Agent-Based Decision Support Framework for Water Supply Infrastructure Rehabilitation and Development

D.N.Davis ☿

Artificial Intelligence & Pattern Recognition Group, Department of Computer Science,

University of Hull, Kingston-upon-Hull, HU6 7RX, UK.

D.N.Davis@dcs.hull.ac.uk

Fax: +44 (0) 1482 466666

Abstract

The broad platform of our research addresses the possibilities offered by the concept of complete agents as a theoretical basis for designing computational architectures of use in modelling intelligent behaviours. One area that we are investigating is the application of the computational metaphor of agents to decision support systems. Here we present the current state of our progress in developing a framework of use in making decisions about water supply infrastructure rehabilitation and development. The framework currently supports database reclamation, data warehousing, water mains pipe-failure prediction and strategic overview information based on customer complaints and chemical analysis of supplied water. Further aspects of the framework have been developed for other problem domains. This ongoing research draws on methodologies from work on agents, knowledge engineering, data-mining and knowledge discovery.

1. Introduction

Research into agents and the development of related technology (Ferber, 1999; Muller & Wooldridge, Jennings, 1996; Wooldridge & Jennings, 1994) has grown dramatically since the 1980s as domains suitable for their application have emerged. One of the reasons is the uptake in information technology across many different domains, for example the business sector, industry, the health and utility services. Allied to this is the realisation that many of the emerging problems in sophisticated information technology frameworks can be solved only

☿ Research commenced while a member of Artificial Intelligence Group, Staffordshire University, UK.

through the development and application of more intelligent systems. Current research into agents is starting to demonstrate many potentially useful qualities, for example in information sorting such as e-mail and Internet applications (Maes, 1994) and patient monitoring (Hayes-Roth, 1995). As AI researchers address the limitations of their communities' earlier work (for example, purely symbolic systems and the lack of rigour in producing AI systems), the issues related to complex and/or intelligent systems change. It is suggested that insights into the nature of solutions for these challenges may be offered through the consideration of what type of artificial system (agent) could possibly augment (or even model) the human qualities studied in cognitive psychology (Port & Van Gelder, 1995; Sloman, 1996).

This paper presents a number of perspectives on our ongoing work in developing MADSS (Multiple Agent Decision Support System). MADSS is an agent-based decision support framework that can be of use in supplying decision-enabling information in a number of domains. We suggest that behaviours useful in solving problems, associated with specific information domains, results from designing specific architectures for particular types of agent communities. The grain of the design and architecture varies with the domain and task. Here we describe how MADSS can be of help to U.K. water companies in their water mains rehabilitation decision making. We highlight the type of challenges that had to be met in one specific project. The initial aim of that project was to develop a prototype expert system that would model the current decision making processes adopted by one U.K. water company and their associates. This prototype system was to be developed within a framework that would allow the addition of further decision support capabilities. We have, since then, investigated further the nature of data-mining for this domain, the application of KADS (Knowledge Analysis and Documentation System) and KQML (Knowledge Query Meta Language) to agents, different types of multi-attribute decision making and issues related to adaptive natural language interfaces for these types of systems.

2. The Domain and Its Challenges

Water mains rehabilitation is a major challenge for water supply companies, with a large annual and protracted budget, as to be expected for any company that services a fundamental resource for a number of counties in the U.K. mainland. There are a number of important constraints that impinge on these financial and resource decisions. For example regulatory requirements which are monitored by government established watch-bodies and factors related to customers and users that affect which areas (i.e. fundamentally which pipes) are to be maintained or replaced in any one year. Furthermore the decisions made in current and previous years affect the decisions to be made about future water supply infrastructure maintenance. This application area is therefore not only very important but also quite complex and is in need of systems that facilitate more consistent and effective strategic decisions.

The framework being developed models and augments different parts of the current human strategy used in making rehabilitation decisions; i.e. it embodies some of the expertise of the water company engineers, their consultants and their information handling models. However, this information handling expertise has been endowed with a degree of extra flexibility that for example allows the manipulation of the ranking processes used in the company's interactive database model. This enables the production of multiple ranked lists of supply regions from any given set of data. This allows alternative rankings of the potential rehabilitation areas, and so allows the impact of the current rehabilitation strategy and pertinent 'what-if' scenarios (e.g. alternative rehabilitation strategies) to be explored. The research and design framework extends this with further capabilities, additional heuristics, statistical models and other decision support facilities. In producing the first delivered system, a number of challenges were encountered which may be indicative of the challenges researchers will have to meet in analogous urban system domains when developing intelligent decision support tools, particularly if these systems are to work within an existing IT infrastructure.

The water company required a flexible system to assist water supply experts in analysing the impact of implementing any changes to the current water mains rehabilitation strategy. The domain encountered seemed, at first, to be a classic scenario for the application of expert systems but due to the emerging complexity of the application domain, an agent based approach to knowledge engineering was identified to be a more fruitful avenue to pursue. The proposed system would need to access a number of existing and new databases held on PCs, plus other potential sources of information, for example a Geographic Information System (GIS) database. In the short term the framework must be able to support a number of functions each of which could be modelled individually using relatively uncomplicated expert systems.

3. Expert Systems and Knowledge Engineering

A classic expert system aims, by definition, to model human decision making expertise in some specific domain. It has long been recognised (Englemore & Morgan, 1988; Medsker, 1995) that to produce such a system is challenging and typically requires much more than the naïve knowledge base plus reasoning mechanism plus interface approach. This section looks at how to model the active nature of knowledge as it is typically used in some problem-solving or decision-making domain. We contend that the appropriate perspective for producing any such system fits within the bounds of knowledge engineering. That is, a knowledge engineering project may contain the processes by which an expert (or perhaps more appropriately a collaborating multi-agent) system is produced, but that the (*mere*) production of an expert system is a less encompassing challenge.

KADS (Schreiber & Wielenga, Breuker, 1993) is knowledge engineering methodology in which the main concern is with the two phases of analysis (for example product definition) and design (such as system and module-level design). The central theme is that of modeling the behaviour of the intended knowledge-based system, the context within which it will work and the framework for more general requirements and constraints. While KADS can be criticized as being large with the resulting tendency of being difficult to understand and apply, it provides a useful set of models for sophisticated knowledge engineering projects.

