S-Paramics Principles

Paramics Microsimulation
SIAS Limited
37 Manor Place
Edinburgh EH3 7EB
UK

www.paramics-support.com
www.sias.com

Copyright 2009, Paramics Microsimulation, SIAS Limited

All rights reserved. This document is subject to copyright protection.

No portion may be copied without the prior written consent of SIAS Limited.

Why Microsimulation?

Road traffic microsimulation is a computer modelling system which represents the behaviour of individual vehicles and their drivers in a road network. These are modelled to observe the rules of the road and to interact with other road users through simple rules. The cumulative effect of modelling individual vehicles is to realistically represent road traffic flow on a physical road network. Microsimulation is a powerful communications tool because it is able to present its outputs as a real-time visual display.

Microsimulation, exemplified by **S-Paramics**, is a very different methodology from *deterministic* traffic modelling, as used by SATURN, TRIPS and VISUM. These ensure that a particular solution will always result from a particular set of input data. In reality, traffic data is rarely constant and repeatable, and the microsimulation methodology has been developed to enable traffic modellers to produce design solutions which encompass variability.

Deterministic systems are unable to address situations for which they have not been specifically designed, and this presents problems for using them as a predictive tool. The interactions in congested road networks produce complex effects, and microsimulation is able to spontaneously represent these and enables traffic modellers to understand their designs in a context of high traffic volumes.

A microsimulation model can test in detail the effect of redesigned road layout, changes to traffic demand and traffic control and intelligent transport systems (ITS). As in real life, modelling small local schemes and effects can have a significant impact on traffic flow over a wide area, an effect that other modelling methodologies have difficulty in reflecting.

Figure 1: S-Paramics linked to ITS

Interfaces from S-Paramics to ITS and Urban Traffic Control systems enable simulation models to be used to develop control strategies for incident and event management and to investigate options for optimising adaptive signals, and urban or motorway control systems. Figure 1 shows an example of this.

The strength of S-Paramics is its ability to apply microsimulation to large area models. The *micro* in microsimulation reflects the level at which the interactions between vehicles is modelled, not a limit to the extent of the geographical area covered. Large models may extend to hundreds of square kilometres, and the ability to apply the inherent fidelity of microsimulation to wide area traffic issues makes S-Paramics an invaluable tool for traffic modellers.

Overview of S-Paramics

S-Paramics is a software system which simulates the individual components of traffic flow and congestion. It presents its output as a real-time visual display for traffic management and road network design. S-Paramics represents the actions and interactions of individual vehicles as they travel through a road network and models the detailed physical road layout, including features such as bus operations, traffic signal settings, driver behavioural characteristics and vehicle kinematics. As a consequence, S-Paramics can portray and evaluate the variable circumstances which lead to congestion in all types and sizes of road network.

The components of **S-Paramics** are:

- **S-Paramics Explorer (SPX)** provides an intuitive single interface for accessing each of the main S-Paramics modules and the various maintenance utilities. It is based on the concepts of Microsoft Windows Explorer. Its primary functions are to launch the S-Paramics modules, and to help users manage and navigate around their S-Paramics models. There are also a number of other functions including licence management, batch farm management and operation, and an interface to associated external applications.
- The **Editor/Simulator/Visualiser (ESV)** module is the primary S-Paramics interface. It is used to create and edit a model, to run that model, and to view the vehicles in it. It can also write summary statistics describing the traffic flows in the network.
- The **Simulator/Visualiser (SV)** module is a cut down version of the ESV module, which does not allow models to be edited or statistics to be generated. It is an add-on module to allow models to be demonstrated without the need for an additional full S-Paramics licence.
- The **Batch Run** module uses the same simulation engine as ESV but has no graphical interface. It is much faster to run to produce summary statistics of the traffic in the model, but offers no opportunity to edit the network or view the vehicle behaviour.
- The **Matrix Estimation (ME)** module takes as input route information from an existing S-Paramics model, a prior matrix estimated from observed data, and survey data comprising roadside counts. It then estimates a trip matrix consistent with the three sources of input data. It also allows sensitivity testing to determine the reliability of the survey data.
- The **Statistics** module is the S-Paramics Data Analysis Tool (DAT), which takes summary statistics produced by the ESV or Batch modules and displays and analyses them in graphical form.
- The **Economic Assessment (PEARS)** module is a component of DAT.
- The **Batch Farm** is an integral part of SPX. It can be used to manage the throughput of large numbers of simulation runs on a network of computers.
- The **Advanced Controller Interface (ACI)** enables adaptive signal control and links directly to intelligent transport systems.

Model Building

Principles

S-Paramics represents the complex and apparently random nature of traffic flow by requiring the user to provide limited and simple components in the form of a description of the road network and the traffic demand. S-Paramics does not require the user to anticipate where problems may exist and to code attributes to address these, because their manifestation is a natural consequence of the microsimulation process. S-Paramics applies real world effects, and does not rely on the use of artificial modelling constructs and arithmetical proxies to represent road traffic.

For example, S-Paramics does not require the user to prescribe the proportion of vehicles using a lane at any point in the road network, for instance at the approach to a junction. It is, however, possible to define where drivers become aware of a junction and hence where they will start to make decisions about their next manoeuvre. The flow of vehicles in each lane occurs as a result of decision making within the model. If the model fails to replicate the observed data the modeller is encouraged to re-examine the opposing flow of vehicles and turning movements as well as the position at which drivers make a decision on lane choice.

S-Paramics uses a descriptive methodology of controlling driver behaviour rather than one of prescribing the desired effect. This gives the model a more robust predictive ability in testing changes to the road network and travel demands.

Network Construction

Building a simulation model comprises a number of key stages. The scope of the study area must first be established to include the area in which both direct and indirect changes may occur. A zoning scheme is created to determine where road users gain access to the network. Each zone forms a collective origin and destination for the demand, which is represented in terms of travel between these zones. The model of the road network is then built, to include links, zone connectors, junctions and traffic control. Detail such as aerial photograph overlays, 3D buildings or street furniture may be added to enhance presentation of the model.

Once a model has been created, its performance is measured against field observations of traffic flow. A process of model calibration follows, which is an iterative procedure required to tune the road network description to ensure that the traffic flows in the model are a proper representation of the observed flows.

Road Network

S-Paramics describes roads as a set of pairs of one way links joined together at nodes. Links represent the roads in the network and have attributes such as speed, width, number of lanes and any lane based restrictions such as bus lanes. A node typically marks an area where links join to form junctions. Nodes can also mark points where links change characteristics. These can be any of their attributes, such as speed limit, number of lanes or a change in curvature or direction. Figure 2 shows how nodes, links and zones are arranged and numbered in an S-Paramics model.

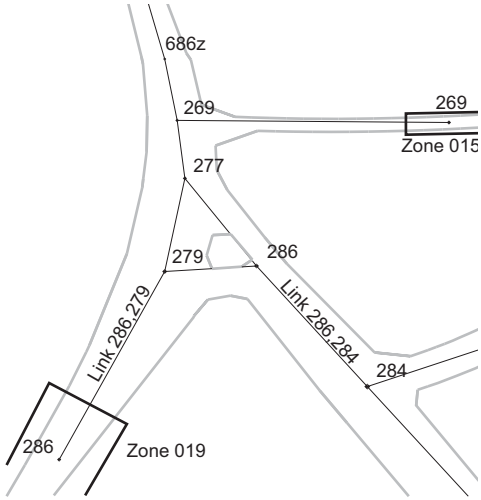

Figure 2: Links, nodes and zones

At junctions, stop lines mark the points where drivers wait for gaps in oncoming traffic and where they move from one link to another. The stop lines also serve as locus points to describe the paths that the vehicles will follow.

S-Paramics uses a hierarchy of major and minor links to control the routeing of vehicles. The major links are in effect the signposted routes and the minor links the secondary network. Drivers are classified as *familiar* and *unfamiliar* with the latter perceiving the minor routes as less attractive than the major routes.

