

A Distributed Traffic Monitoring and Information System

Iisakki Kosonen

*Helsinki University of Technology, Transportation Engineering, P.O.Box 2100, FIN-02015 HUT, Finland
Iisakki.Kosonen@Hut.Fi*

Andrzej Bargiela

*The Nottingham Trent University, Real-time Telemetry Systems, Butron Street, Nottingham, NG1
4BU, United Kingdom
andre@doc.ntu.ac.uk*

Contents

- 1. Introduction***
- 2. Real-time simulation***
- 3. Using the real-time data***
- 4. Technical framework***
 - 4.1 Simulation of Mansfield test area***
 - 4.2 Developing Internet and GIS interface***
- 5. Conclusion***
- References***

ABSTRACT The paper summarizes experiences in the development of a real-time simulator which underlies the implementation of novel traffic information systems. The research project was carried out in collaboration between the Helsinki University of Technology (HUT)/ Laboratory of Transportation Engineering and the Nottingham Trent University (NTU)/ Department of Computing during a one year period from September 1997. The HUTSIM microscopic simulation model was interfaced to the SCOOT-urban traffic control system providing real-time traffic data through the Distributed Shared Memory system (DIME) developed at NTU. The emphasis of this paper is on presenting the principles of interfacing micro-simulation model to real-time field measurements, demonstrating a prototype implementation and outlining the range of information services that can be provided on the basis of such a system.

KEYWORDS:Real-time simulation, distributed processing, information services.

1. Introduction

There is a growing need for the improvement of the efficiency of urban traffic in order to ensure the sustainability of modern cities. It is now recognized that this objective requires not only the improvement of traffic monitoring and management schemes in traffic control centres but also the provision of information services for ordinary road users. The former measure has been widely adopted by many urban traffic control centres during the last decade (deployment of traffic control systems such as SCOOT, SCATS, CARS etc.) and the latter is currently the subject of intensive research and development.

The prerequisite for the development of any traffic telematics application is the availability of real time traffic data. The system reported here makes use of the telemetry data underlying the operation of the SCOOT-urban traffic control system. The SCOOT is an adaptive system optimising the split, cycle and offset times of traffic signals. The optimisation is based on maintaining mezzoscopic models of queues and traffic flows and ballancing saturation flows on all approaches to the controlled intersestions. The system provides both the unprocessed on-line detector and signal data and the various higher order measures derived through its traffic model, thus making it an ideal basis for the development of traffic information systems.

The potential for the amelioration of urban traffic through the combined supervisory control and the road-traffic information, has created the need for a flexible computing environment in which various new applications can be fully integrated with the existing traffic control systems without adversely affecting their performance. The DIME (Distributed Memory Environment) system (Argile *et al.* 1996), allows various distributed traffic telematics applications to communicate with each other through the LAN or WAN networks while maintaining the shared memory logical view of data. The communication harness is based on TCP/IP-protocol and client/server architecture and it is independent of the physical network type. The performance of this system has been verified in extensive on-line tests with the SCOOT system at the Nottingham Traffic Control Centre.

HUTSIM is an object-oriented microscopic urban traffic simulation model developed at Helsinki University of Technology during the nineties (Kosonen 1996). The connectivity to real control systems has been the principal aspect from the very beginning of HUTSIM. The simulation model comprises a flexible and interactive object-oriented framework with a detailed rule-based vehicle dynamics, which has been calibrated based on field measurements. The first Windows version of HUTSIM was developed in 1997. As a result of this collaborative project, HUTSIM has now been equipped with an interface to the DIME system opening, as a result, a path for the development of integrated simulation, monitoring and supervisory control applications based on real-time data.

2. Real-time simulation

The main idea behind the real-time simulation is that of the use of real-time traffic measurements as input data to microscopic simulations. In such a context, various non-measured traffic parameters, that can be deduced from micro-simulations, are deemed to be a good approximation of the reality by virtue of being based on the actual measurements. Although the prototype system presented in this paper makes use of the specific microsimulator (HUTSIM) and the urban traffic control system (SCOOT), the approach is general and is applicable to any traffic control system that supplies the required real-time data.

