

Deployment of State-of-the-Art Technology for Incident Management: *The Gowanus Expressway Project*

By Panos G. Michalopoulos of Image Sensing Systems, Inc.

And

Kevin Samartin of Traffic Control Equipment Corporation, Hackensack, New Jersey

Submitted to the
Transportation Research Board
1998 Annual Meeting

ABSTRACT

Although successful deployment of new technology in Advanced Traffic Management Systems is constantly sought in ITS projects, it is often not realized. One of the limited number of new technologies that is proving successful is wide area video detection. In this paper successful deployment of this technology in freeway traffic management projects is presented. Even though the technology can be used in any traffic management application requiring vehicle detection, incident management appears to be its most popular application on freeways so far. This may be because incident detection, response, and management is still one of the major challenges in urban freeway operations, requiring constant attention and considerable investment in manpower and equipment. However, despite efforts worldwide, fast and reliable automated incident detection has been elusive. In this paper, a new automated incident detection and management system based on wide area video detection (machine vision) is briefly reviewed. The advantage lies in the wide area detection capabilities which allow efficient detection of shock waves and other traffic parameters that cannot be easily obtained by conventional devices. In addition, video detection allows employment of ancillary information such as traffic on the shoulders, stopped vehicles, lane changing and speed differential, traffic slow downs in the opposite direction, etc.

Following many years of development and field implementation at several beta sites, the system has been substantially improved and a more sophisticated version was recently deployed on the Gowanus Expressway in Brooklyn, New York. This paper describes the fieldable version of the system along with results observed up to the time of the writing of this paper. So far the system successfully detected 81% of the incidents, had no false alarms, and reduced incident response time substantially. Deployment of the technology in various forms at several other sites for freeway management is also summarized.

Key Words: Machine Vision, Incident Management, Incident Detection, Wide Area Detection, Advanced Traffic Management Systems (ATMS)

INTRODUCTION

Successful deployment of state-of-the-art technology in the field is essential for supporting the ITS movement. However, as in most programs of a similar nature, progress requires long-term experimentation, risk taking, and commitment—a fact that is often forgotten. It is perhaps for this reason that truly innovative and effective technology deployment is frequently not realized. One of the new technologies which is gaining acceptance among practicing traffic engineers is wide area video detection (machine vision). This paper presents the successful deployment of this technology in a recently completed and operational project for freeway traffic management. Even though the technology can and is employed in any traffic management application requiring vehicle detection and wide area data extraction, incident management appears to be the most popular application of this technology on freeways at this early stage. This is because incident detection, response, and management is one of the major challenges in urban freeway operations requiring constant attention and considerable investment in personnel and equipment. While several methods are currently employed for detecting incidents, automated techniques are becoming increasingly important for decreasing the detection/response time, reducing the operator's fatigue, and increasing reliability. However, despite continuous efforts worldwide, fast and reliable automated incident detection and management has been elusive. Conventional, automated techniques based on computerized algorithms are less effective than desirable for operational use as they generate a high level of false alarms or missed incidents. Manual, operator-assisted methods, on the other hand, minimize the false alarm risk but also suffer from missed or delayed detections, are labor intensive, and restrict the potential benefits of advanced integrated traffic management as they require human attention for detecting incidents rather than confirming, responding, and managing them through computer-aided means.

Perhaps the major handicap of existing automatic incident detection algorithms is that they are designed to operate with the limited data provided by existing conventional vehicle detection devices. This information alone, typically volume and occupancy, has not proven to be sufficient for effective and reliable incident detection, partly because volume is not a dynamic measurement and partly because occupancy is a surrogate rather than a true measurement of a spatial traffic flow variable, namely density.

Furthermore, the measurements upon which current detection algorithms must rely are essentially taken at a point rather than over space.

Video detection is receiving much attention for incident detection applications because of its ability to detect over a wide area. Video detection has now been available commercially for several years and is gaining acceptance as a more effective tool than conventional inductive loop detectors. The additional benefits of using video are many. Key is the ability to cover many lanes with one camera and extract wide area measurements, such as density, queue length, speed profiles and others. Additionally, lane closures are typically not needed during installation which results in increased driver and traffic personnel safety and minimal traffic disruption. In fact, once installed, they are typically used during subsequent road construction or resurfacing by repositioning cameras, as needed, as the road geometry varies.

