A distributed strategy for eliminating incident-based traffic jams from urban networks

by Penina Roberg, *Road Traffic Research Centre, Middlesex University**

New techniques are needed to deal with traffic jams that have formed in urban areas. As part of a study supported by EPSRC, the author has investigated dynamic strategies for controlling and dispersing jams that originate from temporary obstructions at particular locations in an idealised network. The strategies involve the application of bans to a number of critical junctions in the network. The bans come in two forms: turn or ahead. Turn bans are imposed on selected links to break gridlock cores at the nucleus of the jam. Ahead bans are implemented around the envelope of the traffic jam to reduce input into critical sections of the road. Under appropriate circumstances, certain combinations of these control measures have been found to have an appreciable impact on jam development and dispersion.

In this paper, the effects of these core and gating strategies are investigated with the aid of a simulation model. A qualitative explanation of the performance and effect of the strategies under a variety of road network conditions is put forward. A quantitative estimate of the delays incurred over the period of control is also provided.

The paper discusses the application of the ideas presented here to the two-way idealised network and concludes with a brief review of future research directions aimed at applying the experience gained to real networks.

1. INTRODUCTION

The form of an incident-based traffic jam in a simple, idealised, one-way road network is characterised by irreversible cores or knots which develop at specific locations within the jam area. These knots persist even when traffic demand falls away at the end of the peak period. External measures are required to break the interlocking queues apart in order to restart vehicular movement.

Previous research[1] has highlighted two approaches for the successful control of traffic queues. The first approach considers how features in the road layout, such as the allocation of queue storage space between ahead and turning vehicles, can be exploited so that the rate of jam expansion may be curbed. The second approach suggests responsive measures which can alleviate congested situations which have already assumed a gridlock characteristic.

There are several ways to deter the spread of such jams and eliminate them once they have formed. One such way is via the vehicle movement ban, so that movement is restricted to ahead only or turn only. Two variations of the ahead ban have been investigated:

(1) vehicles are barred from entering the congested region (gating); and
(2) vehicles moving into the congested region are forced onto nearby routes (rerouteing).

In real life, these measures take the form of a set of rules which channel vehicles away from the sensitive locations of the network. Although these rules may temporarily restrict the motorists' freedom, they can easily be applied through traffic signals.

2. THE TRAFFIC SIMULATION MODEL

The construction of a simple, idealised simulation model may lead to the understanding of some of the fundamental characteristics of traffic-jam growth as well as to the complementary process of decay. The model used in the context of this research is simple, but the principles involved in its formulation have already been extended to incorporate more realistic network features. Here, the salient features of the model are introduced along with specific details relating to the experiment described later in the paper.

2.1. The environment

The network consists of equally-spaced streets at right-angles. It forms a grid-like structure typical of town centres, for example Manhattan in New York. A sequence of alternating sources and sinks is constructed at the edges of the square grid. Vehicles travel from any source towards any sink.

Generation of vehicles from each of the sources is determined by a Poisson model reflecting demand for access to the network of roads. The expected mean of the demand at each access point is described by μ, which has the same value for all sources to create a steady flow of vehicles through the system as a whole.

The mechanism of vehicle movement mimics simple traffic control. Vehicles are partially segregated into traffic lanes and progress, subject to:

- — space being available downstream;
- — capacity constraints at junctions; and
- — interactions between queueing vehicles.

2.2. Vehicle progression

Vehicles can either go straight ahead or turn in prescribed directions. The propensity for a vehicle to turn is controlled by a fixed probability value which is the same for each intersection and is used to determine the proportion of turning vehicles at each intersection. This proportion is subject to numerical rounding to ensure that vehicular flow is always integer.

Queues are maintained on all the network's links. The size of each individual queue is reviewed each time-slice and updated according to the movement of vehicles away downstream. Vehicles requiring access to a link on which a queue has already formed may join the queue providing there is space to accommodate them. If the link becomes full, vehicles may no longer enter and must wait in the link upstream until space appears in the blocked link. This space becomes available after vehicles begin to discharge across the downstream stop-line.

2.3. Spillback and queue propagation

The propagation of queues throughout a road network can be described in relatively simple terms[2]. The pattern of development focuses on the consequence of any single link becoming blocked. The link can block as a result of an incident on the network which causes queues to evolve around the original source of obstruction. The interaction between the individual queues on the link governs the nature of the spillback mechanism involved. More details of the process of queue propagation and spillback can be found in Roberg[1] and Abbess and Roberg[3].

**The author's address: Road Traffic Research Centre, Middlesex University, Queensway, Enfield, Middlesex EN4 4SF.*

3. EXPERIMENT DESIGN

The effect of the model's parameters on the total delay incurred as a result of congestion has been examined in a detailed experiment. The experiment focuses on the performance of a particular control strategy when applied to treat jam formations on the network.