To manage the task of creating links, S-Paramics has a set of link categories. The concept is similar to that of a word processor's style sheet. Categories can be defined and applied as the model is built and a set of predefined categories is also available. Links can later be modified either by category or individually.

Signals are added at the signalised junctions in the model. These may be formed from single nodes or may have more complex geometry and include several nodes. Phases and stages are assigned to the movements and stage times allocated. Signals may be further controlled with a simple *plans* language to adjust stage times when required, or may be controlled by an external signal controller linked to the simulation.

Zoning Scheme

Zones represent the network entry and exit points for vehicles. They are the logical network access points for vehicles and may be of a particular land use. At the edges of the network zones represent external demand entering/leaving the network from/to a specific direction. Simple, single junction or corridor models may only require external zones.

Vehicles and Demand

Vehicles

S-Paramics can model many vehicle types ranging from cars to HGVs and light or heavy rail and buses. Vehicle types may be used to differentiate between:

- vehicles of different types, e.g. cars, LGVs, HGVs and taxis
- vehicles with different physical characteristics, e.g. small, medium or large cars
- vehicles with different journey purposes e.g. commuting, business, leisure
- vehicles with different emission characteristics
- passenger transport vehicles

Each vehicle has a set of basic physical properties such as size and number of sections (for articulated vehicles), maximum speed, acceleration and deceleration. Other parameters are also defined, including the allocated demand matrix, and engine type, which governs the quantities of emissions the vehicle will generate.

Vehicles in the simulation have a notional driver who is ascribed levels of awareness and aggression. The awareness parameter controls the likelihood that a driver will collaborate with others on the road, e.g. by adjusting headways to allow others to make lane changes. Aggression controls how a driver behaves with respect to speed selection and lane use. A more aggressive driver will tend to travel faster and delay lane changes required for a turn, preferring to use lanes with faster moving traffic.

Vehicle types also define a physical appearance. While this has no effect on the results of the simulation, it has a significant impact when presenting the model to a wider audience, perhaps at a public inquiry. Presentation options range from complex and realistic shapes to simple cuboids. The decision about which to use depends on the target audience and the speed of visual rendering required.

Trip Demand

S-Paramics uses origin-destination (OD) demand matrices to control the loading of vehicles into the network. OD demand matrices may be disaggregated into a number of matrix levels which can represent different vehicle or journey types or be used to add development specific demands to an existing model. For example, a model may include a matrix of current demand and a second one which contains additional demand relating to new developments. It may also contain a car based demand matrix and a separate HGV matrix.

Profiles

User defined departure time profiles control vehicle release rates, and may be applied separately to each matrix. These profiles play a large part in reflecting the constantly changing demand throughout the day and enable the deployment of microsimulation models across wide time spans.

Consider the example in Figures 3a and 3b. The two graphs each represent traffic demand over the same three hour period. The total hourly demand is the same in both cases, but the demand profiles are expressed to different levels of detail. The one hour profile in Figure 3a overestimates the number of vehicles in the initial half hour of the simulation and underestimates the peak in the middle of the three hour period. The one hour profile would not allow the model to reflect the traffic conditions in the peaks and the troughs.

S-Paramics allows the modeller to adjust the number of vehicles released in each five minute interval to represent the prevailing traffic flows.

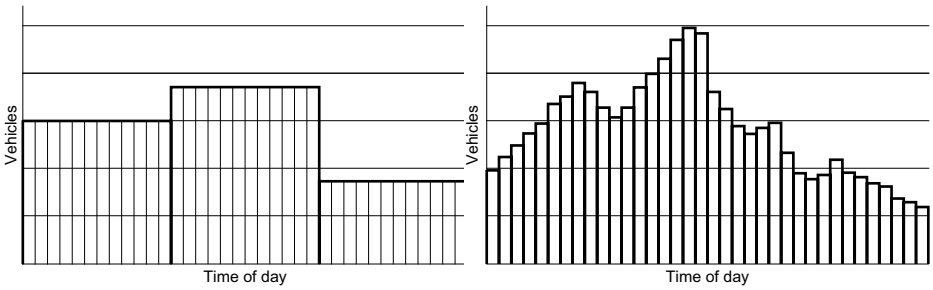

Figure 3: Demand profile detail 1 hour demand (a) v. 5 minute profiled demand (b)

Passenger Transport

Scheduled buses are programmed into the simulation by specifying their fixed routes, bus stops and timetables. Bus routes are created automatically along the shortest path between specified start and end links. Intermediate links are specified where the route does not follow the shortest path. When a variant of the model is made with changes to the road network, the bus routes are automatically re-created with the same start, intermediate and end links and will automatically find a path through the altered network.

Time Periods

Time periods are used to segment the overall period for which the model runs. Typically a period boundary marks where a significant time-based change occurs in the model. This could relate to:

- lane restrictions, e.g. bus only lanes change to become traffic lanes outside the peak period
- junction priorities change, e.g. a turn is barred in a peak period
- vehicle type proportions, e.g. ratio of HGVs to cars

Time periods may also be used to implement changes in demand or changes in signal timings but often it is better to implement these using the demand profiling facility or signal time plans.

Presentation

Annotation, aerial photographs, buildings and other 3D shapes may be incorporated into the model to enhance its presentation quality. The amount of effort worth investing depends on the use and intended audience for the model and the budget allocation for presentation.

Figure 4: Presentation options

The quality can vary from a simple schematic network representation to a full 3D virtual reality model. Simpler background graphics can use appropriate generic buildings, landscape objects and street furniture that are available with a standard S-Paramics installation. A high end presentation may require development of a 3D landscape or cityscape of the modelled area. Figure 4 shows three images of the same simulation scene with different presentation levels, increasing in complexity from left to right. Figure 5 shows S-Paramics simulation data included in a complex 3D model.

Figure 5: 3D presentation (images courtesy of baseplusworld & Truescape)

Route Choice

Assignment is an area where microsimulation differs from deterministic modelling systems. In a deterministic model, assignment refers to the loading of traffic demand (in the form of vehicles) onto the network along all links simultaneously along routes predetermined during the process of building *trees*. In microsimulation, assignment is the aggregation of vehicles as they travel.

S-Paramics has a range of algorithms to find routes for vehicles in a simulation model. The static route finding approach allows a driver to find a route based on the static information about the network. This includes:

- the driver's perception of the cost of travel between O-D pairs and waypoints. This is a combination of time, distance and monetary tolls
- the variability of this perception of cost
- the road hierarchy, which influences some road users' propensity to use certain links

Dynamic route finding adds the effect of congestion in different areas of the model to the driver’s knowledge and enables him to adjust his route as he learns about congestion.

Road Network

An S-Paramics road network is built with a two level hierarchy of links. The *major* links correspond to the main roads, often characterised as the signposted routes. The *minor* roads correspond to the secondary routes. The incremental cost of using minor routes is increased for those drivers marked as unfamiliar. In Figure 6 the major roads are shown in black and the minor roads in grey.

Figure 6: Road hierarchy

At each junction there are tables of costs for familiar and unfamiliar drivers from each of the exits of the junction to each of the destination zones in the network. Drivers approaching a junction consult these tables and select the appropriate exit depending on the subsequent cost to the destination.

S-Paramics can accommodate physical, or statutory, restrictions in the network, such as height restrictions and the banning of HGVs from city centre streets during peak periods.

Car Parks

S-Paramics has the capability to refine the start and end points of trips through the use of car parks. Car parks can be coded to reflect actual car parks or be used to define multiple trip origin/destination points within an area. The car parks are associated with single or multiple zones to connect them to the origin/destination data in the demand matrices. Walking times are assigned between each car park and the centre of associated zones and contribute towards the overall trip cost. The car park with the lowest overall trip cost will be selected when deciding which to use to begin or end a journey.

