from $\mu = \frac{p}{(1-p)\eta}$, where p is given from measurement data [13] and η

is calculated from the distribution function for open valve time. The resistance/capacitance model described above is used also for simulation of the energy needed for household heating (radiators). The variance over time of the outdoor temperature is simulated by the following model [14]; $T_o = T_m + T_v * e^{-((i*s-4) \bmod 24-12)^2 / 20}$, where T_m is the lowest temperature to expect, T_v is the maximum temperature to expect, and S is the time interval expressed in hours.

Distribution. The distribution network, D , is assumed to be a tree structure with the production plant as the root and the distribution pipes corresponding to the edges. A relaxing constraint in district heating systems is that the distribution time and cost between customers physically close to each other (situated in the same branch of the distribution tree) is negligible. We model this by internal cluster nodes, I , one for each cluster of neighboring customers, see Figure 5.

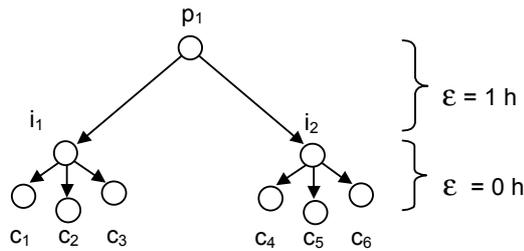


Figure 5. Customers close to each other form clusters where redistribution is instantaneous.

Moreover we assume that:

- the distribution time, ε_e , is 1 h for edges, e , between the producer, p_1 , and internal nodes, $n \in I$, and negligible, i.e., 0 for edges between internal nodes, $n \in I$, and customers, $n \in C$,
- the distribution cost, ϕ_e , is 0 for all edges.
- the distribution capacity, γ_e , is greater than the total demand for the cluster to which e is pointing,
- the buffer capacity, λ_n , is 0 for $n \in C$ (the customer has no potential to keep inventories of hot water) and for p_1 . If, $n \in I$ the capacity is greater than the total demand for the cluster to which n belongs,
- the buffer cost, μ_n , is negligible, i.e., 0, and
- the deterioration rate, θ , is computed from the common resistance/capacitance model;

$$T_{xi} = \frac{1}{1 + \frac{1}{TR_x TC_x}} \left(T_{x(i-1)} + \frac{\frac{T_{outi}}{TR_x}}{TC_x} \right), \text{ where } T_{xi} \text{ is the temperature of a}$$

object, x , at the time i which had temperature $T_{x(i-1)}$ one time unit ago with thermal resistance TR_x and thermal capacity TC_x in an surrounding environment with temperature T_{outi} .

4.2 Multi-Agent System Approaches to JIT Production and Distribution in District Heating Systems

There are a number of different approaches to solve the JIT production and distribution problem outlined above. The most basic approach (and probably the most used) is strictly centralized, where the producer based on experience makes predictions of how much resource to produce in order to satisfy the total demand. These resources are then distributed directly on request from the customers. A bit more sophisticated is the approach where each customer makes predictions of future consumption and informs the producer about these predictions. Since local predictions typically are more informed than global predictions, this approach should give better results. The multi-agent system (MAS) architecture we suggest below partly builds upon this insight but also introduces a means for automatic redistribution of resources. In order to solve the problem of producing the right amount of resources at the right time, each customer is equipped with an agent that makes predictions of future needs that it sends to a production agent. The other problem, to distribute the produced resources to the right customer, is approached by forming clusters of customers within which it is possible to redistribute resources fast and at a low cost. This usually means that the customers within a cluster are closely located to each other. In this way it is

possible to cope with discrepancies between predicted demand and the actual consumption. For instance, this happens if the demand of a customer changes while the resources are delivered. The customer would then be faced with either a lack or an excessive amount of resources, thus leaving the system in an undesired state.

Based on the above insights, we used the GAIA methodology [15] to design the MAS. This led us to the architecture outlined in Figure 6, that has the following three types of agents:

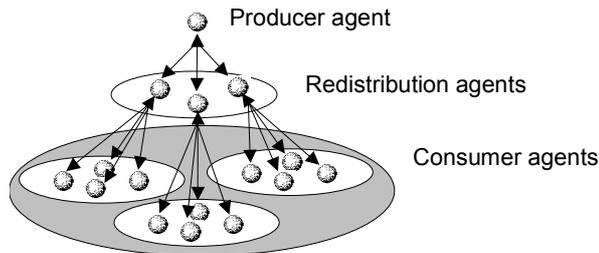


Figure 6. Each redistribution agent manages a cluster of consumer agents.

