Delft University of Technology

Delft Center for Systems and Control

Technical report 03-006

Real-time traffic management scenario evaluation*

S.P. Hoogendoorn, H. Schuurman, and B. De Schutter

If you want to cite this report, please use the following reference instead:
S.P. Hoogendoorn, H. Schuurman, and B. De Schutter, "Real-time traffic management scenario evaluation," *Proceedings of the 10th IFAC Symposium on Control in Transportation Systems (CTS 2003)*, Tokyo, Japan, pp. 343–348, Aug. 2003.

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.24.73 (secretary)
URL: https://www.dcsc.tudelft.nl

REAL-TIME TRAFFIC MANAGEMENT SCENARIO EVALUATION

S.P. Hoogendoorn * **H. Schuurman** **
B. De Schutter ***

* *Faculty of Civil Engineering and Geosciences, Delft University of Technology*
** *AVV Transport Research Centre, Ministry of Transport, The Netherlands*
*** *Faculty of Information Technology and Systems, Delft University of Technology*

Abstract: This paper described a prediction system called BSES (Boss Scenario Evaluation System), which can evaluate control scenarios in real time, predicting their effects in terms of various measures of effectiveness. The system is case-based, i.e. it uses either synthetic or real-life examples of the effect of control scenarios under different circumstances. It determines the similarity of the current situation to different examples in the case-base using fuzzy logic, and uses agent-technology to predict the effects of the different measures for small sub-networks and to combine these predictions afterwards. The test results illustrate the workings of the system, and show that the system can provide the operator with real-time predictions. It is shown that the method is applicable to generalise the data it uses.

Keywords: Decision Support Systems, Dynamic Traffic Management, Fuzzy Logic, Traffic Control, Traffic Prediction.

1. INTRODUCTION

Dynamic Traffic Management (DTM) is moving into a new era. Contemporary DTM focuses on the integrated (opposed to isolated) and coordinated (combination of different measures) deployment of measures, anticipating on future changes in traffic conditions. Controlling and guiding traffic flows are the core tasks of the Regional Traffic Management Centres (RTMC). In the RTMC, traffic operators decide when and which DTM measures are to be deployed in case of recurrent and non-recurrent conditions. These decisions are based on information from available measurement systems, pertaining to the current state of the network. It turns out that the operators require support in order to come up with optimal control decisions, mainly due to the fact that prediction the effect of deploying different combinations of control measures to different parts of the network is very important, but also a very difficult task, especially (but not only!) under non-recurrent conditions such as incidents and special events.

This is why the Traffic Research Centre is developing a Decision Support System (DSS) called BOSS. In broad terms, the objective of BOSS is to provide the operators with conditional predictions on the future state of the traffic network under their supervision, given the current state of the network and conditional on the candidate control scenarios. The operator can use these predictions to efficiently intervene at the network level by deciding which control scenario is to be used. At the present time, two systems are being developed:

(1) BOSS on-line (implemented in RTMC, uses real-time traffic data to provide decision support to the network operator), and;
(2) BOSS off-line (used by traffic engineers to prepare candidate control scenarios, used by BOSS on-line).

BOSS uses traffic predictions of 1 hour ahead. The system will become operational at the end of 2003 in the RTMC De Wijde Blik.

Given the complexity of the networks to be controlled as well as the large number of control scenarios that may be used, using traffic flow models in the on-line simulation of these control scenarios will not be feasible due to limited computational resources. This is why on behalf of the AVV Transport Research Centre, Ministry of Transport, Delft University of Technology has developed an alternative method for the on-line evaluation of these control scenarios. The method that has been developed is referred to as Fuzzy Multi-Agent Case-Base Reasoning (FMA-CBR). The system is based on generalising examples (so-called cases) that describe the effects of deploying control measures under specific recurrent or non-recurrent conditions. These cases may either be represented by real data or synthetic data (from simulation). The latter option was chosen here, due to the limited data availability at the time of writing. It is believed that the system will perform satisfactory with real-data as well. The resulting prediction method is called BSES (BOSS Scenario Evaluation System).