Car park occupancy can be constantly monitored during the course of the simulation. When a car park is full, vehicles may queue at the entrance until spaces become available or re-route to alternative car parks with spare capacity. ITS systems in the simulation may change the destination car park for individual vehicles.

Driver Knowledge

Drivers are classified as *familiar* or *unfamiliar* in the routeing system, corresponding with their knowledge of the conditions of the network, and how they perceive journey time costs on major and minor routes. Unfamiliar drivers tend to keep to the main (major) roads, while familiar drivers use both main and secondary (minor) roads. The ratio of familiar to unfamiliar can be set for each vehicle type. For example, it could be assumed that taxi drivers are 100% familiar. The driver attributes of *aggression* and *awareness* are also used to determine how they react to congestion, whether they accept the extra delay or whether they opt to save as much time as possible by using every available rat run.

Static Route Choice

Static routes are deduced by minimising time, distance and monetary tolls. Different drivers may place different emphasis on them when combining them to form a perceived minimum route *cost*.

In evaluating the cost of a route, the incremental cost of minor links is increased for the unfamiliar drivers. Hence, when there is route choice, they are more likely to select the routes made up of major links and less likely to use the minor links as rat runs.

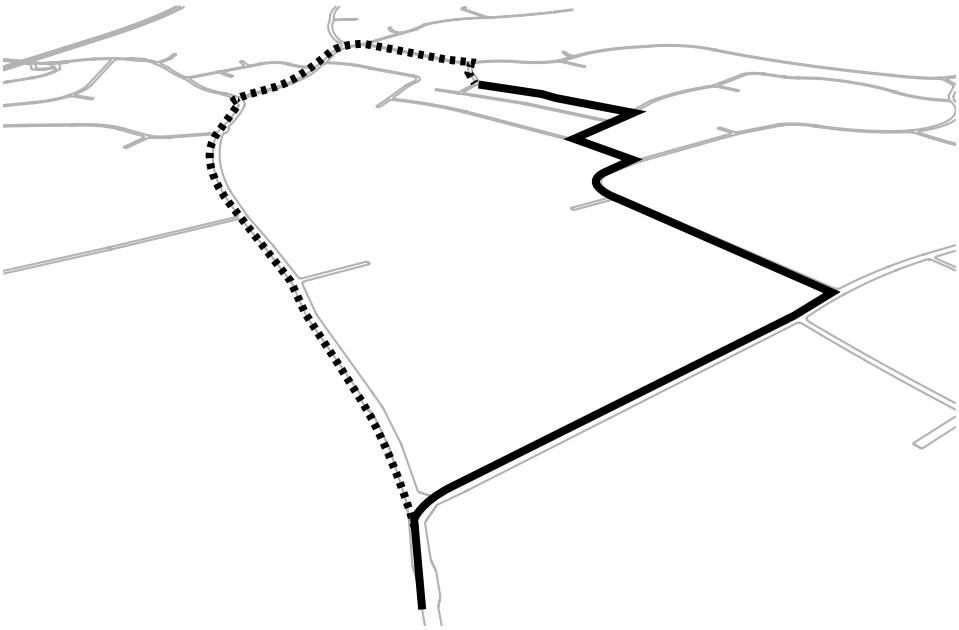

Figure 7: Route choice due to perturbation

When there is a choice of routes from a junction to a destination a driver must determine the best exit to take. To reflect the driver's imprecise perception of the true route cost from each exit, before the decision is taken, values are randomly varied up to a predetermined level. This process, known as *perturbation*, ensures that a number of different routes of comparable cost may be used.

In Figure 7 two routes are shown to a destination. One has a journey time of 10 minutes, and the other 10 minutes 30 seconds. With 10% perturbation the relative costs overlap and both routes are used by drivers at this junction.

The selection process for the next turn based on route cost is repeated at every junction. Route cost perturbation ensures that a range of routes may be used by vehicles travelling between the same origin and destination. The lowest cost route will be preferred, although other plausible routes may also be used.

Dynamic Route Choice

Dynamic assignment includes the route finding algorithms used in the static assignment method and adds the ability for drivers to learn about the congestion they will encounter on their journey. Drivers may adjust their routes in reaction to this acquired knowledge. As each individual reacts, the location and severity of the congestion will vary. Over time, drivers will try to avoid congestion and learn the quickest routes to their destinations, collectively minimising the aggregated journey costs.

A pragmatic approach to dynamic assignment is taken by S-Paramics. The concept of a unique solution or perfect optimisation has little meaning in the context of variable driver behaviour, travel demand and trip cost perception. Dynamic feedback in S-Paramics adjusts the transit times for each link at regular intervals throughout the simulation, which simulates the driver learning process. Familiar drivers apply their understanding of the congestion on the road network by using revised times to minimise their journey costs. Unfamiliar drivers will continue to follow static routes, with their perception of journey time based on the advisory speed limits only.

Drivers who receive feedback of link based journey times may vary in their responses to it. The degree of journey optimisation relates to the level of driver aggression and awareness. S-Paramics allows modellers to vary driver responses to knowledge of the network congestion based on the distribution of driver aggression and awareness.

Over-prescription of routeing is not encouraged in S-Paramics. For example, turning proportions or lane use cannot be directly programmed into the simulation, because to allow this could potentially mask a misunderstanding of behaviour within the model. Variation of individual link costs is possible, but should only be undertaken when the changes can be justified following the observation of on-street conditions.

Multiple Level Routeing

As drivers in the real world who have asked for directions will be aware, the form of the response varies according to the distance involved. If the destination is nearby then very detailed instructions may be given, but if it is further away, the instructions will tend to be based on key points on the journey rather than a long list of individual turns. Typically, the detailed directions are only given for the immediate segment of this overall journey.

S-Paramics embodies these two modes of route understanding by using *waypoints* at strategic locations in the model. Waypoints can be thought of as key route decision points, which may be large junctions or identifiable areas which serve to define long distance routes.

A vehicle in a model using waypoints will have a *macro* level route linking its origin zone to its destination zone. This may pass through a number of waypoints. The overall journey is therefore segmented and each smaller segment is traversed by vehicles using the *micro* level routes. Macro level routeing between waypoints may be controlled with cost equations, perturbation and dynamic feedback in the same way as the micro level routeing between junctions.

A consequence of using the single level route algorithms is that at the start of a long trip, the application of perturbation may cause a driver to select sub-optimal routes for the first part of the journey. When waypoints are used, the micro route for each segment of the overall trip is independently derived and hence the initial routeing decisions are not masked by the perturbation calculations for the longer journey.

Calibration

Calibration is the process of adjusting the parameters used in the model to ensure that it reflects the input data. (*Validation* is the process of comparing the model outputs with observed data that has not been used to adjust the model parameters in the calibration process. If the model is validated using this independent data then the modeller may be confident that it is capable of predicting effects that were not explicitly programmed into it.)

Assignment Calibration

If a model is of a simple corridor or single junction, assignment is achieved with an *all or nothing* approach. If the model has route choice, then correct assignment of vehicles to the model is a key aspect to calibration. Assignment is calibrated by modelling the road network to reflect the true levels of congestion and hence delays in the system, and to providing good quality demand matrices and release profiles. Assignment calibration is an iterative process. Adjustment to the demand matrices and the road network description will affect flows at junctions and the resulting delays, which in turn affect the assignment. S-Paramics includes a Data Analysis Tool to enable the modeller to compare modelled data with surveyed data in order to check on the progress of the model calibration.

Demand Calibration

Good quality travel demand data is typically collected during roadside surveys, because it provides the information of where journeys start and end. Such surveys are expensive and inconvenient to conduct, and automatic number plate recognition (APNR) technology is often used as a substitute to provide data from which trip origins and destinations can be inferred. Manual classified vehicle counts may be sufficient for very small (single node) models, but ideally data should come from all these sources, and be supplemented by, or at least checked against, information extracted from the national census.

