SCOOT MMX (SCOOT Multi Modal 2010)

David Bretherton
Principal Consultant, TRL
Crowthorne House, Nine Mile Ride, Wokingham, Berks, RG40 3GA
Tel: 01344 770733, Fax: 01344 770356, dbretherton@trl.co.uk

Mark Bodger
Systems Marketing Manager, Siemens Mobility, Traffic Solutions
Sopers Lane, Poole, Dorset, BH17 7ER
01202 782848, mark.bodger@siemens.com

Jason Robinson
ITS Chief Engineer, Transport for London,
Palestra Building, 197 Blackfriars Road, London, SE1 8NJ
020 3054 2640, jason.robinson@tfl.gov.uk

Ian Skeoch
Software Manager, Peek Traffic Ltd
Hazelwood House, Lime Tree Way, Chineham Business Park, Basingstoke, RG24 8WZ
01256 338301, Ian.Skeoch@peek-traffic.co.uk

Helen Gibson
Researcher, TRL
Crowthorne House, Nine Mile Ride, Wokingham, Berks, RG40 3GA
Tel: 01344 770585, Fax: 01344 770356, hgibson@trl.co.uk

Gavin Jackman
Head of Software, TRL
Crowthorne House, Nine Mile Ride, Wokingham, Berks, RG40 3GA
Tel: 01344 770411, Fax: 01344 770356, gjackman@trl.co.uk

ABSTRACT

The SCOOT Urban Traffic Control system is now operating in over 250 cities and towns worldwide. Since the first system was installed there has been a continuous program of research and development to provide new facilities which take into account new technology and policies, and meet the multimodal requirements of the traffic manager. Recent releases have incorporated many new features aimed at providing priority to public transport and in the last release SCOOT MC3 new strategies to provide benefit to pedestrians at pedestrian crossings were included. In this new version, SCOOT MMX, the multi modal theme is continued with additional facilities to prioritise pedestrians at junctions, a significant update and enhancement of emissions estimates as well as features to improve operation during low flow periods. The latter includes cycle time independence, allowing individual junctions to operate at their optimal cycle time when it is

more beneficial than being coordinated and the introduction of ‘ghost’ stages to reduce cycle times when demand dependant stages are infrequently called.

PAPER

The SCOOT Urban Traffic Control system is now operating in over 250 cities and towns worldwide. Since the first system was installed there has been a continuous program of research and development to provide new facilities which take into account new technology and policies, and meet the multimodal requirements of the traffic manager. Recent releases have incorporated many new features aimed at providing priority to public transport and in the last release, SCOOT MC3 (1), new strategies to provide benefit to pedestrians at pedestrian crossings were included. In this new version SCOOT MMX the multi modal theme is continued with additional facilities to prioritise pedestrians at junctions as well as features to improve operation during low flow periods. This paper focuses on the following new features which have been incorporated into SCOOT MMX:

- Pedestrians: variable invitation to cross
- Cycle time sub-region independence
- Reduced cycle time at quiet times
- Updated and enhanced emissions estimates

PEDESTRIANS: VARIABLE INVITATION TO CROSS

In the last release SCOOT MC3 Service Pack 1 a facility was included that allowed users to give greater priority to pedestrians at signalised pedestrian crossings (mid-block crossings). For SCOOT MMX, a new facility targeted primarily at traffic junctions with pedestrian crossing facilities has been developed. In particular it enables priority to be given to pedestrians at sites where the pedestrian demand is high. Where there are large numbers of pedestrians waiting to cross the green man invitation period (and hence the overall time available to pedestrians to cross) can now be extended.

BACKGROUND

SCOOT is generally based on improving road based traffic performance. The region cycle time is determined through a philosophy of operating the most heavily saturated node in a coordinated network at no greater than a preset target degree of saturation. This target degree of saturation parameter which is normally 90% is determined by the user using the ‘isat’ value. Setting ‘isat’ to 90 ensures that, if possible, the most heavily saturated node in the coordinated region operates at no greater than 90% saturation. When flow levels increase, the cycle time throughout the network will rise trying to satisfy this criterion until the ‘isat’ limit is reached. Pedestrians have not been considered at all in the calculations, the time they receive as an invitation to cross period or the duration they have to wait is an outcome of road based traffic performance. As the cycle time increases in response to greater traffic demand there will be longer intervals between successive pedestrian stages and hence increased pedestrian waiting times.