2. DECISION SUPPORT FOR DTM

In The Netherlands, planning and instalment of measures aiming to solve recurrent local and network problems are designed according to the guidelines provided in the Architecture for Traffic Management. The framework prescribes how policy related issues are translated into a so-called *frame-of-reference*, describing the desired traffic state (in terms of average speeds, queues, waiting times, etc.). The objective of *operational* DTM is to control the system towards this desired state, considering weighting factors indicating which parts of the network are to be prioritised, or if on certain parts of the network, congestion would be acceptable to the policy-makers.

2.1 Operator tasks

The RTMC operates in a multi-level control framework: at the lowest level, we have semi-autonomous local traffic controllers. An active traffic light or metering installation, operates using only local traffic conditions. At a higher level the operation of several local traffic controllers is coordinated and synchronised by the supervisory operators in the RTMCs.

The traffic operators in RTMCs have a variety of task, amongst which are

(1) Monitoring functioning subsystems / measures;
(2) Monitoring the state of the network, recognising irregularities and other problems, and diagnosing their causes;
(3) Setting up candidate solution scenarios to solve identified problems, choosing scenario, and implementing it in practise;
(4) Monitoring resulting developments in the system;
(5) Informing other actors.

Especially the state monitoring, prediction and control tasks are complicated. This is caused by among other things the following issues:

- Data interpretation problems caused by the large amount information received by the operator
- Lack of insight into the network dynamics, in particular under non-recurrent circumstances
- Diversity and complex interactions between the measures

As a result of these complications, expert knowledge and experience are often not sufficient to adequately determine the cause of the problem at hand, or to determine the most efficient control scenario. This is why decision support is needed.

2.2 Tasks of a Decision Support System

The tasks of a Decision Support System (DSS) are:

(1) *Identification*, consisting of *monitoring* (automatic collecting and summarising data from the monitoring system), and *diagnosis* (identification of the cause of the problem, given the data collected during monitoring).
(2) *Prediction*, i.e. conditionally forecasting the traffic conditions in the network, given prevailing traffic conditions, the predicted traffic demands, and the candidate control scenarios.
(3) *Providing advise*, i.e. presenting to the operator the control scenario that yields the optimal predicted traffic conditions, as determined by comparing the predicted situation with the frame of reference using weight factors.

The remainder of the paper focuses on the support of prediction and advise tasks.

3. STATE-OF-THE-ART

Several authors have described decision support systems for traffic management, such as FRED (Freeway Real-Time Expert System Demonstration) (Ritchie, 1990; Ritchie and Prosser, 1991; Zhang and Ritchie, 1994), or the Santa Monica Smart Corridor Demonstration Project (Karimi and Gupta, 1993; Roseman and Tvedten, 1997). Fuzzy decision support systems for traffic control have been developed in (Cuena *et al.*, 1995; Krause and von Altrock, 1997; Molina *et al.*, 1998). The TRYS system described in (Cuena *et al.*, 1995; Molina *et al.*, 1998) is an agent-based system for urban motorway control. The network is divided in overlapping regions and to each region an agent is assigned. These agents have to detect and diagnose traffic problems in their regions and subsequently suggest possible control measures to a higher-level coordinator, taking care of negotiations, and deciding which action will be taken. The decision process in the TRYS system is based on knowledge frames, and some of these frames use fuzzy logic. The paper (Krause and von Altrock, 1997) describes a fuzzy logic control architecture that can be applied in existing traffic control systems on a multi-lane motorway with VMS's. This system uses fuzzy logic to incorporate the experience of human traffic operators.

The main aim of the system presented here is to make the process of on-line, real-time evaluation and selection of the traffic management measures more efficient. To this end, fuzzy case-based interpolation was used to evaluate the effects of traffic control measures. In that way, a large set of possible traffic control measures for a given traffic situation can be rapidly evaluated, and the best control scenarios can then be simulated in more detail using microscopic or macroscopic traffic simulation.

3.1 Case-based reasoning

A common approach to decision support is so-called Case-Based Reasoning (CBR). Case-based reasoning is the process of solving new problems based on the solutions of similar past problems. The main characteristics of CBR are (Aamodt and Plaza, 1994):

- Actual knowledge describing what has happened in the past (domain knowledge) can be used directly.
- After implementing the control scenario, the resulting situation can be added to the case-base (*continuous step-wise learning*).

