Simple models for traffic jams and congestion control

C. Wright, BSc, PhD, ACGI, CEng, MICE, FIHT, CMath, FIMA and
P. Roberg-Orenstein, BSc, PhD

Proc. Instn Civ. Engrs Transp., 1999, **135**, Aug., 123–130

Paper 11942

Written discussion closes 1 February 2000

It would be useful to know more about the spatial structure of traffic jams and how they propagate, so that techniques can be developed for inhibiting and dispersing them. The authors have analysed the development of traffic jams caused by blockages or incidents on idealized rectangular grid networks, and although the models are not very realistic, they reveal some basic features of jam formation and point to methods for controlling them. With hindsight, the principles seem to apply qualitatively to real traffic situations as well. The models show that the structure of jams on idealized grid networks can follow either of two regular patterns. Also, jam growth can be inhibited by reallocating queue storage space between turning traffic and ahead traffic, and by shortening the length of the segregated queue storage space. Containment, however, and the prevention of gridlock, are two distinct objectives that conflict with one another. Once a jam has formed, in many cases it can be dispersed by applying a gating strategy around the affected area in combination with temporary turn bans at selected junctions near to the site of the blockage. Recently, the models have been extended to two-way rectangular grid networks. The aim of this paper is to summarize the results and draw together some conclusions that have emerged from the project as a whole. Finally, the paper reviews the practical implications of the model and assesses the potential for inhibiting and dispersing jams on real networks using low-cost methods.

Keywords: management; mathematical modelling; traffic engineering

Introduction

Traffic jams are growing in frequency and severity in major cities throughout the world. It would be useful to know more about their spatial structure and how they propagate, so that something can be done about them. The potential for low-cost measures such as channelization and turn restrictions has scarcely been explored, yet they can have appreciable effects.[1,2]

2. The authors have explored the development of traffic jams caused by blockages or incidents on idealized rectangular grid networks.[3,4] Although the models do not simulate the movements of vehicles in detail, and are therefore not as 'realistic' as the micro-simulation models that are commonly used to visualize traffic movements at the local level in traffic research, they are designed to run at high speed and to capture the main features of a jam extending over a wide area of network. They reveal some interesting phenomena and point to methods for deterring jams and controlling them once they have formed. The principles apply qualitatively to real traffic situations, and would not have been easily accessible in any other way.

3. Recently, the models have been extended to two-way rectangular grid networks. The aim of this paper is to summarize the results and draw together some conclusions that have emerged from the project as a whole. Finally, the paper reviews the practical implications and assesses the potential for inhibiting and dispersing jams on real networks using low-cost methods.

The models

4. A traffic jam becomes an interesting spatial phenomenon when a one-dimensional queue develops into a two-dimensional queue branching system through the merging process. Additionally, queues on a particular route may block one or more junctions in such a way as to obstruct traffic movements on other routes that *cross* that stream, a development that is called *cross-blocking*. When a traffic queue propagates around all sides of a city block (*gridlock*), vehicles at the head of the queue will have difficulty in dispersing even if the original obstruction is removed.

The networks

5. The model networks are idealized representations of city centres built on a rectangular grid pattern. They consist of two types: one-way and two-way. Both have two sets of parallel roads at right angles with the roads at equal spacing, and in the case of the one-way network, the direction of traffic alternates between neighbouring roads. There are no intermediate origins or destinations within any of the links, and only the severed links at the

Chris Wright, Head, Transport Management Research Centre, Middlesex University

Penina Roberg-Orenstein, Research Analyst, Bellcore, Network Traffic and Performance Division, New Jersey

boundary act as traffic sources or sinks. In the absence of a traffic jam, the models stipulate uniform demands on all the links in the network. Hence the flow between any origin—destination pair is constant; furthermore, the route used by each driver is effectively fixed in advance—there are no diversions in response to congestion.

6. For both types of network, the entry directions have equal priority, that is, the opportunities for discharge are shared between the approaches in proportion to the demands. Hence each junction is seen as controlled by a traffic signal that has very short cycle times, and is undersaturated for all approaches.

