# WaRP - A Reactive Planner integrated in an environmental decision-support system for Wastewater treatment plant management

**Luigi Ceccaroni**[1] and **Dave Robertson**[2]

**Abstract.** In this paper we present an environmental decision-support system (named DAI-DEPUR+) and a reactive linear planner (named WaRP[3]), and their application to the wastewater domain. The environmental decision-support system:

- receives and processes on- and off-line information about a wastewater treatment process.
- uses this information for selecting and revising models of the treatment-system dynamics.
- applies these models for autonomously planning the control of a wastewater treatment plant (WWTP) and for supporting its management by operators.

## 1 INTRODUCTION

The DAI-DEPUR+ system is an environmental decision-support system (EDSS) which receives on-line inputs from sensors all over a *wastewater treatment plant* (WWTP) as well as off-line inputs from the WWTP's laboratories and human operators. The system uses its internal knowledge bases and inference mechanisms to process and understand this information, to diagnose the ongoing WWTP-state, and to predict the evolution of that WWTP state. Finally, the output of the system is represented by statements about actions to be taken, or statements to support human decisions in future actuations, or direct control signals to WWTP devices in order to maintain the plant's operating state.

The general issue we address with this work is the one of optimizing wastewater treatment operation by more reliable management and automatic control.

The general operation of a WWTP always includes various internal pre-designed standard units whose sub-operation is already optimized to accomplish a single task[4]. However, each sub-operation usually has effects on other downstream treatment processes, and tradeoffs between increasing the efficiency of one process or another are necessary, taking into account major constraints such as water characteristics, effluent quality and costs of each operation.

The process of wastewater treatment is so complex that it is difficult to develop a reliable supervisory technology based only on a chemical-engineering classic-control approach. It is hoped that by using AI systems we can obtain better results in wastewater management by incorporating more of the modelling of human decision procedures within the architecture of automated control systems. Classic knowledge-based expert systems proved able to cope with some known difficulties and to face several WWTP-domain problems, even if they are not the definitive solution to the treatment problem as a whole. The architecture in which WaRP is embedded integrates different reasoning systems, such as ontologies, rule-based reasoning, case-based reasoning and reactive planning, and it is flexible enough to deal with the complexity of the wastewater treatment process, given an adequate amount and kind of data.

### 1.1 General overview

In this paper we present WaRP, a reactive planner for wastewater management and we discuss the approach it takes to some important issues in reactive planning. This planner belongs to a more general decision support system for the supervision of wastewater treatment plants, which, in turn, is part of the knowledge and technology needed for the rational management of water resources.

We start by describing the state of the art of reactive planning and *environmental decision-support systems* (EDSSs) (section 2). Section 3 explains how the decision support system (DAI-DEPUR+) has been designed; it includes a brief description of its layered architecture. Finally, section 4 introduces the WaRP planner; its architecture and features are considered, as are the associated problems, and the functioning of its components is discussed.

## 2 REACTIVE PLANNING AND ENVIRONMENTAL DECISION-SUPPORT SYSTEMS

### 2.1 Planning

As control and decision-support systems become more sophisticated, they almost always involve planning of some sort. Above all, planning nowadays plays a role in expert systems and decision support systems when they have to reason about events occurring over time.

Classic planners (the term is due to Wilkins [13]) rely on perfect knowledge assumptions:

1. The planner has full knowledge of the initial conditions in which the plan will be executed, e.g. whether there is a bulking sludge problem;

---

[1] Universitat Politecnica de Catalunya, Departament de Llenguatges i Sistemes Informatics, Campus Nord, Modul C6, Jordi Girona 3, 08034 Barcelona, Spain; email: luigic@lsi.upc.es

[2] The University of Edinburgh, Division of Informatics, 80 South Bridge, Edinburgh EH1 1HN, UK; email: dr@dai.ed.ac.uk

[3] *Wastewater Reactive Planner*

[4] A task in this context is generally the removal / remediation of contaminant substances or pathogenic microorganisms [3].

2. All actions have fully predictable outcomes, e.g. reducing the oxygen input will definitely either solve or not solve the bulking sludge problem;
3. All change in the world either occurs through actions performed by the planner or is otherwise defined with no uncertainty.

