



KNOWLEDGE-BASED MODELS FOR ADAPTIVE TRAFFIC MANAGEMENT SYSTEMS

JOSÉ CUENA, JOSEFA HERNÁNDEZ and MARTÍN MOLINA

Department of Artificial Intelligence, Technical University of Madrid,
Campus de Montegancedo s/n, 28660—Boadilla del Monte, Madrid, Spain

(Received 10 July 1995; in revised form 7 September 1995)

Abstract—This paper describes a general approach for real time traffic management support using knowledge based models. Recognizing that human intervention is usually required to apply the current automatic traffic control systems, it is argued that there is a need for an additional intelligent layer to help operators to understand traffic problems and to make the best choice of strategic control actions that modify the assumption framework of the existing systems. The need for an open architecture is stated, in order to allow users to modify decision criteria according to their experience, given that no skills are available yet to deal with real time strategy decision making. An architecture of knowledge is described that is oriented towards traffic management strategic advice applied in the TRYS system developed by the authors. This system has been installed for urban motorway control in several Spanish cities. Finally, an example of knowledge-based modeling, using TRYS, is presented in a case study where both the TRYS model and its operation are described. It is concluded that such an approach is feasible, and is compatible with existing state of the art traffic control systems.

1. INTRODUCTION

The concept of *knowledge-based systems* was introduced at the end of the 1970s following several well known applications such as DENDRAL (Lindsay *et al.*, 1981), MYCIN (Buchanan and Shortliffe, 1984) and PROSPECTOR (Duda *et al.*, 1979). The idea of the knowledge-based architecture is to organize a computer application in declarative components describing the static knowledge about a problem domain together with inference procedures modeling the reasoning processes to solve different problems using the different domain declarative knowledge components.

In the first generation of knowledge-based systems, a limited set of basic standard reasoning procedures was defined using declarative knowledge formulated using rules: production rules as in OPS5 (Brownston *et al.*, 1985), clause rules like the ones used in the PROLOG language (Lloyd, 1987) and semantic networks or frame networks (Minsky, 1975; Brachman and Schmolze, 1985; Goldstein and Roberts, 1977).

A second generation of knowledge-based systems was proposed in the 1980s supporting a more general view of the architecture where different problem solvers, more adapted to the specificity of each problem, could be defined (Chandrasekaran, 1983, 1987; McDermott, 1988). The organisation was based on task architectures where it was possible to structure a knowledge model by integration of problem solvers at different levels. The different modules may use diverse symbolic representations including conventional modules formulated using a procedural language without explicit declarative knowledge (Cuenca and Molina, 1993).

Organising an application with a knowledge oriented approach provides three main advantages:

- The possibility of modeling, using declarative representations (rules, frames, etc.), the knowledge which is not well formalised in algorithms, such as heuristics used in decision processes. Conventional software engineering methods are not adequate to model such domains.

- The possibility of easy management of the knowledge by non-programmer users because it is easy for them to understand the declarative content of the application. This advantage allows the system to be easily adapted to new requirements.
- The possibility of getting explanations for every answer provided by the system. An explanation is a description of the knowledge components used by the reasoning problem solver to reach an answer to one problem. If the user is familiar with the procedures applied by the problem solver (which is usual in the different areas of expertise), he/she may understand very well the role played by the different pieces of domain knowledge existing in the model. Getting explanations from one application is an important feature for some systems, such as real time decision support systems where the system may propose actions that should be implemented by the user.

The issue of an adequate design of the user interface to ensure good human–computer interaction is now a subject of growing attention. Traditionally, two possible extreme models may be considered:

- *Minimizing the role of the system.* The system provides the right information and the user manages 100% of the decision making.
- *Minimizing the work and responsibility of the user.* The user is at the service of the application simply providing data and the application automatically assumes 100% of decisions.

The first one minimizes too much the potential service of a computer application and, when the time available for decisions is short, it may create too much stress in the human operator resulting in inefficient performance. The second one minimizes too much the role of the operators who will have a tendency to reduce their responsibility for the decisions suggested by the system. A good design must be based on intermediate concepts capable of integrating the human capacity for assuming moral options and the computer efficiency for data interpretation and decision evaluation.

A typical intermediate approach is:

- The user receives proposals and explanations from the system justifying questions such as:
 - *why* do you need this premise,
 - *how* did you meet this conclusion,
 - *what will happen* if I decide: D1,D2,...,D4,
 - *why* will it happen,
 -
- The user implements options after a user–system conversation in the previous terms.

The knowledge based architecture permits creation of adequate areas of knowledge to be inspected by the user and adequate patterns of general and local explanations supporting a *model of conversation* that can optimize the efficiency of the user–system couple. Traffic management systems represent a class of decisions support systems where this approach may be productive.

Work in knowledge-based systems applied to traffic management started at the end of the 1980s (Scemama, 1989; Cuenca, 1989). At this time, the OECD created an expert group about expert systems in transport that produced two symposia, OECD (1990, 1992). Also, the European ATT/DRIVE Program promoted different experimental projects such as CLAIRE (Scemama, 1992) and KITS (Boero *et al.*, 1993, 1994; Cuenca *et al.*, 1992a,b).

This paper proposes a general structure for traffic control knowledge modeling at the strategic level to represent criteria for problem understanding and problem solving in areas controlled by traffic lights and/or variable message sign (VMS) panels. This model has been implemented in the TRYS project funded by the Spanish Directorate for Traffic (DGT, 1994).

The contents of the paper are organised as follows. First, the paper provides a rationale for the introduction of knowledge-based models into the current urban traffic control (UTC) technology and traffic management practice. Then, a general conceptual model for traffic analysis and management is described, followed by a formulation and discussion of the components of the model at the symbol level. Finally, an example, including model formulation and operation for illustrative purposes, is described in detail.

The presentation in this paper is oriented towards the understanding of this type of system by users or designers with traffic management expertise, but not necessarily experts in AI based models.

2. THE ROLE OF KNOWLEDGE-BASED SYSTEMS IN TRAFFIC MANAGEMENT

Traffic management systems must be reactive to the different states of traffic flow in the controlled network. In the early systems, the approach was based on a library of signal plans applied on-line in different predefined situations according to some time-based criteria (fixed-time systems) or to the traffic data collected by roadside sensors (traffic actuated systems).

However, this precalculated-plan approach usually lacked the conceptual granularity required by the system to be adaptive enough to the variety of situations, in time and space, that may occur in the network. As a result, this type of system was usually operated in a complementary way: a plan was selected from the library but, in certain situations, the operator had to introduce modifications according to her/his perception of reality together with her/his knowledge of traffic behavior in the network.

In the 1980s more adaptive systems were introduced starting from the SCOOT model (Hunt *et al.*, 1981; Bretherton and Bowen, 1990) where, for the first time, an intelligence for understanding traffic situations in real time was designed and integrated with a model for decision making. Several other responsive systems were developed in the past decade, including SCATS (Lowrie, 1982), OPAC (Gartner, 1983), PROLYN (Henry *et al.*, 1983) and UTOPIA (Mauro and Di Taranto, 1989).

Dynamic systems like SCOOT signaled a very relevant change in traffic control systems, since in such systems on-line intelligence for traffic control was implemented. This intelligence was organized as an evaluation step of local analysis at the level of single junctions together with a more general view evaluating a sequence of junctions. The analysis was performed over time to act on signal phase modifications during time slices of few seconds. The reasoning model, however, was hidden to the user although some parameters remained externally modifiable.

The experience in using SCOOT-like systems shows that good performance is attained when traffic situations are not critical, or congested conditions occur only for limited time intervals. In persistently congested situations, as in the precalculated plans approach, operator intervention is usually required because:

- (i) The evaluation of the traffic situation by the system may be biased due to the fact that sensors are sometimes insufficient for understanding the congestion process (e.g. the length of the queues at the stop line may overflow past the sensors placed for queue detection).
- (ii) For the sake of stability in the network, the decisions taken by the system are small adjustments to the current signal plan, while more drastic changes might often be required.

Studies performed in a number of sites have shown that operator intervention is almost customary in most UTC installations. The knowledge used by the human operator to decide on these changes is mostly based on her/his understanding of the reasons causing the problem (i.e. the structure of the traffic demand, the existence of alternative paths to some conflicting nodes etc.).

The same need for taking into account specific aspects in problematic areas appears in motorway control through VMS panels. The advice displayed to drivers should be

sufficient to: (i) allow the users to take some of the available options; and (ii) not produce excessive flow outside the motorway that may congest alternative options. Every problem area therefore requires analysis of the situation using knowledge about traffic behavior and control criteria specific for that area.

The above considerations suggest a need to complement existing systems for traffic control (including pre-calculated plans systems, dynamic systems and VMS systems) with an additional layer where strategic knowledge is applied to understand the specific processes of congestion development, and corresponding actions for alleviating the problem may be modeled. As the knowledge to be introduced depends on the characteristics of the area, general models such as SCOOT may be not open enough since the form of the computer application supporting this type of model should allow the user to introduce, modify and maintain over time some conceptual formulations of his/her knowledge related to every area of interest in the network. This may require more than a set of parameter values to be tuned to the specific application site, as used by existing traffic control systems. The technology of knowledge-based systems may help in designing and implementing suitable knowledge structures to formulate conceptual models for traffic analysis and management and to use such models for on-line strategic traffic management operations.

