Comparative Performance Evaluation of Incident Detection Algorithms

YORGOS J. STEPHANEDES, ATHANASIOS P. CHASSIAKOS, AND PANOS G. MICHALOPOULOS

A critical review of the most widely accepted conventional incident detection algorithms is presented. An improved algorithm, employing short-term time averaging to filter the traffic data, is proposed and compared with the best of the existing algorithms. The filtering addresses the problem from false alarms that are due to short-term traffic inhomogeneities. On the basis of data collected from a typical freeway in the Twin Cities metropolitan area, evaluation results suggest that the new algorithm achieves a 30 to 80 percent false alarm reduction compared with existing algorithms.

Fast and reliable freeway incident detection is instrumental in reducing traffic delay and increasing safety. In particular, with the information from incident detection, optimal control strategies guide the traffic flow toward smooth operation by preventing additional vehicles from entering the freeway upstream of the incident and by communicating relevant information to travelers. In addition, incident detection constitutes the cornerstone for prompt incident management and safety improvement near the incident location.

Existing techniques for the detection of freeway incidents do not provide the necessary reliability for freeway operations. Because they generate a high level of false alarms, conventional automated techniques based on computerized algorithms are less effective than is desirable for operational use. Operator-assisted methods minimize the false alarm risk, but they suffer from missed or delayed detections, they are labor-intensive, and they restrict the potential benefits from advanced, integrated traffic management schemes.

Responding to the need for effective and reliable detection of freeway incidents, an essential element for improved traffic management and control in freeway corridors (*1*), the authors initiated this research to investigate the performance limitations of conventional automatic incident detection systems and define the specifications for a new algorithmic logic that can lead to improved detection performance. The research initially focused on assessing the ultimate detection performance that can be accomplished with existing and new incident detection systems that use traffic data from presence detectors. A new algorithm was developed and tested against the major existing ones with promising results towards the development of a more-sophisticated detection structure. All tests employed a unified system of performance assessment (*2*), suitable for direct algorithm evaluation.

Department of Civil and Mineral Engineering, University of Minnesota, Minneapolis, Minn. 55455.

BACKGROUND

Automatic incident detection (AID) involves two major elements: a traffic detection system that provides the traffic information necessary for detection and an incident detection algorithm that interprets the information and ascertains the presence or absence of a capacity-reducing incident. Presence detectors imbedded in the freeway pavement are used extensively to obtain traffic data, primarily on occupancy and volume. Video detectors can also be used for the same purpose.

A number of AID algorithms can be found in the literature. Their structure varies in the degree of sophistication, complexity, and data requirements. The most important include the comparative algorithms (California logic) (*3–5*), the type employing statistical forecasting of traffic behavior (time-series algorithms) (*6–8*), and the McMaster algorithm (*9*). These algorithms operate on typical detector outputs of 30- to 60-sec occupancy and volume data. Equally important is the HIOCC algorithm (*10*), which uses 1-sec occupancy data to detect stationary or slow-moving cars; however, such data are not always available with existing surveillance systems. Additional detection methods (*11–13*) involve macroscopic traffic flow modeling to describe more fully the evolution of traffic variables. Table 1 presents the major AID algorithms, including the proposed algorithm, on the basis of their data requirements. For instance, the table indicates that the McMaster algorithm employs volume occupancy, and (optionally) speed data, averaged over 30-sec periods from one or two adjacent stations.

Algorithms of the comparative type, developed by Payne et al. (*3*) and later by Levin et al. (*4*), rely on the principle that an incident is likely to increase significantly occupancy upstream while reducing the occupancy downstream. To detect incidents, actual values of occupancy are compared against preselected thresholds. A typical algorithm includes tests to distinguish between incident and bottleneck congestion, compression wave tests to isolate compression shock-waves, and a persistence test to ensure that exceeding a threshold is not due to random fluctuations in the data. In an attempt to improve the significance of incident alarms, Levin and Krause (*5*) observed historical probability distributions of traffic variables under both incident and incident-free conditions; then they proposed the use of Bayes' rule to derive the optimal thresholds.