KADS identifies three levels of models. At the process level, the process model identifies the tasks involved in the domain, the nature of data flows and stores, and the assignment of ownership of tasks and data stores to agents. At the system level, the co-operation model describes in detail the interactions between the system and external agents, and how the *internal agents* (i.e. the components of the system) interact. The co-operation model can be used to separate the user task model and the system task model. It allows the knowledge engineer to distinguish between interactive knowledge and the different types of knowledge wholly internal to the system. The expertise level corresponds to an expertise model. This divides the task of describing the expertise level of the system into a number of supportive tasks. The four-layer framework for expertise consists of the domain, inference, task and strategic layers. The domain layer is comprised of static or slowly changing knowledge describing concepts, relations, and structures in the domain. The inference layer reformulates the domain layer in terms of the different types of inferences that can be made. The task layer defines knowledge about how to apply the knowledge in these two layers to problem-solving activities in the domain. The strategic layer defines how to select appropriate problem-solving capabilities for specific types of task.

Within this framework knowledge engineering becomes a structured search for appropriate task, inference and knowledge models. This is performed through the use of analysis (the process of generating abstract descriptions of the inference patterns) and design and coding (a process of implementing the details of those patterns).

4. Agents and Knowledge Engineering

Agents offer the flexibility to integrate many different categories of processing within a single system in much the same way as blackboard systems. However, the system designer need not necessarily require a full specification of intended final system in the initial stages. The adoption of an appropriate design methodology gives the application builder the opportunity to design a very flexible and incremental system, that should be responsive to the changing and developing requirements and needs of the end user. There are a plethora of agent definitions (Ferber, 1996; Franklin & Graesser, 1996; Muller & Wooldridge, Jennings, 1996; Wooldridge & Jennings, 1994). These range from descriptions based on a functional analysis of how agents are used in technology to far more ranging expositions based on different interpretations of the role and objectives of artificial intelligence and cognitive science (Sloman, 1996). Artificial intelligence is a very diverse field and agents are used as metaphors for work in many areas. This can lead to confusion and devalue the term agent, as other authors have noted (Franklin & Graesser, 1996). This section clarifies the framework of the agent research being highlighted here.

A seemingly endless list of agent attributes are possible reflecting these different definitions and the intentions of their creators, and include factors such as intentionality, autonomy reactivity, flexibility, communication, learning, self-actuation etc. Two plausible but very different agent definitions are:

Definition 1: *An agent is an integrated computational entity with intentionality and some degree of autonomy (Franklin & Graesser, 1996).*

Definition 2: *An agent is a synthetic entity that enables us to study, at a computational, design and theoretical level, what a mind could and can be (Davis, 1998).*

The first definition equates to the idea of *weak agents*, i.e. agents as (intelligent) information processing systems and is quite open to extension and interpretation. The second definition equates to the notion of *strong agents*, i.e. agents as computational cognitive models that explain and/or simulate, to some degree, reported findings and theories in cognitive psychology, or some other study of minds (or life). The work, reported on here, is very much application driven and therefore relates to the first definition. However, the second more encompassing definition has a bearing if we consider an expert system to be an artificial system capable of reasoning about a category of problems in a way that is analogous to the way a human expert resolves similar problems.

A particular class of agent system that is of relevance here are those described as tightly coupled. These can be seen as building on the ideas expounded by Medskar (1995) but firmly rooted in agent technology. Such systems can be viewed as a development of the ideas embodied in blackboard systems, where multiple agencies, each with a set of designated tasks contribute in the formation of solutions to complex tasks requiring many kinds of representation and processing categories. For example through a consideration of the processes outlined within the KADS methodology, we can define for any specific task, local to one agent four types of knowledge: static knowledge defining the ontological decomposition of the agent's task and its problem domain; dynamic knowledge which provides the heuristic delineation of the agent's task and problem domain; strategic knowledge which provides the control definition of the agent; and interactive knowledge whereby the communication capabilities of the agent are described. By identifying the appropriate (active) representations for this task-defined knowledge, we can then identify the appropriate task-specific agent architecture. At a more global level, a similar type of analysis provides a means to identify the most appropriate framework (or the society) of agents to meet the system's overall requirements.

5. Analysis and Design

There are a number of qualitatively different tasks that the agent framework must be capable of supporting. Some of these tasks are generic to decision support systems, while others are specific to the application and its domain. The system must be able to accept tasks from and present the results of its reasoning to managers and executives who may have a well defined task in mind but may be incapable of specifying this in a (formal) computational language. Typically the decision supporting information they require can be described in terms of a combination of everyday language and a domain specific language; i.e. there is a need for an intelligent interface agent capable of learning the language of the user(s) of the framework. The system needs to be able to manage several databases and combine information held in those databases in a form of use to the interface agent and other reasoning mechanisms; i.e. there is a need for a data-warehousing agent. The framework needs to be able to supply summary information related to different categories of information, available from the data warehousing agent, whether specified in terms of geographic categories or more domain specific terms; i.e. there is a need for a data mining agent. There is a range of domain specific reasoning tasks required of the framework. This type of task includes, for example, identifying general patterns of information (as generated by a data-mining agent) and also identifying the valid (and probably quite important) exceptions to these patterns. In statistical terms, these out-layers are important indicators of highly specific infrastructure problems, such as the identification of faults in one water supply pipe to a small number of houses within a region of well-maintained supply infrastructure. Other types of task involve producing supply profiles by allowing the user to define ranking criteria, related to problematic water supplies, based on information present in decision-enabling databases and then applying this standard across the information pertaining to the entire supply zone.

Following many meetings and exchanges of documents, a detailed specification was produced identifying the types of data, knowledge, internal and external factors likely to be of relevance in a full decision support tool. Figure 1 provides a high level summary of these factors. The four headings used to group the identified factors and constraints are:

- **Engineering factors**, which relate to the company's policy on issues such as quality of supplied water, response to customer complaints, water mains leakage, decisions to reline or replace sub-standard water mains, and geographic factors within the area supplied by the water company. These factors can be used in assessing whether rapidly expanding urban areas with few current water supply problems are more important sites for infrastructure development than relatively uninhabited (and population static) rural areas with a higher complaint rate per household.

- **Technical factors** typically define the viability of a water mains rehabilitation plan allowing a more consistent decision-making process. Case libraries and archive databases related to prior and existing rehabilitation decisions, which allow case studies extending into the company's history, influence current decisions. For example, if an urban community is to be expanded, does the water company wait on the infrastructure investment until the needs of that community and the effect of that community upon the existing water distribution infrastructure are known, or does the company take a more proactive but more uncertain stance. Archive databases are of use in determining the effectiveness of past solutions, and again allow for a more consistent water mains rehabilitation and development policy.

