Development and Evaluation of a Breadboard Video Imaging System for Wide Area Vehicle Detection

PANOS G. MICHALOPOULOS, ROBERT FITCH, AND BLAKE WOLF

Traffic engineers are constantly seeking new technology and equipment to deal with the problem of urban congestion. Among the most promising concepts available today is the use of video imaging for vehicle detection, automatic surveillance, and extraction of data needed for developing advanced control concepts. A recently developed video detection system is presented in this paper. This system operates on real time, can simultaneously detect traffic at multiple points within the camera's view, and emulates loop detectors. The system was installed and tested both off-line and in real time through taped data and field installations, respectively, and was directly compared to loops. The results suggest similar accuracy levels. In speed measurements, higher accuracies are expected for video systems than for loops. Finally, software is being developed for real-time extraction of traffic parameters, state variables (i.e. queue lengths and size), and measures of effectiveness (delays, stops, energy consumption, etc.) by the same device.

Vehicle detection appears to be the weakest link in traffic surveillance and control. Although accurate equipment is available for detecting vehicle presence on the roadway, it essentially employs technology of the late 1950s, has limited capabilities, presents reliability problems, and often requires massive and expensive installation for true traffic-responsive control. The latter is particularly true in state-of-the-art surveillance and control systems, which often involve large-scale street or freeway corridor networks. Regarding reliability, most cities with mature systems in the United States report that, at any time, 25 to 30 percent of their detectors are not functional or operating properly. Furthermore, discussions with suppliers and manufacturers suggest that often loop detectors, the most widely used detection device, seem to be active but actually produce false or inaccurate actuations. Finally, adverse weather conditions or pavement reconstruction present additional challenges for maintaining these detectors.

Perhaps the most important drawback of existing detectors is their limitation in measuring important traffic parameters and accurately assessing traffic conditions. This is because the technology employed represents a "blind" type of detection; only the presence or absence of vehicles over the detectors can be assessed with high accuracy. Traffic parameters, such as speed and traffic composition queue length, must be derived from presence or passage and require multiple detection, which increases cost and exacerbates the reliability problems mentioned earlier. Furthermore, common detectors (such as loops) do not have surveillance or sufficient vehicle recognition capabilities; most importantly, they are not flexible—they detect traffic only at fixed points. This is an important drawback for traffic control since the detection points should vary with speed, volume, and control objective.

Despite the aforementioned problems, existing detectors cannot be casually dismissed as they represent proven technology that will continue to serve its purpose in the foreseeable future. However, recent advances in image processing, electronic cameras, special-purpose computer architectures, and microprocessor technology have made the machine vision alternative for vehicle detection attractive, economical, and promising. A machine vision system for vehicle detection consists of an electronic camera overlooking a long section of the roadway. A microprocessor or a larger computer determines vehicle presence or passage from the images received by the camera, and derives other traffic parameters, preferably in real time. Vehicle detection can be obtained at specific points of the roadway while other traffic parameters can be derived by analyzing the images of the entire roadway scene. In the system described here, the microprocessor alternative was selected.

The advantages of vehicle detection through image processing are many, as a video detection system (VIDS) has multitasking capabilities. While performing its basic detection functions, it could simultaneously derive traffic measurements locally (using a microprocessor) or at a central location, perform surveillance functions, act as a vehicle counting and classification station, detect incidents and alert a human operator, and recognize special vehicles (ambulances, fire trucks, buses, etc.). There are, of course, other secondary tasks that a VIDS system can perform, such as (a) collecting and preprocessing data to be used in conjunction with existing traffic software packages, (b) revealing the nature of an incident by transmitting images of the scene, and (c) recording data for accident analysis, reconstruction, etc. Finally, it can be used as an evaluation device for measuring and assessing the flow quality or deriving measures of effectiveness for traffic studies.

An imaging detection system does not disturb the pavement and should, therefore, improve reliability, especially during reconstruction operations. Additionally, it can detect traffic at multiple spots of the roadway, within the camera's view, thereby becoming cost effective. For instance, in a previous

---

P. G. Michalopoulos and B. Wolf, Department of Civil Engineering, University of Minnesota, Minneapolis, Minn. 55455. R. Fitch, Triple Vision, Inc., 777 Harding St., Suite 4, Minneapolis, Minn. 55413.

feasibility study performed in Minnesota (*1*), it was estimated that fully instrumenting an intersection with a VIDS system (4 cameras, 1 microprocessor) would cost less than loop detectors, assuming that at least 3 loops/approach are required. Furthermore, simultaneous detection at 30–40 points using one or more cameras is possible. It was also estimated that the VIDS system design presented here would save 35 percent in maintenance costs and reduce the man-hours required by about 70 percent. Further savings could be realized if the same microprocessor also performed control functions (a viable alternative), thereby eliminating the need for a separate controller.

The flexible detection configuration of VIDS, combined with its ability to extract traffic variables difficult to obtain by conventional detection devices, suggests that the system should be particularly effective for automatic surveillance and control of saturated networks.

Because of these advantages, research on a cost-effective image-processing system for vehicle detection began to evolve during the mid-1970s in the United States, Europe, and Japan. In 1984 research at the University of Minnesota started through projects funded by the Minnesota Department of Transportation and later by the Federal Highway Administration (FHWA). As a result, a real-time multispot breadboard system was just completed, installed, and tested in several actual situations. Placement of the detectors with a mouse device can easily be accomplished by the user in minutes, by placing detection lines on a television monitor in any desirable configuration. Once these "pseudo-detectors" are placed, the system generates presence and passage signals compatible to loops, measures speeds, and generates essential traffic parameters such as volumes, headways, and occupancy. Furthermore, the system allows visual inspection of detection results along with actual traffic conditions for validation purposes and optimization of detector placement. The latter can easily be changed as often as desired either manually or automatically. Special algorithms for treating artifacts such as rain, snow, shadows, pavement reflections, etc., were developed. Also, the system can operate under both day and night conditions. Finally, any ordinary video camera used for surveillance purposes can be hooked to the breadboard system, i.e., no special-purpose cameras are required, although it should be evident that better-quality cameras, without blooming or streaking characteristics, improve the system's accuracy and effectiveness. Unlike earlier experimental units, the breadboard system not only operates in real time and deals effectively with all the aforementioned artifacts (rain, shadows, snow, etc.), but also operates under all traffic conditions. This is an important attribute since background compensation, when congestion sets in, is very difficult and has not been previously researched.