Case-based reasoning is a four-step process:

(1) Retrieve (retrieve cases from memory that are relevant to solving it)
(2) Reuse (map the solution from the previous case to the target problem, for instance using fuzzy reasoning)
(3) Revise (test solution and, if necessary, revise)
(4) Retain (learning, i.e. store resulting experience as a new case in memory)

CBR starts with a set of cases or training examples; it forms generalisations of these examples, albeit implicit ones, by identifying commonalities between a retrieved case and the target problem.

3.2 Case-base size and approach motivation

The advantages of using a case-based approach are clear. However, given the high-dimensionality of the prediction problem addressed here, setting up a case-base that has sufficient coverage is unfeasible. In illustration, the conditions in a network are typically described by the period of the day, densities of its links, traffic demands on the network boundaries, control measures that have been deployed, and the incident status. Using these to describe the conditions on a 25 link network would yield a case-base with a magnitude in the order of 10^{24}. Clearly, it is impossible to collect and story such a number of cases. Even if the number of cases can be reduced, by considering less links, or applying some other form of aggregations, the standard case-based approach will thus yield considerable problems with respect to the high number of cases (maintenance when network is changed, ability to upgrade to larger networks, etc.).

4. FMA-CBR APPROACH TO SCENARIO EVALUATION

For the problem at hand, a case (either simulated or measured on a real network) contains the following information: description of the situation, including both the state in the network at the initial time (average densities on a set of network links referred to as subsubnetworks), and the conditions at the boundary (inflows and outflow restrictions) of the network during the considered time period; the control scenario that was used during the period, and finally the result of applying the control scenario in terms of traffic conditions (average flows, densities, speeds, etc.) and criteria (travel times, fuel consumption).

To resolve the computational / memory issues described in the previous section, as well as to keep maintenance of the system possible, two aspects are introduced into the CBR-framework. For one, fuzzy logic is used to combine different cases in the case-base (F-CBR: Fuzzy Case-Based Reasoning). By doing so, a precise match between the current situation in the network and the example situations in the cases is not required. This approach has been successfully applied to small-scale networks (Hegyi *et al.*, 2001). For larger networks, the dimensionality of the vector describing the situation in the network yields too many combinations that need to be stored in the case-base. Secondly, the network to be controlled is divided in n partially independent subnetworks for which the aforementioned F-CBR approach can be applied. For each subnetwork j, a case-base is established. Except for the situation in the subnetworks itself (the state, described by prevailing and future densities), also the outflows to the other subnetworks are predicted using F-CBR. The n subnetworks are of course interdependent: traffic conditions in subnetwork j will be dependent on the outflows from subnetworks $j' \neq j$. In turn, the traffic conditions in subnetworks j' may be dependent on the outflows from subnetwork j. To attain consistency between the predicted traffic conditions and the subnetwork outflows, prediction of subnetwork conditions are iterated until a situation results in which all flows are consistent with each other.

In the remainder, we will describe the case-base for a single subnetwork j. Next, the fuzzy case-base reasoning approach predicting the subnetworks' traffic conditions will be described, followed by the iterative approach applied to assure consistency between the inter-subnetwork flows. Finally, we will discuss the approach to determine the network performance.

4.1 Specification of cases for a subnetwork

It was mentioned that for each of the subnetworks $j = 1, \ldots, n$ case-bases are determined. These case-bases contain specific situations that have occurred in the subnetwork, and describe the relation between the input of the subnetwork and the output of the subnetwork for these situations. These 'situations' are determined either from real-data or from simulations pertaining to the entire network.

Let T_{pred} denote the prediction horizon. A case for subnetwork j is described by the following input characteristics x:

- period of the day (morning, evening, off-peak)
- current state (i.e. average densities) on all subsubnetworks $K = 1, \ldots, K_j$ of subnetwork j at time t
- average external traffic demands (traffic flowing into network) and internal traffic demands (traffic flowing from other subnetworks j' to subnetwork j) during $[t, t+T_{\text{pred}})$
- average external supply restrictions (for traffic flowing out of network) and internal supply restrictions (for traffic flowing from subnetwork j to other subnetworks j') during $[t, t+T_{\text{pred}})$
- local measures deployed in current subnetwork j (ramp-metering, speed-limit control) and global measurements deployed in other network parts (route information) during $[t, t+T_{\text{pred}})$
- time-average incident conditions in subnetwork (location, duration, severity) during $[t, t+T_{\text{pred}})$

and output characteristics y:

