Improving Water Quality by Coordinating Industries Schedules and Treatment Plants

Víctor Muñoz, Javier Murillo, Dídac Busquets and Beatriz López
Institut d'Informàtica i Aplicacions
Universitat de Girona
Campus Montilivi
17071 Girona, Spain
{vmunozs,jmurillo,busquets,blopez}@eia.udg.es

ABSTRACT

Having a proper waste water treatment system is crucial for making a good use of water resources. Current regulations enforce some restrictions to the industries producing waste, according to the capacities of waste water treatment plants. However, these are usually not sufficient to ensure that these capacities are not exceeded. In this paper we present a coordination system that provides a more integrated view of the problem, taking into account all the elements involved in the treatment system. The goal of the system is to coordinate the individual industries' discharge schedules over time to help the treatment plant in the cleaning process of the water. The system we propose is based on an auction mechanism, in which the industries can bid for the right to perform a discharge. We have extended it with a priority mechanism in order to provide a fair solution to the problem. A prototype of the system has been implemented and tested on simulation, obtaining successful results.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence-Multiagent Systems, Coherence and coordination

General Terms

Algorithms

Keywords

Schedule Coordination, Auctions, Multi-agent Systems, River-basin Management

1. INTRODUCTION

Water is a vital natural resource, not only for urban and industrial consumption, but also as the main element to maintain any natural environment. Thus, the need of having a good treatment system is basic in order to account for the high demand it suffers. Moreover, the increase in population growth and industrial activity results in the harmful effect of having more contaminated waters, therefore making the treatment task harder.

The most important element in the treatment system is the *Waste Water Treatment Plant* (WWTP). Its job is to remove contaminants from sewage and produce an (up to a certain degree) clean waterstream that can be put back into the river. In order to ensure that the treatment process is correctly performed two conditions must hold:

- Keep the *incoming water flow* below the WWTP hydraulic capacity (that is, the amount of water the plant can absorb at any instant in time); otherwise, the overflown water goes directly back to the river without receiving any treatment, increasing its contamination level.

- Keep the *contamination level* of the incoming water below the WWTP treatment capacity. The contamination level is defined by a set of quality variables (oxygen demand, nitrogen level, etc.). If the level of any of these variables is above the WWTP capacity, the water cannot be fully treated, and it increases the contamination of the river. Moreover, if the levels were too high, the microorganisms used to treat the water may be damaged and the whole process could be stopped until these were regenerated. During this time, the plant could not accept any incoming water and it would be redirected to the river without any treatment.

These capacities, or thresholds, are usually referred to as the WWTP's *design parameters* and depends on the amount of water to be treated in relationship with the surrounding industries and cities.

The water entering the WWTP comes from three different sources: domestic use, rainfall and industries. Current regulations and legislations are in place so as to minimize the contaminating effects of industrial waste discharges. However these are not sufficient to guarantee the proper treatment of the water. The problem is that, although these regulations enforce industries to respect the WWTP capacity thresholds, they do not take into account that simultaneous discharges by different industries may exceed these thresholds. In such a case, no industry would be breaking the rules, but the effect would be to have overflow or overcontaminated water going to the WWTP.

CAPS AAMAS'07 Workshop, Honolulu, Hawaii, USA

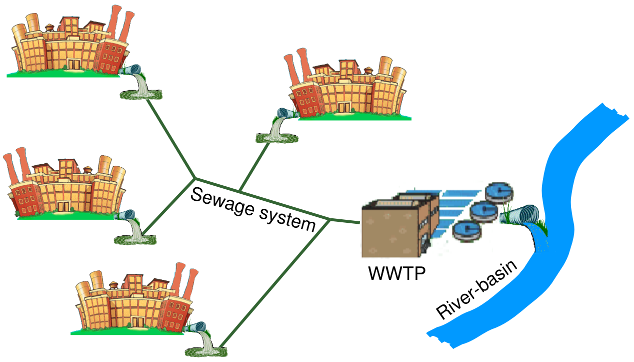

Figure 1: Water treatment system

Thus, in order to ensure that the thresholds are not exceeded, some coordination between the industries and the WWTP is needed. Through this coordination, the different discharges would be temporally distributed so that the WWTP is capable of processing all the incoming water. This would be beneficial at an environmental and health level, and is in the direction of having a more integrated management of the river-basin and all the involved systems. Moreover, it could also go together with economic incentives to those industries collaborating the most (such as discounts in the discharge fees).