3. THE CONCEPTUAL MODEL OF TRYS

The ideas presented above have been put into practice in the TRYS system (DGT, 1994). TRYS is a knowledge representation environment supporting models to perform traffic management at a strategic level in urban, interurban or mixed areas. The city or traffic network where traffic has to be supervised is divided in several sections called problem areas. The decomposition of the city into problem areas allows a better analysis and understanding of the causes and evolution of traffic problems than if performed from a global perspective. This split does not define a set of disjointed areas whose sum is the whole city, but every area represents a part of the city where a determined traffic behavior is usually present and where a set of signal elements can be managed to influence this behavior. Then, a problem area may overlap with surrounding areas sharing, for instance, some signals but using them from different points of view. So, a problem area is a part of a city where traffic behavior is locally studied and suitable control actions may be defined to improve the traffic state.

Every problem area is supervised by an agent which understands the traffic conflicts that may appear, the usual behavior of vehicles in the area and the signal and/or VMS actions that can improve the traffic. The control proposals generated by every agent are received by a higher level agent, called the co-ordinator, whose aim is to produce global proposals for the whole city by putting together the local proposals provided by the agents and removing the inconsistencies among them. Figure 1 shows the organization of a set of agents.

Based on this concept, TRYS organizes the control knowledge in a set of local control entities named agents and a co-ordinator to synthesize the proposals. In the next sections, the knowledge used by each one of these units is detailed.

3.1. Model of an agent

The goal of an agent is to provide two types of information: diagnosis of the traffic problems present in a local area together with an explanation justifying such a diagnosis, and proposed control actions for the available signal devices to improve traffic conditions using the diagnosis information.

To support this functionality, the knowledge of an agent is distributed in the following way:

- *Physical structure:* Knowledge about the behavior of the traffic network of the problem area. It includes procedures for interpreting sensor data (data abstraction) and understanding the network physical structure at different levels.

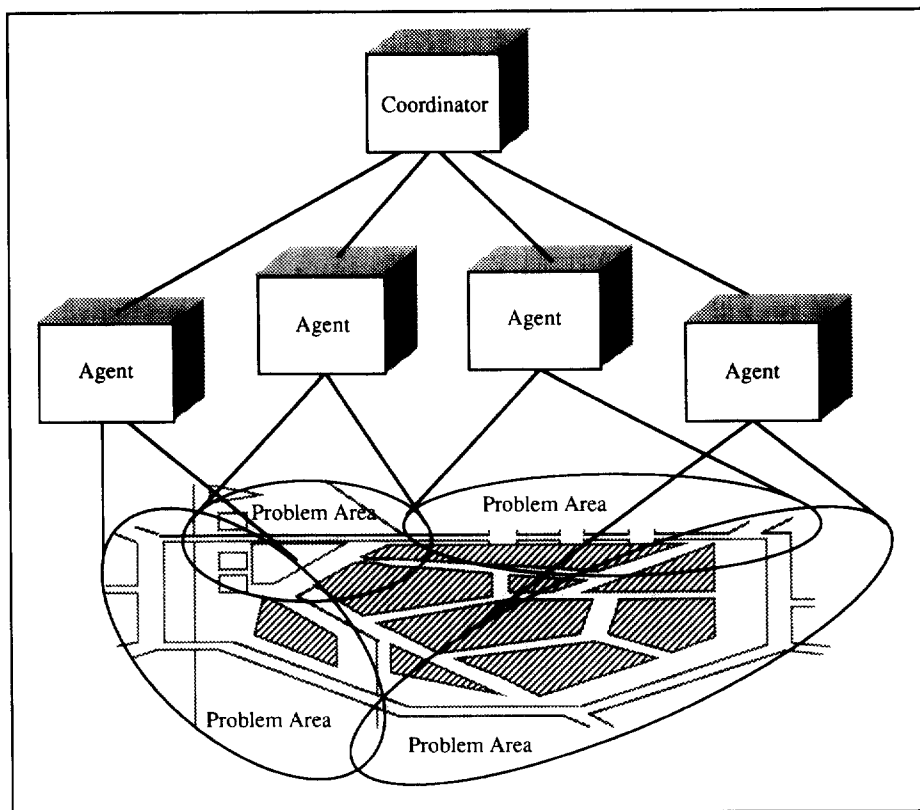


Fig. 1. Basic structure of the control model.

- *Traffic problems*: Knowledge about the detection and diagnosis of the presence of incidents or congestion.
- *Control actions*: Knowledge about the definition of control strategies adequate to solve the different problems.

These knowledge areas are presented in detail in the next three sections.

3.1.1. Physical structure. The aim of the physical structure knowledge is to represent both static information about the network structure of the problem area (components and their relationships) and the dynamic aspects of this structure. The procedures used by this block are the following:

- Methods for data abstraction that perform:
 - (a) Elaboration of values of basic parameters starting from values provided by sensors (i.e. speed, occupancy, traffic volume).
 - (b) Interpretation of the state of the different control devices according to the local view of the agent.
 - (c) Computation of values for variables associated with higher level traffic concepts. For instance, the traffic volume generated in an entrance node or the spatial gradient of speed in certain area.
- Derived structural information, i.e. concepts about the structure that can be automatically derived from the basic information. This information includes the possible routes between an entrance node and an exit node.

3.1.2. Traffic problems. The second knowledge area specializes in the detection and explanation of traffic problems that may appear in its problem area, like congestion or incidents. The goal is to analyze the data recorded by sensors to; (i) point out the current or foreseeable presence of problems taking into account the recent trends of traffic

conditions; and (ii) explain the severity and possible cause of these problems to be used later in the selection of the appropriate control actions.

In general, an agent understands traffic problems as an imbalance between capacity and demand which generates an increase in density that impedes the fluidity of traffic. In a city network, the problems appear when a queue of vehicles propagates to the surrounding streets, blocking intersections and generating a so-called congestion tree. In a motorway, the problems are a consequence of a loss of capacity due to an unexpected incident or to the characteristics of the network.

The problems are analyzed not only by observing the instant where the problem has been detected but with certain deepening in time. In this way, the presence of short-term congestion can be foreseen. This capability of short term prediction is very important for a system willing to propose control actions that contribute to the prevention of problems, and not only the solution of currently detected problems. This aspect is important because the effort needed to avoid congestion is significantly less than that necessary to eliminate existing congestion.

The approach made to describe a problem and to explain its causes is based on associating a traffic problem with an excess traffic volume, in such a way that a problem appears because the traffic demand received by a section of the network is bigger than the capacity offered by that section. According to this, the characterisation of a problem is based on the following issues:

— *Where is the problem?*

A problem is located in a so-called critical section, i.e. the point where the congestion starts due to a lack of capacity. Examples of critical sections are: works on the road disabling a lane, an entry ramp to a highway with a very high traffic demand, a turn movement controlled by a signal with a short green time, an accident that blocks a lane etc.

— *How severe is it?*

The size of the imbalance between the network capacity and the traffic demand is a measure of problem severity. This imbalance is characterised by the excess of arrivals (measured in veh/h), which is understood as the minimum demand decrease (or capacity increase) needed to solve the actual problem. It may be computed as the difference between the flow demand arriving at the section where the problem is located, obtained upstream of the congested area, and the capacity in that section. An accurate estimation of problem severity is a key aspect to pose an effective control plan.

— *What is the cause of the problem?*

Finally, the concept of participation is also managed to distribute the causes of the problem among the paths that cross the critical section carrying a significant traffic flow. Those paths are called involved paths. This information is important to understand the origin of the problem and to pose adequate actions on the points of the network that generate the traffic.

The goal of problem identification reasoning is to analyze data from sensors to detect which are the active critical sections, indicating the excess of arrivals and the paths involved. In order to carry out this task, the necessary knowledge is represented using frames. The frames language provides a traffic language that can be used by an expert to express his experience in detecting and explaining problems. A frame is a collection of state variables for traffic and signals, characterising the situation in an area, whose values are included in intervals representing prototypical situations (each one of these variables is called a slot).

The state variables of a frame are organized in the following sections:

- *Control state.* This includes a collection of slots representing the state of the VMS panels or traffic lights that influence the traffic entering problem areas.
- *Entrances and exits.* This is a collection of slots representing the flow structure in the problem area, so every frame defines a class of flow scenarios. It includes the in

and out flows during a short time period. This time period must be long enough for a vehicle to cross the problem area using the longest path.

- *Significant sections.* This is a collection of slots representing internal traffic states, where flow and occupancy values are given.
- *Detectors state.* This is a set of variables for basic traffic magnitudes directly measured by sensors.
- *Detection areas.* This includes variables to detect decreases in traffic fluidity due to loss of speed or/and increase in occupancy.
- *Critical sections.* This is a set of slots characterising the state of potentially problematic points. Every focus is defined by a list of slots describing a problem area that due to the demand state on specific hours of the day, to unexpected events, or to its topology may present imbalances between incoming demand and capacity for a time period.

The designed modeling approach aims not only to detect traffic problems but also to explain conflicting situations based on the differences between incoming traffic demand and infrastructure capacity. Obviously, congestion can be identified with the occupancy value provided by a detector in the congested area. However, with the aim of defining a reasonable theory for different traffic problems, these have been described using a behavior model. This behavior model is a conceptual framework that can be used to detect inconsistencies in data or to identify plausible explanations of aspects of the situation useful to justify control actions.

Every frame is defined with a representation language designed for this purpose. Figure 7 shows an example of problem frame written with such a language.

The reasoning model checks which problem prototypes in the frame knowledge base match the current situation described by sensor data. For every matched frame, a potential explanation of the problem is obtained, providing criteria to later select suitable control actions.

3.1.3. Control actions. The goal of the traffic control process is to choose proposals of VMS panel displays which induce drivers to take paths that do not pass through congested areas. The starting point is the set of paths with the greatest influence on the state of the critical sections, called problem paths. The objective is to find recommendations (warnings of slow traffic, warnings of congestion, recommended paths etc.) that decrease traffic on the set of problem paths. The input of this process is the state of every problem focus indicating if it is free, overloaded or an incident.