- traffic conditions in subnetwork and in particular on boundaries (outflows, inflow restrictions) during $[t, t+T_{\text{pred}})$
- average performance expressed via Measures-of-Effectiveness (e.g. queue lengths, travel times, delay times, etc.) during $[t, t+T_{\text{pred}})$

The case base for subnetwork j consist of cases $i = 1, \ldots, m_j$ that link the input $x(j)$ for subnetwork j to the output $y(j)$ from subnetwork j. These cases can be written as rules R_i, which look like

> IF period $= \tilde{t}_i$ AND currstate $= \tilde{r}_i$ AND demand $= \tilde{d}_i$ AND supplyrestr $= \tilde{s}_i$ AND controlscenario $= \tilde{c}_i$ AND incidentscenario $= \tilde{i}_i$
> THEN outflowspred $= Q_i$ AND inflowrestrpred $= S_i$ AND performance $= P_i$

for $i = 1, \ldots, m_j$. We have used the tilde-notation to show that the elements in the antecedent part of the rule are in fact fuzzy numbers, while the elements of the consequent part are crisp. The fuzzy numbers (period, densities describing the current state, demands, etc.) are determined based on the (crisp) examples in the case-base. These are fuzzified using either bell-shaped or triangular membership functions, the centre of which lie at the crisp values that describe a case. In illustration, case (or rule) $i = 2$ may be represented by a time-averaged density of 30 veh/km/lane on subsubnetwork j_1 and of 20 veh/km/lane of subsubnetwork j_2, and the average external traffic demand of 3000 veh/h flowing into the subnetwork during a specific time period. These values will then represent the centres of the bell-shaped or triangular membership functions used to represent the current state and the demand. The width of the membership functions is chosen relative to the domain of the respective variable, and can be specified by the end-user.

4.2 Fuzzy Case-Base Reasoning to determining predictions for a subnetwork

To determine which cases correspond best to the current situation, and to combine different cases in order to provide a prediction of the traffic conditions and performance indicators for subnetwork j, fuzzy inference is used to describe the similarity between the current state (period, state, demand, supply restrictions) and the fuzzy antecedent part of case i, which can be reflected by rules as was shown in the previous section. For each subnetwork j, this entails:

- Determining similarity of current situation with each of the cases i in the case-base and
- Combining the consequences of the cases using these similarities to determine total prediction.

For each case i, the similarity (or degree of membership) with the current situation (current period, state, demand, etc.) in subnetwork j is determined by considering the mean membership of the elements of the antecedent part of the rule, i.e.

$$\mu_i(x) = \frac{1}{N_j}\sum_{l=1}^{N_j} \mu_{i,l}(x_l) \ , \tag{1}$$

where N_j denotes the number of elements in the antecedent part of rule i for subnetwork j. In other words, the mean fuzzy membership was used to quantify the fuzzy 'AND' operator used in the rule-representation shown in the previous section.

When $\mu_i(x)$ has been determined for all cases i, the prediction $y = (Q, S, P)$ for the conditions in subnetwork j is determined by taking the weighted sum of the consequent part of all rules i, using $\mu_i(x)$ as weights, i.e.

$$y = \frac{\sum_{i=1}^{m_j} \mu_i(x) y_i}{\sum_{i=1}^{m_j} \mu_i(x)} \ , \tag{2}$$

where $y_i = (Q_i, S_i, P_i)$ denotes the crisp consequent part of rule i (i.e. the output flows, inflow restrictions, and the performance); m_j denotes the number of cases in the case-base of subnetwork j. In turn, this operator describes the fuzzy 'OR' operator, used to aggregate consequences of the specific prediction rules.

The approach will yield a conditional prediction of the output (outflow and performance) of subnetwork j. The prediction is conditional, since part of the state x (and thus also the prediction) depends on the endogenous demands from other subnetworks j' as well as the supply restrictions limiting the flows to these subnetworks j'.

To determine the cases for subnetwork j, real life measurements can be used where traffic demands and traffic conditions are monitored using for instance inductive loops. The case-base used in the prototype application have been determined using simulation software. Using this software, network traffic conditions for the entire network (so not just the subnetwork) where determined using various prespecified input and control settings. For each subnetwork j, the traffic state was determined from the simulation results.

