

This is a repository copy of *INTELLIGENT DECISION SUPPORT FOR TRAFFIC MANAGEMENT*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/89502/>

Version: Submitted Version

Conference or Workshop Item:

Krishnan, Rajesh, Hodge, Victoria orcid.org/0000-0002-2469-0224, Austin, Jim orcid.org/0000-0001-5762-8614 et al. (4 more authors) (2010) *INTELLIGENT DECISION SUPPORT FOR TRAFFIC MANAGEMENT*. In: *Proceedings of 17th ITS World Congress*, 25-29 Oct 2010.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

INTELLIGENT DECISION SUPPORT FOR TRAFFIC MANAGEMENT

Rajesh Krishnan

Research Associate, Imperial College London
Centre for Transport Studies, Department of Civil and Environmental Engineering, Imperial
College London, Exhibition Road, London SW7 2AZ, United Kingdom
TEL: +44 (0) 20 7594 6111, FAX: +44 (0) 20 7594 6102, E-mail: rajesh.k@imperial.ac.uk

Victoria Hodge

Research Associate, University of York
Department of Computer Science, University of York, York YO10 5DD, United Kingdom
TEL: +44 (0) 1904 567 739, FAX: +44 (0) 1904 432 767, E-mail: vicky@cs.york.ac.uk

Jim Austin

Professor, University of York
Department of Computer Science, University of York, York YO10 5DD, United Kingdom
TEL: +44 (0) 1904 567 681, FAX: +44 (0) 1904 432 767, E-mail: austin@cs.york.ac.uk

John Polak

Professor, Imperial College London
Centre for Transport Studies, Department of Civil and Environmental Engineering, Imperial
College London, Exhibition Road, London SW7 2AZ, United Kingdom
TEL: +44 (0) 20 7594 6089, FAX: +44 (0) 20 7594 6102, E-mail: j.polak@imperial.ac.uk

Tom Jackson

Project Co-ordinator, University of York
Department of Computer Science, University of York, York YO10 5DD, United Kingdom
TEL: +44 (0) 1904 433 382, FAX: +44 (0) 1904 432 767, E-mail: tom.jackson@cs.york.ac.uk

Mike Smith

Professor Emeritus, University of York
Department of Mathematics, University of York, York YO10 5DD, United Kingdom
TEL: +44 (0) 1904 628 275, FAX: +44 (0) 1904 432 767, E-mail: mjs7@york.ac.uk

Tzu-Chang Lee

Research Associate, Imperial College London
Centre for Transport Studies, Department of Civil and Environmental Engineering, Imperial
College London, Exhibition Road, London SW7 2AZ, United Kingdom
TEL: +44 (0) 20 7594 6022, FAX: +44 (0) 20 7594 6102, E-mail: tlee@imperial.ac.uk

ABSTRACT

Urban traffic control systems such as the widely deployed SCOOT system, incrementally respond to changing traffic conditions. Such systems are often complemented by traffic

control centres where road network managers intervene manually to mitigate rapidly developing congestion events. An Intelligent Decision Support (IDS) system developed by the authors within the UK FREEFLOW (FF) project to aid the network managers is presented in this paper. The primary objective of the FF IDS system is to identify traffic congestion in near-real-time and to recommend appropriate traffic control intervention measures. The FF-IDS consists of multiple internal components. A state estimation component monitors live traffic sensor data and determines if there is a congestion problem on the road network. If a problem is identified, a binary neural pattern-matching component is used to identify past time periods with similar congestion events. This is able to rapidly search large historic traffic datasets finding sets of traffic control interventions carried out during similar historical time periods. The effectiveness of each intervention is evaluated using a Performance Index (PI) and the intervention that resulted in the highest improvement in PI is recommended to network managers. The FF-IDS system can also present traffic incidents and equipment faults that occurred during these historical time periods to the network manager as potential causes of the problem. This paper describes the FF-IDS system in detail. The system is currently under development. An early version of the FF-IDS system was trialled using off-line data from London, yielding encouraging preliminary results.