Finally, if desired, the video can be used to provide or supplement existing video surveillance. For these and other practical and theoretical reasons video detection systems have generated much interest for those involved in advanced traffic management systems. Unfortunately, video-based incident detection systems proposed or implemented so far are rather simplistic. They, in effect, have degenerated to simply detecting the onset of congestion while they are unable to detect incidents beyond the effective field of the camera's view (only a few hundred feet for effective image processing), as they lack sophisticated algorithms or incident detection logic capable of capitalizing on the real advantages of machine vision.

The incident detection system described here was initially developed as part of the Minnesota Department of Transportation's (Mn/DOT) Guidestar project to deploy and evaluate IVHS technology in the I-394 corridor. The video detection deployment system on I-394 in Minneapolis has been designed to be a live laboratory for operational and research purposes. The cameras, 39 in all, have been closely spaced over a 3.5 mile stretch of road to supplement traffic simulation efforts with extensive measurements in time and space and to evaluate camera placement alternatives and determine optimal camera spacing for incident detection.

Additionally, the detection data are being provided to the live ITS laboratory of the Center for Transportation Studies at the University of Minnesota along with loop detector data and video for congestion monitoring, incident verification, and ramp control. The incident detection system discussed here is an outgrowth of an earlier system described in the next section. It is worth observing, however, that following many years of development and experimentation in Minnesota, the system was enhanced and incident management features were added. In addition, the system was installed experimentally in Montgomery County, Maryland prior to large-scale deployment on the Gowanus Expressway, on Interstate-75 (Atlanta Olympics), on the Olympic Road Freeway in Korea, and in several other projects. The I-75 (Atlanta) and Korean deployments are described in References (1) and (2). The Gowanus deployment is presented here. It should be noted that in spite of the fact that the enabling technology used in all projects is the same, there are significant differences in

the specific system deployments at each site. For instance, Atlanta is using the machine vision technology in connection with a different incident detection/management scheme of its own design, while Gowanus expanded on the basic system described here. Other deployments, briefly presented as well, have their own idiosyncrasies.

BACKGROUND

The initial development of the machine vision-based Automated Incident Detection Algorithm (AIDA) was completed using an 80386-based PC, interfaced to a video detection system. The video detection system used was the Autoscope 2002, first developed by the University of Minnesota (3). Because only single, isolated, video detection stations were available, the detection algorithm was required to use a single station rather than multiple stations. There are several advantages to a single station algorithm. To begin with, no sensor data from an adjacent detector station is required, thus reducing algorithm complexity and deployment concerns regarding spacing and communications issues. Furthermore, the single station can be used as stand-alone in isolated locations. Several algorithms were developed and compared and are discussed in Reference (4).

However, since the AIDA algorithm was believed to be the most promising, it was selected for continued development and evaluation. The algorithm initially used speed, occupancy, and volume data as provided by the video detection system. These were averaged into 30-second intervals, and thresholds were applied to look for rapid changes indicative of a capacity reducing event. Later AIDA was improved to include ancillary information provided by video detection such as stopped vehicles and shock wave signature recognition.

The initial AIDA algorithm was evaluated on continuous, around-the-clock, real-time data collected over a four-month period from December 1991 to March 1992 using video detection station data from I-35W at 26th Street in Minneapolis. The 26th Street test site is a very complex one. It is just upstream of a point where I-35W splits into two freeways and a downtown exit ramp. The freeway that splits to the right has another exit ramp to a third freeway. The complexity provided an abundance of incidents especially during the winter season. Performance evaluation was accomplished by comparing the AIDA alarms to handwritten incident logs maintained by the operators at Mn/DOT's Traffic Management Center (TMC). Eighteen logged incidents were judged to be relevant to the 26th Street location.

The following convention was used to calculate algorithm alarm rates. The alarm rates are normalized to one multi-lane station and for a one-day time period. We define one alarm to consist of the entire time period from when an alarm turns on to when it turns off. Thus whether an alarm has a duration of five minutes or one hour, it is weighted as one alarm. The reason for this is that once an alarm is turned on, it will receive an operator's attention thereafter, rendering the duration irrelevant. The operator will either identify the incident causing the alarm or determine that it is a false alarm, or in the case of sudden and severe recurrent congestion, continue to monitor for incidents. In the

case of false alarms, the operator will turn the alarm off immediately and not have to deal with it any further.