Whatever the sources of data, these cannot hope to provide a complete picture of movements between all origins and destinations (model zones) for anything but the smallest of study areas. To overcome the inevitability of partial data, S-Paramics includes a *matrix estimation* module to modify matrices so that they represent the full observed traffic conditions. Matrix estimation is an iterative process using the current assignment in the model and a *prior matrix*, which represents a good estimate of OD travel demand obtained from the observed data. Unobserved cells in the prior matrix are first seeded where, with local knowledge, such a trip is seen to be possible. Other constraints are applied to ensure that the numbers of vehicles flowing into and out of zones are consistent with their size and demographics. The matrix estimation process adjusts the prior matrix in a manner which matches the assigned predicted flows to the surveyed volumes.

Traffic demand is rarely, if ever, constant for long periods of time. Flows build and decay during the day and the level of congestion in a real world network is heavily dependent on

the rate at which this happens. S-Paramics reflects this through the use of time based *flow profiles*, which control the release of vehicles into the simulation at five minute intervals. Profiled demand is a key factor in model calibration, and is essential to the understanding and representation of traffic flow in congested networks. It provides for realistic queueing at junctions which in turn has a significant effect on dynamic assignment.

Driver Behaviour Calibration

The development of an S-Paramics model should not normally require alterations to the global parameters which affect driver behaviour. In general, driver behaviour fluctuates in response to specific road circumstances, and network wide changes should not be introduced to override S-Paramics unless a sound case can be made that drivers behave differently across the entire modelled area. The key overall driver behaviour parameters are:

- driver aggression and awareness distribution
- network headway factor

Network Calibration

A key element in model calibration is the proper functioning of each junction in the network. This is achieved by adjusting the paths that vehicles follow through the positions of stop and give way lines. Link attributes are also adjusted to define permitted manoeuvres from individual lanes and to identify where drivers assess opposing traffic flows, particularly on junction approaches.

Stop lines control the paths followed by vehicles through junctions and where drivers stop to wait for gaps in opposing traffic. The curvature of a path controls the speed at which the vehicle can traverse the junction, a tight turn resulting in a lower speed than a shallow curve.

Permitted turning movements are defined for each lane entering a junction, and may also be mapped to exit lanes. A junction must be calibrated for correct use of turning lanes and movements. Turning movements are prioritised in S-Paramics as *major*, *medium* or *minor*.

Stop line positions, permitted turning movements and movement priorities are all estimated by S-Paramics from the network description when the model is built, but complex junctions require additional refinement by the modeller.

Hazards and Signposts

Away from the influence of junctions and changes in road layout, a vehicle's speed and lane are set with reference to the road layout and the proximity of other vehicles. Closer to a junction, or a node where the road layout changes (e.g. a lane gain or a lane drop), a

driver must reassess the lane choice in the context of the forthcoming manoeuvre. A location where any action may be required is referred to as a *hazard* and the point at which drivers become aware of this is referred to as the signpost distance. Figure 8 shows a *diverge* hazard at a motorway off-ramp.

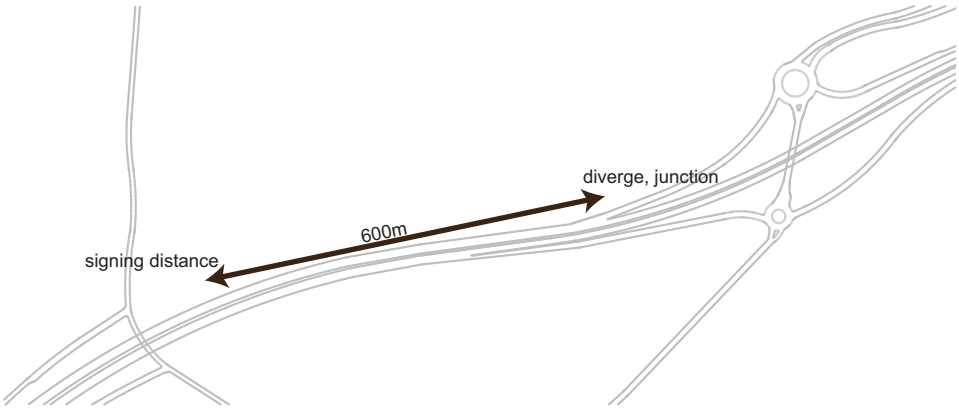

Figure 8: Diverge hazard

The aggression and awareness of each driver control the setting of vehicle speeds and headways, and the speed at which they react to hazards when first notified at the signpost. When there are no active signposts or hazards, less aggressive drivers tend to stay in lane and only change when the speed differential is high. More aggressive drivers select lanes with faster traffic based on smaller differentials in vehicle speed. On becoming aware of a hazard, the less aggressive drivers re-evaluate their lane choice and change lane to make the manoeuvre if required to do so. More aggressive drivers re-evaluate their lane choice later.

Visibility

Visibility is a key calibration parameter at junctions for medium and minor movements. If the visibility is set to 0m, drivers will stop at the stop line before they determine if there is a gap to move into. If visibility is more than 0m, then at this distance from the stop line, drivers will start to determine if there is a gap in opposing traffic as they slow on their approach to the stop line. If a gap is present, they may not need to stop and may proceed to cross the junction.

The level of visibility is determined by the geographical features of the junction. Figure 9 shows two roundabouts. The first has low visibility due to high walls and gradients approaching the roundabout. Drivers inevitably have to come to a stop as they reach the roundabout before being able to move safely. The second is more open with good visibility. Here, drivers are less likely to come to a stop before moving out onto the roundabout. Visibility has a considerable effect on queueing at roundabouts.

Figure 9: Comparison of visibility at roundabouts

Gap Acceptance

There are three gap acceptance parameters: lane merge, lane cross and path cross. These control the gaps that drivers will accept as they merge or cross into traffic streams. They are normally only changed if there are special circumstances at a junction, or drivers are observed to regularly accept small gaps at congested junctions, or those with poor visibility. Similarly, default vehicle headways may be modified for individual links if evidence is available to support these local modifications.

Link Attributes

While S-Paramics enables the modeller to code speeds and wait times at the end of links, this is not intended to aid the calibration by influencing the queueing behaviour. This facility is provided for the modelling of speed bumps and toll booths. Junction throughput should be calibrated by correctly modelling the opposing flows through accurate assignment and by adjusting vehicle swept paths, stop line positions, and visibility for manoeuvres. It should not be calibrated by including artificial delays or speeds.

Simulation

The simulate mode of ESV runs the model, while the visualiser works in tandem to display vehicles travelling through the modelled network in real time.

When the simulation is running, an internal clock synchronises the sequence of events. At every second, the demand and demand profile in each zone is examined to see if a vehicle is to be released into the network. Each vehicle has a destination zone marking the end of its trip, and travels through the network to reach it.

At each simulation time step (typically 0.5 second) every driver determines what to do next, based on four simple decisions:

- which lane should I be in?
- at what speed should I travel?
- if I am waiting for a gap in traffic, can I move yet?
- which exit should I take at the next junction?

Drivers' decisions take into account other vehicles. The speed is set to ensure that drivers follow at a safe distance and choose a lane that allows for the next turn, and to only change lane when there is a suitable gap. The cumulative effect of the decisions by all the drivers of individual vehicles and their interactions with each other combine to simulate the traffic flows on the road network.

Each driver decision has random elements, for instance a variation in route choice, the precise time at which the vehicle is released into the network and the place at which a lane choice decision is made. Each simulation run can start from a different random position and result in a different set of flows. While each separate run simulates the traffic behaviour in the model, only by analysing several runs is it possible to determine the average road network condition and its inherent variability. Multiple runs represent the variability of day to day traffic flows on a typical road network.

