UDC 004.75, DOI:10.2298/CSIS100118011J

A Case Study on Availability of Sensor Data in Agent Cooperation

Christian Johansson, Fredrik Wernstedt, Paul Davidsson

School of Computing, Blekinge Institute of Technology
PO Box 520, SE-372 25, Ronneby, Sweden
christian.johansson@bth.se
fredrik.wernstedt@bth.se
paul.davidsson@bth.se

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. Another defining problem characteristic is that the supply is usually limited as well as consisting of several layers of differentiating production costs. 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.

Keywords: agent co-operation, team work

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. In this paper we focus on a problem domain where multi-agent task sharing cooperative behavior is applied. However, as practical implementations within this domain often are spread geographically over a wide area and lack dedicated network communication infrastructure, there are often practical limitations on the availability and quality of sensor data which in turn limits the effectiveness of the multi-agent system cooperative behavior.

For agents to effectively coordinate their actions, the agents normally need to share information. Information sharing, i.e. communication and its effect on overall performance is a well established area and has been studied by several researchers [6], [7] and [13]. Also,the area of multi-sensor networks and sensor data quality and fusion has received a fair amount of interest [4], [10] and [8]. However, the quality of information in combination with information sharing has so far, to our knowledge, only received little attention.

1.1. Problem Domain

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. For example, although it is possible to foresee the overall heating demand within a building being higher tomorrow as the weather prognosis shows a drop in outdoor temperature, it is not possible to predict when individual residents will take a shower, thus creating peaks in the total energy demand when combining usage of hot tap water and space heating. Basically these processes are driven by one or more producers supplying some kind of utility and one or more consumers acting to satisfy their own demand of the utility.

When optimizing the operational production one tries to determine the financially and operationally most efficient way to combine the production resources, while satisfying the consumer needs. This problem is often formalized by using the Economic Dispatch Problem (EDP) and the Unit Commitment Problem (UCP) [5]. By solving the EDP we find out how much load to generate by each of the different available production units at a given time. Since most production units in real life settings cannot be turned on and off at the blink of an eye, it is important to plan ahead of time and determine which units need to be started, when they need to be started and how long they should be committed to being in use, i.e. solving the UCP. These problems are related to each other and are solved using similar optimization techniques. A complicating factor when optimizing production is that the production costs usually display non-linear patterns, due to physical processes like valve effects and the usage of differently priced fuels in production. This leads to objective functions with discontinuous and non-differentiable points, which means that it is generally more appropriate to treat the cost function as a set of piecewise quadratic functions [9], [3]. 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.

By solving the UCP and EDP the producer finds an optimal production strategy for a given time frame, e.g. the next twenty-four hours. This means that the producer wants the consumption to be as close to this strategy as possible; if the consumption falls to low it will result in unnecessarily low income, while a too high consumption will necessitate starting costly peak-load production units. In order to achieve and uphold the desired production strategy multi agent systems and other distributed systems can be used to manage the consumption. We evaluate the success of such systems by measuring their ability to stick to the production strategy in question, while at the same time satisfying consumer demand.

Typically the consumer entity has a wanted state which it tries to uphold at all times. This wanted state is dependent on the physical environment in which the system is functioning, e.g. maintaining comfortable indoor climate in a district heating system. Often, however, a consumer agent can accept smaller deviations from this wanted state during shorter periods of time. This is what makes it possible for a control system to manage the society of consumers in order to achieve some local or global goals. By measuring the deviation from the wanted state it is possible to evaluate the impact of change in the overall system caused by individual consumers.

1.2. Problem Description

The consumption, and thus the production, follows 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], [2] and [11]. The reason agent technology is useful in DSM, is that there is no need for any centralized entity supervising the Quality of Service (QoS) among the consumers as each consumer is assigned an agent responsible for this task. Each agent will participate in achieving the overall goal set by the DSM strategy, only as long as sufficient QoS can be upheld. This makes the system highly scalable and easy to maintain.

In theory this a school book example for an agent system to solve. 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. That an agent system will perform at its best in a domain were high quality sensor data and communication solutions are readily available is not in question, and it is not the intent of this paper to compare different agent based resource allocation schemes within such a high quality domain. The point of this paper is instead to develop an understanding of how different levels of availability of sensor data influence the behaviour of the agent system in a practical setting. Normally there are practical limitations on the sensor data infrastructure and communication set-up which leads to situations far from any high quality scenario. Investing in modern sensor data and communication solutions can be expensive and there is an apparent need to quantify the performance benefits within different scenarios. In Figure 1 this is visualized.