- *Consumer agents*: (one for each customer) which continuously (i) make predictions of future consumption by the corresponding customer and (ii) monitor the actual consumption, and send information about this to their redistribution agent.
- *Redistribution agents*: (one for each cluster of customers) which continuously (i) make predictions for the cluster and send these to the producer agent, and (ii) monitor the consumption of resources of the customers in the cluster. If some customer(s) use more resources than predicted, it redistributes resources within the cluster. If this is not possible, i.e., the total consumption in the cluster is larger than predicted, it will redistribute the resources available within the cluster according to some criteria, such as, fairness or priority, or it may take some other action, depending on the application.
- *Producer agents*: (one for each producer, however, we will here only regard systems with one producer) receives predictions of consumption and monitors the actual consumption of customers through the information it receives from the redistribution agents. If necessary, e.g., if the producer cannot produce the amount of resources demanded by the customers, the producer agent may notify the customers about this (via the redistribution agents).

The suggested approach makes use of two types of time intervals: (i) prediction intervals and (ii) redistribution intervals. A prediction interval is larger than a redistribution interval, i.e., during each prediction interval there

is a number of redistribution intervals. Each consumer agent produces one prediction during a prediction interval and sends this to its redistribution agent, who sums the predictions of all consumer agents belonging to the cluster and informs the producer agent about this. The predictions made by the consumer agent must reach the producer at least $\alpha + \varepsilon$ before the resources are actually consumed. Typically, there is also a production planning time that also should be taken into account (i.e., added to the sum).

The coordination technique that we use is *organizational structure*, i.e., the responsibilities, capabilities, connectivity and control flow are defined *a priori* [16]. Organizational structure can be seen as a long term, strategic load balancing technique [17]. Malone and Crowstone define coordination as “managing the dependencies between activities” [18]. The basic coordination process to manage in district heating systems is the producer/consumer relationship, i.e., the main dependencies are prerequisite constraints (some activity must be finished before another can begin), transfer (something needs to be transported), usability (one part of the system needs information produced by another part), and simultaneity constraints (some activities need to occur at the same time).

4.3 Simulation Results

The focus during the initial simulation experiments [19] was to see how the quality of service (QoS), measured in terms of the number of restrictions issued, varied with the amount of excessive production (in relation to the predicted consumption).

We found that the MAS performed well, coping with faulty predictions (even though the discrepancy between the predicted and the actual consumption sometimes was quite large). We also discovered that in order to achieve the same QoS between the two sets of experiments, the centralized approach required an additional 3% excessive production.

Figure 7 shows the total number of restrictions to tap water and the number of restrictions to water for household heating (radiator) during one day for different degrees of surplus production in a cluster of ten consumer agents. We see that there is a clear trade-off between the quality of service (number of restrictions) and the amount of surplus production and that there are almost no restrictions of any kind when 4% more hot water than the predicted consumption is produced. Moreover, this approach allows large fluctuations in customer demand, which is something that has been argued not to suit JIT approaches [20].

We have compared this approach to a more centralized approach without redistribution agents and where all computation is carried out at the producer side. The only task of the agents on the customer side was to read sensor

data and send those to the producer agent. Simulation results showed that more than 7% overproduction was needed in order to avoid any restrictions/shortages. For more information about this approach and simulation results, see [19].

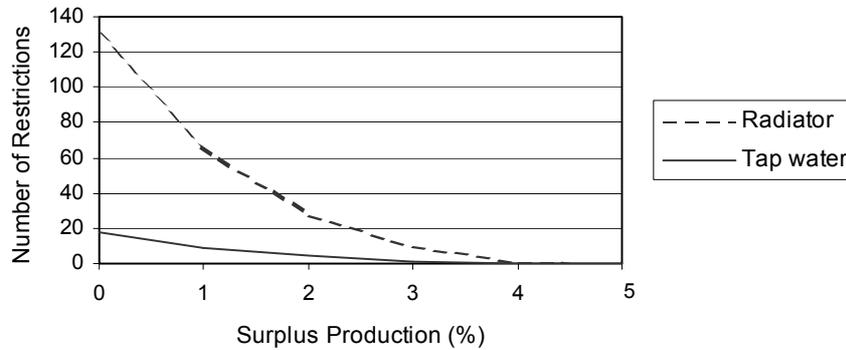


Figure 7. Trade-off between quality of service and surplus production for the semi-distributed approach. The dashed line corresponds to the number of restrictions for the radiator water and the other line to the number of restrictions for the tap water.