Now, if we consider an example of a plan in the domain of wastewater, we realize that the perfect knowledge assumptions are an idealization of the planning context and are not realistic, because the world of wastewater is to some extent unpredictable:

"To solve bulking sludge problems, try to reduce the oxygen input by one and then two units, and, if bulking sludge is still there after three days and the weather is dry, block the recirculated sludge flow completely."

In general those assumptions may lead the planner to drop options that would have been useful if potential problems had been anticipated. For example, on the assumption that the weather will be sunny, as forecast, you may reduce the water treatment capacity (and the costs) of a WWTP for the weekend, when nobody works there; if the forecast later turns out to be erroneous, it is then impossible to treat all the water entering the plant (incurring high fines).

#### *2.1.1 Frame problem*

Planners can specify actions at a higher level (e.g., in term of their effect on the world, such as "increase the oxygen content of the reactor") or a lower level (e.g., "turn on the second oxygen-pump motor"), but high level actions have to include the specification of the micro control steps to actually make a system perform the plans. An important issue in planning is the need to characterize what is not changed by a particular action. Increasing the oxygen flow of the reactor of a WWTP does not change the recirculation flow of the sludge. The specification of what is true in one state of the world and exactly what is changed by performing some actions in the world is known as the frame problem [7].

#### *2.1.2 Reactive planning*

A problem that can arise is that, if the planner takes control of the processor, the EDSS can no longer respond to events in the environment. A solution to this problem is for the planner to work incrementally in a temporally integrated system[5], doing a few computation steps during each state transition cycle, then storing its state until the following cycle. Other parts of the EDSS that can react more quickly to changes in the environment run in parallel with the planner and therefore are able to act (between cycles) even if the planner has not finished its global computation. At each tick, if the planner is finished, it issues the plan, otherwise it emits a signal that indicates it is not yet ready [6].

### 2.2 Environmental decision-support systems

An EDSS is an integrated *knowledge-based system* (KBS), applied to an environmental issue, that reduces the time in which decisions are made and improves the consistency and quality of those decisions [5].

An EDSS usually includes the following features:

[5] Temporal integration of a number of processes that work at different rates is to define a constant minimum-cycle time for the entire system. At each tick of this time the inputs are read or calculated, some computation is done and the outputs are set.

- The ability to assist the user in deciding when and how the different available tools have to be applied.
- A structured framework for the assessment which draws information from the user and the environmental system about domain-characteristics and processes in a logical manner. This framework, besides acquiring the domain knowledge, has to be able to organize and represent it.
- Specific knowledge-bases pertinent to the type of domain being considered or to the process being carried out at the site. These knowledge bases contain general data on environmental parameters and processes that are relevant to the domain (e.g. what toxic materials are used in the processes; which kinds of physical, chemical and biological samples need to be collected; which is the relative importance of the features in play; which are the requirements of the local legislation).
- A general environmental knowledge which is used to deduce the relative significance of different environmental impacts given appropriate data about the specific domain and processes.
- The ability to assist the user during the interpretation of the results and the selection of the solution.

The study of this paper is based on DAI-DEPUR [10], which is a distributed and integrated supervisory multi-level agent-based architecture for WWTP management. It combines in a single framework several cognitive tasks and techniques, such as learning, reasoning, knowledge acquisition, distributed problem-solving, and different AI techniques, such as rule-based reasoning and case-based reasoning. Four levels are distinguished from the domain-model point of view [12]: data, knowledge, situations and plans. On the other hand, from the supervision-task point of view, five levels are considered: evaluation, diagnosis, supervision, actuation and learning. This system was developed for the WWTP domain, but it can represent a general framework for complex-process supervision [11].

## 3 THE DAI-DEPUR+ SYSTEM

An environmental decision-support system for wastewater treatment plants has been developed. This system, called DAI-DEPUR+, is an evolution of DAI-DEPUR and has an architecture in which several artificial intelligence techniques integrate and operate in real time. This integration has, as the main new element, a reactive planner for the supervision of the wastewater treatment process.

### 3.1 Architecture

The architecture of the system has a modular design, to improve modifiability, understandability and reliability. We chose an architecture with basically a standard vertical decomposition approach (see Brooks [2] and Kaelbling [6]): a division is made into many specialized subsystems, such as data interpretation, diagnosis, modeling, planning, execution and effector control modules.