The basic data obtained through the traffic telemetry systems is the lane-occupancy data from detectors embedded in the road surface. Each detector provides information about the presence or absence of a vehicle in a discrete location. The rest of the information about traffic situation must be derived from the general knowledge of system layout, statistics of the traffic patterns and the estimate of vehicles' dynamics. In a simulation model all these factors are methodically combined. The simulation model provides also an engine for creating hypothetical traffic situations and for deriving higher order measures to be used by traffic information services.

In the absence of the actual traffic data from the detectors, simulation model generates vehicles on a statistical basis. Ideally this should be accurate enough to produce reliable average measures i.e. in off-line simulation mode. However, in real-time operation average measures are not always relevant, so the micro-simulation model is made realistic by replacing the time headway distribution with real-time arrivals. Since the simulation model is to mirror the operation of the actual traffic control system it requires also the real-time signal status data.

The real-time simulation approach postulated here extends significantly the monitoring capabilities of telemetry systems by extrapolating the traffic occurrences in discrete locations through to the simulation of realistic traffic flows in the whole of the network.

3. Using the real-time data

The accuracy of real-time simulations clearly depends on the accuracy and the availability of the relevant detector data. Within the limits of the accuracy of the traffic model, the presence of discrepancies between the real and the simulated traffic can be seen as an indication of the potential for the improvement of the telemetry system itself. If the measurement data is inaccurate or incomplete the simulation will highlight this fact by demonstrating the cumulative effect of these errors.

In a real-time microscopic simulation model, individual vehicles are generated according to the lane-occupancy detector data. Vehicle arrivals are recognized from the edges of detector signal i.e. the changes of signal status from passive to active or vice versa. Usually a single vehicle corresponds to the occurrence of one pair of edges (pulse). However, some detectors cover two lanes and therefore it is possible that two vehicles travelling side-by-side produce only one pulse. This introduces an error into real-time simulations which may be significant when traffic volume is high. Another source of discrepancies between the simulations and the reality is the absence of information about parking and minor streets traffic and the inaccuracy of the estimates of the turning movement percentages.

Consequently the lane-occupancy data has to be augmented by some additional measurements that facilitate resetting of the cumulative errors in simulations. In the SCOOT environment such a resetting of errors can be facilitated with special queue length detectors. When these detectors indicate continuous occupancy over several seconds this is considered to be an indication of a queue extending from the stop-line up to the detector. The SCOOT-model employs this type of detectors in resetting its own queues so that the >back-of-queue= data reported in the SCOOT messages is a dynamically validated indirect measurement that can be used by the real-time simulator.

Accepting that the simulated traffic will always differ somewhat from reality, because even comparatively small errors can accumulate over long periods of simulation, it is important that a real-time simulator employs some mechanism for resetting the errors. In the HUTSIM model this is facilitated by corrective generation and/or removal of vehicles in the links. This type of procedure has been shown to be effective in correcting discrepancies, on a continuous basis, thus preventing the build-up of errors.

Within the context of real-time simulations, detector data can be used not only to generate the instances of vehicles in specific locations but also to improve the accuracy of the model of driver dynamics, the estimates of turning movement percentages, the estimates of discharge flow rates etc. For instance if lane-occupancy detectors are deployed at each exit from an intersection, turning movements can be calculated on the moving average basis and be compared with the values used by the simulator. If the corresponding queues in the "downstream" links indicate a systematic error that is positive for some links and negative for others, then the turning movement coefficients are adjusted in the model. On the other hand, if the errors are all positive or negative in the "downstream" links, then the discharge flow rates can be adjusted accordingly.

When detector data is available immediately after the stop line, this data could be used for on-line adjustment of the discharge flow rate. This also allows dynamic adaptation of discharge flows which may vary during the day. Furthermore, if separate detector data is available for multiple lanes, this data could be used for tuning the lane change parameters. Also, when lanes determine the direction of turning it is possible to adapt turning movements on the basis of the incoming link data.

It should be pointed out that while HUTSIM normally operates on OD-matrix basis, it can be set to deal with turning percentages by adding route generators to each internal link. A route generator will randomise a new turning direction for every passing vehicle.

4. Technical framework

A prototype of the real-time simulation system has been constructed using the distributed shared memory environment (DIME) developed at NTU. This enables several pieces of software to execute on networked computers while cooperating in performing the simulation task (Figure 1).

