An Agent-Based Approach to Monitoring and Control of District Heating Systems

Fredrik Wernstedt and Paul Davidsson

Department of Software Engineering and Computer Science
Blekinge Institue of Technology
Soft Center, 372 25 Ronneby, Sweden
{fredrik.wernstedt,paul.davidsson}@bth.se

Abstract. The aim is to improve the monitoring and control of district heating systems through the use of agent technology. In order to increase the knowledge about the current and future state in a district heating system at the producer side, each substation is equipped with an agent that makes predictions of future consumption and monitors current consumption. The contributions to the consumers, will be higher quality of service, e.g., better ways to deal with major shortages of heat water, which is facilitated by the introduction of redistribution agents, and lower costs since less energy is needed for the heat production. Current substations are purely reactive devices and have no communication capabilities. Thus, they are restricted to making local decisions without taking into account the global situation. However, a new type of "open" substation has been developed which makes the suggested agent-based approach possible.

1 Introduction

Agent technology is currently a very active area of research and is widely applied in research labs all around the world. However, few industrial applications exist, in particular where the problem domain is truly distributed and heterogeneous, i.e., the type of domains in which agent technology is supposed to excel. We will here present some initial results from a project concerned with such a domain, namely district heating (which is very similar to district cooling).

The control of district heating systems can be seen as a just-in-time [9] production and distribution problem where there is a considerable delay between the production and the consumption of resources. The reason for the delay may be either long production time or, as in the case of district heating, long distribution time. Another characteristic of this class of problems is that resources need to be consumed relatively quickly after they have arrived to the consumer. In order to cope with these problems it is essential to plan the production and distribution so that the right amount of resources is produced at the right time.

T. Hendtlass and M. Ali (Eds.): IEA/AIE 2002, LNAI 2358, pp. 801-811, 2002.

District heating systems are inherently distributed both spatially and with respect to control. A customer, or more commonly, a set of customers, is represented by a substation embedded within the district heating network. Currently, the substation instantaneously tries to satisfy the demands of its customers without considering the amount of available resources or the demands of other substations. Each substation can be viewed as a "black-box" without communication capabilities, making local decisions without taking into account the global situation. Thus, today a district heating network is basically a collection of autonomous entities, which may result in behaviour that is only locally optimal. For instance, during a shortage in the network, resource allocation is unfair since consumers close to the production source will have sufficient amount of heat, while those distantly located will suffer.

Another consequence of the way that current substations work is that the producers only have very limited information concerning the current state of the district heating system. This together with the considerable distribution time, results in that the amount of heat to produce and deliver to the substations is typically based on uninformed estimates of the future heat demand, made by the control engineer at the production plant. In order to ensure sufficient heat supply, the tendency has been to produce more heat than necessary and hence an important waste of energy [1, 3].

The ABSINTHE (Agent-based monitoring and control of district heating systems) project is an effort aimed at dealing with these problems. It is a collaboration with Cetetherm AB, one of the world-leading producers of substations (heat exchanger systems). ABSINTHE is part of a larger venture "Intelligent District Heating", which aims at revolutionising the district heating industry, mainly by developing a new type of "open" substation, which is based on a communication and computation platform developed by Siemens. The main goal of ABSINTHE is to develop a decision support system for the district heating system operators that makes it possible to reduce the surplus production by increasing the knowledge about the current and future state of the system. To implement this we equip each substation with an agent that continually makes predictions of future consumption and monitors current consumption. However, in order to deal with situations where there is a shortage of heat in the network, there is also a fully automated part of the system supporting cooperation between substations. To do this we introduce redistribution agents that are able to impose minor, for the customer unnoticeable, restrictions on a set of substations. Another goal of the project is to more fairly deal with situations where there is a global shortage of heat by using different modes when issuing restrictions.

We begin by briefly describing the district heating domain and the involved hardware technology. This is followed by a description of the multi-agent system that has been developed within the project in order to solve the problems discussed above. Finally, we provide conclusions and pointers to future work.

2 District Heating Systems

The basic idea behind district heating is to use cheap local heat production plants to produce hot water (in some countries steam is used instead of water). The water is then distributed by using pumps at approximately 1-3 m/s through pipes to the customers where it may be used for heating both tap water and the radiator water. The cooled water then returns to the production plant forming a closed system (see Fig. 1).