INTRODUCTION

Adaptive traffic control systems such as SCOOT (1) have been in existence for the past three decades. Such systems dynamically optimise the performance of the road network by changing traffic signal times in response to changing traffic conditions measured using traffic sensors. These systems are adaptive in nature and make gradual changes to signal settings in response to changing traffic demands. Small incremental changes provide control stability against stochastic fluctuations in traffic demand. Moreover, the small changes avoid the extra delay caused due to sudden changes in signal timings. However, the incremental approach is not sufficient when responding to rapidly developing congestion due to incidents. Automated approaches have been developed in the past to improve SCOOT operations such as changing the signal time to the desired value directly or allowing SCOOT to respond more rapidly (2). Moreover, the incremental nature of optimisation (3) may lead the system to a locally optimal solution than a globally optimal solution.

Manual traffic control interventions aim to overcome the limitations of adaptive traffic control systems. They are typically made based on the experience of individual traffic controllers. A drawback of the manual system is that generally the effectiveness of the interventions are not measured quantitatively, and thus they are subjective. Rule-based interventions are devised using the past experience of traffic controllers and applied to recurring traffic problems. For example, UTMC (4) compliant systems in the UK provide functionality to change signal

plans based on traffic sensor readings. However, the rule-based interventions are not adaptive and are only triggered if the traffic variables satisfy the strict criteria defined in the rules.

This paper provides an overview of an intelligent decision support (IDS) that has been developed as part of the FREEFLOW (FF), a large UK based project involving collaboration between academic, industrial and public sector partners (5). The objective of the FF-IDS system is to automatically detect traffic conditions when interventions are necessary and recommend appropriate intervention actions based on past experience. The problem identification within FF-IDS is carried out using state-estimation models, which are more flexible than rule-based triggers. Candidate interventions are selected using pattern match and the best is recommended based on the past performance of the candidate interventions, quantitatively measured using a Performance Index (PI). The details of the IDS system are presented in this paper. The overall system is described in the next section, and individual components are described in subsequent sections.

FREEFLOW IDS OVERVIEW

Several authors have described decision support systems for traffic management. Ritchie (6) uses an integrated set of expert systems to process real-time data. The system has a similar high-level architecture but our system aims to be more automated. Zhang and Ritchie (7) proposed an advancement of their system to integrate human knowledge with automatic operations. Hernandez et al. (8) use multiple agents and a coordinator module to perform conflict resolution between agents to optimize traffic management. The system described by Hegyi et al. (9) uses fuzzy logic based traffic control to manage congestion. This system has since been extended by de Schutter et al. (10) who combine case-based reasoning and fuzzy logic to develop a multi-agent evaluation system that can be used by traffic operators to analyze the expected performance of several potential interventions. The system, described in this paper, has many similarities with case-based reasoning: using historical interventions, intervention similarity measures and intervention performance to propose the most effective intervention for the current situation.

The overall design of FF-IDS is shown in Figure [1]. The FF-IDS is designed to monitor traffic sensor data at regular intervals, determine if there are traffic problems and recommend interventions if problems are identified. The data access component of FF-IDS is designed to monitor the traffic sensor data and filter out erroneous readings. The state-estimation component uses the filtered data and determines if there is a traffic problem. The FF-IDS also provides functionality to specify rule-based interventions. A rules-engine within FF-IDS will combine the outputs from the rule-set and the state-estimation component to determine the method of intervention identification. In case of the rule-set is triggered, the intervention is

identified from a look-up table. On the other hand, a problem identified by the state-estimator will invoke the pattern-matcher component. The pattern-matcher component, based on the AURA technology (11), uses an optimally stored historic traffic sensor dataset to rapidly identify past time periods when the intensity and spatial distribution of congestion are similar to the currently observed traffic situation. The effectiveness of interventions carried out during time periods with similar congestion pattern are evaluated using a Performance Indicator (PI). The intervention that was able to improve the PI the most is recommended by FF-IDS. In addition to intervention recommendation, the FF-IDS will also identify incidents and equipment faults during time periods with similar congestion. This information will be presented to road network managers as potential causes of the current traffic congestion. The FF-IDS is typically configured to run every 5 or 15 minutes. The next sections will describe the main components of FF-IDS in more detail.

