

An Intelligent Assistant for Public Transport Management

Martin Molina

Department of Artificial Intelligence, Universidad Politécnica de Madrid
Campus de Montegancedo s/n 28660, Boadilla del Monte, Madrid, SPAIN
mmolina@fi.upm.es

Abstract. This paper describes the architecture of a computer system conceived as an intelligent assistant for public transport management. The goal of the system is to help operators of a control center in making strategic decisions about how to solve problems of a fleet of buses in an urban network. The system uses artificial intelligence techniques to simulate the decision processes. In particular, a complex knowledge model has been designed by using advanced knowledge engineering methods that integrates three main tasks: diagnosis, prediction and planning. Finally, the paper describes two particular applications developed following this architecture for the cities of Torino (Italy) and Vitoria (Spain).

1 Introduction

The problem of management of a fleet of vehicles has been recently facilitated by the current telecommunication and information technology. In this field, a new generation of information systems have been proposed for a wide range of services in public transport and fleet management [1], [2]. In particular, the recent advances of this technology allow operators in a control center to monitor the location and state of each particular vehicle, in order to apply on real time global transport strategies. This is especially useful to quickly react in the presence of incidents produced by unexpected situations such as vehicles with malfunctions, road blocked in the transport network, etc. In this context, advanced computer systems can help operators in improving their answer [3].

The development of this type of systems requires formulating advanced models that capture the different facets of the strategic knowledge for public transport management. In order to provide an integrated solution for this problem, we describe in this paper the architecture of a computer system that follows the idea of *intelligent assistant* [4]. An intelligent assistant is a concept derived from artificial intelligence that identifies a kind of systems whose role is to *assist* the user in decision-making processes. This type of systems emphasizes that the operator is the final responsible of decisions, so the system is not designed to substitute the operator but, on the contrary, its goal is to provide services for assistance such as: information filtering and interpretation to identify significant data, *what if* analysis, justification of conclusions, etc.

According to this, the paper describes first the problem of public transport management corresponding to a fleet of buses of a urban network, identifying the main tasks to be provided: diagnosis, prediction and planning. Then, the paper describes the knowledge model designed to support these tasks, showing the selected knowledge-based methods for this purpose. Finally, the paper describes two of the applications developed following this approach. The first one was developed for the city of Torino (Italy) in the context of an international European Project (funded by the European Commission within the Telematics Applications Programme). The second application was developed in a national project in Spain for the public transport management in the city of Vitoria.

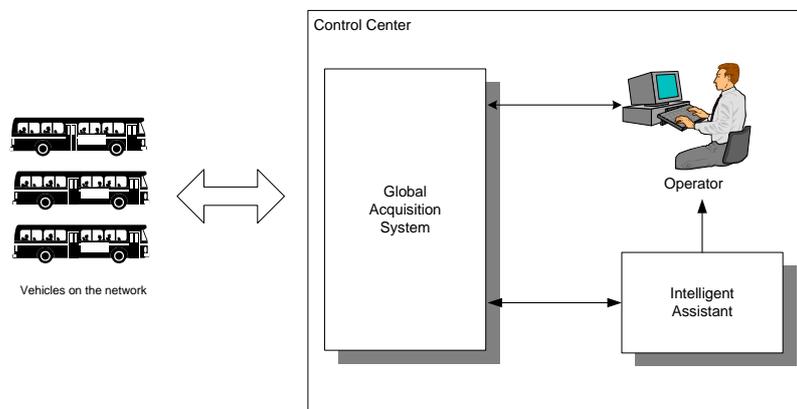


Fig. 1. Basic control scenario in public transport management.

2 The Problem of Public Transport Management

A typical public transport management scenario is composed of a fleet of buses that in real time send information about their location and state to a control center (figure 1). The goal of the public transport operators at the control center is to give to the passengers an adequate service in order to: (1) guarantee as much as possible the normal service according to the initial planning, (2) respect the timetable of drivers, (3) guarantee the security of the service, and (3) avoid discomfort of passengers and drivers.

The information available every minute in the control center about every vehicle includes: location, previous and next vehicles, actual headway/delay measure and its trend, etc. Other information are alarms (overloaded vehicle, broken vehicle, vehicle in traffic-jam, lack or relief driver, etc.), in-line drivers, current timetables, depots state, etc. In a typical decision scenario of this type of control centers, operators perform three basic tasks: diagnosis, prediction and planning.

