TRANSPORT and ROAD
RESEARCH LABORATORY

Department of the Environment
Department of Transport

SUPPLEMENTARY REPORT 526

AUTOMATIC INCIDENT DETECTION – TRRL ALGORITHMS HIOCC AND PATREG

by

J F Collins, C M Hopkins and J A Martin

Highway Traffic Division
Traffic Engineering Department
Transport and Road Research Laboratory
Crowthorne, Berkshire
1979
ISSN 0305–1315

CONTENTS

AUTOMATIC INCIDENT DETECTION – TRRL ALGORITHMS HIOCC AND PATREG

ABSTRACT

An Automatic Incident Detection (AID) system, using computer based algorithms to identify the traffic disturbances following an incident, is being developed at TRRL. It incorporates two new algorithms, HIOCC and PATREG.

The HIOCC algorithm operates by detecting stationary or slow moving vehicles over individual vehicle detectors. It looks for several consecutive seconds of high detector occupancy to detect queues and incidents in high traffic flows. Off-line tests with recorded data show HIOCC's detection ability to be very good and a low false alarm rate can be achieved.

The PATREG algorithm measures the average traffic speed between upstream and downstream detector stations using a pattern recognition technique, detecting incidents by identifying significant speed changes. Preliminary tests suggest that it performs satisfactorily in free-flow conditions up to about 1500 veh/h per lane, and has an acceptably low false alarm rate.

By incorporating both HIOCC and PATREG in an AID system a wide range of traffic conditions can be catered for. A new AID site on the M1 in Bedfordshire will be used to complete the development and testing of the algorithms.

1. INTRODUCTION

The rapid detection of traffic incidents is very desirable on a motorway so that warnings can be transmitted and traffic control measures taken before secondary accidents occur.

The automatic incident detection system under study at TRRL relies upon the rapid identification of the disturbances to the traffic flow which invariably follow an incident on the carriageway. This is also the approach being followed in other countries and, at present, no one system is available which is entirely satisfactory in operation under all traffic conditions – particularly with regard to the occurrence of false alarms in incident-free traffic. The key to the problem lies in the detection of the traffic perturbations that follow an incident whilst ignoring the normal variations in the traffic flow which occur continually. The basic detection system includes a small computer (or a number of microprocessors distributed throughout the system) which continuously examines data from a sequence of vehicle detectors spaced along the carriageway, typically 500 metres apart in each lane. Mathematical processes, or algorithms, are used by the computer to detect the occurrence of an incident by looking for particular disturbances in the data stream; obviously, the performance of a system depends mainly upon the efficiency of its algorithms in detecting traffic incidents without creating false alarms.

TRRL have developed two algorithms: HIOCC and PATREG. This report describes the operation of these two algorithms and gives some details of their performance during preliminary testing. More comprehensive testing, followed by an operational trial, will be carried out on an experimental site on the northbound carriageway of the M1 motorway in Bedfordshire.

2. THE HIOCC ALGORITHM

The HIOCC (HIgh OCCupancy) algorithm operates by identifying the presence of stationary or slow moving vehicles over individual vehicle detectors, usually inductive loops. This is achieved by looking for several consecutive seconds of very high detector occupancy and, when this occurs, initiating an alarm and giving the location.

The output from each detector in an installation is scanned by the computer at one-tenth second intervals to determine whether or not the detector is occupied. At the end of each second the number of times a detector has been found to be occupied is taken as a measure of the one-second occupancy for that detector, with values in unit steps from 0 to 10 (approximating to 0 and 100 per cent occupancy respectively). This value of one-second occupancy is termed the instantaneous occupancy. Figure 1 illustrates how the instantaneous occupancy can vary second by second for a detector in normal traffic and during an incident – this example uses data from a busy Paris motorway. To find out when stationary or slow moving vehicles are present these consecutive values of instantaneous occupancy are checked to see if a pre-determined threshold has been reached; a threshold value of 10, ie 100 per cent occupancy, lasting for two consecutive seconds is often appropriate.

In practice, because of the random occurrence in time of the detector output pulses with respect to the regular one-tenth second sampling process and its associated one second measuring interval – both controlled by the computer – only some of the vehicles giving occupancy pulses of between 2.0 and 2.8 seconds will reach the threshold, whereas all vehicles occupying the detector for 2.9 seconds or more will be 'caught'. In the case of a 3.7 m (12 ft) long detector loop this means that all small cars crossing the detector loop at a speed of 10 km/h (6 miles/h) or less, and all large articulated lorries passing at 23 km/h (14 miles/h) or less will cause an alarm with this threshold.

Each second, after monitoring the instantaneous occupancy values, smoothed values of occupancy are recalculated for each detector. These are used to determine the end of the alarm, as described later. The single-stage exponential smoothing of the instantaneous occupancy values is computed by the method:

New smoothed value = P x present instantaneous value + (1–P) x previous smoothed value

where P, the smoothing factor, is positive and less than one. A typical value for P is 1/64, (using a power of 2 simplifies the computing process). The effect of such smoothing on the instantaneous values of occupancy is shown in Figure 1.

Several important features are included in the HIOCC algorithm. The first is the artificial raising of the smoothed occupancy value to a high level, usually 90 per cent, at the onset of an alarm; this prevents the switching on and off of the alarm signal due to fluctuation of the smoothed occupancy value around the level used to terminate the alarm state. This raising of the smoothed occupancy value biases subsequent values and helps to ensure that the alarm is not terminated prematurely. Thereafter, continuing from this high value, a new smoothed value is recalculated every second in the normal manner.

The second feature concerns the termination of the alarm state. In principle, the HIOCC algorithm terminates an alarm when the value of the smoothed occupancy falls to the level which existed before the incident occurred. This pre-alarm level is taken as the average of the smoothed occupancy values at the end of each of the five oneminute periods preceding the alarm; this is done in case the smoothed occupancy reached an abnormal value immediately before the alarm was first signalled. In certain situations, the alarm state can also be terminated when the smoothed occupancy reverts to a predetermined level. An example is the case when a queue forms in a phase of below-average traffic flow; in such a situation the duration of the alarm may be unnecessarily prolonged by waiting for the smoothed occupancy to drop to the lower-than-normal pre-incident level. This predetermined level is chosen to correspond to the normal traffic flow for the particular site.

The third feature incorporated in the algorithm is designed to maintain the continuity of the alarm during the stop-go movement of a queue. In this situation vehicles can often stop clear of the detectors, leading to a reduction in the smoothed occupancy value and sometimes causing a premature termination of the alarm. This is avoided by suspending the calculation of the new value of smoothed occupancy whenever the instantaneous occupancy is zero for a continuous period of, say, 8 seconds during an alarm state. Up-dating of the smoothed occupancy is restarted when a vehicle passes over the detector again.

Figure 2 shows the operation of the algorithm on data collected during an incident on the Boulevard Périphérique in Paris. It illustrates the raising of the smoothed occupancy value to 90 per cent at the onset of the alarm, and the suspension of its recalculation during a period of zero occupancy. The values of the real smoothed occupancy during the incident are shown by the dashed line.

A flow diagram for the HIOCC algorithm is given in Figure 3.

3. THE PATREG ALGORITHM

The PATREG algorithm has been developed to monitor the traffic speed in each lane between pairs of detector stations in an incident detection installation, and to indicate the occurrence of an incident by sensing significant changes in speed. As for the HIOCC algorithm, each detector station is equipped with a single detector per lane. The measurement of the traffic speed is based on a pattern recognition technique using data from sequential detectors spaced along each traffic lane, the separation being about 500 metres.

3.1 *Calculation of traffic speed*

The speed calculation technique is based on the presumption that under steady, incident-free conditions, there is some recognisable continuity in the traffic patterns at two consecutive detector stations. In the ideal case, the patterns of flow (or occupancy) with time, obtained simultaneously from an upstream and from a downstream detector, would be found to match exactly if the downstream pattern were advanced by a time interval equal to the journey time between the detectors. In practice, allowance must be made for some change in the flow pattern between detector stations.

Given the flow patterns from the upstream and downstream detector stations in a lane, expressed as instantaneous flow*, the following cross-correlation procedure is implemented every second on the two series of flow

* Note: 'Instantaneous flow' is used here to indicate the flow parameter measured over short, successive periods of time, such as a second, with the latest value being available at the end of its measurement period.

values to determine the relative time displacement giving the best match between them:

1. The 40 most recent values of upstream instantaneous flows are stored for use in the calculation. Every second these are updated as a new upstream flow becomes available.

2. The quantity DU(I), the product of the latest downstream instantaneous flow and the upstream instantaneous flow for I seconds ago, is calculated for each of the 40 stored upstream values, giving 40 values of DU(I) corresponding to a range of I from 1 to 40 seconds ago.

3. Single exponential smoothing is then used on the DU(I) values to give the weighted sums, MATCH(I)s, of the present and past values of DU(I) for each time I, ie

   New MATCH(I) = Q x present value DU(I) + (1–Q) x old MATCH(I)

   where Q, the smoothing factor, has a typical value of 1/128.