7. The area of carriageway occupied by each link is divided into two distinct zones: the approach to the downstream junction stopline where traffic is channelized into separate turning areas, one for the ahead movement and one for each of the turning movements; and upstream, a non-channelized 'reservoir' where all vehicles are mixed together regardless of their intended movement at the downstream junction. Each storage area is of uniform width, and the storage arrangements are identical for all links.

8. When unobstructed, the ahead and turning discharges per unit time across the stopline at the exit of any link are fixed, and proportional to the width of road occupied by the lanes under consideration, unless the upstream queue is exhausted. But if a queue spills back over the stopline from downstream, the flow falls instantaneously to a lower level until the queue spillback disappears. The density of vehicles per unit area of storage space in a spillback queue is the same for all queues.

9. There is an additional complication for two-way networks in that the right-turning vehicles discharge in two stages: from the segregated right-turn queue storage area they progress into a relatively small 'bin' and then discharge separately at the end of each simulation cycle. A representation of the traffic flowing into and out of a typical link on a two-way network is shown in Fig. 1, while the layout of the individual queue storage areas is shown for both one-way and two-way networks in Figs. 2(a) and (b) respectively.

10. The speeds of all vehicles are assumed to be the same everywhere on the network until they enter spillback queues. Speeds within a queue are determined by the rate of discharge of vehicles at its head.

Spillback and queue propagation

11. The distinction between the channelized queuing areas and the upstream reservoir is important for two reasons. First, it reflects the way road markings typically are applied on city streets, and second, it allows the system to model explicitly the interference between the different turning movements when spillback occurs. The different turning movements interfere because even if there is a separate turning lane for each exit, drivers do not necessarily position themselves in the correct lane as soon as they enter the link. Hence if a particular exit (say, the left-turn exit) becomes blocked, not only will vehicles intending to turn left form a queue that spills back along the link, but at some point the queue will spread across the full width of the road, blocking other traffic.

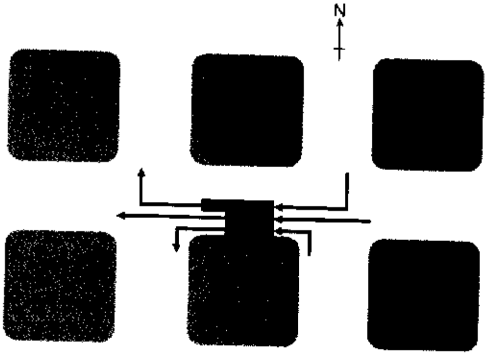

Fig. 1. Arrangeme of links and juncti on the two-way networks: the black area represents a queue in the right-hand lane on a westbound link tha has spilled back int the non-channelize area upstream

12. The authors visualize a jam starting because of a temporary obstruction or a permanent bottleneck. What happens subsequently is the same in either case: a traffic jam propagates upstream from link to link, each queue generating new branches at each junction in turn. Initially, the topology of the branches resembles that of a tree, but at some stage a queue will spill back around four sides of a block to form a closed loop—'gridlock'. This may be repeated elsewhere until the tree evolves into something that would be more accurately

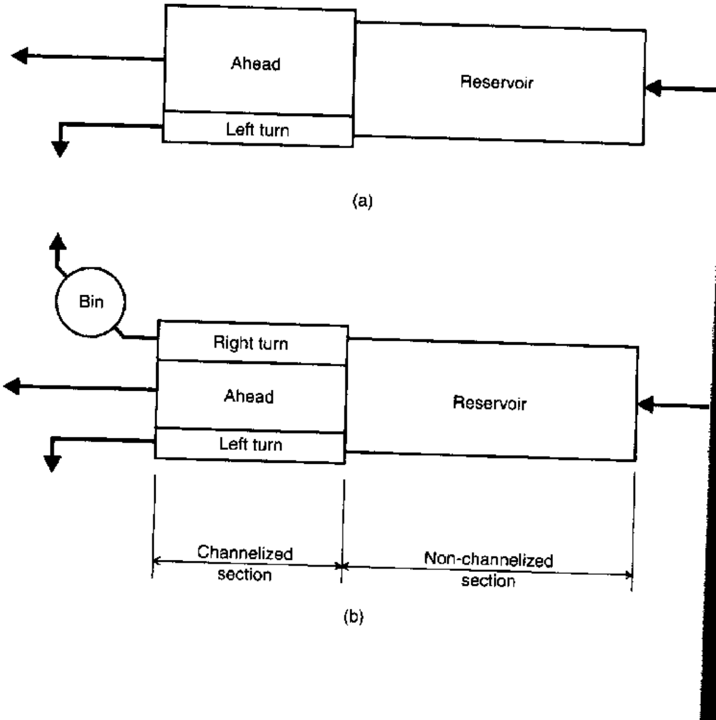

Fig. 2. The layout of queue storage areas for: (a) a one-way network; (b) a two-way network

