Improved Data Screening Techniques for Freeway Traffic Management Systems

DON CLEGHORN, FRED L. HALL, AND DAVID GARBUIO

Proper functioning of computer-based freeway traffic control algorithms depends on good data being received from roadway detectors. At present, only elementary screening procedures for these data are in use on most freeway traffic management systems. There has recently been some work to improve such procedures. That work has been extended, and several new screening approaches are proposed. The first approach develops a theoretical upper bound for part of the flow-occupancy data and is useful for single-detector systems. The other approaches have been developed for paired-loop systems and involve comparisons of loop data as well as the development of boundaries defining acceptable combinations of speed, flow, and occupancy data.

The ability to monitor the conditions on a freeway is the fundamental feature of a freeway traffic management system (FTMS), on which all other functions depend. Most FTMSs measure flow and occupancy, and some also measure average speed. Using these data, the central computer of an FTMS can perform many complex functions, such as incident detection and travel time estimation. Unfortunately, the most sophisticated and elaborate computer algorithms inevitably fall short of their goals if the information they receive is of poor quality. Some new procedures for the identification and removal of bad FTMS data in an on-line computer environment are described. The resulting screened FTMS data is a more accurate representation of freeway conditions and thus should improve system performance.

The first section provides background, describing current data screening approaches. The second section describes the FTMSs from which data have been taken for analysis. Because the objective is to demonstrate the effectiveness of some proposed procedures, not to test hypotheses, the development of the procedures is explained as the results are presented, rather than discussing development separately from applications. Hence, the third section contains both the discussion of the procedures and a demonstration of the type of results that can be obtained from each of the proposed screening methods. The final section contains conclusions.

BACKGROUND

The effects of bad data on freeway monitoring and control algorithms have been documented in the past. Dudek et al. (*1*) noted that hardware problems resulted in an increase in the number of false alarms observed in the Gulf freeway warning system. Courage et al. (*2*) found that design and operational parameters can have significant effects on accuracy, even without hardware failures. For example, the detector scan rate can have a large effect, especially in a system in which the central computer monitors the detectors directly. Courage et al. determined that the minimum expected error in speed observations (for the smallest scanning interval examined) from such systems is 19 percent. Chen and May (*3*) also found that slight misadjustments in the detecting equipment (which are bound to occur from time to time) led to significant differences in output values.

Although the screening of data from vehicle detectors is not a new concept, little research has been done on the topic, and the testing methods that exist are still very simple (if used at all). In a 1984 survey of 32 freeway management projects in North America, only 19 had data checking procedures (*4*). Most of these 19 systems checked for stuck-on or stuck-off conditions and chattering in the detector amplifiers. Only 11 systems checked the plausibility of values received from the detector stations. Jacobson et al. (*5*) divided data screening tests into two categories: microscopic and macroscopic. Microscopic tests are defined as those tests done by the controller microprocessor itself as pulses are received from the detectors. Macroscopic tests are performed by the central computer after the data have been aggregated over time.

The macroscopic plausibility testing done most often (*4*) is comparison of a single variable against upper or lower threshold values. For example, the California Department of Transportation (Caltrans) algorithm checked averages of values over 5-min periods; the Chicago system checked data averaged over 1 min (*6*). The thresholds are usually constants input by system personnel on the basis of an examination of the reasonable range of the variables. For example, occupancy is only defined from 0 to 100 percent and both extremes can be observed in traffic, so its reasonable limits are clear. Some system operators choose to reduce this range, judging that the extremes are rare enough to be unimportant (*5*). Values of speed or flow cannot be less than 0 regardless of the units used. Their maxima are not easy to determine, although historical data may be used to produce estimates.

These threshold tests implicitly assume that the acceptable range for a variable is independent of the values of the other variables. Because combinations of variables are not tested, single-variable tests using such thresholds cannot catch completely impossible combinations, such as an occupancy of 3 percent together with a flow of 2,000 vehicles per hour (vph). This observation is clearly not reasonable, but the only published test for such combinations is that reported recently in Washington State (*5*). The Washington State approach in-

D. Cleghorn, UMA Engineering, 5080 Commerce Boulevard, Mississauga, Ontario, Canada L4W 4P2. F. L. Hall, Department of Civil Engineering, and D. Garbuio, Department of Geography, McMaster University, Hamilton, Ontario, Canada L8S 4L7.

volves setting limits for acceptable values of volumes for any given occupancy on the basis of plausible ratios between the two within specific occupancy ranges (see Figure 1). Although ingenious, their method provides wider ranges than are necessary, due to the stepped nature of the boundaries. Testing on the basis of combinations of variables is clearly on the right track, but a tighter limit should be possible.