In the experiment, vehicles progress through the idealised one-way grid network. This process is repeated until the system has reached a steady state, whereupon an obstruction is installed at the centre of the network. As a result, a traffic jam begins to evolve. The growth period of the simulation is limited to 10 time-slices to prevent the jam overflowing the system boundary. Following the period of growth, the obstruction is removed. This prompts a slight rearrangement of the queues located close to the original source of obstruction. A further three time-slices are allowed to enable the traffic jam to settle. The control strategy is then implemented for a fixed time period. When the congestion is deemed to have sufficiently cleared, the restrictions enforced by the control measures are lifted. The simulation program is terminated as soon as the network re-achieves an isotropic state.

3.1. Factors influencing the delay

The effects of three factors on the incurred delay (defined below) are evaluated in the experiment. These factors can be represented in terms of the following parameters:

- μ — the number of vehicles per entry point per minute;
- p — the turning proportion: a fixed probability value in the range 0 . . . 1; and
- α — the stop-line width allocated between the segregated queues.

(In the context of this research, the proportion of stop-line width devoted to the *ahead* queues has been denoted by α, with the remainder of the width $1 - \alpha$ allocated to the *turning* queues. For example $\alpha \pm 0.66$ denotes the situation of two lanes for ahead traffic and one for turning vehicles.) The possible values of α have been limited to the parameter set $\alpha \, \varepsilon \, \{0.33, 0.50, 0.66, 0.75, 0.80\}$, as these represent typical lane configurations.

In the experiment, traffic jams have been simulated under a range of conditions. Demand varied from $\mu = 18$ to $\mu = 27$ vehicle per minute per entry point whilst the proportion of turning vehicles ranged from $p = 0.1$ to $p = 0.9$. Different queue storage allocations were investigated by varying α as described above. The effect of these parameters on the total delay incurred during the dispersion period was then measured and evaluated.

4. ESTIMATING THE DELAY IN THE NETWORK

At the end of each time-slice a record is made of the number of blocked links enclosed in the traffic-jam structure, the total number of vehicles present in the system, the total number of vehicles which have entered the system and the corresponding total which have left the road network. These statistics are used to calculate the delay encountered by vehicles in the congested system.

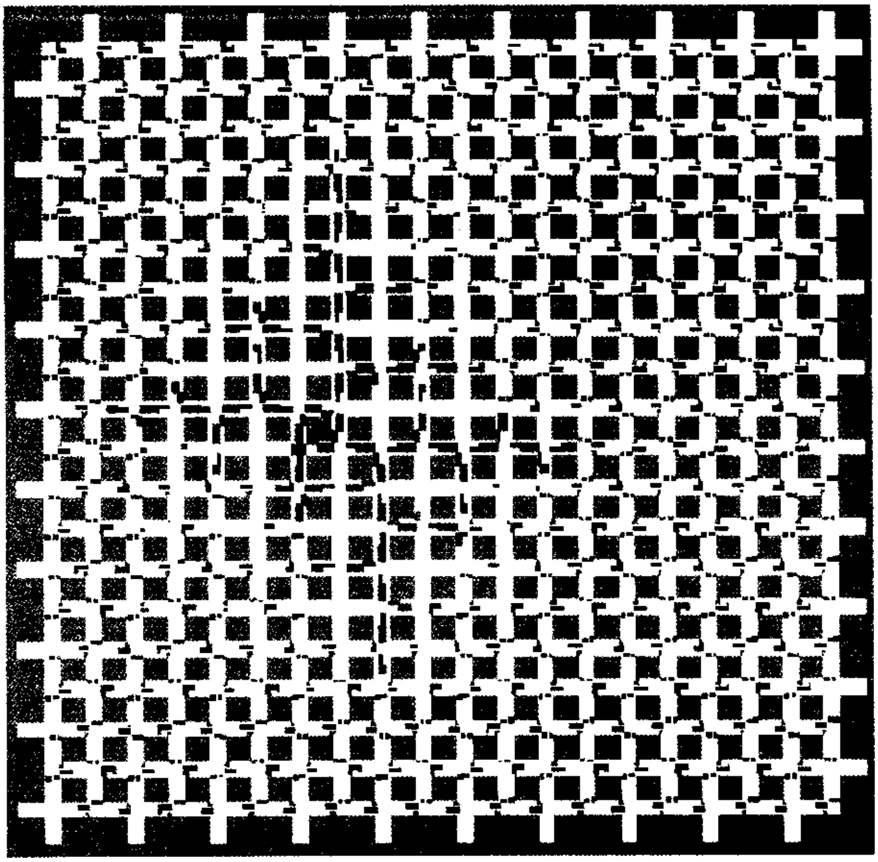

Fig 1. Spatial form of traffic jam just before treatment.

4.1. Steady state conditions

The simulation is brought to a steady state before an obstacle is introduced into the network. The steady state is isotropic — vehicles are evenly distributed among the road links, with some random variation caused by stochastic inputs.

The simulation is run in two phases. During the first phase, known as the RUNUP period, no obstruction is present on the network. The total input to the road network is recorded at the beginning of each time-slice. Similarly, the total output is recorded at the end of each time-slice. The difference between the two variables forms a sequence which can be tested for stationarity using time-series methods[3]. Once the road network has reached a steady state, an estimate is made of the mean and standard deviation of the number of waiting vehicles. Following this, the second phase of the simulation begins. During the simulation period, the experiment described in the previous section is carried out. Summary statistics, such as the number of blocked links and the total number of waiting vehicles, are gathered each time-slice.