NEW DEVELOPMENT

The objective of the new development has been to modify SCOOT to enable the length of the pedestrian invitation to cross period to be increased in response to higher pedestrian demand, without increasing the region cycle time. In particular it enables priority to be given to pedestrians at sites where the pedestrian demand is high.

Care was required to ensure that the longer pedestrian period does not result in longer cycle times due to the reduced vehicle green times. Longer cycle times will result in longer waiting times for pedestrians and this would defeat the objective. To ensure that the cycle time is not affected, an increase in the pedestrian green time is coupled with an increase in the ‘isat’ value. This will have the effect of potentially increasing the delay to vehicles on the busiest links of the junction.

Users are able to set up to 4 priority levels. For each pedestrian priority level users can specify the increase in pedestrian green time (in seconds) and the ‘isat’ value to which the vehicle degree of saturation at the junction is allowed to rise.

MEASURING PEDESTRIAN DEMAND AND CAPACITY

Pedestrian capacity has at least two elements, the capacity of the waiting area and the effective capacity of the crossing itself. At present local authorities do not have detectors installed which can determine the demand in either of these areas.

One measure that is currently available is the knowledge of when a push button is first pressed to demand a pedestrian stage or phase. The rolling average, or other smoothed measure, of the time relative to the end of pedestrian green can be used as a proxy for pedestrian demand. It does not provide a measure for a particular signal cycle, but provides an average level over a longer period. This measure has been included as part of this development and can be used to automatically determine when to increase the length of the pedestrian variable invitation to cross period. Its use is optional.

When detectors are installed that can count pedestrians it should be possible to provide a measure of demand relative to capacity. However, the measure of pedestrian demand will need careful consideration. Some pedestrians cross the road against the traffic signals, therefore, they would not be included in any calculation of the time required to service those who do cross in the green man. However, such pedestrians should not be ignored. The number crossing against the signals could be taken as a measure of failure to meet the requirements of pedestrians. At least they represent some measure of vehicle green that is not fully utilised.

BENEFITS

The main benefit of the development is that it will enable a pedestrian priority policy to be implemented. At sites where pedestrian demand exceeds current pedestrian capacity there should be an improvement in safety. Monetary benefits are difficult to quantify, but the new development should ensure that pedestrian friendly policies are implemented in an efficient and controlled way.

CYCLE TIME SUB-REGION INDEPENDENCE

Under normal operation, SCOOT maintains coordination between nodes in a region. In order to do this, the cycle time of all the nodes must be the same, or strictly half in the case of double cycling. At quiet times, however, the benefits of coordination can be reduced. If one node in the region is experiencing heavy traffic it will increase the cycle time of the whole region, increasing the delay at otherwise lightly loaded junctions. These lightly loaded junctions would benefit from a lower cycle time, and there would be an overall region benefit if the cycle time of these junctions was lower than the other junctions, as long as the disbenefits due to lack of coordination do not outweigh this benefit.

This new development ‘cycle time sub-region independence’ makes it possible for nodes to operate at a cycle time that is independent of the regional cycle time. Whilst coordination will be lost, operating at lower cycle time will reduce waiting times at red signals. The mechanism is set up within SCOOT to enable the independence, and triggers and thresholds established such that the process can also occur automatically if desired.

Often a particular set of nodes work best if they have the same cycle time irrespective of other conditions. Therefore the concept of a sub-region has been established. A sub-region is a set of nodes that will always be considered for cycle time independence together. None of the nodes within the sub-region can operate at a different cycle time from the others. If the sub-region is selected to operate independently from the region, the internal links within the sub-region will still be coordinated.

OVERVIEW OF OPERATION

If a node is allowed to operate independently, it will operate at its minimum practical cycle time (MPCY). A sub-region will operate at the highest minimum practical cycle time of the nodes in the sub-region.

Whether or not a node’s or sub-region’s cycle time is independent can be either user-defined (using operator commands or timetables) or automatic.

If it is automatic, the decision as to whether a node’s or sub-region’s cycle time may operate independently is taken on the basis of a comparison between the estimated delays when independent and when operating at the regional cycle time (non-independent).

DELAY CALCULATION

SCOOT determines whether the node or sub-region should cycle independently or not every five minutes. This is done through a trigger based on the ratio of the predicted delays calculated at the regional cycle time and at the independent cycle time.