The system receives raw data from the sensors and the laboratory, and emits commands to the sensors and effectors. The action component takes the output of the perception component as input and it is the one which generates commands to both the sensors and effectors.

Excepting cases of failure, there is a continuous sensory data stream from all sensors, which goes directly into the perception component, along with the results of laboratory analyses and the commands that were last sent to the effectors.

In this section we explain the three layers of DAI-DEPUR+'s architecture: perception (or data interpretation), diagnosis and decision support. In next section we will focus on the WaRP planner module.

## 3.2 Perception layer

The DAI-DEPUR+ system operates in a domain which physically consists of a wastewater treatment plant. In particular, all the physical, chemical and biological measurements are gathered in treatment plants located in Catalunya. Some parameters are measured on-line by sensors, while other ones are measured off-line in laboratories.

### *3.2.1 Awareness*

The DAI-DEPUR+ system is designed in such a way that there is a constant control on the time the sensors and the laboratory analyses get a particular reading and the time the effectors can react to that information. The WWTP environment is very slowly evolving compared to the speed of the reasoning of the decision support system: even if a WWTP is a truly dynamic domain, it never changes to such extent that the results of relatively long calculation would no longer be useful.

### *3.2.2 Temporal integration*

In a WWTP, sample intervals range from a few seconds to a few days. Our approach to the temporal integration of a number of processes that work at different rates is to define a constant minimum-cycle time for the entire system. This time is equal to one hour and at each tick of this time the inputs are read or calculated, some computation is done and the outputs are set (by the action component of DAI-DEPUR+). If a process, such as a laboratory analysis, cannot complete (or even cannot be started) by the tick of the time, either because its scheduling is non-constant or there is failure, its outputs are inferred, if possible, in an alternative way (often just reproducing the outputs of the previous hour) and its execution is replanned for the following tick.

Once obtained, the data are arranged according to different criteria: separations are made between physical-chemical and microbiological features, and between quantitative and qualitative ones. To have a preliminary idea of the quality of the treatment process, a classification with LINNEO+ [1] and AutoClass algorithms is made, which gives, as result, an estimation of the corresponding state of the plant.

### *3.2.3 Physical and chemical features*

Among the available physical and chemical features, the most relevant ones used by the DAI-DEPUR+ system are selected on the basis of human experience, tradition and utility measures. These features are not problematic and their modeling and application both in chemical engineering and artificial-intelligence systems are well documented in bibliography.

### *3.2.4 Microbiological features*

On the other hand, there are the microbiological features, whose modeling exists in the scope of biological disciplines, but has not yet been integrated into a decision-support system dedicated to environmental issues, such as the DAI-DEPUR+ system.

Initially, the identification of the microorganisms existing in the activated sludge is carried out. This is generally done in the laboratories of the plant and generates qualitative off-line data (e.g., presence of Paramecia species or diversity of Ciliate). Subsequently, a comparative study about microorganism communities of different treatment-plants is accomplished, to understand which can be the influence of biological variability at a geographical level. A set of microbiological features is then selected, to be used by system, and, for a high performance to be maintained throughout the domain (the different WWTPs), these features' measurements need to be widespread enough to have a representational data-set with a relatively abundant amount of instances. Portability-wise, parameters available only in the minority of the treatment plants are not very useful in the development of the ontological knowledge-base of the system, but they can be employed as specific-domain knowledge by specially developed modules.

Missing and incomplete information does not represent a problem in principle, but raises the degree of uncertainty in planning.

## 3.3 Diagnosis layer

The knowledge bases of DAI-DEPUR+ model the particular kind of WWTP the data are coming from. The knowledge bases make use of different reasoning models, whose integration seems to be necessary to obtain good results in wastewater process management.

In the DAI-DEPUR+ system there are a numerical control module and three knowledge bases (related to an ontology a case-based reasoner and a rule-based expert system) which manage the general wastewater-treatment operation when the plant is in a normal state or in a standard abnormal state, such as bulking, storm or foaming states.