4.2. Method for estimating the total delay

When gridlock occurs, it is not possible to measure delay using conventional methods. However, it is possible to compute the total delay (an important statistic in that it gives a direct measure of the financial cost due to congestion) caused by the traffic jam using a different technique.

The delay incurred as a result of congestion can be estimated, either by considering the shortfall of vehicles leaving the system or by comparing the excess number of vehicles in the congested system with the expected number in the steady state. These two techniques are equivalent — see Roberg and Abbess[4]. Here, the latter method is described.

The delay incurred during a single cycle of the simulation can be approximated by comparing the difference between the number of vehicles waiting in the congested system with the corresponding number present in the steady state. This excess, when multiplied by the cycle-time, describes the delay incurred during this particular cycle. In other words, if w_i denotes the number of vehicles present in the system at the end of time-slice i, $\overline{w}$ represents the total number of vehicles waiting in a free-flow, isotropic system and τ corresponds to the length (in minutes) of a single cycle of the simulation, then the contribution to delay *within* time-slice i (denoted by $\delta\Delta_i$) can be written as $\delta\Delta_i = \tau(w_i - \overline{w})$.

The total delay, Δ_j^*, incurred throughout the period of congestion can be found by summing these individual contributions to delay. Thus, if the implementation of control measures in time-slice s successfully eliminates the traffic jam so that the system recovers its steady state at time-slice k, then the total delay incurred as a result of congestion is estimated by

$$\Delta_j^* = r \sum_{j=s}^{s+k} (w_i - \overline{w}) \quad (\text{vch} - \text{min.}) \quad \ldots\ldots (1)$$

The shape of a typical traffic jam just before control measures have been implemented is shown in Fig 1. An obstruction has been set at the junction shaded in purple. In the experiment, external countermeasures are introduced in the 13th cycle. In some situations the traffic jam has been successfully treated, yet in others it continues to grow irrespective of the control measure introduced. Alternatively, the control measure may clear the original jam but cause a new one to build elsewhere in the network. In the context of this research, this phenomenon is referred to as 'jam migration'.

If the strategy completely eliminates the traffic jam from the road network, it is possible to estimate the total delay incurred by the presence of the jam. However, where the traffic jam continues to grow regardless of the control measures which have been introduced, the total delay incurred is not finite, it grows each time-slice with no bound.

5. A COMBINED DISPERSAL STRATEGY

In their discussion of the treatment of 'catastrophic' or severe urban congestion, Huddart and Wright[5] and Quinn[6] propose a number of approaches for tackling the problem. They suggest that:

- *(1)* the control system be altered to disperse or free critical queues; or
- *(2)* reserve capacity be provided to relieve congested links; or
- *(3)* the level of demand be reduced albeit temporarily.

Rathi[7] categorises the treatment of congestion in terms of internal and external metering. Internal metering is implemented to cope with queues along critical intersections, whereas external metering is applied along the periphery of a control area to reduce flow into the congested region. The paper deals extensively with various external metering control-test scenarios.

The implementation of vehicle bans follows a methodology which is similar to the ones proposed above. The bans can be applied in two forms: turn or ahead. Turn bans can be imposed on selected links to break gridlock cores at the heart of the jam, thus dislodging critical queues. Ahead bans around the envelope of the jam reduce input into critical sections of the road. These ahead bans can be used in two ways. Vehicles using the banned junctions can either be queued outside the congested region or else be rerouted away from the jam. Due to the current nature of the software, vehicles which have been re-routed will not necessarily track their original destinations.

Treatment which is solely aimed at the jam boundary will only work providing the gridlock core has not already formed and sufficient storage area is available for the vehicles waiting at the control area's periphery[4].

On the other hand, the success of core strategies (turn bans directed at the gridlock cores in the jam) on their own is not always guaranteed. Results have shown[1] that if the demand is fairly light or if the jam has not been active for more than a few time-slices, it may be sufficient to ban a small number of critical turning movements to clear the whole jam. However, experiments have demonstrated that this simple expedient often forces the core to migrate to a nearby location. This problem has been addressed by using compound *core strategies*. Multiple core strategies become less effective when demand increases and when the jam has already assumed area-wide proportions.

This paper considers the effect of *simultaneously* applying the two types of control available for the treatment of severe urban congestion.

5.1. The control strategy used in the experiment

Experiments with bans have resulted in the formulation of a comprehensive dispersal strategy which will be referred to as the distributed strategy. The distributed strategy combines two control plans which operate in tandem: the flow of vehicles into the congested region is reduced using ahead bans whilst the gridlock cores located at the centre of the jam are fragmented using turn bans.