gement
unctions
ay
black
ts a
right-
a
k that
ck into
alized
a

ut of
reas
ay
wo-

described as a lattice. Provided it is assumed that drivers do not reroute to avoid congestion, the pattern of queue propagation depends entirely on the interference between turning movements, and not on the origin–destination pattern.

13. To start with, the authors assume there is no 'cross-blocking', that is spillback does not block crossing traffic at the upstream junction. Instead, vehicles wait on their respective approaches until there is space for them to proceed. Later, the authors will try to assess the effects of cross-blocking in qualitative terms.

Traffic jam structure and growth

14. The paper can now review the main features of jam structure and growth on idealized networks as determined from the analytical model and the simulation models, and draw some broader conclusions about jams on real networks.

Idealized one-way networks

15. The main results concerning the one-way networks have been described by Wright and Orenstein,[4] and only a brief summary will be given here. Suppose that a blockage restricts the throughput of traffic along a link to a value less than the demand. It does not matter exactly where the blockage is located: a queue will propagate upstream towards the beginning of the link and spread across all lanes, affecting entry from both of the two links feeding into the upstream junction. Suppose that the total entry flow is now restricted to the value cq, where q is the total demand for entry to the link under unobstructed conditions in vehicles per unit time, and c is the 'constriction factor' ($c < 1$). The discharge of ahead vehicles into the blocked link from the upstream junction, and the discharge of turning vehicles, will both be reduced by the same factor c. Queues will now propagate along both upstream links that feed into the blocked link, and to other links in turn.

16. The spatial structure of the resulting queue system can take either of two distinct forms. If the ahead queues spill back more quickly than the turning queues, propagation will take place predominantly along the four axes of the grid, leading to a diamond-shaped jam boundary (type I jam). If the turning queues spill back more quickly than the ahead queues, the queue propagation will take place predominantly in tightly packed 'curls' around neighbouring city blocks, leading to a convex eight-sided jam boundary (type II jam).

17. The relative rates of spillback of ahead and turning queues, and hence the type of jam structure, can be determined by comparing (a) the ratio ϕ of ahead to turning traffic demand under unobstructed conditions with (b) the ratio α of stopline width allocated to ahead and turning traffic. If the two are equal ($\alpha = \phi$), the layout is 'balanced' and the jam conforms to an intermediate structure on the threshold between types I and II. If the stopline width is allocated more favourably to turning traffic ($\alpha < \phi$), the ahead queues will spill back more quickly than the turning queues and the jam will develop a type I structure. If it is allocated more favourably to ahead traffic ($\alpha > \phi$) the turning queues will spill back more quickly than the ahead queues and the jam will develop a type II structure.

18. The analytical model does not take into account any randomness in the vehicle arrival times, nor the phenomenon of 'starvation', in which the flows on links both within and around the traffic jam are reduced because the vehicles that would normally be moving along those links are held up elsewhere within the jam itself. Nor can it handle two-way networks. The authors have therefore used a simulation model to see how these factors might affect jam structure. It is based on the same assumptions about queuing behaviour, the same network configuration, and the same structure of traffic inputs and outputs, but there are differences of detail which are described in an earlier paper.[4] Stochastic variation in the level of demand is modelled in terms of Poisson inputs with mean demand μ vehicles per unit time at each access point.