In this sense, there is another kind of frame used to describe control recommendations. These frames are called path-use frames. Every path-use frame has the following organization of slots:

- *Panel state.* It includes a slot for every VMS panel and its value is one of the possible states of the panel.
- *Path uses.* It includes a slot for every origin–destination pair with sub-slots for the different routes. The value of a sub-slot is the percentage of vehicles that go from the origin to the destination taking a particular route.
- *Critical sections.* It includes a slot for every critical section. The value of the slot is the state of the section (free, overloaded, incident). The existence of this type of slot maintains consistency between the state of the VMS panels and the state of the problem area.

Path-use frames are written in a language similar to the one used to write problem frames. Figure 8 shows an example of a path-use frame.

The diversion reasoning process goes over the path-use frame knowledge base analyzing every frame which is impacted by the signal proposal defined in the frame on the problem improvement (i.e. how much decrease in the incoming flows to congested or almost congested areas can be obtained). Finally, the frames selected are those whose signal removes or reduces the excess arriving traffic to the congestion's focus.

3.2. Model of the co-ordinator

The result of the individual reasoning performed by an agent is a set of local control proposals for its problem area. All the local control proposals defined by the agents are analyzed by the so-called co-ordinator to build global proposals by coherently synthesizing the local proposals. The co-ordinator specializes in integrating the different local proposals provided by a set of agents to build signal recommendations for the whole network according to a specific traffic state. Its reasoning is supported by two knowledge areas:

- knowledge to determine the compatibility of different proposals,
- knowledge to define a model of agent priority that is used to take decisions when conflicts appear.

3.2.1. Compatibility of control actions. This is the knowledge about coherence between control actions which detects when two local proposals are incompatible during the process of building a global proposal. The incoherent situations between signal proposals include two possibilities:

- Physical conflicts caused by different actions on the same signal device. For instance, two agents propose displaying different messages on the same VMS panel.
- Semantic incoherence between proposals for different signal devices that may be incoherent from the point of view of a user traversing a given path. For instance, a panel recommending a certain speed and another one, in the same or a neighbouring area, suggesting a significantly different speed.

The representation used to establish incompatible situations is rule based, expressing the inconsistency between signal actions. For example, a message suggesting a certain route is incompatible with a message of road works in the same path. Rules are managed by a procedure whose inputs are pairs of control actions and whose output is their qualification as compatible or incompatible. Rules are formulated as logic clauses.

3.2.2. Priorities between agents. When incompatibility problems appear, a model of agent priority is used to decide which agent must change its signal proposal. The knowledge required to do this task includes control strategy rules which decide what agent must change its proposal according to several criteria, such as the importance of the area, the value of the state variable, etc.

This rule-based representation provides a powerful capacity to include more specific control strategies in such a way that this module is capable of retaining a wide range of control strategies learnt by the operators in their work in the control centre.

3.3. Summary of the TRYS model

The proposed model is summarized in Fig. 2, based on:

- A set of data bases, each one describing the physical structure of a problematic area.
- A collection of frames representing prototypical problematic situations, using the general format described in Fig. 7, which the expert user will be able to formulate, consult and evaluate according to his/her experience through the fulfilment of typical computer operations for text editing.
- A knowledge base to support control decisions organized as a collection of control frames where, according to the format described in Fig. 8, different patterns of traffic behavior are described in terms of the impact provoked on path uses and critical sections. This set of frames is also easy to consult and modify by expert users. It is interesting to remark on the way of using this frame base: it may be applied in two reasoning modes: (i) in the prediction mode to evaluate the possible impact of a set of messages on the traffic state; and (ii) in the control mode to identify which frames produce a significant change in the state of the problem focus with respect to the current situation. In the prediction mode, messages are premises and focus states are conclusions. In the control mode, the change in the state of the

problem focus is the input and the messages to be displayed are the output. So, the same declarative knowledge is used to achieve two different goals.

- The coordination knowledge is based on two rule bases:
 - One that models the different unacceptable displays of signals or VMS panels, that are incompatible both from the point of view of control and from the point of view of consistency for the users traversing the different paths in the network (i.e. the user must find messages along a path that are consistent between each other and consistent with the traffic situation).
 - The rules for solving the different detected conflicts propose alternative messages for the conflicting VMS panels starting with general messages that maintain consistency.

These rule bases may be inspected by traffic experts in such a way that the model is viewed by the user as a collection of text that is quite understandable and thus, easy to maintain and improve. The possibility of improvement derives both from the personal analysis of the expert and by inspection of the explanations provided by the system for specific decisions during on-line operation.

4. EXAMPLE OF THE REASONING MODEL OF TRY5

This section illustrates, through the use of two examples, the type of reasoning performed by the TRY5 system. The first example shows the way a single agent detects congestion in its problem area and how that agent specifies control action proposals to

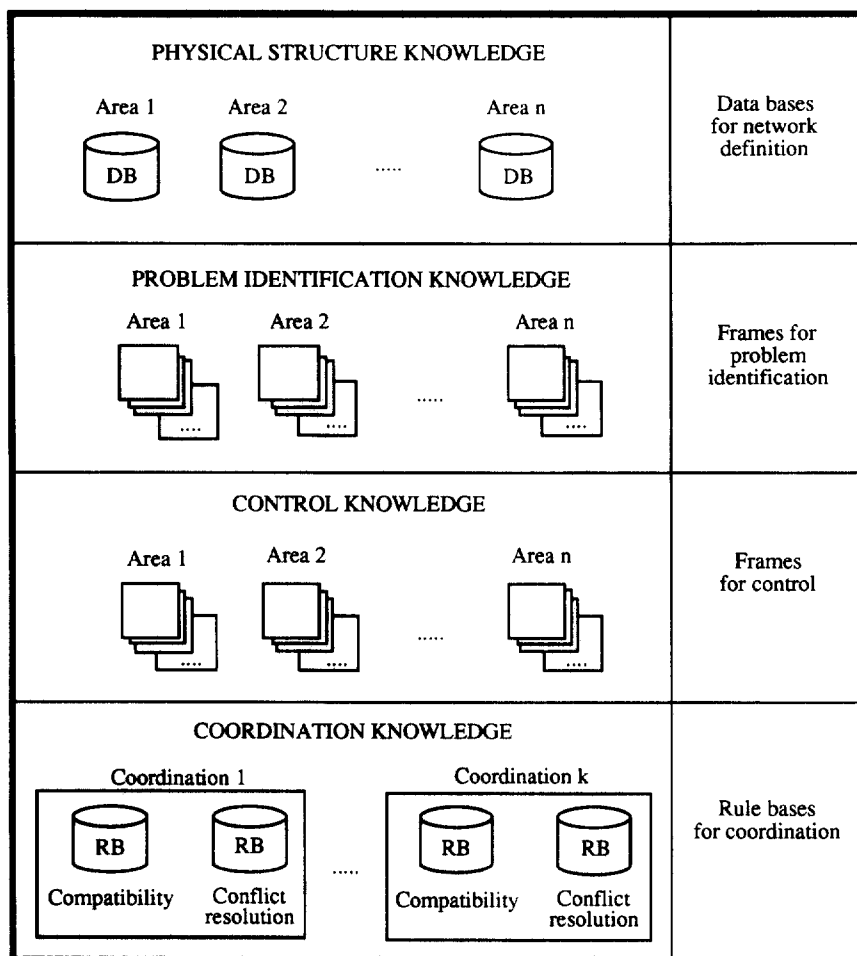


Fig. 2. Structure of the TRY5 model.

solve or improve the situation. The second example shows the way a co-ordinator manages the local control actions proposed by different agents to generate several consistent global proposals.

Both examples and the network used are taken from the TRYS model developed for a wide area of Madrid city. The network covers about 10 km of the M30 urban ring, and a similar length of two main access roads to downtown: N-III and N-II roads. Along the M30 and the two motorway accesses loop detectors and variable message sign panels (VMS) are available. Those devices are connected with the Traffic Control Center through fiber optic communications, which makes it possible to receive data from sensors (every minute) and to display messages on VMS panels in real time. Figure 3 shows a map of this area.

Six agents are considered in this area:

- two for the M30 urban ring, one for M30 southbound traffic and the other one for M30 northbound traffic,
- two for the N-II access, for inbound and outbound traffic,
- two for the N-III access, for inbound and outbound traffic.

In addition, there is a co-ordinator which controls this group of agents. The first example shows the reasoning made by the agent controlling the M30 northbound problem area, and the second one presents a situation where the co-ordinator generates global proposals from different local proposals obtained, with analogous reasoning to that performed by the M30 northbound agent, by the other agents.

4.1. Example of the agent's reasoning

The M30 northbound agent receives data from 11 detectors and may act on 9 VMS panels (Fig. 4). It performs three tasks:

- data abstraction,
- problem identification, and
- control proposal selection.




Fig. 3. Madrid test site.

4.1.1. *Data abstraction.* The information received from the detectors comprises the last 5 min of data. Three traffic characteristics are defined for every detector: speed (km/h), flow (veh/h) and occupancy (%). Since each detector records its values every minute, the information associated with a detector includes a time series of five values for each characteristic. Figure 5 shows an example of data sent by some of the detectors in the area controlled by this M30 northbound agent.

The goal of this stage is to generate an abstract view of the traffic state, close enough to the level of reasoning where the problem detection is performed. The knowledge used for




Fig. 4. Scheme of the M30 northbound problem area.

this operation includes a description of the problematic area's physical structure. That structure is formulated with the following traffic concepts:

- node: entrance and exits points to the area,
- section: part of a road or street, characterized by its capacity,
- detection area: type of section specialized in detecting strong variations in the traffic behaviour,
- detector: sensor on the road,
- panel: a VMS panel, with the set of possible states,
- group of traffic lights, with its set of possible states.