5.1.1. Turn bans. The arrangement of the bans at the centre of the jam focus on and near the point of the original obstruction. When four sides of a city-block in the network become jammed with turning vehicles, an irreversible knot or core is formed. The traffic can be released from this core using a set of turn bans which are superimposed on the four locking turns in the network. The shaded regions in Fig 2 describe the location of the four cores within the jam structure. Movement along any of the links which bound the shaded regions has been restricted to *ahead only*.

5.1.2. Ahead bans. A cordon of ahead bans is set up around the boundary of the jam. These bans deflect traffic away from the traffic-jam area. The size of the cordon can vary and is usually fixed so that it neatly encompasses the jam without impinging too much on the jam structure. In the experiment, a cordon is set at a fixed distance of six links from each of the four cores at the centre of the jam. Figure 2 shows the ban scheme proposed by the distributed strategy.

5.1.3. Gating. Instead of forcing vehicles away from the congested region via ahead bans, vehicles can be queued on the approaches to the jam without being diverted. This option is described in more detail in Section 7.1.

5.2. Qualitative performance of the control strategy

Applying the *core* and either of the *boundary* (re-routeing or gating) strategies in parallel seems to disperse jams on the idealised network quite quickly. The cordon of restrictions protects the jam from excessive demand whilst the compound core strategies are applied directly to knots at the heart of the jam. The shielding of the cordon minimises the tendency of the cores to reform at nearby locations, and the reduction in the overall jam size and delay time can be dramatic within a few time-slices of the simulation process.

6. EXPERIMENT RESULTS

The experiment has outlined three critical parameters which have an appreciable impact on jam development and dispersion. These factors are not the sole contributors to the underlying processes, indeed the proportion of segregation within a link is associated with minimal jam growth[2]. Here, we deal only with the effect of the level of demand, the proportion of turning vehicles and the space allocated to the segregated queues on the dispersal process of a traffic jam.

6.1. Increasing demands

The experimental results confirm that the rate of jam expansion accelerates with each increase in demand. As a result, the total delay incurred during the period of dispersion also increases proportionately with demand. The distributed control strategy is an effective one. With low, intermediate and often high levels of demand, the strategy successfully eliminates the queue structures.

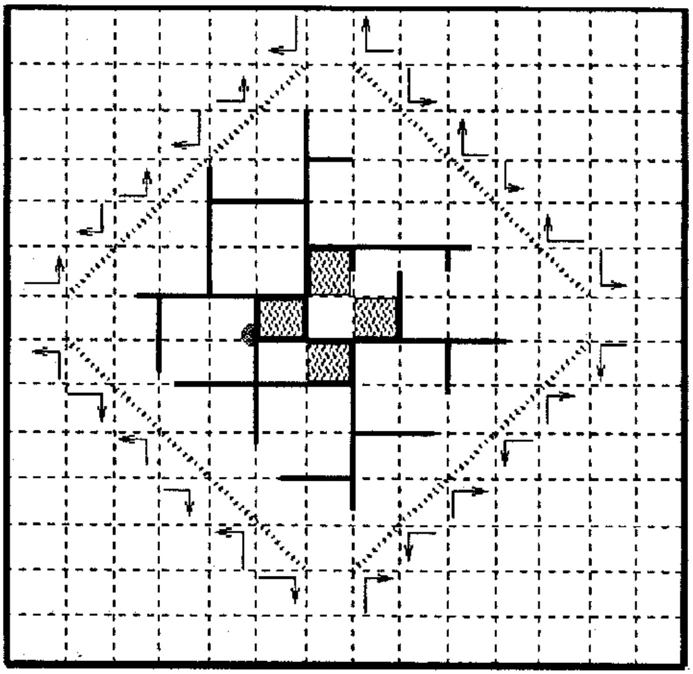

Fig 2. A distrbuted control scheme for jam elimination.

6.2. Increasing turning proportions

The effect of the turning proportion, p, on the performance of the distributed control strategy cannot be easily explained. The results generally show that increasing the turning proportion reduces the performance of the control strategy.

Careful analysis of the results has shown that the turning proportion of vehicles has a crucial effect on the dispersal process. The control strategy is more effective and induces smaller delays when p is low. As p assumes higher values, the control strategy is less successful and often fails to eliminate the jam. For values of p occurring between these two upper and lower bounds, the total delay incurred during the dispersal period does not follow a systematic trend. Instead the total delay is extremely sensitive to changes in p and changes markedly with each variation. It is not possible to predict how the pattern of fluctuations will vary with respect to changes in the turning proportion.

This non-smooth behaviour may be linked with the discrete nature of the software, in which the variable for vehicular flow always remains integer. Thus, certain values of p, when associated with other parameters, such as the level of demand, may induce the dispersing jam into one critical state, whereas another combination may produce a different, albeit related state. Each state will exhibit similar yet distinct patterns of delay. Furthermore, other permutations of model parameters could produce a jam which cannot be controlled by the distributed strategy described in this paper.

