What if a wastewater treatment plant were a town of agents

Luigi Ceccaroni
Universitat Politècnica de Catalunya
Departament de Llenguatges i Sistemes Informàtics
Campus Nord, Mòdul C6, Jordi Girona 3
08034 Barcelona, Spain
luigic@lsi.upc.es

ABSTRACT

We present a potential application of a multiagent infrastructure to the domain of wastewater treatment. In this domain, each task represents an effort that contributes to the process, but the whole process and the tasks belong to separate conceptual-layers. We employ several agents, who make use of case-based reasoning, rule-based reasoning and reactive planning, and who decouple between interaction syntax and local knowledge. The discussion about their behavior leads us to take into consideration the following issues: the dynamic modification of agents' acting, and the communication based on a domain ontology, whose content varies in time. The ontology which we employ provides a set of concepts that can be queried, advertised and used to control agent cooperation. It is also the basis of the definition and the exchange of the semantics of the domain model.

1. INTRODUCTION

We used to think about processes like wastewater treatment as something under a strict supervisory control, and we developed a decision support system based on this principle for the management of this specific domain [3]. But what if a wastewater treatment plant were *a town of software agents*, capable of semi-independent action in their open and unpredictable environment?

We propose, as a supervisory system, an autonomous structure where agents generate expectations, act, make mistakes and react to the violation of their expectations, without too much control from humans. With an exception: some goals can be imposed by the plant manager to one of the agents. In this section we introduce some basic concepts about the domain of application (the wastewater treatment) and ontologies in general, while in next sections we will deal with more specific topics related to the design of the multiagent environment and with some of the issues raised by its functioning.

Wastewater treatment plants (WWTPs) are the physical element of the application domain. They serve to decontaminate wastewaters prior to their discharge into a natural body of water. The process of wastewater treatment is very complex and difficult to be reliably supervised by a technology based only on a chemical-engineering classic-control [8]. Plant supervision typically involves the following activities: diagnosis of the quality of water input, identification of treatment requirements, evaluation of economic and environmental costs, control of water and sludge cycles, diagnosis of the state of the plant, planning a course of action, management of emergencies, and interaction with users. Today, all these activities are in general executed by humans or classic-AI systems, but in the future we may see them being conducted by software agents. This shift mainly depends on the magnitude of the improvement in performance, if there is any, due to the introduction of agent technology with respect to present techniques, and it is still too early for this kind of assessment.

Ontologies have been developed in artificial intelligence to facilitate knowledge sharing and reuse. In general, ontologies provide [6]: a shared and common understanding of a domain that can be communicated among agents and across application systems, and an explicit conceptualization that describes the semantics of the data. Such common understanding facilitates an accurate communication and a sharing of meaning among software agents, and leads to other benefits such as knowledge reuse.

2. AGENTS IN TOWN

Let us label our town of agents as WaWAT[1]: a multiagent cooperation infrastructure for supervision and decision-making in WWTPs. In the distributed environment of WaWAT, agents deal with a large number of heterogeneous information sources with dynamic availability. After retrieval, this information is processed under partial human control. The main characteristics of WaWAT's agents are:

- Their *attitudes* follow the classification used in [1]. Intended goals are in some cases pre-designed and built-in by human experts or managers. In general, these imposed goals consist in the diagnosis of the state of the plant or reflects changes in treatment standards. The existence of externally established goals and the allowance of user intervention at any time are the reasons of the *semi-autonomy* of the agents.

- They are *customized*, that is, they are adapted to the particular environment of WWTPs.

- They are *continuously running* and constantly advertise information which may be relevant to other agents or to the users.

[1] WasteWater Agent Town. Historically, Wawat was a nation of fertile land between the first and second cataract of the Nile river.

- They express their knowledge, abilities and intentions through *messages*. In this way they can interact with other agents, present requests and answer queries. This cooperation requires the utilization of common concepts from a domain ontology and raises the following issues: (1) the need of a mechanism to exchange ontology knowledge, and (2) the need of a standard message-format, with meaningful structure and semantics.
- They are *dynamic* entities. Agents in WaWAT are capable of carrying data, knowledge and code as objects, and of executing code (such as rules or axioms). They can delete these objects and add new ones; they send and receive information; they adjust their capabilities to accommodate changes in the surrounding environment and in the requirements imposed by users.

