A Case Study on Availability of Sensor Data in Agent Cooperation

Christian Johansson, Fredrik Wernstedt, Paul Davidsson

Abstract Multi-agent cooperation can in several cases be used in order to mitigate problems relating to task sharing within physical processes. In this paper we apply agent based solutions to a class of problems defined by their property of being predictable from a macroscopic perspective while being highly stochastic when viewed at a microscopic level. These characteristic properties can be found in several industrial processes and applications, e.g. within the energy market where the production and distribution of electricity follow this pattern. We evaluate and compare the performance of the agent system in three different scenarios, and for each such scenario it is shown to what degree the optimization system is dependent on the level of availability of sensor data.

1 Introduction

Schemes for sustaining cooperative behavior among agents are often dependent on a certain level of communication in order to establish and maintain a reciprocal sense of trust. However, in real-life applications it is not always possible to uphold the desired level of availability and quality of data being communicated among the agents, thus causing suboptimal cooperative behavior.

Information sharing, i.e. communication and its effect on overall performance is a well established area and has been studied by several researchers [5, 6, 13]. Also, the area of multi-sensor networks and sensor data quality and fusion has received a

---

Christian Johansson
Blekinge Institute of Technology, 372 25, Ronneby, Sweden e-mail: christian.johansson@bth.se

Fredrik Wernstedt
Blekinge Institute of Technology, 372 25, Ronneby, Sweden e-mail: fredrik.wernstedt@bth.se

Paul Davidsson
Blekinge Institute of Technology, 372 25, Ronneby, Sweden e-mail: paul.davidsson@bth.se

fair amount of interest [3, 9, 7]. However, the quality of information in combination with information sharing has so far, to our knowledge, only received little attention.

The problem domain is characterised by being predictable from a macroscopic perspective while being stochastic when viewed at a microscopic level. As the macroscopic behaviour is a reflection of a collection of highly stochastic microscopic events which in themselves cannot be predicted, it follows that although a process control system is able to foresee general trends and tendencies within the process, it must be able to handle the stochastic behaviour in order to actually manipulate the process.

When optimizing the operational production one tries to determine the financially and operationally most efficient way to combine the production resources, while satisfying consumer needs. This problem is often formalized by using the Economic Dispatch Problem (EDP) and the Unit Commitment Problem (UCP) [4]. By solving the EDP we find out how much load to generate by each of the different available production units at a given time, while the solving the UCP shows when to start and how long each unit should be committed to being in use.

The consumption, and thus the production, follow certain patterns which are predictable to some extent from a system wide perspective. These patterns are generated by a composition of highly stochastic microscopic behaviour among consumer entities, which, as long as their demand is fulfilled, are oblivious to their surroundings or any other part of the larger system. By reacting on these individual microscopic events and controlling and limiting the effect of them, the overall system can achieve several benefits for both the consumers and the suppliers of the utility. Trying to control the consumption in such a way is generally called Demand Side Management (DSM), and can in many cases be achieved by using agent technology or other distributed control schemes [14, 11, 15].

The problem is that the agent based solutions proposed for solving DSM in such environments are dependent on the availability of high-quality sensor data, which in practice can be hard to achieve due to limitations in underlying hardware and communication solutions. By using different levels of availability and quality of communicated sensor data among the agent system we try to quantify the impact on overall system performance.

2 The Agent System

The agent system we study in this paper is used to implement DSM strategies within district heating systems and its function has been described in previous work [14]. The agent system is based on distributed cooperative entities with an overall goal of combining the production and consumption in an optimal manner.