The delays can be calculated by two methods: complex and simple. The complex method calculates the delay directly from the flow profiles, providing accurate estimates of delay, but for uncoordinated delays this involves repeating the calculation at every offset and then taking the average which is very processor intensive. The simple method estimates the delay through a single calculation based on average flows, and reduces the resultant delay if the link is coordinated. The simple calculation is much less processor-intensive and is provided as an alternative to the complex calculation if processor power is an issue. However it is less reliable as the benefit of coordination may be site and/or time specific.

The complex delay calculation

SCOOT has four link types: *entry*, *filter*, *exit* and *normal*. When calculating the delays experienced on links whilst independent and non-independent, only *entry* and *normal* links need to be considered.

In order to do the delay calculations, links are categorised into *entry*, *upstream*, *downstream*, *internal* or *ignored*.

The correspondence between link types and these new categorisations is as follows:

Link type	Link category
entry	*entry*
filter	*ignored*
exit	*ignored*
normal	*entry*, *upstream*, *downstream*, or *internal*.

The categorization of *normal* links is described with reference to the example shown in Figure 1. Region R123 is shown as containing the sub-region R523 and a number of individual nodes.

Links of type *normal* are categorised as follows:

If the calculation is taking place for a sub-region R523, the links are categorised as:

entry	if its down node is in the region R123 and its up node is not.
upstream	if its down node is the sub-region R523 and its up node is not.
downstream	if its up node is in the sub-region R523 and its down node is not.
internal	if both its up node and its down node are in the sub-region R523.

If the calculation is taking place for a node N11, the links are categorised as:

entry	if its down node is in the region R123 and its up node is not.
upstream	if the node N11 is the link's down node
downstream	if the node N11 is the link's up node

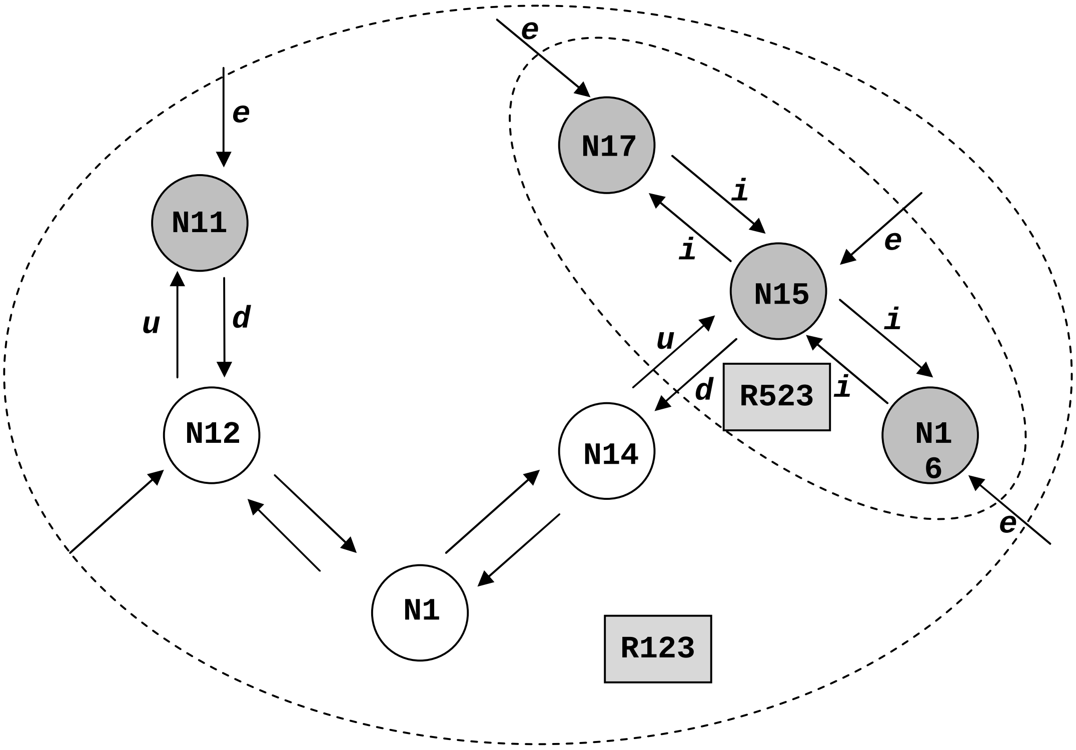

Fig. 1 Categorisation of links

When a node or sub-region is non-independent, *upstream*, *downstream* and *internal* links will all be coordinated. These are marked by u, d and i respectively in figure 1. *upstream* and *downstream* links are categorised relative to the node or sub-region that is being considered for independence, in this case node N11 and sub-region R523. The delay taken for those links is simply the current delay, calculated using the current signal timings and flow profile. Uncoordinated links, i.e. *entry* links (marked by e in figure 1), are represented by the average delay.