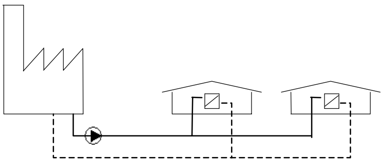

Fig. 1. A simple district heating network containing one heat producer and two consumers

At the customer side, there is a *substation* (see Fig. 2). It is normally composed of two or three heat exchangers and a control unit, which receives hot water from the district heating network. The substation heats both cold tap water and the water in the radiator circuit by exchanging the required heat indirectly from the primary flow of the distribution network. The hot network water is returned to the network at a somewhat lower temperature. Both the temperature of the returning water and the flow rate in the network are dependent on the consumption of substations. When the water, returned by substations, arrives at the heat production plant it is heated and again pumped into the distribution network.

Several different energy sources may be used for heating, e.g., waste energy, byproduct from industrial processes, geothermal reservoirs, otherwise combustion of fuels as oil, natural gas etc. is used. If the demand from the customers is high several heat producing units must be used. A district heating system in a large city can be very complex, containing thousands of substations and hundreds of kilometers of distribution pipes. In addition, they are dynamical as new substations may be added or old substations may be replaced by new ones with different characteristics.

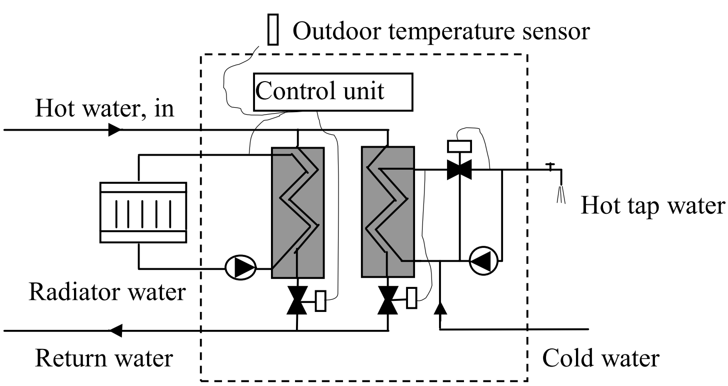

Fig. 2. A substation consisting of heat exchangers (the shaded boxes), control valves, pumps and a control unit. The radiator system (household heating) is controlled by the control unit, using information about actual outdoor temperature

Most district heating control systems of today are strictly reactive, i.e., they only consider the current state and do not predict what is likely to happen in the future. As the distribution time from the heat production plant to the customers is large, the decision on how much heat to produce becomes complicated. Ideally, the control engineer knows the amount to produce several hours ahead of consumption. Load prediction is difficult since many factors are unknown or uncertain which force the operators to make coarse estimations of future consumption based mainly on experience and simple rules-of-thumb [13]. As a consequence, and in order to be sure to satisfy the consumers, district heating systems are typically run with large margins producing more heat than necessary [1, 3]. Furthermore, operators are usually busy with keeping the heating plants running and the time available for making production decisions is therefore limited [1].

Consumption in a district heating network is mainly composed of two parts [14]:

- The heating of buildings, which mainly is a linear function of the outdoor temperature.
- The consumption of tap water, which mainly is dependent on consumption patterns, e.g., social factors.

The tap water consumption of a substation is very "bursty" even in large buildings, and therefore very difficult to predict, whereas the radiator water consumption is "smoother" and therefore relatively easy to predict assuming that reliable weather predictions are available.

Due to the rising demand of automation of building services (heating, ventilation, and air-conditioning etc.) Siemens have developed the Saphir, an extendable I/O platform with an expansion slot for a communication card, suitable for equipment control. Easy and quick access to sensor data is provided by a Rainbow communication card in the expansion slot (see Fig. 3). It has previously not been possible to develop, or make commercially available, such an advanced platform due to high costs.

The Saphir containes a database that continuously is updated with sensor data from the I/O channels by a small real-time operating system, which is directly accessible from the Rainbow card. On the Rainbow card a small computational platform (a handheld PC) makes it possible to easily deploy software and by that providing the possibility to host an agent. Hence, an agent deployed on such a platform could potentially read all connected sensor input as well as send commands over the I/O channel to actuators on the hardware, e.g., valves on a heat exchanger.