In this paper we propose a coordination system to mediate between the different industries willing to discharge waste at the same period of time. More specifically, we have developed an auction mechanism to coordinate the individual industries' discharge schedules. The paper is organized as follows. In Section 2 we present the coordination system, describing in detail each of the steps involved. The implementation of a first prototype is described in Section 3. In Section 4 we discuss the experimental results obtained through simulation. Some related work is presented in Section 5, and Section 6 concludes the paper.

2. COORDINATION SYSTEM

A typical water treatment system is depicted in Figure 1. The industries discharge their wastes to a sewage system, which directs the water to the WWTP. The plant, once the water has been treated, puts it back to the river. The main goal of our system is to ensure that the water flow entering the WWTP and its contamination levels are below some given thresholds, so that it can be correctly treated. As mentioned previously, we propose to achieve this goal by coordinating the discharges performed by the industries. In this section we first describe the assumptions made about the industries performing the discharges, and then we present how the coordination system works.

2.1 Industry Model

Although there are two other sources of waste water (rainfall and domestic use), in this first approach to the problem we focus our attention to industries.

We assume that the industries have some kind of working plan that allows them to foresee what discharges will be necessary in the near future, according to their production strategy. This knowledge would permit the industries to inform the WWTP about the characteristics of their discharges (starting time, duration, flow and contaminants levels), so that coordination can be achieved.

```
while not all discharges authorized do
    receive schedules from industries
    loop through discharges
        if no conflict then
            authorize current discharges
        else
            solve conflict
            inform industries about resolution
            break loop
        end if
    end loop
end while
```

Figure 2: WWTP agent pseudo-code

We also assume that each industry has a tank where it can store its waste in case a discharge is not authorized. Without this tank, an industry would be forced to perform a discharge in the case it were not authorized, making the whole coordination process useless. Obviously, if the industry is denied to discharge and its tank is full, it will be forced to realize the discharge anyway. This could seriously affect the process in the WWTP, and therefore should be avoided at all costs. The coordination mechanism that we propose tackles this problem by distributing the authorizations among the industries, trying to avoid that any of them has to perform unauthorized discharges.

An industry may also reschedule its discharges when it is denied the right to perform them at a given time. However, this would depend on the kind of production process taking place in that industry, which may or may not allow this kind of change in discharges. In our industry model we assume that the industries can only delay their denied discharges. Regardless of the rescheduling behavior of the industry, we also assume that at most, an industry could perform two discharges at the same time: one coming from the production process, and another one coming from the retention tank.

In addition, the industries are equipped with sensors in their outgoing pipes, which are used to control the amount of waste being discharged. This information can then be used for two purposes: to compute the fee each industry has to pay (according to the volume and contamination levels of its discharges) and to control whether the industries perform any unauthorized discharge.

2.2 System Overview

As we can see, there are two main elements in the treatment system: the WWTP and the industries. We have designed our system as a Multiagent System to reflect the physical separation between these elements and also to support privacy in the decision making process of each of the involved agents. Thus, there is an agent for the WWTP and one agent for each industry.

The WWTP agent, upon reception of all the industries' discharge schedules one day (or any different predefined period of time), checks for conflicts between them. That is, it

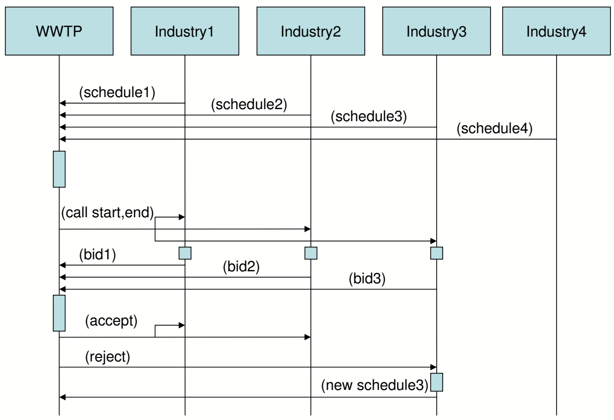

Figure 3: Communication between agents

checks whether two or more simultaneous discharges would exceed the WWTP capacity thresholds. When a conflict is detected, the system has to select which industries are allowed to discharge and which should be delayed. The resolution is done in a sequential way, treating one conflict at a time in chronological order. Once the current conflict is solved, the involved industry agents are informed about the resolution, and then each agent updates its discharge schedule (depending on whether it has been authorized or denied to discharge). Those agents modifying their schedules inform the WWTP agent, and it can then check for the next conflict. This iterative process is performed until all discharges are authorized by the WWTP agent. A pseudocode of the coordination algorithm is depicted in Figure 2. The diagram of the communication between industry agents and the WWTP agent is shown in Figure 3.