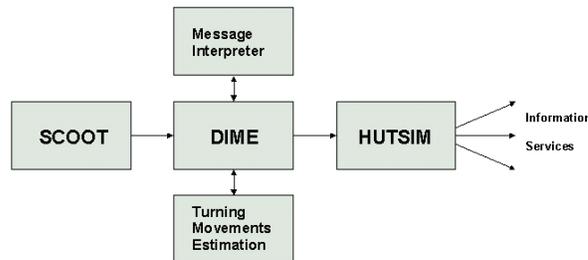


Figure 1 The implementation of HUTSIM / DIME real-time simulation

Each process connects to the shared memory manager which runs in either UNIX or Windows-PC environment. The memory manager acts as a server to which multiple clients can connect. Each

client can have a read and/or write access to several areas (buffers) of the memory manager. A client can also create a memory area with an exclusive write access so that other clients, that connect to that area, are allowed a read access only. Two types of areas are supported in DIME namely buffers for passing messages and arrays for sharing static data.

One of the clients connected to the memory manager is the SCOOT-system which provides real-time telemetry data. The real-time simulation application makes use of two types of messages from the SCOOT system: M19 and M14. The M19-messages supply the detector status data. These messages are generated once per second for each detector and they contain the last four states of the lane-occupancy detector with 0.25 second resolution. The M14 messages are generated every four seconds and supply the link flow and the queue length data together with the signal head status over the last four seconds. Although only the >vehicle occurrence= and the signal head status are mandatory for the simulations, the queue length data from the M14-message is also used for resetting of the cumulative error.

The M19 and M14 buffers of the DIME memory manager are read by the message interpreter client which converts them into a suitable format. The interpreter turns the SCOOT detector data (M19) into HUTSIM vehicle arrival messages by identifying changes in the detector status. Changes in the signal heads status are identified in the M14 data and are converted into HUTSIM signal change messages. Because only the changes of the status are recorded the amount of output data is much less than the amount

of input data. The output messages are written into another DIME buffer readable by other clients like HUTSIM. The message interpreter successfully isolates the simulation task from the SCOOT specific interface issues.

The turning movement estimation client was originally developed for the predictive macroscopic simulation model (PADSIM) but it could also be used here in the context of micro-simulations to provide adaptively updated turning movement coefficients. The estimation results are stored in a static memory array of DIME since the coefficients are updated only every 20 minutes. The message interpreter reads the array and generates HUTSIM-messages that update the turning percentages of traffic generators of boundary links and route generators of internal links. These messages are stored in the same DIME buffer as the vehicle arrival and signal messages. In the current version however, the turning movement coefficients are kept constant throughout the simulation.

Finally the fourth client of the system is the simulation program itself. HUTSIM is operating normally and is running with real-time speed processing external events provided by an input stream. The input stream can be supplied from various sources, but most commonly from an input file. In the case of real-time simulation the input stream is supplied by the message interpreter through the DIME. The present implementation involves three types of messages namely the vehicle arrival-, signal change- and turning percentage messages.

The whole communication link is quite long. From the Mansfield test area to the Nottingham Traffic Control Centre and from there to the memory manager run at the Nottingham Trent University. Finally, after being processed by the message interpreter, the messages reach HUTSIM. Since both the local- and the wide area network are subject to random fluctuations of the communications load, the maximum delay of the whole chain can be several seconds and messages can appear in bursts. Therefore, in the interest of realism the message processing has been based on time stamping. By offsetting the simulation, by a time representing the maximum communications delay, the random delays in communications are prevented from affecting the simulation. In the test system, a 15 seconds offset has been found sufficient under most circumstances.

4.1 Simulation of Mansfield test area

The HUTSIM / DIME real-time simulation system described here has been implemented at the Nottingham Trent University and is connected to the SCOOT-system of the Nottingham Traffic Control Centre. The system typically runs as a fully distributed application with each of the four cooperating tasks executing on a separate computer. However the DIME framework does not place any constraints on the hardware configuration and e.g. HUTSIM and the message interpreter can run on the same PC while the memory manager and the turning movements estimator can execute either on a UNIX workstation or a PC. The HUTSIM / DIME simulation is running successfully and the framework has shown itself to be reliable.