Every producer and consumer entity in the system is represented by an agent. A producer agent will try to minimize its own supply cost function while supplying enough utility to satisfy consumer demand. When a producer agent deems it necessary to implement an DSM action it will try to do so by sharing the task among the consumer agents in order to minimize the side effects of DSM on any individual consumer agent. A consumer agent will seek to implement these requests as long as its internal comfort constraints allow for this. The producer agent is responsible for supervising the continuous utility consumption and also for instigating and distributing DSM tasks when the measured consumption deviates from the desired DSM level. The task sharing is done by first decomposing the initial task into smaller tasks. This is done since the optimization action as a whole is usually too large for one single consumer agent to handle. The tasks are then allocated through a series of auctions. The DSM level is found beforehand by solving the optimization problem relating to the production units, and this is then used as input to the production agent. The producer agent needs to know the wanted consumption level in order to implement DSM. This is found by solving the EDP and the UCP. These solutions are then used as decision basis for the DSM strategy for the following time frame, normally the next twenty-four hour period. In order to solve the EDP the agent uses an objective function which is found in the smooth function described in Equation 1 and 2.

$$\text{Minimize} \sum_{i \in I} F_i(P_i) \tag{1}$$

$$F_i(P_i) = \alpha_i + \beta_i P_i + \gamma P_i \tag{2}$$

This is simply a summation of the utility cost in all supply units [1]. The value of α describes a fixed cost for starting and running the production unit, while the values of β and γ describe costs dependant on the level of production. The accompanying equality constraint is the utility balance which should be satisfied accordingly:

$$\sum_{i \in I} P_i = D + P_{loss} \tag{3}$$

where D represent the utility demand and P_{loss} indicates any production and distribution losses. The inequality constraints describes the production units working within their respective limits:

$$P_i, min \leq P_i \leq, max \qquad\qquad \forall i \in I \tag{4}$$

In practical settings these functions are normally not sufficient to describe many situations in utility production. Normally the production entity will have to treat the cost function as a set of piecewise quadratic functions which are defined as [8, 10]:

$$F_i(P_i) = \begin{cases} \alpha_{i1} + \beta_{i1} P_i + \gamma_{i1} P_i & if P_i^{min} \leq P_i \leq P_{i1} \\ \alpha_{i2} + \beta_{i2} P_i + \gamma_{i2} P_i & if P_{i2} \leq P_i \leq P_{i2} \\ . & . \\ . & . \\ . & . \\ \alpha_{im} + \beta_{im} P_i + \gamma_{im} P_i & if P_{im} - 1 \leq P_i \leq P_i^{max} \end{cases} \tag{5}$$

This behaviour is due to the fact that a utility provider usually has a range of different production units, using differently priced fuels. From a economical point of view there is no smooth transition when switching between the different fuels, which makes the resulting function non-differentiable. As demand rises the producing entity is forced to engage increasingly costly production units, and eventually the production costs exceed the possible sale price of the utility. The only way for the producer to mitigate such a situation is to manipulate consumption in order to lower the demand. The UCP is interconnected with the EDP and uses similar optimization methods.

Each consumer unit is controlled by a consumer agent which is responsible for contributing to achieving the overall DSM strategies while maintaining a sufficient level of local comfort. The amount of deviation from the optimal comfort state is used as currency when a consumer agent participates in an auction process, i.e the more the consumer agent is straying from its desired comfort state, the less likely it will be to win any auction. The consumer agents are cooperative in the sense that they do not lie about their cost for participating in a DSM task, since this could possibly jeopardize their internal comfort levels.

The goal for the agent system is then; for each point in time achieve a total actual consumption as close as possible to the total wanted consumption while keeping all local comfort levels within their individual constraints. In a steady state system this could be seen as a traditional optimization problem, i.e. to find a optimum between two conflicting objective functions. However, since we are dealing with a dynamic system the aspects of adaptation and re-planning becomes important, which requires a more sophisticated solution. Whenever a producer agent needs to implement a DSM action it will distribute this by using a private value first priced, sealed bid auction process. This type of auction based multi agent system has previously been successfully implemented in district heating networks in order to achieve DSM [15]. Strategic decisions are made based on global or local views within the environment, and the specific optimization actions rely on continuous sensor data.