19. The simulation model confirms that the boundary of a type I traffic jam is roughly diamond shaped (see Fig. 3 for a typical example). However, alternate roads inside the

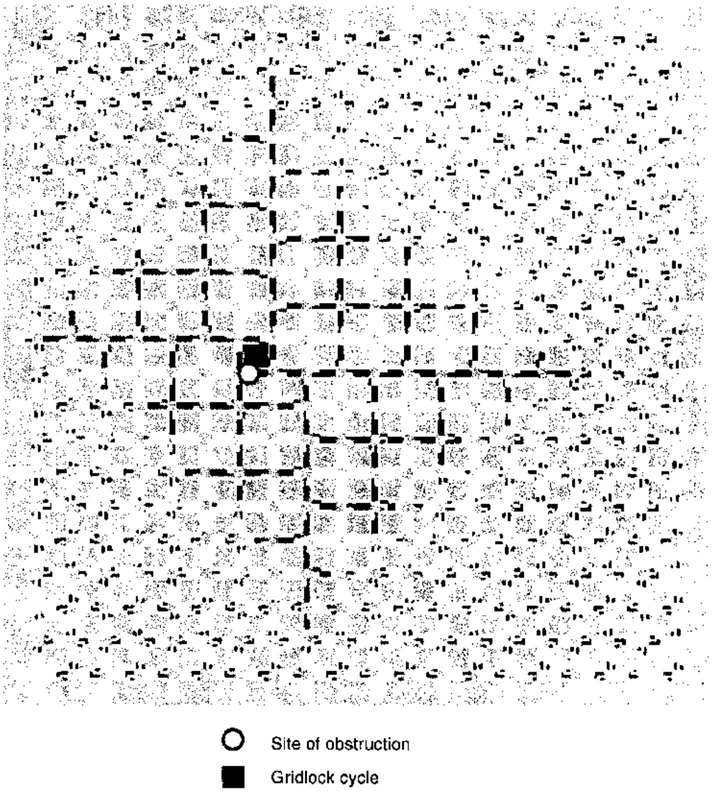

Fig. 3. A simulated type I jam on a one-way network

boundary that lead away from the centre of the jam are starved of vehicles, and the absence of traffic on these roads is referred to as 'anti-queues'. Starvation affects the shape of the boundary as well, but this only becomes apparent when the jam is allowed to grow to a larger size, when the jam boundary becomes indented like a four-petalled flower; in fact it is a fractal and its structure has been analysed in detail by Abbess and Roberg.[5]

20. There is a third starvation effect that is less obvious than the other two. Vehicles are not able to leave the jam at the rate they would have done in the non-congested state, and so the outbound flows are depleted for some distance outside. The 'halo' of reduced traffic flows around the boundary is just perceptible in some of the computer visualizations (particularly for the two-way network shown later on in Fig. 4).

21. Unexpectedly, when the model parameters are set to values that in theory correspond to a type II jam, the simulation model does not generate a type II structure except under extreme conditions. Again, this appears to be a result of traffic starvation. The tightly packed 'curls' predicted by the analytical model cannot build up as quickly as they would otherwise do because of the restricted supply of vehicles within the jam boundary.

22. Finally, stochastic fluctuations in the input flows seem to have little effect on the overall shape and size of the jam after the first few cycles. On the physical scale of an area-wide jam, the random effects are swamped by the deterministic component.

Idealized two-way networks

23. So far, the authors have not been able to develop an analytical model for two-way networks. However, the simulation model has been extended to handle two-way grid networks with additional routines for handling both right-turning traffic and left-turning traffic at each junction. No attempt has been made to model the effects of right-turning vehicles on other turning movements when they are 'frozen' in the centre of the junction because of a blocked exit.

24. The results obtained with the two-way simulation model are similar in many respects to those obtained with the one-way model, with four main differences. First, the queues propagate more quickly, in the sense that when a link fills up with queuing vehicles, the queue blocks entry from three connecting links at the upstream junction rather than two. Second, there are no empty road links within the jam boundary, although the queues in the outbound direction are usually very small. They form a pattern of little-used traffic movements that echo the 'anti-queues' observed on the one-way networks. Third, for type I jams, the boundary shape is altered somewhat with one of the four quadrants relatively free of queues in either direction on each link (Fig. 4). As might be expected, jam growth rates on two-way networks are particularly sensitive to the amount of storage space allocated to right-turning vehicles. Fourth, gridlock can occur either in the clockwise direction or the anticlockwise direction, depending on the rate of spillback of left- and right-turning queues.