The *diagnosis* task is oriented to identify the presence of abnormal situations based on the received information in real time about the location and state of the vehicles. An abnormal situation in this context is a situation where the planned service is

altered due to a set of possible causes. These situations have different degrees of severity that go from a slight delay of a single vehicle to a severe disruption of the service due to a blockade of the line or a malfunction of one vehicle. In general, abnormal situations include different types of problems for each line such as deviations of vehicles from their scheduled services (slight or severe), service disruption, unreachable relief point, slight or severe malfunction of a vehicle, blockade of the line, etc.

The goal of the *prediction* task is to estimate the future impact of the detected problems. For this purpose, operators can estimate the number of passengers affected and the average delay assumed for the next period of time. This is performed using knowledge about the historical demand of the lines and knowledge about the behavior of the vehicles (e.g., average travel times between stops).

Finally, the goal of the *planning* task is to find appropriate actions according to the detected problems that should be taken in order to improve the transport service. These actions in general need to be performed in a particular order and they can cover more than one bus line. In particular, some of the types of basic actions are: (1) *skipping*, a bus does not take passengers during several stops, (2) *detour*, a bus takes a path that goes out of the line in order to avoid an obstacle (for example, traffic incident, constructions, etc.), (3) *limitation*, a vehicle takes the opposite direction of the line, without covering the complete path, (4) *reinforcement*, a reserve vehicle (a bus that is not performing a service in a particular line) is sent to cover an unattended area in the line with the objective of restoring the timetable of delayed vehicles, (5) *shift*, the drivers of two vehicles are exchanged to restore the right order in the duty, and (6) *rotation*, the schedule of buses is shifted one step to recover delays.

3 The knowledge model of the intelligent assistant

In order to provide an appropriate level of support for public transport management, the intelligent assistant system should simulate the natural thinking process followed by operators using their strategic criteria at the same levels of abstraction. For this purpose, we applied advanced knowledge engineering techniques following a *model-based* approach. This modeling approach considers the existence of a conceptual level, at which the knowledge is first described without considering implementation issues. Some of the recent methodologies for knowledge system development follow this model-based approach (e.g., CommonKADS [5]). These methodologies basically organize the whole knowledge using the following concepts: (1) a *task* that identifies a goal to be achieved (for instance, diagnosis or prediction), (2) a *method* indicates how a task is achieved, by describing the different reasoning steps (subtasks) by which its inputs are transformed into outputs, and (3) a *type of knowledge base* that identifies explicitly the type of domain knowledge that supports a task.

According to this, figure 2 shows a global view of the knowledge model that we designed for the public transport problem. This figure shows a hierarchy of tasks (circles) and methods (rectangles) with types of knowledge bases at the bottom (cylinders). The figure shows that the global task, *public transport management*, is divided into the three main subtasks (*diagnosis*, *prediction* and *planning*) that

correspond to the three tasks described in the previous section. The figure also shows how each subtasks is decomposed in simpler subtasks and how they are supported by different types of domain knowledge.

Thus, the goal of the first of the three main tasks, *diagnosis*, is to interpret the current situation in order to detect the presence of problems. This type of reasoning includes a qualitative interpretation of the raw information received in the control center together with a classification of the situation according to a prefixed set of types of problems. For this purpose it is useful to use an adaptation of the heuristic classification method [6], with a data-driven control regime that considers two simpler subtasks: *abstract* and *match*. The abstract subtask is a primary task whose goal is to produce qualitative values from numerical data about the state of the system. For example, this task interprets the delay (in minutes) of each vehicle producing one of three qualitative values (slight, medium or severe). The goal of the *match* subtask is to determine the type of problem that is present at each line, using a set of classes of problems. Two types of knowledge bases support the *diagnosis* task: one that includes abstraction knowledge and another one with problem types.