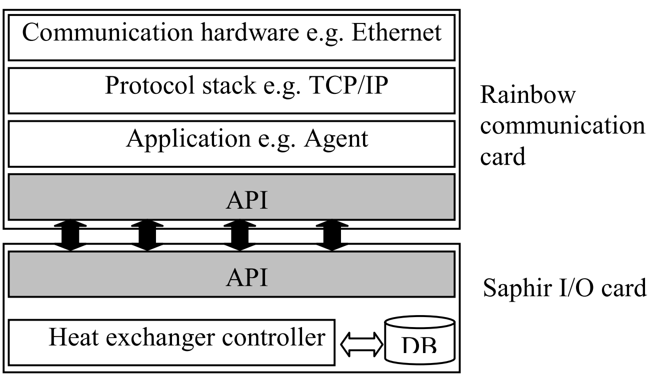

Fig. 3. The Rainbow communication and computation card is here shown on top of the Saphir hardware interface card

3 Software System Architecture

In this chapter we will first discuss the advantages and disadvantages of a distibuted compared to a centralized solution to the problem of district heating monitoring and control. We will then argue for an agent-based approach, and eventually suggest a multi-agent architecture.

3.1 Distributed versus Centralized Approaches

District heating systems are by their very nature physically distributed. Thus, if we aim at a monitoring and control approach based on knowledge about the current state of the system, at least sensor data must be collected via a distributed system. The question is whether also computation and control should, or, need to, be distributed.

In principle, it is possible to continuously collect all sensor data at each substation, do all computations necessary for the control of the system at a single central computer, and then send control signals to each of the substations. In some problem domains, it is possible to increase the utilization of resources if a global picture of the system state is available. In the district heating domain, this is certainly true for the part of the problem related to the production, which we implement as a decision support system (basically a monitoring system) for the network and production plant operators. However, we argue that one subproblem, the construction of the local consumption predictions, should be performed locally. The reason is that these computations may involve substantial amounts of sensor data, which otherwise need to be communicated to a central computer. By doing the consumption predictions locally, less communication is necessary without reducing the quality of the predictions. Also the computations involved in making the forecasts can be computationally resource demanding, and performing these in a centralized fashion would hardly be feasible. One could also argue that some of the sensor data used should be constrained for local usage only, due to its potentially sensitive nature.

Furthermore, to assume that it is easy to collect and use the sensor information from each entity in the network in a centralized fashion is somewhat wrong. Since substations, pumps and valves, etc., often are manufactured by different organisations, it would be a complex task to keep track of all these aspects centrally. Developing local monitoring and control software adopted for each type of substation, but with the same inteface to the rest of the software system, seems as a much more natural approach.

A more general argument against centralized approaches for problems as complex as the management of district heating systems (where a large number of parameters and constraints should be taken into account), is that when the problems are too extensive to be analysed as a whole, solutions based on local approaches often allow them to be solved more quickly [12].

Regarding, the part of the problem that concerns how to fairly deal with shortages of heat in the system, also a semi-distributed approach, were the control is distributed to clusters of substations, has been considered. A completely centralized approach may result in severe communication problems without achieving greater fairness. It would use the same number of messages as the semi-distributed approach, but with a possible communication bottleneck at the central computer. Also, each message would need to travel a longer route which would increase the total network load. A completely distributed approach, on the other hand, would result in a larger number of messages being sent than the semi-distributed approach without achieving greater fairness.

3.2 Why an Agent-Based Approach?

A general advice on when to consider agent technology is to be requirement-driven rather than technology-driven. Thus, we should investigate whether the characteristics of the target domain match the characteristics of the domains for which agent-based systems has been found useful. Parunak [10] argues that agents are appropriate for applications that are modular, decentralised, changeable, ill-structured, and complex. We argue that district heating systems has all these characteristics:

- *Modular:* Each entity of a district heating system, i.e., substations, heat production plants, pumps etc., can be described using a well-defined set of state variables that is distinct from those of its environment. Also, the interface to the environment can be clearly identified for each entity.
- *Descentralised:* The entities of a district heating system can be decomposed into stand-alone geographically distributed autonomous nodes capable of performing useful tasks without continuous direction from some other entity .
- *Changeable:* The structure of a district heating system may change as new entities are added or old enitites are replaced. In addtion, there are short-term changes in the system when individual substation or parts of the network are malfunctioning.
- *Ill-structured:* All information about a district heating system is not available when the monitoring and control system is being designed.
- *Complex:* District heating systems are considered to be very complex systems [3]. The entities of a district heating system exhibits a large number of different behaviours which may interact in sophisticated ways. In addition, the number of entities in a district heating system can be very large, up to a couple of thousands.