The California algorithm (the most widely known comparative algorithm) consists of three simple comparisons to preset thresholds (*3*). An incident is detected when upstream oc-

TABLE 1 Summary of Automatic Incident Detection Algorithms

Algorithm	Traffic Variables			Time (sec) Discretization			Number of Stations	
	Vol.	Occ.	Spd.	1	5–15	30–60	Single	Adjacent
PROPOSED ALG		x				x		x
COMPARATIVE		x				x		x
TIME SERIES	x	x				x	x	x
McMASTER	x	x	x[1]			x	x	x
HIOCC		x		x[2]		x	x	
WILLSKY	x	x			x[2]			x[3]
CREMER	x		x		x[2]			x

[1]Optional
[2]Not typically obtained from existing loop detector systems
[3]Requires closely spaced stations along the freeway

cupancy is significantly higher than downstream occupancy both in absolute value (Test 1) relative to upstream occupancy (Test 2) and at the same time, downstream occupancy has adequately decreased during the past 2 min (Test 3). Test 3 distinguishes an incident from a bottleneck situation by indicating that a reduction in downstream occupancy has occurred over a short period of time as a result of the incident. An alarm persistence and an incident termination test may be added to improve the algorithm performance. The persistence test requires the relative spatial occupancy difference (Test 2) to remain larger than the corresponding threshold for two consecutive time periods before signaling an incident. Moreover, a termination test seeks to signal the end of the incident. In particular, after an incident has been detected, the relative spatial occupancy difference is tested during the succeeding time periods; a drop in the value of this parameter below its threshold indicates that the incident effect has terminated.

Within the family of comparative algorithms, Algorithm #7 (*3*) has attracted the attention of previous reviewers. It is similar to the California algorithm, but it replaces the temporal downstream occupancy difference in the third test of the California model with the present downstream occupancy measurement. This replacement seeks to reduce the false alarms produced by compression waves, a common feature of the Los Angeles freeway traffic, where the algorithm was originally tested.

A second general class of algorithms considers the recent history of a traffic variable and employs a time-series model to provide short-term traffic forecasts. Significant deviation between observed and forecast values are attributed to incidents. The first of three algorithms in this class, the standard normal deviation algorithm (*6*), calculates the mean and standard deviation of occupancy for the last 3 to 5 min and detects an incident when the present value differs significantly from the mean in units of standard deviation.

The double exponential algorithm (*7*) performs a double exponential smoothing of traffic occupancy to forecast occupancy and identifies as incidents the calibrated deviations (the algebraic sum of all previous-estimate errors divided by the current estimate of the mean absolute deviation). The same idea characterizes the autoregressive integrated moving average (ARIMA) algorithm, in which an ARIMA model provides short-term forecasts of the state variable (traffic occupancy) and the associated 95 percent confidence limits (*8*). An incident is detected when the observed occupancy values lies outside the confidence limits.

Unlike the other algorithms that use mainly occupancy data, the McMaster algorithm (*9*) is based on a two-dimensional analysis of the traffic data. In particular, it proposes separating the flow-occupancy diagram into four areas corresponding to different states of traffic conditions. Incidents are detected after observing specific changes of the traffic state in a short time period. This approach requires calibration of the boundaries separating different traffic conditions—algorithm thresholds—individually for each station, as volume-occupancy characteristics vary across stations. Simplicity of design and a potential for improved detection performance are the major advantages of the algorithm.

Besides the previous approaches, which use aggregate traffic data averaged over 30 to 60 sec, Collins et al. (*10*) developed the HIOCC algorithm for the British U.K. Transport and Road Research Laboratory (TRRL) on the basis of 1 sec instantaneous occupancy data. The algorithm seeks several consecutive seconds of high detector occupancy in order to identify the presence of stationary or slow-moving vehicles over individual detectors. A weakness of this method is the lack of an effort to distinguish incidents from other congestion-producing traffic phenomena (e.g., compression waves).

The above algorithms use only detector output to make a decision, but other methods take advantage of insights gained from research in traffic flow modeling. Willsky et al. (*11*) proposed using macroscopic traffic modeling to describe the evolution of spatial-average traffic variables (velocities, flows, and densities), thus capturing the dynamic aspect of the traffic phenomena to alleviate the false alarm problem. Although scientifically appealing, the two methods resulting from that research did not attract the interest of practitioners. The lack of interest was owed to the complexity of the methods and the strong data requirements, in terms of type of variable (density, space mean speed) and short time-space discretization. These restrictions limited testing of the methods to a small number of simulated incident patterns.

Cremer has proposed a similar approach applicable to congested cross-country freeways in Europe, where detectors are located several kilometers apart (*12*). Whereas Willsky models an incident as having a capacity reduction effect, Cremer proposes that detection can be improved by modeling the attenuation of the road capacity with an additional (fictitious) volume input at the location of the incident.