Statistics are gathered by the simulation as it runs, and can be used for subsequent analysis. They include:

- link and turn flows
- journey times between zones or along predetermined routes
- queue lengths
- emissions
- events such as lane changing

To ease the task of undertaking multiple runs, S-Paramics has a batch simulation mode in which the statistical outputs are collected without the real time visualisation. As a result the model runs much faster.

Analysis

Data Analysis Tool (DAT)

DAT is used to aggregate and compare data from multiple model runs.

Flows

DAT performs its aggregate analysis in terms of *partitions*, which may be defined as screen lines, cordons, or sets of links associated with a junction. A partition may also be a set of links for which the modeller has survey data. DAT analyses the results of multiple model runs and presents them in graphical form. Figure 10 shows a partition comparison between two model variants in terms of mean flows.

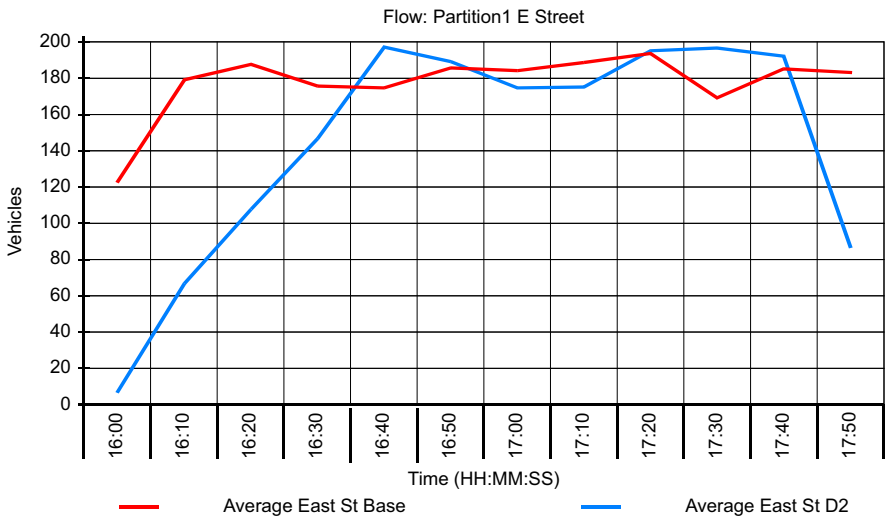

Figure 10: Flow comparison

Comparison of flows between models must include an analysis of the inherent random variance of traffic flow in order to assert if a real change has been made or if the changes are due to random sampling. DAT can be used to plot the mean and confidence intervals for model runs. Examination of the overlap between these intervals provides the modeller with an indication of the statistical significance of the change (Figure 11).

Aggregated flow information from several runs of different model variants can be graphically displayed and compared, as in the example shown in Figure 12.

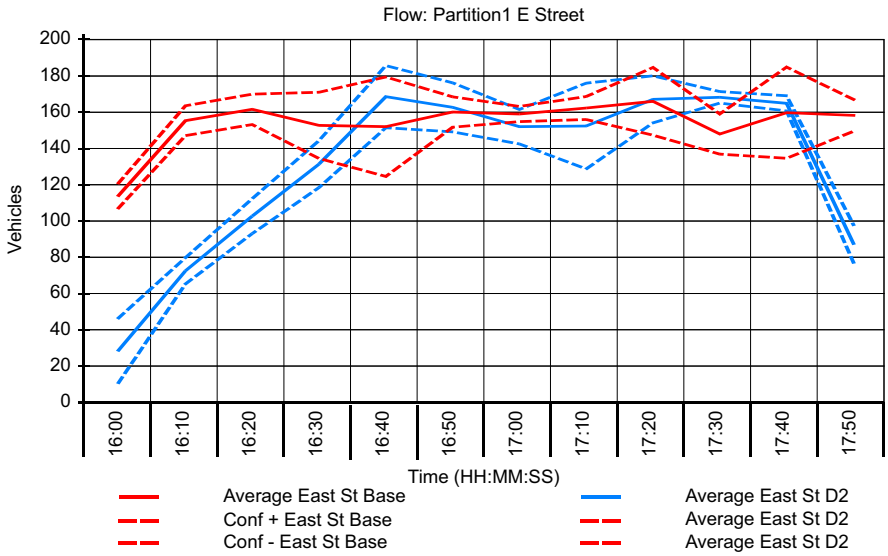

Figure 11: Flow comparison with confidence intervals

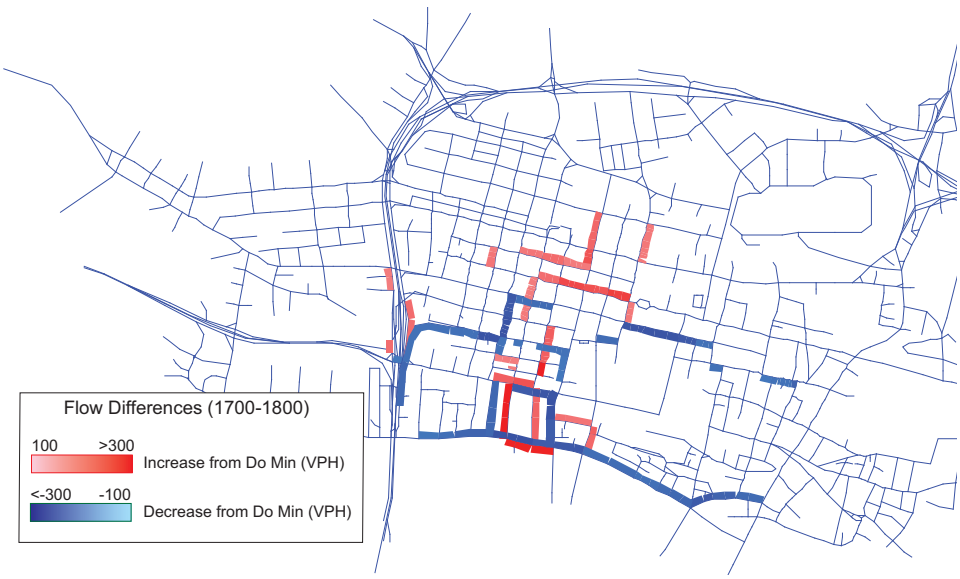

Figure 12: Flow differences

Queues

Queues are measured along paths or links specified by the modeller, who may select minimum, maximum and average queue length over the collection interval in terms of metres, vehicles or PCUs. Statistics are computed for the time at which queues form, and in terms of averages over a specified collection interval.

Figure 13 shows a queue path towards a junction, and the graph shows the maximum, mean and minimum queue length by time of day.

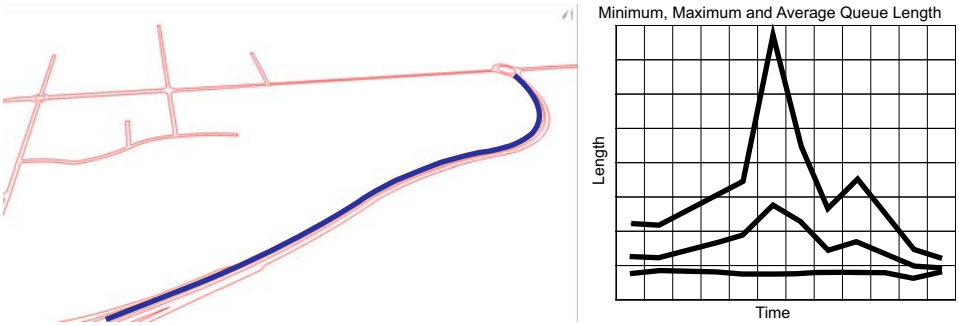

Figure 13: Queue length summary

Journey Times

S-Paramics measures journey times both for a whole trip from origin to destination and along a determined path in the model. DAT's statistical analysis of journey time comparisons between multiple model runs ensures that variability is properly taken into account.