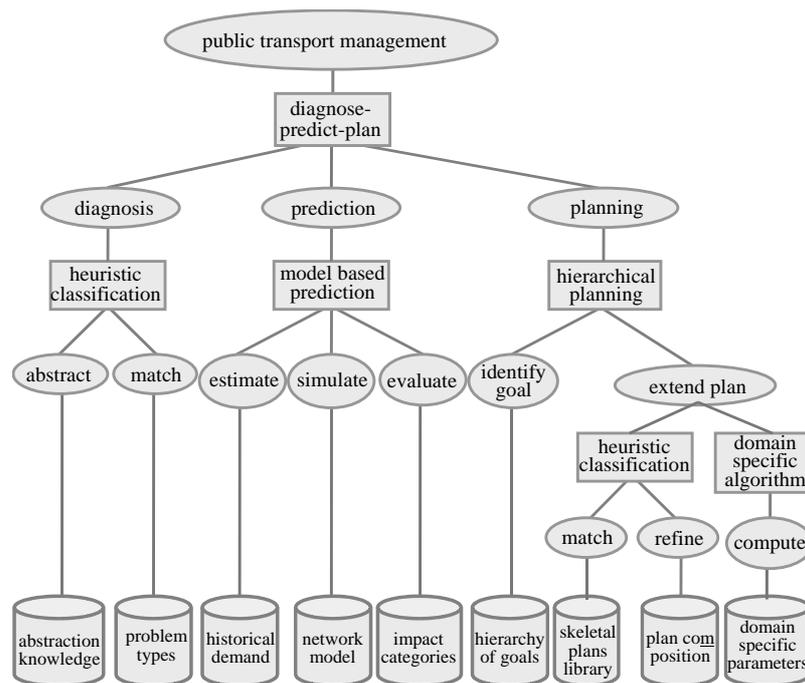


Fig. 2. Knowledge model for public transport management.

On the other hand, the goal of the *prediction* task is to estimate the future impact of the problem. This prediction, estimates the total number of passengers affected by the delay together with the average delay for those passengers. This is calculated for the next T minutes (for instance, $T=15$ minutes). In particular, for example, this method may compute the values for the impact I and total number of affected persons P :

$$I = \sum_{i=1}^n p_i \cdot e_i \qquad P = \sum_{i=1}^n p_i$$

where p_i is the estimated number of persons waiting at the stop i , and e_i is the increase of waiting time due to the delay of the bus. The sum is done for the n stops affected in the T minutes. The value I gives a quantitative measure to compare the severity of different problems. Thus, it provides a global criterion to solve conflicts when different problems in different lines need the same resource to improve the service (e.g., a reserve vehicle). The value $E = I/P$ expresses the average increase of waiting time.

The prediction task is divided into three subtasks: *estimate future demand*, *simulate behavior* and *evaluate impact*. The first subtask estimates the future demand by using a local knowledge base with historical demand. This includes the expected number of passengers at each stop each interval of time. The second subtask simulates the movement of delayed buses for the next T minutes by using a model of the bus-line (stops, distances, travel-times, etc.) and determines the affected stops. Finally, the third subtask estimates global metrics (e.g. I and E) and uses domain knowledge to interpret the severity of the situations.

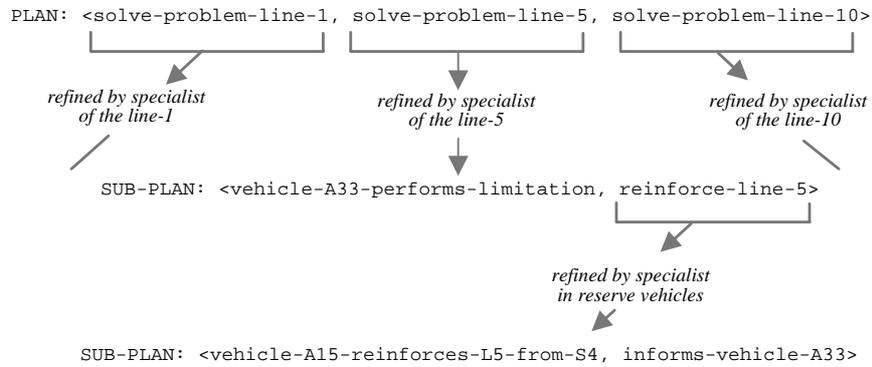


Fig. 3. Example of development of a plan by the intervention of different specialists that, by turn, refine parts of an abstract plan.

The third task, *planning*, recommends a set of control actions that may solve the detected problems. In this type of reasoning, different classes of specialized heuristic knowledge are typically used to dynamically construct the plan. For example, knowledge about alternative paths of bus lines, criteria about reserve vehicle management, specific criteria to exchange drivers, etc. In order to simulate this type of reasoning it is possible to use a method for hierarchical planning that integrates the idea of skeletal planning [7] and the concept of specialists [8]. The method is based

on a search in a hierarchy of goals (specialists) that are knowledgeable about partial abstract plans. Each specialist is responsible of producing a partial plan to solve part of the detected problem.