These delays are then compared against an estimation of the delay whilst independent. The estimation for links that become uncoordinated (*upstream* and *downstream* links) when independent is made by adjusting signal timings and flow profiles to the new cycle time and performing the calculation at every offset to get the average delay. The process for *internal* links, however, involves predicting what the coordinated delay will be at a different cycle time. (See below). If, however, the node or sub-region is already independent, the delay used is the average delay for uncoordinated links (*upstream* and *downstream* and entry links). But when estimating what the delay will be when *upstream* and *downstream* links become coordinated (or indeed, *internal* links remain coordinated) in the non-independent regime, it is not possible to directly calculate what the delay will be when they are coordinated if they are currently under a different cycle time. Therefore a method of predicting the delay under these circumstances has been developed.

Predicting delay under a different cycle time
The method involves calculating the delay on the coordinated links when the node or sub-region is non-independent (D_{cur}) and comparing this with the minimum delay (D_{min}) and the average delay (D_{ave}) across all possible offsets. These values are used to calculate a factor:
The factor is the ratio $f = (\Sigma D_{ave} - \Sigma D_{cur}) / (\Sigma D_{ave} - \Sigma D_{min})$

This is recalculated every time the complex calculation takes place, and each new result is smoothed by a percentage of the previous result as specified by the input variable "factor smoothing factor" (FASM) as follows:
$SF = SF. (FASM/100) + F . (1 - FASM/100)$

This smoothed factor, SF, is then utilised when the node or sub-region is independent to estimate the delay if the node or sub-region were to become non-independent again:
$D_{cur} = D_{ave} - SF. (D_{ave} - D_{min})$

Internal links in sub-regions are also coordinated when independent. Therefore a second factor is required that is calculated whilst the sub-region is independent (only summing the delays on the internal links) and utilised whilst non-independent to estimate the delay whilst independent.

The simple delay calculation
As already mentioned, the simple calculation is based on average flows. A calculation is performed that assumes uniform flows. However this would overpredict the delays on coordinated links. Therefore, on coordinated links the delay is reduced. How much it is reduced by depends on the mean modulus deviation (MMD), which is a measure of how platooned the traffic is. If MMD/flow = 0 the traffic is not at all platooned. In contrast, if MMD/flow = 2 the traffic is totally platooned, and could expect zero delay if the green stage came at the right time. However, as most stages serve vehicles in both directions, it is unlikely that all links will experience zero delay.

Therefore the user input variable "Simple Delay Factor" (SIDF) was introduced. This is used to limit the reduction in delay due to platooning. The graph in figure 2 shows the percentage of the delay assuming uniform flows that is taken as the coordinated delay. If SIDF is 100, there is no reduction in delay due to coordination. On the other hand, if SIDF is 0, fully platooned traffic will receive zero delay.