The rule-based expert system and the case-based reasoner work in parallel and they both produce as output a diagnosis on the state of the plant. If the diagnosis of the two systems is the same, it is passed to the planner. If the diagnoses exist and are different, the system prioritizes as follows:

- If the database contain a predefined minimum historical series and the case distance is smaller than a predefined value, the case-based reasoner's diagnosis prevails.
- Otherwise, the rule-based expert system's diagnosis prevails.

In case of impasse, DAI-DEPUR+ turns to the ontology or the plant manager, demanding an off-line diagnosis based on their knowledge. This external solution is recorded and learned.

## 3.4 Decision support layer

This is the supervisory level of the DAI-DEPUR+ environmental decision-support system. It integrates the diagnosis of the reasoners and the ontology, and runs the plans proposed by the WaRP planner.

This layer runs always in real time and it is very robust in this sense. At least one of the levels of competence of the planning module always knows what to do because, at the lowest level, it works with very weak information and its computing speed in generating plans is very high with respect to the entire system's minimum-cycle time.

# 4 THE WaRP REACTIVE PLANNER

Similarly to SRI's PRS-CL [8] [4], the WaRP reactive planner is a real-time, continuously-active, intelligent system being developed for representing and using experts' procedural knowledge for accomplishing goals. At the present time it is a mainly data-driven, second-principle planner, which does not evaluate the quality of the information in its database. It is our intention to develop an extension of the planner with different levels of sophistication, which activate according to the strength of available information.

The reactive planner has the task of selecting and enacting linear plans for plant control given some environmental conditions. We confine our attention to linear plans because these have so far been adequate for our application. We enact them sequentially, and choice among sequences is obtained by choosing among linear plans.

To give a feel for the sort of reasoning involved in this sort of reactive planning, we give a simplified example here. Suppose that we have two possible sequences of WWTP states as invocation conditions: Plan 1, for which the sequence is $a \Rightarrow b \Rightarrow c \Rightarrow d \Rightarrow e$, and Plan 2, for which the sequence is $b \Rightarrow c \Rightarrow a \Rightarrow c$. Given the observations we make of the real state of the WWTP we would like to be able to choose between plans 1 and 2 as models for operation of the system. We hope that when we choose a plan it will be useful for the entirety of the sequence of states it describes but this may not happen. If the real state and the chosen plan diverge then we want to be able to revise to a new plan in mid-sequence or restart an earlier plan. Suppose that in our example the real state of the system evolves as shown in the first row of the table below. We start with real state 'a' and the best plan appears to be Plan 1. As the real state changes, through states 'b' and 'c', our choice of Plan 1 appears to be confirmed but the next real state is 'a', which does not match to the expected state ('d') of Plan 1. It does, however, match to the third state of the invocation conditions of Plan 2 and we can see that the earlier states of Plan 2 would also have matched the real state sequence had we shifted soon enough, so we revise to Plan 2. The next real state is an 'a', which does not match to the final state of Plan 2 so we abandon it and restart Plan 1. This sequence of plan revision is shown in the lower rows of the table below, with the active transitions between plan states shown as '⇒' symbols and the inactive transitions as '. . .' symbols.

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Real state sequence | a | ⇒ | b | ⇒ | c | ⇒ | a | ⇒ | a |
| Plan 1 | a | ⇒ | b | ⇒ | c | . . . | d | . . . | e |
| Revise to Plan 2 | | | b | . . . | c | ⇒ | a | . . . | c |
| Restart Plan 1 | | | | | | | | ⇒ | a |

## 4.1 Issues

The WaRP planner possesses the following capabilities:

- It is capable of changing its plan of action adaptively.
- It handles exogenous events, such as water input changes.

However, there are several issues that have not been addressed:

- The problem of determining whether it is worth planning for a particular outcome.
- The scheduling of sensing actions that detect the occurrence of a particular contingency.
- The use of any information about the likelihood of any events: WaRP is not a probabilistic planner.
- The possibility of interleaving planning and execution.

In the WWTP domain not all sources of uncertainty and not all possible outcomes of the actions are known. WaRP is not a classic planner: its job is to construct and run plans that are not guaranteed to achieve its goals. It runs in real-time and receive the main goals from plant managers. We believe that WaRP is sound, complete and fast enough to be of practical use.