Figure 4 shows the whole structure of the problematic area, where there are no traffic lights but including the following components:

nodes: Madrid north, NIII, NII and M30 south,
 sections: R1, R2, R3, R4, T1 and T2,
 detection areas: A1, A2, A3, A4, A5 and A6,
 detectors: E16C, E18C, E20C, E22C, E24C, E26C, E28C, E16R, E18R,
 E24R, E26R,
 panels: P1, P2, P3, P4, P5, P6, P7, P8, P9.

In the data abstraction phase the goal is to abstract the data from the sensors. In order to estimate the traffic demand, the following points must be taken into account:

- Traffic volume at entrance and exit points of the problem area. If the five values of a flow time series for detector D_k are F_1, F_2, \dots, F_5 and this detector records the

E16C:					
Flow	41	44	40	39	39
Occupancy	6	7	6	5	5
Speed	75	77	76	76	76
E18C:					
Flow	59	54	58	50	68
Occupancy	9	8	7	8	11
Speed	61	62	63	61	60
E20C:					
Flow	45	26	37	30	24
Occupancy	8	28	34	39	45
Speed	68	40	26	23	18
E22C:					
Flow	47	48	41	59	47
Occupancy	8	8	6	14	8
Speed	76	75	73	66	73
E24C:					
Flow	41	44	40	38	38
Occupancy	6	7	6	5	5
Speed	75	77	76	76	76
E26C:					
Flow	61	53	49	48	46
Occupancy	11	8	8	7	7
Speed	72	73	72	71	72
E28C:					
Flow	41	44	40	38	38
Occupancy	6	7	6	5	5
Speed	75	77	76	76	76
E16R:					
Flow	8	8	6	14	8
Occupancy	7	11	8	8	7
Speed	51	52	53	51	50
E18R:					
Flow	9	8	7	7	6
Occupancy	6	6	9	10	7
Speed	51	48	52	53	51

Fig. 5. Example of data recorded by the detectors.

entrance (or exit) of node N_i , then the traffic volume is the sum of all those values (veh/h instead of veh/5 min):

$$V_i = 12 \sum_{j=1}^5 F_j \text{ veh/h}$$

- Saturation levels of sections. The flow of a section is calculated using the flow recorded by the detectors within the section. The saturation level S_i is the flow F_i divided by the capacity C_i of the section expressed as a percentage:

$$S_i = 100 \frac{F_i}{C_i}$$

This measure gives information about the usage degree of the section so that if a section has a saturation level near 100, it is a possible point where congestion may be generated.

In order to detect variations in the continuity of the traffic flow, the following parameters are also calculated:

- Temporal gradients of detector magnitudes (flow, occupancy and speed) which give a view of the short term evolution. These gradients are calculated by subtracting the first value from the last value of a formula defined for the time series of the last five values (data time interval is 1 min). The position of these five values in the time series is considered $\{-2, -1, 0, 1, 2\}$ so the temporal gradient is $S'_2 - S'_{-2}$. The formulae applied to obtain the S'_n values are the following: $S'_n = c_1 \times n + c_0$

$$c_0 = (1/5) \sum_{n=-2}^2 S_n, \quad c_1 = \frac{\sum_{n=-2}^2 n \times S_n}{\sum_{n=-2}^2 n \times n}$$

S_n is the value of the time series at the moment n .

- Spatial gradients which show differences of state between consecutive detectors. In a detection area there are always two detectors D and D' , the downstream and upstream detectors, respectively. Considering V and V' to be the mean values of the time series for a certain variable (e.g. speed) measured by D and D' , the spatial gradient for this variable is $V - V'$.

The result of applying these formulae to the data presented in Fig. 5, is shown in Fig. 6.

4.1.2. Problem identification. After the previous phase is carried out, the goal now is to match, using the interpreted data, one or more frames representing problems. The result is a detailed description of the problems with their location, severity and causes. Figure 7 shows a frame describing a problem caused by an incident blocking a section. Note that, as described earlier, a frame is part of the conceptual model that may be inspected and updated by the traffic expert in the same way as a data base.

The inference procedure starts once the results of the data abstraction phase have been received (saturation levels of the sections, spatial and temporal gradients etc). When these data sufficiently match the data included in a frame, the characterization of the problem described in the frame is generated using the slots of the critical sections part.

The result of the matching process need not be 100% accurate. Fuzzy logic techniques are used to carry out this process so, after the inference, the frame (or frames, because the model may be non-deterministic) closer to the current traffic state are selected. Note that this process works even when partial information is received; for instance, when one or more detectors are out of service.

For the example developed in this section, the traffic state presented in Fig. 4 matches the frame presented in Fig. 7. The conclusions of this problem detection phase are the slot values deduced for the critical sections. In this case, the values describing the problem are:

SECTIONS				DETECTORS				
(R2)	flow:	444	veh/h	(E16C)	temporal gradient of speed:	1	Km/h	
	speed:	51	Km/h		temporal gradient of occupancy:	-1	%	
	occupancy:	7	%		(E18C)	temporal gradient of speed:	-2	Km/h
	saturation level:	29	%			temporal gradient of occupancy:	2	%
(R3)	flow:	468	veh/h	(E20C)	temporal gradient of speed:	-47	Km/h	
	speed:	51	Km/h		temporal gradient of occupancy:	34	%	
	occupancy:	8	%		(E22C)	temporal gradient of speed:	-6	Km/h
	saturation level:	31	%			temporal gradient of occupancy:	3	%
(A1)	flow:	3468	veh/h	(E16R)	temporal gradient of speed:	-2	Km/h	
	speed:	61	Km/h		temporal gradient of occupancy:	-1	%	
	occupancy:	8	%		(E18R)	temporal gradient of speed:	2	Km/h
	saturation level:	64	%			temporal gradient of occupancy:	2	%
(A2)	flow:	1944	veh/h				
	speed:	35	Km/h	DETECTION AREAS				
	occupancy:	30	%	(A1)	spatial gradient of speed:	15	Km/h	
	saturation level:	36	%		spatial gradient of occupancy:	-3	%	
(A3)	flow:	2904	veh/h	(A2)	spatial gradient of speed:	26	Km/h	
	speed:	72	Km/h		spatial gradient of occupancy:	-22	%	
	occupancy:	8	%	(A3)	spatial gradient of speed:	-37	Km/h	
	saturation level:	53	%		spatial gradient of occupancy:	22	%	
.....							

Fig. 6. Results of data abstraction.

(section A2)

state: incident,
 excess: 1200 veh/h,
 participation:

[50,70] M30 south → Madrid north: by A2,
 [30,50] M30 south → NII: by A5,
 [20,40] NIII → Madrid north: by A5.

The excess of arrivals to section A2 is estimated using the temporal gradient of flow and observing the flow decrease in that part of the network. The participation of the different paths crossing section A2 is explicitly predefined in the frame but it could be obtained by estimating an origin-destination matrix starting from the flow values and the structure of the network. The participation values are minimum-maximum intervals because the aim is to provide a pseudo-qualitative view capable of giving an approximate, common sense, perspective of the traffic state.

Both, temporal gradient flow and excess of arrivals to critical sections allow characterization of congested situations in motorway and city networks. In the motorway, the sudden drop in flow and speed shows the presence of a capacity loss. In the city network, this may be less evident because of the lower speed in streets but a mixed procedure may be used using both the temporal gradient and the estimated excess. In this way, low values of speed and temporal gradients near zero provide signs of congestion that may be confirmed or rejected with the evaluation of arrivals to the critical section from an isocrone of, for instance, -3 min.

In any event, the interest of this open architecture is that several methods may be tuned to model problem behavior in different sites and networks, in such a way that different gradient thresholds and procedures to estimate excessive arrivals may be tested and refined to fit the characteristics of every possible congestion in the network.

PROBLEM FRAME incident congestion at A2			
SECTIONS			
(R2)	saturation level:	in [0, 75]	[%],
(R3)	saturation level:	in [25, 100]	[%],
DETECTORS			
(E18C)	speed:	in [50, 180]	[Km/h],
	occupancy:	in [0, 25]	[%],
(E20C)	speed:	in [0, 50]	[Km/h],
	occupancy:	in [25, 100]	[%],
	temporal gradient of speed:	in [-100, -40]	[Km/h],
	temporal gradient of occupancy:	in [20, 90]	[%],
(E18R)	speed:	in [50, 180]	[Km/h],
	occupancy:	in [0, 25]	[%],
(E22C)	speed:	in [50, 180]	[Km/h],
	occupancy:	in [0, 25]	[%],
DETECTION AREAS			
(A2)	spatial gradient of speed:	in [20, 150]	[Km/h],
	spatial gradient of occupancy:	in [-100, -20]	[%],
CRITICAL SECTIONS			
(A2)	state: incident,		
	excess: (value to be obtained with the temporal gradient of flow)		
	participation:		
	[50, 70]	M-30 South -> Madrid North: by A2,	
	[30, 50]	M-30 South -> NII: by A5,	
	[20, 40]	NIII -> Madrid North: by A5.	

Fig. 7. Problem frame describing congestion caused by an incident.

4.1.3. *Control action proposals.* Finally, after the problem is detected and characterised, the last phase is to determine one or more control action proposals capable of improving the traffic state. In this example the goal is to generate sets of messages to warn about the presence of congestion, to give the drivers the chance to take an alternative route.