On the whole, then, current FTMSs use data from detector stations with minimal examination of credibility. Because the purpose of an FTMS is to monitor a freeway system and thus allow its control, the quality of the data collected is of critical importance to the proper operation of the FTMS. The designs of several new data screening algorithms, which incorporate tests to detect previously overlooked flaws in FTMS data, are presented in the following sections.

SOURCES OF DATA

Data from two FTMSs are used in the analysis, both on the Queen Elizabeth Way (QEW) in Ontario, Canada. The first system is on the Burlington Skyway and its approaches. The Skyway is the portion of the QEW that crosses the entrance to Hamilton Harbour; therefore, it is high enough for ocean-going and Great Lakes shipping to pass under safely. This height is accomplished by a 3 percent grade for about 1.25 km. The second FTMS used is in Mississauga on the main western commuter approach to Toronto, a section that experiences recurrent congestion nearly every weekday.

These systems have been selected because they are representative of many others, not because they are unique. Indeed, the Skyway system has only recently (1986) been put into operation and so contains state-of-the-art components, both hardware and software, as of that time. The Mississauga system was initially installed in 1978 (*7*). It has had some equipment replacements during the past 2 years but still has some older equipment. Hence, the stations in this system offer a range of data error experience. The data collection systems and screening procedures for these systems are described in the following paragraphs.

Most of the data collection stations have two inductance loop detectors embedded in the roadway, 6 m apart, for each lane at a station. The detector loops are monitored by a Type 170 controller, which scans the loops at 60 Hz. The controller summarizes the information gathered from the loops over a 30-sec period for transmission to the central computer. If a value cannot be calculated for a loop during that 30-sec interval, the value is marked as missing by the controller. Missing variables can arise for a variety of reasons. For example, speed cannot be calculated when there has been no traffic during the interval. Before sending data to the central computer, the station controller may flag one or more detector loops as bad because of a stuck-on or stuck-off condition. In the Skyway FTMS, a loop is considered stuck if it is in the same state (either occupied or not occupied) continuously for 2 hr.

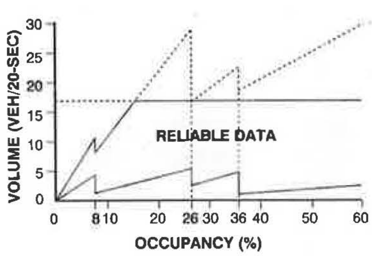

FIGURE 1 Boundaries on reliable data set in the Washington State study (*5*).

The data received by the central computer consist of 30-sec average values of speed, as well as values for volume and occupancy for each lane and station in the system. Basic data screening is done by the central software before the incident detection algorithms are run, and 5-min averages are tested to discover failed detectors (*8* and informal communication, Freeway Traffic Management Systems Section, Ontario Ministry of Transportation, April 1988). The simple range tests are performed every interval and can be summarized as follows:

1. Volume must not be negative and must not exceed a specified threshold, currently 60 vehicles per interval (roughly double the maximum observed to date).
2. Occupancy must not be negative and must not exceed 100 percent.
3. Speed must not be negative and must not exceed 150 km/hr. (Actual speeds in excess of 130 km/hr are occasionally observed, so this limit is only about a 10 percent margin).

If any one of these tests on 30-sec data fails, a flag is set to mark the invalid datum. Any incident detection algorithms that require data that have failed the test are disabled for that lane or station for the current 30-sec interval.

PROPOSED COMPARATIVE TECHNIQUES FOR SCREENING

Systems with paired loops in each lane at each station, such as those on the QEW, have more data-screening options available. Not only can they test the validity of flow-occupancy data pairs, but they can also compare the received speed-flow-occupancy points with a calibrated three-dimensional speed-flow-occupancy region and compare upstream and downstream loop values. The following discussion of improved screening techniques is divided into two main sections, the first covering an improved method for single-loop systems and the second identifying the additional screening that can be done with paired-loop systems.