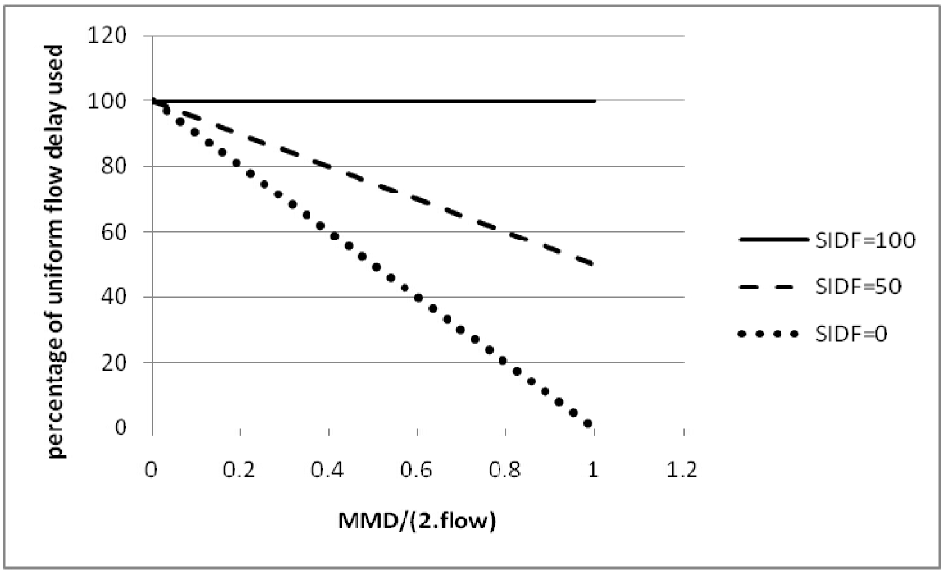

Fig. 2 Impact of SIDF

The relationship is as follows:
‘coordinated’ delay = delay (1 – MMD (1 – SIDF/100)/(2 . flow))

If SIDF is 100 then the delay is identical whether the link is coordinated or not. The lower the value of SIDF, the greater the assumed benefit of coordination and the more the delay is reduced.

THRESHOLDS FOR INDEPENDENCE

The delay for the node (or sub-region) is estimated for when the node (or sub-region) is independent (DI) and also for when the node or sub-region is non-independent (DN). The value of the smoothed delay ratio (DN/DI) determines the independence status.

If a node or sub-region is cycling at the region cycle time, the delay ratio is compared against the percentage input by the user in “independence threshold”. The recommended setting for this is 1.15. If the delay ratio is more than “independence threshold”, the node or sub-region will begin cycling independently.

If the node or sub-region is cycling independently the delay ratio is compared against the percentage input by the user in “non-independence threshold”. The recommended setting for this is 0.95. If the delay ratio is less than “non-independence threshold”, the node or sub-region will begin cycling again at the regional cycle time.

MONETARY BENEFIT OF CYCLE TIME INDEPENDENCE

Prior to its installation in London an attempt has been made to estimate the likely monetary benefit that the new facility will produce.

The criterion for determining when a node or sub-region should become independent is that the delay ratio (DN/DI) exceeds 1.15. When the node becomes independent this corresponds to a reduction in delay of (100*(1 – DI/DN)) which is about 13%. The criterion for determining when a node or sub-region should lose independence is that the delay ratio (DN/DI) is less than 0.95, so there will be some time that the node or sub-region is disadvantageously independent. A reasonable estimate of the benefits of independence is that on average there will be a reduction in delay of at least 12% at the node (sub-region) during the periods of independence.

Data was collected from the London system during the 24 hour period. The data was analysed and for 6 regions the period for which the cycle time was consistently below the maximum (and therefore flows can be assumed to be relatively ‘low’) was identified. From this the annual cost of delay per node in each region during the period when the cycle time was consistently below the maximum was calculated. A cost of £12 per vehicle hour delay was assumed. The results varied considerably from region to region. Averaging over the 6 regions the cost of delay during the period when the cycle time was consistently below the maximum and cycle time independence is potentially of benefit was £257,000 per node per annum.

Assuming a reduction in delay of 12% over this period would give an annual cost saving of approximately £30,000 per junction.

If on average there are 115 junctions (5% of the 2300 SCOOT controlled junctions in London) independent during the ‘low’ flow periods then the annual cost savings would be £3,450,000.

REDUCED CYCLE TIME AT QUIET TIMES

At quiet times it is usually desirable to keep the cycle time as low as possible. In some regions, however, the cycle time is kept higher than necessary because of the minimum cycle time constraint at one of the junctions in the region.

The minimum cycle time of junctions is calculated such as to allow all stages to run for their minimum time. This makes sense where all stages are permanent or where demand dependent stages are frequently called. However, at quiet times when demand dependent stages e.g. pedestrian stages are called infrequently this can result in the cycle time being higher than desired as the minimum cycle time must cater for the pedestrian stage being called.

In SCOOT MMX, logic has been developed to overcome this problem. The concept of a “ghost” stage has been introduced which is defined as a demand dependent stage that has effectively been omitted from contributing to the node minimum cycle time. The user now has the option of allowing demand depend dependent stages to be treated as ghost stages at chosen times of day. The effect of this is that at quiet times the minimum cycle time of affected nodes will be reduced. In many regions this should result in a lower region cycle time.