There are also more general arguments for choosing an agent-based approach [4]. From a methodological perspective the concept of agents introduces a new level of abstraction that provides an easier and more natural conceptualisation of the problem domain. Other advantages are increased, e.g., *robustness*, the distribution of control to a number of agents often implies no single point of failure, *efficiency*, less complex comuptations and communication are necessary if control is distributed, *flexibility*, the use of agent communications languages that support complex interaction between entities provides a flexibility that is difficult to achieve using traditional communication protocols, *openness*, by having a common communication language, agents implemented by different developers are still able to interact with each other, *scalability*, it is easy to add new agents to a multi-agent system, and finally, *economy*, since agent technology provides a natural way to incorporate existing software.

3.3 Multi-Agent System Architecture

The MAS architecture we suggest is composed of one type of agents associated with the producers, responsible for the interaction between the heat production plant and the other agents of the MAS, and another type of agents associated with the consumers, responsible for the interaction between the substation and the other agents of the MAS.

In a domain with limited resources, such as district heating, agents must coordinate their activities with each other to satisfy group goals [6]. To identfy group goals is complicated by two types of conflicts between individual goals. First, there is a conflict between consumer agents, who wants to maximize the comfort of the consumers by taking as much heat from the network as the consumer asks for, and the producer agents, who want to produce as little heat as possible to reduce costs. To deal with this conflict we define the following group goal: produce as little heat as possible while maintaining sufficient level of customer satisfaction. The second type of conflict is between consumer agents when there is a shortage of heat in a part of the network. In this situation each of the consumer agents in that part of the network wants to satify their consumer´s demand, which is impossible. To deal with this conflict we define the following group goal: when there is a shortage, the available heat should be shared fairly between the consumers. Satisfaction of these group goals could be achieved by either :

- competition, where each agent competes for resources, i.e., a self-interested approach (This approach is similar to the current situation, where a substation only consider the demands of their customer.), or
- cooperation, by letting the overall goal of primarily ensuring consumption of tap water, in favour of radiator consumption, affect all consumers, i.e., where one consumer could reduce radiator consumption to benefit tap water consumption for another consumer.

We implement cooperation to achieve redistribution of resources. This is practically possible since pressure propagates by the speed of sound in water and that district heating systems are, at least in principle, parallel coupled, i.e., resources not consumed by one consumer is available for other consumers in the vicinity. Also, we need the capability to restrict consumers if they try to consume more in total than predicted. For these reasons we introduce another type of agent, redistribution agents, which are responsible for a cluster of consumer agents and has both a mediator and decision maker role [5]. The mediator role includes receiving the consumptions and predictions from the consumers in the cluster, summarize them, and distribute the result to the producer agent (see Fig. 4). The decision maker role concerns what actions to take, i.e., impose restrictions upon consumer agents, to maintain an overall acceptable consumption rate (which is defined by the predictions made earlier).

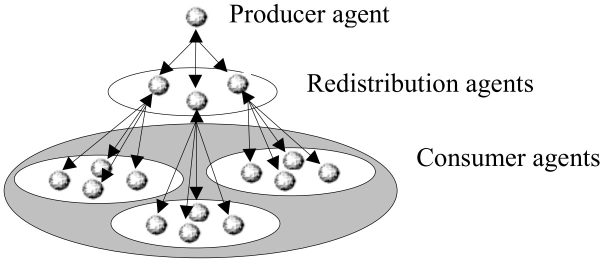

Fig. 4. Each redistribution agent manages a cluster of consumer agents

The *consumer agents* are responsible for the interaction with the substations (see Fig. 5). It monitors the actual heat consumption by reading the substation's sensors and_decides which data to send to the redistribution agent. The agent continually evaluates previously made predictions using historical data and creates new predictions of future consumptions and sends these to the redistribution agent. The consumption is sent rather often, e.g., each minute, whereas the predictions are sent at larger interval, e.g., for each 10 minutes period.