The coordination process is done offline, that is, the whole process is done before the discharges are actually performed. It could be done one or several days in advance, depending on the planning capabilities of the industries. The result of the coordination process is a set of new schedules for each industry, which will have no conflicts. In the next subsections we describe in detail each of the steps the WWTP and industry agents perform during the coordination process.

2.3 Conflict Detection

As mentioned previously, each industry agent informs the WWTP agent about its discharge schedule. Thus, this agent is provided with the following information:

$$
\begin{aligned}
\text{D} =\{\ & (s_1^1, q_1^1, \overline{c}_1^1, d_1^1), (s_1^2, q_1^2, \overline{c}_1^2, d_1^2), ..., (s_1^{n_1}, q_1^{n_1}, \overline{c}_1^{n_1}, d_1^{n_1}),\\
& (s_2^1, q_2^1, \overline{c}_2^1, d_2^1), (s_2^2, q_2^2, \overline{c}_2^2, d_2^2), ..., (s_2^{n_2}, q_2^{n_2}, \overline{c}_2^{n_2}, d_2^{n_2}),\\
& \vdots \\
& (s_{NI}^1, q_{NI}^1, \overline{c}_{NI}^1, d_{NI}^1), \cdots, (s_{NI}^{n_{NI}}, q_{NI}^{n_{NI}}, \overline{c}_{NI}^{n_{NI}}, d_{NI}^{n_{NI}})\}
\end{aligned}
$$

where:

- s_i^k is the start time of the k^{th} discharge of industry i,
- q_i^k is the flow of the discharge,
- $\overline{c}_i^k$ is a vector containing the contaminants level of the discharge per volume unit,

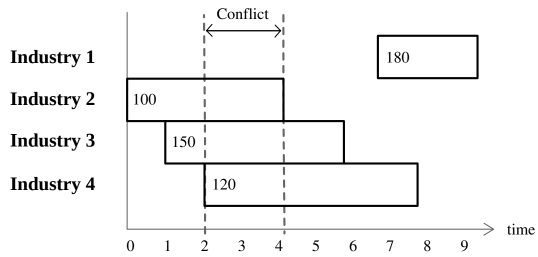

Figure 4: Conflict example

- d_i^k is its duration,
- n_i is the number of discharges of industry i,
- and NI is the total number of industries

A conflict arises when the set of discharges in a given instant exceeds the WWTP hydraulic capacity or the contamination levels. We consider that a conflict begins when an industry starts a discharge that causes any of these thresholds to be exceeded. The conflict ends when an industry finishes a discharge and the WWTP levels go back to be within the allowed limits. The industry that causes the begin of the conflict can be different to the one that causes the end. All the industries that are discharging during the conflict are the ones involved in the conflict.

Figure 4 illustrates an example of a conflict. There are four industries discharging waste with different flows. For example, in timestep 0, the second industry begins to discharge with a flow of 100 m^3/d, finishing at timestep 4. Supposing that the maximum flow capacity of the WWTP is 300 m^3/d, a conflict arises in timestep 2, when industry 4 starts its discharge, because the sum of flows being discharged by the industries (370) exceeds this limit. The conflict ends at timestep 4, when industry 2 finishes its discharge and the sum of the remaining flows (270) falls below the capacity threshold. In this case, the involved industries in the conflict are 2, 3 and 4.

2.4 Conflict Resolution

Once the discharges involved in a conflict are detected, the corresponding industry agents are informed about it, and the coordination process begins. The WWTP agent has to select a subset of the conflicting discharges that will be authorized, while some others will be asked to be delayed. To perform this selection, we have chosen to use an auction mechanism, a well-known mechanism to distribute goods among competing agents when information privacy is a concern [1]. The winners of the auction are allowed to discharge, while the losers have to wait for another opportunity.