Events

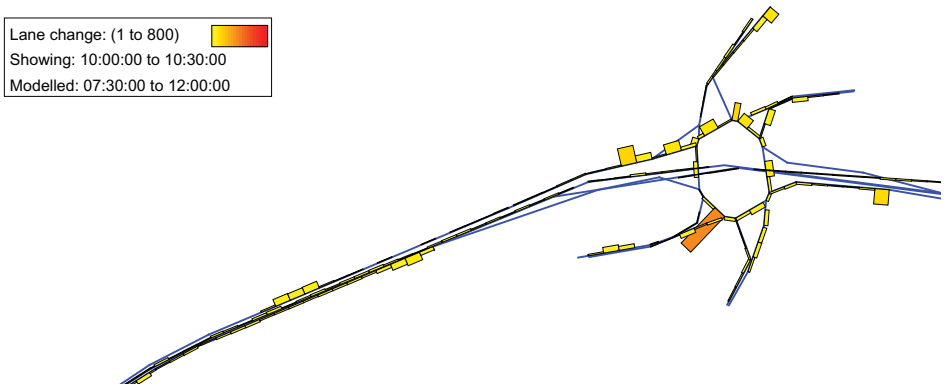

Figure 14: Event plot showing lane change events

S-Paramics logs events such as lane changes and overtaking manoeuvres, and DAT may be used to identify the locations of concern. Figure 14 displays lane change events plotted at 30m intervals, and clearly shows significant weaving sections.

Economic Assessment (PEARS)

PEARS is an acronym for Program for the Economic Assessment of Road Schemes. It has been developed for Transport Scotland as an economic assessment package for microsimulation models. The methodologies and costs are derived from the UK TAG Unit 3.5.6 - Values of Time and Operating Costs (www.webtag.org.uk).

PEARS undertakes trip-based assessments of changes in travel time costs and vehicle operating costs. The costs of a trip-based assessment are derived by aggregating the costs of each individually modelled vehicle in the network. Microsimulation enables a particularly robust economic assessment because it includes:

- a network definition which reflects real world detail
- a representation of the operating characteristics of different vehicle types
- travel demand profiles, which ensure a proper representation of traffic flow fluctuation
- detailed representation of real world phenomena embodied in platooning, overtaking and queueing
- an assessment of emissions and fuel consumption resulting from constantly changing vehicle speeds and accelerations

The output from PEARS is in the form of a set of standard tables of scheme costs and benefits aggregated over time and discounted to a base year, which assist the prioritisation of highway schemes.

Signal Control and ITS

Advanced Control Interface (ACI)

The ACI is a means of gathering data from simulated collection devices within an S-Paramics microsimulation model, and using the data to inform or control the simulated vehicle drivers. It was designed to enable adaptive signal control and ITS *hardware in the loop* to be linked directly to a computer running an S-Paramics model.

The ACI uses the SNMP (Simple Network Management Protocol) communications standard, which is in common use by traffic control systems engineers. The data exchanged by the ACI closely mimics those available to current urban traffic control (UTC) systems. Examples are:

- inductive loop data: speed flow and occupancy from traditional detector loops
- journey time data: flows and speeds along known paths - in effect a proxy for ANPR camera technology
- emissions data: the results of a roadside emissions detection device
- congestion detector: measurement of queue lengths
- car park occupancy: the number of spaces and the capacity of each car park

The actions implemented in the simulation model by the ACI are designed to closely follow those available to traffic control systems engineers.

Traffic Signals

Traffic signals may be controlled in the simulation by the time based controller in S-Paramics which sequences stages. Alternatively, signals may be controlled by a controller external to S-Paramics via the ACI.

An external controller may set stage times for a single use in the next cycle or for continued use over multiple cycles. Hurry calls for stages may terminate an existing stage and move to another, and stages may be run in variable order.

Intelligent Transport Systems (ITS)

Variable message signs (VMS) provide drivers with information as they pass specified locations. Broadcast devices are also available through the ACI which pass information to all drivers in the simulation or to an area subset. The information comprises 3 parts:

- the verbatim message received by drivers, for presentation purposes
- the formal interpretation which is used to influence driver behaviour
- response profile, i.e. a definition of which drivers will act on the information

VMS can be deployed to influence:

- speed, either as a target with the same variability as a normal road speed limit or as a mandatory maximum speed regardless of driver aggression levels
- headway modification
- lane restrictions
- changes in aggression and awareness
- warnings of congestion, to be included as additional delay in route calculation
- diversion instructions to divert via a specified waypoint
- car park advice, to change a destination to a specified car park

The response profiles are selected by:

- vehicle type, e.g. HGVs - selection may be inclusive or exclusive
- driver awareness and aggression levels - selection may be based on a greater than or less than comparison
- random factor or fixed percentage of vehicles
- vehicle destination zone or destination car park

A single message may have multiple formal interpretations and multiple response profiles. For example a message posted as "20 Minute Delay at New Town Junction" may be interpreted by the optimistic drivers as a 10 minute delay, by most as a 20 minute delay, and by the more pessimistic as a 30 minute delay. These delays will then be included in the route choice calculations for each driver.

ACI Examples

Automated Traffic Management

Automated Traffic Management came to prominence in the UK in reaction to a proposal in 2001 to widen the M25 London Orbital motorway to 14 lanes in its busiest sections. This option was considered to be untenable and as a result, more active control of motorways was planned.

The S-Paramics ATM controller uses the ACI to control motorway traffic using algorithms defined by the UK Highways Agency for flow control, incident detection, hard shoulder running, incident management and ramp metering. It may also be used to include the effect of ITS management tools on the normal day to day running of the motorway. It may also be used to test motorway incident management plans, including the effect of varying response times after an incident has occurred.

UTC Signal Control

S-Paramics has links to several Urban Traffic Control (UTC) systems. These enable the simulation to include coordinated adaptive signal control over a wide area.

Data from vehicle detectors in the simulation is relayed to the UTC system which aggregates the information from many detectors. The UTC system uses its own internal model of the road network and determines when signals are to change to optimise flows through the network. The instructions to the signals are relayed back to the simulation model where the stage changes are made.

S-Paramics currently has links to the following systems:

- SCOOT from TRL (Siemens version)
- SCATS from RTA in Australia
- MOVA from TRL
- CCOL from Vialis-TPA in The Netherlands
- Sentinel from Telent

Case Studies

Large Models

Plymouth

SIAS was commissioned by Plymouth City Council to develop an S-Paramics model of Plymouth and the surrounding area. The model includes 230 zones and 454km of road network and covers an area measuring 35km x 20km, as shown in Figure 15. It also includes 1300 bus stops within its description of every scheduled city bus service. Beyond the city boundaries it extends to Liskeard, Tavistock and Ivybridge.

The Plymouth model contains many of the most complex features to be found in other S-Paramics models, and some additional features unique to the city. For instance, the Torpoint ferry service has been simulated, to include both the full operation of the ferry service and the queueing and vehicle marshalling areas. The checkpoints at the naval dockyards are also included, where the effects of occasional periods of heightened security can be assessed for their impact on the rest of the road network.

With its S-Paramics model, Plymouth City Council has taken steps towards holistic transport impact assessment and traffic management. The model has been validated to a high standard to ensure that land developments within the entire Plymouth area can be assessed within the same unified framework. This enables transport planning priorities, phasing of road works, local and remote traffic impacts, and the most appropriate application of ameliorative measures to be assessed and applied.