Thus, the knowledge model for this case is represented by a method that divides the *planning* task into two main subtasks: *identify goal* and *extend plan*. The plan is dynamically composed during the reasoning developing a search in such a way that the first subtask identifies the next goal to reach, analyzing the current plan and using a hierarchy of goals (the hierarchy of specialists). Then, the second subtask extends the current plan with partial sub-plans. This process is repeated step by step until the complete final plan is produced. The second subtask is performed by two alternative methods (depending on the level of the hierarchy of goals): (1) one method that selects a plan using heuristic classification knowledge, and (2) another method that constructs the plan by using a domain-specific algorithm with certain parameters (in the bus transport application, there are several instances of this method, for example, the method for reserve vehicles selection).

Type of KB	Representation	Examples												
Abstraction knowledge	Rules	IF <i>X</i> is-a vehicle, delay of <i>X</i> = <i>D</i> , <i>D</i> > 10 THEN level-of-delay of <i>X</i> = severe												
Problem types	Rules or frames	IF <i>X</i> is-a vehicle, full-vehicle-alarm of <i>X</i> = on, current-stop of <i>X</i> = <i>S</i> THEN new problem <i>P</i> , type of <i>P</i> = overflow, affected-vehicle of <i>P</i> = <i>X</i> , location of <i>P</i> = <i>S</i>												
Historical demand	Table	<table border="1"> <thead> <tr> <th>Stop</th> <th>Time-Interval</th> <th>Type-of-day</th> <th>Hist-Demand</th> </tr> </thead> <tbody> <tr> <td><i>S</i>₁,</td> <td>6:30-9:30,</td> <td>weekday,</td> <td>20,</td> </tr> <tr> <td><i>S</i>₂,</td> <td>6:30-10:00,</td> <td>weekday,</td> <td>15, ...</td> </tr> </tbody> </table>	Stop	Time-Interval	Type-of-day	Hist-Demand	<i>S</i> ₁ ,	6:30-9:30,	weekday,	20,	<i>S</i> ₂ ,	6:30-10:00,	weekday,	15, ...
Stop	Time-Interval	Type-of-day	Hist-Demand											
<i>S</i> ₁ ,	6:30-9:30,	weekday,	20,											
<i>S</i> ₂ ,	6:30-10:00,	weekday,	15, ...											
Network model	Objects	Object <i>L</i> ₃ is a Line. Attributes: Stops: (<i>S</i> ₁ , <i>S</i> ₇ , <i>S</i> ₁₅ , <i>S</i> ₁₈ , <i>S</i> ₂₁ , <i>S</i> ₄₅ , <i>S</i> ₅₆ , <i>S</i> ₇₈), ...												
Impact categories	Rules	IF <i>X</i> is-a line, affected-passengers of <i>X</i> > 50, average-delay of <i>X</i> > 5 THEN impact-level of <i>X</i> = severe												
Hierarchy of goals	Table	<table border="1"> <thead> <tr> <th>Specialist</th> <th>Goals</th> <th>Pre-cond.</th> <th>Post-cond.</th> </tr> </thead> <tbody> <tr> <td>Reinforcement,</td> <td><i>G</i>₁, <i>G</i>₂, ...</td> <td><i>A</i>₁, ...</td> <td><i>A</i>₂, ...</td> </tr> <tr> <td>Limitation,</td> <td><i>H</i>₁, <i>H</i>₂, ...</td> <td><i>A</i>₃, ...</td> <td><i>A</i>₄, ...</td> </tr> </tbody> </table>	Specialist	Goals	Pre-cond.	Post-cond.	Reinforcement,	<i>G</i> ₁ , <i>G</i> ₂ , ...	<i>A</i> ₁ , ...	<i>A</i> ₂ , ...	Limitation,	<i>H</i> ₁ , <i>H</i> ₂ , ...	<i>A</i> ₃ , ...	<i>A</i> ₄ , ...
Specialist	Goals	Pre-cond.	Post-cond.											
Reinforcement,	<i>G</i> ₁ , <i>G</i> ₂ , ...	<i>A</i> ₁ , ...	<i>A</i> ₂ , ...											
Limitation,	<i>H</i> ₁ , <i>H</i> ₂ , ...	<i>A</i> ₃ , ...	<i>A</i> ₄ , ...											
Skeletal plans library	Rules	IF <i>X</i> is-a problem, type of <i>X</i> single-severe-delay, affected-vehicle of <i>X</i> = <i>V</i> THEN new plan <i>P</i> , type of <i>P</i> = reinforced-limitation, affected-vehicle of <i>P</i> = <i>V</i>												
Plan composition	Frames	PLAN reinforced-limitation. ACTION TYPE concrete: vehicle <i>X</i> performs limitation ACTION TYPE abstract: vehicle <i>X</i> is reinforced ...												
Domain specific parameters	Table	<table border="1"> <thead> <tr> <th>Parameter</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>max-number-of-stops-to-regulate,</td> <td>5,</td> </tr> <tr> <td>min-percentage-of-absorption,</td> <td>80%, ...</td> </tr> </tbody> </table>	Parameter	Value	max-number-of-stops-to-regulate,	5,	min-percentage-of-absorption,	80%, ...						
Parameter	Value													
max-number-of-stops-to-regulate,	5,													
min-percentage-of-absorption,	80%, ...													