4.3 Iterative approach for finding consistent solution

For each subnetwork, the approach discussed in the previous section computes among other things the conditional outflow and inflow restrictions. These predictions are conditioned on the internal traffic demands and supply restrictions for the other subnetworks. In turn, these may depend on the outflow and inflow restrictions of the current network. In the end, the solution is sought in which the internal traffic demands and the supply restrictions are consistent. To solve this fixed-point problem, an iterative scheme was developed. Without formally deriving scheme stability criteria, it turns out that in practice the scheme converges within only a few iterations (less than 10).

4.4 Predicting network performance

Predicting the performance of the entire network is an easy task: the different indicator values determined for the subnetworks j are added or averaged. The overlap between different subnetworks is taken into account explicitly. The system will provide both results pertaining to the different subnetworks and the entire network.

5. EXAMPLE BSES APPLICATIONS

To test the concept of the system, an off-line prototype BSES was implemented. The prototype consists of the prediction model, and a simple Graphical User-Interface (GUI). The user of the system must first prepare a number of 'scenarios' or situations, which he or she aims to evaluate. A scenario is defined by the following:

(1) The current state in the network, generally determined by the monitoring system (e.g. inductive loops), consisting of the densities on the subsubnetworks of the network considered;

(2) The predicted network inflows (demands) and network outflow restrictions (i.e. the external boundary conditions), in general determined from historic traffic data;
(3) The control scenario (i.e. the settings of the different control / ITS measures available in the network, such as ramp-metering, speed homogenising control, shoulder lanes opening, lane closures, etc.);
(4) The incident conditions (duration of the incident, severity of the incident, location).

The GUI allows the user to study the evaluation results, to change the membership functions, and to show the network and subnetwork definition. Furthermore, the GUI warns the user when the predictions become unreliable because the examples in the case-base are not representative for the scenario the user wants to evaluate. If this occurs, the user is advised to extend the case-base with additional cases.

Let us remark that the verification results presented in the remainder are not intended to show the expected effects of incidents or of deploying DTM measures, but aims to show how the system is able to predict network conditions in line with the cases in the case-base; the predictions are at best as accurate as the off-line predictions in the case-base (which can be very accurate, when historic data is used!).

5.1 Verification results

To test the BSES, a case-base was set-up using simulation results of the macroscopic simulation model METANET (cf. (Kotsialos *et al.*, 1999)). The initial case-base consists of 1464 cases describing different situations (e.g. control scenarios, incident conditions, etc.) in the motorway network around the Dutch city of Amsterdam (see Fig. 1). The network is divided into 5 subnetworks, which are in turn split up into 3 or 4 subsubnetworks. The definition of subnetworks and subsubnetworks was done manually by identifying which links belong to which (sub)subnetwork. The subsubnetworks were defined such that they contained at most one major link or a major node. A major link may contain several on-ramps, off-ramps, lane-drops, etc. A major node connects two or more major links. It is clear that the way in which the network is divided into sub- and subsubnetworks has an influence on accuracy and reliability. Fine-grained divisions lead to more accurate results, at the expense of larger case-bases and computation time.

METANET computes a number of performance indicators, examples of which are shown in table 1. It is emphasised that the use of a different simulation model would naturally lead to a different set of and values for the performance indicators.