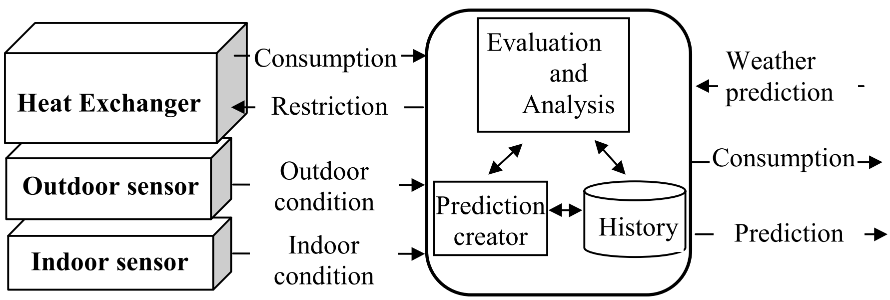

Fig. 5. Consumer agent architecture and its interaction with the environment

The *redistribution agents* are responsible for collecting predictions and monitoring the total consumption of a cluster of consumer agents. If the redistribution agent notices that the cluster is trying to consume more heat than predicted it invoke restrictions of consumption to the cluster. Restrictions can be invoked in two modes, one where each consumer agent in the cluster is to reduce its consumption with the same amount and one where different priority values are used, e.g., there might be reason that a hospital has a higher priority than a university.

A restriction will enforce consumers to not use any radiator water during the next consumption interval. If this is not enough to compensate for the excessive consumption, the redistribution agent will also impose restrictions on tap water consumption. This redistribution strategy (and all other strategies based on radiator water restrictions) will lead to a radiator water deficit that needs to be compensated (otherwise the temperature in the households, eventually, will fall). In order to do this, each substation will individually compensate by using more radiator water than predicted when the tap water consumption is less than predicted. Compensation for tap water deficits work in a similar way, but on the cluster level (see Fig. 6).

Another problem that has to be solved by the redistribution agent is how to cope with the "bursty" consumption of tap water without commanding unnecessary restrictions. The approach used here is to let the cluster use more tap water than predicted (measured in minute averages) in the beginning of a consumption period and then gradually lower the allowed average consumption towards the predicted average consumption.

The *producer agent* receives predictions of consumption from the redistribution agents and is responsible for the interaction with the control system of the heat production plant (possibly including human operators). The producer agent is also responsible for monitoring the actual consumption of consumer agents. The figures of consumption may be used to calculate the returning temperature, i.e., the producer would know in advance the temperature of the water to heat. The producer agent is also capable of imposing proactive restrictions for consumers to reduce the cost of producing heat to cover for short temporary heating needs. This will especially be the case if the predicted consumption marginally exceeds the capacity of the primary heating plant or only exceeds the capacity for a short time period (see Fig. 7).

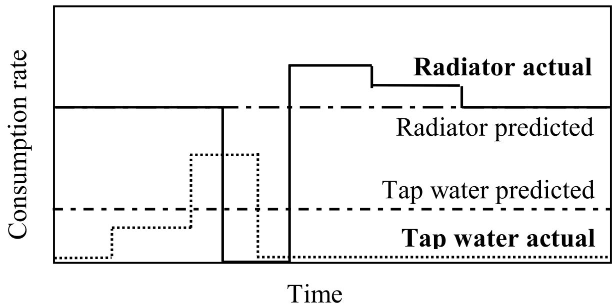

Fig. 6. A schematic view of the consequences of a restriction

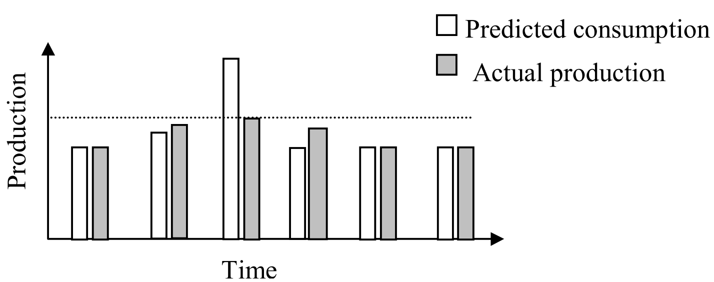

Fig. 7. A predicted consumption for a limited time period above the capacity of the primary production source (the dashed line) can be covered by a temporary imposing restrictions of consumption

The system is dependent on the truthfulness of agents, e.g., individual consumer agents should not lie to gain benefits. However, we believe that this will not pose a problem since the system can be viewed as closed with respect to development and deployment of the agents, for instance by introducing certification procedures.