Single-Loop Systems: Setting Boundaries on Flow-Occupancy Points

The Washington State approach—trying to identify the region of feasible flow-occupancy pairs—represents a greatly improved method for screening FTMS data from single-loop systems. It makes use of knowledge about the relationship between these variables that is ignored by single-variable range tests. However, those boundaries can be improved, given their stepped nature (see Figure 1). To improve those boundaries, emphasis has been placed on the upper bound because experience has shown that nearly all pairs of flow-occupancy values are possible below the normal flow-occupancy function, at the time congestion begins.

Two methods for obtaining a tighter upper bound suggest themselves. The first relies on historical data to set the limits. The second, and the one pursued here, is to try to develop a theoretical approach to an upper bound, thereby eliminating both the Washington State indirect method and the need for historical data. An additional benefit of such an approach would be that the boundary could be manipulated to suit a specific system using values of the variables acquired there.

The relationship between volume and occupancy depends on speed, vehicle length, and detector length. Following Athol (*9*), the relationship is

$$\text{Occupancy} = [\Sigma\ (x_i + d)/u_i]/T$$

where

x_i = length of Vehicle i,
u_i = speed of Vehicle i,
d = detector length, and
T = duration of the time interval.

The summation is taken over the vehicles passing in T.

This equation does not reduce to any simple form except under the assumptions of uniform vehicle length and speed. Fortunately, those assumptions are reasonable for the bulk of uncongested flow, as found by Hall and Persaud (*10*), so it may be possible to establish an upper limit for those conditions from the equation that results from those assumptions:

$$\text{Volume} = (\text{occupancy} * \text{speed} * T) \div (\text{vehicle length} + \text{detector length})$$

Although the calculations were done with the equation in this form, it is easiest to explain with occupancy calculated from volume. The upper bound of interest can be visualized as the minimum value of occupancy allowed for a given volume. If volume is constant, occupancy will decrease as speeds increase and as vehicle lengths decrease. Hence, the upper bound should be calculated using the highest feasible speeds and the shortest reasonable vehicle lengths. The maximum speed used in this analysis was 130 km/hr, but this limit can easily be changed for different systems. A minimum vehicle length of 4 m was used, which is the shortest that might be expected in the median lane. Motorcycles are shorter, but the average vehicle length during 30 sec is of interest. For the shoulder lane, a higher value might be reasonable. The detection zone length used in the calculations was 2 m.

The result of the calculations, up to an occupancy of 20 percent, is shown in Figure 2 along with the observed data for Station 17 on the Mississauga system. The slope of the line is the ratio of speeds to combined vehicle and detector lengths. If lower speeds or longer vehicles are used for the calculation, the slope decreases. Theoretically, volume is zero with an occupancy of zero. However, an observation of up to three volume counts at zero occupancy is not unusual because of truncation of the occupancy values in these systems. Thus, a vertical shift of the boundary upward by three volume counts, as shown in the figure, is needed to ensure that the boundary reflects the way this particular system operates on the data. Similar shifts no doubt are needed for other systems. Although the shifted upper bound would accept all of the data in Figure 2 (as it should), such a boundary would reject the type of data that the Washington State screening test found. Similar data were gathered during testing from one station on the Mississauga system. As shown in Figure 3, this boundary can reject such data.

To continue the upper bound beyond the critical occupancy (i.e., that at which highest volumes occur) poses some difficulty. Observed maximum speeds decline abruptly but follow no obvious pattern. In addition, where previously the highest speed was wanted, to find the lowest possible occupancy, now the task is to find the highest possible occupancy for a given volume, which requires the lowest possible speed. To choose a minimum speed for higher occupancies is difficult. However, the existence of a tight upper bound for uncongested flow may be sufficient. If the computer is programmed to disable a detector when it consistently fails a screening test (e.g., for 1 or 2 hr), then the detector fault would likely have been discovered before congested conditions began. Otherwise, the Washington State boundary may be used within the congested area.

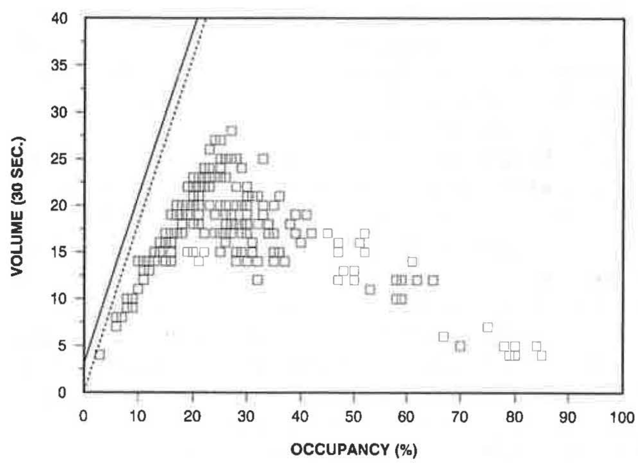

FIGURE 2 Upper bounds at Station 17, 6:00–11:00 a.m., May 15, 1990.