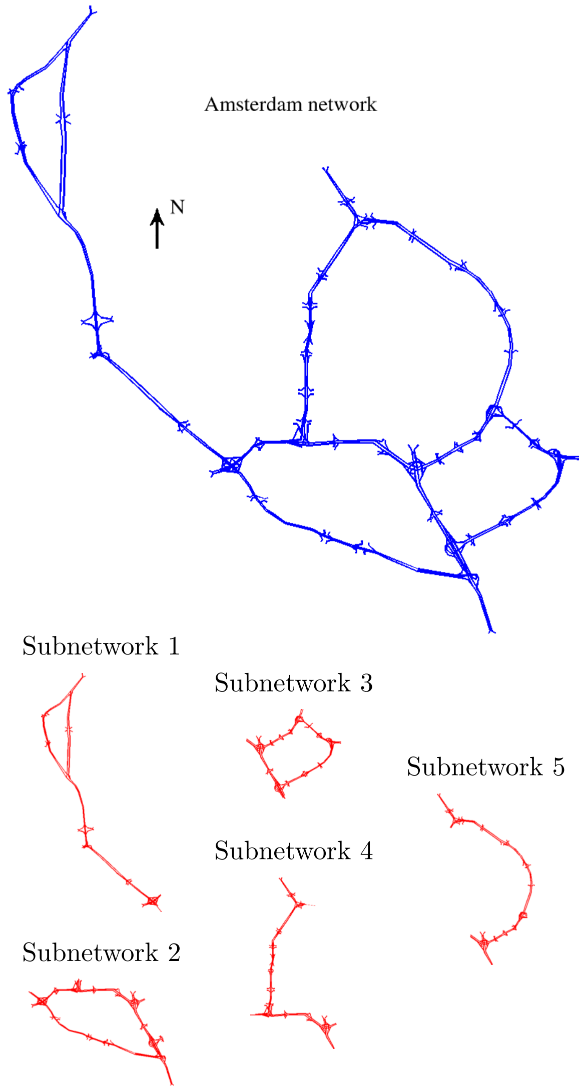

Fig. 1. Amsterdam network and division in subnetworks 1 to 5.

Table 1 shows the prediction results for regular circumstances, i.e. no incidents and no control measures. The table shows how subnetwork 2 is assigned a larger weight than the other subnetworks, indicating the high importance of that network part. Subnetwork 4 has a lesser importance, reflected by a smaller weight. In practical applications, the weights will stem from the frame of reference discussed in the introduction.

Another example is table 2, which shows the BSES prototype prediction results of applying ramp-metering on part of the Amsterdam network. In this particular case, the scenarios 11–14 represent different ramp-metering settings, i.e. which ramp-meters are operational in different parts of the network. In scenarios 11 and 12, all on-ramps to the main arterial of subnetwork 4 (see Fig. 1) in respectively the North-bound direction and the South-bound direction are metered; scenario 13

Table 1. BSES prediction results for subnetworks 1–5 and entire network (reference)

Results scenario 1	subnetwork					total
	1	2	3	4	5	*
Weight	*0.16*	*0.32*	*0.16*	*0.12*	*0.24*	*1.00*
Vehicle loss time	894	301	127	679	634	2108
Total time spent	3278	5215	2911	8811	8286	22847
Total travel time	2960	4983	2750	3492	3029	14837
Total queueing time	318	232	161	5319	5258	8011
Mean queue length	19.55	9.78	6.78	202.62	221.38	81.06
Mean link speed	90.95	88.89	92.70	85.11	90.47	89.80
Mean vehicle speed	78.26	93.39	94.86	80.06	78.77	86.80
Distance travelled ($\times 10^3$)	231.2	465.2	260.7	279.5	238.5	1475.2
Total inflow	22250	30411	19221	22809	17676	95552
Total outflow	20897	35642	19546	24624	14165	98768
Total vehicle number	13111	20861	11643	35242	14165	98768
# vehicles in net	11840	19931	10999	13967	12114	59346
# vehicles in queues	1271	930	644	21275	21031	32042
Total fuel consumption	20853	36302	20689	32394	29307	117858

Table 2. Overview of BSES predictions describing the effects of ramp-metering

Result scenarios	Scenario index				
Performance criteria	1	11	12	13	14
Period	8:00–9:00	8:00–9:00	8:00–9:00	8:00–9:00	8:00–9:00
Reliability	0.847	0.826	0.790	0.768	0.794
Vehicle loss time	-1805	-1806	-1686	-1687	-1799
Total time spent	18443	18463	18320	18340	18437
Total travel time	11308	11316	11196	11204	11303
Total queueing time	7134	7146	7123	7135	7133
Time stamp	0.25	0.25	0.25	0.25	0.25
Mean queue length	87.6	87.8	87.4	87.6	87.6
Mean link speed	89.8	89.8	90.1	96.1	89.8
Mean vehicle speed	85.5	85.5	86.5	86.5	85.5
Total distance travelled	966464	967171	967312	968018	966486
Total inflow	74691	74768	74722	74799	74692
Total outflow	76208	73194	76216	76202	76208
Total vehicle number	73772	73852	73281	73361	73749
# vehicles in net	45234	45266	44787	44819	45214
# vehicles in queues	28537	28585	28493	28541	28535
Total fuel consumption	93109	91389	91400	91479	91312

describes the case where all on-ramps of subnetwork 4 are metered. Scenario 14 describes the case where only 1 on-ramp in the Northbound direction is metered. The results shown in table 2 pertain to the entire network. From table 2, it can be observed that in general, deployment of ramp-metering has a beneficial effect on traffic conditions on the main-roads. Given the expected effects and the effects predicted by BSES, we conclude that the predictions are plausible.