6.3. Channelisation

The results show that the queue storage configuration of a link can dramatically affect the delay incurred during congestion. However, care must be exercised when changing the allocations between ahead and turning proportions. The results show that, as p assumes higher values, it would be unwise to increase the value of α beyond 0.50. The reason for this is that allocating more space to the ahead traffic simultaneously reduces the amount of space available for turning movements. When turning proportions are high, spillback can be triggered more rapidly, thus inducing a premature onset of gridlock-type phenomena. This would make dispersal more difficult.

The results for $\alpha = 0.33$ deserve specific attention. The layout allocates one-third of the stop-line width to the ahead queues and two-thirds to the turning queues, and would almost never be used in practice. Nevertheless, it highlights a number of interesting phenomena. One of these is the spontaneous jam which develops over a network without an obstruction. Although the system has apparently achieved a steady state, local pockets of traffic queues are distributed in one or two places in the system. The effect of these upward fluctuations can be graphically observed in the model. Instead of dissipating freely to create a truly isotropic system, these pockets develop, albeit slowly, often one at a time, until one or more traffic jams have formed and extended over the complete network. This is an example of how inappropriate road layout can induce a lack of stability in the system as a whole.

Furthermore, the delays involved when attempting to disperse the jams in the case of $\alpha = 0.33$ are higher than for the remaining values of α. This suggests that the performance of a control policy may be hampered when their exists a mismatch between road storage configuration and the vehicular flows using them.

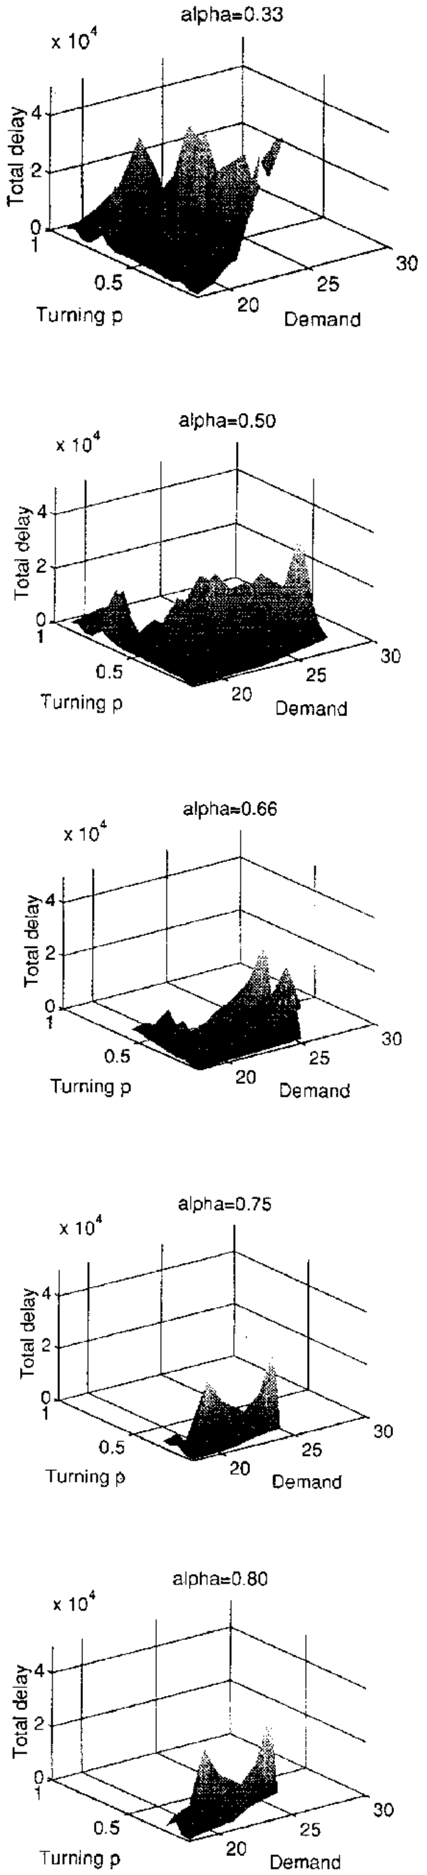

Fig 3 (right). Operative domain of the control strategy with respect to demand, turning proportion and queue storage configuration. (Total delays measured in veh-min.)

6.4. Evaluation of the joint effect

So far, we have seen that the three parameters, μ, ρ and α have a joint impact on the dispersion process of a traffic jam. Figure 3 shows this connection graphically. The figure shows how the total delay, $\wedge_j^*$, varies with respect to demand and the proportion of turning vehicles. The effect of queue storage configuration can also be assessed by considering each case separately. The shaded regions in the respective parts of the figure denote the field or domain in which the chosen strategy has been totally successful in eliminating the jam.

For example, when $\alpha = 0.50$ the volume of the surface covers a large proportion of the plotting area. This indicates that the distributed control strategy is effective for most of the values of μ and ρ which were chosen in the experiment. The degree of shading on the surface demonstrates whether the delays are low (darker shades), or whether they are high (lighter shades). The white regions in the figures demonstrate where the strategy has not been successful at eliminating the jam, for example $\mu = 30$.