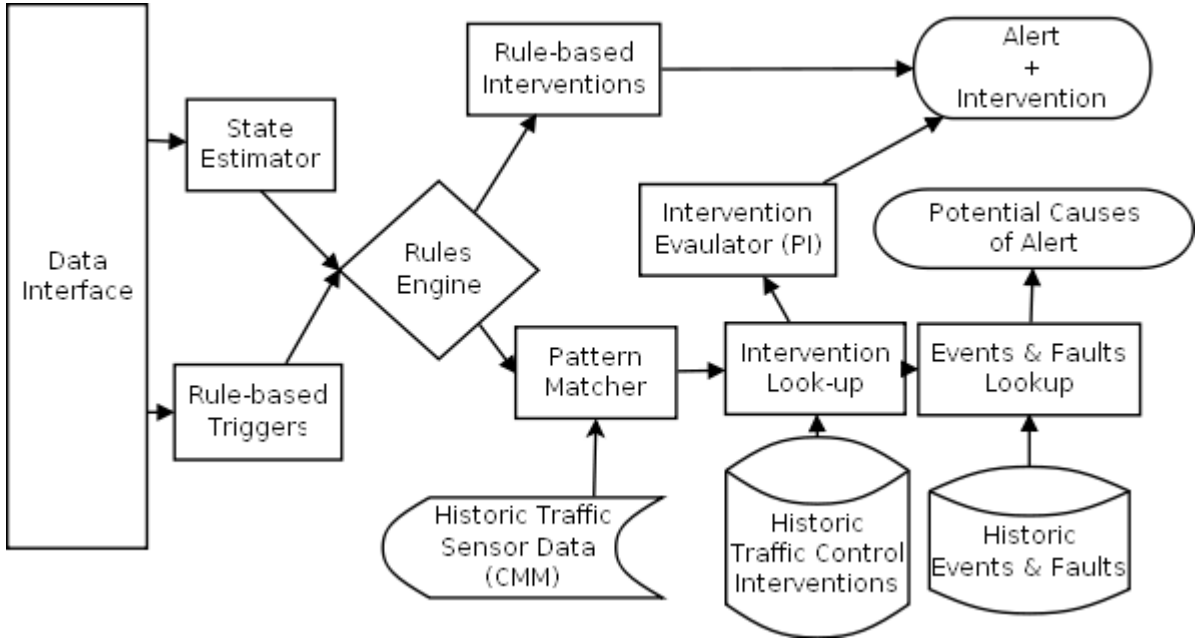


Figure 1: FF-IDS Overview

DATA INTERFACE

The objective of the Data Interface (DI) component is to obtain traffic sensor data in near-real-time. The FF-IDS system will be initially deployed in the UK during the development phase. Hence, the DI component is designed to obtain data from UTMC compliant systems, as UTMC is the open standard for ITS systems in the UK. The DI can process Inductive Loop Detector (ILD) data consisting of flow and occupancy values (12), travel time data from Automatic Number Plate Recognition (ANPR) camera links, bus journey-time data, bus schedule adherence data, signal plan data, data about road-works and incidents, VMS message data as well as other traffic variables such as degree of saturation. The DI component also currently obtains data from non-UTMC compliant systems such as

the Datex-II feed from the UK Highways Agency (13), SIRI incident feeds and the tpegML incident feed from the BBC (14). In addition to data gathering, the DI component filters out erroneous data due to detector faults.

STATE ESTIMATOR

The objective of the State Estimator (SE) component is to determine if there is a traffic problem using flow and occupancy data from the ILD. A simple rule-based error filter ignores data due to a number of known ILD faults (15). The FF-IDS is designed to work with partial data when some ILDs are deemed faulty. During uncongested conditions, occupancy increases linearly with flow. The rate of increase depends on the average speed of vehicles over the ILD and the sensitivity of the ILD, and hence the rate varies between ILDs. On the other hand, once a *critical* occupancy is reached, the flow decreases with further increase in occupancy as the traffic enters the saturated regime from free-flowing conditions. The SE method categorises a given {flow, occupancy} reading into either the congested or the uncongested state (16).