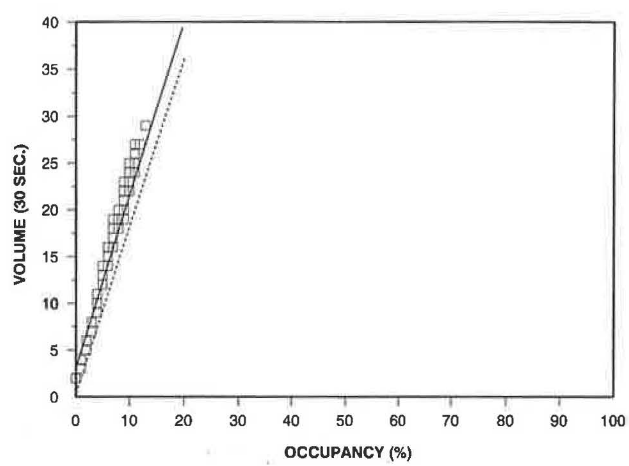

FIGURE 3 Upper bounds at Station 18, 6:00–9:00 a.m., May 15, 1990.

Screening Methods for Paired-Loop Systems

Three methods are discussed for paired-loop systems. The first is a comparison of the values obtained from the upstream loop of the pair with those from the downstream loop, for both volume and occupancy data. The second method identifies unusual combinations of speed, flow, and occupancy that may nonetheless contain valuable data. The final method involves setting feasible speed ranges for flow-occupancy pairs of data, on the basis of historical data.

Variation Between Loops at a Lane-Station

The two loops in a speed trap pair often have different volume and occupancy counts for the same 30-sec interval. Clearly, it is possible to have two volume counts that disagree but both of which pass the simple range test. The question, then, is how much of a difference to allow. A difference of two or less was selected. Because of the discreteness of the variables being measured on a 30-sec basis, it is expected that differences of one vehicle in the volume count will be encountered reasonably often. As well, there is the possibility of lane-changing maneuvers leading to a discrepancy of one or two vehicles, even between loops spaced only 6 m apart. Figure 4 shows that for properly functioning loops a difference of two covers nearly all of the data. There are other stations that regularly report larger differences, which probably indicates a loop pair in need of calibration.

A similar situation might be expected for occupancy. Inspection of the data, as shown in Figure 5, disproves this assumption. For uncongested flow (i.e., occupancies below about 25 percent in the figure), the difference in the two measurements of occupancy is small, but for congested flow the differences become much larger. The explanation is obvious, once the data have been seen. During uncongested travel, each vehicle is moving at roughly constant speed and occupies the two detectors in the pair for nearly the same amount of time. However, during congested flow, each vehicle is continually accelerating or decelerating and, thus, travels at somewhat different speeds as it crosses even these closely spaced detectors.

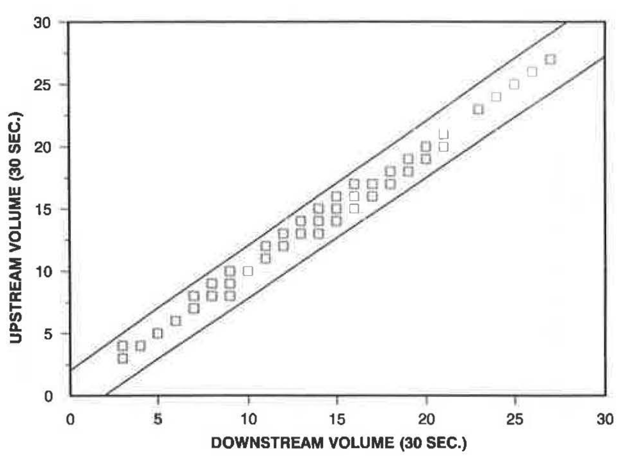

FIGURE 4 Upstream versus downstream volume at a double-loop station (Station 25).