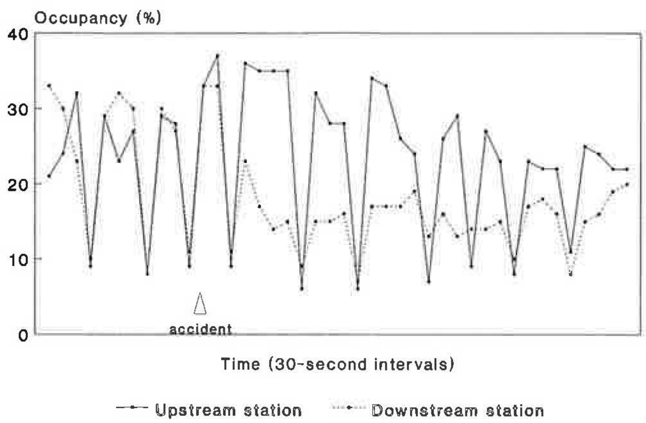

FIGURE 1 Noisy detector occupancy data.

Kühne (*13*) proposed the use of high-order continuum models to calculate the standard deviation of the speed distribution and found that the deviation broadens when density approaches the critical value in which stop-start traffic movement is observed. He concludes that detecting such broadening could be the basis for an early warning strategy and incident detection. However, his work has not produced a practical incident detection algorithm.

Despite several appealing features, existing algorithms are characterized by certain limitations. These limitations result mainly from two sources: (a) the use of raw data with only limited filtering and (b) the lack of effort, or effectiveness of effort, in distinguishing incidents from incident-like traffic situations.

More specifically, our experience suggests that raw detector data are often inappropriate for incident detection if care is not taken to filter out the noise before use—a weakness characterizing comparative algorithms. In particular, when corrupted by noise (see Figure 1) incident patterns in the traffic data may not be detected easily by an algorithm. Similarly, fluctuations produced by noise sources can be detected as incidents.

In addition, comparative algorithms are not very effective in detecting incident patterns. In an attempt to increase the success of differentiating incidents from recurrent congestion, these algorithms use tests that are somewhat restrictive. As a result, the only traffic patterns easily identified are those occurring under severe incident conditions and satisfying every test of an algorithm.

Statistical forecasting algorithms employ filtering as dictated by their design specifications. However, either they use simplified models or their models depend on excessive parameter estimation at each new site; this limits algorithm transferability. The major weakness of these algorithms lies in their lack of effort in distinguishing incidents from similar traffic patterns. In particular, the algorithms focus on detecting abrupt changes in the traffic stream without involving more-specific tests that classify the source of the abrupt change.

ALGORITHM DESCRIPTION

The proposed logic seeks to reduce the number of false alarms that can result from short-term traffic inhomogeneities. More specifically, it uses a moving average to filter the raw data before algorithm implementation and performs detection on the basis of spatial occupancy difference between adjacent stations. Attention is given to keeping the new algorithm structure as simple as possible. In particular, the algorithm does not use highly specific tests that would seek to account for given traffic abnormalities. From the authors' review of existing algorithms, highly specific tests are strongly data-related and can lead to decreased algorithm performance when transferred across data sets.

According to the new algorithm, an incident is likely to create congestion in the upstream and reduce flow in the downstream detector station; this leads to a high spatial occupancy difference between the two stations. Unlike existing algorithms, which consider such difference at single time intervals, the new algorithm uses a 3-min average (y_t, corresponding to six 30-sec measurements) of the spatial occupancy difference, x_t, between adjacent stations for the period following the incident occurrence time t:

$$y_t = 1/6 \sum_{k=0}^{5} x_{t+k} = 1/6 \sum_{k=0}^{5} (\mathrm{OCC}^u_{t+k} - \mathrm{OCC}^d_{t+k}) \quad (1)$$

where OCC^u_t is the upstream station occupancy at time t, and OCC^d_t is the downstream station occupancy at time t.

The averaging technique ensures that high values of y_t are not random but reflect congestion in the freeway segment.

If used alone, the preceding variable may lead to detection errors, since it does not consider the past flow condition. Incorporating the recent past information is crucial, especially in locations with special geometrics. For instance, in a lane-addition location, upstream occupancy is typically higher than the downstream occupancy (even under normal traffic conditions). To account for such permanent traffic abnormalities, y_t is compared with the average spatial occupancy difference, z_t, observed in the 5 min before the incident:

$$z_t = 1/10 \sum_{j=1}^{10} x_{t-j} = 1/10 \sum_{j=1}^{10} (\mathrm{OCC}^u_{t-j} - \mathrm{OCC}^d_{t-j}) \quad (2)$$

Significant difference between y_t and z_t indicates a temporal change in the state of traffic, signaling an incident.