## 4.2 Terminology

To avoid confusion, we describe the relation between the terms we use and the ones used by other authors. We use the term *reactive plan* to refer to a plan whose execution depends on the evolution of surrounding environment. It contains actions that may or may not actually be executed, depending on the circumstances that hold at any particular time. Other authors (e.g. Pryor and Collins [9]) use the term contingency plan to refer to the same concept. We refer to facts and beliefs about the world as *facts*; we refer to plans, Acts and procedures as *plans*; we refer to invocation part and preconditions of a plan as *invocation conditions*; we refer to body and plot of a plan as *body*. We use the term *context dependent effects* simply to refer to action effects that depend on the context in which the action is performed. We use the term *descriptor* when dealing with databases to refer to the terms: descriptor, variable, attribute and parameter. A *database* is usually a descriptor-value matrix.

## 4.3 Architecture

The WaRP reactive planner is intended to simultaneously achieve its goals based on its current beliefs about the world (facts) while noticing and responding to new events. Two important features of this sort of reactive system are its ability to detect new events and its capacity to adapt to them.

### *4.3.1 Awareness*

In comparison to the slow evolution of the WWTP state, the work of WaRP is not very time consuming. Perceptual inputs coming either from the diagnosis layer of DAI-DEPUR+ or directly from the external world are never lost because of slow processing. The standard tick of the time for WaRP is every hour, but particular inputs can trigger a faster response. In this way the WaRP system, in case of emergency, can prepare new plans before the following tick.

### *4.3.2 Robustness*

Robustnesswise, we propose a solution similar to that of Kaelbling [6]. The WaRP system includes different levels of competence in such a way that, if higher levels break down, the lower levels still continue to work acceptably. We concern ourselves with robustness in relation to failed sensors, to lack of diagnosis and to the possibility of general confusion because of new or unusual situations. We refer to the first two types of robustness as perceptual and to the third type as behavioral.

Perceptual robustness is achieved by integrating all sensory information, laboratory analysis results and diagnosis information into a structure (the database of facts) that represents the system's knowledge or lack of knowledge about the WWTP. If a particular sensor (or the reasoners) fails and its failure has been detected, the system's information about the world is weaker than it would have been if all the sensors and reasoners had been working correctly. Thus, with weaker information, the system can make less perfect discrimination among the states of the WWTP, but it can still be the case that the information integrated from the remaining sensors will suffice for reasonable, even if degraded, planning. But, if planning in a particular situation depends entirely on a single sensor or on the reasoners' diagnosis and this fails, there is no room for graceful degradation; the planning operation will simply fail. The problem of detecting sensor failure is a difficult one that we shall not be examining in this paper.

Behavioral robustness depends upon the ability to trigger actions in direct accordance with the strength of available information. If the system's actual information is insufficient for the planner to produce a high-level plan, perhaps because the system has just been switched on or has become confused, that level will simply emit an output indicating its inability to form a plan.

### 4.4 Database

The database is a large collection of descriptors with their values of some type. The database includes facts about properties of the WWTP domain, such as the wastewater characteristics, the presence of microorganisms or the general state of the WWTP. Some facts go continuously from the sensors to the database, other facts are acquired dynamically by WaRP as it executes its plans. The planner as a whole can be visualized as the database and the code that update the descriptors once per tick. The descriptors can be divided into input, state and output. The input descriptors, conceptually connected to the sensors and the laboratory of the WWTP, contain sensory and analytical values at the beginning of each cycle. The state descriptors are updated once per tick, the values coming directly from the diagnosis layer of DAI-DEPUR+ or from the user. If the planner does not receive these data, it calculates them as a combined function of the values of the input descriptors and the old values of the state descriptors. The output descriptors, conceptually connected to the effectors, are updated once per tick as a combined function of the inputs and the old values of the state descriptors.

### 4.5 Goals and intentions

Goals are expressed as conditions over some sequence of world states and are described by applying temporal operators to state descriptors (e.g.: Achieve ”WWTP state = ok” during next 10 days). This enables the representation of a wide variety of goal types, including goals of achievement, goals of maintenance and goals of testing for given conditions. In WaRP, intentions are defined as a particular type of persistent goals, in which the system chooses to perform an action (or a set of actions) and then it actually perform it.