In order to achieve this goal, the knowledge used specifies the way drivers select routes to go from an origin to a destination through the problem area or by passing the problem area, and how the messages on VMS panels affect that decision. A first version of this knowledge was obtained from the experience of the traffic operators in the Traffic Control Center of Madrid. The advantage of this open modeling approach is that according to the new experience gained, the traffic operators can revise and improve this knowledge through direct manipulation of the corresponding slot values in the frames contained in the files of knowledge bases. One of the frames that represents this knowledge, for the congested situation shown in Fig. 4, is described in Fig. 8.

The path use section describes the effect on traffic distribution provoked by the configuration of panels defined in the upper section. For instance, there are three paths (that cross sections A2, R2 and R4) to go from the M30 south node to the Madrid north node. Panels P3 and P4 warn about congestion at A2, so some drivers will avoid the main path (which crosses the A2 section) and will take the path that goes through R2. This effect is represented with numeric intervals.

The inference procedure of the control stage starts with the problem description obtained in the previous stage and looks for frames whose set of messages may decrease

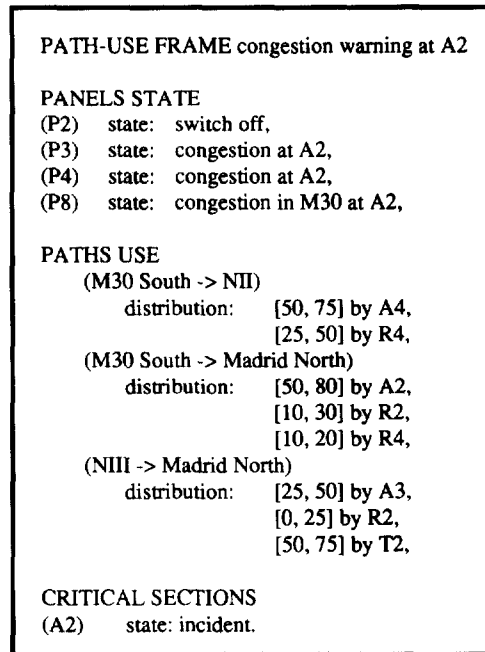


Fig. 8. Path-use frame applicable to an incident congestion.

traffic volumes in the problematic paths. This procedure estimates the reduction provided by every frame using the path selection values and the current incoming traffic to the critical sections. Those frames providing a significant improvement in traffic conditions are selected and their sets of messages are presented as local control proposals of the corresponding agent.

For this example, consider that there are two frames that can reduce traffic arriving at A2. The first one acts on the P3 and P4 panels and the second one furthermore proposes a message on P8 (this second proposal corresponds to the frame presented in Fig. 8), as follows:

Proposal one: "low congestion warning at A2"

P3: congestion at A2,
P4: congestion at A2,

Proposal two: "congestion warning at A2"

P3: congestion at A2,
P4: congestion at A2,
P8: congestion in M30 at A2,

Each proposal has an associated explanation which shows the estimation of the impact it causes on traffic. For instance, the explanation for proposal one may be:

The estimated effect of the control proposal "low congestion warning at A2" is an excess reduction from 1200 to 600 veh/h. The effect on the main paths is:

-450 veh/h FROM M30 south TO Madrid north BY A2
-250 veh/h FROM M30 south TO NII BY A4

The idea is to show how much the flow along the main paths can be decreased using the proposal. It can be observed that the total decrease or excess is less than the sum of the partial reductions in the two paths, this is due to the fact that there is still traffic flowing from NIII to Madrid north crossing the A2 that is not warned. It is important to remark that the numeric values presented must not be understood as exact numbers but as an estimation that allows creation of criteria to select an adequate configuration of VMS panels.

4.2. Example of the co-ordinator reasoning

This second example illustrates how the co-ordinator reasons, and in particular the co-ordinator of the Madrid east problem area. The co-ordinator specializes in generating control plans for the whole traffic network by combining the local control proposals provided by different agents.

The inputs of this task are lists of < agent, state, control proposals > where each agent gives a measure of the traffic state in its area and a list of local proposals to improve the situation. The value of the state variable is qualitative (free, light, medium and severe) and is deduced from the excess values of critical sections. The value of the control proposals argument is a set of lists with panel-message pairs. Each one of these lists define a local proposal and they are sorted in decreasing order of preference.

This example considers a situation where there are three problems in the area described in Fig. 4, as represented in Fig. 9:

- An incident at A2 is detected and evaluated by the M30 northbound agent.
- A pre-congestive situation at T2 is detected and evaluated by the NIII inbound agent.
- A regular problem at T1 is detected and evaluated by the NII inbound agent.

The state and proposals provided by the six agents of this example might be the following:

- Agent M30 Northbound,
State: medium,
Proposal 1: P3: congestion at A2,
P4: congestion at A2,
P8: congestion in M30 at A2.
Proposal 2: P3: congestion at A2,
P4: congestion at A2.
- Agent M30 Southbound,
State: free.
- Agent NII inbound,
State: medium,
Proposal 1: P6: congestion at T1,
P7: congestion at T1.
Proposal 2: P7: congestion at T1.
- Agent NII outbound,
State: free.
- Agent NIII inbound,
State: light
Proposal 1: P8: slow traffic at T2,
P9: slow traffic at T2.
Proposal 2: P9: slow traffic at T2.
- Agent NIII outbound,
State: free.

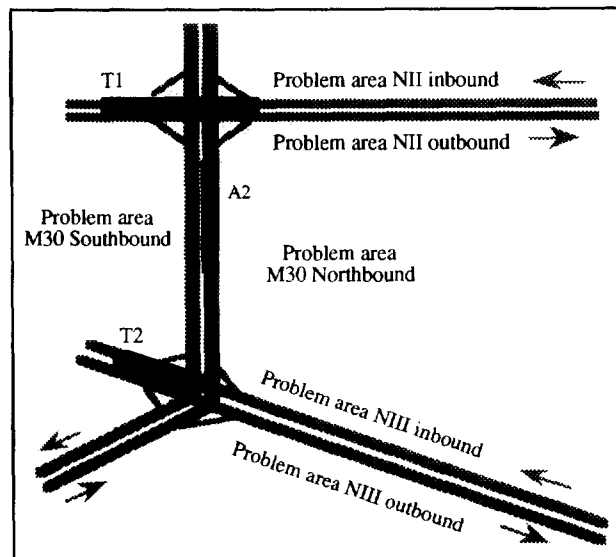


Fig. 9. Screen of Madrid problem area with three problematic situations.

Once the co-ordinator receives this information about the state of the whole network it performs two tasks:

- detecting conflicts among local control plans, and
- solving these conflicts in a way that is coherent with the global traffic state, by finding consistent combinations of proposals.

Given the mentioned proposals, the following conflicts may be detected:

Between the M30 northbound and NIII inbound agents:

- Incompatibility between panels P4 and P8, because the message for panel P4 will make some drivers choose to leave M30 and enter Madrid using the NIII access, and the message proposed for panel P8 will produce a similar effect but in NIII.
- Physical incompatibility of different messages for the same panel (P8).

Between the M30 northbound and NII inbound agents:

- Incompatibility between panels P3 and P6, because the message for panel P3 will make some drivers choose to leave M30 and enter Madrid using the NII access, and the message proposed for panel P6 will produce a similar effect but in NII.

After those conflicts are detected the goal of the co-ordinator is to make up global proposals applying three criteria: (i) taking into account the order of local proposals given by the agents; (ii) using as many proposals as possible to build global proposals avoiding the conflicts; and (iii) when several global proposals are composed with the same number of local proposals the order among these global proposals is established according to the severity of the traffic state in the different problem areas.

When one or more conflicts appear during the development of a global proposal they are overcome using priorities between agents. These priorities are determined by the control strategy rules. For instance, according to the characteristics of the traffic in the example network, the congestion in M30 is usually more critical than that on the accesses, and if there is more than one area of congestion in M30 priority must be given to the agent controlling the area with the most severe problem.

This knowledge is represented with rules such as the following:

```

IF   (agent-1) name = M30 south, state = X
      (agent-2) name = NIII, state = Y
      X=Y,
THEN agent-1 is better than agent-2
IF   (agent-1) state=light or medium
      (agent-2) state = severe
THEN agent-2 is better than agent-1

```

For the previous example, the following priorities would be established:

- Conflict between M30 north and NII: M30 north is given preference because problems in this area are usually worse than problems in NII.
- Conflict between M30 north and NIII: M30 north is given preference because the state of this area is worse than the state of NIII.

Taking into account the previous considerations in building global proposals, the co-ordinator of the Madrid east problem area would generate the global proposals in the following order until it reaches the maximum number of proposals required by the user:

- Proposal 1 of M30 northbound agent
- Proposal 2 of NII inbound agent
- Proposal 2 of NIII inbound agent
- Proposal 2 of M30 northbound agent
- Proposal 2 of NII inbound agent
- Proposal 2 of NIII inbound agent
- Proposal 1 of M30 northbound agent
- Proposal 2 of NII inbound agent
- Proposal 2 of M30 northbound agent
- Proposal 2 of NII inbound agent

.....

Therefore, if the maximum number of global proposals requested is four the answer of the co-ordinator would be the following:

Proposal 1:

P3: congestion at A2,
 P4: congestion at A2,
 P7: congestion at T1,
 P8: congestion in M30 at A2,
 P9: slow traffic at T2.

Proposal 2:

P3: congestion at A2,
 P4: congestion at A2,
 P7: congestion at T1,
 P9: slow traffic at T2.