Figure 2: Typical flow-occupancy data from an ILD

The SE method is based on the principle of clustering. The k-means clustering method was initially tried to historic {flow, occupancy} data from ILDs. Typical output from an ILD on a given day will contain a large number of points from the uncongested regime and a smaller number of points in the congested regime, as shown in Figure [2]. Moreover, the range of occupancies during congestion can be quite large. Hence, a direct application of the k-means clustering lead to a number of congested data points being identified as uncongested. Use of other partitional clustering techniques, such as the fuzzy c-means clustering, could also lead to a similar problem. This is because such methods try to group data points such that an overall criteria is minimised for the clustered dataset, which does not guarantee partition in line with traffic states. But, it gives a starting point; hence a two-step clustering method was used.

This first step is to cluster the data points into two clusters roughly representing congested and uncongested regimes using k-means clustering, where $k=2$. The distance metric used is cosine, which uses the difference between the angles made by two different data points with the origin to determine cluster memberships. The use of the cosine distance metric takes advantage of the fact that the flow vs. occupancy curve is linear in the congested regime, and most of the uncongested data points should be grouped in the same cluster. However, due to the range of occupancy values in the congested regime, some of the congested data points may be classified into the first cluster of uncongested data points. The second step is to fit a linear regression model on the data points in the uncongested cluster identified in the first step. All the data points identified as outliers by the regression model are moved to the second cluster, representing the congested state. Hence, all points are now labelled either congested or uncongested. More details about this approach can be found elsewhere (16).

The method was tested using data from six ILDs from the English Highway network, where flow and occupancy are available at 5-minute intervals through the Datex-II feed from the National Traffic Control Centre (NTCC). The actual state (congested vs. uncongested) of the ILD is determined manually, and is used as the reference against which the accuracy of the SE method is determined. Incorrect state identification is categorised into false positives and false negatives. False positive is a situation where the SE method outputs congested state when the ILD is not congested. Similarly, false negative is when the SE method outputs uncongested state but the ILD is actually congested. Table [1] summarises the results of the SE method.

ILD	No. of observations	No. of congested Links	False Positives	False Negatives
1	1424	7	2	0
2	1898	2	0	0
3	1753	1	0	0
4	931	5	0	0
5	1587	42	0	11
6	1441	21	0	4
All	9034	99	2	15

Table 1: Accuracy of the State Estimation model

RULES ENGINE

In addition to the SE component, specific rules can be put in place to define known patterns of traffic congestion and trigger interventions. For example, an intervention can be set to be triggered if the occupancy at an ILD is greater than 30%. Alternatively, the pattern-matcher can be invoked if the occupancy on a critical link exceeds a given threshold. The objective of

the Rule Engine (RE) is to determine the appropriate course of action when the SE and one or more rules are triggered. The RE will either recommend a pre-defined intervention defined by a rule, or it will invoke the pattern-matcher that will determine the most appropriate intervention action. The logic by which the pattern-matcher will determine the most appropriate intervention is described in the following sections.

PATTERN MATCHER

The pattern-matcher is effectively a k-Nearest Neighbour (k-NN) tool that finds past time periods with similar traffic patterns. K-NN is known to be a robust and flexible method that allows the pattern matcher to be updated continuously. One drawback of conventional k-NN is the speed of the method, which typically becomes very slow for large problems. However, this is overcome using the Advanced Uncertain Reasoning Architecture (AURA) (11). AURA is a family of methods and applications based on binary neural networks aimed at high speed search and match operations in large data sets. AURA is fast, scalable and compact. The foundation of AURA is the Correlation Matrix Memory (CMM). A CMM is a binary matrix used to store and retrieve vectors. In the AURA k-NN (17), each column of the CMM is a tuple or record (set of variable observations) and each row indexes a variable value, as shown in Figure [3]. Sets of rows represent individual variable observations and cover the range of values for the variable. The set of all rows represents all values for all variables.

The CMM is trained by learning associations between an input vector I_n which indexes matrix rows and an output vector O_n which indexes the matrix columns. Thus, an input vector representing the variable values for a record is associated with an output vector uniquely identifying that record for all N records in the data set. The CMMs require binary input vectors for training; so numeric variable values must be quantised (binned). Each variable is quantised over its range of values to a set of bins where each bin indexes a specific and unique row in the CMM. For example, for an integer-valued variable such as vehicle count per 5 minutes with range 0-9 and 5 bins then each bin would have width 2: bin 0 {0,1}, bin 1 {2,3} ... bin 4 {8,9}. For a real-valued variable such as speed with range 0.0-9.999 and 5 bins then each bin would have width 2: bin 0 [0,1.999], bin 1 [2.0,3.999] ... bin 4 [8.0,9.999].