In WaWAT town, taking inspiration from [7], we characterize the following agents: the supervisory agent, the case-based reasoning agent, the rule-based reasoning agent, the pre-treatment agent, the primary settler agent, the aeration tank agent, the secondary settler agent and the sludge agent.

3. MODELLING THE BASICS OF COMMUNICATION

The whole supervision of the treatment process and its multiple tasks are handled by separate agents. At the process-supervision level (handled by the *supervisory agent*), human influence on agents' goals is possible; at the task level, autonomy is allowed. In wastewater treatment, agents are often concurrently involved in multiple task and most tasks require interaction with other agents, with limited knowledge of other agents' internal state. Interactions are critical because, if they are not properly designed, the multiagent system is brittle and susceptible to failures when some part of it is modified [5]. To collaborate, agents need to communicate (data availability, queries) with other agents, who can provide them with needed information to accomplish specific tasks and to manage plans at runtime [2]. They need: (1) to generate or access a *model of the domain*, (2) to avoid *terminological confusion* and (3) to *describe the services and the data* they lack or they can provide. An important step towards developing a successful communication in the WWTP domain is represented by the introduction of a shared and focused ontology, which (1) provides a common vocabulary for describing the domain and the agents' capabilities (involving various kinds of knowledge and services), and (2) allows to codify the data and the services into a set of agreed terms that can be shared among agents and users. In this section, we characterize more exhaustively the issues which stand behind the introduction of an ontology for wastewater treatment [4] and the determination of a syntax for ontology representation.

The separation of domain-specific aspects, such as the semantics of the content, from the syntax related to interactions, such as communication protocol and structure of the messages, facilitates the reuse of standard interaction protocols in our application. Because of the very specific domain-ontology, at the present time the system used in WaWAT is competent in its narrow and well defined domain, which include the quickly changing environment of WWTPs, but it would be of little use if the domain changes.

3.1 Ontology definition, representation and implementation

In WaWAT, many different agents work together. The disciplines of human expertise involved are: artificial intelligence, chemical and environmental engineering, and microbiology. Without defining a precise, unified, complete and consistent terminology, each expert or agent would use a distinct vocabulary. There would be no rules helping in the use of each term. Synonyms could cause confusion. Some terms could be used by different agents with similar, but not identical, meanings. All these issues justify the introduction of an agreed-upon ontology, which can be used in a formal context by agents and humans, as a basis for communication.

The ontology of our case study (WaWO [4]) is based on a model that we defined for WWTPs, and describes the meaning of each term used in the wastewater domain in a precise and unambiguous manner. Our general model of the domain basically grounds the course of action on the evaluation of the current state of the plant; therefore one of the primary task of the agents is the diagnosis of this state. This is accomplished using case-based and rule-based reasoning.

The WaWO ontology is designed and built following the mainstream ideas about ontology construction [10, 9] and is a hierarchically structured set of terms and a set of axioms describing the domain of wastewater treatment. Concepts at the highest level refer to: wastewater, descriptors of environmental variables, treatment plants, states of a plant, biological living objects, input-output devices, planning, actuators, bodies of water.

To easily communicate, agents have to be designed in a way such that they can readily use the domain ontology (or, alternatively, the ontology has to be built in a way such that it can be readily interpreted by agents). Agent capabilities, to be expressed, need the modelling of at least two aspects: the one of the knowledge over which services operate and the one of the syntax of messages and available services.

Having WaWO been defined and implemented at the Ontolingua KSL Server, we implicitly chose a *frame-based formalism* for the Ontology, and this constraints the knowledge representation. The central modelling primitives of frame-based and object-oriented approaches are classes with certain properties called attributes. These attributes do not have a global scope but are only applicable to the classes they are defined for and the *same* attribute name may be associated with different value restrictions when defined for different classes (i.e., inherited attribute values can be overwritten).Frames are well suited for the modelling of control systems and physical-biological processes, and object-oriented design can be used also to model axioms about wastewater treatment in general, causal relations and non-causal energy flows.

The WaWO ontology links classes and individuals at least with taxonomic relations and provides a terminology that each agent involved can understand and use. It includes a large collection of environmental-resources concepts related to wastewater, such as information about various contaminants and the microorganisms living in WWTPs. Ontology objects include classes referring to the general wastewater domain, to decision support, as well as to a detailed microorganisms taxonomy. In this way, the ontology integrates the knowledge about WWTP management with the current knowledge about microbiological processes.