In this study we compare the performance of a fully functional agent system with two other systems displaying increasingly worse availability of sensor data. These three different scenarios are based on the level of system wide knowledge available to the participating agents; global, partial and local. We choose to compare these specific three levels of system wide knowledge because they correspond to infrastructural prerequisites which can normally be found in actual physical systems, and because they display a broad and clear view of the problem discussed.

Global knowledge is the normal operational scenario for the MAS used to operate the DSM strategies. The producer agents are able to continuously supervise the use of production utility and are able to instigate system wide auctions as need arises. The consumer agents are able to uphold their individual QoS by deciding when and how to participate in these auctions, i.e. a DSM task is never forced upon a consumer agent against its will. Partial knowledge means that the producer agents are able to supervise the consumption of production utility, but they are not able to communicate local sensory data with consumer agents or to uphold cooperative behaviour through auctions. A producer agent is, however, still able to instigate uninformed DSM actions. This is normally done by using predefined activation lists, which try to force consumer agents to implement DSM tasks. The local consumer agents might however decide to reject the appointed task without being able to tell the producer about this. In the local scenario the producer agents have little or no knowledge about the continuous consumption of production utility, and they do not have any possibility at all to implement any DSM actions, either by cooperation or force. In such a system the consumer agents are often assigned the task of keeping the local utility use to a minimum while upholding the desired QoS. Depending on the situation such behaviour might or might not be for the good of the global system state, but the consumer agent will never know anything about this.

3 The Experiment

The experiment is based on operational data from an agent based control system operational in a district heating network in the town of Karlshamn in the south of Sweden [14, 15]. This data is used as input when simulating the various scenarios described in the previous sections. District heating networks are good examples of the described problem domain as they display most, if not all, of the mentioned characteristics. The reference data in question is collected during a twenty-four hour period with no DSM strategy active, i.e. no external control is applied to the consumers.

The consumer agents all have different comfort constraints based on a function of size, shape and material of the individual building, i.e. the amount of thermal buffer available [12]. In the operational system each consumer agent has access to sensor and actuator data through an I/O hardware platform, which enables the agent to measure the physical behaviour of the heating system within the building as well as the outdoor temperature.

Each agent has a value of wanted indoor climate, and constantly tries to minimize all deviation from this value. The consumer agent has two basic values to consider, the comfort level and the thermal buffer level. It is possible to adjust the energy buffer during shorter periods of time without the comfort level having the time to react. When a consumer agent responds to an auction it will use its currently available buffer level as the price it is willing pay for implementing a single DSM task. We evaluate the performance of the consumer agents by measuring how they choose to use their individual buffers.

The optimization strategy used in this experiment is that of peak shedding, i.e. at any given moment when the total energy use exceeds a certain threshold the producer agent will try to convince the consumer agents to lower their local energy usage in a coordinated fashion. The success of the system wide optimization is measured by the amount of deviation between the wanted threshold and the resulting actual level.

We use real operational data from the Karlshamn district heating network as input into the simulation model, where actual flow data is used as initial values for the calculations. The implemented agent system is functioning according to the same principles as previously described. In the simulation there are fourteen active agents; one producer agent and thirteen consumer agents. By simulating the described levels of agent knowledge we can evaluate the performance of the agent system during different scenarios.

A simulation run begins by calculating specific solutions to the EDP and the UCP. These solutions yield a wanted system wide consumption level for each time step throughout the day. This wanted consumption level is then used by the producer agent as a decision basis, when deciding when and how to instigate DSM actions throughout each time step. This buffer levels are then adjusted through each time step as the agents perform DSM tasks, which in turn makes it possible to calculate the comfort levels for each time step.