Finally, because the values of y_t and z_t depend not only on the incident severity but also on the local flow level at the time of the incident, the foregoing variables are normalized to reflect relative changes with respect to existing traffic conditions. The normalization is done by m_t, the maximum of upstream and downstream station occupancy averaged over the most recent 5-min period before the incident:

$$m_t = 1/10 \max\left(\sum_{j=1}^{10} \mathrm{OCC}^u_{t-j}; \sum_{j=1}^{10} \mathrm{OCC}^d_{t-j} \right) \quad (3)$$

In summary, the algorithm logic includes two tests: first, when

$$\frac{\sum_{k=0}^{5} x_{t+k}}{m_t}$$

exceeds a given threshold, congestion is detected. If the congestion test is satisfied, a second test examines whether

$$\frac{\sum_{k=0}^{5} x_{t+k} - \sum_{j=1}^{10} x_{t-j}}{m_t}$$

exceeds a second threshold. If true, the test indicates an incident rather than recurrent congestion.

The selection of the pre- and post-incident period lengths for data averaging at 5 and 3 min, respectively, reflects our findings from preliminary analysis of the data. In particular, 5-min averaging removes high-frequency fluctuations of traffic flow without modifying the overall preincident traffic pattern. However, a shorter postincident averaging period of 3 min is needed to prevent high delays in the algorithm response time.

The proposed logic, while retaining simplicity, offers certain advantages over conventional logic. First, because of data filtering, the new algorithm substantially reduces the risk of false alarms from random fluctuations in traffic flow. The filtering also provides a more efficient persistence test than the tests employed by the comparative algorithms, since it accounts for six time periods instead of two. In addition, unlike traditional algorithms of the comparative type, this algorithm does not require the formation of a highly specific incident pattern to ensure detection. Nevertheless, a more sophisticated algorithm could benefit from considering the characteristics of each individual station. Such an algorithm could also benefit from a more effective effort in distinguishing incidents from other traffic patterns (e.g., compression waves) that resemble incidents, an issue currently being addressed by the authors.

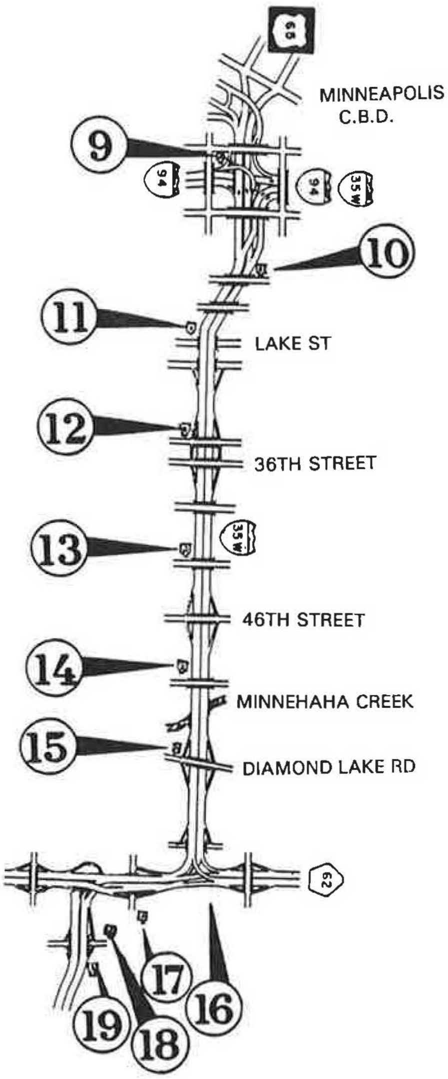

FIGURE 2 Study site in Minneapolis I-35W (Source: Traffic Management Center, Minnesota Department of Transportation).

DATA DESCRIPTION

All data for algorithm testing were collected from Interstate 35W, a heavily traveled and often congested freeway in Minneapolis, Minnesota (Figure 2). The selected freeway segment is fully instrumented with television cameras (camera locations are shown in Figure 2 with a circle) that allow detailed incident information to be gathered. The study was confined to the afternoon peak period (4:00 p.m. to 6:00 p.m.), since incident detection under moderate to heavy traffic conditions is of greatest importance for advanced freeway management. Traffic data, routinely collected by the traffic management center, consist of 1-min volume and occupancy counts that are updated every 30 sec and averaged over all lanes.

Information on actual incidents is available from the incident logs filed by the traffic management center operator daily. These logs report the time and location of each incident, its type, duration, severity, and impact on traffic, as well as the roadway condition and other incident-related information. Detection of incidents is mainly accomplished through observation of television monitors by on-duty traffic personnel, map displays, scanners, and state patrol reports.

Algorithm development and testing used 140 hr of traffic data from 72 weekdays. The data were obtained from 14 single detector stations along a 5.5-mi segment of I-35W. This segment includes most types of geometric configuration usually found in a freeway: entrance and exit ramps (often carrying heavy volumes) lane drop and addition, and freeway intersection. The complete data set consists of about 240,000 data points.