Fig. 4. Symbolic representation of knowledge bases for the model of public transport management.

Figure 3 shows an example about how a plan is developed by the intervention of several specialists. In the example, first, a global plan is produced with three abstract actions that respectively are associated to line-1, line-5 and line-10. The corresponding specialist refines each abstract action. For example, the second action is refined by the specialist of the line-5 and produces a sub-plan with two actions. In addition to that, the second action of this sub-plan needs to be refined by another specialist.

In summary, according to the previous description, the model includes a total of 9 types of knowledge bases. Figure 4 shows the symbolic representation followed by each type of knowledge base. Thus, for instance, the knowledge base that supports the *abstract* task can be represented with *if-then* rules with a forward-chaining inference method. This representation is very flexible and intuitive to formulate a wide range of expressions for qualitative interpretation. The knowledge base for problem types can be represented either with rules or frames. Here, the antecedent of a rule can express the set of conditions for a certain type of problem and the consequent determines the characteristics of the deduced problem. Frames can be also used here, where each frame represents a type of problem with a set of slots that expresses a set of single conditions.

The knowledge base for historical demand can be represented with a table that relates each stop of a line with its historical demand according to the type of day and the time interval. For the knowledge base of the network model, an object-oriented representation can be used with concepts such as stop, line, vehicle, etc. and attributes with values that characterize the concepts and establish relations between concepts. The knowledge base for impact categories can be also represented using *if-then* rules to interpret the quantitative values about the expected delays.

Concerning the knowledge related to the planning task, the knowledge base about the hierarchy of goals can be represented with a table that relates each type of goal with the corresponding specialist and the conditions to be considered. The knowledge base for skeletal plans library includes the existing types of plans. Here, rules can relate conditions of the problem with the corresponding type of plan. The knowledge base for plan composition expresses the set of actions in which a plan is decomposed. For this purpose, frames can be used in such a way that each frame identifies a plan and the set of slots identifies the structure of the plan with abstract or concrete actions. Finally, when a specialist is supported by a set of parameters, the knowledge base with domain specific parameters include the corresponding values.

It is important to note that these 9 types of bases can produce a higher number of specific knowledge bases in a particular model because some of the bases are specified with different contents for each bus line. Thus, in a concrete model for several bus lines, some knowledge bases can be shared for the management of all the lines but other knowledge bases need to be specified for each particular line. Thus, for instance, a particular model with N bus lines could include: 1 abstraction knowledge base, 1 knowledge base for problem types, N knowledge bases for historical demand, N knowledge bases for network model, 1 knowledge base for impact categories, 1 for hierarchy of goals, 1 for skeletal plan library, 1 for plan composition, M for domain specific parameters (where M is the number of specialists that use domain specific algorithms to dynamically construct a partial plan).

4 Applications

The model described in this paper was applied for the development of two different applications. The first application was developed within an international project in Europe, called Fluids [9], funded by the European Commission within the Telematics Application Programme. One of the main goals of this project was to provide advanced methods for user-system interaction in the context of real-time decision support. The design methods developed in the project were demonstrated in different applications for transport. In particular, a realization was developed for the case of public transport management for the city of Torino (Italy). The analysis of this problem produced a first version of the general approach presented in this paper.