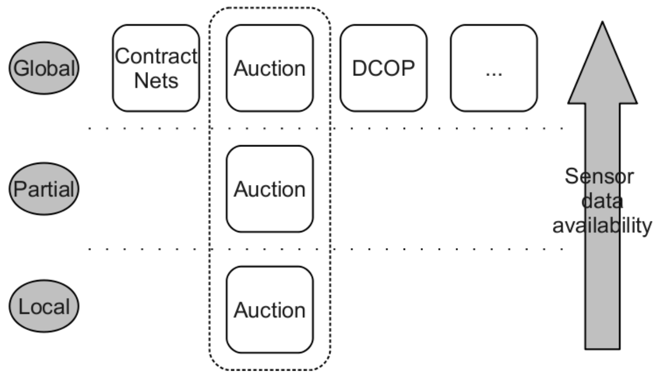

Fig. 1. Comparing availability of sensor data using three different scenarios

Within this study three different scenarios are used to represent different levels of availability of sensor data, i.e global, partial and local. The global level corresponds to a scenario with full access to high quality sensor data while the partial and local scenarios display various levels of deteriorating access to sensor data. There are several ways to coordinate resource allocation within a multi agent system, e.g. contract nets, different auction processes or distributed optimization models. In this study we have used an auction process in order to compare the different scenarios according to Figure 1.

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]. In district heating systems one or several production plants heat water which is then pumped through a pipe network throughout a city in order to heat buildings and tap water. Every building has a substation with heat exchangers which transfer the heat energy from the primary pipe network into the buildings secondary piping system. The agent system is based on distributed cooperative entities with an overall goal of combining the production and consumption in an optimal manner.

2.1. Agents

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.

Producer Agent 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. In larger agent systems it is possible to use cluster agents which act as mediators between a producer agent and a group of consumer agents. This eases the computational load in the producer agent when handling large scale auctions.

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 Equations 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:

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

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 [9], [3]:

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

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.

The producer also has to solve the UCP. The UCP is interconnected with the EDP and uses similar optimization methods. The UCP is used to determine which production units to commit to usage and which ones not to use at any given time. In a real world scenario a production unit cannot be turned on and off with a simple switch. It takes a substantial amount of time to start and stop such units, and the cost related to these processes should be kept at a minimum.

By solving the above systems for each relevant point in time it is possible to identify a wanted system wide consumption level within the studied time frame.

Consumer Agent 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 consumer agents act locally in order to monitor any deviations from the wanted comfort state. 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. A positive side effect from using the comfort state as currency, is that these calculations are made by the consumer agent in any case and thus the computational effort for valuation and information gathering in regards to the auctioning is kept at a minimum.

2.2. Agent Goal

For every consumer agent there is at any time a wanted comfort level which is dependent on the level of local consumption. Since the physical nature of the process introduces a delay in the dependency between the comfort level and the local consumption level a time frame is created within which it is possible to manipulate the local consumption while still keeping the local comfort level within its constraints. An example of this phenomena is that it will take some time before people notice if you shut off the radiators in a building, i.e. there is a delay before the people start to freeze even though the energy consumption is reduced directly. Combining the local consumption from each consumer agent will yield the total actual consumption. The goal for the agent system is then; for each point in time to 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.

2.3. Auction Process

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. For the consumer agent the value is to implement as much DSM tasks as possible, and the currency used is the amount of deviation from the optimal comfort state possible without affecting the local QoS. 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. Global knowledge is needed in order to identify individual consumer agents able and willing to participate in local task accomplishment. Without sufficient communication abilities the auction process is not able to function, thus making it more difficult to distribute DSM tasks while taking into account the local consumer agent comfort state. By using an action process it is possible to distribute the complexity and computational effort, since all reasoning and planning about the delivered QoS is done by the consumer agents. This leads to a more scalable and extendable system.

2.4. Levels of Agent Knowledge

In the described DSM system the agents are able to communicate freely among their peers, in order to continuously propagate system status based on available sensor data and perform coordinated task sharing when needed. In this study we compare the performance of such a fully functional 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 This is the normal operational view 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 DSM actions as need arises. Each DSM action is divided into control tasks which can be distributed throughout the network of consumer agents by system wide auctions. 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 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. This means that cooperative behaviour through auctioning is not possible. A producer agent is, however, still able to instigate uninformed DSM actions. This is normally done by using predefined activation lists, which force consumer agents to implement DSM tasks in a round-robin fashion. Since no communication of consumer sensor data is available it is not possible for the producer agents to gain any feedback about the impact of these DSM tasks on the QoS delivered to the local consumer. The local consumer agent is still working to uphold its own QoS, and it might decide not to implement the DSM task appointed to it. Either way, as the consumer sensor data communication is impaired, the producer agent will never gain any knowledge about what decision the consumer agent takes.