Proposal 3:

P3: congestion at A2,
 P4: congestion at A2,
 P7: congestion at T1,
 P8: congestion in M30 at A2,

Proposal 4:

P3: congestion at A2,
 P4: congestion at A2,
 P7: congestion at T1,

5. IMPLEMENTATION ISSUES

This section gives a brief overview of the main implementation issues concerning the TRYS system and the traffic models developed with it. One of the main difficulties of the implementation of an application like TRYS is the management, in a structured way, of different types of knowledge such as declarative and procedural, generic and domain specific, using several knowledge representation formalisms (rules, tables, frames) together with modules based on algorithmic approaches. In addition, all this complexity must be presented to the final user in such a way that it is understandable and flexible in accepting changes.

In order to cope with these requirements, TRYS has been implemented using an innovative software environment for application development using a knowledge modeling approach. The environment is called KSM (knowledge structure manager) (Cuenca and Molina 1994; Molina and Cuenca, 1995). KSM understands the development of an application as a knowledge modeling activity instead of the traditional knowledge acquisition of prefixed types of representation. The whole KSM environment is conceived to support models organized as a structured collection of knowledge blocks, the so-called knowledge units, each with its own knowledge representation and inference procedures.

A knowledge unit is described in two parts:

- what it knows; represented by a collection of knowledge areas that may be also described as knowledge units.
- what the unit does, represented by a collection of tasks carried out with methods.

The whole knowledge of a system built using KSM constitutes a network of knowledge units. There is one knowledge unit that represents the entire application that is decomposed via other units at lower levels that may also be decomposed into other units and so on. The units at the lowest level are called primary knowledge units and usually contain a knowledge base of rules, frames or constraints, and a set of tasks that use the knowledge base to generate responses. A primary knowledge unit does not have to be a knowledge-based module. It may be a neural network, a conventional database or even a conventional program with an algorithmic approach.

The knowledge structure of any traffic control model of the TRYS system, like the ones described in the previous sections, presents the following types of knowledge units: the highest level unit representing the whole application includes the co-ordinator of the whole traffic network and the local controllers named agents. Every agent has knowledge about the characteristics of the physical structure in the local area it controls, about the description and diagnosis of problematic situations and, about the formulation of adequate local control proposals. These three knowledge areas are represented as primary knowledge units whose knowledge bases include elements of a traffic network or frames, and whose tasks perform the navigation across the network, the identification of the frame describing the current problem or the selection of the set of frames with the most suitable control proposals. Figure 10 shows a set of windows of the KSM environment for a small TRYS model with part of its knowledge unit structure.

The important advantage provided by this type of approach is that the distance from the informal way of understanding, expressed at the knowledge level, and that which is more formalised, expressed at a symbolic level (such as the one provided by TRYS), is significantly less than the typical distance to the symbolic level provided by standard programming environments such as the C language. This decrease in conceptual distance increases reliability in model formulation and, to non-programmer traffic experts, the possibility of developing and maintaining their own models.

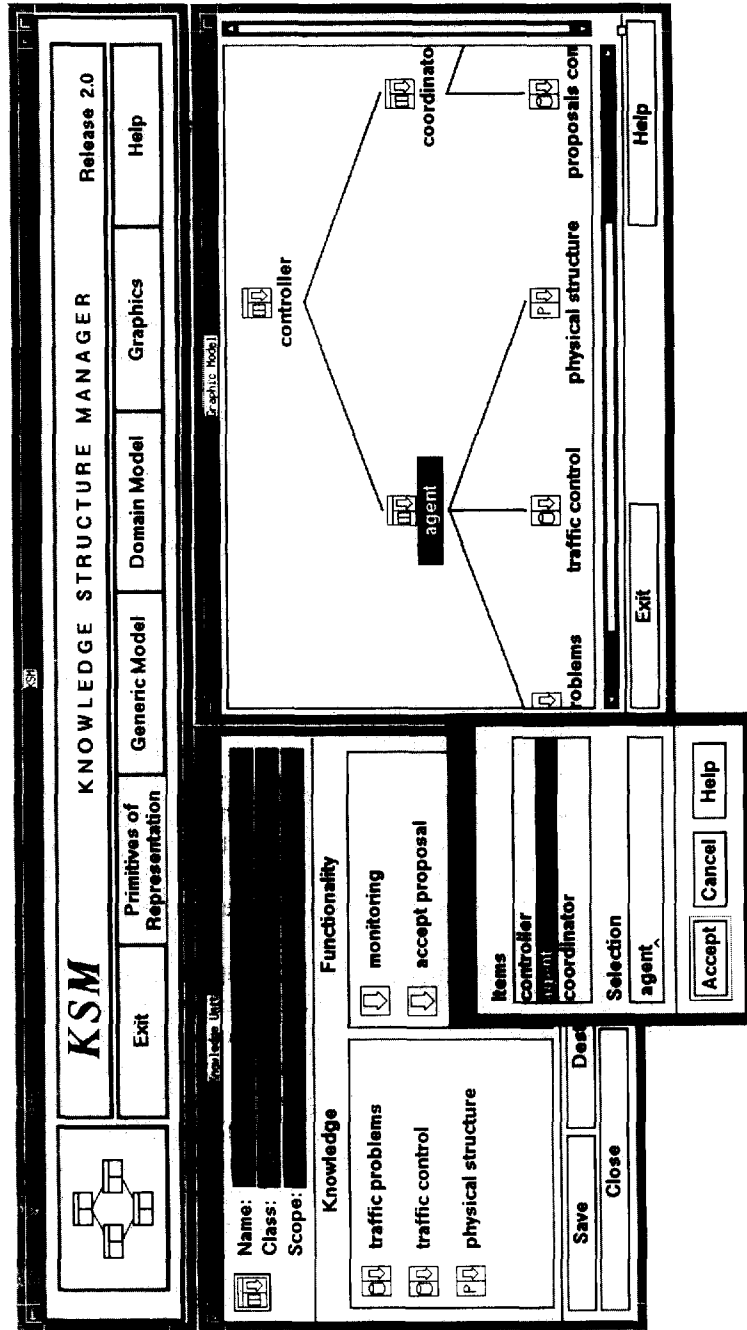


Fig. 10. Knowledge structure of a model built with KSM.




Fig. 11. Initial window of TRYS system.

6. OPERATION OF THE TRYS SYSTEM

The TRYS system can operate in two modes: an on-line mode, where the system works connected to the road sensors and proposes suitable signal plans moment by moment and, an off-line mode, used when the operator is building traffic models through the KSM environment and a traffic simulator. In this section, both operation modes are described presenting some examples of the screens displayed to the user.

TRYS's initial window, necessary to perform any kind of work, is shown in Fig. 11. Using the buttons that appear in this window, the following actions can be performed (read from left to right).

- Exit TRYS
- Go into the KSM environment
- Start the reasoning of the expert system
- Consult the signal proposals of the expert system
- Install a signal proposal on the road
- Consult a help text of TRYS operation

Next, the operation modes are explained: exploitation work (on-line mode) and modeling work (off-line mode).

6.1. Exploitation

In this working mode, TRYS receives every minute the traffic state measured by the detectors and information on the state of the VMS panels. The reasoning of the system is started by clicking the start button in the main menu bar (see Fig. 11). In this way, the system will remain in a closed loop reasoning about the traffic state and proposing signal proposals (if necessary) every 2 – 3 min. The operator can stop this loop by clicking the same button.

Figure 12 shows an example of a screen presented by TRYS at the end of a reasoning cycle where two problems have been detected in the M30. The maps display the results for the M30 ring road (the bigger window), for the accesses and for the exits of the city. The location of a problem is indicated with a colored area (shaded area in Fig. 12), and the presence of new control proposals is pointed out with a sign icon.

Both the shaded areas and the icons are active graphical elements on the map that can be selected with the mouse. Clicking a shadowed area displays new windows that give a detailed description of the problems detected in that area. The information contained in these windows is the following:

- Problem characterisation: where, how severe and what is the approximate participation of every significant path crossing the problem area.
- Justification of problem detection, indicating the elements of the knowledge base that have been used in the deduction.

Clicking a sign icon, TRYS displays new windows with the proposed changes in the current signal state and their associated explanations.

Besides the local consultation of the results obtained in a particular problem area, the global control proposals built by TRYS, through the synthesis of the local proposals, can also be consulted. Clicking on the corresponding button of the main menu bar (see Fig. 11), the global proposals and their explanations are shown. An explanation includes the order of the local proposals that have been used in the synthesis for the global one.

Figure 13 shows in the two windows on the right of the screen a local proposal for one of the two problems detected in M-30 and an associated explanation for the local proposal. The proposal is formulated with a list of VMS panels and the corresponding messages to be displayed. The explanation describes the estimated impact on traffic that the proposal may provoke in terms of decreasing the excessive arrivals and how this excess is redistributed along the main paths. In the bottom of the figure, a small window shows one of the global proposals generated by the co-ordinator where three different messages can be distinguished, two of them for the two problems detected in M-30.