Modelling queue propagation on more general networks

25. Many of the characteristics of jam growth that the authors have observed in their models seem to apply qualitatively to real networks. One important difference, however, is that on real networks, the traffic demands and the queue storage arrangements will vary from link to link, and so will the mode of queue propagation. A jam can be a mixture of types I and II. Nevertheless, for any given network it is not difficult in principle to build a model that will predict the branching structure of queues propagating from an obstruction occurring at any given point, provided that the design carries over some of the simplifying assumptions that underlie the idealized models. These include constant demands over time, together with the postulated mechanism of queue interference at the interface between the channelized and non-channelized segments within each link, and no cross-blocking.

Fig. 4. A simulated type I jam on a two-way network

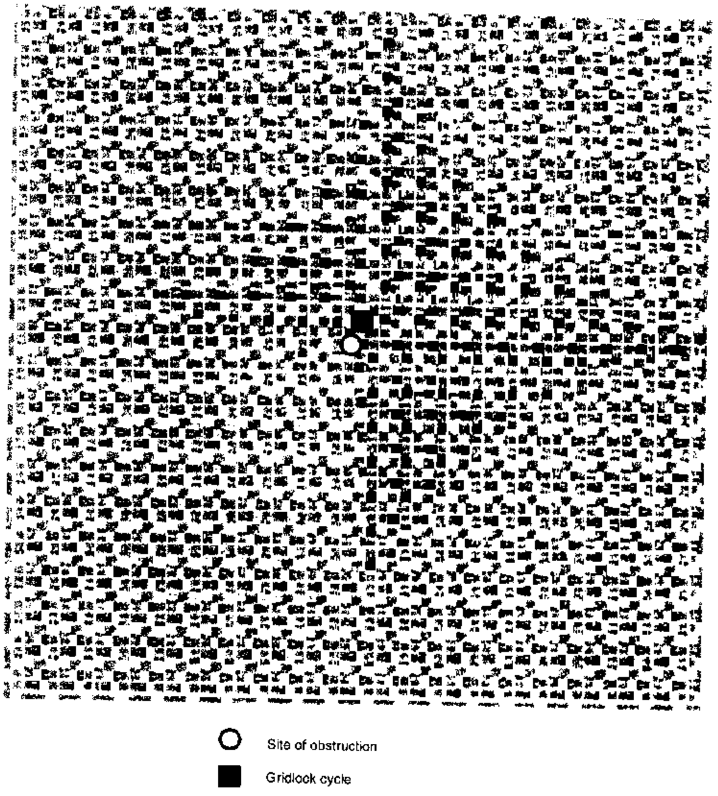

26. The essential idea is that each link of the network has a characteristic spillback time. It has two components: the time taken for a queue to fill up one or other of the channelized queue storage areas, and the time taken for it subsequently to fill up the non-channelized area upstream. Suppose that link j has M exits, numbered $i = 1, 2, ..., M$. Any one of these exits may become blocked. Denote the area of channelized queue storage area associated with exit i by A_{ij}. The spillback time $T_{ij}(c)$ for exit i of the channelized section of link j is defined as the time taken for a queue to fill up the area A_{ij} when exit i is obstructed, such that the discharge is reduced to a value c times the corresponding demand q_{ij} under unobstructed conditions. Let the maximum number of vehicles that can be stored in the area A_{ij} be K_{ij}. Then $T_{ij}(c)$ is given by

$$T_{ij}(c) = \frac{K_{ij}}{q_{ij}(1-c)} \qquad (1)$$

27. Similarly, denote the spillback time for the upstream reservoir by $T_{uj}(c)$, its storage capacity by K_{uj}, and the total demand for entry to the link under unobstructed conditions by q_{uj}. Then

$$T_{uj}(c) = \frac{K_{uj}}{q_{uj}(1-c)} \qquad (2)$$

28. The next step is to construct a network consisting of possible queue propagation paths. It will resemble the original network except that all the channelized queue storage areas are coded as separate links, as shown in Fig. 5. The process of tracing the queue propagation tree is identical to that of finding a minimum-cost tree in traffic assignment, except that in this case the process involves working against the direction of traffic flow and using the spillback times rather than journey times as the cost criterion.

29. Such a model would be useful in identifying 'critical' links within a network. It could be run several times, with an obstruction located in a different place each time. A comparison of the results would reveal the locations that are most sensitive to obstruction, and also those links that play an important role in queue propagation.