Following the initial algorithm laboratory development, testing was performed for algorithm optimization using videotaped data. Subsequently, the system was installed at the freeway surveillance and control center of the Minnesota Department of Transportation in Minneapolis and tested against live data from several cameras. The results are very encouraging and they suggest performance comparable to loops. For this reason, following additional testing, prototype fabrication is planned in early 1989.

In this paper, the breadboard system is described along with the facility that was also developed for quick algorithm testing and optimization. A brief description of the detection methodology is also presented followed by test results.

BACKGROUND

Research on image processing for vehicle detection began to evolve during the 1970s in the United States, Europe, Japan, and Australia (*2*). In the United States, research on this topic was initiated by the FHWA and conducted by the Jet Propulsion Laboratory (JPL) (*3–5*). Although the major objective of this project was individual vehicle tracking, algorithms for vehicle detection and speed measurement were also developed. The imaging system developed by the JPL, called Wide Area Detection System (WADS), was recently evaluated by Sperry Corporation (*6*). In this study, recommendations were made for improving the hardware and software design of the WADS system (*7*). Briefly, although the work performed by the JPL was pioneering, the WADS system was too primitive for practical applications; however, this should be expected at the initial stages of new technological developments.

Several countries are currently funding research and development on this subject including (a) work in England on image processing applied to traffic at the University of Manchester Institute of Science and Technology (UMIST) (*8*), the University of Sheffield (*9*), and the University College, University of London (*10*), (b) a vehicle tracking system being developed in France by the National Research Institute for Transportation and Security (INRETS) (*11*), and (c) a real-time multispot detection system being developed in Belgium by Devlonics, Ltd.

The UMIST project utilized a solid-state camera generating a 100 × 100 pixel/frame image at 8 frames/sec. The camera was mounted at a height of 22.5 m above a two-way highway and data was collected during a period in which illumination varied by a factor of 4. The output was digitized and averaged. An image corresponding to road background in the absence of any vehicle was stored in the memory of the digital processor. During operation, the digitized image was subtracted from the reference image to generate the road background. In the absence of vehicles, the two images should be similar and therefore their difference was due to noise and changes in illumination. A threshold was then used to compare the differences of the two images. The resulting binary image was compressed and stored on video cassette and processed in the laboratory. This system was not implemented in real time and would only work in ideal conditions where the background did not change significantly and where there were no common artifacts, such as shadows and reflections, to cause false detections.

The system currently being developed at the University of Sheffield operates under the assumption that the roadway background does not change significantly over a period of 1 min, which is considered to represent ideal conditions. This approach is highly prone to errors due to illumination changes, shadows, and reflections. At the University College, University of London, the focus is on implementing vehicle tracking on real-time, parallel image processing computer architectures. This vehicle detection approach requires a background to be manually sampled, which is impractical in field situations, so work is under way to automate this estimation. Once objects are separated from the background, features needed for vehicle tracking are extracted. These features would result in tracking not only vehicles but also common artifacts which would generate a substantial number of false detections.

In France, INRETS is also developing a real-time vehicle tracking system. The system automatically determines the roadway lane positions and then tracks vehicles down each lane through the entire camera field of view. The major problems with this system are that it can lock onto common artifacts, as well as have problems tracking vehicles through various background changes (e.g., asphalt/cement boundaries, building shadows) and in congested situations where the tracking mechanism breaks down. Finally, some observed test sequences indicated that the primary objects being tracked were dark shadow areas under vehicles and not the vehicle itself.

Recently, Devlonics, Ltd. of Belgium has advertised a real-time system that can accomodate up to 4 detection spots, each covering a 10-m lane area. The approach taken, which originated in cooperation with the Catholic University of Louvain (*12*), was to detect vehicles relative to an automatically determined reference background and track their movement through the 10-m area so as to also determine vehicle speed. Little detailed information about the approach taken is available; however, it was learned that vehicles must move through the 10-m area in less than 2.5 sec or they become part of the background signal. Furthermore, a microcomputer is needed to implement the detection for each spot, so the full 4-detection-spot system requires four microcomputers. Preliminary testing of the system in the Netherlands revealed detection problems in the presence of rain, shadows, congestion, and other artifacts. Additionally, the system does not seem to operate in real time but with a 5-sec constant decision delay, which is too long for critical intersection control applications.

The Japanese government sponsored the Institute of Industrial Science, University of Tokyo, research on measuring traffic flow using real-time video processing (*13–16*). The non-imaging sensor designed by Shigeta and Ooyama is of interest (*17*, *18*). The sensor is an array of photoelectric elements with geometry designed to match the perspective distortion produced by the camera installed at a specific height and angle of view. The photoelectric elements have a spectral response with a maximum of 930 nm that is thought to be optimum during the complete 24 hr day/night cycle. Detection is produced by illumination differences which are discovered by pairs of sensors. The distance between these sensors is known, and by measuring the time difference between detection by the first and the second element in a pair, the speed of the vehicle can be estimated. This system was tested in Tokyo for two years. The Shigeta-Ooyama system, which is the most cost effective, is not truly an imaging system and cannot be extended beyond simple detection as it requires fixed roadway placement geometries and has only fixed and discrete detection points in the field of view.

The Australian Research Board has developed a real-time vehicle presence system (*19*) that allows placement of up to 16 detection spots at any position in the camera field of view via front panel thumbwheel switches. To determine the background level, an additional reference detector is required, which must be placed in an area free of vehicles. This reference is compared with the detection spot outputs and, when fixed thresholds are exceeded, a vehicle is detected. Each detection spot has a manual offset adjustment to compensate for the difference in road surfaces between the reference and detector areas. The approach works adequately for ideal situations, but the system cannot distinguish the difference between vehicles and major artifacts such as vehicle shadows, reflections, and building shadows. Also, since the detection algorithms are hard-wired, there is no flexibility to reprogram and improve the system.