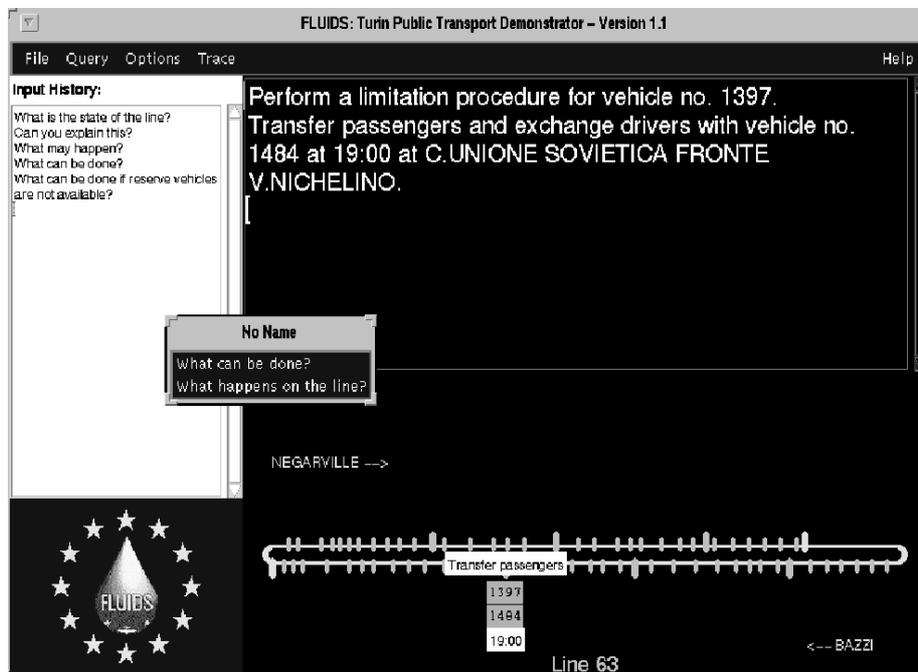


Fig. 5. User interface of the Fluids application for public transport management.

Figure 5 shows an example of the user interface developed for the Fluids project (it was tested on-line in the control center in 1998). The user interface combines a dialogue based on a prefixed set of questions for decision support together with a presentation created dynamically that integrates text and an animation to clarify with illustrations how to carry out specific control actions for vehicle management.

From the point of view of implementation, the development of this type of systems is a complex task because it requires to integrate different types of problem solvers with different types of knowledge bases. The final architecture must be both efficient for real time operation and flexible to accept changes according to the identification of new strategic knowledge. Thus, in order to help in the final implementation of the

system, for the case of the Fluids application, we used a software environment called KSM [11] that has been already used for the development of other applications in the field of transport [12]. This environment is an advanced knowledge engineering tool that provides a model-based approach and a set of software components to facilitate the development.

The Fluids application was followed by another application for the city of Vitoria (Spain) which included some improvements compared to the initial version. It was conceived in a general way, according to the model described in this paper, to be reused for different public transport networks. This realization was developed and integrated on real-time in 2001 with the rest of the information system for the management of the bus lines of the control center of the city of Vitoria (Spain). The model includes the management criteria for a total of 15 bus lines. Figure 6 shows the global user interface of this application.