The above procedure gives a set of 40 values of MATCH(I); each value corresponds to a particular time displacement I, between the two flow patterns, with I ranging from 1 to 40 seconds. Ideally, according to this cross-correlation technique, the maximum value of MATCH(I) in this set would correspond to that value of I which is the relative time displacement between the two flow patterns, ie the journey time. However, as vehicle speeds are not usually the same, or constant, between measuring sites, the resulting spread in journey times is allowed for in the computation by an additional process using a set of 13 weighting factors, which are derived from journey time data.

These weighting factors, spread over 13 consecutive MATCH(I)s, have a symmetrical distribution with values of 1, 2, 4, 6, 8, 9, 9, 9, 8, 6, 4, 2, and 1. This group of weights is stepped along the 40 MATCH(I) values and at each step the 13 products (weight x MATCH(I) value) are totalled to give a SUM(I). The actual journey time, I seconds, is derived from the highest value of this series of SUM(I)s because this occurs when the weighting factors are distributed symmetrically about the MATCH(I) value that corresponds to the journey time. Because MATCH(I) has 40 values and the weighting factors 13, the range of journey times covered at present is 7 to 34 seconds. The journey times are converted to speeds for convenience, and can be displayed at any chosen time interval from one second upwards.

3.2 *Alarm thresholds*

The second part of the PATREG algorithm deals with the inspection of the speed values for indication of a traffic incident. It compares the traffic speeds calculated for each lane with predetermined lower and upper thresholds for that lane. If the speed lies outside these threshold values for a pre-set number of consecutive seconds – the persistence time – an alarm is given and the identification numbers of the two detectors involved are displayed to show the incident location.

The speed thresholds are chosen by using PATREG to measure traffic speeds in each lane, for different traffic flows during incident-free conditions. Having established the normal, or incident-free, range of speeds, the thresholds are set to 'catch' speeds falling outside this range. In principle it should not be necessary to have an upper threshold as vehicles would not be expected to increase their speed because of an incident. However, the upper speed threshold is required because PATREG, when it is unable to obtain a definite pattern match between the upstream and downstream data, goes either to the top or bottom limit of its journey time range and gives the corresponding speed – 111 miles/h (179 km/h) or 26 miles/h (42 km/h) for a station separation of 400 metres. The inability of PATREG to find a pattern match over a period of time is likely to be an indication of congestion or of some incident disturbing the traffic sufficiently to disrupt the pattern between measuring stations.

Typical threshold speeds found by PATREG on an urban motorway near London are:

Nearside lane	Lower limit	=	40 miles/h
	Upper limit	=	59 miles/h
Centre lane	Lower limit	=	49 miles/h
	Upper limit	=	74 miles/h
Offside lane	Lower limit	=	55 miles/h
	Upper limit	=	89 miles/h

In practice, an upper threshold of 89 miles/h can be used for all three lanes. The introduction of the persistence time helps to ensure that the alarm signal is genuine. Preliminary off-line tests on data from free-running motorway traffic suggests that, for a range of flows, the persistence time needed to avoid false alarms is about 20 seconds.

A basic flow diagram for PATREG is shown in Figure 4. Unlike the HIOCC algorithm PATREG has no features for signalling the continuation or termination of an incident. It is envisaged that PATREG would be used together with HIOCC, or other algorithms, so that after the initial alarm due to the operation of PATREG, these would subsequently take over and cater for the later parts of the incident. The combination of algorithms is discussed in Section 4.4.

4. DEVELOPMENT AND TESTING OF ALGORITHMS

The principle operational requirements of an incident detection algorithm are:

(i) The reliable indication of a significant traffic incident and its location.

(ii) An acceptable false alarm rate; a false alarm may occur during free flowing, incident-free traffic if the algorithm identifies a chance distribution of traffic parameters as an incident.

(iii) A short response time, the response time being the interval between the start of an incident and its detection by the algorithm.

During algorithm development the extent to which a particular algorithm meets these requirements can be conveniently measured by running the algorithm off-line with recorded motorway traffic data. Sets of recorded data corresponding to different traffic situations are, therefore, needed to make repeated runs with different values of algorithm parameters in order to optimise the latter. Unfortunately, if it is not possible to stage incidents on motorways, the acquisition of real incident data of the right kind, and in sufficient quantity, can be very difficult and time-consuming unless an instrumented length of motorway is available with equipment for the continuous automatic recording of traffic data. Such a facility is now being installed at the TRRL experimental site on the M1 motorway in Bedfordshire, in order to obtain the necessary data to complete algorithm development. Subsequently, the algorithms will be tested on-line in an experimental incident detection system on the same site.

4.1 *Existing test data*

The present algorithm development and preliminary testing has been done using two sets of traffic data, the first from the M4 motorway near London and the second from the Boulevard Périphérique in Paris.

The M4 motorway site at Osterley was instrumented with four detection stations, spaced at 400 metres, each station having one detector in each of the three lanes. Because the monitoring equipment required the supervision of an operator, continuous recording of data was not possible. The site was monitored at intervals over a period of several months for a total of more than 300 hours. Unfortunately, on this short site, no traffic incidents occurred during the measuring periods and it was not possible to arrange for staged incidents. However, two kinds of useful data were obtained; on a number of occasions recordings were made as vehicles queued through the monitoring site because of downstream congestion; and also data from free-running traffic at various flow levels was recorded. The queueing data was used to test the algorithms for their ability to detect queues, whilst the free-flow data was necessary for studies concerned with the occurrence of false alarms. The latter data covered a range of carriageway flows from 128 to 3200 veh/h.

The Boulevard Périphérique traffic data[1] was obtained by a team from TRRL co-operating with two French organisations, the Institut de Recherche des Transports, and the Préfecture de Paris, who were able to organise a number of staged incidents. Data was recorded by TRRL personnel before, during and after 12 incidents, from two sets of detectors spaced 530 metres apart. The incidents were located in four different positions relative to the detectors and each lasted between 3 and 12 minutes, always causing queues to form in the heavy traffic. These incidents have so far provided the only available data for the off-line testing of algorithms in relation to their detection of incidents; unfortunately the data covers a limited range of traffic conditions. A small amount of free-flow data was also recorded from three of the four traffic lanes. The combined flows of these three lanes was approximately 5000 veh/h.

4.2 *Off-line testing of HIOCC*

To examine its likelihood of giving false alarms, off-line tests were carried out on HIOCC with the sets of incident-free traffic data, to determine the numbers of false alarms that occurred for different thresholds of the two major algorithm parameters, ie the percentage of instantaneous occupancy and the number of consecutive seconds at this percentage occupancy before an alarm signal is given. For example, Figure 5 shows the effect, given a persistence time of 2 seconds, of reducing the occupancy from 100 per cent to 50 per cent, on the number of false alarms occurring in the M4 incident-free data for six levels of flow.

For the above sets of data the detection threshold could be set at 70 per cent occupancy for 2 seconds without giving any false alarms. However, more recent tests with some preliminary data from the experimental site on the M1, have indicated that the threshold value should be raised to 100 per cent occupancy for 2 seconds to ensure that false alarms are avoided on this site. This was not unexpected as the traffic conditions are different on the M1; it has a slight uphill gradient and, in addition, there is a higher proportion of heavy goods vehicles, both factors tending to increase the one-second occupancy values. It is appropriate, therefore, to use a threshold of 100 per cent for 2 seconds in all cases.

Regarding the detection of queues or incidents, it was found that the HIOCC algorithm, using this threshold, gave an alarm signal for all the occasions when queues formed through the M4 site and for each of the 12 staged incidents on the Boulevard Périphérique. The response times for the detection of the staged incidents ranged from 20 seconds to 2 minutes 10 seconds.

4.3 *Off-line testing of PATREG*

The incident-free traffic data from the M4 was used also with PATREG to examine the effect on false alarms of changing the smoothing factor, Q, in the MATCH calculation. To exaggerate the occurrence of false alarms for test purposes the upper and lower speed thresholds for each lane were set at levels which would be certain to give false alarms and, in addition, the persistence time was set to one second so that no alarms were suppressed. Figure 6 shows the result of varying Q for a range of flow values. A smoothing factor of 1/128 seemed to provide a satisfactory compromise between the number of alarms produced and a reasonable response time for the algorithm; the confirmation of this value will be the subject of further work when incident data has been received from the M1 site.

Further off-line testing of PATREG has shown that its pattern recognition technique breaks down when the flow exceeds about 1500 veh/h per lane. It is presumed that under heavy flow conditions the traffic pattern is too random in character to provide a recognisable feature between measuring stations. Because of this, PATREG did not detect the staged incidents in the heavy traffic on the Boulevard Périphérique, although the HIOCC algorithm was successful in all cases. Because no suitable incidents have been recorded from the M4, PATREG has not been adequately tested over a range of flows for its ability to detect incidents and such tests are included in the future work programme.

4.4 *Combination of algorithms*

The results of work to date, both abroad and at TRRL, suggest that no single algorithm with fixed parameters will operate efficiently in an automatic incident detection system under all conditions of traffic from low to high flow levels. It seems likely that a combination of two or three different algorithms will be necessary, the appropriate one coming into action according to the traffic conditions. It is believed that the two TRRL algorithms, HIOCC and PATREG, will cater for a wide range of traffic conditions. The best technique for combining the two algorithms is the subject of further study.