Multiple attributes are mapped onto the input vector by pre-allocating a specific and unique section of the input vector to each variable. Thus each variable maps to a specific and unique section of the input vector and each variable value maps to a specific and unique subsection of the vector (within its corresponding variable's section).

For AURA k-NN, the input vector X_q is represented by a set of parabolic kernels with one kernel ($Kernel_x$) for each variable x in X for each bin bin_{xj} as shown in Figure [3].

$$Kernel(x_n) = \left[\left[\left(\frac{\max(b)}{2} \right) \right]^2 - \left(|bin(x_n^q) - bin(x_n^h)| \right)^2 \alpha(x_n) \right] \text{ where } \alpha(x_n) = \frac{(\max(b))^2}{(b(x_n))^2} \quad (1)$$

Where, $\max(b)$ is the maximum number of bins across all variables, $|bin(x_n^q) - bin(x_n^h)|$ is the number of bins separating the bin mapped to by the variable value for the query vector (x_n^q) from the bin mapped to by the variable value for the stored historical vector (x_n^h), and $b(x_n)$ is the total number of bins for variable x_n .




Figure 3: Figure showing the application of kernels to the CMM to score records and identify the k-nearest neighbours

Each kernel is centred on the bin representing the variable value for the query record. $X'_q = X_q \oplus (Kernel_x + offset)$, for all variables x in X , where $offset$ indexes variable x 's section of vector X . The values in X_q , multiply the rows of the matrix as shown in Figure [3], where the first column represents the kernel values corresponding to the query record. If the bit is set to one in a particular column, then the column will receive the kernel score for the corresponding row as given in Equations [1] and [2]. The process is illustrated in Figure [3]. For *Loop_I_Flow*, columns 1 and 2 (indexing from 0 on the left) receive a score of 9 as the set bit in the respective columns aligns with the score of 9 in the input kernel. In contrast, the right column receives a score of 5 as the set bit in the right column aligns with the score of 5 in the input kernel.

$$S^T = \sum X_q \bullet CMM \quad (2)$$

The AURA k-NN can function as a standard k-NN or in a two-step mode. In the latter mode, AURA k-NN can be used to preselect a set of candidate matches which are then ranked by an alternative ranking measure. AURA k-NN can perform up to four times faster than the standard k-NN in this mode (17), and provides a huge potential for efficiently implementing for non-parametric traffic models and ITS applications that operate against large traffic data archives. However, AURA k-NN introduces a degree of approximation due to quantisation of the variable values. Hence, the speed of AURA k-NN is exploited to preselect candidates and a finer-grained and computationally intensive measure is then used to rank the records in the candidate set.

Within FF-IDS, the feature vector currently consists of flow and occupancy values from all the FF-ILDs in the target road sub-network. The CMM stores historic flow and occupancy values from the same set of ILDs. When current flow and occupancy values from the ILDs are presented to the pattern-matcher, the top k time periods with closest traffic patterns are returned. FF-IDS bases its further processing on the k time periods identified. More details about the pattern-matcher component can be found in (18).

Tests of this system have been performed using data from Hyde Park Corner (HPC) in London, UK for 01/04/08 to 31/03/09 (19). There are 32 ILDs in the HPC area. Data are aggregated at 15-minute intervals and linked to traffic control interventions. The objective is to determine if the AURA k-NN can identify time periods with similar traffic patterns in the historic data. Hence, five serious or severe congestion events in HPC area were identified to validate the pattern matcher: congestion in the northern/western arms of HPC, congestion in all arms of HPC, an equipment fault, a spillage and a broken down vehicle (19). AURA matches historical time periods ranging from when congestion was developing until when congestion had dissipated. False positives (FP) indicate sensors that are congested during the historical time period but are not congested during the current time period. False negatives (FN) indicate sensors that are congested during the current time period but not during the historical time period. A data vector comprising flow and occupancy values for the 32 ILDs for each of the 5 incidents in turn was input to the AURA k-NN and the top 5 matches per incident were retrieved. The data labels (congested / uncongested) for all 32 ILDs from the input data were then compared against the labels for the corresponding ILDs of each of the top 5 matches in turn which gives 800 comparisons of ILD labels in total ($5*5*32$). The AURA k-NN produced 18 FP and 25 FN across the 800 comparisons. Note that this analysis is to some extent data dependent. There may not be any perfect matches in the historical data for a novel situation; so some FPs and FNs are inevitable.