Detailed incident data included 27 capacity-reducing incidents, primarily accidents and stalls in the moving lanes and shoulders, reported during the study period by the traffic management center on the I-35W segment. The data impact on traffic operations varied from limited to very severe. Two additional incidents (stalls on the shoulder) classified by traffic operators as causing no impact on traffic were excluded from the analysis, because detection is based solely on observable changes in the traffic flow stream.

ALGORITHM EVALUATION

Evaluation of incident-detection algorithm performance generally uses three major indicators: detection rate, false alarm rate, and mean time to detect.

- Detection rate is the ratio of incidents detected out of all incidents that occur during a specified time period.
- False alarm rate is the ratio of false alarms out of all decisions (incident and nonincident) made by the system during a specific time period. Certain authors also employ an online definition of false alarm rate as the percent of false decisions out of all incident decisions during a certain period of time (*5*).
- Mean time to detect is the average amount of time required by the system to make a detection.

The preceding measures of effectiveness are related because, at least in the single-test algorithms, increasing the detection rate causes the false alarm rate to increase. No standards have yet been adopted for determining the best combination of detection and false alarm rates. The lack of standards lies primarily in the difficulty of reconciling the difference between the consequences of a missed incident and those of a false alarm, so that the average cost of incident detection is minimized.

The need for evaluating existing algorithms is evident in the literature, which does not address conclusively the potential for transferability of evaluation results. Our findings indicate that evaluation results are not always transferable within an acceptable error range, probably because of varying traffic conditions, weather, and driver characteristics across application sites. Differences across sets of incident data are an additional reason for the lack of transferability. These differences result primarily from varying assumptions on whether incidents with very limited impact on the traffic should be included. Therefore, unbiased evaluation of a new algorithm requires concurrent evaluation of major existing ones on a common data set.

We can increase the potential for robustness and transferability of comparative-evaluation results by employing a unified system for the evaluation through an operating characteristic curve (*2,3,7,8,14*). Such a curve is a plot of detection probability (or rate) P_D versus false alarm probability (or rate) P_F that an algorithm can achieve. To construct this curve, the threshold parameters are allowed to vary over a wide range of values. Every threshold (or threshold set) produces a performance point (P_D, P_F) on the curve. The main advantage of this technique for comparing detection algorithms lies in its independence from algorithm structure. Consequently, algorithms can be compared to each other without regard to the number of tests, traffic variable involved, or type of algorithm. Additionally, the operating characteristic curve covers the whole range of detection rate (0 to 100 percent), allowing the algorithm user to establish the thresholds that best meet the requirements for a traffic operation.

Existing Algorithms

In this work, we performed detailed testing of two general algorithm types: comparative (California logic) and time-series. We did not include other approaches in the evaluation because they either require information currently not obtained routinely from loop detectors at most U.S. test sites (HIOCC algorithm) or do not provide explicit algorithms to be implemented with real data (Willsky, Kühne). We are in the process of testing the McMaster algorithm and plan to incorporate its most appealing features in more sophisticated detection algorithms under development.

The performance of two comparative algorithms—the California algorithm and Algorithm #7—is shown in Figure 3 in terms of their operating characteristics. A common feature of these algorithms is that each employs three traffic variables in its structure. Thus, increasing some thresholds while reducing others can maintain the same detection rate at different false alarm rates and vice versa. The lack of one-to-one correspondence between detection and false alarm rates in the case of multiple variables leads to the scattered performance pairs (P_D, P_F) as shown in Figure 3. In practice, high sensitivity to threshold selection reduces the algorithm attractiveness, since performance can easily deteriorate if the threshold set is poorly selected.

Three time-series algorithms, the standard deviation, the double exponential, and the ARIMA, are compared in Figure 4 employing station occupancy data. The standard deviation algorithm uses a 5-min time base to calculate averages and standard deviations and requires persistence of alarm for two consecutive time periods to signal an incident. The values for the smoothing constants in the double exponential and the ARIMA algorithms were taken as recommended (*7,8*). Updated values of these parameters could lead to improved algorithm performance. The evaluation reveals that at low detection rates the standard deviation exhibits the best performance, whereas at high detection rates the best performance is exhibited by the double exponential algorithm.

Test results also indicate that spatial occupancy difference (the key variable in comparative algorithms) is more sensitive to incidents than single station occupancy and leads to improved algorithm performance. This sensitivity differential, exhibited by all tested algorithms, is shown in Figure 5 for the double exponential algorithm.