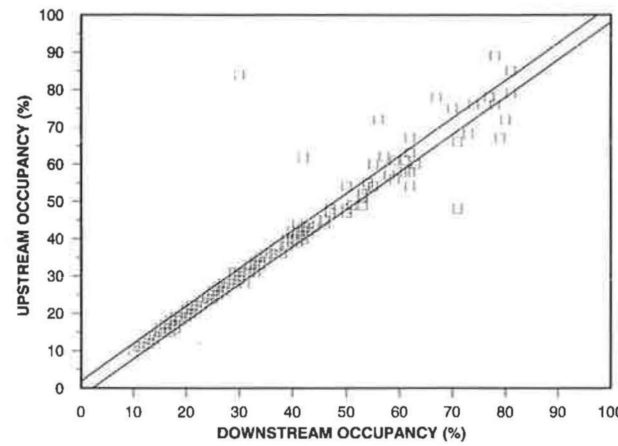

FIGURE 5 Upstream versus downstream occupancies at Station 14.

Unusual Values of Speed, Flow, and Occupancy Data

Not all apparently impossible combinations of the variables are actually bad data: the technicalities of data collection can result in some strange values. For example, observations with small positive occupancies (such as 1 or 2 percent) and zero volumes are occasionally found. These observations can be explained by the ways in which volume and occupancy are detected by a typical inductance loop detector and amplifier system.

Consider a signal generated by a vehicle passing over a detector loop. As long as the entire signal is contained within the same 30-sec interval, the output will be as expected (one volume count defined by the signal's leading edge, and a small positive occupancy count, the magnitude of which depends on the speed of the vehicle). If the 30-sec interval ends just after the leading edge of the signal, however, the subsequent interval has a small positive occupancy but no volume count (assuming no more vehicles pass over the detector in this interval).

Another seemingly nonsensical observation, that of zero occupancy and small positive volume, occurs several orders of magnitude more often than the zero volume observation. This observation can also be explained by the data collection process. Occupancy is reported as an integer value, so some information is lost in the truncation that produces the integer. This truncation means that a detector must be occupied for at least 0.30 sec in one interval to register an occupancy of 1 percent. A vehicle traveling at 100 km/hr must have length (as perceived by the detector) of at least 8.33 m to register 1 percent occupancy. For higher vehicle speeds, the perceived vehicle length over the 30-sec interval must be larger (e.g., 10.8 m at 130 km/hr). It is possible that two vehicles together will not exceed this perceived length threshold. The result is a positive volume accompanied by zero occupancy.

The final two types of data to be considered bring in the speed variable. In discussing them, it is important to recall that speed is measured across the two closely spaced loops in these systems and is not calculated from the flow and occupancy data. In the first case, speed is missing, but the other two variables are present. This observation is valid and should not call into question the other two variables. Speed is missing when the controller logic is unable to sort out vehicle arrivals at the two loops (e.g., lane changes or stop-and-go traffic), but this problem does not affect the other two variables. In the second case, speed is present and positive, but both volume and occupancy are zero. This condition caused numerous false alarms in an incident detection algorithm on the Skyway system before the error was discovered. Each variable on its own passed the threshold tests, so the additional test was needed to uncover the error.

Screening of Speed, Flow, and Occupancy Data on the Basis of Historical Data

Figure 6 shows 3 hr of FTMS data, representing both congested and uncongested operation. Several points in the speed-occupancy plot represent unrealistic conditions, such as 100 percent occupancy at 35 km/hr and 60 percent occupancy at 85 km/hr. There are, of course, many other possible combinations of the variables that are unreasonable. This type of error seems to occur at random intervals in the data set examined and has been observed in different lanes and stations. To catch the error, a screening mechanism that uses all three variables is needed.

The first step in designing such a screening method was to determine how a computer algorithm could describe the range of possible observations of the variables. In other words, the problem is to represent the region occupied in three dimensions by the set of all observable combinations of the traffic variables. A fitted mathematical model would be the ideal tool for representing the data in three-space because an equation is easy to evaluate and the decision on the validity of an observation would be trivial. Such a model would define a boundary about the region occupied by the data in three dimensions. Unfortunately, a simple mathematical model could not easily be applied to enclose such an irregular shape.