The WWTP agent calls for an auction in which the conflicting industry agents can place their bids for the right to perform their discharges. Each bid sent by the industry agents has the following form:

$$b_i = \{(s_i, q_i, \overline{c}_i, d_i), v_i\}$$

The first component of this tuple contains the characteristics of the discharge (as explained in the previous section). The second component, $v_i \in I\!R^+$, is the value the corresponding industry agent gives to the discharge. A description of how this value is computed is given in the next section.

The WWTP agent will receive as many bids as conflicting discharges are. Note that an industry agent could have at most two discharges involved in a conflict (one coming from the industry's production process and another one from the retention tank). Thus, the number of participating agents in the auction can be less than the number of bids. However, each bid is considered to be independent, so there is no restriction on the number of bids an agent can be awarded (it could either be both discharges, one of them or none).

The goal of the WWTP agent is then to select those discharges that maximize a given objective function, subject to the capacity restrictions (hydraulic and contaminants). Formally, to find the winners of the auction, the clearing algorithm must solve the following optimization problem:

$$\begin{aligned} \max \sum_{i=1}^{ND} x_i \cdot g(\_i) & \\ \text{s.t.} \sum_{i=1}^{ND} x_i \cdot q_i & \leq Q \\ \overline{K} & \leq \overline{C} \end{aligned}$$

where:

- ND is the number of conflicting discharges,
- $x_i \in \{0, 1\}$ represents whether discharge i is denied (0) or authorized (1),
- $g(\_i)$ is the contribution of discharge i to the objective function to be maximized. With $\_i$ we refer to all information associated to bid i (start time, duration, flow...). This function can vary depending on the goal to achieve; possible candidates are:

$$g(\_i) = \begin{cases} v_i & \text{maximize discharges' values} \\ q_i & \text{maximize incoming flow} \\ q_i \cdot v_i & \text{tradeoff between previous criteria} \end{cases}$$

- Q is the maximum hydraulic capacity of the WWTP
- $\overline{K}$ is a vector containing the contaminants levels given the authorized discharges:

$$\overline{K} = \frac{\sum_{i=1}^{ND} x_i \cdot q_i \cdot \overline{c}_i}{\sum_{i=1}^{ND} x_i \cdot q_i}$$

- and $\overline{C}$ is a vector containing the maximum contaminants levels accepted by the WWTP

This formulation is similar to a multi-unit combinatorial auction [9], in which the auctioneer offers multiple (but limited) units of different goods and bidders submit bids for a certain number of units of each good. In our case, the goods would be the flow and the contamination levels entering the WWTP, and the available units would be defined by Q and $\overline{C}$. Moreover, since in our case each bidder is allowed to place only one bid per discharge, it could also be seen as a multi-dimensional knapsack problem [10], the optimization problem of selecting a subset of valued objects that can fit into a bag with restrictions on its dimensions, with the goal of maximizing the stored value.

To solve the winner determination problem we have used a Linear Programming approach, since the size of our problems is not too large. However, if the size were to be intractable by linear programming, other algorithms could be used, such as Genetic Algorithms or any of the existing efficient combinatorial auctions algorithms, such as those presented in [5]. Note that the auction is repeated each time a conflict is detected, so we are dealing with a recurrent auction.

2.5 Bidding Policies

One of the key points in auctions (besides the winner determination algorithm) is the bidding policy of the bidders. This policy determines how an agent generates its bids, i.e. how it chooses the goods to bid for and the value (price) associated to each good (or set of goods).

In our case, the agents do not have to choose among different goods, since these are already defined by the discharge characteristics. However, the agent still has to compute the value v_i for each bid. We have considered two different alternatives for such computation. In the first one, the value represents the urgency the industry has for performing the discharge. This urgency would depend on the production process of the industry and the state of its resources (e.g. is the retention tank available? How much can I store there?). Thus, the more urgent a discharge is, the higher its associated value should be.

The other alternative is to give a more economical view to the value. In this case, it could represent the price the industry is willing to pay for performing a discharge at a given time. This alternative has somehow the urgency degree embedded, since probably an industry needing to perform a discharge immediately is inclined to paying much more than when the urgency is low.

2.6 Adding Priorities

The auction mechanism presented previously has a major drawback: it can be *unfair*. The problem is that the main goal is to maximize the objective function, without taking into account what industries are authorized or denied the right to discharge. This happens because the identity of the bidders is not used when determining the winners of the auctions. Thus, if an industry wanted never to be denied a discharge, it could always bid very high, pretending that its discharges are very urgent, guaranteeing its success in the auction. This behavior could prevent other industries from getting any discharge authorization, which could highly affect its production processes, or even force them to perform a discharge without having been authorized (act that would be detected by the industry's sensors and could be penalized by the regulatory authorities).