The performance of the system has been validated through real-time simulations of the Mansfield-south SCOOT area. This test area covers six intersections in the of Mansfield. The rush hour traffic in this area is very busy. A detailed HUTSIM-model was constructed (Figure 2) including the geometry, lane organisation, detector positions etc. The boundary link detectors were replaced with vehicle generators and internal links were equipped with route generators. All signals and generators were labeled according

to the link and detector numbers of SCOOT so that that all messages could be delivered to the correct object.

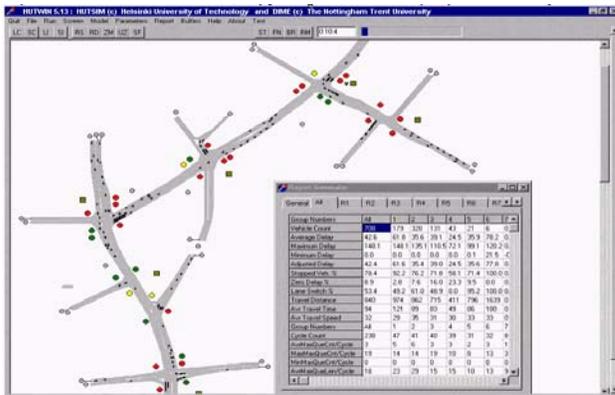


Figure 2 Real time simulation Mansfield test area

In order to compare the simulated and the real traffic, a "bird's eye" video recording of the entire SCOOT area would be most useful. However, such an recording was beyond the means of this project. An alternative video recording of individual approaches to all simulated intersections was also considered uneconomical and

exceedingly tedious when attempting to time-synchronise 24 video recordings and the real-time simulations.

As a practical approach we have used therefore the queue measure from the M14-SCOOT message and have compared it with the queue simulated by HUTSIM. The SCOOT queue is expressed in LPU:s (Link Profile Units) and has been found to provide a good estimate of vehicle count. The results of the field validation of the SCOOT queue as a measure of the number of vehicles on the stop-line has been performed for a number of intersections and a representative set of results is given in Figure 3. It can be seen that the majority of queue readings give an unambiguous vehicle count and the statistical variability of readings is a realistic reflection of the variability of vehicle types.

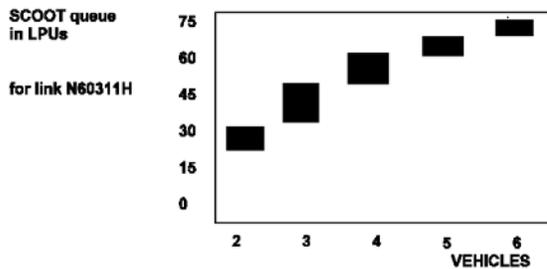


Figure 3. Validation of the SCOOT-queues as a measure of vehicle count

Figure 3 Validation of the SCOOT-queues as a measure of vehicle count.

Preliminary simulations show that the matching of vehicle counts on boundary links was good. In internal links the matching between simulation and field data can be improved by tuning the model parameters. Especially the discharge flow parameters may need to be adjusted properly for each link. The results show that there is also a need for

refinement of the calculation of dynamic turning movement coefficients which are adversely affected by the missing data in some of the links. Consequently, static turning movement percentages need to be used in some cases.

Because the number of settings and parameters can be quite large, the manual tuning of the simulator can easily becomes tedious. The preliminary experiences indicate

that there is a need for adaptive tuning of the parameters. This is one of the areas to be addressed in future research.

The computational performance of the simulation system executing on a 200MHz Pentium PC has been found to be sufficient for this size of model. For larger models the performance issues may need to be reconsidered. The critical factors are the size of the simulation model and the maximum rate of data communication.

4.2 Developing Internet and GIS interface

The HUTSIM / DIME real-time simulation system can already be used by several parties who already have the simulation program itself and have access to the Internet connection. However, it is believed that many more users and general public could benefit from the results of real-time simulations of urban traffic if these are made available in an accessible form on the World Wide Web (WWW).