Techniques for inhibiting and dispersing jams

30. The severity of a traffic jam can be quantified in terms of aggregate journey times and delays, but the authors have found it more convenient to use two simpler measures with their model. The first is the rate of growth of the traffic jam, its size at any moment being represented by the number of blocked links. The second is the *area* enclosed within the jam boundary, where 'boundary' means an imaginary polygon joining the outer extremities of the queues at any particular moment. This is not a rigorous concept, but an intuitively useful one when dealing with jams that are spread out over a large part of the network.

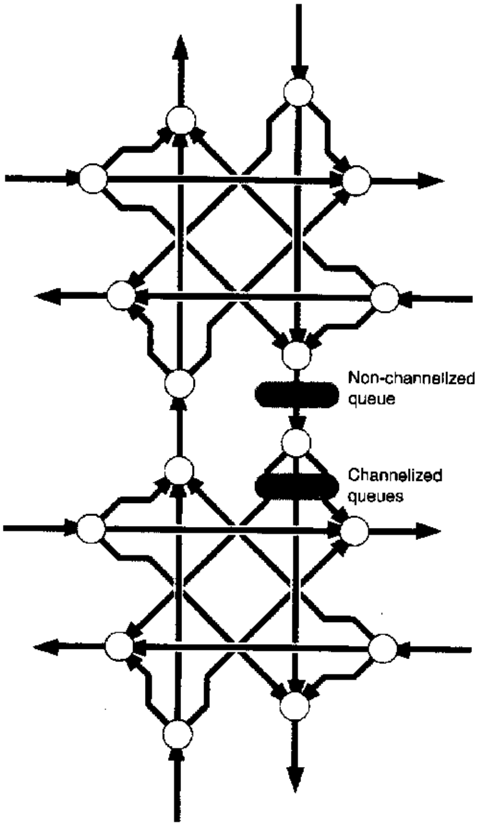

Fig. 5. Representation of network for tracing queue propagation paths

31. Congestion control measures fall into two main categories

(*a*) those that inhibit jam formation
(*b*) those that disperse them once they have formed.

Inhibiting jams itself involves two distinct objectives: minimizing the rate of growth, and deferring the onset of gridlock. These will be dealt with separately in turn, and the dispersal of jams will be considered later.

Inhibiting jams

32. Using the analytical model described earlier, Wright and Orenstein[4] have shown that the rate of growth of the area of a one-way jam is approximately minimized when the allocation of queue storage area between the ahead and turning movements is balanced ($\alpha = \phi$). Unless a link is channelized along most of its length, the effect of varying either of the parameters α or ϕ is not very large, and on real networks the allocation will be roughly balanced anyway for capacity reasons. However, for any given values of α and ϕ the rate of growth can also be reduced by decreasing the proportion σ of the total length of each link that is channelized. The reason is that when a particular exit

becomes blocked and a queue spills back to the upstream reservoir, the supply of vehicles to the other exits is cut off. The channelized storage areas associated with the other exits will become empty of vehicles, and remain empty while the jam propagates to other links upstream. Hence, the shorter the channelized length (σ small), the greater the proportion of carriageway area that can be used for storing vehicles, and the more compact the jam, whether or not the allocation of stopline width is balanced.

33. These results have been confirmed using the simulation model described in ¶15–20. It has not been possible to extend them to two-way networks at the same level of detail, but preliminary results obtained with the two-way simulation model suggest that the same principles apply—it would be surprising if they did not.

34. Another important consideration in congestion control is the time taken for queues to propagate around four sides of a city block, a quantity referred to as the 'gridlock time'. It depends on the rate of spillback of turning queues only. For a given channelization proportion σ, it increases if the stopline width is allocated favourably towards the relevant turning movement ($\alpha < \phi$), because the turning queue storage space is thereby increased, and the turning queues spill back less quickly than they would otherwise have done. Alternatively, for a given allocation of stopline width, the gridlock time increases as σ increases.

35. Hence, the models point towards two quite different strategies in terms of queue storage layout, according to whether it is required to minimize jam growth rate or defer gridlock. The strategies are summarized in the first two rows of Table 1. Neither strategy is entirely compatible with current traffic engineering practice (third row), where a balanced stopline width allocation and maximum channelization are the norm.