INTERVENTION, EVENTS AND FAULTS LOOKUP

INTERVENTION RECOMMENDATION

The FF-IDS requires a historic database of all traffic control interventions corresponding to the historic traffic sensor dataset. This information can be typically obtained from the UTC or UTMC systems. For example, all traffic control interventions are logged by the UTC system in London. The log file is parsed and is stored in a searchable format in a relational database for FF-IDS. FF-IDS identifies all traffic control interventions made during the k time periods identified by the pattern-matcher in the target sub-network.

These interventions are quantitatively ranked using a PI. Any PI metric that is in line with a local authority's policy objectives can be used in FF-IDS. However, historic PI data should be available, or the FF-IDS should be able to calculate the PI using available sensor and associated traffic control data. The PI is calculated (or obtained) n time periods after each intervention was applied. The PI is compared against the daily PI profile for the day of the week, and percentage improvement in the PI due to the intervention is calculated. The intervention that resulted in the highest percentage of PI increase is recommended by FF-IDS.

POTENTIAL CAUSES OF TRAFFIC PROBLEMS

The pattern-matcher component, in combination with the PI evaluator, recommends a traffic control intervention that maximises the performance of the road network in spite of the event. However, the traffic managers need to remove the root cause of the traffic congestion in order to improve network performance. Typical supply-side factors that cause road congestion are incidents such as broken down or illegally parked vehicles, accidents, road works or equipment faults such as traffic light outage or sensor defects. The FF-IDS provides hints about the root cause of the problem using historic data.

Historic incident data are generally available from UTMC systems at local authorities for specific areas in the UK. In addition, the BBC (14) and the NTCC disseminates incident information in XML format over the internet for the whole of UK. Similarly, equipment fault data are archived by many local authorities in the UK. FF-IDS provides information about incidents and faults during the k time periods identified by the pattern-matcher relevant to the sub-network as a hint for network managers about probable causes of the problem. It is hoped that this information will help the network managers to identify the actual root cause of the problem in an expedient fashion.

The output of the FF-IDS thus consists of recommended intervention actions along with events and faults that caused similar incidents in the past. The network operators will use this information to confirm the cause and nature of the incident, usually using CCTV, and decide whether to implement the recommended intervention.

CONCLUSION AND FUTURE WORK

An Intelligent Decision Support (IDS) platform is being built to automatically identify problems in the road network and recommend traffic control intervention actions. The FF-IDS system is designed to run at regular intervals, e.g. every 5 minutes. When invoked, FF-IDS will monitor current traffic data, such as sensor data, and determine if there is a problem in the road network. If a problem is identified, using pre-defined rules, using the state estimation component or both, FF-IDS will recommend a traffic control intervention to alleviate the problem based on the past success of intervention actions during similar congestion events. In addition, FF-IDS will provide information about incidents and equipment faults during similar congestion events to help network managers identify the cause of congestion.

An initial version of IDS was tested offline against historic data in London. Plans are in place to test the system against live data in London, Kent and York in the UK in 2010. The authors plan to report on the accuracy and performance of the IDS system at a later date after live trials. The future work will also include extending IDS to other data sources such as journey times, bus schedule adherence and fusing information from multiple sensor sources.

ACKNOWLEDGEMENTS

The work reported in this paper forms part of the FREEFLOW project (18), which is supported by the UK Engineering and Physical Sciences Research Council, the UK Department for Transport and the UK Technology Strategy Board. The project consortium consists of partners including QinetiQ, Mindsheet, ACIS, Kizoom, Trakm8, City of York Council, Kent County Council and Transport for London.