Overall performance summaries for the AIDA algorithm during the initial four-month operation are discussed in detail in Reference (4). For this period the AIDA algorithm produced a total of 73 alarms, resulting in a station alarm rate of 0.6 alarms per day. A major difficulty with the evaluation was that alarms were generated off-line and compared to the operator's log periodically rather than evaluated in real-time as they occurred. Of the 73 alarms only 14 could be matched to confirmed incidents and two were identified as false alarms. The remaining 57 alarms could not be traced to any known cause. Eleven occurred after hours when no operators were present to log incident activity. Thirty-four occurred during rush hours and may have been caused by recurrent congestion, lane closures, unrecorded minor incidents or even missed incidents--there is no way of knowing. The sudden capacity reduction in all 23 off-peak hour incidents appeared serious enough to merit an operator's attention, whether they were caused by actual incidents or merely a heavily congested freeway system. Instant review and response by an operator or on-site personnel would have proven invaluable for algorithm evaluation.

This incident detection system was subsequently operated on-line continuously since the initial evaluation until October 1993, collecting 30-second interval statistics data and also alarm data. In addition, video tapes of traffic were automatically recorded when incident alarms were declared. The 30-second interval data and video data were archived and used in off-line testing to refine and evaluate further algorithm changes as discussed in the next section.

VIDEO-BASED INCIDENT DETECTION SYSTEM DEVELOPMENT STATUS AND PRELIMINARY EVALUATION

Subsequent to the initial four-month evaluation, a number of improvements to the original system and algorithm have been made. First, the video detection system itself was greatly improved by the introduction of the third generation Autoscope 2003, which was designed for outdoor installations in traffic cabinets and met 170/179 and NEMA TS1/TS2 environmental requirements. Improved features include a more robust and accurate speed detector, which was redesigned to use a field of view calibration. Also, the calculation of interval statistics for time intervals ranging from ten seconds to one hour was added. In 1994, a new stopped vehicle wide area detector was developed which can detect the presence of stopped vehicles within the camera's field of view. This detector type can cover all lanes including the shoulders over a long range of the roadway (1000 ft. or more).

Second, the 80X86 PC platform video detection server has been streamlined and improved. It has been designed to talk to many detection devices either directly, or on multidrop communication lines, or any combination of the two as shown in Figure 1. The communications bandwidth requirements were greatly reduced by using the distributed processing afforded by the new video detection system. Interval statistics,

no longer calculated on the server, are obtained by simply polling each video detection system. A new level of service (LOS) congestion grade is displayed using user-selected thresholds either from the Highway Capacity Manual or from custom-defined thresholds. An application programming interface was developed to enable end users to write their own applications (such as their own incident detection logic, ramp control schemes, speed alarms, congestion level indicators, etc.) using data collected by the server. The new AIDA algorithm now runs as an application itself on the server. In fact, several different incident detection algorithms can be run concurrently if desired. Finally, to provide quick verification of incident alarms, integrated camera management capabilities have been added to automatically call the appropriate video camera output to the operator's computer screen or a user-selected monitor.

Third, a number of changes have been made to improve the AIDA algorithm by adjusting the turn-off logic, streamlining the turn-on logic, and adding logic to create an alarm when stopped vehicles on the shoulders or in the traveled lanes are detected. As mentioned earlier, in order to achieve this latter feature, a new type of detector was developed. This is called the stopped vehicle detector. It simply detects stopped vehicles within the field of the camera's view. This feature cannot be duplicated by loop detectors. The user can place many virtual stopped vehicle detectors interactively on the video monitor and set location and lane-specific thresholds which, if exceeded, will generate an alarm.

A typical incident detector configuration for a four-lane roadway is shown in Figure 2. To test the effectiveness of these changes and compare to the previous version's performance, the new algorithm was tested off-line using the archived data from the I-35W test site discussed in the previous section. Data from 30-second intervals from January 1993 through August 1993 were used to evaluate the effect of the changes. The resultant alarms were compared to TMC operator's handwritten incident logs as was done for the initial evaluation. The initial test results did not include stopped vehicle detector alarms, simply because this feature was added later (in 1994).