5. M1 MOTORWAY TEST SITE

As referred to previously, an experimental incident detection installation is nearing completion on the M1 in Bedfordshire. This is part of an instrumented site prepared for other experiments. The detection installation will provide data from various kinds of traffic incident over a range of vehicle flows. Seven monitoring stations are installed along the northbound carriageway at a spacing of approximately 550 metres. The layout is shown in Figure 7. A pair of closely spaced detectors is installed in each lane at each station as part of the equipment for the overall site, where there is a requirement for the measurement of individual vehicle spot speeds in the driver behaviour studies. However, data from only one detector of each pair is used by the TRRL algorithms. The data recording system is automatic and operates unattended. The output from the 21 detectors is logged on a magnetic tape cartridge which can store up to three days' traffic data. Data will therefore be available for a wide range of traffic flows and include the periods before, during and after traffic incidents.

6. CONCLUSIONS

The HIOCC algorithm, which detects stationary or slow-moving vehicles to indicate a traffic queue caused by an incident or by congestion, has shown good results at two very different sites. Because its operation depends upon a queue being formed, the traffic flow has to be sufficiently high for incidents to cause queueing; however, the queue need only cover a single detector loop.

The PATREG algorithm, which indicates incidents by identifying significant speed changes, does not depend upon the formation of a queue, and appears suitable for use in traffic flows up to about 1500 veh/h per lane. The ability of PATREG to detect incidents needs to be investigated more thoroughly and the experimental site on the M1 motorway in Bedfordshire will provide data for this purpose. Tests of sensitivity of the algorithm parameters with respect to the occurrence of false alarms, have suggested that, over a range of traffic conditions, it is possible to select parameter values which keep false alarms at a low level.

A combination of these two algorithms in an incident detection system should permit most traffic conditions to be covered.

7. ACKNOWLEDGEMENTS

The work described in this report was carried out in the Highway Traffic Division (Division Head: Mr K Russam) of the Traffic Engineering Department of TRRL.

8. REFERENCE

1. COLLINS, J F. Incident detection experiments on the Boulevard Périphérique in Paris. *Department of the Environment Department of Transport, TRRL Report* SR 362. Crowthorne, 1977 (Transport and Road Research Laboratory).