The next logical choice was a graphical model, that is, a method of defining the acceptable region of observations without using mathematical curve-fitting techniques. Because the shape in three dimensions was of no known polyhedron or other simple object, slicing the object into two-dimensional planes seemed to be a good way to simplify the problem. Volume was selected as the axis to divide the region because it is the only variable of the three that must be an integer by definition. It is also conveniently the variable with the smallest range, from 0 to perhaps 25 vehicles per 30-sec interval, whereas both occupancy and speed are observed over at least four times this range. Using volume results in the fewest slices to deal with later in the algorithm. The slices were made parallel to the speed versus occupancy projection so that each slice contained points representing various speed and occupancy pairs observed at a given flow rate.

The next task was to design a perimeter about the observations on each slice of the three-dimensional region. Several methods were tried. The best approach started with the further division of each speed-occupancy slice of the three-dimensional region into columns of constant occupancy. The choice of occupancy over speed was arbitrary, but it has the advantage of a constant range of observations (from 0 to 100 percent), which simplifies coding. The mean and standard deviation of the speed observations at each occupancy were calculated. The vertices of the perimeter were the mean plus or minus a number of standard deviations for each occupancy value (see Figure 7). This method provided a simple way to test points against the calibrated perimeter. New observations could be compared with the maximum and minimum of the speed range (derived from the mean and standard deviation) for the appropriate volume slice and occupancy column.

One simplification for the algorithm, reducing computation time and memory requirements, involves the assumption that the standard deviation of speed for a given flow-occupancy pair is constant over time. Only the means need to be updated as the algorithm runs operationally. To test this assumption, 3 days of FTMS data from 1989 at one station on the QEW-Mississauga system were examined. The days covered a range of weather and traffic conditions. For selected volume-occupancy combinations, the F-test was used to compare the variance of speed on each day with that on the other 2 days. The results do not strongly support the assumption of constant standard deviation, but they do show that the most consistent variances occur in the mid-range of volume-occupancy points. Hence, it is probably acceptable to retain the assumption.

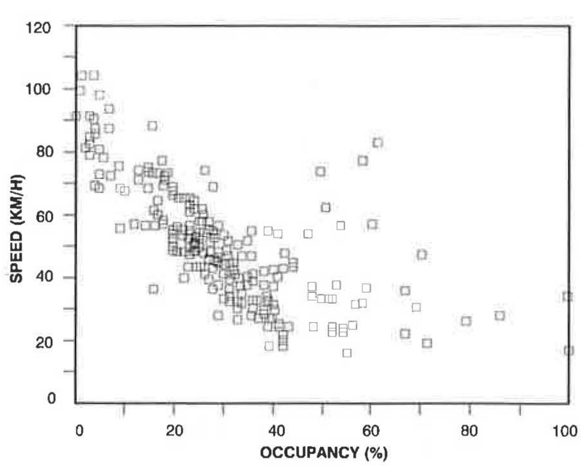

FIGURE 6 Three hours of FTMS data illustrating invalid combinations of speed and occupancy.

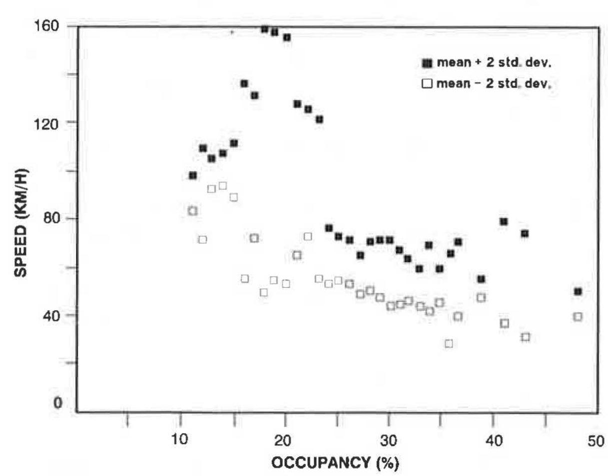

FIGURE 7 Calibrated speed ranges based on 7 days of data from QEW-Mississauga FTMS in 1989 (standard deviation calculated using 25 observations).

A final question about speed distribution remains: How many standard deviations should be used for the algorithm to accept good values and reject bad ones? It is known that a range of plus or minus three standard deviations will contain 99.7 percent of a normally distributed population. To check the assumption that speeds are distributed normally at each volume-occupancy pair, 4 days of weekday data from the median lane at three different stations on the QEW-Mississauga FTMS were examined. For each speed distribution, the coefficients of skewness and kurtosis were calculated. These statistics are qualitative indicators of the normality of a distribution. Skewness indicates whether or not one tail of a distribution is thicker than the other, with a zero value denoting a symmetric distribution. Kurtosis indicates the peakedness or flatness of a distribution relative to the normal distribution. A negative kurtosis represents a broad-peaked distribution, whereas a positive kurtosis indicates a sharp-peaked distribution. As with skewness, a zero value occurs when the peak of the distribution is normal.