4 Results

The control strategy is evaluated by measuring the flow of hot water into the area. Energy usage in a district heating network is measured by combining the temperature of the water with the flow. Since the supply water temperature in the primary network is more or less stable throughout a single day the flow in itself gives a good estimation of the energy usage within all the buildings. In Figures 1, 2 and 3 we show the flow data achieved during the three different scenarios in relation to the wanted DSM strategy. The straight dashed line is the wanted DSM level. This level of consumption is based on a solution of the Economic Dispatch Problem and the Unit Commitment Problem. The peaks above the dashed line represents peak loads which would need to be satisfied by using financially and environmentally unsound fossil fuel. In other words, the global goal of the agent system is to keep the consumption as close to the straight dashed line as possible.

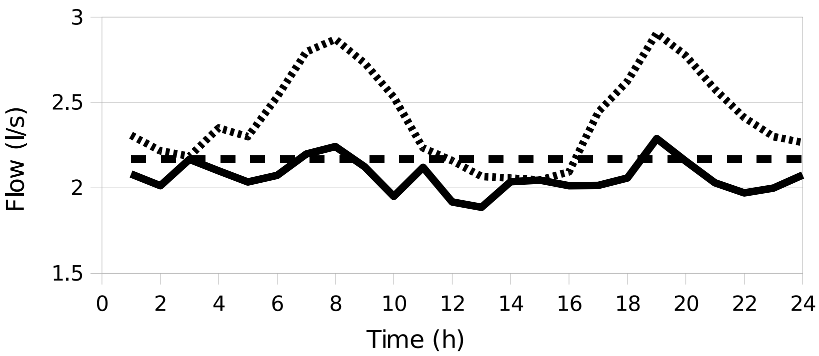

Fig. 1 Global scenario. Agent performance (continuous), reference data (dotted) and wanted DSM level (dashed)

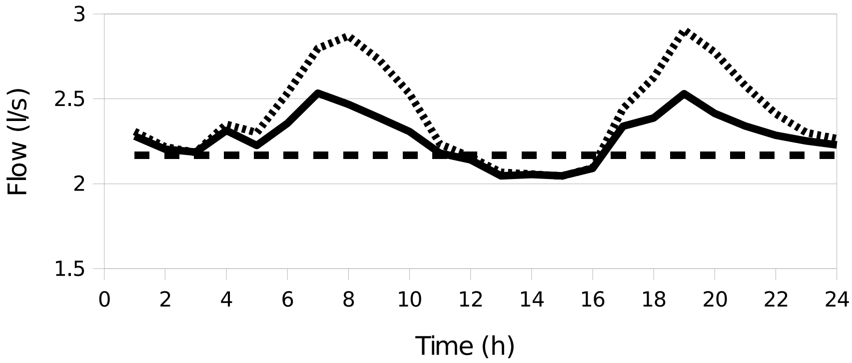

Fig. 2 Partial scenario. Agent performance (continuous), reference data (dotted) and wanted DSM level (dashed)

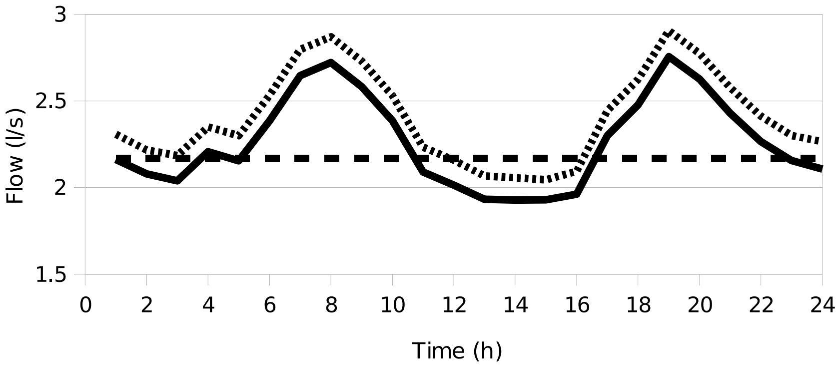

Fig. 3 Local scenario. Agent performance (continuous), reference data (dotted) and wanted DSM level (dashed)

It is clearly visible that the flow value in the global scenario, Figure 1, most closely resembles the desired DSM strategy, with the partial scenario, Figure 2, being somewhat worse, and finally the local scenario, Figure 3, showing a distinct lack in ability to achieve the desired level of consumption.