Dispersing jams

36. The measures that have been outlined for inhibiting jams can loosely be described as 'static'. They are concerned with the allocation of queue storage space on each road link, and once they have been implemented, the lane markings stay that way until the paint wears out. When a traffic jam occurs, however, more 'dynamic' measures are needed to disperse it.

37. Experiments with the one-way model networks have shown that the most effective strategy has three strands that must be applied in parallel. The first strand involves removing the original obstruction. The second involves 'gating' the flow of traffic at a suitably placed cordon around the jam in order to reduce its growth rate. The third involves the application of turn bans at selected junctions close to the site of the original blockage in such a way as to disentangle the gridlock cycles at the heart of the jam. In principle, the last two could both be accomplished via centrally controlled traffic signals. Roberg and Abbess[6] have given details of the strategy, and the conditions under which it is most effective. The salient points are as follows.

(*a*) The response time of the control system is important: the quicker the response, the more likely it is that the jam will disperse.

(*b*) The time required for dispersal increases with each of the following three quantities
  (i) increasing traffic demand
  (ii) increasing proportion of turning vehicles
  (iii) increasing proportion α of stopline capacity devoted to ahead traffic.

(*c*) If vehicles are deflected away from the jam area altogether at the surrounding cordon, as opposed to merely being held up for a while before being allowed to proceed, the dispersal times observed with the authors' model are very roughly halved.

(*d*) The cordon gating process can be phased out before the core of the traffic jam has actually cleared, and indeed it is important not to build up long queues at the cordon because they will themselves develop into traffic jams if held for too long.

38. Of course, the outcome will not always be successful. Depending on the demand conditions and other factors, the models show that a jam may contract briefly then resume growth, oscillate in size, migrate to another location, or even split into two.

Table 1. Alternative strategies for allocating queue storage space

Objective	Strategy	
	Allocation of stopline width	Extent of channelization
Minimize jam growth rate	Balanced	Minimal
Defer gridlock	Turning movements favoured	Maximal
Minimize journey time under non-congested conditions	Balanced	Maximal

Limitations of the models

39. Real networks have at least three features that the authors have not taken into account. First, they are geometrically irregular. Second, the traffic sources and sinks are scattered at intervals along each street within the network, not just on the network boundary. Third, traffic signals break up the flow into platoons, in other words the flow pattern is cyclical. Much depends on whether these cycles are second-order effects superimposed on the deterministic pattern of jam growth as pictured in the model: if they are, the model might still usefully predict the growth of the jam envelope on a coarse physical scale over a long period of time. However, it is suspected that dynamic 'coupling' effects between neighbouring signals can have significant effects.

40. In order to investigate the sensitivity of the model to the interaction among vehicles at the interface between the upstream reservoir and downstream storage areas, an additional parameter s has been incorporated. When its value is zero, all vehicles are blocked at the interface from the moment when any one of the channelized queues fills up its available space. However, when its value is 1, spillback of the ahead queue has no effect on turning vehicles, and spillback of a turning queue has no effect on ahead vehicles. The authors imagine that, somehow, vehicles in the reservoir whose downstream exit is clear are able to filter past the baulked vehicles (perhaps by driving on the footway!).

41. According to the model output, relaxing the interaction mechanism in this way has a beneficial effect on the jam growth rate (a reduction of up to 25% can be observed). Unless the proportion of turning vehicles is more than about 20%, the benefits are modest, because the jam spills back along the principal axes at much the same rate as before. Relaxation has an *undesirable* impact on the gridlock time, which was observed to fall by up to 25% in the model runs. The authors suspect this is because the interaction between ahead and turning queues is largely responsible for the 'starvation' effect noted earlier. When the interaction is removed, the turning queue storage areas are less likely to be 'starved', and fill up more quickly than they would otherwise do.

42. It is not possible to include the effects of cross-blocking within the computer model, but it is possible to make an educated guess about the effects it might have on jam development. In a type I jam, an ahead queue spills back more quickly than the turning queue, so cross-blocking effectively propagates the jam to the link concerned more quickly than would otherwise be the case, and the average delay to vehicles will increase as a result. However, according to the analytical model, this has no effect on the rate of growth of the jam *boundary*, whose shape is determined by the positions of its corners on the four principal axes. These positions in turn are defined by the rate at which queuing propagates in the *ahead* direction from link to link. On the other hand, it does speed up gridlock.