4 Conclusions and Work in Progress

In the short term the suggested project will adapt and introduce agent technology in an industrial application were the use of leading edge information technology currently is very low. To our knowledge, agent technology has never been applied to monitoring and control of district heating systems. It will provide a novel combination and integration of existing technologies, which will open up new possibilities.

The contributions to the final users, i.e., the operators and the consumers, will be higher quality of service, e.g., better ways to deal with major shortages of heat water, and lower costs, i.e., less energy is needed to produce the heat water. Since the heating of water often is associated with burning fuel that pollutes the air in one way or another, the project obviously contributes to increase the quality of life for the inhabitants. We also believe that the introduction of advanced information and communication technology will enhance the work situation for the network operator staff, e.g., through the new possibility for remote diagnosis of heat exchanger systems.

We have implemented a simplified version of the MAS described in this paper using Jade [2] and applied the Gaia methodology [15] for the design. Initial simulation ex-periments have shown that it is possible to reduce the production by at least 4 percent while maintaining the overall consumer satisfaction, or quality of service. However, there are indications that this figure can be improved considerably.

The implemented system is not a finished product, however, it could at this stage be considered an emulated system. To reach the next maturity level [11], prototype, we are in progress of testing the approach on real domain hardware, but under laboratory conditions. Furthermore, the prediction mechanism of the consumer agents needs to be improved and we are looking at, e.g., neural network-based approaches [3, 7].

Acknowledgements

This work has been financially supported by VINNOVA (The Swedish Agency for Innovation Systems). The authors also wish to acknowledge the valuable input from the ABSINTHE project members employed at Cetetherm AB.

References

1. Arvastsson, L.: Stochastic Modeling and Operational Optimization in District Heating Sys-tems, Doctoral Thesis, Lund Institute of Technology, Sweden (2001).
2. Bellifemine, F., Poggi, A., and Rimassa, G.: Developing multi-agent systems with a FIPA-compliant agent framework. Software: Practice and Experience, Vol. 31(2), John Wiley & Sons, Ltd, New York (2001) 103-128.
3. Canu, S., Duran, M., and Ding, X.: District Heating Forecast using Artificial Neural Net-works, International Journal of Engineering, Vol. 2(4) (1994).

4. Davidsson, P. and Wernstedt, F.: Software Agents for Process Monitoring and Control, Journal of Chemical Technology and Biotechnology, 2002 (To appear).
5. Ferber, J.: Multi-Agent System, An introduction to Distributed Artificial Intelligence, Addison Wesley (1999).
6. Huhns, M., Stephens, L.: Multiagent Systems and Societies of Agents. In Weiss, G. (ed.) Multiagent Systems, MIT Press (1999).
7. Malmström, B., Ernfors, P., Nilsson, D., and Vallgren, H.: Short-term forecasts of district heating load and outdoor temperature by use of on-line-connected computers. Värmeforsk, 1996:589, ISSN 0282-3772 (in Swedish).
8. Nwana, H., Lee, L., and Jennings, N.R.: Coordination in Software Agent Systems, The British Telecom Technical Journal, 14 (4) (1996) 79-88.
9. Ohno, T.: Toyota Production System. Productivity P, US (1988).
10. Paranak, H.: Industrial and Practical Applications of DAI. In Weiss, G. (ed.) Multiagent Systems, MIT Press (1999).
11. Parunak, H.: Agents in Overalls: Experiences and Issues in the Deployment of Industrial Agent-Based Systems. International Journal of Cooperative Information Systems Vol. 9(3) (2000) 209-227.
12. Rinaldo, J. and Ungar, L.: Auction-Driven Coordination for Plantwide Optimization. Foundations of Computer-Aided Process Operation FOCAPO (1998).
13. Weinspach, P.M.: Advanced energy transmission fluids. International Energy Agency, NOVEM, Annex IV, Sittard, Netherlands (1996).
14. Werner, S.: Dynamic heat loads from fictive heat demands. Fjärrvärmeföreningen, FOU 1997:10, ISSN 1402-5191 (in Swedish).
15. Wooldridge, M., Jennings, N.J., and Kinny, D.: The Gaia Methodology for Agent-Oriented Analysis and Design. Journal of Autonomous Agents and Multi-Agent Systems, Vol. 3(3) (2000) 285-312.