Experience with machine vision over the past 5 years suggests that despite the impressions generated throughout the literature, a reliable, fieldable, real-time multispot vehicle detection system is still lacking. The major problems with existing systems that have been addressed and resolved by the breadboard system are as follows:

1. The inability of existing systems to automatically adapt to a wide variety of backgrounds without reference marks prevents them from running reliably or autonomously. A unique approach to estimating the background at the detection spot was therefore developed; this allows automatic adjustment to any uniform or nonuniform road surface without operator intervention at startup or while running.
2. The operation of prior approaches in the presence of common artifacts such as shadows, illumination changes, and reflections has resulted in these systems having high false alarm rates. In the system presented here, these problems were resolved using a vehicle-signature-based detection approach that can differentiate vehicles from these artifacts.
3. Congested traffic conditions and stopped vehicles have caused the loss of the vehicle and erroneous background estimation in prior approaches. The VIDS system allows vehicles to stop for much longer periods of time without "blending" into the background.
4. Most existing systems only support a small number of fixed position detectors, and not the arbitrary placement of any type of detector in any configuration within the camera's field of view. In contrast, using the VIDS system one is able to place detection spots of any number, size, and shape anywhere in the camera's field of view, and one can reposition these spots dynamically under software control. This is accomplished without requiring the camera to be placed at a fixed geometry (e.g., height or angle).
5. Existing approaches to cost-effective real-time implementations have resulted in oversimplification of the sensor, hard-wiring the detection processing, or using prohibitively costly processors. Cost effectiveness was a major consideration in the development of the VIDS detection system. The system can operate with standard video cameras; no specialized sensors are needed. The approach taken in developing the VIDS breadboard system allows operation in real time while still being fully programmable. By using an IBM AT-compatible personal computer for the breadboard rather than an expensive image processing platform, it is demonstrated that the final system implementation is cost effective.

DEVELOPMENT OF THE BREADBOARD SYSTEM

The objective of this breadboard system was to fully develop and implement, in real time, the functions of vehicle presence, passage detection, and speed estimation with performance comparable to magnetic loop-based systems. Emulation of loops was considered an essential step for fully demonstrating concept feasibility. Derivation of other traffic parameters and measures of effectiveness can easily follow from these basic detection functions, considering that this is being accomplished at multiple spots within the camera's field of view. This latter limitation, including the range of operation (currently up to 500 ft), is mainly a function of the camera placement which is nominally assumed to be 40 ft; however, these limits will increase during prototype development. In addition to the multispot detection capabilities, the system is applicable to both freeway and intersections. The detection spots can also be dynamically positioned (without having to reposition the camera) by the system's software or by the user. Thus, the system is not locked into a fixed geometry, i.e., it is a true imaging system. To demonstrate cost effectiveness and feasibility, the breadboard system is based on an IBM AT-compatible personal computer which maximizes the use of existing software while minimizing custom hardware. The final prototype will be substantially more compact.

Detection and speed algorithm development were based on videotaped data recorded in real traffic situations, under a variety of environmental conditions, at both intersection and freeway locations. The main purpose in collecting this data was to capture as many different conditions as possible in order to achieve a high confidence in the algorithm performance. The data was recorded over three years at the Minnesota Department of Transportation (MnDOT) Traffic Management Center (TMC) in Minneapolis from video camera sites in the MnDOT freeway surveillance system. Data were also collected at sites in Michigan, Florida, California, and Maryland. The site selection was based on close proximity to both intersections and freeways; in this manner, collection of both types of traffic data was possible.

Over 70 hr of videotaped data sequences were collected covering a number of lighting, vehicle, traffic, weather, and other conditions, as summarized in Table 1. For initial algorithm development, 50 representative sequences were selected from these videotapes (79,000 images) and transferred to optical video disks. Video disks were used because of their rapid, random access, excellent image stability, step framing capabilities, compact archiving, and low cost.

These recorded data were used in the development and evaluation of the presence, passage, and speed algorithms. The development and evaluation of these algorithms were done on an IBM AT-compatible personal computer. The development facility was first used for off-line algorithm development and evaluation (i.e. verifying the presence, passage, and speed algorithm performance on the recorded video data). The same system was then used to evaluate the algorithms in real time while connected on-line to video cameras in the field at the MnDOT TMC. By using this common facility for both off-line and on-line development and evaluation, the time to transition from the lab to the field was greatly reduced.

TABLE 1 SUMMARY OF TYPES OF DATA COLLECTED

Conditions	Description
Lighting	Dawn, day, dusk, night
Traffic composition	Cars, trucks, semi-trucks, buses, ambulances, motorcycles, bicycles
Traffic flow conditions	Normal, congestion, queues, turning, stopped, multi-lane
Environmental factors	Clear, overcast, fog, rain, snow, haze, abrupt lighting change, hot/cold temps, high humidity
Artifacts	Vehicle shadows, building/sign shadows, cloud shadows, sun glare, wind motion, reflection, occlusion, lens spots

DEVELOPMENT FACILITY AND REAL-TIME SYSTEM

Once again, it should be stressed that although the development facility is based on an IBM AT-compatible personal computer, the final design will be a much smaller, less expensive, and self-contained device. The primary reasons for choosing an AT compatible system were the low cost and high availability of software and peripheral hardware (e.g., video digitizers), the high throughput of the processors, the excellent software support, and the dedication to continuing software compatibility; for instance, the system is fully compatible with the newer 80386-based machines such as the Compaq 386, even though initial development began on an IBM PC based on the 8088 microprocessor.

The development facility hardware system is shown in Figure 1. The personal computer was equipped with a real-time video digitizer that accepts video from any standard video source and converts it into a digital format that can be processed by the PC-AT. The video sources include an optical video disk player, a videocassette recorder (VCR), and a video camera or demodulator. Video output from the digitizer was used to display (or record on a VCR) the digitized results overlayed with graphic information. More output to a digital tape recorder was used to record real-time results from the real-time detection algorithms. The PC-AT also contains a 12 Mhz 80286 microprocessor, 3 megabytes of processor memory, two 1.2-megabyte floppy disks, and a 40-megabyte hard disk, along with graphic display terminals for both monochrome and color graphics and text outputs. For real-time

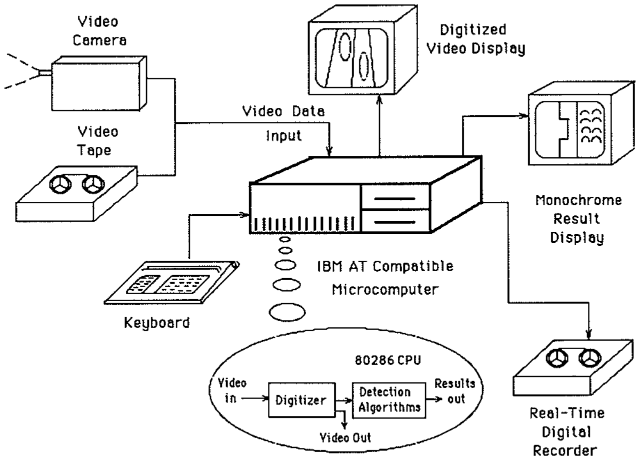

FIGURE 1 Development facility hardware system.

operation, only the real-time video digitizer and one other video image preprocessing circuit card (formatter) are needed in addition to the standard computer.