The performance evaluation results reveal a superiority of the comparative algorithms to time-series. Figure 6 shows the operating characteristic curves, summarizing the performance of algorithms from both classes (California, Algorithm #7, standard deviation, and double exponential), as well as the new algorithm; its detection performance is discussed in the next section. The curves of the time-series algorithms are based on spatial occupancy difference data so that only their improved performance is considered. For both comparative

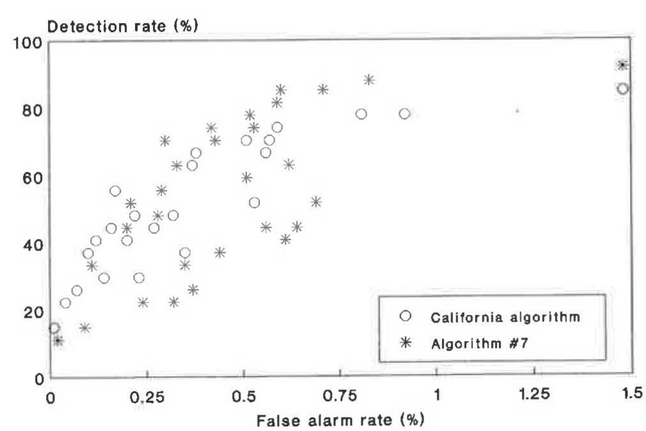

FIGURE 3 Operating characteristics of comparative algorithms.

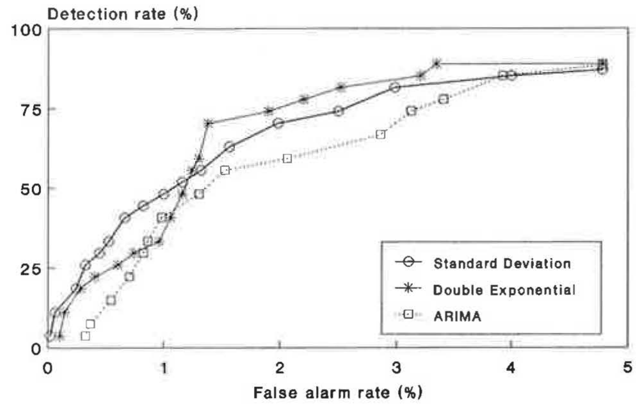

FIGURE 4 Operating characteristic curves: time-series algorithms.

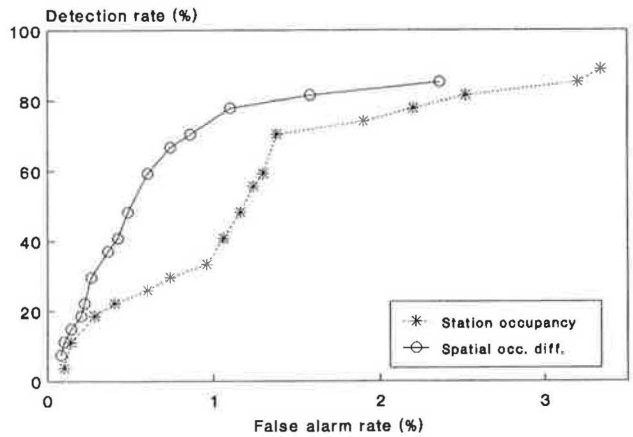

FIGURE 5 Performance comparison between two detection parameters.

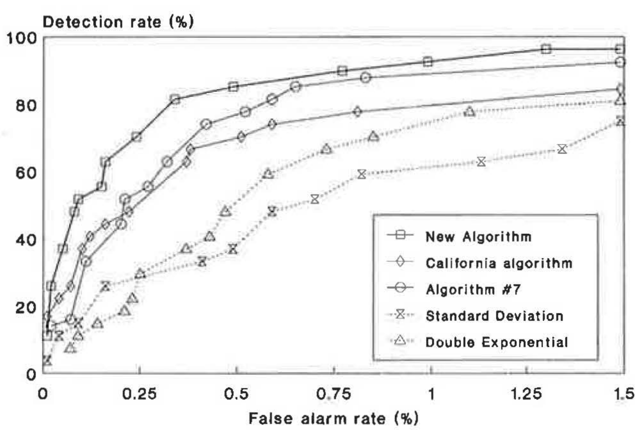

FIGURE 6 Algorithm performance comparison.

algorithms, when a specific detection rate can lead to several false alarm rates (see scattered plots in Figure 3), only the superior operating point is included in Figure 6; the assumption is that, in practice, the thresholds that lead to superior performance can be obtained through testing. As the figure illustrates, the false alarm rates of the comparative algorithms are 30 to 60 percent lower than those of the time-series algorithms.

All conventional algorithms exhibit performance that at first glance may appear to be acceptable for implementation. However, a better feeling of this performance can be gained by considering the absolute number of false alarms instead of the false alarm rate. For example, among conventional algorithms, at 50 percent detection rate, the lowest false alarm rate equals 0.21 percent—achieved by Algorithm #7. Such a false alarm rate in a sample of 140 hr of 30-sec traffic data from 14 detector stations corresponds to approximately 3.5 false alarms per hour for the relatively small 5.5-mi freeway segment of our application.