To avoid this situation, and in order to find a *fair solution* to the conflict, we have added a priority mechanism that takes into account the history of each agent having been authorized or denied to discharge over time. The mechanism assigns a priority to each agent, $W = \{w_1, w_2, \cdots, w_{NI}\}$, that is used in the auction clearing algorithm to find the solution. High priority values indicate that the agent should be authorized to discharge, while low priority values indicate that it would not be unfair to deny a discharge to that agent. These values are updated after each auction, according to its outcome. If an agent wins an auction, its priority is decreased, while if the agent loses, it is increased.

These priorities could be used in very different ways, such as defining new constraints that should be satisfied by the solution, or directly designating some or all of the winning agents, among others. We have chosen to use the priorities as a modifier of the bids sent by the agents. Formally, given a bid b_i submitted by agent k, a new bid is computed as:

$$b'_i = f(b_i, w_k)$$

This function f is a parameter of the system. The use of the priorities introduces an *egalitarian* view of the auction, in contrast with an utilitarian view, which is used more often but does not take into account how fair is the system. Thus, even if an agent bid very high, given it had a low priority, its bid should be decreased somehow. Similarly, an agent bidding low could still be the winner of the auction if it had a high priority.

3. IMPLEMENTATION

To evaluate the coordination mechanism we have implemented a prototype of the system. For programming it we have chosen Repast [14], a free open source software framework for creating agent based simulations using the Java language. The simulation reproduces the process and the communication between the WWTP and the industries performing waste discharges. We have created an agent to represent the WWTP and another one for each one of the industries. So far we have only considered the hydraulic capacity. In the near future we will consider different contaminant components.

To calculate the bid, the industry agent takes into account the urgency for performing the discharge, based on the retention tank occupation of the industry:

$$v_i = \frac{\text{tank occupation}_i}{\text{total tank capacity}_i}$$

In case an industry agent has to reschedule its discharges, its behavior is the following: it first tries to store the rejected discharge into the tank. The discharge of the tank is then scheduled as the first activity of the agent after the current conflict finishes. The remaining discharges are shifted so that they do not overlap.

To calculate the priority of agent k, w_k, we take into account the number of lost and won auctions:

$$w_k = \frac{\text{lost auctions}_k + 1}{\text{total participated auctions}_k + 2}$$

The initial priority of each agent is 0.5. If an agent loses an auction its priority is increased, otherwise it is decreased.

The function chosen to modify the industry's bid according to its priority changes the value sent by the industry. So, given v_i, the value bid by agent k and its priority, w_k, the actual value used is $v'_i = v_i \cdot w_k$.

The objective function $(g(_{-i}))$ to maximize in the auction clearing is the sum of discharge values. The linear programming toolkit GLPK [8] has been used to solve the winner determination problem.

As a first evaluation of the system, we have supposed that the industries always obey the WWTP decisions, as long as they have enough tank capacity. We will introduce different industry behaviors in future experiments to have more realistic scenarios.

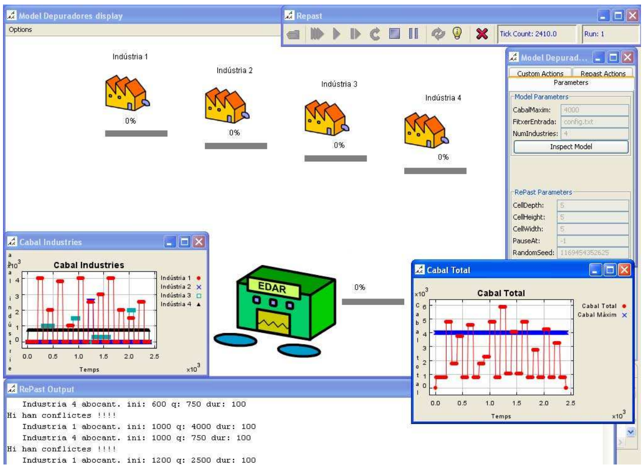

Figure 5: User interface

Figure 5 shows the application user interface. The graphical representation shows the tank occupation levels of the industries and the occupation degree of the WWTP.