As Table 1 indicates, roughly half of the distributions were skewed to each side, although the values were close to zero in almost all cases. This finding suggests that speed is distributed almost symmetrically about the mean for all volume-occupancy pairs. The kurtosis, however, indicates that 71 to 88 percent of the speed distributions had broader peaks than those of the normal distribution. Thus, the frequency of observations is relatively constant about the mean, instead of dropping off as does the normal distribution. For this reason, the choice of the number of standard deviations is made at the calibration stage, after the data have been inspected.

The final aspect of the algorithm design is the provision for on-line updating of the pattern to compensate for basic changes in freeway operations caused by such conditions as unfavorable weather (*11*). If the data are acceptable, they are incorporated into the mean for the volume-occupancy pair. Even though the standard deviation is held constant, the movement of the mean serves to move the range of acceptable points as freeway conditions change.

Calibration and testing of the three-dimensional screening algorithm were accomplished off-line on a microcomputer. The program calibrated each volume-occupancy pair according to a specified number of observations required. Once a volume-occupancy pair was calibrated, the program immediately started screening with the calibrated mean and standard deviation, using a user-supplied multiplier for the standard deviation.

Two parameters are needed before calibration: (a) the number of observations required to build the distribution of speed at each volume-occupancy pair and (b) the range about the mean speed that will be accepted as valid, expressed as a number of standard deviations. Limited analysis of QEW FTMS data has shown that, to obtain an accuracy of ± 10 km/h at the 95 percent confidence level, the number of observations required is between roughly 15 and 40, assuming that the sample standard deviation is a valid approximation of the true standard deviation of the population. To obtain an estimate of the suitable number of standard deviations to use, screening trials were run using the 3 days of 1989 FTMS data from the QEW-Mississauga system. The 3 days of data were run as one contiguous piece of data, simulating 3 consecutive days.

Overall, performance was as expected when the number of observations and the standard deviation multiple were varied (see Table 2). Larger numbers of observations used for calibration result in fewer points calibrated with the same number of observations, so the time to full calibration is increased. Increasing the multiple of the standard deviation results in reduced rates of data rejection. The rejection rate drops dramatically when *M* changes from 1 to 2, but the rate is reduced only slightly as *M* changes from 2 to 3.

On more than one occasion, the acceptable range for speed given the volume and occupancy reported by one loop did not include the observed speed, but, when the volume and occupancy from the second loop were used, the observed speed fell within the acceptable range. For this reason, the secondary loop data should be used for a confirmation check if available, before an observation is rejected. In this way, speed data can be used to select between volume-occupancy data that are not consistent for the two loops at a station.

TABLE 1 SUMMARY OF SKEWNESS AND KURTOSIS RESULTS

Station	Number of Distributions	Skewness: Number Positive	Skewness: Number Negative	Kurtosis: Number Positive	Kurtosis: Number Negative
17	146	75	71	18	128
21	325	155	170	62	263
25	209	111	98	61	148

CONCLUSIONS

The quality of FTMS data is an issue critical to the operation of the entire system. Currently accepted screening methods

TABLE 2 SCREENING RESULTS FOR VARYING COMBINATIONS OF NUMBERS OF OBSERVATIONS USED FOR CALIBRATION AND NUMBER OF STANDARD DEVIATIONS AWAY FROM THE MEAN

	N = 20					N = 25				
M	# obs	ok	fast	slow	% fail	# obs	ok	fast	slow	% fail
1	2317	1630	430	257	29.7	2055	1495	329	231	27.3
2	2317	2226	6	85	3.9	2055	1989	0	66	3.2
3	2317	2285	0	32	1.4	2055	2030	0	25	1.2

	N = 35					N = 40				
M	# obs	ok	fast	slow	% fail	# obs	ok	fast	slow	% fail
1	1657	1218	254	185	26.5	1496	1095	235	166	26.8
2	1657	1605	0	52	3.1	1496	1443	0	53	3.5
3	1657	635	0	22	1.3	1496	1477	0	19	1.3