Figure 7 shows the mean time-to-detect performance of the above algorithms. The standard deviation presents the lowest response time in signaling incidents, 1 to 1.5 min earlier than comparative algorithms. The double exponential tends to respond slowly to traffic changes; this causes excessive delays in detecting incidents. The negative time-to-detect values in the figure result from the fact that time of incident is the instant when the operator identifies that incident, usually a few minutes after its occurrence. For instance, a negative value of mean time to detect implies that the algorithm detects the incident before the operator does.

For examining the transferability of evaluation results across application sites, we compared our findings to those from previous performance studies. In particular, we considered earlier findings on the California algorithm and Algorithm #7. The analysis revealed that the operating characteristic curves of the California algorithm produced by the earlier (*3*) and present studies almost coincide, indicating a potential for transferability. However, the performance of Algorithm #7 in the original study was different from that in our tests. Algorithm #7 was designed specifically with emphasis on eliminating false alarms resulting from compression waves, as, for instance, in the Los Angeles freeway system. Performance may be expected to deteriorate in the Minneapolis–St. Paul freeways on which compression waves are fewer and less severe.

New Algorithm

The new algorithm is compared with the best-performing existing ones from the previous section via their operating char-

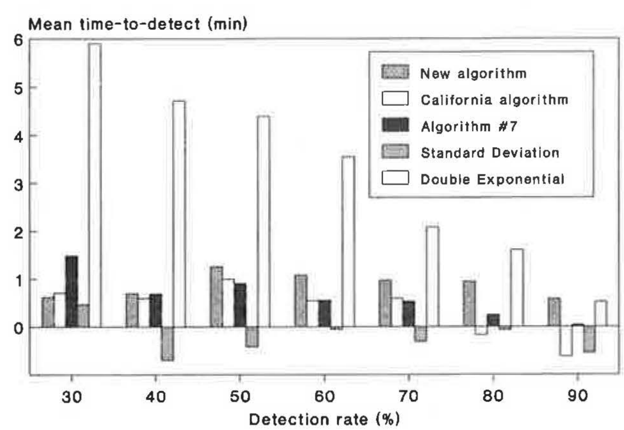

FIGURE 7 Mean time-to-detect performance.

acteristic curves (see Figure 6). At all levels of detection rates, the new algorithm is superior to all existing algorithms. As shown in the figure, the new algorithm yields detection rate of 50 percent at a 0.1 percent false alarm rate, whereas best-performing Algorithm #7 produces more than twice the number of false alarms at the same detection rate. Similarly, the new algorithm achieves a 0.3 percent false alarm rate at 80 percent detection rate, whereas Algorithm #7 results in a 0.6 percent false alarm rate at the same detection rate. As the figure suggests, the new algorithm produces 50 to 70 percent fewer false alarms than the California algorithm, 30 to 60 percent fewer than Algorithm #7, and 70 to 80 percent fewer than the standard deviation and double exponential. In practical terms, although better than that of the existing algorithms the performance of the new algorithm may not be operationally acceptable. For instance, at a 50 percent detection rate the new algorithm produces about 1.5 false alarms per hour within the 5.5-mi segment. The false alarm rate could decrease further through improved traffic flow modeling that could distinguish more effectively incidents from incident-like traffic patterns, an issue being addressed by the authors.

The new algorithm also exhibits acceptable mean time to detect. As Figure 7 indicates, although its structure imposes a 3-min detection delay, its mean time to detect remains close (by less than 1 min) to the detection times required of the comparative algorithms. The competitive response time of the new algorithm can be expected, as the algorithm begins sensing an incident from the onset of congestion (with a 3-min delay); on the other hand, comparative algorithms initiate an alarm only after the incident-generated congestion reaches a level that is sufficiently high for the test variables to exceed the thresholds.

CONCLUSIONS

Conventional automatic incident detection algorithms and the new algorithm were evaluated on the basis of their operating characteristic curves. The evaluation revealed that comparative algorithms, employing three test variables, can distinguish incidents from other traffic phenomena more effectively than single-variable time-series algorithms that use statistical forecasting of traffic. At all detection levels, the comparative algorithms produce 30 to 50 percent fewer false alarms than time-series algorithms.

Test results with the new algorithm indicate a decrease of 30 to 70 percent in false alarm rates compared to comparative algorithms and a 70 to 80 percent reduction compared to time-series algorithms. In addition, the new algorithm presents a mean time to detect comparable to that of existing algorithms.