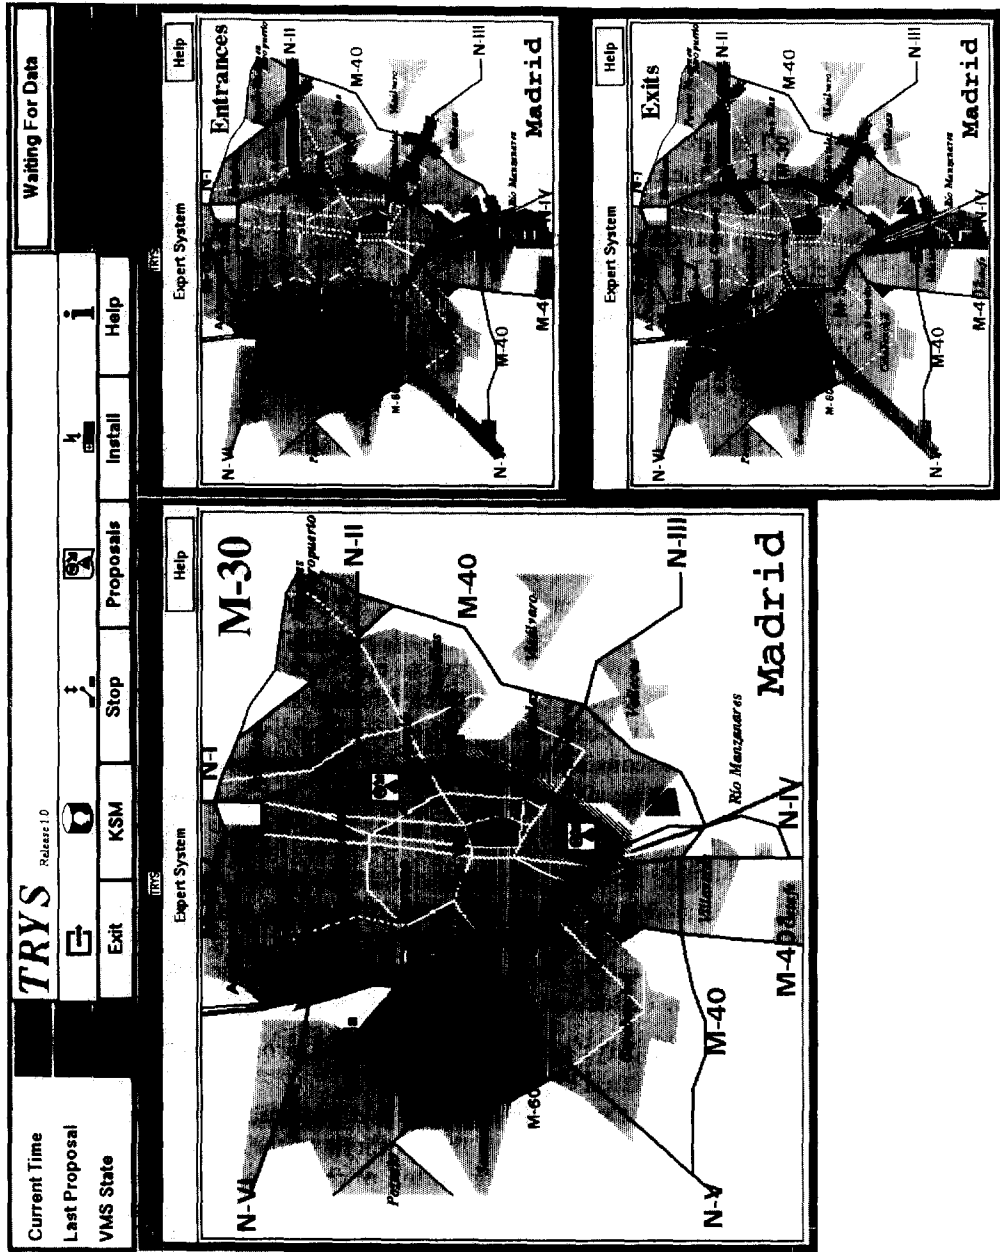


Fig. 12. TRYS presentation of the state of Madrid problem area.




Fig. 13. TRYS presentation of signal proposals for two congestions in M30.

If the system operator wants to send some of these signal proposals to the road, he/she has to click the “*Install*” button in the main menu bar (see Fig. 11). This action will make the system show the list of global proposals. After the selection of one of these proposals, TRYS communicates to an associated Information Management System the need to modify the current control state in the terms specified in the proposal. When these systems modify the state of the signals, a confirmation is sent to TRYS.

6.2. Modeling

The aim of this mode of operation is to access the internal knowledge model of the system with two purposes:

- to consult the knowledge bases to understand the behavior of the system,
- to modify (adding or deleting) elements of the knowledge bases or general parameters to improve system performance.

These activities are carried out with the aid of the KSM environment. KSM works like a knowledge management automatic system that provides the user facilities to navigate through the knowledge of the system, as was described in Section 5.

The whole knowledge of the TRYS system is organized into knowledge units, making up a structure that can be traversed using the KSM environment. Figure 14 shows the result obtained when the knowledge unit “*M30 from ODonnell to NIV*” is consulted. On the left side of the screen can be seen a menu with all the existing units, where any of them can be selected. In the center and right side of the figure, two windows show the selected unit, its components and the structure of the problem scenarios knowledge unit.

7. CONCLUSIONS

Recently, there has been an important increase in the use of advanced information technology applied to adaptive traffic management in urban areas. In addition to the usual signals at intersections whose state may be changed from a remote control center, there are other control resources such as ramp metering, variable message signs (VMS) or advisory radio systems, which provide a high capability for traffic management from control center. In large cities, the number of these control devices is so great that automatic tools are necessary to help control center personnel consistently determine the best control strategy for each traffic situation, for example, by ensuring that traffic signal timing is consistent with messages on VMS panels.

The central difficulty of having automatic tools to manage together all those devices in a co-ordinated and adaptive way is that their management is mainly supported by strategic human intervention based on heuristic and imprecise knowledge, which is difficult or impossible to be captured in algorithms. Accordingly, one solution is to use artificial intelligence techniques which are capable of modeling reasoning procedures similar to the way persons solve problems. This line of work has created a new generation of systems and tools in traffic control using these techniques.

Knowledge-based systems have several characteristics that make them suitable to develop and maintain applications for complex process management environments, as in traffic applications. Users can profit from these advantages in the following sense:

- they can build up and make operative models of knowledge and expertise in the traffic domain,
- the adaptation and maintenance of the models over time with new knowledge or results from experience is feasible because knowledge is expressed in terms familiar to the user,
- the interaction between the user and system is improved due to the transparency and accessibility of the application models (knowledge representation) and of system behavior (exploration of alternative options of reasoning, explanations etc.).

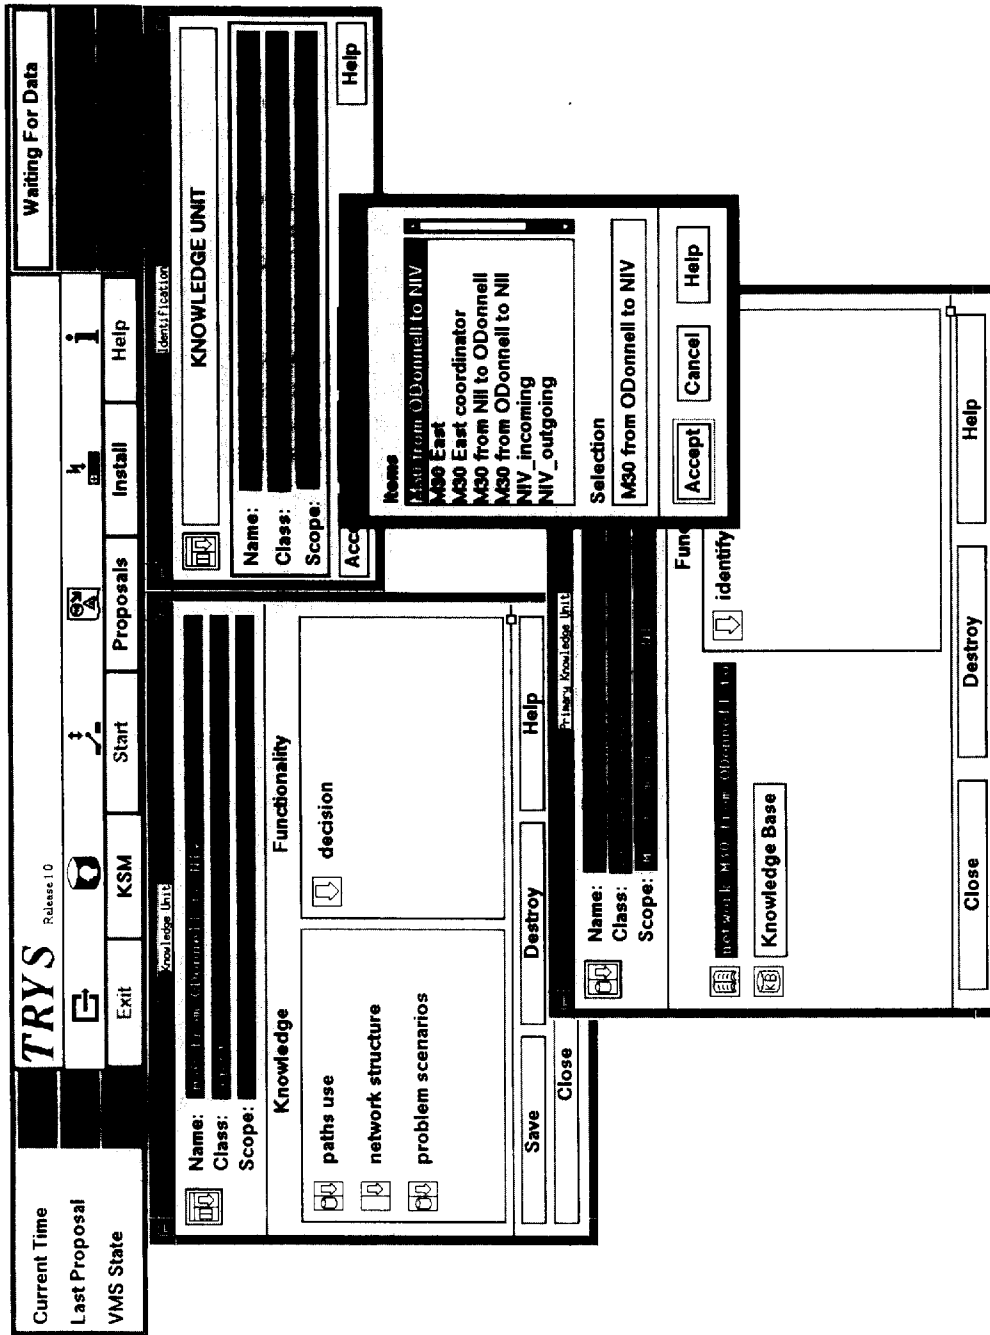


Fig. 14. Access to the knowledge structure of the TRYS system using KSM.