- **Organisational policies** involve such factors as asset and risk management. For example, what happens if the company does not invest in the water supply infrastructure of a rapidly expanding urban community? Demographic policies - for example does the company respond more favourably to certain levels of rehabilitation in certain areas. Operational costs and policy and budgetary constraints - can the company afford to let its contractors perform certain levels of rehabilitation? Organisational policies related to the rehabilitation criteria - for example does the company respond more favourable to taste, colour and contamination measures given the constraints of other factors.

- **External factors** affecting the company's rehabilitation strategy such as overriding policy statements by the Department of Water Industries, OFWAT (the regulatory body for the water industry) which prescribe legal constraints and company priorities, customers, the company’s consultants and contractors responsible for carrying out any rehabilitation.

The computational modelling of all these aspects requires several types of functionality and knowledge handling beyond an orthodox (single or multiple knowledge base) expert system, and more appropriate to the sophistication associated with blackboard systems and multiple agent systems. Furthermore as an incremental approach to developing a full system was required and the initial implementation did not require an orthodox rule-based system, we thought it wise to consider alternative technologies based on the distributed approach to blackboard systems, such as multiple agent systems. This was also driven by an interest in the dynamic systems approach to cognition. From this perspective, cognition is viewed as the changing focus of a number of interacting sometimes extant, sometimes dormant processes. Any particular reasoning or problem-solving capability can be tuned to a number of related, but different tasks, and that the achievement of complex tasks requires the use of many qualitatively different capabilities. KADS and agents can be combined to achieve a computational correspondence to this metaphor within a domain tractable to knowledge engineering. Through the type of analysis summarised at the start of this section, we can begin to identify the categories of tasks required to support decision making in the domain, and for any one specific task the static, dynamic, strategic and interactive knowledge required to perform that task.

6. Explanation of the framework

MADSS is a developing generic framework for decision support. As we apply the knowledge engineering philosophy behind MADSS to specific scenarios, specific designs result. Figure 2 depicts a collaborative society of agent based processes for managing decision support tasks in this domain. It also enables an incremental progress from a partial (but functioning) system to a full implementation based on the KADS driven analysis and design process. The figure differentiates between data exchange, for example the transfer of data from the databases to the task specific agents, and agent communication. For example a Strategy Agent request that the DataWarehouse Agent update the Decision Enabling Database. Agent communication is more encompassing than data exchange, and encapsulates the information of a data exchange between any of the agents and the databases. Within this overall framework exist a number of task specific agents.

The Interface Agent accepts tasks from and supplies solutions to the user in a readily understandable format. The Information Agent is responsible for the internal description of any given task and ensuring that the DataWarehouse Agent responds to the needs of the other agents. Information, relevant to the system's current task, is held by the Information Agent which exchanges information with all the other agents, requesting further information from or depositing new information with the DataWarehouse Agent as required. The Information Agent must have knowledge that models the functions and task capabilities of the other agents. The DataWarehouse Agent is responsible for ensuring that the data used within the framework is reliable, and available as the integrated information held in the Decision Enabling Database. From the perspectives offered by blackboard system design, the Decision Enabling Database together with the Information Agent act as a dynamic (i.e. active) blackboard. The Strategy Agent performs specific ranking processes based on information about water supply zones within the company's regions. While this agent may make use of one of a number of default decision strategies, a custom strategy can be adopted based upon a design arising from the interactions of the user with the decision support tool via the Interface Agent. The Predictor Agent can be used to supply information to support water mains rehabilitation based on structural information about supply pipes. The Constraints Agent reasons about the viability of any tentative rehabilitation decision on the basis of the geographical and policy type factors listed above (and shown in figure 1). The influence of cognitive science theories is also evident in the grouping of agent-based processes into the three broad categories of: sensing and acting (the Interface Agent); task-related cognitive memory (e.g. the Strategy Agent); and Long Term Memory Management, i.e. the agent(s) responsible for managing the database systems used by the decision support tool.

Early development work (Sharp & Edwards, Dean, Davis, Bancroft, 1998) required extensive data cleaning on the strategic database which contains demographic and customer related information. Third-party developed databases were not compatible with standard data communication protocols (e.g. ODBC), requiring unexpected and prolonged involvement in data cleaning - one of the more time-consuming processes involved in data mining. Subsequent work (Agar, 1999), on developing the Data Cleanser Agent, has ensured that the information present in this database is not only ODBC compliant but consistent and complete at design and implementation levels. We have so far failed to develop an adequate generic automatic data-cleansing agent. In this and other data-mining exercises we are finding that large portions of the supporting databases need to be cleansed interactively. Automated data cleansers typically disallow a substantial amount of data from becoming available to the data-mining processes. This is in fact one of the major challenges that confronts us in developing intelligent systems that need to work with legacy systems. We suggest that as AI practitioners look to other computational domains related to the environment and urban systems, such challenges will need to be confronted as they try to retrieve expertise locked away in archive and legacy systems. This will be particularly so where these knowledge engineering projects address large scale decision making situations such as utility management and consistent decision making over time in large organisations.

Similar problems are also to be found in the area of data-mining and knowledge discovery, and various authors have suggested the use of appropriate methodologies. One such methodology (Fayyad & Piatetsky-Shapiro, Smith, Uthurusamy, 1996) comprises of six stages. The first, *problem analysis*, requires the answering of questions such as: Is the domain or application area suitable for tackling using data-mining techniques? And what are the requirements of the data mining exercise? The second, *data preparation*, requires the transformation of data, of whatever format and nature, into the format required for the subsequent stages. The perspectives originally placed on the domain and its data models by the designers of a legacy system are typically very different from the perspectives placed on the domain and the legacy system when involved in data mining and new decision making scenarios. *Data exploration* relates to the exploration and visualisation of the patterns in data in readiness for *pattern generation*. The pattern generation stage includes not only for example the generation of decision trees, or the training of classification mechanisms (whether automatic or interactive), but also the validation and interpretation of discovered patterns. The fifth stage, *pattern deployment*, involves the integration of generated patterns into decision support systems, the production of reports or the filtering of data for further processing. The last stage, *pattern monitoring*, is a proactive ongoing situation that serves to detect shifts in the applicability of currently generated patterns at the earliest possible time. The methodology, as a whole, is incremental and iterative, with many possible cycles through the different stages over the lifetime of a data-mining project.