The development facility software capabilities include

(a) interactively placing multiple detection spots in any number of lanes and at any position along the roadway (in the camera field of view),

(b) ground-truthing sequences of images that are on the video disk (ground truthing is manually determining if a car is present in a detection area for repeated algorithm performance "scoring"),

(c) iterating various processing algorithms on image sequences,

(d) determining algorithm performance automatically,

(e) examining results of processing and experimenting with alternative vehicle detection techniques,

(f) selecting and processing specific video disk image frames or sequences, and

(g) calibrating detector positions to correct for roadway perspective.

ALGORITHM DEVELOPMENT

Presence and passage detection signals were generated at all the detection spots within the field of view. These "pseudo-detectors" were interactively placed by the user at any position in the camera field of view and at any orientation. Examples of possible detector placements are in Figure 2. Pairs of closely spaced detectors were used to estimate vehicle speeds. Detectors across lanes were primarily for vehicle passage, while downlane (or longitudinal) detectors sense vehicle presence. Multiple crosslane detectors were also used for area presence.

Certain spatial and temporal features needed to be extracted for detection and speed estimation for each detector. Spatial features provide information on vehicle signature regardless of the vehicle's speed, while temporal features respond to vehicle motion. Spatial features are the relationships between intensity values across a detector at any instant in time. Temporal features were taken for each detector over a number of time samples (i.e. over a number of image frame times).

The extracted spatial and temporal features were combined using sequential decision processing to generate the background detection and vehicle presence and passage detection signals. Reliable background detection, and its adaption to a wide range of both uniform and nonuniform backgrounds, is a key improvement over earlier approaches; it should be noted that the background is automatically determined by the system—no assumption is made about the road surface signature or its uniformity. Background values were continuously updated and a special logic was developed for updating when a vehicle is present. This logic prevents the background from being "lost" or falsely determined in congested or stopped vehicle traffic situations. Given reliable background estimation, vehicle detection is determined by differences relative to the background level.

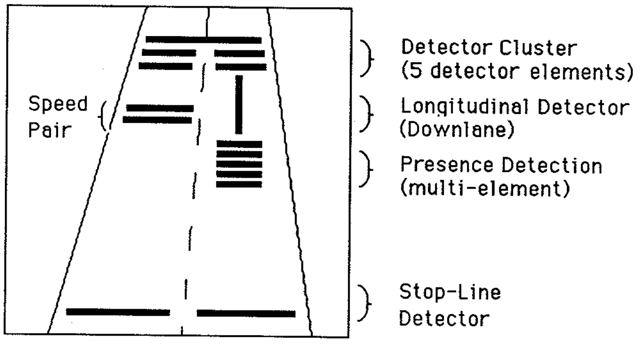

FIGURE 2 Example detector configurations.

Edge-based features are customarily used in research to detect motion. They provide good separation between vehicle signatures and those of common false-alarm-generating artifacts such as shadows (vehicles, clouds, fixed objects), illumination changes (camera AGC, transition periods, lightning), and reflections (headlights, sun glint). As a result, it should be possible to suppress most of the false alarms associated with these common artifacts.

In addition to vehicle detection, individual vehicle speed was also measured. This was accomplished by using pairs of closely spaced detectors and measuring the time it takes the vehicle to move between the detectors. This is shown conceptually in Figure 3. By estimating the time (*t*) that it takes the vehicle to travel from the first detector (*D*1) to the second (*D*2) and knowing the distance between the detectors (*d*), the speed (*s*) was easily estimated. For higher speeds, this time can be reliably measured using the difference in time between the passage signals generated by the vehicle detection algorithms (similar to a speed trap used with loops). But in some situations, such as in congestion, the passage signal generation was not reliable enough to generate an accurate speed measurement. In fact, at lower speeds the accuracy of passage signal generation was increased if speed was used by the vehicle detection algorithms. As a result, a speed estimation technique that works independently of vehicle detection had to be developed.

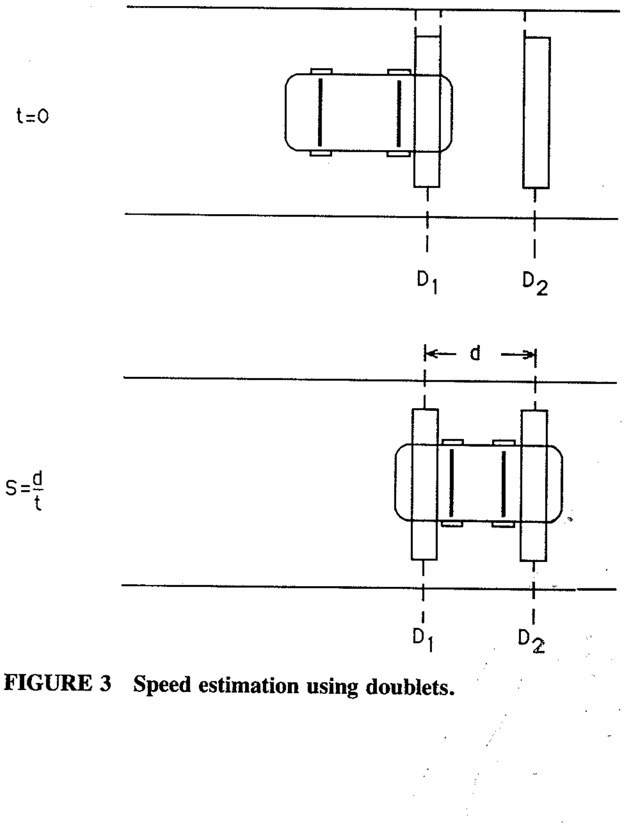

FIGURE 3 Speed estimation using doublets.

Briefly, the technique compares features generated by each detector and registers them in time using a signal correlation technique. This yields estimation of instantaneous speed as opposed to average speed so that the speed estimate can be used directly for improved control purposes. This technique does not have to rely on either the vehicle detection or background detection outputs.