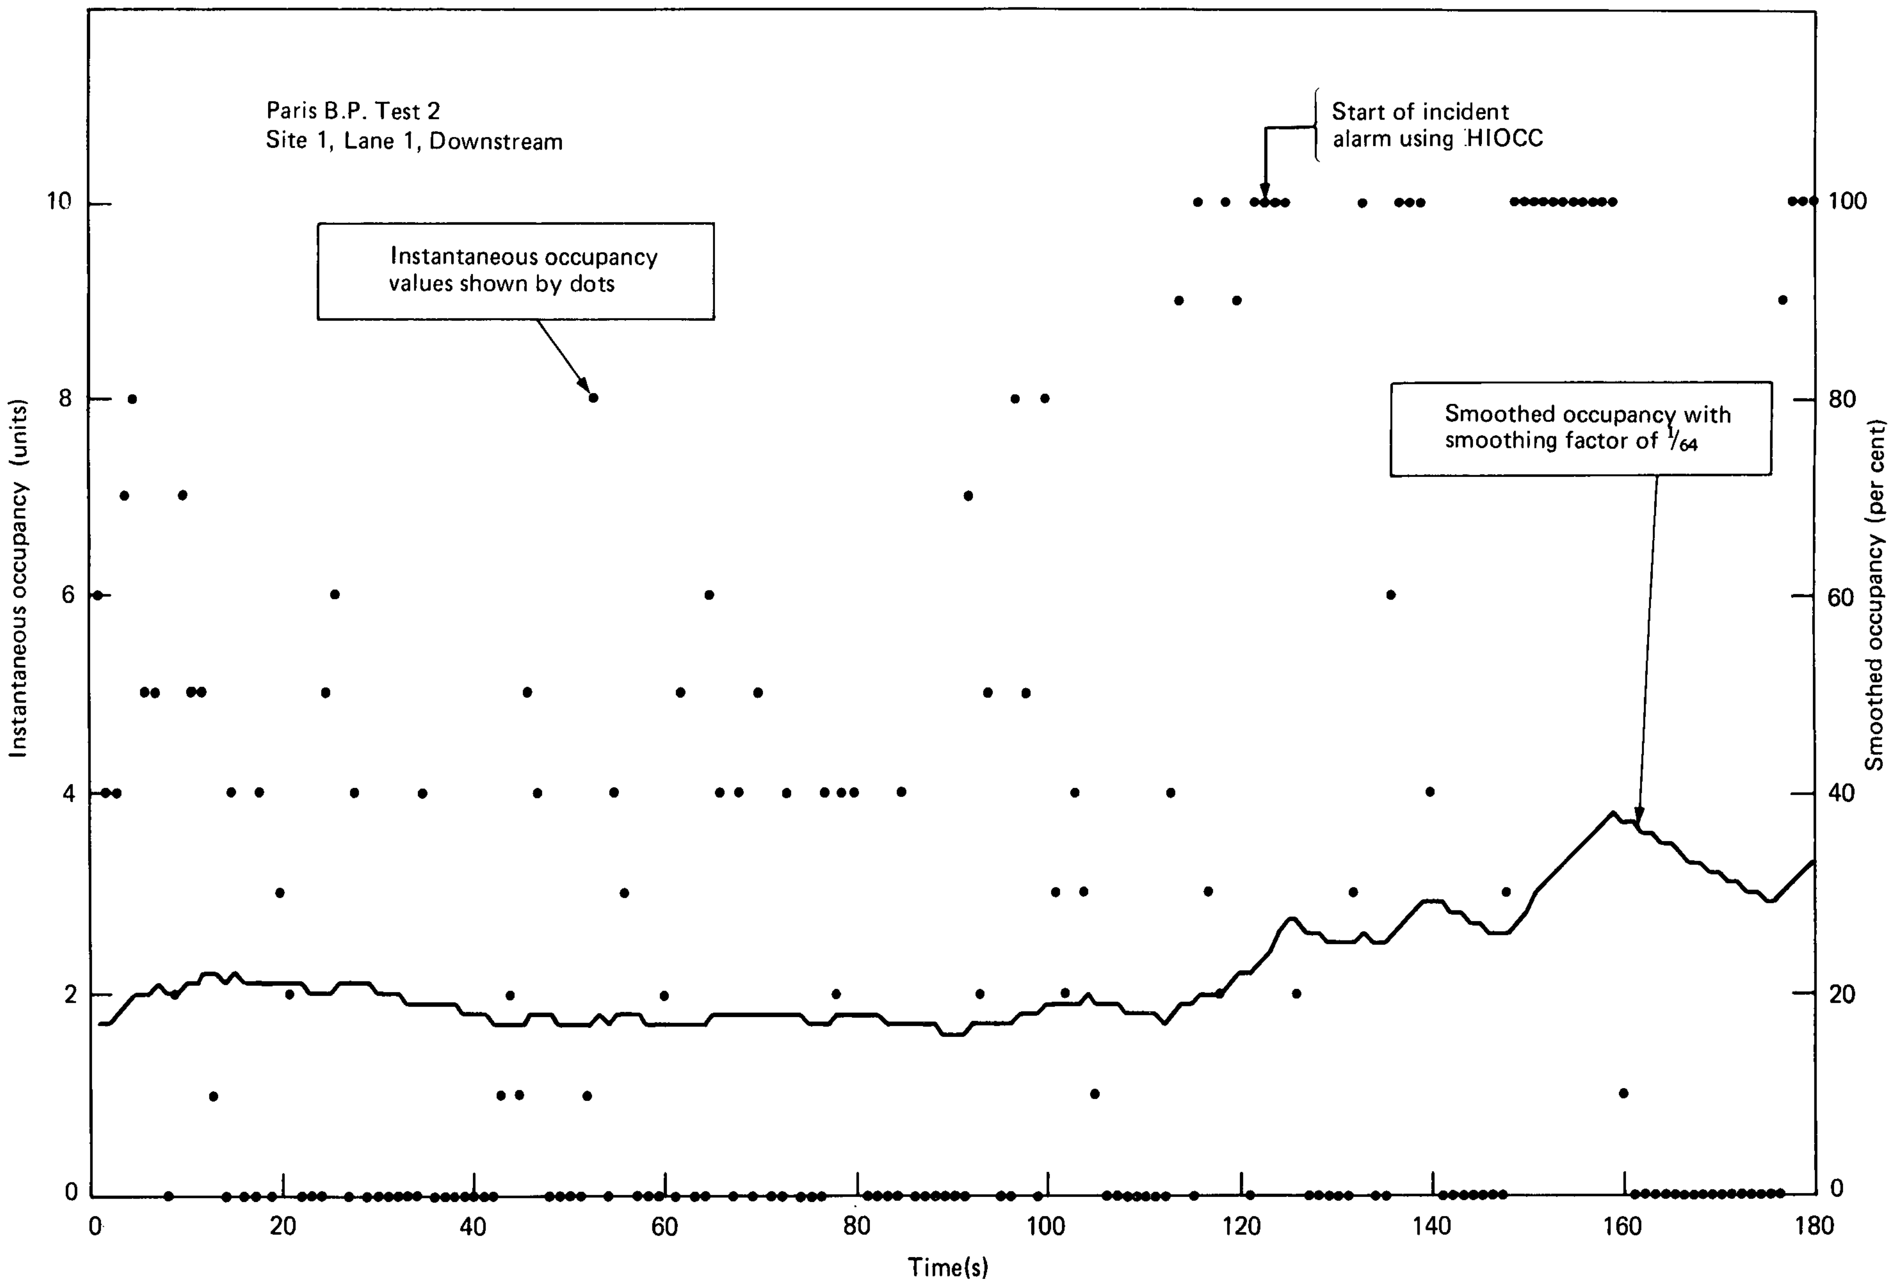

Fig. 1 CORRESPONDING VALUES OF INSTANTANEOUS OCCUPANCY AND SMOOTHED OCCUPANCY

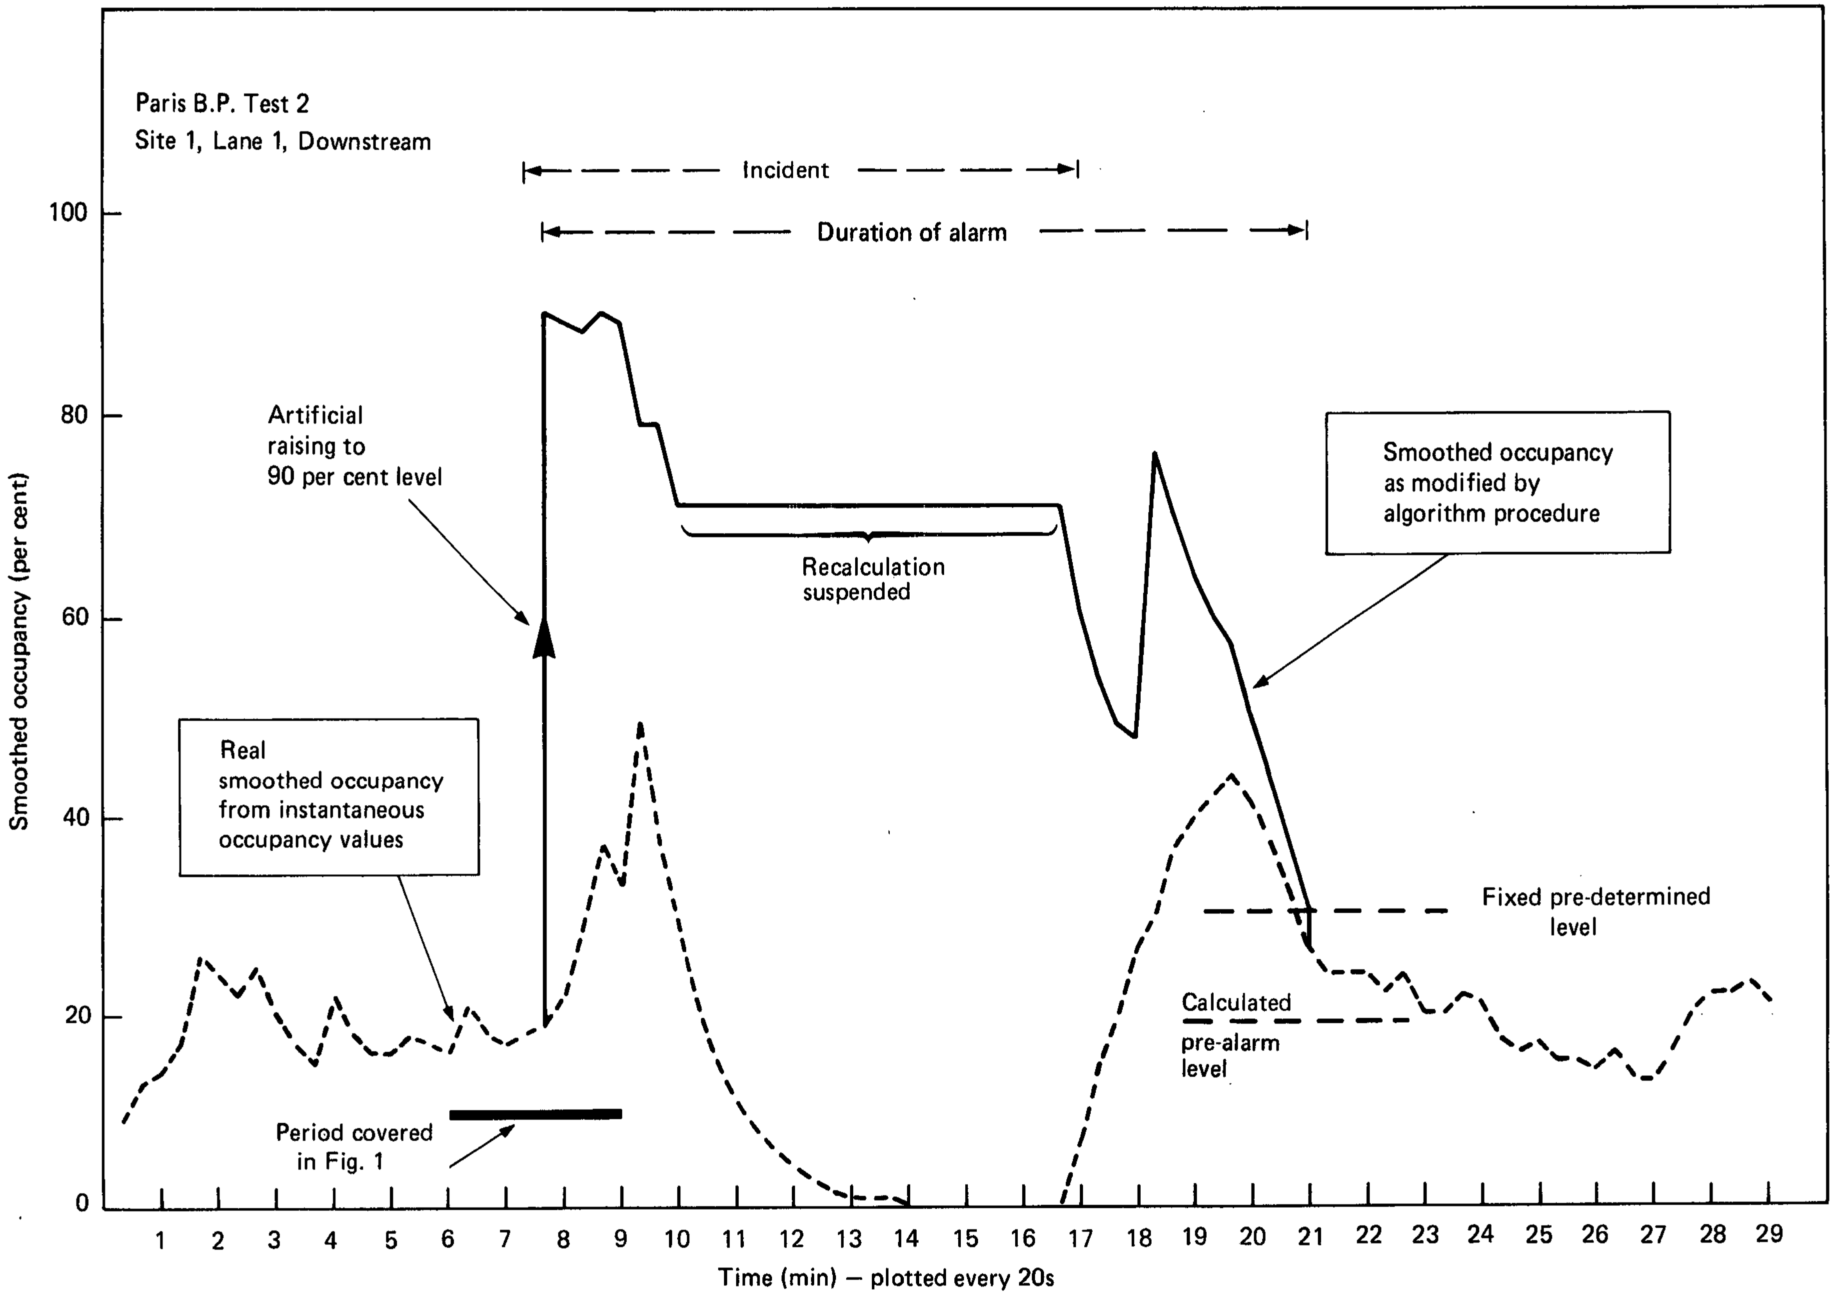

Fig. 2 SMOOTHED OCCUPANCY – REAL VALUE AND VALUE AFTER MODIFICATION IN HIOCC ALGORITHM

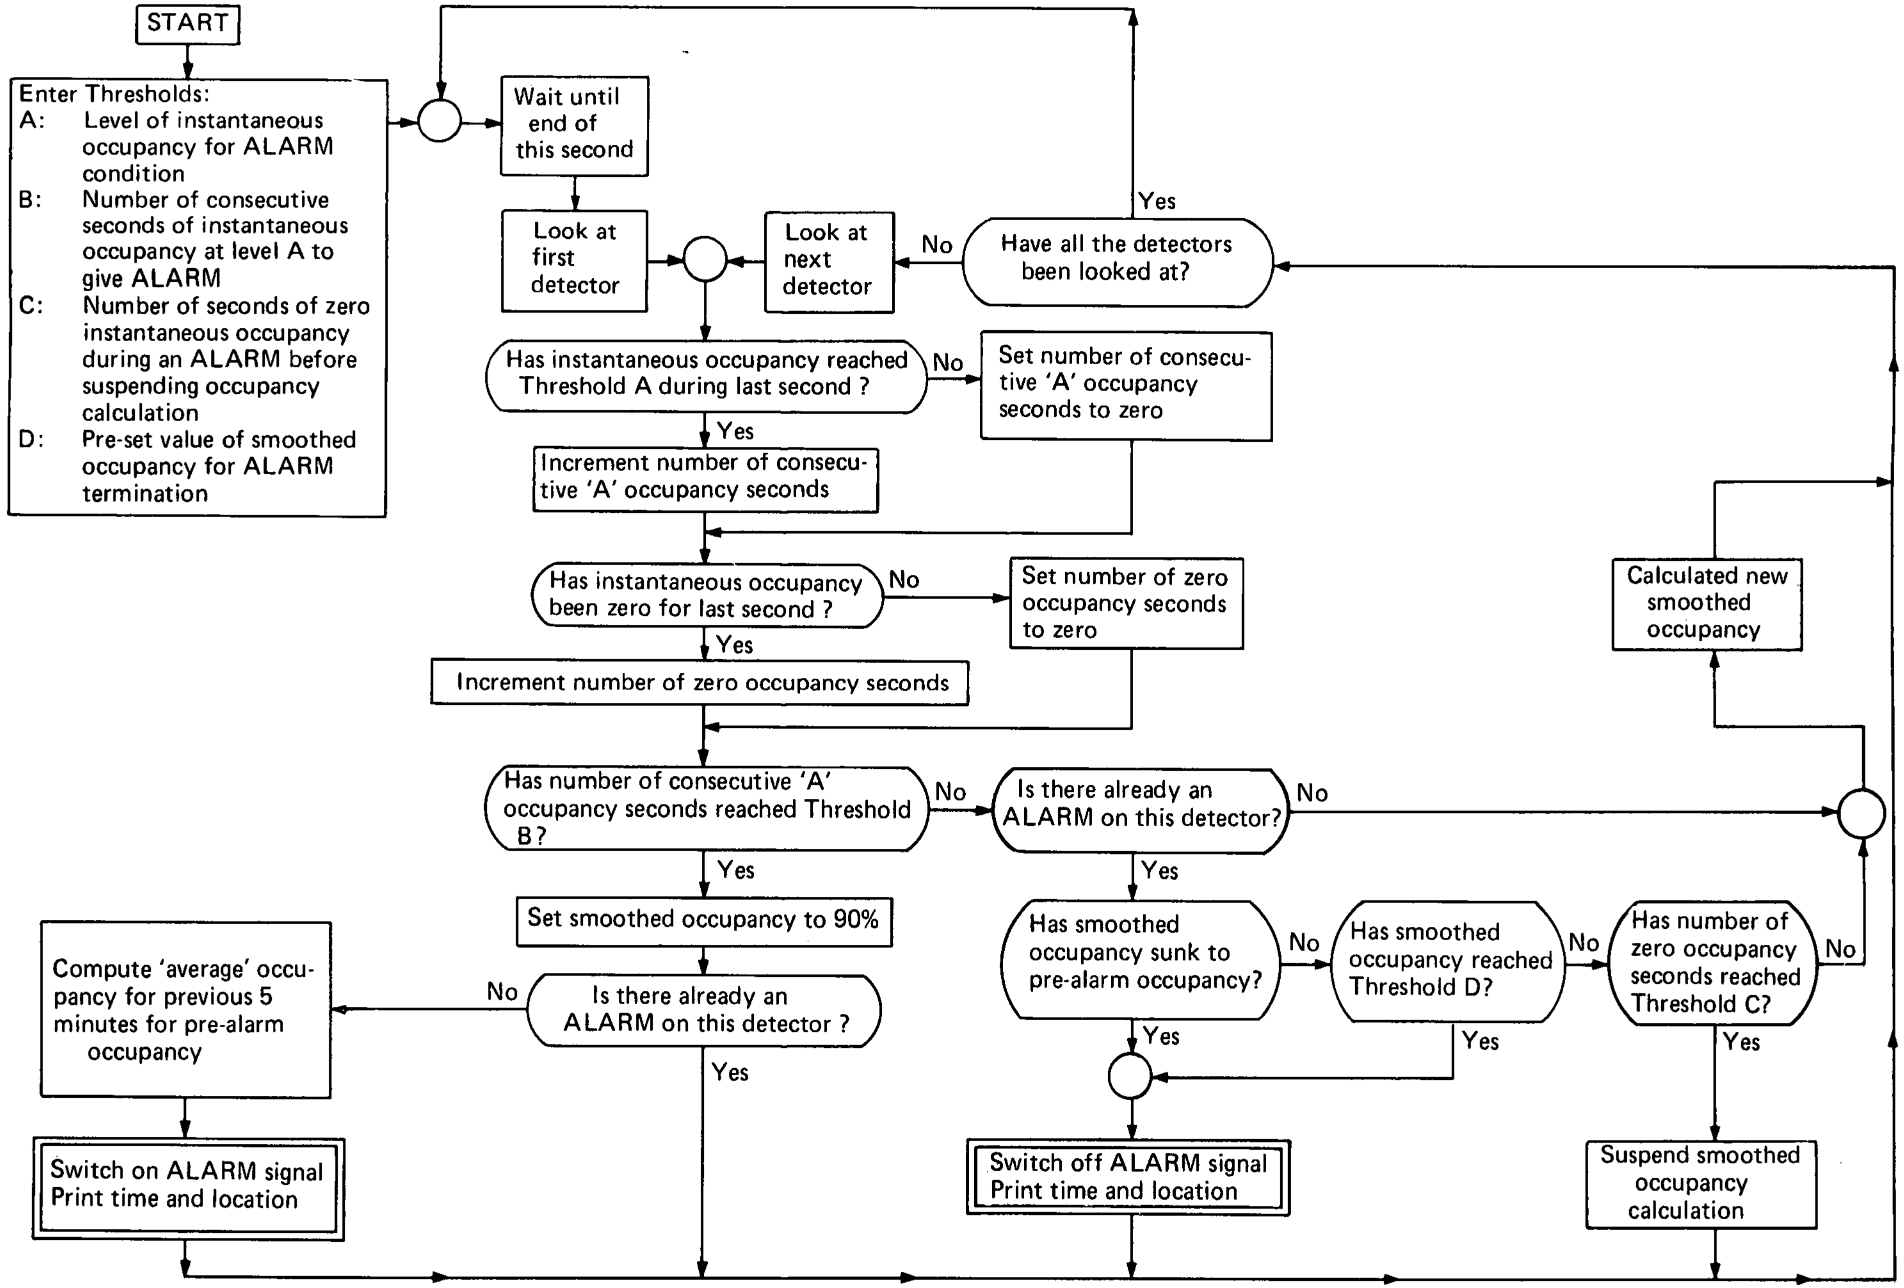

Fig. 3 HIOCC ALGORITHM FLOW DIAGRAM

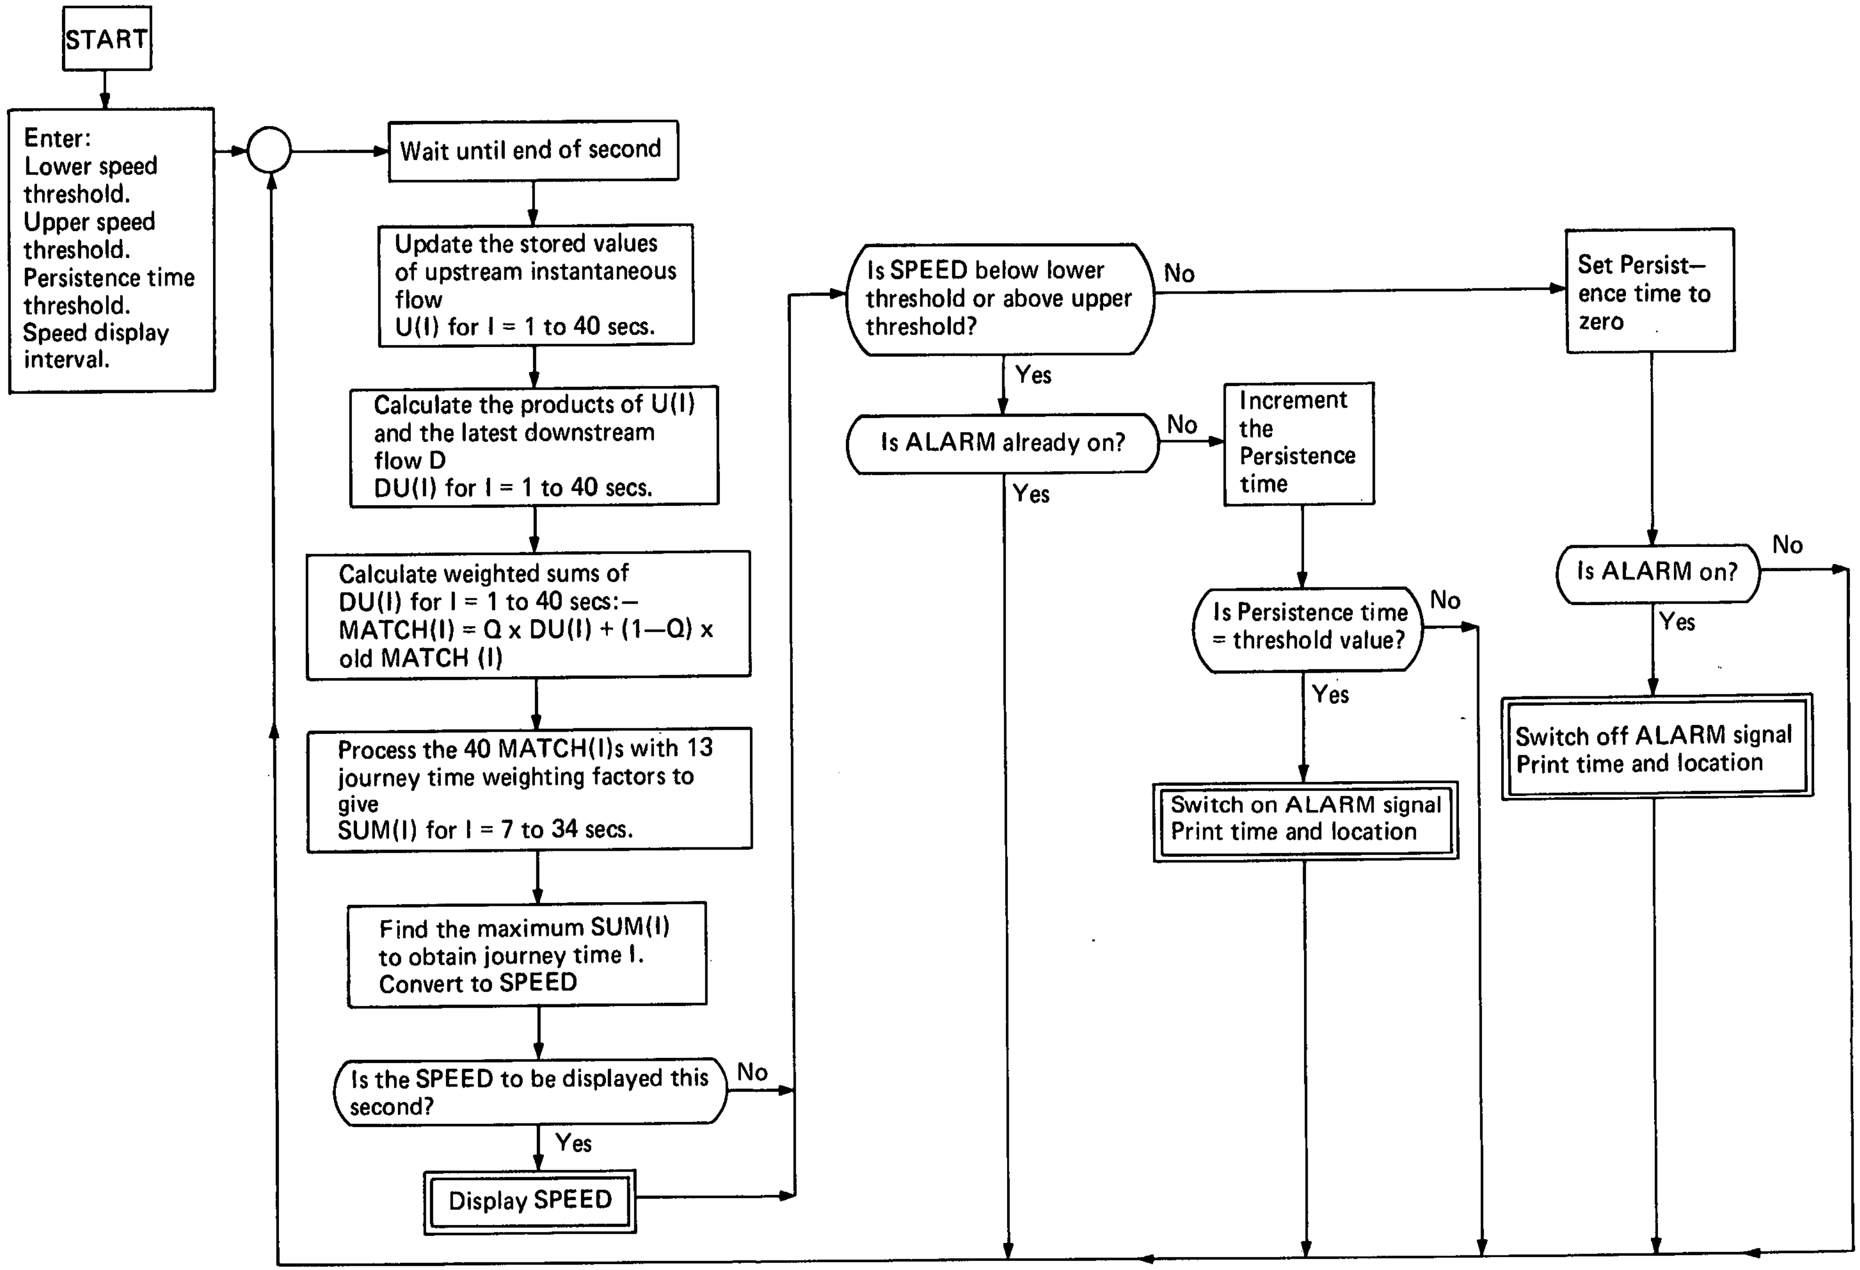

Fig. 4 PATREG ALGORITHM FLOW DIAGRAM

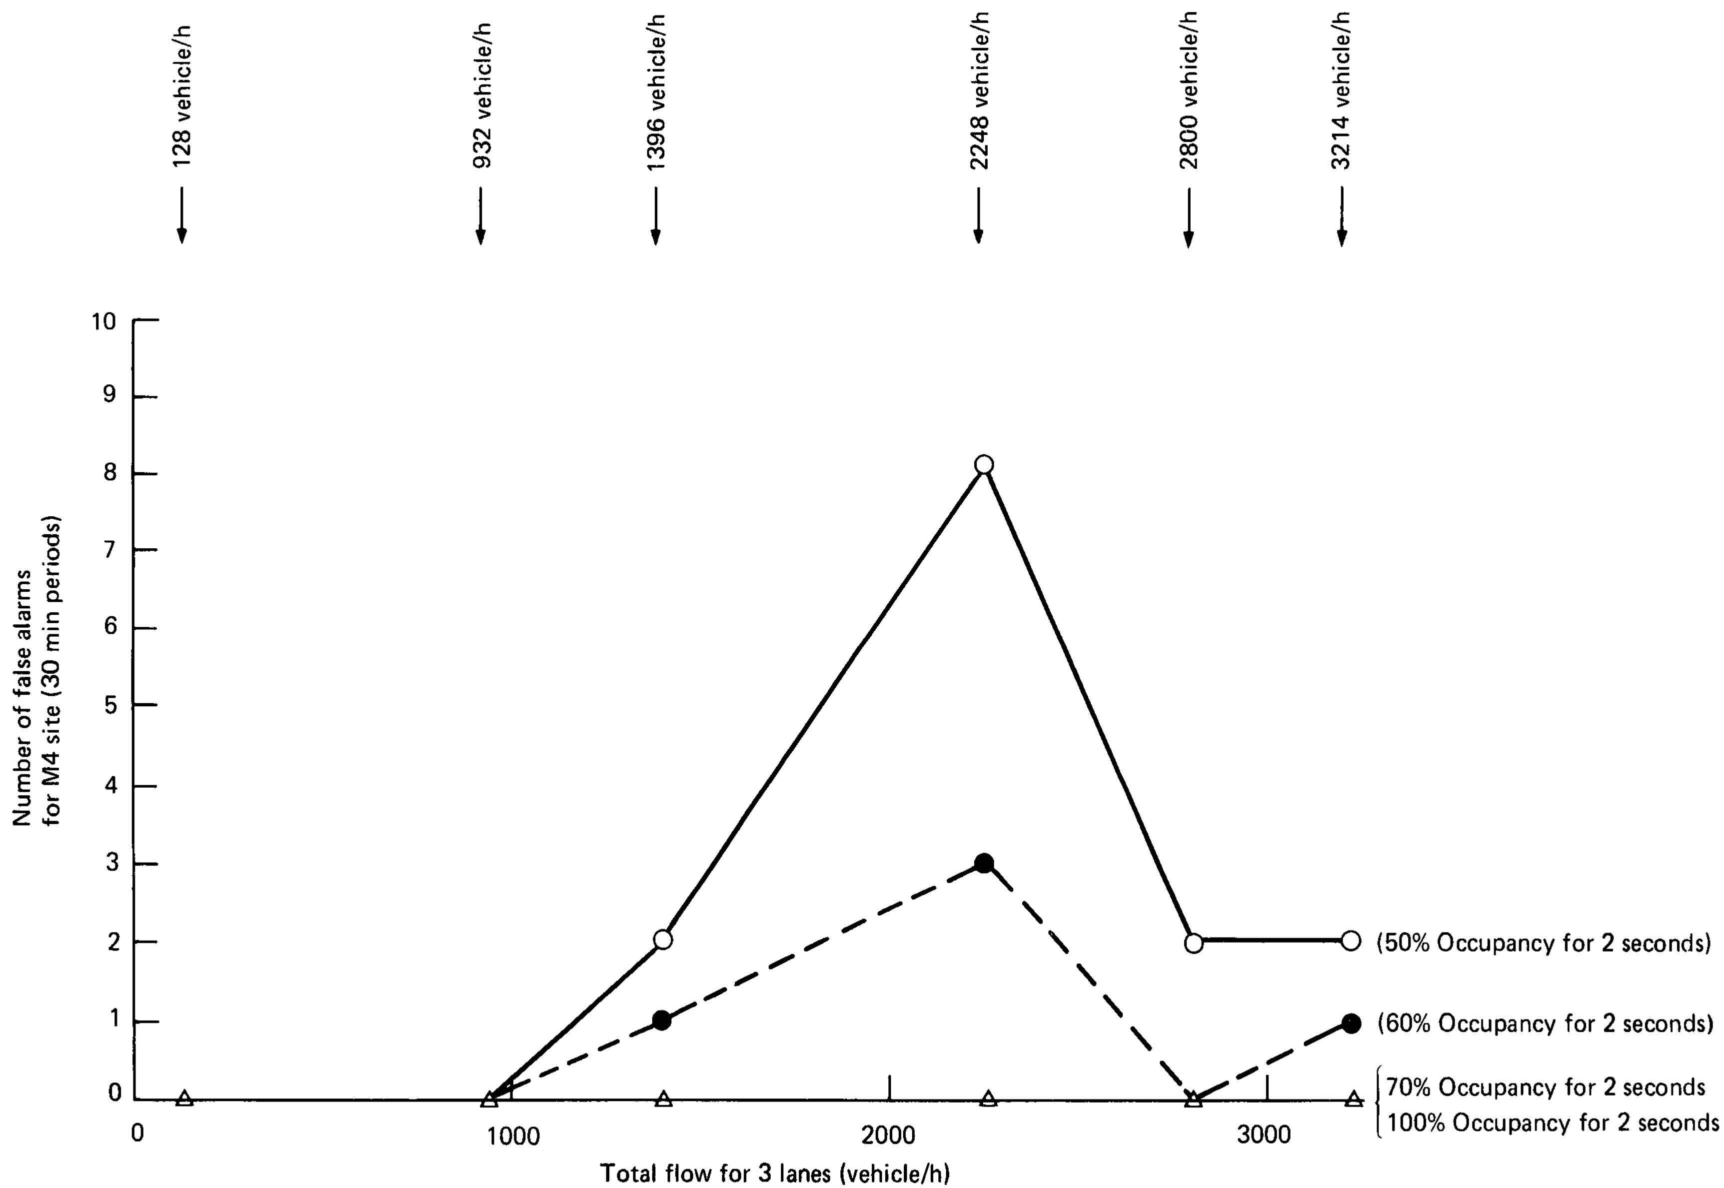

Fig. 5 HIOCC FALSE ALARM TESTS – THE EFFECT OF VARYING THE PERCENTAGE OCCUPANCY THRESHOLD ON THE NUMBER OF FALSE ALARMS

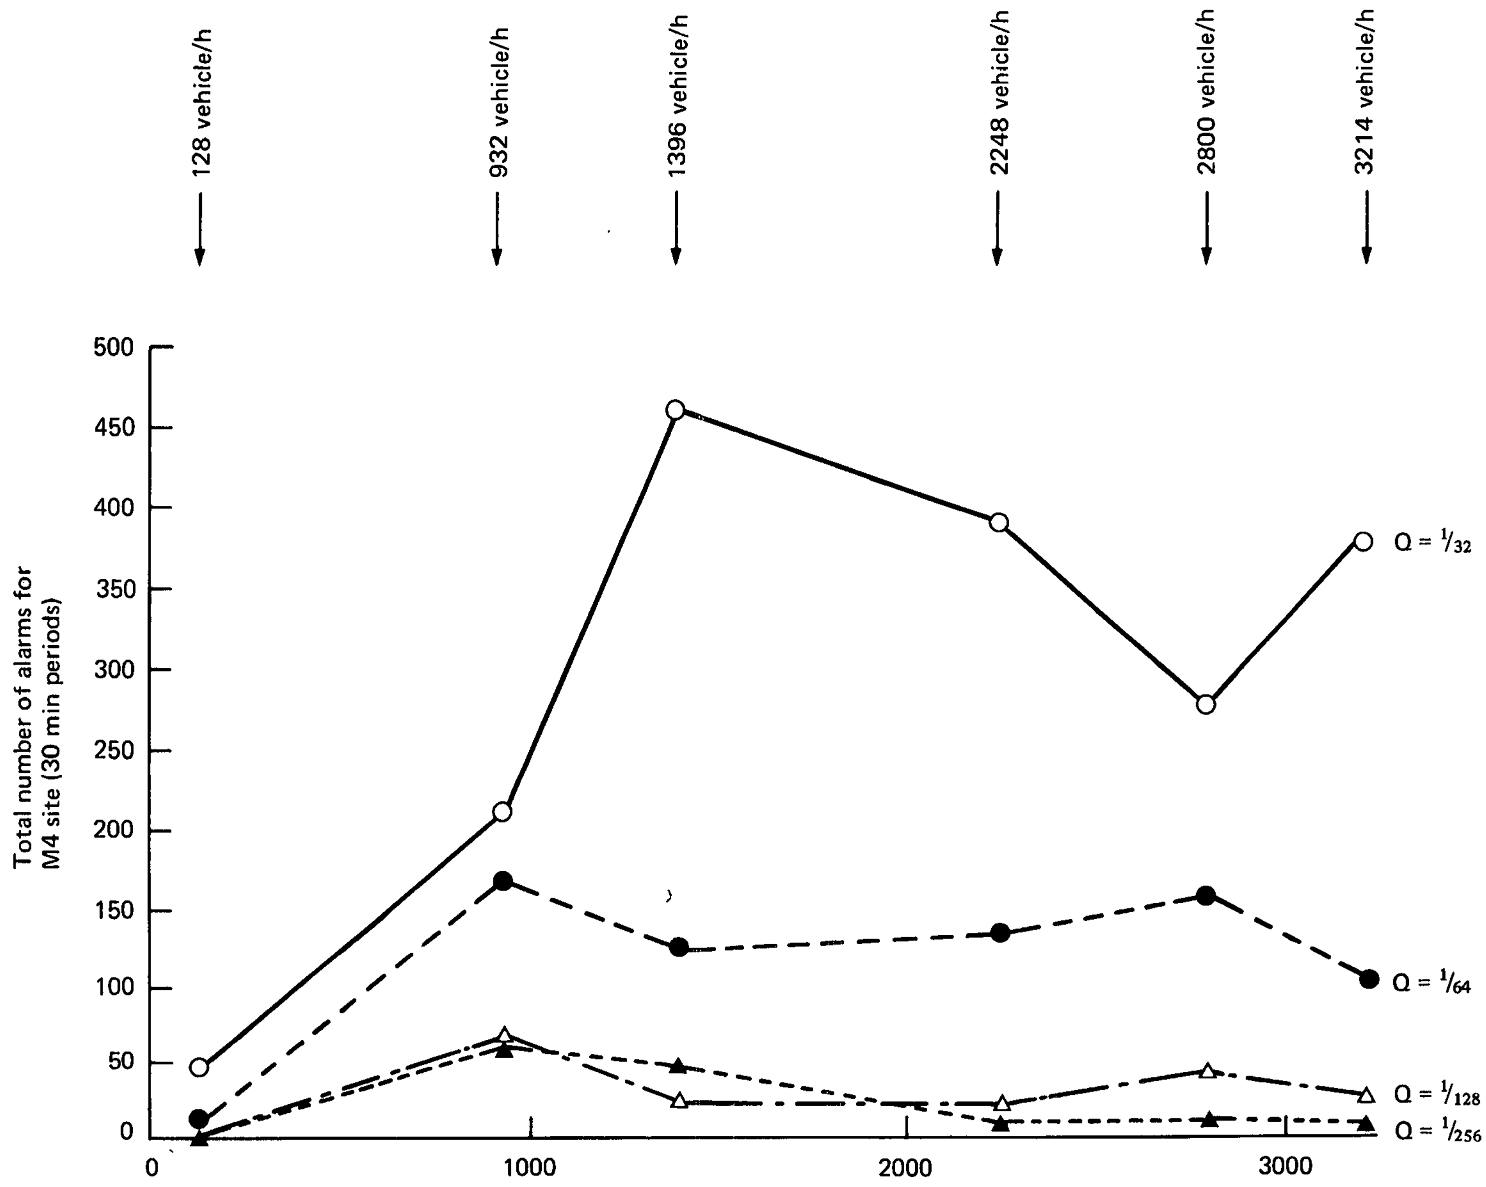

Fig. 6 PATREG FALSE ALARM TESTS – THE EFFECT OF VARYING THE SMOOTHING FACTOR Q OF MATCH

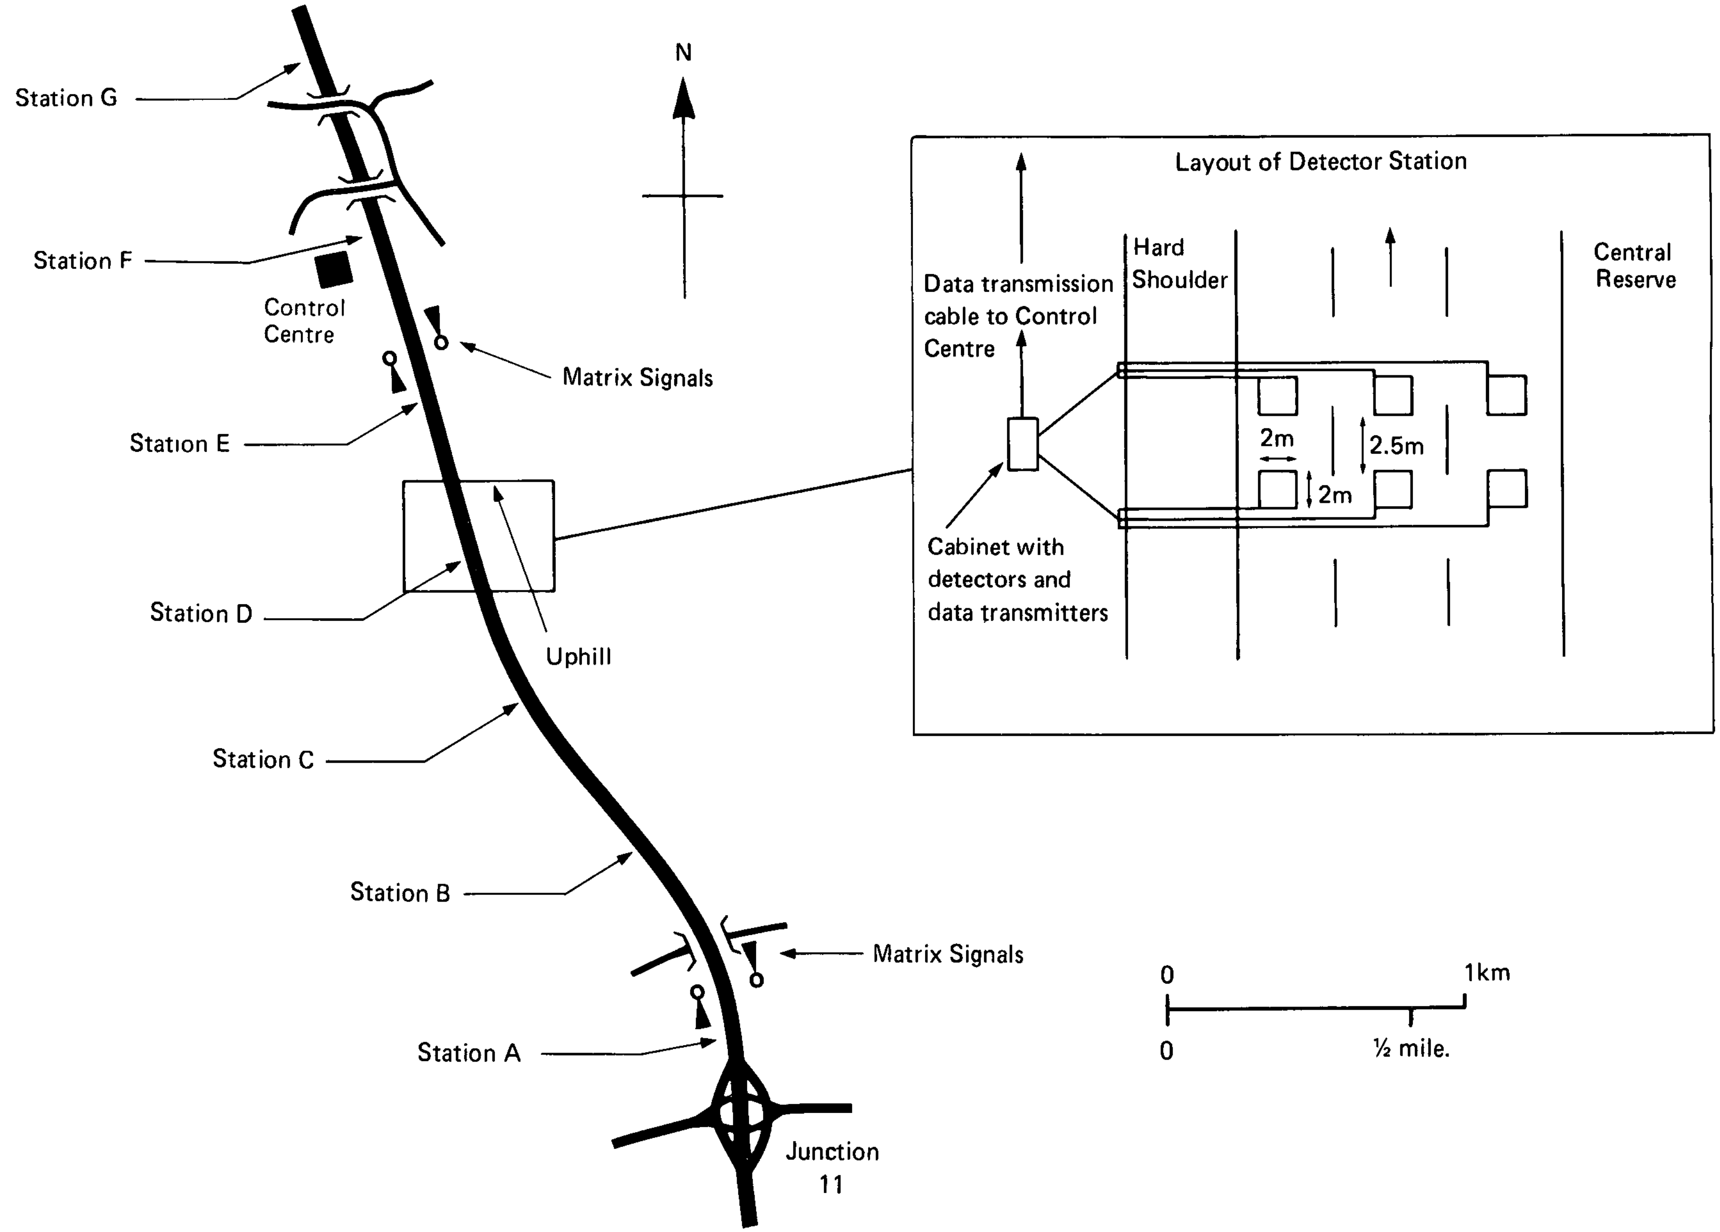