Local Knowledge In this 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. The consumer agents still have access to their own local sensor data but they cannot successfully communicate this to other agents within the MAS. This basically means that they act oblivious to the state of any other agent. 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

In this study we investigate the effects of different levels of availability of sensor data within an operational agent based control system. Under normal circumstances the agent system is based on cooperative behaviour which is in turn heavily dependent on functioning and reliable communication of high quality sensor data. Performance of the multi-agent system will deteriorate if the availability or quality of sensor data declines. We have studied how extensive this deterioration will be in a practical setting, i.e. how will the quality of the communication underlying the decision-making affect the overall performance from a system wide perspective. The case study 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.

3.1. Reference Data

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 data is representative of normal usage during wintertime when the heating demand is substantial. The energy consumption in a district heating system is measured by combining the flow of the water with the primary supply temperature of the water. During the course of a single day the primary supply temperature in this district heating network is rather constant so the flow is a good estimate of the total energy use.

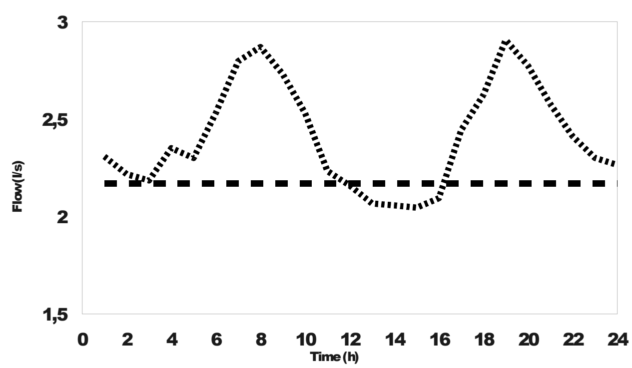

Fig. 2. Reference data (dotted) and wanted DSM level (dashed)

Figure 2 shows the flow data from the Karlshamn district heating network during a full twenty-four hour time period. The straight dashed line shows the calculated wanted level of energy consumption which the producer agent uses as a benchmark during this specific day. 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.

3.2. Utility Evalution

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. Based on this data the agent can calculate the indoor temperature which is then used as the basis for the agents own comfort evaluation. Each agent has a value of wanted indoor climate, and constantly tries to minimize all deviation from this value. However, in order to participate in achieving the global system goal an individual consumer agent can accept smaller deviations from this value under shorter periods of time, as this will not affect the indoor climate to a degree where the inhabitants will notice it. The consumer agent has two basic values to consider, namely the comfort level and the buffer level. These are connected with a delay, so that it is possible to adjust the level of the energy buffer during shorter periods of time without the comfort level having the time to react. It is possible for the consumer agent to use more than the available buffer, but then the comfort level will start to fall. Under normal circumstances a consumer agent will be very unwilling to use more the available buffer, although it might do so during shorter periods of time in order to achieve some global goal. 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. This process will ensure that only the consumer agent which is best suited at any given time, i.e will be least in risk of compromising its comfort level, will be appointed the DSM task in question. We evaluate the performance of the consumer agents by measuring how they choose to use their individual buffers.

The producer agent in the system use the energy delivered to the area as input for its calculations concerning the need for DSM actions. 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. When implementing this strategy the producer wants to limit the utility consumption down to the wanted threshold, as forcing the consumption down even further than the threshold will reduce sale of utility produced by economically viable production supplies. Therefore we measure the success of these system wide optimization actions by measuring any deviations between the wanted fluid flow value and the resulting actual flow level. By analyzing these values it is possible to evaluate to what extent the overall agent system accomplishes its objective, i.e. to uphold the wanted DSM strategy level without jeopardizing the comfort among the consumer agents. If the agent system is to be considered successful in its endeavors it will have to fulfill both these requirements.