Building on the HUTSIM / DIME system the research has been extended to develop techniques facilitating the use and diffusion of traffic data within a geographical context. We intend to provide, at the interface level, an open and flexible system for managing the integration, manipulation and display of geographic information for the benefit of traffic management applications. The system will allow to display and query geographical and traffic data in an integrated graphical interface. Traffic data will be overlaid on urban maps for location reference. These multiple data views allow the user to examine different aspect of data. The identification of potential users and respective technical skill and needs is an important objective of our interface development (McGraw 1992). Modularity allows a progressive application and to accept different user-interaction levels from engineers to decision-makers. In particular, our system is seen as contributing to a broader issue of investigating new solutions for the diffusion of traffic simulation results.

From a database point of view, the traffic system domain exhibits both static and dynamic characteristics. Static components are data objects such as street layouts and topological features such as roads. These features are considered as static within a traffic system that examines short time scales in object transformations. Dynamic components (measured or simulated) are objects that change on a relatively short time scale from an application point of view (e.g. cars, traffic queues, traffic signals). Our objective is to identify significant properties about the dynamic traffic mechanisms, to explicitly record relationships among entities involved in traffic processes, and to model causal relationships when they are identified from an application point of view. We consider different abstraction levels in both space and time in order to integrate different levels of spatial and temporal details. At each abstraction level, computing models generate a set of database primitives from measured and simulated measurements (Figure 4).

The proposed interface integrates geographical data and traffic data and supports adaptive transmission depending on the final user requirements and capabilities. The database integrates the static and dynamic components of a traffic system including the network infrastructure, real-time traffic data (e.g. queues, traffic flow), geographical data, descriptive and statistical data. The geographical component of the interface is used as a context background that complements the visual environment of the graphic interface.

The simulation results can also be made accessible directly through the Internet as discussed by (Gul, 1998). Graphic and numerical data provided by HUTSIM and SCOOT can be made visible for ordinary web users. In Figure 5 alternative methods are outlined.

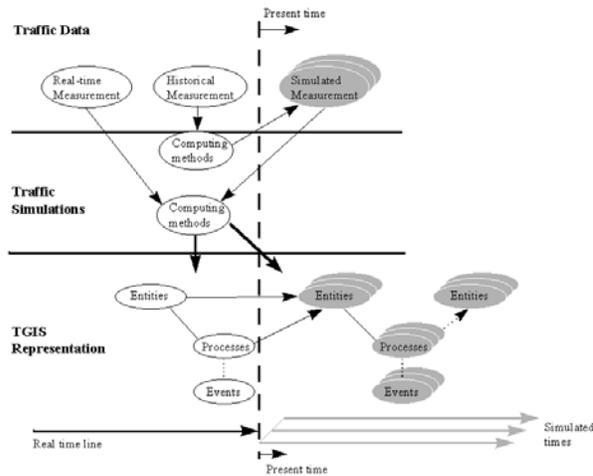


Figure 4 Temporal geographic information system (TGIS) - traffic model principles.

The most straight forward method is to update HTML-files on the web-server. The simulation could produce an HTML-file as an output, which is then transferred to the server for replacement of the existing file. Initially a BMP file is generated by the HUTSIM application every minute this is a screenshot of the simulation depicting the current traffic situation. Once the BMP image is

captured it is converted into a GIF image which then dynamically updates an HTML file at a given time interval. This allows users with net connections to view the traffic simulation in a static form.

Further output tables can be generated by the HUTSIM application to display the delay, traffic queue, and travel times statistics as shown in Figure 2. An output file containing this report data can also be generated to be used in web pages.

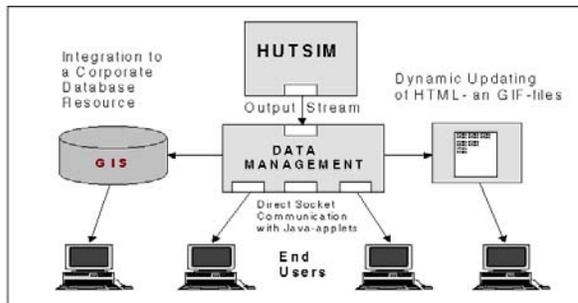


Figure 5. Outline of the information services through the Internet

Figure 5 Accessibility of the simulation results via the Internet

More advanced solutions can be obtained by utilising the HUTSIM output stream (Table 1). The output stream can be directed to a socket interface that allows other clients only to read the data in real-time.