An analysis of the water domain framework and this methodology show a correspondence between the tasks of the individual agents and the individual stages of the methodology. The Data Cleanser and DataWarehouse Agents effectively function as a data preparation stage, ensuring that the data entering the Decision Enabling Database is fit for purpose. The Strategy, Data Mining and Hot Spot Agents allow pattern generation whether general patterns about a supply area, or valid exceptions to such patterns. The Information Agent together with the DataWarehouse and Interface Agents allow the user to opportunity for data exploration. The Strategy Agent together with the Predictor and Hot Spot Agents can make use of generated patterns to support infrastructure rehabilitation and development decisions. Finally various combinations of these agents can be used to monitor incoming information and so support proactive strategic decision making.

7. Data and Knowledge Representation

The data as originally supplied by the water company was held in a Microsoft Access Database. This data, after cleaning provides us with a further Access database. This is accessed from an object-oriented expert system shell (Kappa-PC) via ODBC. Requested entries from the database are collated as objects in Kappa-PC. The domain ontology comprises of 13 object classes as shown in figure 3. As a database is opened, a new instance of the class database is created. All information about the state of the database and the nature of the communication with that database is kept there. As information about supplied properties is read into the framework, it is associated with an instance of the object class ParishZone. All information from the water company's database is initially kept in instances of these two object classes. If the user requests that information be presented in terms of a different supply zone (e.g. DMA or WIS) then the information held in the instances of ParishZone objects is summarised within newly created instances of the required object classes.

If the user selects to create a custom ranking process (see section 9), a new sub-class of the Custom object is created. When the system performs a ranking process, an instance of the selected Process is created, making use of information in all the ParishZone instances, keeping all generated information about the analysis. This information is subsequently displayed to the user (see figure 6). All algorithms, and heuristics used in the ranking process are kept as methods and monitors associated with the parent class Process. The Interface or Strategy or Information Agents run a specific ranking process by triggering the appropriate method associated with a specific Process instance. Constraints about data type and value range for entries in the database are held as monitors which effectively act as filters on incoming information (retrieved from the database) as it is read into the appropriate ParishZone instance.

8. Communication and Knowledge Exchange

Agents in MADSS communicate with each other using a number of mechanisms. The communication protocol depends upon the type of agents involved and the information and knowledge being exchanged. Agent to database communication is ODBC based using SQL requests and commands. This protocol is relatively straightforward. Agents connect to an active database via ODBC, and phrase their information requests using SQL. The existing water rehabilitation system makes use of this, and we have also investigated how it can be used in other domains with a natural language interface (Simmonds, 1999). We are currently investigating different media access protocols in e-commerce scenarios.

Agent to agent intercommunication for the water domain remains at the design level. However we have investigated the nature of this type of communication in a related (mobile-agent) project that allows multiple attribute decision making about air-flight ticket reservation (Davis and Cailleteau, 2000). Here we rephrase that work in terms of the water domain. An agent communication language is a protocol that enables agents to co-operate and communicate with each other. No real standard is defined at this time but two inter-related communication languages (KQML and KIF) are used world-wide.

KQML (DARPA, 1993) is an agent language and protocol for exchanging information and knowledge between software agents. KQML is both a message format and a message-handling protocol to support run-time knowledge sharing among agents. Agents can be heterogeneous, each with their own particular data representation. KQML allows them to exchange information in spite of this potential incompatibility. KQML operates at several levels: in terms of transport, i.e. how agents send and receive messages; the language lexicon with an agreement about the meaning of the messages; policy, such as what is the structure of the conversation; and architecture, i.e. how to connect systems with the appropriate constituent protocols. An important component of a KQML message is the performative, which defines the focus of a message. Example performatives include *tell, perform*, *ask-if* and *reply*.

KIF provides a syntax for the message content. It has declarative semantics - the meaning of expressions in the representation can be understood without appeal to an interpreter for manipulating those expressions. It is logically comprehensive and provides for the expression of arbitrary sentences in first-order predicate calculus. It also provides an explicit means to describe knowledge representation schemes, the representation of non-monotonic reasoning rules, and the definition of objects, functions, and relations.

A full implementation of KQML and KIF would be too complicated for this framework. Based on our preliminary work with mobile agents, a limited language set would be required. We retain the readability of KIF-KQML by using understandable performative names and structures. In a full implementation of the decision support framework, the message shown in figure 4 would cause the Data Warehouse Agent to search its databases for the supply-interruption risk associated with a specific water zone, and to respond with one of the specified responses. If that agent responded with a *no-answer-to-your-request* message (using the *null* response), the Information Agent would need to communicate with some other agent (e.g. the Predictor Agent) to reason about the specified zone.

Figure 5 shows how the agents interact for the scenario described below in section 9, given that the Decision Enabling Database has already been built. For this particular situation a sequence of successful communication events occurs over three levels. In step1 the user requests the

Interface Agent for a ranking of three zones on three named ranking processes and a further process that combines them. Step 8 is the final response with which the Interface Agent displays the requested information for the user. At one level of nesting, in step 2, the Interface Agent asks the Information Agent to deal with the request and to reply when the request is complete. In step 7, the Information Agent collects the required data from the Decision Enabling Database and replies to the Interface Agent message of step 2. As the requested information is not initially in the database, a further level of communication and tasks transpires. In step 3, the Information Agent asks the Strategy Agent to perform and store the three standard ranking processes on all the zones. The Strategy Agent performs and stores the three standard ranking processes on all the zones, and responds with request complete. In step 5, the Information Agent asks the Strategy Agent to perform and store the custom ranking process on the three named zones. In step 6, the Strategy Agent performs and stores the custom ranking process on the named zones, and responds with request complete. If the user were now to request a further custom ranking over the data produced through steps 3 and 4, a similar communication sequence minus those two steps would occur. If the user were to request information that already existed in the database, only steps 1, 2, 7 and 8 would occur. In other situations more asynchronous communication and processing can occur.

9. Example Urban Scenario

In this section, we give an example of how the data, information and knowledge are used in the prototype SWIMM system. For reasons of confidentiality, the names of the water company and the UK areas used in the example have been altered; all other information remains factual.