Every agent has an maximum allowed buffer usage of one, with a minimum of zero. The level of comfort will not be negatively effected by a usage between one and zero. If the buffer usage is above one the consumer agent has used more than the allowed buffer and the comfort can be in jeopardy if such a status is allowed to continue for a longer period of time. In other words a consumer agent has an optimal buffer usage of one, i.e. the agent participates in achieving the global goal as much as possible but does this without sacrificing its desired comfort level.

Figure 4 shows the dynamic system wide buffer usage during the whole time period. The range on the y axis is dependent on the amount of consumer agents, since every such agent has a optimal buffer usage of one. In this case study we have thirteen agents, so an optimal usage of the system wide buffer would be thirteen. In the global and partial scenarios the buffer usage clearly follows the reference data as the agents continuously try to counter the varying consumption.

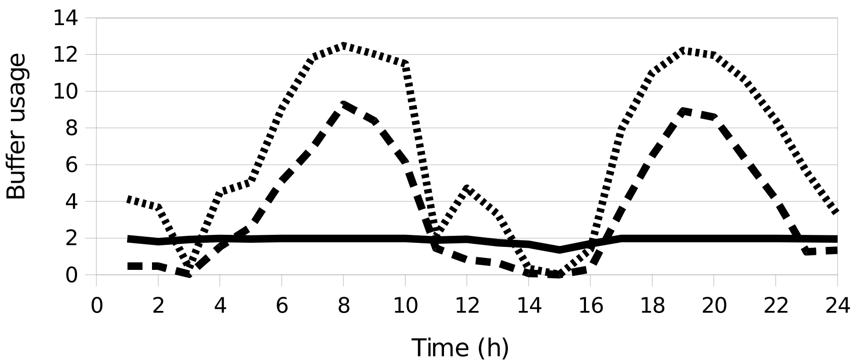

Fig. 4 Buffer usage. Global scenario (dotted), Partial scenario (dashed) and Local scenario (continuous)

5 Conclusions

Multi-agent system solutions being applied to the physical processes described in this paper are heavily dependent on the availability of high-quality sensor data to function properly. This study quantifies the way system performance rapidly deteriorates as the availability of high-quality sensor data is reduced. It is important to factor in both the DSM strategy and the consumer agent comfort value when evaluating an implementation for handling DSM within the problem domain. If a system is only evaluated on the basis of its ability to adhere to the DSM strategy it might give rise to problems on the consumer side as no consideration is given to upholding a sufficient level of QoS.

The local scenario is similar to a type of control system that is often implemented in both electrical grids and district heating networks, as a local uninformed optimization technique. This study indicates that such systems have little global effect in regards to overall production optimization strategies. The reason that the local scenario never goes beyond a certain level in Figure 4 is that the consumer agents are only reacting to their own local peak loads, which are well beyond their own capacity to handle. This is due to the fact that individual peaks are much larger than any individual buffer, so in the local scenario some agents are always maximizing their use of their individual buffer, but without the ability to somehow distribute the load through the producer agents their efforts will always fall short on a system wide scale.

Figure 4 also shows that producer agent knowledge is needed in order to dynamically counter the user demand in regards to the DSM strategy. This is also the buffer usage, which shows that the partial scenario is not able to fully use the available buffer. This is due to the fact that the agents cannot perform cooperative work. The lower use of available buffer of the partial scenario is caused by the fact that although the consumer agent is handed a DSM task, it can choose not to implement the task if the agent considers it to jeopardize its internal QoS level. Since the producer agent never receives any feedback about this, it will not be able to distribute the task to another consumer better suited for the task, and hence the system will on average not utilize the maximum available buffer.