PERFORMANCE TESTING

The performance of the detection and speed estimation algorithms has been evaluated continuously during the development stage as well as after the completion of the real-time breadboard system. The evaluation was accomplished off-line using the algorithm development facility, and was performed using the video disk sequences; this was key to developing vehicle detection and speed estimation algorithm concepts which could be quickly tested on a large data set (over 79,000 image frames on the video disks). Off-line evaluation is currently being performed to optimize the speed-estimation algorithms, but the vehicle-detection algorithms are now past this stage and are being evaluated on-line in real time. On-line evaluation was made possible when real-time implementation of the detection algorithms was completed. This evaluation was performed in real time on both videotapes and used live data from cameras at the MnDOT TMC. The latter required installation of the breadboard system, which monitors freeway traffic in Minneapolis and St. Paul through 36 camera installations. The breadboard system was connected with several of these cameras to allow visual inspection of the detection outputs. The on-line evaluation allowed quick determination of the system's performance over long time periods on many image frames. On-line evaluation is ongoing as the detection algorithms are being improved; evaluation of other traffic parameters and measures of effectiveness, such as occupancy, time headways, queue lengths, stops, and delay, is planned.

VOLUME PERFORMANCE

Three primary evaluation measures were used to determine the performance of the vehicle detection algorithms: detection accuracy, miss error rate, and false alarm error rate. Detection accuracy is the ratio of the number of vehicles correctly detected by the system to the total number of vehicles (determined by visual inspection). This measure indicates how many vehicles actually present were detected by the system. Conversely, the miss error rate is simply 100 percent minus the detection accuracy and is often referred to as just the "error." The false alarm error rate is the ratio of the number of vehicles falsely detected to the total number of vehicles. This measure indicates how many times the system indicated a vehicle was present when there was none. Note that the detection accuracy and the false alarm error rate do not add to 100 percent as do the detection accuracy and miss error rate.

The real-time vehicle detection performance evaluation allowed direct processing of videotape and live camera data that resulted in extremely fast evaluation on a large data set. At the time of this writing, an all-day live evaluation of the system was just completed at two locations monitored by the MnDOT TMC. At the first location, the camera monitored a freeway section, while at the second an intersection was monitored. At the freeway site, passage (counting) detection was evaluated and presence detection was evaluated at the intersection. The system was run from 7:00 a.m. until 10:00 p.m. on July 15, 1988. Every ½ hr, the system was run continuously for 5 min for both cameras. Over 350,000 image frames containing 8,299 vehicles were processed and scored. These images included a number of significant artifacts including shadows, clouds, rain, congestion, and occlusion; over 1,000 shadows were counted. For 987 vehicles at the intersection, the performance over the entire day was a 98.5 percent detection accuracy (972 out of 987 vehicles), a 1.5 percent miss error rate (15 out of 987), and a false alarm error rate of 5.1 percent (detected 50 vehicles when none were present). For the 7,312 vehicles on the freeway section, the performance was a 91.8 percent counting accuracy (6,712 out of 7,312), an 8.2 percent miss error rate (600 out of 7,312), and a false alarm error rate of 2.3 percent (detected 168 vehicles when none were present).

A plot showing the performance for each 5-min time slice is shown in Figure 4. On each plot the vertical axis represents percentage and the horizontal axis the time of day from 07:00 (7:00 a.m.) to 22:00 (10:00 p.m.). Each plot shows the detection accuracy and false alarm error percentages during each 5-min time-slice every ½ hr; to avoid misunderstanding, it should be reiterated that the miss error rate is defined as 100 percent minus the detection accuracy.

As shown in the plot for the intersection, the detection accuracy ranged from 95 percent to 100 percent throughout the entire day. The false alarm error rate remained below 10 percent except between 08:00 and 09:00 when 60-mph winds caused the camera to move the detection area over a lane marker resulting in false vehicle detections. Wind similarly affected performance around both 10:30 and 17:30. This wind problem can be significantly reduced by using smaller and more aerodynamic camera housings (the one used was an older, bulkier design with approximately a 3 ft^2 "sail" area). Recent camera installations are more rigid and compact and do not seem to move significantly. This wind effect can also be compensated for electronically by measuring the scene shift and translating the digitized image; the prototype will have this capability if camera motion is determined to be a problem.

For the freeway case, Figure 4 indicates that the detection accuracy remained above 90 percent for the entire day, except between 15:00 and 17:00 when congestion caused some pairs of vehicles to appear as one vehicle, which resulted in lower vehicle counts. It should be reiterated that scoring on the freeway was done by counting vehicles (passage). If vehicle presence had been scored instead, the performance would have exceeded 90 percent, even during this period, since presence accuracy remains high even in congested situations. The algorithms for counting in congestion are still being optimized. The false alarm errors shown in the plot were less than 10 percent throughout the day. The majority of the errors occurred between 08:00 and 10:00 (peaked at 10 percent) and were due to large numbers of vehicle shadows on the road.

In addition to evaluating the system all day for 5-min periods every hour, the performance of the system operating for an entire hour at two additional freeway locations with near-capacity traffic was determined. In the first location (Case 1 in Table 2) evaluation was performed for four detection spots placed as shown in Figure 5. The camera was viewing I-35W