3.3. Availability of Sensor Data

The agents within the Karlshamn district heating area communicate through a LAN network and have direct access to high quality sensor data. In this sense they are extremely spoiled, as this kind of communication solution is rarely part of the hardware set-up in similar real-world environments. As building a physical network and sensory infrastructure can be costly, similar agent systems are usually limited to using communication techniques such as GSM-modems, radio link or standard limited master/slave networks to evaluate operational data. With such solutions there is often limitations in regards to communication bandwidth and temporary sensor breakdowns are not uncommon. In this experiment we evaluate the impact of such system deterioration by simulating different levels of availability of the sensor data. In order to do this we model the three previously described scenarios, i.e. global knowledge, partial knowledge and local knowledge.

In the global scenario the producer agent and the consumer agents are allowed to communicate freely throughout every time step in the simulation. The producer agent can instigate auction processes according to its own desires, and the consumer agents are able to respond to this.

During the partial scenario there is only one-way communication possible from the producer agent to the consumer agents. The producer agent can distribute DSM tasks, and does so according to a previously defined static list. The producer agent can distribute several DSM individual tasks during each time step. A consumer agent might implement such a DSM task or it might not, depending on the current level of its internal buffer. Any which way, the producer agent will not receive any response about this.

In the local scenario there is no communication what so ever between the agents. The consumer agents can perform local load control, but this is done purely based on local knowledge. The local load is made up of a combination of energy used for space heating and tap water heating. During local load control, the consumer agent will try to limit local space heating when there is local tap water usage. The tap water usage is randomized within the simulation.

3.4. Simulation

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 Economic Dispatch Problem and the Unit Commitment Problem. 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. The consumer agents starts the simulation with full available buffer levels. This buffer level is then adjusted through each time step as they perform DSM tasks, which in turn makes it possible to calculate the comfort levels for each time step.

For each time step the producer agent then decides if there is any need for DSM actions based on the current flow level in relation to the wanted flow level. If it deems this necessary it will try to distribute individual DSM tasks. Depending on the specific scenario this will be achieved differently. The system wide consumption is then calculated and used as input into the next time step.

4. Results

We evaluate the different scenarios according to how well they manage to achieve the DSM strategy in question while staying within the comfort constraints set by the consumer agents. We present how well the agent system upholds the DSM strategy within the three different scenarios, and then we show how well the system manages to keep itself within the available energy buffer, and thus the consumer comfort constraints during these same scenarios.

4.1. Control Strategy

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. Figure 2 in the previous section shows the reference data without any DSM strategy active, i.e. this is what the overall consumption pattern looks like in a district heating network without any agent system installed. In Figure 3, Figure 4 and Figure 5 we show the flow data achieved during the three different scenarios in relation to the wanted DSM strategy.

It is clearly visible that the flow value in the global scenario (Figure 3) most closely resembles the desired DSM strategy, with the partial scenario (Figure 4) being somewhat worse, and finally the local scenario (Figure 5) showing a distinct lack in ability to achieve the desired level of consumption. To make these results clearer we also summarize the deviation of the scenarios for every time frame throughout the simulation. This value has a theoretical optimum at zero, i.e no deviation from the desired DSM level what so ever. The values are then normalized around the value achieved by the global scenario. These results are shown graphically in Figure 6 and the actual numbers in Table 1.

4.2. Agent Buffer Usage

The level of comfort is dependent on the buffer used by each individual consumer agent. Every agent has an maximum allowed buffer usage of 1, with a

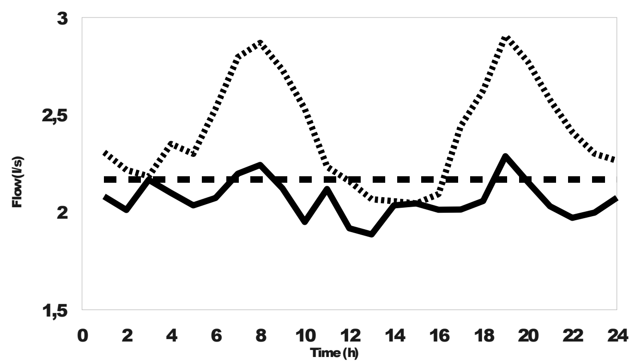

Fig. 3. Global scenario. Agent performance (continuous), reference data (dotted) and wanted DSM level (dashed)

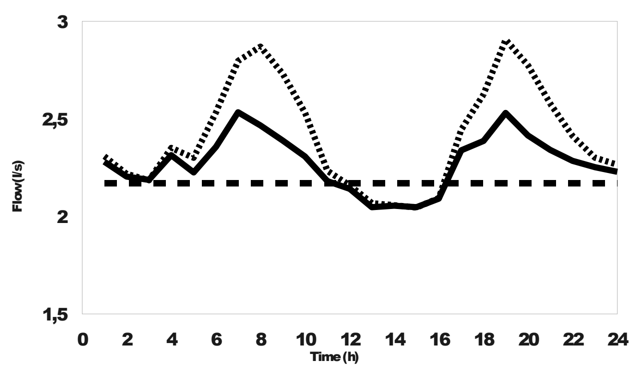

Fig. 4. Partial scenario. Agent performance (continuous), reference data (dotted) and wanted DSM level (dashed)