The shapes of the surfaces with respect to the change in α appear very similar. However, as α increases the domain in which the distributed strategy is effective becomes smaller. This effect can be seen in the reduction in volume of the surfaces with respect to each increase in α. This means that, regardless of the values of other parameters, allocating more space to the turning part of the segregated queues is not to be encouraged if the overall aim is to disperse traffic jams that have formed as a result of an incident. Indeed, road layouts corresponding to values of $\alpha < 0.50$ are not generally conducive to this purpose. However, in the region where control policies successfully treat jams, the results show that, providing the flows are in proportion to the storage configuration capacities, then the total delay measured over the dispersion period, may be reduced slightly by carefully increasing α beyond 0.50. The benefits of such changes are not large and the sensitivity of the total delay to the arrangement of queues on the link suggest that such interference is not advisable.

7. FURTHER COMMENTS ON EXPERIMENTAL RESULTS

The simulation experiment has revealed that the application of control measures to disperse a traffic jam may lead to one of the following outcomes:

- *(1)* Traffic jam clears completely within a few time-slices.
- *(2)* Traffic jam assumes a static nature — it neither grows nor contracts.
- *(3)* Single jam splits into two mini-jams which:
  - — resonate until one dies out followed by the other; or
  - — resonate until one grows and other dies out; or
  - — resonate until both become gridlocked and extend over the network.
- *(4)* Traffic jam migrates to a fresh location.
- *(5)* Traffic jam continues to grow regardless of the control measures introduced.

The variety of outcomes highlights the sensitivity of a traffic jam to external forms of control and it would be desirable if some of the unusual features, such as resonance and migration, could be explored further.

7.1. Gating instead of re-routeing

The distributed control strategy which has been described so far involves core as well as boundary intervention. Vehicles arriving at the jam envelope are deflected away from the jam using ahead bans. The vehicles are forced to turn on to other routes near the jam. This form of control may in some situations be a little drastic.

An alternative measure can be implemented on the jam boundary whereby, instead of being re-routed, vehicles queue on the envelope of the congested region which acts as a gate. As a result, queues will begin to propagate along the boundary. The longer the restrictions remain in place, the more likely that the queues themselves form the cores of new traffic jams. This implies that the queueing option involves heavier delay to drivers caught in the congestion. However, re-routeing takes drivers away from their destinations. In reality they would drive around until allowed to go where they wanted and the delays might be large. The queueing option represents more realistically what might happen in real life.

Figure 4 graphically compares the difference in delay between the two strategies for one particular set of network parameters. A more detailed comparison between the respective control policies may be found in Roberg and Abbess[4].

8. CONCLUSIONS

In this paper, a control strategy that eliminates traffic jams that have formed as a result of an incident has been introduced. The research has established that two elements are generally necessary to control the gridlock phenomenon. They are:

- *(1)* fragmentation of the gridlock cores of the traffic jam; and
- *(2)* protection of the treatment under *(1)* using a re-routeing or gating policy outside the envelope generated by the traffic queues.

Although each of these measures can be applied singly, the paper has examined the effect of implementing these policies simultaneously and described how one particular strategy performed in a wide range of network conditions.

The implementation of these strategies in the model was represented via banning mechanisms. The gridlock cores, clustered at the heart of the jam, were broken by implementing core strategies at the centre of the jam. This action re-started vehicular movement in the area. To shield the cores from further build-up, vehicles were deflected away from the jam or else simply stopped them from entering the traffic-jam area. The boundary restrictions usually remained in place for a short time, to minimise the effects of new jams appearing on the envelope of the original congested region. The core restrictions stayed in place until the traffic jam was sufficiently broken apart, thus ensuring that the dispersion process would be completely successful.

A detailed experiment considered the effect of one particular strategy (four cores and a re-routeing mechanism constructed near the jam envelope) on the *total delay* incurred for the period of dispersion. The total delay was measured by comparing the total number of vehicles waiting in the system in the congested state with the average number of vehicles present in the system when the network was operating in a free-flow *régime*. This method gave a rough measure of the delay incurred.

The performance of the distributed control strategy was analysed under a variety of road network conditions. Simulation experiments confirmed that under increasing levels of demand and turning proportions it became more difficult to achieve total dispersal. The total delay increased with the level of demand and with the proportion of turning vehicles. However, the dispersion mechanism appears to be rather delicate and interference with the queue storage configuration can be damaging. The results also showed that changes in queue storage configuration often limited the performance of the control strategy.

9. Future research directions

Other work[8] has investigated the effects of incident-based congestion in real urban areas using the CONTRAM suite of packages. However, the work has reported that some of the underlying assumptions in CONTRAM restrict the research from considering the effects of queue interaction and blocking-back on the overall process of jam growth and decay. This deficiency is tackled by the model described in this paper in that it can simulate the mechanism underlying the onset and development of jams. However, the network and the demand pattern are both not very representative of reality. The simulation technique is currently being extended for use on a representation of a real-life network, with different link lengths, capacities and a directional demand profile.

A simple two-way grid network has recently been constructed and, although the mechanism underlying congestion growth is similar to the one-way system, the pattern of queue development is slightly changed. The principles of congestion control remain similar, namely treatment of the critical internal parts of the traffic jam whilst simultaneously reducing flow into the congested region. These ideas will be tailored to suit the needs of the specific jam under consideration.