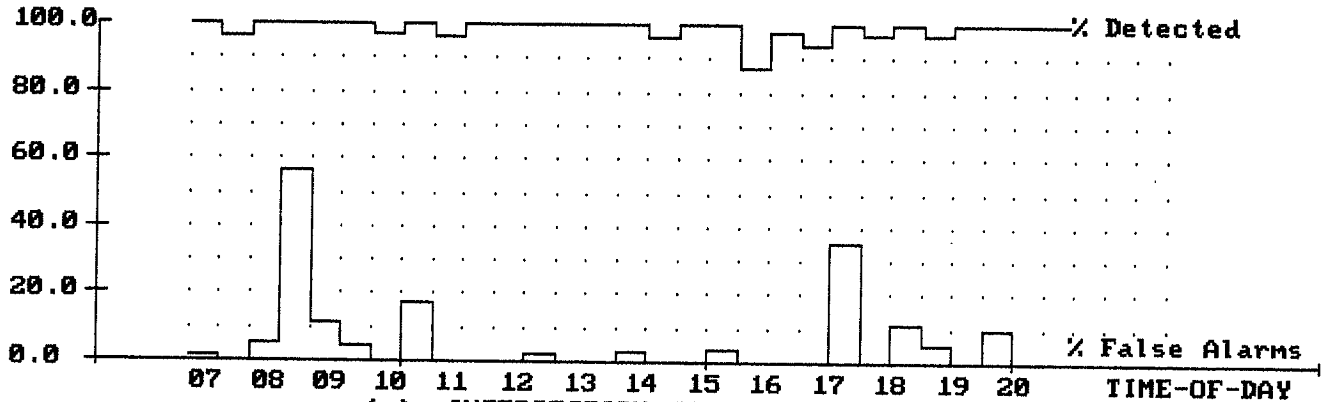

(a) INTERSECTION PRESENCE

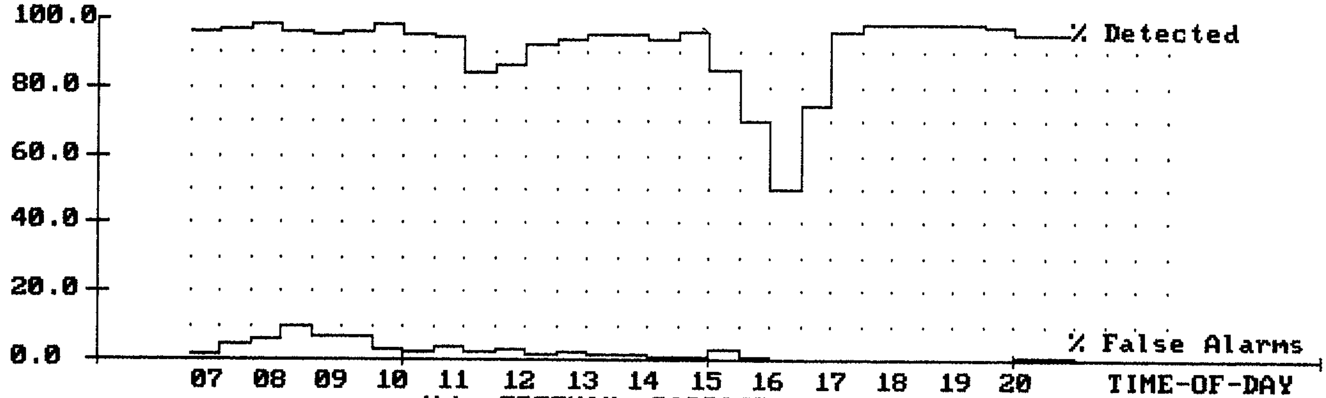

(b) FREEWAY PASSAGE

FIGURE 4 All day performance (7 a.m. to 10 p.m.).

in Minneapolis at 60 ft above the road and approximately 150 ground ft from the first detection spot. Vehicle volumes (counts) were determined automatically by the VIDS system for twelve 5-min time slices from 16:15 to 17:15 (rush hour). Actual volumes were determined by manual counting. The average volume to capacity ratio was approximately .86 assuming a capacity of 2,400 veh/hr/lane.

The performance results for each of the four detectors is shown in Table 2 (Case 1), which indicates the actual and measured volumes on each detector, the root mean square error (RMSE), and the mean absolute error. The RMSE for the entire hour was computed from

$$\left[\frac{\Sigma\ (\text{manual count} - \text{VIDS count})^2}{\text{no. of 5-minute time slices}}\right]^{1/2}$$

TABLE 2 PERFORMANCE IN NEAR CAPACITY TRAFFIC

	Volume (Veh.)	RMSE (Veh.)	Mean Absolute Error (%)
Case 1: Manual/VIDS (1 hour)			
Detector 1	1893/1838	5.68	2.85
Detector 2	1888/1850	8.89	4.78
Detector 3	2256/2205	8.04	3.91
Detector 4	2252/2246	7.48	3.18
Case 2: Loop/VIDS (1.5 hour)			
Detector 1	1754/1737	1.68	1.42
Detector 2	2293/2287	6.13	2.91

As Table 2 suggests, the RMSE ranged from 5.68 to 8.89 vehicles per 5-min interval while the mean absolute error ranged from 2.85 percent for Detector 1 to 4.78 percent for Detector 2. These error levels are negligible despite the high percentage of trucks during the period of data collection. Naturally, traffic composition affects performance since trucks and tall vehicles tend to occlude the camera's view, thereby reducing detection accuracy. In the second test site, the error levels were further reduced when truck composition was closer to normal levels. In general, errors increase with distance from the camera and in the farther lanes due to taller vehicles occluding smaller vehicles. Increased range also causes consecutive vehicles to appear contiguous—two vehicles counted as one. However, algorithms dealing with this problem will be developed, and further testing will be performed to determine whether the errors due to this artifact are significant enough to justify further algorithm development.

The current detection performance evaluation approach requires someone to validate all results manually. This is a time-consuming and error-prone process that can only be used in a limited number of cases. To eliminate this manual step, an automated scoring process is being developed to directly

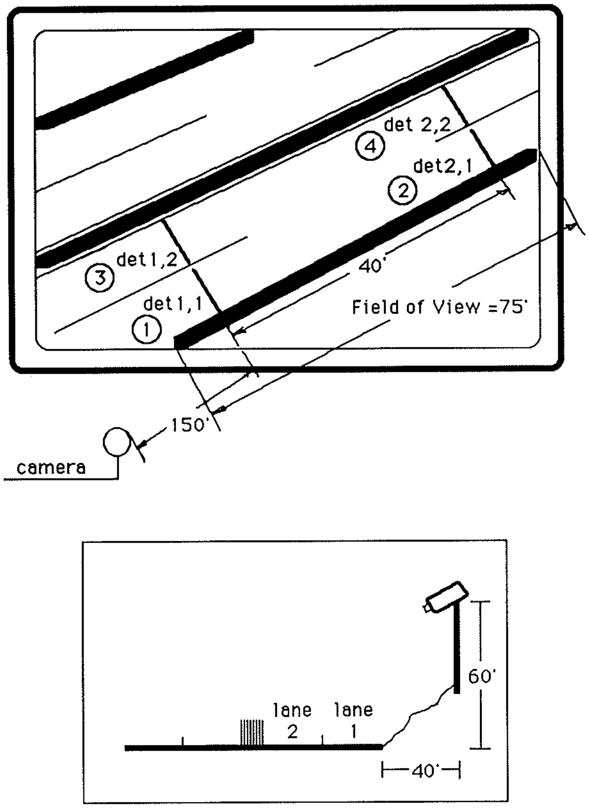

FIGURE 5 One hour evaluation geometry.