### 4.6 Plans

Knowledge about how to accomplish given goals or how to react to certain situations is represented in WaRP by plans, which are declarative procedure specifications consisting of an invocation condition and a body. Together, the invocation condition and the body of a plan specify under what situations the plan is useful (applicable) and express a declarative belief about the results and utility of performing certain sequences of actions under certain conditions. The system selects eligible plans based on their priority (depending on resource requirements, resource availability and costs).

The WaRP planner is still under development and contains a limited number of plans in its library.

### 4.7 Reactivity

The WaRP system needs to be able to change its structure during execution and to be able to respond in real-time to additional conditions (facts and goals). In the WaRP planner, changes in the environment may lead to changes in the system’s goals or beliefs, which in turn may result in new sub-goals (testing, monitoring, diagnosis and recovery procedures). WaRP is therefore able to change its focus completely and pursue new goals when the situation warrants it. In the WWTP domain, such switches happen frequently as emergencies of various degrees arise in the process of handling less critical tasks.

## 5 CONCLUSIONS

We presented an environmental decision-support system (DAI-DEPUR+), which includes a reactive planning module, applied to the wastewater treatment domain. DAI-DEPUR+ assists the user to decide which goals are to be reached, and when and how the different available reasoning tools have to be applied in order to control a wastewater treatment plant. It represents a structured framework for the assessment of the treatment process, and it draws information from the user and the WWTP system about biological processes, domain-characteristics, cost constraints and environmental regulation. This framework, besides acquiring the domain knowledge, organizes and represents it.

The planner (WaRP) embedded in DAI-DEPUR+ is the first application of this kind to the domain of wastewater treatment. Of particular interest is its integration with an ontology, a case-based reasoner and a rule-based expert system, from which WaRP receives the diagnosis that constitutes the basic part of the planner’s database.

## REFERENCES


[1] Béjar, J., 1995. *Knowledge acquisition in ill-structured domains*, Ph.D. Thesis, Software Department, UPC, Barcelona, Spain.

[2] Brooks, R.A., 1986. *A robust layered control system for a mobile robot*, in IEEE J. Rob. Aut. (RA) 2(1) (1986), pp. 14-23.

[3] Droste, R. L., 1997. *Theory and practice of water and wastewater treatment*, Wiley 1997.

[4] Georgeff, M.P. and Ingrand, F.F., 1989. *Decision-making in an embedded reasoning system*, in Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit (Michigan), US, August 1989.

[5] Guariso, G. and Werthner, H. (editors), 1989. *Environmental Decision Support Systems*, Ellis Horwood - Wiley 1989.

[6] Kaelbling, L.P., 1990. *An architecture for intelligent reactive systems*, in Allen, J., Hendler, J. and Tate, A. (editors) *Readings in Planning*, pp. 713-28, Morgan Kaufmann 1990.

[7] Mc Carthy, J. and Hayes, P.J., 1969. *Some philosophical problems from the standpoint of artificial intelligence*, in Meltzer, B. and Michie, D. (editors) *Machine Intelligence 4*. Edinburgh University Press, Edinburgh, UK.

[8] Myers, K.L., 1996. *A procedural knowledge approach to task-level control*, in Proceedings of the Third International Conference on AI Planning Systems, AAAI Press, 1996.

[9] Pryor, L. and Collins, G., 1996. *Planning for Contingencies: A Decision-based Approach*, in Journal of Artificial Intelligence Research 4 (1996), pp. 287-339.

[10] Sànchez-Marrè, M., Cortés, U., Lafuente, J., R-Roda, I. and Poch, M., 1996. *DAI-DEPUR: an integrated and distributed architecture for wastewater treatment plants supervision*, in Artificial Intelligence in Engineering 10 (1996), pp. 275-85.

[11] Sànchez-Marrè, M., Cortés, U., R-Roda, I. and Poch, M., 1999. *Sustainable case learning for continuous domains*, in Environmental Modelling & Software 14 (1999), pp. 349-57.

[12] Steels, L., 1990. *Components of expertise*, in AI Magazine 11(2) (1990), 28-49.

[13] Wilkins, D.E., 1988. *Practical planning: extending the classical AI planning paradigm*, Morgan Kaufmann, San Mateo, CA, USA.