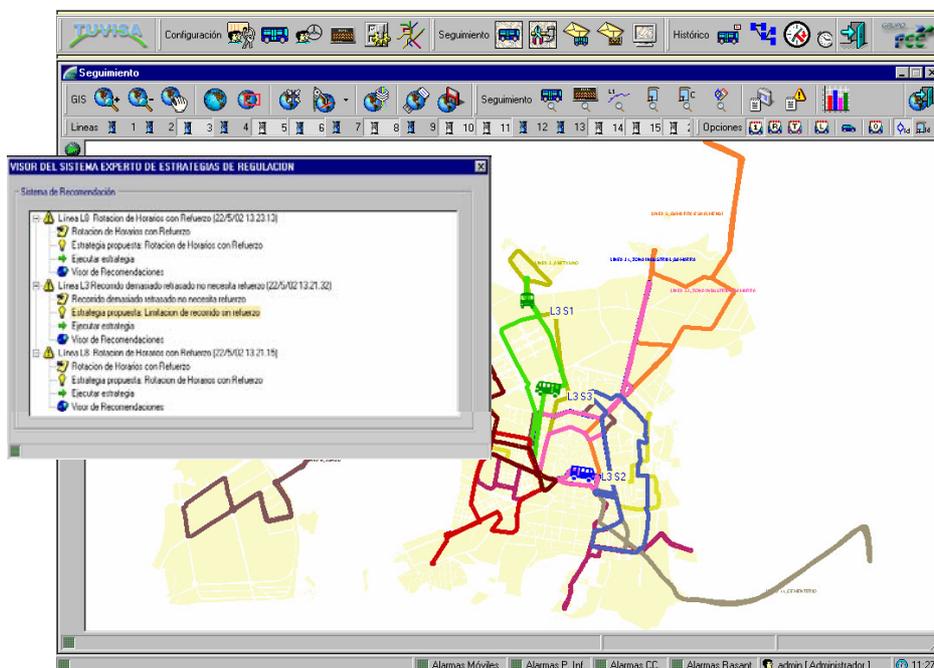


Fig. 6. User interface of the application for public transport management for the city of Vitoria.

The implementation of this system was performed following a particular advanced software architecture that integrates the set of inference procedures and knowledge bases using a common working memory together with an explicit control mechanism. The user can access to the knowledge model using a particular user interface developed for this purpose in order to provide the required flexibility for knowledge model maintenance.

5 Conclusions

In summary, the model described in this paper can be an adequate solution to construct an intelligent assistant for public transport management. The model is a global solution that integrates different types of methods and problem-solvers to support three basic tasks: diagnosis, prediction and planning. The model-based approach followed in knowledge engineering provided the appropriate description level to formulate the complete model, according to the current state of the art of knowledge engineering. The model presented in the paper has been validated with the development of two real-world applications for the cities of Torino (Italy) and Vitoria (Spain).

References

1. Witulski K.: "Knowledge Based Route Selection in Public Transport" in "Operational Experience and Perspectives" Proc. 2nd OECD Workshop on Knowledge-based Expert Systems in Transportation, Vol.1. Montreal (Canada), June 1992.
2. Saint Laurent B. de: "An Information System for Public Transport: The Cassiope Architecture Example of Passenger Information" in "Advanced Telematics in Road Transport" Edited by the Commission of the European Communities, Directorate-General Telecommunications, Information Industries and Innovation. Elsevier, 1991.
3. Cepeda M.: "New Generation of Vehicle Scheduling and Control Systems" in "Advanced Telematics in Road Transport" Edited by the Commission of the European Communities, Directorate-General Telecommunications, Information Industries and Innovation. Elsevier, 1991.
4. Boy, G., Gruber T.R.: "Intelligent Assistant Systems: Support for Integrated Human-Machine Systems" Technical Report KSL 90-61, Knowledge Systems Laboratory, Computer Science Department, Stanford University, 1990. And published also in the proceedings of 1990 AAAI Spring Symposium on Knowledge-Based Human-Computer Communication, March 1990, Stanford University.
5. Schreiber G., Akkermans H., Anjewierden A., de Hoog R., Shadbolt N., Van de Velde W., Wielinga B.: "Knowledge Engineering and Management. The CommonKADS Methodology" MIT Press, 2000.
6. Clancey W.J.: "Heuristic Classification". Artificial Intelligence, Vol 27, pp. 289-350, 1985.
7. Friedland P.E.: "Knowledge-based experiment design in molecular genetics". Proc. Sixth International Joint Conference on Artificial Intelligence, 285-287, Menlo Park, California. También: report STAN-CS-79-771, Stanford University. 1979.
8. Brown D., Chandrasekaran B.: *Design Problem-solving: Knowledge Structures and Control Strategies*, Morgan Kaufman, 1989.
9. Hernández J., Molina M., Cuenca J.: "Towards an Advanced HCI Through Knowledge Modelling Techniques" in "Knowledge Engineering and Agent Technology" Cuenca J., Demazeau Y., García-Serrano A., Treur J. (eds.) IOS Press, 2004.
10. Cuenca J., Molina M.: "The role of knowledge modelling techniques in software development: a general approach based on a knowledge management tool" International Journal of Human-Computer Studies. No. 52. pp 385-421. Academic Press, 2000.
11. Molina M., Hernández J., Cuenca J.: "A Structure of Problem-solving Methods for Real-time Decision Support in Traffic Control". Journal of Human and Computer Studies (Academic Press) N.49, 577-600, 1998.