	N = 65				
M	# obs	ok	fast	slow	% fail
1	863	621	131	111	28.0
2	863	821	0	42	4.9
3	863	847	0	16	1.9

Notes: M is the multiple of standard deviations used in screening
N is the number of observations used for calibration
obs is the number of observations screened
ok - observation passed screening
fast - speed was judged too high by screening
slow - speed was judged too low by screening

leave considerable room for improvement. Systems with paired loops in each lane at each station have more data screening options than those with single loops. Not only can the paired-loop systems compare the received speed-flow-occupancy points with a calibrated three-dimensional speed-flow-occupancy region, following on the Washington State idea, but they can also compare upstream and downstream loop values. A further advantage of tests based on two- or three-dimensional regions is that the single-variable threshold tests would be unnecessary and could be eliminated. Hence, the overall data screening could be much more critical of the data.

The two- and three-dimensional tests provide better data screening and, at the same time, monitor the system for defects more effectively. A detector that consistently produces unreasonable combinations should be investigated and recalibrated. The improved monitoring resulting from this increased sensitivity would result in a system that produces better data overall and thus more successful operation of algorithms for such purposes as incident detection.

ACKNOWLEDGMENTS

This work was supported by a contract from the Ontario Ministry of Transportation and by an operating grant from the Natural Sciences and Engineering Research Council of Canada (NSERC). Funding was also provided through a NSERC undergraduate student research award. The assistance of the Freeway Traffic Management Section of the Ontario Ministry of Transportation in making data available is greatly appreciated. Within that section, special mention should go to Phil Masters, Paul Lim-Hing, and Mark Fox.

REFERENCES

1. C. L. Dudek, C. J. Messer, and A. K. Dutt. Study of Detector Reliability for a Motorist Information System on the Gulf Freeway. In *Transportation Research Record 495*, TRB, National Research Council, Washington, D.C., 1974, pp. 35–43.
2. K. G. Courage, C. S. Bauer, and D. W. Ross. Operating Parameters for Main-Line Sensors in Freeway Surveillance Systems. In *Transportation Research Record 601*, TRB, National Research Council, Washington, D.C., 1976, pp. 19–26.
3. L. Chen and A. D. May. Traffic Detector Errors and Diagnostics. In *Transportation Research Record 1132*, TRB, National Research Council, Washington, D.C., 1987, pp. 82–93.
4. C.-Y. Jeng and A. D. May. *Monitoring Traffic Detector Information and Incident Control Strategies: A Survey*. Working Paper UCB-ITS-WP-84-2. Institute of Transportation Studies, University of California, Berkeley, June 1984.
5. L. N. Jacobson, N. L. Nihan, and J. D. Bender. Detecting Erroneous Loop Detector Data in a Freeway Traffic Management System. In *Transportation Research Record 1287*, TRB, National Research Council, Washington, D.C., 1990, pp. 151–166.
6. H. J. Payne, E. D. Helfenbein, and H. C. Knobel. *Development and Testing of Incident Detection Algorithms, Volume 2: Research Methodology and Detailed Results*. Report FHWA-RD-76-20. FHWA, U.S. Department of Transportation, April 1976.
7. E. R. Case and K. M. Williams. Queen Elizabeth Way Freeway Surveillance and Control System Demonstration Project. In *Transportation Research Record 682*, TRB, National Research Council, Washington, D.C., 1978, pp. 84–93.
8. *Burlington Skyway Freeway Traffic Management System Software Functional Description*. Report WP-85-81-02. IBI Group, Toronto, Ontario, Canada, and JHK & Associates, Atlanta, Ga., July 1983, rev. Feb. 1984.
9. P. Athol. Interdependence of Certain Operational Characteristics Within a Moving Traffic Stream. In *Highway Research Record 72*, HRB, National Research Council, Washington, D.C., 1965, pp. 58–87.
10. F. L. Hall and B. N. Persaud. An Evaluation of Speed Estimates Made with Single-Detector Data from Freeway Traffic Management Systems. In *Transportation Research Record 1232*, TRB, National Research Council, Washington, D.C., 1989, pp. 9–16.
11. F. L. Hall and D. Barrow. Effect of Weather on the Relationship Between Flow and Occupancy on Freeways. In *Transportation Research Record 1194*, TRB, National Research Council, Washington, D.C., 1988, pp. 55–63.

Publication of this paper sponsored by Committee on Freeway Operations.