In order to extend some of the suggestions to a real urban network, a number of issues will be addressed. So far, origins and destinations have not been specified for individual vehicles. An interface between O-D specific information and the simulation model will have to be developed so that variable turning proportions may be calculated for all the links comprising the network. It will also be

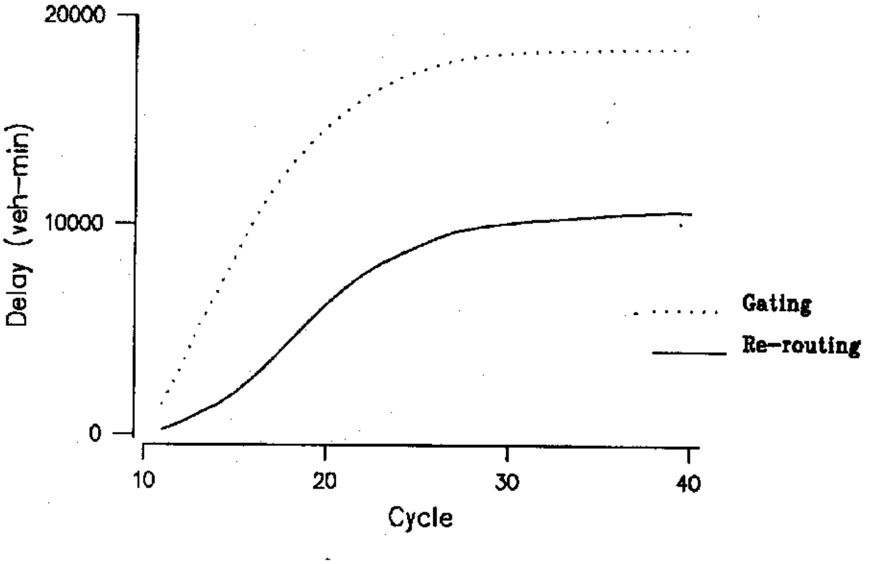

Fig 4. Gating versus re-routeing: a comparison in delay times.

necessary to extend the model so that the junction release mechanism would be able to cope with an irregular network. This could be achieved by devising a strategy for junction processing in a non-symmetrical environment.

The envelope of the traffic queues is likely to depend on the position of the original obstruction. The traffic locking mechanism may be different from the simple patterns simulated in this project. As a result, experiments will be necessary to determine how queues propagate. Strategies for breaking the gridlock will be developed using this information.

This will be tested using a model of a particular real urban network, with and without protection being provided by a shield of appropriately-placed traffic redirections or gates. Ultimately, it is envisaged that the model could be used as a tool to help select appropriate control strategies to clear cases of incident-based congestion.

ACKNOWLEDGMENTS

The author would like to thank her supervisors, Chris Abbess and Professor Chris Wright, for their comments about this paper, and Beth Abbess for her assistance in conducting some of the numerous simulations.

REFERENCES

[1] ROBERG, P. Development and dispersal of area-wide traffic jams. *Traff. Engng Control*, **35**(6), June 1994, 379-385.

[2] WRIGHT, C.C. and P. ROBERG. Traffic jams on idealised rectangular grid networks: A queue propagation model. Submitted to *Transportation Research*.

[3] ABBESS, C.R. and P. ROBERG. Fractual structure of traffic jam images. *Proc., IMA Conference: Complex Stochastic Systems and Engineering*, **54,** 1993, pp 251-281.

[4] ROBERG, P. and C.R. ABBESS. Diagnosis and treatment of congestion problems in central urban areas. Presented at EURO13/OR Conference, Strathclyde, 1994.

[5] HUDDART, K.W. and C.C. WRIGHT. Catastrophic congestion and some possible ways of preventing it. *Proc. TRAFFEX Seminar on 'Congestion, control and parking enforcement'* (Brighton, April 1989). PTRC Education and Research Services Ltd, London, 1989.

[6] QUINN, D.J. A review of queue management strategies. *Traff. Engng Control*, **33**(11), November 1992, 600-605.

[7] RATHI, A.K. Traffic metering: an effectiveness study. *Transp Q.* **45**, 1991, 421-440.

[8] VAN VUREN, T. and D. LEONARD. Urban congestion caused by incidents. *Traff. Engng Control*, **35**(7/8), July/August 1994, 422-429.

Bursaries for transport studies

To encourage graduates to specialise, the following consultants have agreed to provide bursaries for postgraduate study in the 1995-96 academic year: Halcrow Fox and Associates, Ove Arup and Partners, WS Atkins Planning Consultants and The MVA Consultancy.

The bursaries, worth £1 000 to £1 500, can be used to supplement SERC/ESRC or similar advanced course studentships at the following institutions: Cranfield Institute of Technology, the University of Leeds, the University of London Centre for Transport Studies, the University of Newcastle upon Tyne, the University of Salford and the University of Westminster (and others).