4. EXPERIMENTAL RESULTS

In order to evaluate the system we have considered some quality measures based on different characteristics of the solution. These characteristics are the following:

- **number of overflows (NO)**: number of overflows occurred during the execution of the discharge schedules.
- **maximum flow overflown (MFO)**: measured in m^3/d.
- **volume overflown (VO)**: total liters overflown.
- **modifications time (MT)**: sum of differences between the initially proposed discharges times and the actual times after coordinating the schedules, measured in minutes.
- **total delay time (TDT)**: difference in minutes between the final time of the execution and the final time when no coordination is used.
- **% authorized discharges (%A)**: percentage of conflicting discharges that have been authorized.
- **minimum % won auctions (%MWA)**: minimum percentage of won auctions among all agents.

The experiments consisted of ten simulations using a set of real data of ten industries in ten different days. The industries can have the same discharge schedule each day (if they produce the same products every day) or different (if they have changes in the production process).

We have tested the system with three different scenarios. In the first scenario there is no coordination between the industries and the WWTP. The second has coordination, and the third has coordination and uses the priorities.

The results shown in Table 1 are the average and the standard deviation of the ten executions (dashes indicate that

	NO	MFO	VO	MT	TDT	%A	%MWA
without coord	5.7 (1.16)	6515 (2069.36)	1388950 (664854.43)	—	—	—	—
coord w/o prios	0 (0)	0 (0)	0 (0)	10492.1 (2845.67)	335.5 (193.84)	74 (6.2)	25.75 (18.22)
coord w prios	0 (0)	0 (0)	0 (0)	7069.3 (2787.56)	243.4 (176.96)	78 (5.8)	48.73 (18.42)

Table 1: Experimental results (average and standard deviation)

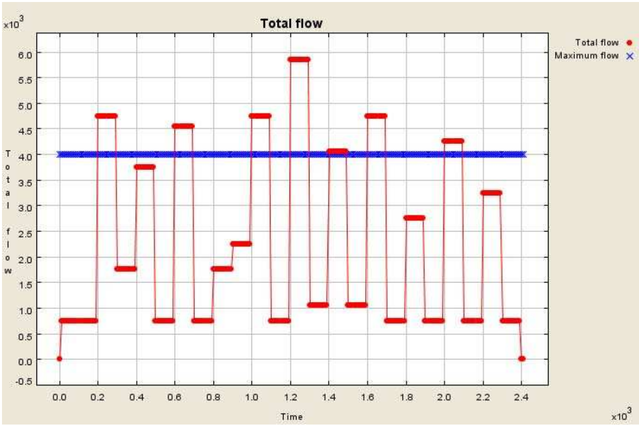

Figure 6: Behavior without coordination

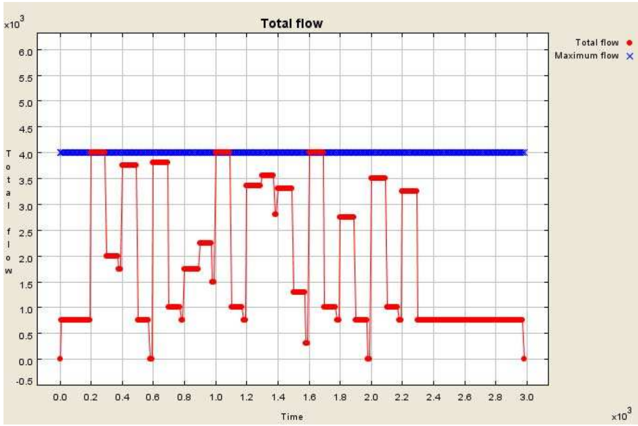

Figure 7: Behavior with coordination and priorities

the measure is not applicable). The results show that with schedule coordination (no matter with or without priorities) we eliminate the overflows. Figure 6 illustrates an example of the behavior of the system without any coordination, while Figure 7 shows the behavior of the system with the same example when using the coordination mechanism with priorities. In the first figure we can observe that the WWTP capacity is being exceeded seven times, while in the second the maximum capacity is never exceeded.

However, the final time of the execution (TDT) is increased by about 4 and 5 hours respectively (243.4 and 335.5 minutes as shown in Table 1). This could cause some problems with the scheduling of the following day. Although with priorities the average and standard deviation of this time are lower than without priorities, it is still a considerable delay. This is also the case with the modification time, for which the increase is lower with priorities than without. Looking at Figure 7 we can observe that sometimes the WWTP flow is underused. If industries were allowed to perform multiple discharges (from the buffer and from the production process) at the same time, the reschedule delays could be shortened. We need to deal with this possibility in future work.