Table 1. DSM strategy deviation values

Scenario	Value
Global	1
Partial	6.75
Local	11.07

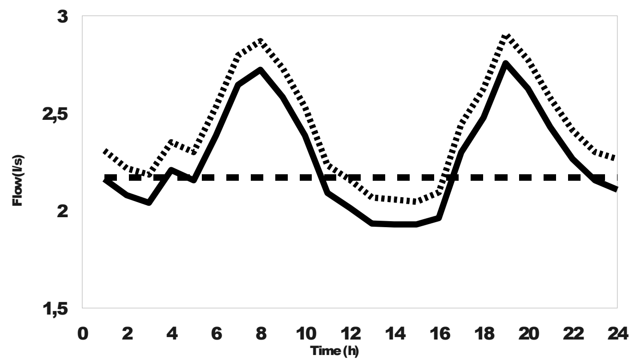

Fig. 5. Local scenario. Agent performance (continuous), reference data (dotted) and wanted DSM level (dashed)

Fig. 6. DSM strategy deviation. Global scenario (black), partial scenario (dark grey) and local scenario (light grey)

minimum of 0. The level of comfort will not be negatively effected by a usage between 1 and 0. If the buffer usage is above 1 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 1, i.e. the agent participates in achieving the global goal as much as possible but does this without sacrificing its desired comfort level. The values for the individual consumer agents are shown in Figure 7, together with a theoretical optimum of 1. The numerical values are then showed in Table 2.

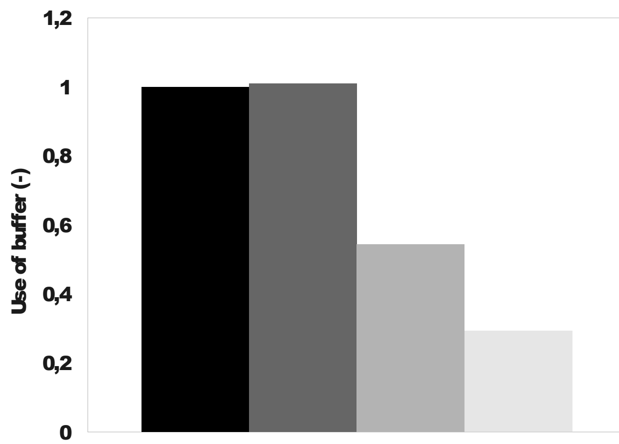

Fig. 7. Individual buffer usage. Theoretical optimum (black), global scenario (dark grey), partial scenario (grey) and local scenario (light grey)

Table 2. Agent comfort and buffer usage for each individual consumer agent

Scenario	Value
Optimum	1
Global	1.01
Partial	0.54
Local	0.29

Figure 8 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 obviously be 13. In the global and partial scenarios the buffer usage clearly follows the reference data as the agents continuously try to counter the varying consumption.

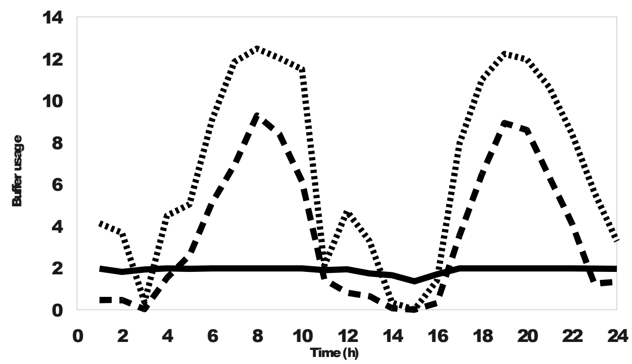

Fig. 8. Total buffer usage. Global scenario (dotted), partial scenario (dashed) and local scenario (continuous)

5. Discussion

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. The evaluation is based on the systems ability to adhere to the wanted control level while maintaining an acceptable level of agent comfort. By combining the control strategy and agent comfort values it is then possible to evaluate the performance.