The development allows some nodes to operate at a cycle time that is only long enough to accommodate their non-demand dependent stages (plus 4 seconds to enable the offset optimiser to make an advance change). Whilst over the region this will be greatly beneficial, when a demand dependent stage is run at such a node, the cycle time at the node will exceed the SCOOT region cycle time and it will become out of step with the other nodes in the region. In practice the running of the demand dependent stage will have created a bad offset at the node. The offset will be displaced by the length of the demand dependent stage.

When the offset optimiser next runs, it will recognize the bad offset caused by the occurrence of the demand dependent stage and make an appropriate change. Because of the large change in offset that has been imposed, the offset optimiser would be expected to make a change to return towards the "good" (pre-existing) offset at the nodes at both ends of affected links. Assuming a normal offset authority of 4 seconds, that would be a change of 8 seconds towards a good offset. If the optimiser does not make such changes, then the offset on that link is not important and the occurrence of the demand dependent stage will not have caused much extra delay. In subsequent cycles the offset will not be as bad and it will be safer to assume only one offset change per link per cycle towards restoring the pre-existing offset.

Assuming a stage length of 20 seconds and a 4 second offset authority, the offset will be expected to be fully restored after 4 cycles (8 seconds restored in the first cycle and 4 seconds in each of the three subsequent cycles). For a 16 second demand dependent stage the offset would be restored after 3 cycles.

The calculations above are for the worst case; the difference between the length of the demand dependent stage and the cycle time may mean that the disruption is less than in the simple worst case calculation.

This new logic is only expected to be active at quiet times, when cycle times would be low, say between 32 and 60s. Therefore, the effect of the occurrence of a demand dependent stage would not be expected to last for more than a few minutes. However, frequent occurrences of the demand dependent stage will have the potential to cause significant disruption and it is probably not appropriate to ignore a frequently occurring stage when calculating the node minimum cycle time.

To prevent this happening, the frequency of occurrence is assessed over a period of 15 cycles. The demand dependent stage is included in a node's minimum cycle time if the stage has occurred 4 or more times in the last 15 cycles. Similarly, it is removed from the calculation of minimum cycle time, if removal is allowed at this time, when it has not occurred more than once in the last 15 cycles.

MONETARY BENEFIT OF REDUCED CYCLE TIME AT QUIET TIMES

The benefits of the development will vary from region to region depending primarily on how much the region cycle time is reduced. Prior to its introduction a rough estimate has been made of the likely benefits in London.

In low flow conditions when junctions are well below saturation the delay at the junction is proportional to the cycle time. If the cycle time c is reduced by x seconds then the delay d at the junction will be reduced by: dx/c.

Similarly, if the cycle time c of the region is reduced x seconds then the delay D in the region will be reduced by: Dx/c.

Data was collected from the London SCOOT UTC system during a 24 hour period. The data was analysed and for 6 regions the following information derived:

- current minimum cycle time
- period for which minimum cycle is in operation
- delay in region during period for which minimum cycle is in operation
- delay over the 24 hour period has been derived

From this the annual cost of the delay per node in each region and the annual cost of delay per node in each region during the period when the cycle time is on its minimum were calculated. A cost of £12 per vehicle hour delay was assumed. The results varied considerably from region to region. Averaging over the 6 regions the cost of delay during the period when the cycle time is on a minimum is £53,500 per node per annum.

Assuming that on average the reduced cycle time facility allows the minimum cycle time to be reduced by 10%, the annual cost saving per node per annum will be approximately £5,000.

Over the approximately 2300 nodes controlled by SCOOT in the London UTC system the cost saving in terms of the delay to vehicles is estimated to be £11,500,000.

UPDATED EMISSION ESTIMATES

Emissions from road traffic are an important source of air pollution in urban areas. Estimates of emissions from vehicles in a SCOOT controlled area were included in Version 4 based on research undertaken in 1995 to 1996. Considerable changes have taken place since 1995:

- A better understanding of the emissions performance of different classes of vehicles as a function of driving conditions through the major EU research project ARTEMIS and work funded by DfT.
- Changes in the performance of vehicles as new regulations EUII, EUIII, EUIV have been introduced.
- Changes in the mix of vehicles on the road, an increasing proportion of cleaner vehicles as new vehicles meeting stricter emissions standards are introduced and old dirty vehicles are retired. Recently the large increase in the popularity of diesel powered cars has changed the emissions pattern.