In any one financial year, Corniche Water Supplies have £1.25million to invest on their water supply infrastructure. The factors and pressures involved in the decision-making process have already been covered, at least superficially, in section 5 and figure 1. The water supply company have to offset the costs of unexpected emergency cover (e.g. a juggernaut causing a severely fractured urban water mains affecting hundreds or even thousands of homes) against such factors as regular water supply maintenance (for example, the flushing of water supply mains directly into the fluid discharge system – a direct loss of a fundamental resource), and the rehabilitation of a predominantly one hundred year old water supply infrastructure. Within this scenario, the company has to balance the needs of non-urban customers (for example the five farms on the Moorlands) against the needs of a rapidly expanding town community (for the example 26% population increase over the last five years in the coastal resort of Pleasantville) with the associated increasing demands on that water supply infrastructure. This situation is further compounded by the ever-increasing number of complaints about water quality from a prominently senior-citizen populated town (the slowly expanding town of Quietville). The complaints centred on water quality related to the latter settlement arise from the fact that the taste of the local water differs from that in the region from which the translocated population recently moved. A consistent balance of these factors over time, with a changing company personnel, provides a fundamental reason for the use of the type of system we have described.

Say for example, that over a period of several months Corniche Water Supplies receive repeated complaints from the five farmers on the Moorland, a considerable number of water quality complaints from Quietville and a proportionally larger number (per population head) of interruption supplies from Pleasantville. Assuming low priority pressures elsewhere, after all the government body water supply regulations (as regulated by OFWAT and DWI) need to be adhered to otherwise the company, in principle, ceases to trade, the managers involved in the decision about the relative importance of these three regions and how they distribute available financial resources, face a relatively complex decision-making situation. The attributes used in judging the importance and urgency of each case differ. What our framework supplies (even in its prototype form) is a means of ranking each of these three regions on the attributes appropriate to each case. The three cases can be ranked on each case and their relative case analysed across their ranking over the three separate cases.

Figure 6 shows an abridged example report, while table 1 provides a ranking for each water supply region over each case. Ranking process 1, used for water quality purposes, relies on just three equally weighted factors: water discolouration, turbidity, and taste complaints. Ranking process 2, used for supply interruption purposes, relies on two equally weighted factors: supply interruptions and low water pressure complaints. The third ranking process is a standard process, relying on eight factors: aluminium count (9%), iron count (27%), manganese count (9%), turbidity (18%), discolouration (9%), supply interruptions (10%), taste complaints (9%) and low pressure complaints (9%). These ranking processes operate over all supply zones of the water company. When normalised the cases for refurbishment of PleasantVille and QuietVille are very much stronger than for the farmers of Moorland. However, over the last five years both these places have received substantial investment in their water supply infrastructure, while Moorland has received very little. On further investigation it is found that the supply interruptions in PleasantVille are localised to water mains near highways used by heavy freight lorries. While many of the water quality complaints in QuietVille relate to newly arrived residents, there is evidence of water contamination arising from old pipes. As the water mains of Moorland are very old, accessible and relatively short in length, they can be readily replaced. Certain sections of water mains in QuietVille are replaced and a greater proportion of water mains are relined. The pipes identified as being at risk from breakage from roadside vibration in

PleasantVille are re-laid using more advanced vibration deadening materials, and the local council informed of the problem related to continued heavy road traffic.

10. Future Work

Currently MADSS is at the computational level, a set of related but almost independent systems, i.e. it is loosely coupled and not an integrated entity. The same is true for the architecture developed for the water domain project. The Interface Agent exists in several forms allied to other implemented systems, and in its most ambitious form as a hybrid Interface-Information agent with natural language and learning capabilities (Simmonds, 1999). The Data Mining Agent currently works only on the database used by the Strategy Agent. A separate database is integrated into the Predictor Agent. We have yet to start implementation work on the Constraint Agent. This will require considerable work on the supporting data and information in its underlying GIS database. More research is required to develop more autonomous DataWarehouse and Data Cleanser Agents; the present system requires time-consuming interaction (Agar; 1999). We have yet to implement the HotSpot Agent. An alternative is to continue to develop these agents as loosely coupled systems within a decision-making framework with a limited amount of user-lead run-time system-interaction. Only future collaboration will ascertain the appropriate path for this domain. These issues will, however, be addressed in applying this framework to other domains, where other constraints may require a different avenue of research and subsequently a different development path.

A major challenge in integrating these systems lies with the communication that will underpin the framework. We have used a portion of the KIF-KQML framework in related work, but the design and implementation overheads can be excessive, given that we do not require the full functionality of that framework for any specific application. Only a small set of the possible performatives and syntax are required for the tasks that any specific application is intended to support. This raises problems for the generic MADSS system. It may well be that MADDS will remain a theoretical framework that allows us to design and implement specific architectures for specific domains and tasks within those domains.

11. References

Agar, F., (1999). Data Reclamation For Data Mining, MSc Thesis, School of Computing, Staffordshire University.

Darpa Knowledge Sharing Effort, (1993). *Specification of the KQML Agent-Communication Language.*

Davis, D.N., (1998). Synthetic Agents: Synthetic Minds*? Frontiers of Cognitive Agents, IEEE Symposium on Systems, Man and Cybernetics*, San Diego, IEEE Press.

Davis, D.N., & Sharp, B. (1999). *An Agent Framework for Decision Support in the Water Industry*, New Review of Applied Expert Systems, Volume 5. Taylor-Graham Press.

Davis, D.N., & Cailleteau, L, (2000), Designing Agents for Electronic Commerce, (To be published).

Englemore, R., & Morgan, T., (1988). Blackboard Systems, Addison-Wesley.

Fayyad, U.M., Piatetsky-Shapiro, G., Smith, P., & Uthurusamy, R. (1996), *Advances in Knowledge Discovery and Data Mining*, MIT Press.

Ferber, J., (1999), *Multi-Agent Systems,* Addison-Wesley, ISBN 0-201-36048-9.

Franklin, S.P., & Graesser, A.G., (1996). Is It an Agent, or Just a Program?: A Taxonomy for Autonomous Agents, In: *Intelligent Agents III*, J.P. Muller, M.J. Wooldridge and N.R. Jennings (Editors.), Springer-Verlag, Heidelberg.

Hayes-Roth, B., (1995). An architecture for adaptive intelligent systems. *Artificial Intelligence*, 72.

Maes, P., (1994). Agents that reduce work and information overload. *Communications of the ACM* 37(7):31-40.