The TRYS system, presented in this paper, proves that knowledge-based models designed for on-line use can be developed to be integrated as components of advanced traffic control centers (TCCs) in order to improve traffic monitoring and management capabilities of current traffic control architectures. Such models identify an additional strategic level in current TCCs, a *traffic knowledge processing layer* developed on top of the available traffic control facilities. This layer allows the development and use of models of knowledge and expertise of traffic operators in the controlled area, providing additional traffic analysis and management criteria which can complement currently applied, optimisation-based automatic traffic control. Furthermore, they also allow modification and improvement over time of traffic control strategies as new knowledge is gained about long-term modifications of traffic behavior through day-by-day operation of the traffic network.

TRYS has been developed as a modeling environment to support development and application of knowledge-based models for network management. A TRYS model provides traffic monitoring functions and control actions. The interface between TRYS and the control system allows the TRYS model to accept input data (i.e. speed and occupancy measurements) from the real-time data collection facilities (via the traffic control computer) and to send back control actions to the traffic control computer. Depending on the traffic control system available at the application site, control actions can range from a set of constraints limiting selection to a library of predefined signal plans (or a library of predefined messages in VMS applications) to a set of constraints on signal setting parameters (i.e. cycle time, phase split and offsets) in a fully adaptive system.

The design of the knowledge representation solutions provided by TRYS has been guided by the double objective of: (1) providing a model for an accurate representation of both the relevant aspects of traffic behavior and the generation of the associated control decisions; and (2) keeping the model within an acceptable level of understandability by the users. Obviously, alternative approaches could be used but any design must try to keep an adequate balance between the models complexity and the user understandability required for reliable maintenance and on-line operation.

The TRYS system and other initiatives carried out in both European programs [EEC's DRIVE program with projects like KITS (Cuenca *et al.*, 1994; Boero *et al.*, 1994) and in US projects such as ARTIST (Deeter and Ritchie, 1993) and FIM (Ritchie and Stack, 1993)], emphasize the increasing interest in the application of artificial intelligence techniques in the transport domain, and in particular, in decision support systems for real time traffic control.

Acknowledgements—This research was sponsored by the Dirección General de Tráfico, Spain. The authors acknowledge Alberto Centeno, Pablo Fernández and José Luis Sierra, our colleagues from the Intelligent System Group (ISYS) Laboratory, for the support provided in software development, and Miguel Luna who was responsible for model operation and maintenance in the control room at the Dirección General de Tráfico. Additional thanks go to the LISITT Laboratory of Valencia University for the support provided in the earlier versions of the Problem Identification Model. Thanks also to Reyes Riera who took care of the editing aspects of the paper.

REFERENCES

- Boero M., Cuenca J., Kirschfink H. and Krogh C. (1993) KITS: a general approach for knowledge-based traffic control models. *Proc. Technical Days on Advanced Transport Telematics*, Brussels.
- Boero, M., Cuenca J., Kirschfink H., Traetteberg H. and Wild D. (1994) The role of knowledge-based models in traffic management and their design. In *Towards an Intelligent Transport System (Proc. First World Congr. on Applications of Transport Telematics and Intelligent Vehicle-Highways Systems)* Vol. 2, pp. 572–579, ed ERTICO. Paris.
- Brachman R. J. and Schmolze J. G. (1995) An overview of the KL-one knowledge representation system. *Cognitive Sci.* 9 171–216.
- Bretherton R. D. and Bowen G. T. (1990) Recent enhancements to SCOOT — SCOOT version 2.4. *Proc. 3rd Int. Conf. on Road Traffic Control*, IEE, London.
- Brownston *et al.* (1985) *Programming Expert Systems in OPS-5*. Addison Wesley, Reading, MA.
- Buchanan B. and Shortliffe E. (1984) *Rule Based Expert Systems: The MYCIN Experiments*. Addison Wesley, Reading, MA.
- Chandrasekaran B. (1983) Towards a taxonomy of problem solving types. *A. I. Mag.* 4, 9–17.
- Chandrasekaran B. (1987) Towards a functional architecture for intelligence based on generic information processing tasks. *Proc IJCAI-87*, pp. 1183–1192. Morgan Kaufmann, CA.

- Cuena J. (1989) AURA: a second generation expert system for traffic control in urban motorways. *IX Int. Workshop on Expert Systems*. EC2, Avignon.
- Cuena J. (1993). Contributions to a knowledge oriented view of software design. In *Knowledge Oriented Software Design* (Cuena J. Ed.), Elsevier, Oxford.
- Cuena J. and Molina M. (1993) A tool concept for knowledge oriented design of applications. *Technical Report FIM 78.1 IA/93*, Technical University of Madrid.
- Cuena, J. and Molina M., (1994) KSM: an environment for knowledge orientated design of applications using structured knowledge architectures. In *Applications and Impacts. Information Processing '94*, Vol. 2. Brunstein K. and Raubold, E., Eds). Elsevier Science B.V.,
- Cuena J., Ambrosino G. and Boero M. (1992a) A general knowledge-based architecture for traffic control: KITS approach. *Proc. Int. Conf. on Artificial Intelligence Applications in Transportation Engineering*. San Buenaventura, CA.
- Cuena J., Martin G. and Molina M. (1992b) An architecture for knowledge-based traffic management for the Expo-92, Sevilla Urban Ring, *OECD 1992*.
- Cuena J., Hernández J. and Molina M. (1994) Case presentation of the use of knowledge based models for traffic management—Madrid. In *Towards an Intelligent Transport System (Proc. First World Congr. on Applications of Transport Telematics and Intelligent Vehicle-Highways Systems)*, Vol. 2. pp. 564–571. ERTICO, Paris.
- Deeter D. L. and Ritchie S. G. (1993) A prototype real-time expert system for surface street traffic management and control. *ASCE, Third Int. Conf. on Applications of Advanced Technologies in Transportation Engineering*, Seattle, WA.
- Dirección General de Tráfico (DGT) (1994) Documento I: Descripción del Sistema; Documento II: Experimentación del modelo de Madrid; Documento III. Volumen I: Modelos de Madrid, Memoria Descriptiva; Documento III, Anexo I: Bases de conocimiento de los agentes de la M30 del modelo TRYS para la DGT; Documento III, Volumen II: Anexo II.
- Duda R. O., Hart P. E., Nilsson N. J. and Reboh R. (1979) A computer based consultant for mineral exploration. *SRI International Artificial Intelligence Center*.
- Gartner N. H. (1983) OPAC: a demand responsive strategy for traffic signal control. *Transpn Res. Record* 906.
- Goldstein I. P. and Roberts R. B. (1977) *The FRL Manual*. MIT, AI Laboratory Memo 409.
- Henry J. J., Farges J. L. and Tuffal J. (1983) The PRODYN real-time traffic algorithm. *Proc. 4th IFAC/IFORS Conf. on Control in Transportation Systems*, Baden Baden, Germany.
- Hunt P. B., Robertson D. I., Bretherton R. D. and Winton R. I. (1981) SCOOT: a traffic responsive method of coordinating signals. Department of the Environment, Department of Transport, TRRL Report No. LR 1014. Transport and Road Research Laboratory, Crowthorne, U.K.
- Lindsay R. K., Buchanan B. G., Feigenbaum E. A. and Lederberg J. (1981) *Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project*. McGraw Hill, New York.
- Lloyd J. W. (1987) *Foundations of Logic Programming*. 2nd Edn. Springer Verlag, New York.
- Lowrie P. R. (1982) The Sydney co-ordinated adaptive traffic system. Principles, methodology, algorithms. *Proc. IEE Int. Conf. on Road Signalling*. London.
- Mauro, V. and Di Taranto (1989) UTOPIA. *Proc. 6th IFAC/IFORS Conf. on Control, Computers and Communications in Transport*, Paris.
- McDermott J. (1988) Preliminary steps toward a taxonomy of problem solving methods. *Automating Knowledge Acquisition for Expert Systems*. (Marcus S. Ed). Kluwer Academic, Boston.
- Minsky M. (1975) A framework for representing knowledge *The Psychology of Computer Vision*. pp. 211–217. McGraw-Hill, New York.
- Molina M. and Cuena J. (1995) Knowledge oriented and object oriented design: the experience of KSM. *Proc. of Knowledge Acquisition Workshop, KAW-95*, Banff, Canada.
- OECD Workshop on Knowledge-Based Expert Systems in Transportation* (1990) Heikki Jämsä, Espoo, Finland.
- 2nd OECD Workshop on Knowledge-Based Expert Systems in Transportation: Operational Experience and Perspectives* (1992) Montreal, Canada.
- Ritchie S. G. and Stack R. (1993) A real time expert system for freeway incident management in Orange County, California. *ASCE, Fifth Int. Conf. on Computing in Civil Engineering*. Anaheim, CA.
- Scemama G. (1989) Creation of a knowledge-based system for urban traffic control: methodology. *Recherche Transport Sécurité*, English Issue 5, June 1990 (translation from RTS, No. 22.)
- Scemama G. (1992) Developpement d'un Systeme a base de connaissance historique pour la gestion du trafic. *Recherche Transport Sécurité*, English Issue 5, June 90 (translation from RTS, No. 22).