Table 1 Data Output Stream

Time	Time
RT 01	RT 02
Object_ID	Light_ID
X1	X1
Y1	Y1
X2	Light Status
Y2	
CR/LF	CR/LF

The Java application can work by opening a client socket to receive the data output stream from the HUTSIM application. The socket can use TCP/IP-protocols to establish

the link between the two applications. A Java application can be programmed to receive the vehicle coordinate data as shown in [Table 1](#), and to show on-line animation of the simulated traffic.

A carriage return (CR)/ line feed (LF) command follows each string this determines the next set of data to read. This string is received for each vehicle on the system and traffic light on the simulation. As this is received the data is tokenised, so that the sets of data can be grouped for plotting the cars and lights, within the memory buffer.

With the coordinates being received into memory, these can then be plotted onto the screen ensuring the vehicle coordinates and the screen image of the Mansfield area coincide, so that when the cars are drawn on the screen they are actually on the roads and not on the verge. This animation would run in a web-browser showing the overview of traffic situation to the user.

The application is compliant with different user-needs. For instance, very fast Internet connections may use on-line animation while slow connections may display a relevant graphic state of the traffic system using appropriate time scales that depend on the local system configuration and on the application objectives. Different Java application levels offer different views depending on the user needs and computing configurations (e.g. terminal types, screen sizes).

Users of such a system would be all organizations involved with traffic system operation such as emergency services, network maintenance, planning and environmental research. Such an application provides an adaptive environment that opens new perspectives to the use of traffic simulation systems in the new emerging information society.

Conclusion

The HUTSIM / DIME system has been implemented and is running successfully on a network of distributed computing nodes. The system demonstrates the principle of real-time simulation and its benefits. Further research is needed to study in depth the achievable accuracy of real-time simulations. This will involve full scale comparison of field data with simulated measures. The research is also likely to result in improvements in the telemetry system by suggesting the optimal number, type and location of detectors in support of real-time simulations. The motivation for further research is twofold: to provide a reference data for the various information services and to develop a reliable predictive model of urban traffic in support of future traffic control strategies.

Increasingly, the combination of simulations and the real traffic data are providing essential information for management and decision making. In this context the capability to communicate information in an intuitive way through a geographic information system can lead to involving general public in aspects of urban traffic planning that touches upon their daily lives. Traffic information systems can be used for the analysis of a wide range of practical environmental issues such as noise and air pollution management. Integration of traffic simulation information within graphical systems will provide crucial information to environmental monitoring and planning applications. An increase in the use of such systems in the context of urban traffic management is likely to enhance the sustainability of cities, which has been identified as a priority research task.

References

- Abel D. J., Yap. R., Ackland R., Cameron, M.A., Smith, D. F. and Walker G.** (1992) Environmental Decision Support System Project: An Exploration of Alternative Architectures for GIS, *International Journal of GIS*, 6(3); 193-204.
- Argile A., Peytchev E., Bargiela A., and Kosonen I.** (1996) DIME: A shared memory environment for distributed simulation, monitoring and control of urban areas, Proc. ESS'96, Genoa, 1996, Vol.1, pp.152-156.
- Etches A., Claramunt C., Bargiela A., and Kosonen I.** (1998) An Integrated Temporal GIS Model for Traffic Systems. Geographic Information Systems Research in UK (GISRUK). April 1998, Edinburgh, UK.
- Goodchild M. F.** (1998) The Geolibary, In: *Innovations in GIS 5*, S. Carver (ed.), London: Taylor and Francis, pp. 59-68.
- Gul I.** (1998) A Java Interface for Traffic Simulation, MSc Industrial Computing Systems Dissertation, The Nottingham Trent University.
- Kosonen I.** (1996) HUTSIM - A Simulation Tool for Traffic Signal Control Planning. Helsinki University of Technology, Laboratory of Transportation engineering. Licentiate Thesis. pp. 150.
- McGraw K. L.** (1992) *Designing and evaluating user interfaces for knowledge-based systems*. Chichester: Ellis Horward.
- Moss et al.** (1998) The KINDS project: providing effective tools for spatial data accessibility and usability over WWW, In: *Innovations in GIS 5*, S. Carver (ed.), London: Taylor and Francis, pp. 69-77.
- Peytchev E., Bargiela A. and Gessing R.** (1996) A predictive macroscopic city traffic simulation model, Proc. ESS'96, Genoa, 1996, Vol.2, pp.38-42.
-