Medsker, L.R., (1995). *Hybrid Intelligent Systems*. Kluwer Academic Publishers, Boston.

Muller, J.P., Wooldridge, M.J., & Jennings, N.R., (1996), *Intelligent Agents III*, Springer-Verlag, Heidelberg.

Port, R.F., & Van Gelder, T., (1995). Mind As Motion : Explorations in the Dynamics of Cognition, MIT Press.

Schreiber, G., Wielinga, B., & Breuker, J., (1993). *KADS: A Principled Approach To Knowledge-Based Systems*. Academic Press.

Sharp, B., Edwards, B., Dean, A., Davis, D.N., & Bancroft, G., (1998). A knowledge base environment to support water mains rehabilitation decision making, *ERA Technology Conference*, London.

Simmonds, R.J., (1999). Intelligent Agents and Natural Language Processing, *MSc Thesis, School of Computing, Staffordshire University*.

Sloman. A., (1996). What sort of architecture is required for a human-like agent? *Cognitive*

Modelling Workshop, American Association of Artificial Intelligence Conference 1996.

Wooldridge, M., & Jennings, N.R. (1994), *Intelligent Agents*. Springer-Verlag.

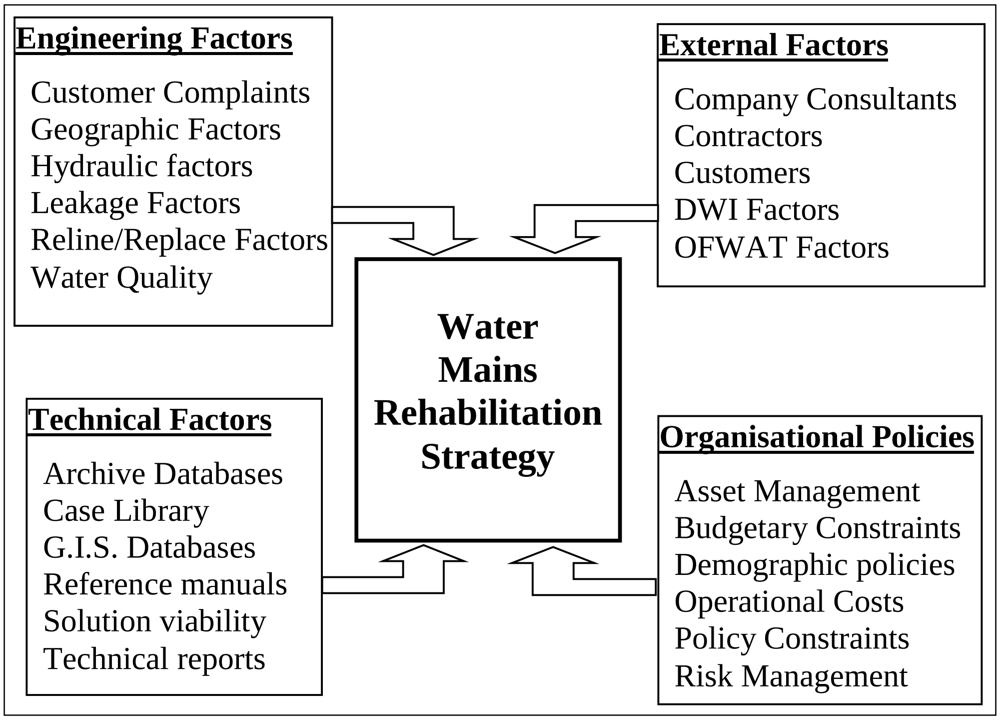

Figure 1. Influences to be captured in a complete decision support too.

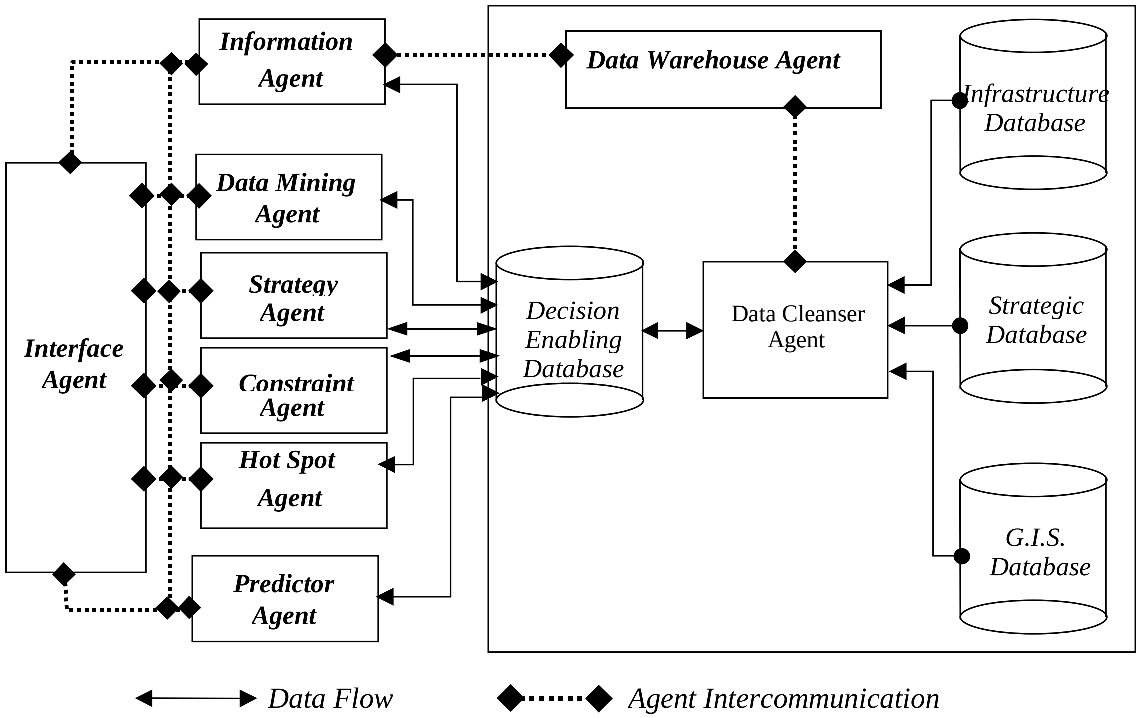

Figure 2. Schematic perspective on the loosely coupled agent-based framework.