REFERENCES

- (1) D. Robertson and R. Bretherton, "Optimizing networks of traffic signals in real time-the SCOOT method," *IEEE Transactions on Vehicular Technology*, 40,(1): 11-15, 1991.
- (2) D. Bretherton, K. Wood, K. Baker, and B. Radia, "Congestion and incident management using the SCOOT UTC system", in Tenth International Conference on Road Transport Information and Control, 2000. (Conf. Publ. No. 472), pp. 96-100, 2000.
- (3) P. B. Hunt, D. I. Robertson, R. D. Bretherton, and R. I. Winton, "SCOOT - a traffic responsive method of coordinating Signals", TRRL Laboratory Report 1014, Transport and Road Research Laboratory, Crowthorne, Berkshire, UK, 1981.
- (4) Department for Transport, "Urban Traffic Management & Control Web Pages", <http://www.utmc.uk.com>, Accessed on 30th July 2010.
- (5) P. Glover, A. Rooke, and A. Graham, "Flow diagram", *Thinking Highways*, 3(3): 20-23, 2008
- (6) S.G.. Ritchie,. "A knowledge-based decision support architecture for advanced traffic

management”, *Transportation Research Part A: General*, 24(1): 27-37, 1990.

(7) H. Zhang, and S.G. Ritchie, “Real-Time Decision-Support System for Freeway Management and Control”, *Journal of Computing in Civil Engineering*, vol. 8, pp. 35-51, 1994.

(8) J. Hernandez, J. Cuena, and M. Molina,. “Real-time Traffic Management Through Knowledge-based Models: The TRYS Approach”. In, Tutorial on Intelligent Traffic Management Models (www.erudit.de), Helsinki 1999, Finland.

(9) A. Hegyi, et al. “A fuzzy decision support system for traffic control centers”, *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*, pp. 358-363.

(10) B. De Schutter, et al. “A multi-agent case-based traffic control scenario evaluation system”, *Intelligent Transportation Systems, 2003. Proceedings. 2003 IEEE*, vol. 1, pp. 678-683.

(11) Advanced Computer Architectures Group, "The AURA Web Pages" <http://www.cs.york.ac.uk/arch/neural-networks/technologies/aura>, Accessed on 26th January 2010.

(12) T. Cherrett, H. Bell, and M. McDonald, "Traffic management parameters from single inductive loop detectors". *Transportation Research Record: Journal of the Transportation Research Board*, 1719: 112-120, 2000.

(13) Highways Agency, "National Traffic Control Centre", Tech. Rep., The Highways Agency, UK., Available at <http://www.highways.gov.uk/knowledge/1298.aspx>, 2009. Accessed on 15th November 2009.

(14) J. Hu, R. Krishnan, and M. G. H. Bell, "TPEG feed from the BBC: A potential source of ITS data?", in Proceedings of the 13th International Conference on Road Transport Information and Control, Manchester, UK, 20-22 May 2008.

(15) S. Robinson and J. W. Polak, "ILD data cleaning treatments and their effect on the performance of Urban Link Travel Time models", in Proceedings of the 85th Annual Meeting of the Transportation Research Board, CD-ROM, 2006.

(16) J. Han, R. Krishnan, and J. Polak, "Traffic state identification using loop detector data", in International Conference on Models and Technologies for Intelligent Transportation Systems, Sapienza University of Rome, Italy, 22-23 June 2009.

(17) V. J. Hodge and J. Austin, "A Binary Neural k-Nearest Neighbour Technique", *Knowledge and Information Systems*, 8(3): 276-292, 2005.

(18) R. Krishnan, V. Hodge, J. Austin, J. Polak, and T.-C. Lee, "On Identifying Spatial Traffic Patterns using Advanced Pattern Matching Techniques", in Proceedings of the 89th Annual Meeting of the Transportation Research Board, CD-ROM, 2010.

(19) R. Krishnan, V. Hodge, J. Austin and J. Polak. “A Computationally Efficient Method for Online Identification of Traffic Control Intervention Measures”. Presented at, 42nd Annual UTSG Conference, Centre for Sustainable Transport, University of Plymouth, UK: January 5-7, 2010.