Fig. 7 MI EXPERIMENTAL SITE FOR AUTOMATIC INCIDENT DETECTION

(2112) Dd0536361 1,500 10/79 HPLtd So'ton G1915
PRINTED IN ENGLAND

ABSTRACT

AUTOMATIC INCIDENT DETECTION – TRRL ALGORITHMS HIOCC AND PATREG: *J F Collins, C M Hopkins and J A Martin:* Department of the Environment Department of Transport, TRRL Supplementary Report 526: Crowthorne, 1979 (Transport and Road Research Laboratory). An Automatic Incident Detection (AID) system, using computer based algorithms to identify the traffic disturbances following an incident, is being developed at TRRL. It incorporates two new algorithms, HIOCC and PATREG.

The HIOCC algorithm operates by detecting stationary or slow moving vehicles over individual vehicle detectors. It looks for several consecutive seconds of high detector occupancy to detect queues and incidents in high traffic flows. Off-line tests with recorded data show HIOCC's detection ability to be very good and a low false alarm rate can be achieved.

The PATREG algorithm measures the average traffic speed between upstream and downstream detector stations using a pattern recognition technique, detecting incidents by identifying significant speed changes. Preliminary tests suggest that it performs satisfactorily in free-flow conditions up to about 1500 veh/h per lane, and has an acceptably low false alarm rate.

By incorporating both HIOCC and PATREG in an AID system a wide range of traffic conditions can be catered for. A new AID site on the M1 in Bedfordshire will be used to complete the development and testing of the algorithms.