Regarding the rate of authorized discharges (%A), the percentages obtained with or without priorities are similar. However, as shown in Figure 8, the standard deviation in each of the experiments is smaller when using priorities. Actually, the average of this standard deviation is 27% without priorities, and 19% with priorities. This means that the difference between the agents is reduced with the use of priorities, increasing the fairness of the system. Moreover, the results also show that the minimum percentage of won auctions (%MWA) is significantly increased with priorities (48,73% against 25,75% when priorities are not used,

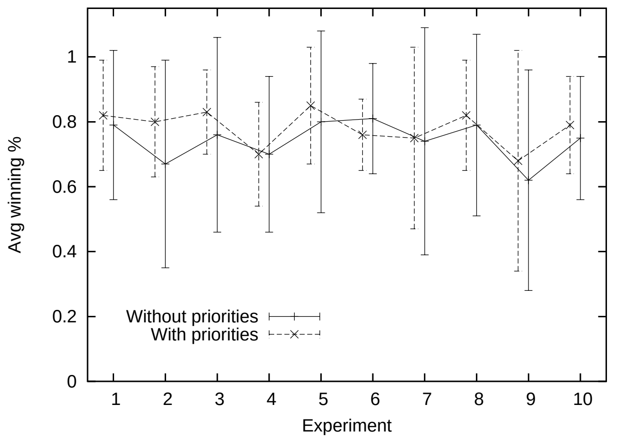

Figure 8: Average and standard deviation of won auctions percentage for each experiment

as shown in Table 1). This indicates that an agent has more chances of being authorized to discharge when the priority mechanism is used.

5. RELATED WORK

Coordination in multi-agent systems is a very important issue since it directly affects the overall system performance. The coordination can be performed at execution time or at planning and scheduling time. In our work we have focused on the latter option, coordinating agents schedules. Schedule coordination can be further divided depending on whether the goal is to coordinate existing schedules or to create new schedules for each of the agents. Since our agents (industries) have their individual schedules, we are faced with the former problem of coordinating schedules that have already been generated.

There are many approaches to handle schedule coordination, from a divide-and-conquer strategy [2, 3, 15], to solving it as a constraint optimization [4], or using auctions [6, 7], among others. We have followed this option, use a market-approach to coordinate the agents' schedules. The characteristics of our problem makes the auction to be continuously repeated, so we are dealing with a recurrent auction. This kind of auction is recently being used for e-services markets, such as assigning advertising time in public displays [12] or in networking markets [11].

This latter work is closely related to our problem, since it tackles the bidder drop problem. This problem arises when bidders are frustrated with the outcome of the auctions (usually because they are constantly losing) and decide to leave the marketplace. We are also very interested in this problem, since we need to incentivize the agents to participate in the coordination process. In [11] the problem is solved by defining a more flexible winner determination algorithm, which takes into account the bidder's outcome history in past auctions. The goal of their work is to incentivize the bidders to stay in the market place, so that the prices do not collapse. We also use this history in order to compute the agents' priorities, but our objective is not economic, but to obtain a fair distribution of the discharge authorizations.

Regarding work on water treatment systems, there has been much research on the internal treatment processes, but very little on coordinating the different systems involved. An example is [13], where a negotiation approach to coordinate different WWTPs treating the same river basin is presented. However, the elements being coordinated in this work are the WWTPs, leaving the industries aside.

6. CONCLUSIONS AND FUTURE WORK

The pressure exerted to water resources is rapidly increasing due to the high demand by industries and domestic consumers. To ensure a proper supply, water must be treated to remove contaminants so that it can be put back to the river. Current regulations prevent industries from discharging waste with high contamination load. However, these are deficient and do not fully attain the goal of maintaining a low contamination level in the water so that it can be successfully treated. Instead of focusing on individual industries, a more integrated view of the problem is needed. Taking into account all the elements involved in the treatment system and their interactions could help improving its performance.

In this paper we have presented a mechanism to coordinate the discharge schedules of a set of industries so that the WWTP is capable of treating all the waste they produce. Through the coordination, the individual schedules are refined whenever a conflicting situation is detected. The new schedules contain a sequence of discharges that are distributed over time and decrease the risk of possible treatment failures by the WWTP. The core of the coordination mechanism is a recurrent auction, in which industries can bid for the right to perform a discharge at a given time.