compare the VIDS system outputs to those from colocated magnetic loops. The MnDOT TMC has a number of surveillance cameras which can view positions on the freeway where loops are present. The counts from these loops are transmitted back to the TMC where they are logged every 30 sec by the TMC central computer. At this writing, a second camera site was selected (Case 2), and the accuracy of the loop-generated counts was confirmed by manually verifying the loop counts as opposed to counting vehicles from a videotape. This manual verification is necessary since some of the loops are nonfunctional. Subsequently, the VIDS system was run on this same videotaped sequence and the results to the loop counts were compared. The performance for this site is shown in Table 2 (Case 2). This second location is similar to the first location (2 lanes, camera mounted on side of road, Detector 1 in near lane, Detector 2 in far lane) except the camera pole is positioned closer to the roadway (10 ft) and the evaluation period increased to 1.5 hr. Also, the traffic composition contained a normal amount of truck traffic whereas Case 1 contains heavy truck traffic; as a result crosslane occlusion was reduced and system error was lower than expected. The system performance is further evidenced in Figure 6, which presents counts for comparison purposes measured from this second test site. In addition to loop (short-dashed line) and VIDS (long-dashed line) volumes, the manually derived volumes (solid line) that were used for loop verification are shown for both Detectors 1 and 2.

For further testing, the VIDS system is currently being interfaced to the TMC central computer to directly read the loop counts for real-time comparison with the VIDS system output. This will allow the system to automatically score its performance while running continuously and unattended all day, which will allow more extensive performance evaluations.

SPEED EVALUATION

The speed estimation algorithms have been evaluated off-line on the video disk sequences by percent errors and percent misses. The percent error measure is the comparison of the estimated speed to the actual speed; the actual speed is measured visually by counting the number of video frame-times it takes the vehicle to traverse the two detectors. The percent misses is the percent of vehicles that did not register a speed measurement (i.e., the algorithm was not able to estimate a speed). The evaluation was performed on 66 video disk sequences containing 391 vehicles. These sequences included vehicles travelling at speeds from 0 (stopped) to 70 mph in all weather conditions during both day and night. The overall error was 12 percent and the misses were 17 percent. It should be noted that these performance numbers are for instantaneous speeds and not average speeds; the average speed performance would improve with the amount of time the instantaneous speeds are averaged.

The majority of the speed estimation errors were caused by having an insufficient number of samples of vehicles moving at high speed. Since the spacing of the two detectors used for all of these test sequences was 10 ft, a vehicle travelling at 60 mph (88 ft/sec) traverses a 10 ft trap in approximately 4 video frame-times; this could result in an average estimation error of 25 percent (a 1 frame-time estimation error). This error can be decreased by increasing the spacing between detectors. For a 40 ft detector spacing, error rates of 6–7 percent have been measured for vehicles moving approximately 60 mph. In fact, the system's dynamic detector placement capability allows this spacing to be automatically adjusted as the speed estimate changes. The speed estimation misses primarily occurred in sequences with extremely heavy fog (too much noise in the video signal), stopped vehicles (the vehicle reached the first detector but not the second), and sequences where one of the detectors was in a fixed shadow (from a building) and the other in the sun (this caused poor signal correlation). The speed estimation algorithms were only recently improved to deal with these problems and integrated into the real-time system for purposes of performing more extensive on-line evaluation. Preliminary test results obtained suggest accuracies of 94–96 percent. Finally, the signal correlation technique described herein does not require presence and passage signal extraction and in preliminary testing resulted in speed measurement accuracy of 90 percent or higher.

CONCLUSIONS

Despite major worldwide efforts to develop a machine vision system for traffic surveillance and control, a real-time device having the capabilities and performance required for practical applications has been elusive. The system presented here may not have the requirements of a commercial product but, at least in terms of detection performance, it is compatible with existing devices such as loops. Speed measurements, developed only recently, are already very satisfactory and should

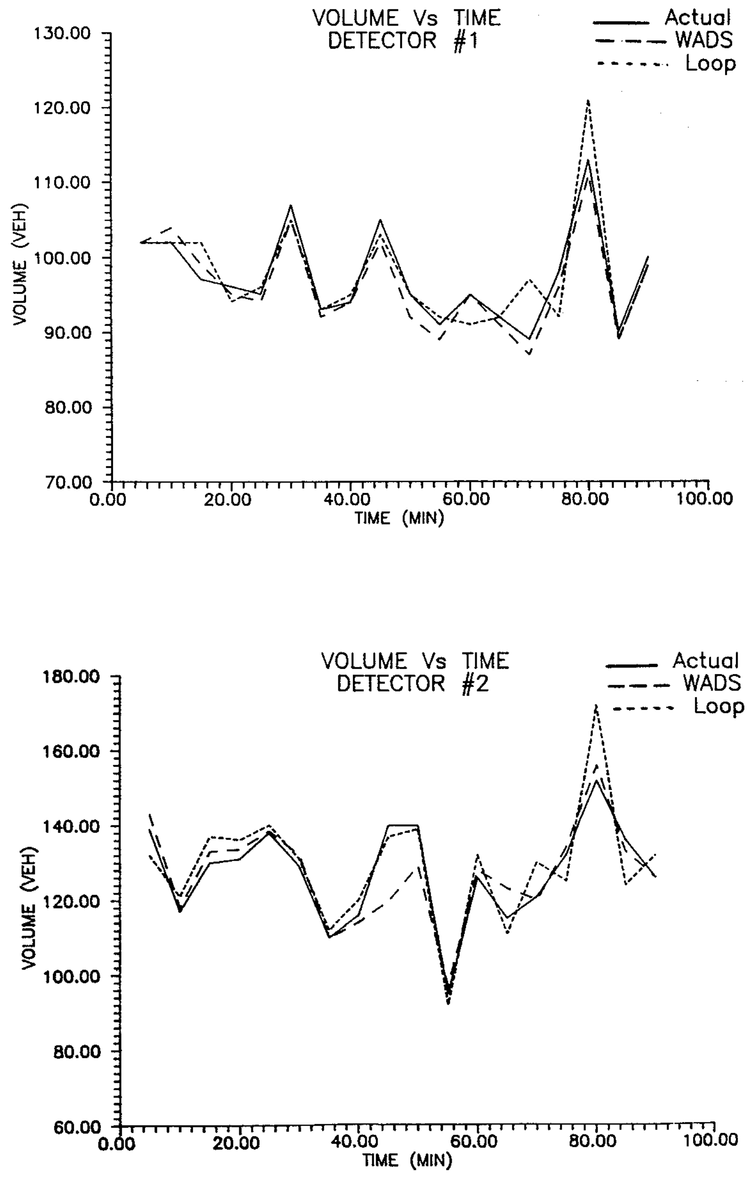

FIGURE 6 Volume comparisons for one-hour period for data derived manually, from loop detectors, and from the WADS system.

improve with further experimentation. It is worth noting that the detection test results presented here refer to individual vehicles rather than averages. Naturally, averaging the measurements over time would result in lower errors, but this method would only mask true system performance and therefore it was not attempted here.