Figure 4 shows that the global scenario is close to using the maximum available buffer on several occasions, while neither the partial or the local scenarios are close to utilizing their full DSM potential.

In this paper we have shown that distributed multi agent systems based on cooperative auctioning are able to achieve the studied DSM strategy, while maintaining an acceptable level of QoS. As the availability and quality of the sensor data diminishes the system performance deteriorates, first into the equivalence of static distributed models and then into the equivalence of simple local optimization models.

This paper is the result of an initial case study in regards to sensor data utilization within industrial multi-agent system applications. In the future we will use this as groundwork while incorporating the financial factors underlying the discussion, in order to further study the economical effects found within such systems.

Acknowledgements The operational data used in this study was supplied by NODA Intelligent Systems. The project has also been supported by Karlshamns Energi and Karlshamnsbostäder.

References

1. Arvastsson, L.: Stochastic Modeling and Operational Optimization in District Heating Systems. PhD Thesis, Lund Institute of Technology (2001)
2. Aune, M.: Energy Technology and Everyday Life - The domestication of Ebox in Norwegian households. In Proceedings of ECEEE Summer Study, 2001
3. Dash, R., Rogers, A., Reece. S., Roberts, S. and Jennings, N.R.: Constrained Bandwidth Allocation in Multi-Sensor Information Fusion: A Mechanism Design Approach. in Proceedings of The Eight International Conference on Information Fusion, Philadelphia, PA, USA, 2005
4. Dotzauer, E.: Simple Model for Prediction of Loads in District Heating Systems. Applied Energy, Vol 73, no 3-4, pp. 277-284, 2002
5. Dutta, P.S., Goldman, C. and Jennings, N.R.: Communicating Effectively in Resource Constrained Multi-Agent Systems. in 20th International Joint Conference on Artificial Intelligence (IJCAI), Hyderabad, India, 2007
6. Goldman, C.V. And Zilberstein, S.: Optimizing Information Exchange in Cooperative Multi-Agent Systems. in Proceedings of Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 137-144, 2003

7. Jayasima, D.N.: Fault Tolerance in Multisensor Networks. IEEE Transactions on Reliability, vol. 45, no. 2, pp. 308-315, 1996
8. Koay, C. A., Lai, L.L., Lee, K.Y., Lu, H., Park, J.B., Song, Y.H., Srinivasan, D., Vlachogiannis, J.G. and Yu, I.K.: Applications to Power System Scheduling. Modern Heuristic Optimization Techniques, Edited by K. Y. Lee and M. A. El-Sharkawi. The Institute of Electrical and Electronics Engineers, Inc, (2008)
9. Lesser, V., Ortiz, C. and Tambe, M., Eds.: Distributed Sensor Networks: a multiagent perspective. Kluwer Publishing, (2003)
10. Lin C.E., Viviani G.L.: Hierarchical Economic Dispatch for Piecewise Quadratic Cost Functions. IEEE Trans Power Apparatus Systems, 103(6):1170-1175, 1984
11. Nordvik, H. and Lund, P.E.: How to Achieve Energy Efficient Actions as an Alternative to Grid Reinforcement. in Proceedings of ECEEE Summer Study, 2003
12. Olsson Ingvarsson, L.C., Werner, S.: Building Mass Used as Short Term Heat Storage. in Proceedings of The 11th International Symposium on District Heating and Cooling, Reykjavik, Iceland, 2008
13. Shen, J., Lesser, V. and Carver, N.: Minimizing Communication Cost in a Distributed Bayesian Network Using a Decentralised MDP. in Proceedings of Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 678-685, 2003
14. Wernstedt, F., Davidsson, P. and Johansson, C.: Demand Side Management in District Heating Systems. in Proceedings of Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Honolulu, Hawaii, 2007
15. Wernstedt, F., Johansson, C.: Intelligent Distributed Load Control. in Proceedings of The 11th International Symposium on District Heating and Cooling, Reykjavik, Iceland, 2008