The emissions estimates have now been updated to provide more realistic estimates of:

- Carbon monoxide
- Hydrocarbons
- Oxides of nitrogen
- Particulate matter
- Carbon dioxide.

In addition to these pollutants that were originally estimated, two new estimates are provided:

- Total carbon
- Fuel consumption

Fuel consumption is estimated from the total carbon emission using the carbon balance approach.

No new infrastructure is required to enable the estimation of emissions—all the data is taken from the normal SCOOT detection and processed in the SCOOT model. The emissions information can be stored and processed in the ASTRID database with the other traffic data, thus providing an historic record as well as current estimate. The availability of improved information on vehicle emissions is useful for general monitoring as well as being the first stage in developing methods to reduce the pollution.

In general there are two stages to modelling air quality, or pollutant concentrations in an area. The first stage is to estimate the emission of pollutants from various sources. The second stage is to model how the emitted pollutants will disperse and chemically react to predict the temporal and spatial variation in pollutant concentrations. Dispersion models require information on the local geographic conditions, and details of the roadside built environment as well as knowledge of the meteorological conditions. It must be stressed that the output from SCOOT relates only to the first stage of the air quality modelling process.

Consideration of the availability of data from SCOOT and corresponding emissions data has led to the development of a model based on the following data:

- The total flow (in vehicles) on each link. This is determined in real time from SCOOT data.

- The mean speed (averaged across all vehicles) on each link. This is determined in real time from SCOOT data. This average speed is an unweighted mean of speeds of all vehicles and includes time spent stationary in queues. The average speed of each vehicle is calculated by dividing the link length by the link traversal time, which is in turn derived from SCOOT delay data.

- The proportions of seven classes of vehicle on each link: petrol cars, diesel cars, petrol large goods vehicles, diesel large goods vehicles, other goods vehicles 1, other goods vehicles 2 and buses. Default values are provided for these quantities based on a standard set of proportions from the national fleet mix, but the operator may specify different proportions on a link by link basis if preferred. It is likely that for most links the default set of proportions will be adequate, the obvious exception being streets open only to particular vehicles such as buses and taxis.

- Distance travelled. Allowance is made for the fact that the SCOOT data only covers the lengths of road between in-road loops and their associated downstream stoplines. The emissions produced by traffic between the upstream stoplines and the loops is allowed for by factoring up the results for the extra distance travelled. SCOOT loops are sited in the optimum position for traffic control, usually 10 to 15 meters downstream from the previous junction, so only a small proportion of each road section is not included in the SCOOT model.

EMISSIONS MODEL

Emissions of carbon dioxide (CO_2), carbon monoxide (CO), oxides of nitrogen (NO_x), volatile organic compounds (VOC) and particulates (PM) are estimated using emission-rate data in the form of a set of functions linking vehicle speed to emission rate. For each vehicle class and each pollutant, there is an emission function in the form of a polynomial function of speed. Each of the emission functions produces valid results for the entire range of average speeds encountered on SCOOT-controlled links. The coefficients which specify the functions are based on complete urban driving cycles, each with a characteristic average speed. These coefficients have been generated from a large database of emissions data.

CONCLUSIONS

SCOOT MMX continues the multi modal theme of recent releases with additional facilities to prioritise pedestrians at junctions. The main benefit of the development is that it will enable a pedestrian priority policy to be implemented and at sites where pedestrian demand exceeds current pedestrian capacity there should be an improvement in safety.

New features have been introduced to improve operation during low flow periods. This includes cycle time independence, allowing individual junctions to operate at their optimal cycle time when it is more beneficial than being coordinated and the introduction of ‘ghost’ stages to reduce cycle times when demand dependant stages are infrequently called.

There has been a significant update and enhancement of emissions estimates which now includes estimates of fuel consumption and total carbon.

The system is currently being installed in London where a preliminary evaluation indicates substantial benefits in terms of delay and cost savings.

REFERENCES

(1) Bretherton, D, Bodger, M and Cowling, J. ‘SCOOT – Managing Congestion, Communications and Control’ in Proceedings of 12th World Congress on ITS, San Francisco, 6-10 November 2005