Enquiries should be addressed to Professor Tony Ridley at ULCTS, Imperial College Road, London SW7 2BU (*Tel:* 0171-594 6100; *Fax:* 0171-594 6102).

The relationship between loss-of-control accidents and impact protection standards: a case-study

Steve Proctor, TMS Consultancy, writes: During the planning of new road schemes, highway designers will apply Department of Transport standards for the provision of safety fences on motorways and high-speed dual-carriageways. But do these standards adequately protect road-users who leave the carriageway in a loss-of-control accident?

Safety audits are now compulsory on new national road schemes. The safety audit involves checking the road design for possible safety problems and recommending improvements.

During the safety audit process on motorways and high-speed dual-carriageways there are a number of frequently-cited problems. One safety problem often highlighted is the potential for vehicles to lose control on high-speed roads, leading to injuries from collisions with fixed road objects. However, the recommendations often suggested by safety auditors for eliminating some of these problems require provisions beyond the current standards for safety-fence installation.

In order to establish whether the problems foreseen by safety auditors are justified — and, if so, how these shortcomings in current safety standards could be overcome — an independent study was undertaken during which the details of over 250 single-vehicle, non-pedestrian (SVNP) accidents in Warwickshire were examined. The accidents occurred during a three-year period (1991-1993) on motorways and derestricted dual-carriageways in the county. The purpose of the study was to examine the consequences of these accidents, in particular the relationship between current safety standards and what happens in real crash situations*.

In addition, an analysis of six accident sites in Birmingham where crash-cushions have been installed was carried out to see whether this form of crash protection system should be more widely used. These crash-cushions have been struck on at least 47 separate occasions, giving a substantial amount of information relating to performance in real crash situations.

The SVNP study showed that a number of scenarios suggested by safety auditors do occur in real accident situations. In particular, it is estimated that collisions with wooden fences, end terminals and collisions after descending embankments could account for around 25 per cent of the SVNP accidents on motorways and derestricted dual-carriageways in Great Britain, costing approximately £28M per year.

The proportion of SVNP accidents occurring on Warwickshire motorways (40 per cent) is significantly higher than the national average (28 per cent). During the three-year period studied, the number of SVNP accidents rose by 37 per cent compared to a rise of 12 per cent for non-SVNP accidents.

What was the cause of these accidents? The most common human factors were drivers falling asleep and excessive speed. An analysis of vehicle factors showed that 24 per cent of accidents involved a vehicle's tyre bursting; and in 6 per cent of accidents, vehicles were towing a caravan or trailer.

But what happened after the vehicles lost control? The majority of vehicles (83 per cent) left the carriageway immediately after losing control, around half of them rolled over and just under half of all vehicles hit a safety barrier (Table I). One-third of accidents involved collisions with objects off the highway; the most commonly struck items were wooden fences, road signs and French drains (Table II).

A detailed study of accidents on the M40 highlighted examples of vehicles leaving the carriageway where the gaps between adjacent sections of safety fence was less than 150 metres. A number of accidents occurred whereby a vehicle descended an embankment through a short gap in adjacent sections of fence. The average gap through which vehicles passed in the study was 120 m. In terms of construction cost, it is not economically viable to leave a gap of less than 60 m between sections of fence.

The study also suggested that there are problems for vehicles both ascending and descending embankments, particularly where a wooden fence is situated close to the carriageway: 15 per cent of vehicles descended an embankment and a further 18 per cent of vehicles ascended an embankment on leaving the carriageway. The M40 study revealed that injuries occurred where the average height of the embankments descended during the accidents was 2.25 m. Existing standards do not provide for the protection of descending embankments of less than 6 m and there is no provision at all for the protection of ascents.

Two-thirds of vehicles that ascended embankments and a half of vehicles that descended embankments struck a wooden fence. Where vehicles hit a wooden fence, there was a higher severity rate (42 per cent fatal and serious, compared to 32 per cent fatal and serious for the remainder of the accidents).

The study looked at situations where impact protection systems may have been inadequate despite being constructed to current standards, e.g. where the leading edges of the safety fence were inadequate in advance of a safety structure. Of the vehicles that hit a safety fence after losing control, 15 per cent involved an impact with a ramped end terminal. In at least five cases, vehicles were 'launched' off an end terminal.

An analysis of crash-cushion terminals in Birmingham showed that there has been a significant reduction in accident severity at these sites (Table III). This fact, together with the 200 per cent rate of return on each installation, implies that the widespread use of crash-cushion systems on derestricted dual-carriageways and motorways (in particular at ramped end terminals) would result in a considerable reduction in the number and severity of casualties. It is estimated that around 125 injury accidents occur each year throughout the country at sites that could be protected by crash-cushions.

Table III. Number of accidents and severity index at crash-cushion sites before and after installation

Severity	Before	After
Fatal	3	0
Serious	5	1
Slight	4	6
Severity Index	67%	14%

**The TMS Consultancy is an independent firm specialising in traffic management and road safety engineering. This note is an abridged version of a paper given at the European Road Safety Conference held in Lille in 1994.*