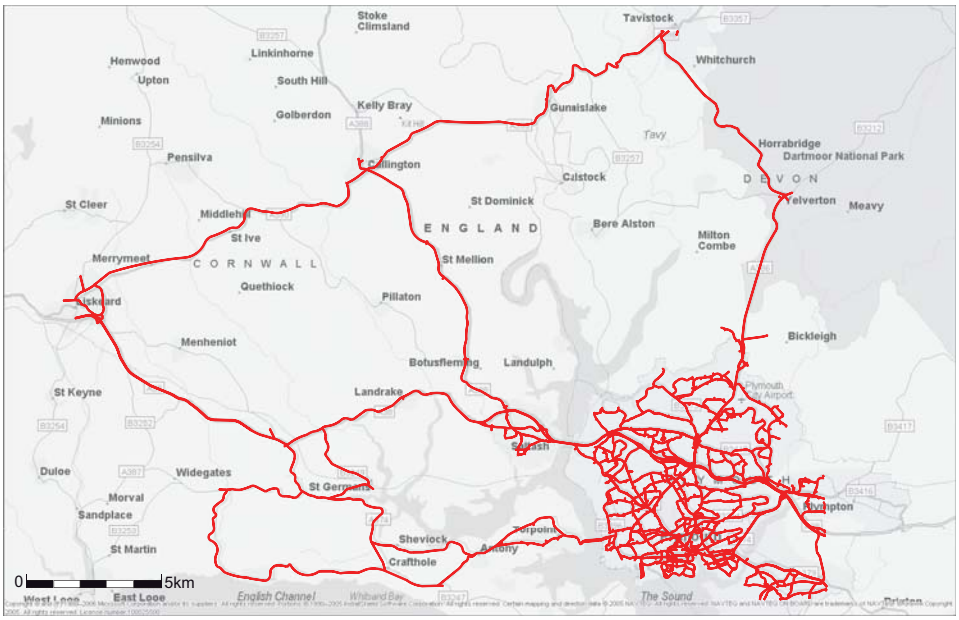

Figure 15: Plymouth model area

Chelmsford

The Chelmsford S-Paramics model, part of which is shown in Figure 16, is one of the largest microsimulation models in the UK, extending from the A120 in the north to the Dartford crossing in the south. The original purpose of the model was to test options for the north eastern bypass, but it has now been used to test the effect of various traffic management design options within Chelmsford, including the installation of pedestrian signalised crossings. More than 90,000 vehicles are represented in each of the weekday peak periods, with up to 12,000 vehicles simultaneously modelled. The model includes a directional flyover at the Army & Navy roundabout where the flow direction reverses at 14:30. It includes the highly successful Park & Ride site at Sandon, opened in April 2006 and now extended to 900 parking spaces.

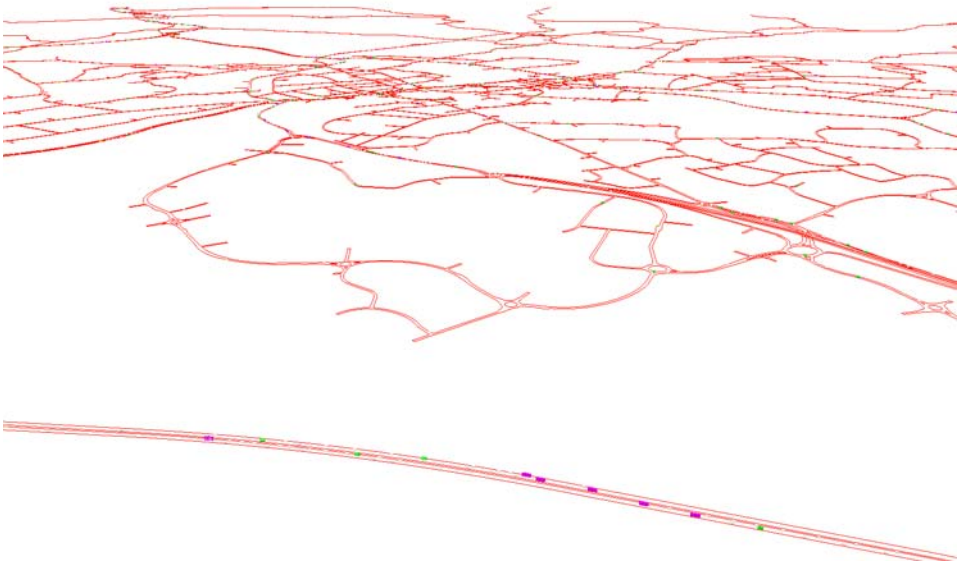

Figure 16: Chelmsford model area

Alkmaar

A wide area simulation model was commissioned in the Alkmaar region of northern Holland to assist the central coordination of a large number of projects and events. S-Paramics was used because of its capabilities for simulating temporary road conditions in a realistic manner over a wide area, such as lane closures, changes to junctions and dynamic route choice through ITS. The ability of microsimulation to record vehicle-specific travel times was required to establish police and fire service travel times during construction and maintenance periods.

The Alkmaar model, as shown on the back cover, provided detailed insight into the effects of road works and events. The forecast effects and the duration of road works are stored in a database. This data includes the re-routeing observed in the model, the initial mitigation measures planned for an event and any extra measures found to be required during testing. The database is used to assist in planning decisions and is continually updated and expanded.

UTC and ITS

Hampton Court Flower Show

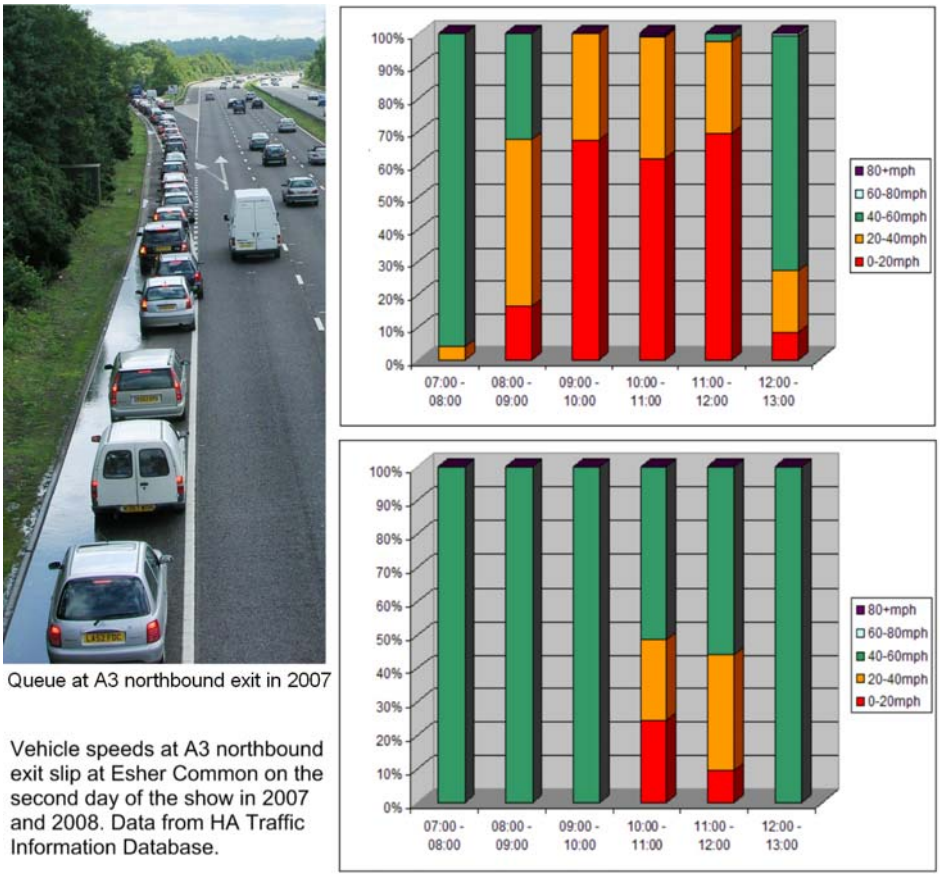

Figure 17: A3 at Esher (images courtesy of Surrey County Council)

Hampton Court Palace Flower Show is an annual event organised by the Royal Horticultural Society (RHS). Sited at Hampton Court Palace in the south west of Greater London, it is the world's largest annual flower show and takes place in early July over a period of 6 days. An average 3,500 vehicles per day visit the show, which causes difficulties because of the peaked nature of the traffic flow, at times generating a queue

extending from the nearby town of Esher to the junction with the A3 trunk road, as shown in Figure 17. This becomes critical when the queue extends onto the A3 main dual carriageway, resulting in erratic driver behaviour and dangerous driving conditions. In 2008 a set of SCOOT UTC strategies was developed by Surrey County Council to reduce the queuing. These were each tested using the S-Paramics model of Esher linked to Esher's UTC system.