43. In a type II jam, each turning queue spills back more quickly than the ahead queue, and cross-blocking of the ahead queue has no effect on the propagation time for that link at all, nor on the gridlock time.

44. The final problem with the model is that it does not address the possibility that drivers will reroute when they see queues ahead. Much depends on the physical scale of the jam in relation to the length of the journey each driver wishes to make. If the driver's destination lies within the jam boundary there may be little incentive to divert. On the other hand, a driver making a through journey may be tempted to find a way round it. In either case, the driver can only guess how far the jam extends, and his or her behaviour may depend on local factors of the sort that are not susceptible to modelling along the lines attempted here.

Conclusions

45. This paper has highlighted some of the factors that influence the development of a jam triggered by an isolated obstruction, for idealized grid networks under artificial conditions. While the approach sacrifices realism, some general conclusions emerge that seem to apply to real networks. Some of them are contrary to the authors' expectations. With the benefit of hindsight, they are perhaps simple and obvious, but it would have been difficult to approach them in any other way: tracing the progress of real traffic jams through time is almost impossible in practice. Only by simulation is it possible to observe, and make sense of, the regularities of structure that emerge consistently on a large physical scale.

46. First, the allocation of queue storage space between ahead and turning traffic at the exit from each link affects

(*a*) the rate of growth of a jam
(*b*) the time required for gridlock to develop.

The results are not very sensitive to the assumptions made about how ahead and turning vehicles interfere with one another, nor to the presence of cross-blocking. In practice, cyclical patterns of traffic movement generated by traffic signals may be just as important as the deterministic component of growth considered here.

47. However, other things being equal, the results suggest that a balanced allocation of channelized storage space over a relatively short length of road on each junction approach will reduce the rate of growth of a jam but also speed up gridlock. Once gridlock has occurred,

a jam may be more difficult to disperse. So what should the traffic engineer do? In practice, total gridlock seems to be quite rare, because drivers seem able to keep moving even in quite congested situations. Hence the engineer might reasonably attempt to

(*a*) contain the overall *size* of the jam through a balanced allocation of junction stopline width together with the minimum practicable degree of channelization

(*b*) develop a capability for preventing turning movements using variable traffic signal aspects when gridlock threatens.

48. In order to make further progress, models are needed for highlighting potential trouble spots on real networks with irregular street patterns. This paper has sketched out a possible formulation, but if it is to be implemented effectively, further research will be needed into the cyclical effects of traffic signal control, particularly on short links, and the precise way in which traffic queues interfere with one another under jam conditions.

Acknowledgements

49. The authors are grateful to the Rees Jeffreys Road Fund and the Engineering and Physical Sciences Research Council for supporting this research.

References

1. Wright C. C. and Huddart K. W. Strategies for urban traffic control. *Transport and Society Discussion Paper 6*, Rees Jeffreys Road Fund, Oxford, 1989 (unpublished).
2. Wright C. C. and Huddart K. W. Breaking the grip of gridlock. *Surveyor*, 1989, **172** (5052), 16–17.
3. Roberg P. Development and dispersal of area-wide traffic jams. *Traffic Engineering & Control*, 1994, **35**, No. 6, 379–385.
4. Wright C. C. and Orenstein P. The conceptual structure of traffic jams. *Transport Policy*, 1998, **5**, 23–35.
5. Abbess C. R. and Roberg P. Fractal structure of traffic jam images. *Proceedings of an IMA Conference on Complex Stochastic Systems in Engineering* (Titterington D. M. (ed.)). 1995, vol. 54, Clarendon Press, Oxford, pp. 251–281.
6. Roberg P. R. and Abbess C. R. Diagnosis and treatment of congestion problems in central urban areas. *European Journal of Operational Research*, 1998, **104**, No. 1, 218–230.

Please email, fax or post your discussion contributions to The Publisher: email: Wilson_l@ice.org.uk; fax: 0171 799 1325; or post to Lesley Wilson, Journals Department, Institution of Civil Engineers, 1–7 Great George Street, London SW1P 3AA.