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 notion of what is an acceptable level of the control strategy value is dependant on the process in question. In a district heating network there are several benefits of having the ability to perform load control, the most apparent being the ability to shed peak loads in order to avoid using expensive and environmentally unsound peak load fuels. In our case the global scenario would be considered acceptable since it manages to shed the peaks.

6. Conclusions

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. As Figure 5 and Figure 8 clearly shows the local scenario is inadequate in order to handle any system wide DSM strategy. The reason that the local scenario never goes beyond a certain level in Figure 8 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. This shows a clear advantage of the two distributed DMS solutions in relation to any local efforts, which can never hope to counter system wide peaks.

Figure 8 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 difference between the global and partial scenarios in Figure 8 basically shows the superiority of agent cooperation versus centralized enforcement. 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 of the available buffer.

Figure 8 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. The global scenario is rather close to achieving the DSM strategy, but it does not manage to fully adhere to the wanted level. To achieve this would require the system to continuously foresee and counter highly stochastic microscopic behaviour within the process, which is not feasible in a practical setting.

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 shows that real-time cooperative behaviour among communicating agent nodes is needed in order to successfully implement DMS in real world applications, and that indirect methods, like differentiable taxation, or uninformed local optimization is not able to produce the coordinated system-wide behaviour required.

7. Future Work

This paper is the result of an 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.

According to Figure 1 we have only used an action process in order to evaluate the the different scenarios. In the future we intend to expand this study by using other collaboration techniques within the different scenarios.

Another issue is the formalization of a model for the follow-up of DSM actions. Sometimes actions are taken when there is no need for them, and other times actions are needed without them being implemented. By improving our understanding and control of this process it should be possible to better utilize the individual consumer agent buffer. This could be used when combining the multi agent system with a continuous optimization model in order to dynamically follow the process.

8. Acknowledgements

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

References

1. Arvastson, L.: Stochastic Modeling and Operational Optimization in District Heating Systems. 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. C.E., L., G.L., V.: Hierarchical economic dispatch for piecewise quadratic cost functions. In: Trans Power Apparatus Systems, vol. 103-6, pp. 1170–1175. The Institute of Electrical and Electronics Engineers, Inc (1984)
4. Dash, R., Rogers, A., S., R., Roberts, S., 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)
5. Dotzauer, E.: Simple model for prediction of loads in district-heating systems. Applied Energy 73(3–4), 277–284 (2002)
6. Dutta, P.S., Goldman, C., Jennings, N.R.: Communicating effectively in resource-constrained multi-agent systems. In: Proceedings of the 20th International Joint Conference on AI (IJCAI). Hyderabad, India (2007)

7. Goldman, C.V., Zilberstein, S.: Optimizing information exchange in cooperative multi-agent systems. In: Proceedings of Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 137–144 (2003)
8. Jayasima, D.N.: Fault tolerance in multisensor networks. IEEE Transactions on Reliability 45(2), 308–315 (1996)
9. Koay, C.A., Lai, L.L., Lee, K., Lu, H., Park, J.B., Song, Y.H., Srinivasan, D., Vlachogiannis, J.G., Yu, I.K.: Applications to power system scheduling. In: Lee, K.Y., El-Sharkawi, M.A. (eds.) Modern Heuristic Optimization Techniques. The Institute of Electrical and Electronics Engineers, Inc (2008)
10. Lesser, V., Ortiz, C., Tambe, M.: Distributed Sensor Networks: a multiagent perspective. Kluwer Publishing (2003)
11. Nordvik, H., 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., 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., 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). pp. 678–685 (2003)
14. Wernstedt, F., Davidsson, P., Johansson, C.: Demand side management in district heating systems. In: Proceedings of Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS). Honolulu, Hawaii, USA (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)

Christian Johansson is a Ph.D. student in Computer Science at Blekinge Institute of Technology (BTH). His main focus is optimization and system-wide coordination in distributed systems with applications in energy systems.

Fredrik Wernstedt has a BSc degree in Computer Science from Blekinge Institute of Technology and a BSc in Nautical Science from University of Kalmar. He received his Licentiate degree in Computer Science from BTH in 2003 and his Ph.D. in Computer Science from BTH in 2005. His main interest is autonomous agents and multi-agent systems, particularly for optimization of operational systems.

Paul Davidsson is a Professor in Computer Science at Blekinge Institute of Technology. Paul is the manager of the Distributed and Intelligent Systems Laboratory at BTH and received his Ph.D. in Computer Science from Lund University 1996. His research interests includes the theory and application of multiagent systems and machine learning.

Received: January 18, 2010; Accepted: May 3, 2010.