To avoid false alarms that are primarily due to short-term traffic inhomogeneities, the new algorithm uses filtered detector output (i.e., values averaged over short time periods). It is simple to implement and less sensitive to random fluctuations of traffic than existing algorithms, and it requires only the minimum amount of data routinely provided by current presence detector systems.

Comparison of our results with findings from the literature indicates that algorithm performance may exhibit varying degrees of transferability across test locations. Whereas transferability potential increases for algorithms designed with general traffic behavior in mind, it deteriorates when algorithms involve specific tests to account for traffic phenomena (e.g., compression waves) that do not exhibit the same frequency of occurrence and severity across test sites.

Although the new algorithm improves on the performance of conventional detection methods, it is still restricted by modeling and hardware limitations. In order to achieve a lower false alarm rate appropriate for operational use, further research is being pursued by the authors, addressing the need for improved traffic modeling, more effective distinction between incident and nonincident alarms, and use of machine vision techniques for the derivation of traffic data.

ACKNOWLEDGMENTS

This study was supported by the National Science Foundation. The Minnesota Supercomputer Institute provided partial support. The Center for Transportation Studies, Department of Civil and Mineral Engineering, University of Minnesota is acknowledged for its support. The Traffic Management Center, Minnesota Department of Transportation, cooperated in the study by providing the necessary data.

REFERENCES

1. Y. J. Stephanedes and K. Chang. Optimal Ramp-Metering Control for Freeway Corridors. *Proc., 2nd International Conference on the Applications of Advanced Technologies in Transportation Engineering*, Minneapolis, Minn., Aug. 1991, pp. 172–176.
2. Y. J. Stephanedes and A. P. Chassiakos. A Low Pass Filter for Incident Detection. *Proc., 2nd International Conference on the Applications of Advanced Technologies in Transportation Engineering*, Minneapolis, Minn., Aug. 1991, pp. 378–382.
3. H. J. Payne and S. C. Tignor. Freeway Incident Detection Algorithms Based on Decision Trees with States. In *Transportation Research Record 682*, TRB, National Research Council, Washington, D.C., 1978, pp. 30–37.
4. M. Levin and G. M. Krause. Incident Detection Algorithms, Part 1: Off-Line Evaluation; Part 2: On-Line Evaluation. In *Transportation Research Record 722*, TRB, National Research Council, Washington, D.C., 1979, pp. 49–64.
5. M. Levin and G. M. Krause. Incident Detection: A Bayesian Approach. In *Transportation Research Record 682*, TRB, National Research Council, Washington, D.C., 1978, pp. 52–58.
6. C. L. Dudek and C. J. Messer. Incident Detection on Urban Freeways. In *Transportation Research Record 495*, TRB, National Research Council, Washington, D.C., 1974, pp. 12–24.
7. A. R. Cook and D. E. Cleveland. Detection of Freeway Capacity-Reducing Incidents by Traffic-Stream Measurements. In *Transportation Research Record 495*, TRB, National Research Council, Washington, D.C., 1974, pp. 1–11.
8. S. A. Ahmed and A. R. Cook. Discrete Dynamic Models for Freeway Incident Detection Systems. *Transportation Planning and Technology*, Vol. 7, 1982, pp. 231–242.
9. B. Persaud, F. Hall, and L. Hall. Development and Testing of the McMaster Incident Detection Algorithm. Presented at the 69th Annual Meeting of Transportation Research Board, Washington, D.C., Jan. 1990.
10. J. F. Collins, C. M. Hopkins, and J. A. Martin. *Automatic Incident Detection—TRRL Algorithms HIOCC and PATREG*. TRRL Supplementary Report 526. U.K. Transport and Road Research Laboratory. Crowthorne, Berkshire, England, 1979.
11. A. S. Willsky, E. Y. Chow, S. B. Gershwin, C. S. Greene, P. Houpt, and A. L. Kurkjan. Dynamic Model-Based Techniques for the Detection of Incidents on Freeways. *IEEE Transactions on Automatic Control*, Vol. AC-25, No. 3, June 1980, pp. 347–360.

12. M. Cremer. Incident Detection on Freeways by Filtering Techniques. *Proc., 8th IFAC Congress*, Vol. 17, Kyoto, Japan, 1981, pp. 96–101.
13. R. D. Kühne. Freeway Control and Incident Detection Using a Stochastic Continuum Theory of Traffic Flow. *Proc., 1st International Conference on the Applications of Advanced Technologies in Transportation Engineering*, San Diego, Calif., Feb. 1989, pp. 287–292.
14. A. P. Chassiakos. *An Improved Time-Space Incident Detection Algorithm*. M.Sc. thesis. Department of Civil and Mineral Engineering, University of Minnesota, Minneapolis, 1990.

Publication of this paper sponsored by Committee on Freeway Operations.