The auction has been extended with a priority mechanism to introduce fairness in the assignment of authorizations. With this mechanism, the authorizations to perform discharges are evenly distributed among the industries, meaning that their original schedules are modified the least possible. This is a very desirable property of a coordination mechanism, since it incentivizes agents to participate in it. Otherwise, if the mechanism were to drastically modify their schedules, agents would be reluctant to participate. This could lead to an overall failure of the system, since each agent would be acting on its own. This is specially important in environments where agents are self-interested and do not pursue a common goal. In these cases, cooperation must offer an add-on so that agents are attracted. In our water treatment domain, this add-on would be discounted fees for those industries complying with the discharge authorizations.

The results obtained through simulation show that the coordination mechanism accomplishes the goal of maintaining the incoming flow below the WWTP hydraulic threshold. The results have also shown that the use of priorities provides a fairer solution of the auctions. However, we need to further study how to reduce the delay produced by rescheduling. We also need to incorporate the contaminants levels restriction, since we have taken into account only the hydraulic capacity of the plant.

The system presented in this paper does not yet capture the complexity of a real water treatment system. As a first step towards getting more realistic scenarios, we plan to introduce disobedience behavior to the industry agents. This implies that agents may perform discharges even when they are not authorized. We plan to use a trust mechanism so that the WWTP agent may foresee this kind of situtation and can react accordingly. Some other open issues include the introduction of a more economic view on the bid value, and study whether the method used to compute the priorities is the most appropriate.

7. ACKNOWLEDGMENTS

We thank the LEQUIA research group (Laboratory of Chemical and Environmental Engineering) for having introduced us to the water treatment problem and for the help given during our research. This research project has been partially funded by the Spanish MEC project TIN2004-063540C02-02 and DURSI AGAUR SGR 00296 (AEDS).

8. REFERENCES

[1] Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[2] B. Clement and E. Durfee. Top-down search for coordinating the hierarchical plans of multiple agents. In *Proceedings of the Third International Conference on Autonomous Agents*, pages 252–259, 1999.

[3] J. Cox and E. Durfee. Discovering and exploiting synergy between hierarchical planning agents. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 281–288, 2003.

[4] J. S. Cox, E. H. Durfee, and T. Bartold. A distributed framework for solving the multiagent plan coordination problem. In *AAMAS '05: Proceedings of the fourth international joint conference on*

Autonomous agents and multiagent systems, pages 821–827, New York, NY, USA, 2005. ACM Press.

[5] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

[6] M. B. Dias and A. Stentz. A Free Market Architecture for Distributed Control of a Multirobot System. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems*, pages 115–122, 2000.

[7] B. Gerkey and M. Matarić. Sold! Auction methods for multi-robot control. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.

[8] GLPK. GLPK (GNU Linear Programming Kit), http://www.gnu.org/software/glpk/.

[9] J. Kalagnanam and D. Parkes. *Handbook of Supply Chain Analysis in the E-Business Era*, chapter Auctions, bidding, and exchange design. Kluwer Academic Publishers, 2005.

[10] T. Kelly. Generalized Knapsack Solvers for Multi-unit Combinatorial Auctions: Analysis and Application to Computational Resource Allocation. *LNAI*, 3435:73–86, 2005.

[11] J.-S. Lee and B. Szymanski. *Service Enterprise Integration*, chapter Auctions as a Dynamic Pricing Mechanism for e-Services, pages 131–156. Kluwer, 2006.

[12] T. Payne, E. David, N. R. Jennings, and M. Sharifi. Auction mechanisms for efficient advertisement selection on public displays. In G. B. et al., editor, *ECAI*, pages 285–289. IOS Press, 2006.

[13] T. Rendón-Sallard, M. Sànchez-Marrè, F. Devesa, and M. Poch. Simulating scenarios for decision-making in river basin systems through a multi-agent system. In *Proccedings of CCIA 2006*, 2006.

[14] REPAST. Repast agent simulation toolkit, http://repast.sourceforge.net.

[15] H. Tonino, A. Bos, M. D. Weerdt, and C. Witteveen. Plan coordination by revision in collective agent-based systems. *Articial Intelligence*, 142(2):121–145, 2002.