5.2 Comparative analysis

It turns out that the predictions made by BSES are in line with the predictions of METANET. However, the time BSES needs to compute a prediction is much less than the time needed to do a METANET simulation (factor between 30 and 3000, depending on the mode of simulation of METANET), showing the potential for the system to be applied in an on-line system setting.

It is clear that the accuracy of the BSES predictions is directly determined by the accuracy of the underlying METANET model simulation, and that in fact verification only proves that the system is able to reproduce the predictions of the METANET model. However, the results obtained so far indicate that the system work equally satisfactory if the case-base is filled with either real-life data or with results of more accurate simulation models.

6. CONCLUSIONS AND FUTURE RESEARCH

This paper describes a new approach to the on-line prediction of the effect of control scenarios under a variety of circumstances in the network. It describes the developed approach, based on combining fuzzy logic, case-based reasoning, and multi-agent approaches. The main advantages are the speed of computation (compared to using traffic flow models), ability to use actual knowledge directly (rather than general knowledge or simulated data), and the ability to learn from previous experiences. It turns out that the system is able to very quickly produce predictions on the impact of different control scenarios to the traffic operations in the network, and can thus support operators in their decision tasks in a real-time decision environment. These predictions are in line with expectations of the effects of DTM and with the control simulations used to test the system. It can therefore be concluded that the system indeed functions properly.

Acknowledgement – The research described here was performed by the Delft University of Technology on behalf of the AVV Transport Research Centre, Ministry of Transport.

REFERENCES

Aamodt, A. and E. Plaza (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications* **7**(1), 39–59.

Cuena, J., J. Hernández and M. Molina (1995). Knowledge-based models for adaptive traffic management systems. *Transportation Research Part C* **3**(5), 311–337.

Hegyi, A., B. De Schutter, S. Hoogendoorn, R. Babuška, H. van Zuylen and H. Schuurman (2001). A fuzzy decision support system for traffic control centers. In: *Proceedings of the 2001 IEEE Intelligent Transportation Systems Conference (ITSC'01)*. Oakland, California. pp. 358–363.

Karimi, A. and A. Gupta (1993). Incident management system for Santa Monica Smart Corridor. In: *Proceedings of the 63rd Annual Meeting of the Institute of Transportation Engineers*. The Hague, The Netherlands. pp. 180–185.

Kotsialos, A., M. Papageorgiou, C. Diakaki, Y. Pavlis and F. Middelham (1999). Traffic flow modelling of large-scale motorway networks using the macroscopic modelling tool METANET. In: *Proceedings of the TRAIL Expert Seminar on Recent Advances in Traffic Flow Modelling and Control* (P.H.L. Bovy and S.P. Hoogendoorn, Eds.). number P99/2 In: TRAIL Conference Proceedings Series. Delft, The Netherlands. Chapter 6.

Krause, B. and C. von Altrock (1997). A complete fuzzy logic control approach for existing traffic control systems. In: *Mobility for Everyone, Proceedings of the 4th World Congress on Intelligent Transportation Systems*. Berlin, Germany. Paper no. 2045.

Molina, M., J. Hernández and J. Cuena (1998). A structure of problem-solving methods for real-time decision support in traffic control. *International Journal of Human-Computer Studies* **49**(4), 577–600.

Ritchie, S.G. (1990). A knowledge-based decision support architecture for advanced traffic management. *Transportation Research Part A* **24**(1), 27–37.

Ritchie, S.G. and N.A. Prosser (1991). Real-time expert system approach to freeway incident management. *Transportation Research Record* (1320), 7–16.

Roseman, D. and S.S. Tvedten (1997). Santa Monica Freeway Smart Corridor project operational multi-agency traffic management and expert system. In: *Mobility for Everyone, Proceedings of the 4th World Congress on Intelligent Transport Systems*. Berlin, Germany. Paper no. 1032.

Zhang, H. and S.G. Ritchie (1994). Real-time decision-support system for freeway management and control. *Journal of Computing in Civil Engineering* **8**(1), 35–51.