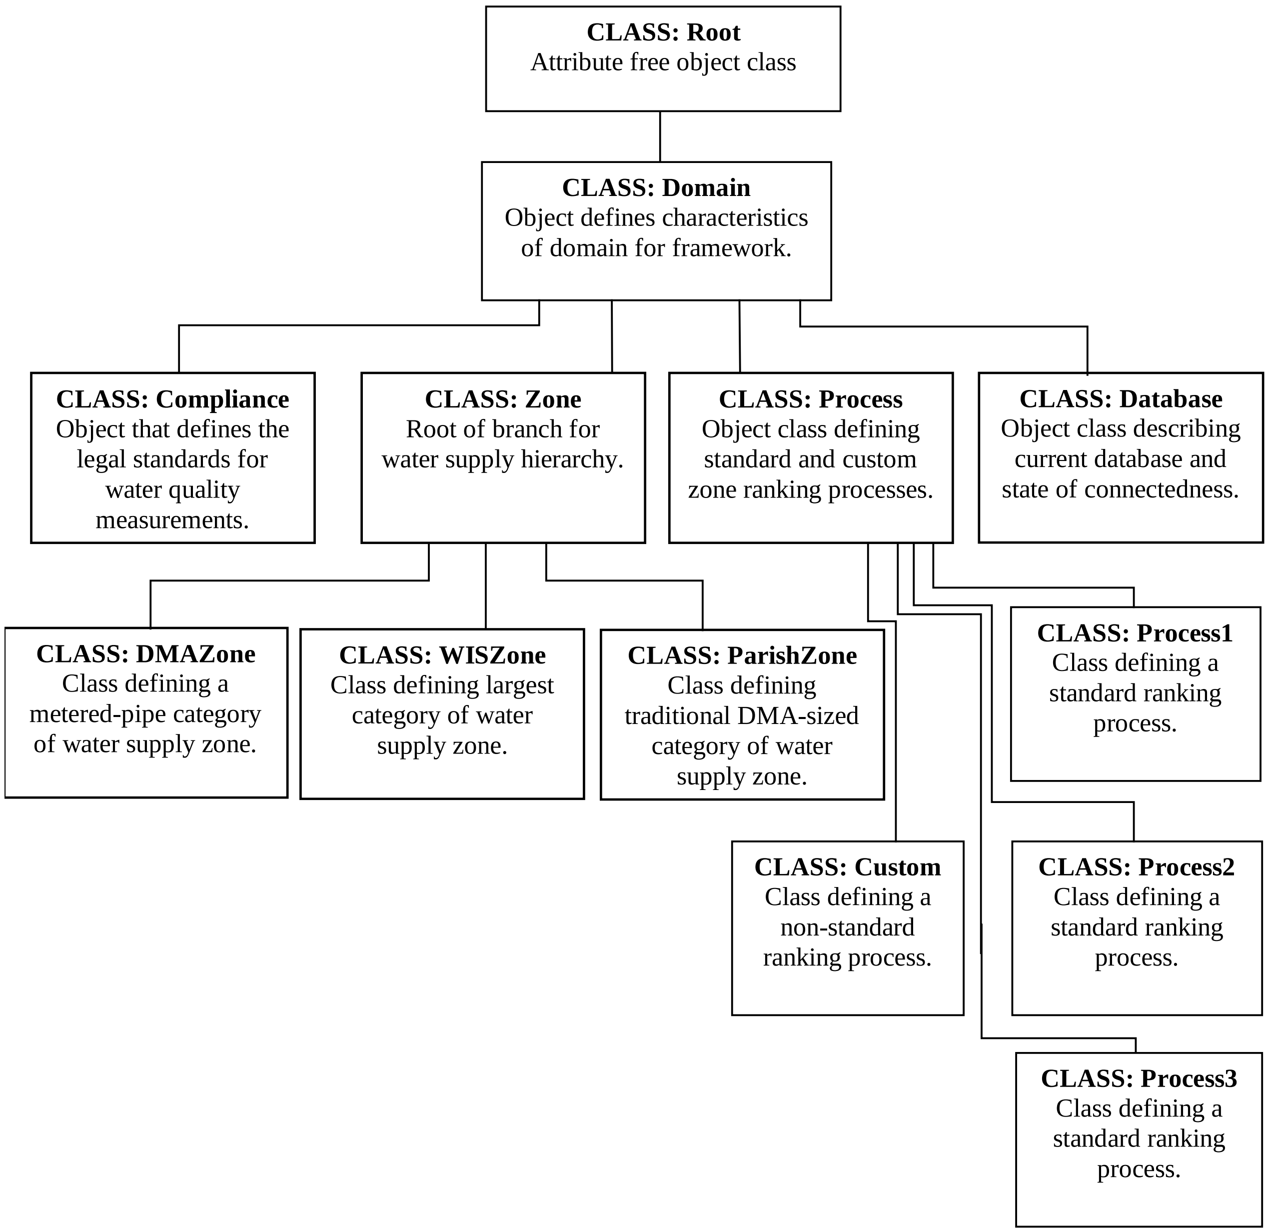

Figure 3:Domain Ontology And Object Class Hierarchy

(ask if	: Language	water-agent-framework
	: Ontology	water-supplier-X
	: Reply-with	one-of(low, medium, high, fail, null)
	: Sender	Information-Agent
	: Receiver	Data-Warehouse-Agent
	: MessageId	IAM-2
	: Content	"?risk(water-zone-1,supply-interruption)")

Figure 4: KQML performative example.

The Information Agent asks the Data Warehousing Agent about the risk of a water supply interruption in some specific water zone associated with a water supply company.