Live demonstrations of the system in its present form to professional engineers and potential users generated favorable comments and lead to the conclusion that the research team should proceed with prototype development and demonstration projects. At least partial funding for prototype development is already available and is expected to be supplemented by early 1989. At this time the prototype development phase should begin; in the mean time, demonstration projects in several states and cities are being considered.

The serious consideration of the system in these locations at this early stage of development is primarily owed to the system's expected impact in traffic surveillance and control. Indeed, the major advantages of this machine vision system lie in the multispot, multilane, wireless detection capabilities. Along with recent advances in image understanding, the system should essentially be transformed to an "electronic eye" for computerized surveillance and control or for automating time consuming and expensive functions (performance evaluation, derivation of measures of effectiveness, etc.). Finally, an imaging detection system can measure traffic variables that cannot easily or accurately be measured by conventional detection devices. For example, queue length and size can be extracted by VIDS without much difficulty; measurement of these parameters requires many loop detectors and cannot be obtained if the queue extends beyond the last detector. Similarly, at this time, density can only be approximated from occupancy; this variable can be measured and more accurately by VIDS. In short, the system should provide the machine vision link needed to take advantage of recent technological innovations in microprocessors, artificial intelligence, and telecommunications. Clearly, the research effort described here suggests that off-the-shelf technology and equipment for developing a cost-effective video detection system are available today.

ACKNOWLEDGMENTS

The authors would like to acknowledge the cooperation of the technical staff of MnDOT's Traffic Management Center; especially valuable was the assistance of D. Stregge, G. Carlson, R. Lau, C. Harens and G. Bodoczy.

REFERENCES

1. P. Michalopoulos, R. Fitch, and M. Geokezas. *Development of a Visible/Infrared Vehicle Detection System: Feasibility Study.* Department of Civil and Mineral Engineering, University of Minnesota, 1986.
2. R. M. Inigo. Traffic Monitoring and Control Using Machine Vision: A Survey. *IEEE Transactions on Industrial Electronics*, Vol. IE-32, No. 3, 1985, pp. 177–185.
3. E. E. Hilbert et al. *Wide Area Detection System—Conceptual Design Study*. JPL Interim Report, FHWA-RD-77-86. FHWA, U.S. Department of Transportation, 1978.
4. A. P. Schlutmeyer et al. *Wide Area Detection System (WADS)*. JPL publication 83-15, FHWA, U.S. Department of Transportation, 1982.
5. E. E. Hilbert et al. A Sensor for Control Arterials and Networks. *IEEE Transactions of Vehicular Technology*, Vol. VT-29, 1980, pp. 208–215.
6. Sperry Corporation. *Wide Area Detection System (WADS)—Applications*. Interim Reports, FHWA, U.S. Department of Transportation, 1984 and 1985.
7. P. Michalopoulos, R. Fitch, and B. Wolf. *Development of Wide Area Vehicle Detection System (WADS): Image Recognition Algorithms*. Center for Transportation Studies, Department of Civil and Mineral Engineering, University of Minnesota, FHWA, U.S. Department of Transportation, 1986.
8. R. C. Waterfall and K. W. Dickinson. Image Processing Applied to Traffic: 2. Practical Experience. *Traffic Engineering and Control*, 1984, pp. 60–67.
9. A. Houghton, G. S. Hobson, L. Seed, and R. C. Tozer. Automatic Monitoring of Vehicles at Road Junctions. *Traffic Engineering and Control*, Vol. 28, No. 10, 1987, pp. 541–543.
10. N. Hoose and L. G. Willumsen. *Real-Time Vehicle Tracking Using Advanced Image Processing Techniques*. Technical Report. University College, University of London, England, 1987.
11. S. Deucher, J. M. Blosseville, and F. Lenoir. Traffic Spatial Measurements Using Video Image Processing. *SPIE Proc., Intelligent Robots and Computer Vision*, Vol. 848, 1987.
12. W. Maes, A. Barbe, and A. Vits. Development of a Detection System Based on the Use of a CCD Camera. In *Proceedings, Seminar on Micro-Electronics for Road and Traffic Management*, Organization for Economic Cooperation and Development, Paris, France, 1984, pp. 78–86.
13. S. Takaba et al. Measurement of Traffic Flow Using Real-Time Processing of Moving Pictures. In *Proceedings, 32nd Conference on Vehicular Technology*, IEEE, New York, 1982, pp. 488–494.
14. S. Takaba and N. Koishi. Real Time Simulation of Road Traffic Flow. Summer Computer Conference, Vancouver, Canada, 1983, pp. 585–590.
15. S. Takaba. Measurement of Flow of Vehicles and Pedestrians Using Real Time Processing of Moving Pictures. In *Proceedings, International Symposium of Processing and Applications*, Institute of Industrial Science, University of Tokyo, Japan, 1984.
16. S. Takaba, T. Sekine, and R. Hwaug. A Traffic Flow Measuring System Using a Solid-State Image Sensor. In *Proceedings, International Conference on Road Traffic Data Collection*, IEEE, London, England, 1984.
17. K. Shigeta, N. Ooyama, and S. Matsumoto. A New Traffic Flow Sensor Using Photoelectric Converting Elements—Operating Principle and Practical Experiments. *Mechanical Engineering Laboratory*, Technical Reports, Namiki Sakura-mura, Japan, Vol. 35, No. 1, 1981, pp. 21–48.
18. N. Ooyama and K. Shigeta. Area Measurement of Traffic Flow Using Photoelectric Elements. *Proc., 4th IFAC/Int. Fed. for Info. Processing/Int. Fed. Operational Res. Soc. Conference*, Baden-Baden, FRG, 1983, pp. 281–287.
19. J. S. Dods. The Australian Road Research Board Video Based Vehicle Presence Detector. In *Proc., International Conference on Road Traffic Data Collection*, London, England, 1984, pp. 96–100.

Publication of this paper sponsored by Committee on Applications of Emerging Technology.