5. CONCLUSIONS AND FUTURE WORK

We have presented an outline of a general simulator based on a formal model of production and logistics networks. The goal is to implement such a simulator and use it for systematical evaluation of both new and existing agent-based approaches to supply chain management. We also described the first step towards this goal in the form of a case study.

In the near future we plan to generalize the current version of the simulator to cover more types of production and distribution domains. We would also like to be able to consider environmental aspects of agent-based systems in supply chains, e.g., dynamical support for alternative transportation routes. Also, we plan to implement other agent-based approaches as well as improving the one presented here. Furthermore, an important item on the research agenda is to identify and analyze alternative coordination processes, e.g., mechanisms that support JIT strategies.

Acknowledgement. This work has been financially supported by VINNOVA (The Swedish Agency for Innovation Systems).

REFERENCES

1. Knoors, F.: Establish the Door-To-Door Management System, European Commission, DG TREN 5th FP, D2D, WP3, Sequoyah International Restructuring, 2002.
2. Fan, M., Stallaert, J., Whinston, A.B.: Decentralized Mechanism Design for Supply Chain Organizations Using an Auction Market, *To appear in Information Systems Research*.
3. Shapiro, J.F.: *Modeling the Supply Chain*, Duxbury, 2001.
4. Hinkkanen, A., Kalkota, R., Saengcharoenrat, P., Whinston, A.B.: Distributed Decision Support Systems for Real-Time Supply Chain Management Using Agent Technologies, *Readings in Electronic Commerce*, 275-291, 1997.
5. Fox, M.S., Barbuceanu, M., Teigen, R.: Agent-Oriented Supply-Chain Management, *The International Journal of Flexible Manufacturing Systems*, 12: 165-188, 2000.
6. Fu, Y., Souza, R., Wu, J.: Multi-agent Enabled Modeling and Simulation Towards Collaborative Inventory Management in Supply Chains, *Winter Simulation Conference*, 2000.
7. Parunak, H.V.D., Savit, R., Riolo, R.L.: Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide, *Multi-Agent Systems and Agent-Based Simulation*, LNAI 1534, 10-25, Springer Verlag, 1998.
8. Swaminathan, J., Smith, S.F., Sadeh, N.M.: Modeling Supply Chain Dynamics: A Multiagent Approach, *Decision Sciences*, 29(3), 1998.
9. Monden, Y.: Toyota Production System: An Integrated Approach to Just-In-Time, *Industrial Engineering and Management Press*, Georgia, 1993.
10. Davidsson, P. and Johansson, S.: Evaluating Multi-Agent System Architectures: A Case Study Concerning Dynamic Resource Allocation, *Third International Workshop on Engineering Societies in the Agents' World*, Madrid, Spain, 2002.
11. Wernstedt, F., and Davidsson, P.: An Agent-Based Approach to Monitoring and Control of District Heating Systems, *Developments in Applied Artificial Intelligence*, LNAI 2358, 801-812, Springer Verlag, 2002.
12. Arvastson, L., and Wollerstrand, J.: On Sizing of Domestic Hot Water Heaters of Instantaneous Type, *6th International Symposium on Automation of District Heating Systems*, 1997.
13. Holmberg, S.: Norrköpingsprojektet – en detaljerad pilotundersökning av hushållens vattenförbrukning, *Technical report M81:5*, Departement of Heating and Ventilation Technology, Royal Institute of Technology, Sweden, 1981 (in Swedish).
14. Ygge, F., and Akkermans, H.: Decentralized Markets versus Central Control: A Comparative Study, *Journal of Artificial Intelligence Research*, 11: 301-333, 1999.
15. Wooldridge, M., Jennings, N.J., Kinny, D.: The Gaia Methodology for Agent-Oriented Analysis and Design, *Journal of Autonomous Agents and Multi-Agent Systems*, 3(3): 285-312, 2000.
16. Nwana, H., Lee, L., Jennings, N.R.: Coordination in Software Agent Systems, *The British Telecom Technical Journal*, 14(4): 79-88, 1996.
17. Durfee E.H., Lesser V.R.: Planning Coordinated Actions in Dynamic Domain, *Technical report COINS-TR-87-130*, Department of computer and information science, University of Massachusetts, Amhest, 1987.
18. Malone, T.W., and Crowstone, K.: The Interdisciplinary Study of Coordination, *ACM Computing Surveys*, 26(1), 1994.
19. Davidsson, P., and Wernstedt, F.: A Multi-Agent System Architecture to Monitoring and Control of District Heating Systems, *To appear in the Knowledge Engineering Review*.
20. Hall, R.W.: *Zero Inventories*, Dow Jones-Irwin, Illinois, 1983.