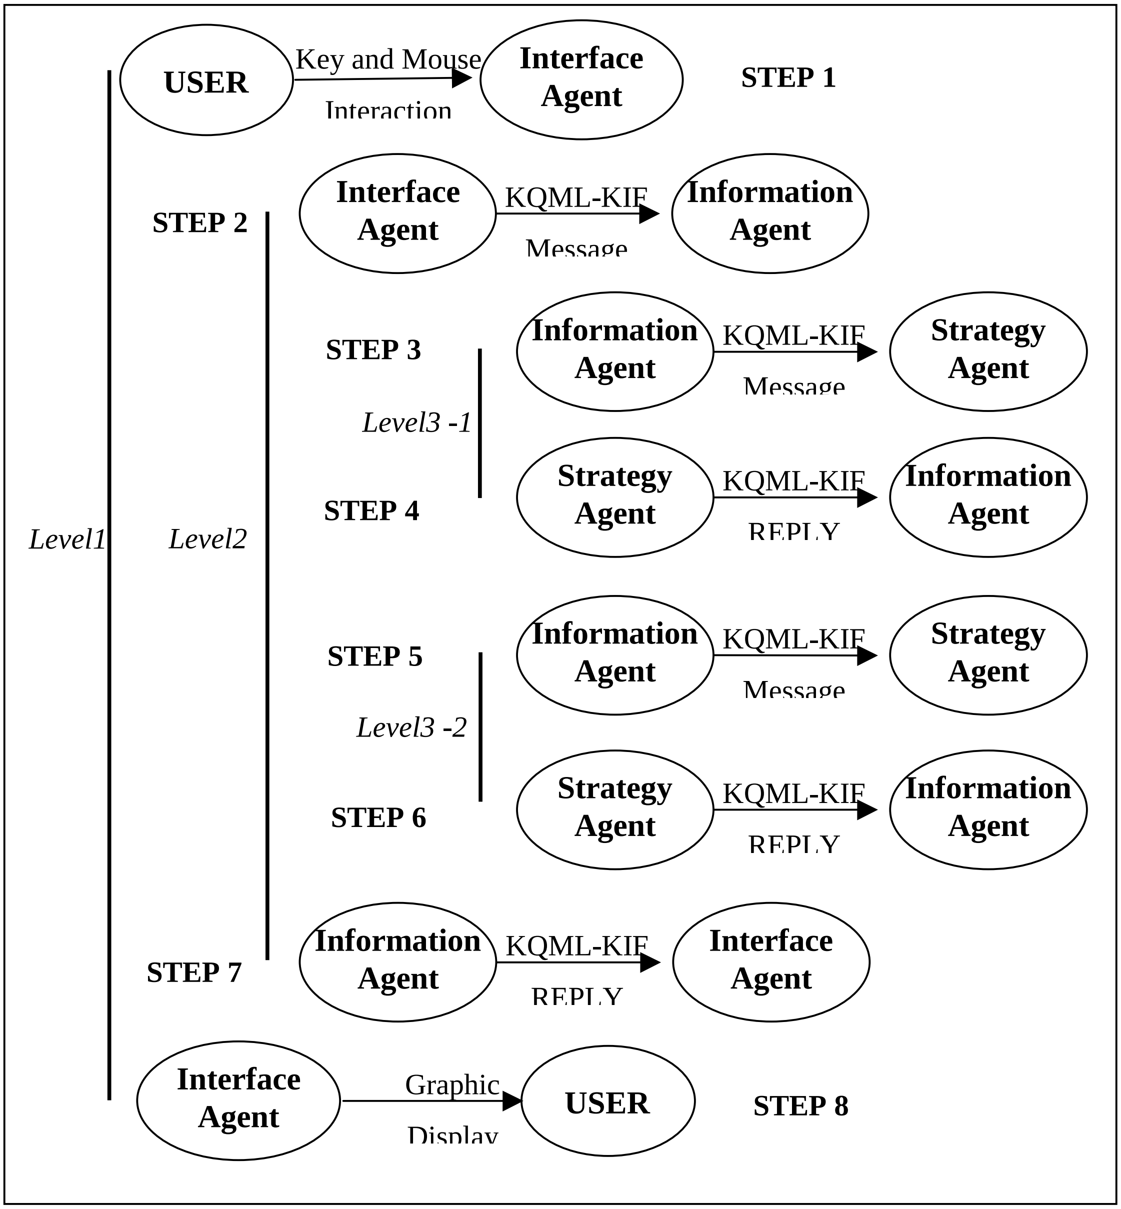

Figure 5: Agent interaction through KQML-KIF communication for the example scenario.

```
************************************************
S.W.I.M. Analysis On 05/09/99 at 14:29:39PM
************************************************
Ranking Process:
PCV Values Used
Aluminium  3573  Iron  3574  Manganese  3577  Turbidity  2  PAH  1
Sample Factors:
   Aluminium Weight :  0.09   Iron Weight       : 0.27
   Manganese Weight :  0.09   Turbidity Weight  : 0.18
   PAH Weight       :  0
Mains Supply and Other Factors:
   Supply Weight             : 0.10      DisColouration Weight     : 0.09
   Taste Complaint Weights   : 0.09      Chemical Dosing Weight    : 0
   Mains Filter Weight       : 0         Mains Flush Weight        : 0
   Pressure Weight           : 0.09      Population Density Weight : 0
************************************************
   Ranking Process Summary Information
   Number of Parishes In Ranking Process  :     633
   Mean Ranking Weight For These Parishes :     0.695
   Standard Deviation For The Weight      :     3.758
   Number of Parishes With Non-Zero Weight:     348
   Mean Ranking Weight For These Parishes :     1.265
   Standard Deviation For The Weight      :     4.996

************************************************
****** Parish Information:
 WISZONE:  309    WISNAME:  XXX     Parish:  311
 No of Properties             :  193
 Mains Length (km)            :  24456
 Population Density           :  463.2
 Compliance Violation Ranking :  1
 Ranking Value                :  67.596
 Interruptions                :  0
 Number of Bursts             :  71
 Pressure Complaints          :  0
 DisColouration Complaints    :  723
 Taste or Odour Complaints    :  8
 Properties Dosed             :  0
 Properties Filtered          :  4
 Properties Flushed           :  16
   Samples:  Failed
     Aluminium:  4  Tested With  0  Failed
     Iron:       4  Tested With  0  Failed
     Manganese:  4  Tested With  0  Failed
     Turbidity:  4  Tested With  0  Failed
     PAH:        4  Tested With  0  Failed
```

Figure 6:Exampleof a text report produced by the framework
(only the first parish zone is given – a further 632 parishes exist)

Zone	Rank1	Weight1	Rank2	Weight2	Rank3	Weight 3	Relative Rank	Overall Weight
Moorland	124	0.83	334	0	198	0.045	3	0.0033
Pleasant Ville	244	0	16	1.4	108	0.25	2	0.2497
Quiet Ville	23	16.59	74	0.48	23	4.39	1	0.2510

Table 1. Example Scenario. Each zone ranked over all 633 zones of the water company. The weights on each ranking process normalised (divided by maximum for that process) and then summed.