ISSN 0305–1315

ABSTRACT

AUTOMATIC INCIDENT DETECTION – TRRL ALGORITHMS HIOCC AND PATREG: *J F Collins, C M Hopkins and J A Martin:* Department of the Environment Department of Transport, TRRL Supplementary Report 526: Crowthorne, 1979 (Transport and Road Research Laboratory). An Automatic Incident Detection (AID) system, using computer based algorithms to identify the traffic disturbances following an incident, is being developed at TRRL. It incorporates two new algorithms, HIOCC and PATREG.

The HIOCC algorithm operates by detecting stationary or slow moving vehicles over individual vehicle detectors. It looks for several consecutive seconds of high detector occupancy to detect queues and incidents in high traffic flows. Off-line tests with recorded data show HIOCC's detection ability to be very good and a low false alarm rate can be achieved.

The PATREG algorithm measures the average traffic speed between upstream and downstream detector stations using a pattern recognition technique, detecting incidents by identifying significant speed changes. Preliminary tests suggest that it performs satisfactorily in free-flow conditions up to about 1500 veh/h per lane, and has an acceptably low false alarm rate.

By incorporating both HIOCC and PATREG in an AID system a wide range of traffic conditions can be catered for. A new AID site on the M1 in Bedfordshire will be used to complete the development and testing of the algorithms.

ISSN 0305–1315