Each component of the plant has its own local agent that makes the information collected at sampling points available to other agents in terms of the WaWO model. Agents can selectively save data locally and become themselves dynamic direct information sources. Apart from making information accessible, each agent advertises several other capabilities (such as the possibility to be monitored or to be queried) specifying the classes and the slots of the ontology that have data in the agent, and whether the services offered are direct (information accessed locally within the agent) or indirect (information retrieved from other agents).

In order to automate agent cooperation and to attain a satisfactory supervision of the plant, it is *necessary* to provide advertisement messages with a meaningful semantics. WaWO implements the semantics in a set of *axioms* that enable: (1) to describe wastewater processes, (2) to automatically deduce the answer to several questions about the wastewater domain, such as cause-effect questions, and (3) to solve diagnosis impasses. Through the use of axioms, agents are able to represent two kinds of *cause-effect relations*: they can associate microorganisms to the problematic situations that these cause, and they can associate the actual state of the plant to the actions that need to be performed in order to reach the normal state from that actual state. In WaWO, axioms connect the descriptors to the states of the WWTP. More general states are linked to the descriptors by more general axioms, while less general states or substates are linked to the descriptors by more specific axioms. In the case of microorganism descriptors, this means linking the state of the plant to higher or lower levels of microorganism taxonomy. (Example: in case of presence of two microorganism *species* and lack of axioms (or discordant axioms), the connection to the state of the WWTP is done at the taxonomic-*class* level.) In this way, we endow the ontology with *inference processes* and we improve the diagnosis capabilities of the agents using it.

4. CONCLUSIONS

We designed a multiagent system applied to the domain of wastewater treatment. The system involves interactions among a relatively small number of agents. Knowledge about the domain and expertise about process management are modeled by a specific ontology, which was developed under the frame-based formalism of the Ontolingua Server. This ontology is used to enable the agents (1) to generate or apply a *model of the domain*, (2) to prevent *terminological confusion* and (3) to *characterize the services and the data* they lack or they can provide.

Acronyms résumé

WWTP: WasteWater Treatment Plant.
WaWAT: WasteWater Agent Town, the multiagent system.
WaWO: WasteWater Ontology.

5. ACKNOWLEDGMENTS

This research is supported by the Generalitat de Catalunya (grant 1997FI 00630 UPC PG). The author wishes to thank the reviewers for their comments, critics and proposals.

6. REFERENCES

[1] F. Brazier, B. Dunin-Keplicz, J. Treur, and L. Verbrugge. *Modelling Internal Dynamic Behaviour of BDI agents*, volume 1760 of *Lecture Notes in Artificial Intelligence, LNAI*, pages 36–56. Springer-Verlag, 1999.

[2] A. Cassandra, D. Chandrasekara, and M. Nodine. Capability-based agent matchmaking. In C. Sierra, M. Gini, and J. Rosenshein, editors, *Proceedings of the Fourth International Conference on Autonomous Agents*, 2000. Barcelona, Spain.

[3] L. Ceccaroni. Integration of a rule-based expert system, a case-based reasoner and an ontological knowledge-base in the wastewater domain. In *ECAI 2000 - W07: Binding Environmental Sciences and Artificial Intelligence (BESAI 2000)*, 2000. Berlin, Germany.

[4] L. Ceccaroni, U. Cortés, and M. Sànchez-Marrè. WaWO - An ontology embedded into an environmental decision-support system for wastewater treatment plant management. In *ECAI 2000 - W09: Applications of ontologies and problem-solving methods*, 2000. Berlin, Germany.

[5] K. Erol, J. Lang, and R. Levy. Designing agents from reusable components. In C. Sierra, M. Gini, and J. Rosenshein, editors, *Proceedings of the Fourth International Conference on Autonomous Agents*, 2000. Barcelona, Spain.

[6] D. Fensel, I. Horrocks, F. Van Harmelen, S. Decker, M. Erdmann, and M. Klein. *OIL in a nutshell*. Lecture Notes in Artificial Intelligence, LNAI. Springer-Verlag, 2000.

[7] M. Sànchez-Marrè. *DAI-DEPUR: an integrated supervisory multi-level architecture for wastewater treatment plants*. PhD thesis, Software Department, UPC, Barcelona, Spain, 1995.

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

[9] M. Uschold. Knowledge level modelling: concepts and terminology. *The Knowledge Engineering Review*, 13(1):5–29, 1998.

[10] M. Uschold and M. Gruninger. Ontologies: principles, methods and applications. *The Knowledge Engineering Review*, 11(2):93–136, 1996.