A lot of careful planning was undertaken, which included monitoring and data collection during the 2007 show. The approach taken followed guidance in *A Best Practice Manual for Innovative UTC Schemes (Biora et al. 1995).* Four strategies were tested, and output from S-Paramics and SCOOT was used to evaluate the effectiveness of each change. The overall effectiveness of each strategy was evaluated using a composite measure including air quality and safety. The latter was determined by a proxy of how many vehicles were queueing at the A3 junction.

Surrey County Council concluded that one of the greatest advantages of testing strategies in a microsimulation environment was the ability to observe and quickly assess effects over a wide area. Strategies could be quickly eliminated or identified for further refinement and testing in a controlled environment. During the event the chosen strategy, a SCOOT technique involving the gating of traffic through Esher, met the safety and throughput objectives. A high correlation between observed flows and model predictions was demonstrated. During the event, the operation of the network was monitored closely by CCTV to see if the strategy needed adjusting, but no changes were required and no complaints from the public were received. Figure 17 shows the situation in 2007 with the queue extending on to the A3 trunk road. The graphs compare vehicle speeds at the top of the slip road in 2007 (top) and 2008 (bottom) and show the reduction in queueing in this critical region due to the changes in the UTC system.

M25

Figure 18: M25 Section 1 S-Paramics model (image courtesy of Hyder Consulting)

An S-Paramics model of the controlled section of the M25 between junctions 10 (A3) and 16 (M40) was developed to enable the testing of motorway control strategies. Figure 18 shows the model superimposed on an aerial photograph. This was done through modifications to the Motorway Incident Detection and Automatic Signalling (MIDAS) system that currently manages traffic on the M25.

The operation of MIDAS was incorporated into the model by linking it to an ATM controller. This links to overhead gantries and traffic detector loops in the simulation to replicate the MIDAS system through the generation of speed controls and automatic signalling in response to congestion and incidents. Changes to the MIDAS specification were tested and the results showed an improvement in journey times but more importantly an improvement in journey time reliability, a key performance indicator for the M25.

Car Park Guidance

The Dutch town of Nieuwegein is experiencing rapid growth and plans were drafted in 2008 to develop the city centre, involving an increase in city centre parking from 2759 places to 4669. An S-Paramics model was commissioned to investigate the effect of the increased flows into the town centre and the creation of a city centre pedestrian zone. This investigation included the design of an ITS methodology to direct vehicles to car parks on the basis of capacity and priority, thus reducing the volume of traffic within the town by removing the journeys between car parks.

The S-Paramics model includes multiple entries and exits from each car park. Car parks serve multiple origin and destination zones and central zones can be served by multiple car parks. Vehicles arriving at a car park are constrained to leave from the same car park on their return journey. The base model validation included comparisons between observed and modelled car park occupancies.

Figure 19: Car park occupancy

Figure 19 shows the effect of the ITS system on the car park occupancy near the end of the simulation run with and without the ITS guidance system. The image on the left is

without the ITS system and shows two car parks near capacity. In the image on the right, with the ITS system in operation, the vehicles are more evenly distributed and spaces are available in all car parks. There is an overall benefit to drivers, both observing and not observing the advice from the ITS car park advisor, in a reduction in the need to hunt for a parking space.

Road Design Studies

Overtaking Study

On a single carriageway, a number of vehicles may be slowed by a single slow lead vehicle, typically an HGV or an agricultural vehicle. Consequently, vehicles will form a platoon behind the lead vehicle for as long as they are unable to pass it. If an improvement scheme includes carriageway widening to facilitate dedicated or opportunistic overtaking, then the benefit of the scheme will be felt by the vehicles that can now pass the lead vehicle, an effect known as platoon dispersion. This benefit is not felt just over the length of the scheme but for some distance downstream.

S-Paramics includes an overtaking model in which a driver will first assess its desire to overtake based on its target speed with respect to the speed of the vehicle ahead. It will then assess its ability to overtake based on the gap available ahead of the vehicle to be overtaken and also the visibility of the road ahead and absence of an oncoming vehicle in that space. The overtaking manoeuvre is initiated when both desire and gaps are present.

Transport Scotland, in association with DfT and HA, commissioned a study to assist with the preparation of a new technical advice note for provision of overtaking lanes. The study examined platoon formation, the effectiveness of different lengths of overtaking lane, and the extent of the downstream benefits.

An S-Paramics model was developed to replicate a 15km section of single carriageway and was coded to include three sections; an approach length, a central study section, and a run-out section. The run-out, or downstream, section of 10km allowed the model to measure the downstream benefits. Calibration included adjusting the overtaking parameters such that the change in vehicle order in the model matched that measured from a set of number plate matching surveys. The model was used to investigate the effect of adding WS2+1 sections (wide single carriageway, 2 lanes plus 1 overtaking lane) and comparing this with existing opposed overtaking.

The graphs in Figure 21 show the average speed of vehicles under three different flow rates, with and without the overtaking section. The results indicate that:

- at 10,000 AADT, the benefit in increased vehicle speed is accrued for around 5km downstream
- at 5,000 AADT, the benefit is of shorter duration as the overtaking demand is satisfied earlier
- at 15,000 AADT, the benefit is again of short duration as overtaking vehicles soon catch up on next platoon

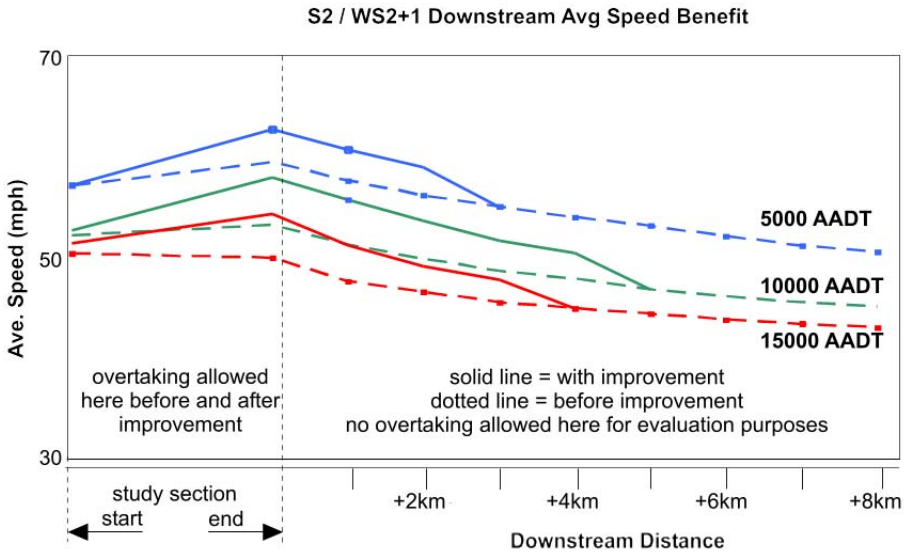

Figure 20: Overtaking downstream benefits

S-Paramics is able to capture the economic benefits derived from overtaking lanes by detailed simulation of individual vehicle behaviour, specifically by varying the aggression component of driver behaviour which leads to speed differentials and overtaking. Platoon formation is implicit under such circumstances and essential to the analysis of the economic benefits of the provision of overtaking sections.