Again, because the comparison was made many days after the incidents were logged, it was not possible to verify whether alarms were justified. However, the detection accuracy compared to documented incidents that occurred within two miles downstream of the 26th St. test site, was 98.5%. The "likely" alarms were judged to be serious capacity reduction of the roadway with similar signal characteristics to other confirmed incidents. That is, the detection signals were of the same magnitude and duration as actual, confirmed incidents, and occurred at off-peak hours. The "uncertain" alarms were few and typically of a lesser signal strength, were very short in duration, and mostly occurred during rush hour congestion. It is critical to any such evaluation to correctly assess the "truth" of all alarms and tune the algorithms accordingly. This is best done in real time at the time an alarm is declared--an operator or on-site person can examine video in the vicinity of the alarm and confirm that an incident has occurred or report that a false alarm has occurred.

There was roughly one alarm per day per station generated on the average during this eight-month period with the new AIDA algorithm. This is higher than the station alarm rate of 0.6 alarms/day that the old algorithm produced during the earlier four-month evaluation period. Reasons for this increase in station alarm rate are in part due to changed alarm turn-off logic which has in some cases led to two alarms where before there was only one alarm of long duration. Secondary shock waves produced either during clearing or by perhaps minor secondary incidents appear to have caused the secondary threshold crossings related to the first alarm.

Alarms of unknown cause provided an upper limit to a false alarm rate for the new AIDA algorithm. Over the 243-day test period, the "unknown cause" alarm rate was 0.69 alarms/day/station and work continued to reduce this rate. It should be noted that because the evaluation was made off-line in batch mode, there was no opportunity to tune the incident detection logic as yet due to the time involved to analyze the data. Furthermore, the false alarm rate in this test should be less than this rate since some of these alarms will have been justifiable events requiring operator attention. Further field evaluation and analysis followed an AIDA deployment in Montgomery County, Maryland to reduce and characterize these alarms. This development started in 1993, and was completed in 1995. It involved close cooperation with the engineers in charge and followed a similar tedious approach for algorithm improvement. Recurrent congestion proved to be the main cause of false alarms and AIDA was modified to filter this out. Because of this improvement the subsequent deployment of AIDA in the Gowanus site resulted in significantly lower false alarm rates as it will be seen later.

Before concluding this section, another significant improvement is worth mentioning. Specifically in 1995, the incident detection logic was moved from the centrally located PC-based system server to the Autoscope unit in the field. This ensured that at large-scale installations a software failure would contain the problem to only a few cameras rather than the entire system. Of equal importance was the development of a more compact and reliable as well as considerably less expensive fourth generation Autoscope unit in 1995 (the 2004). This unit has also been substantially improved since then through several software releases to achieve substantially more effective shadow treatment, image stabilization, video frame compression, and numerous other performance and functional features. The incident detection system was also incorporated into the 2004.

The 2004 processes the traffic data (speed, volume, occupancy, etc.) in ten-second increments enabling incident alarms within 30 seconds, a 66% improvement. Along with this, as large-scale deployment of the video-based technology began to take hold, it became necessary to communicate with many machine vision units and manage the processed data in real time. For this purpose the ScopeServer Interface Developer's Kit (IDK) was developed. This IDK is a Microsoft Windows® application that provides communication between a central computer and over 100 Autoscope units (processing over 400 cameras). In this manner the integrated data can be used for developing other applications, such as ramp control, signal optimization, adaptive control, variable message signs, incident management, and others. The data (in up to ten-second

increments) can also be stored or graphically presented on demand and passed to multiple users. This allows many applications to run in parallel. For instance, ramp control, variable message signs, incident detection, and tunnel control can all run in parallel. Figure 3 shows how the ScopeServer, in conjunction with machine vision, can be integrated in a TMC for generating real-time traffic databases and implementing the most common incident and freeway management tasks. Finally, Figure 4 demonstrates the system's architecture for general TMC client-server applications.

GOWANUS EXPRESSWAY DEPLOYMENT

The Gowanus Expressway rehabilitation project in Brooklyn, NY is the first site in which the complete incident detection system described earlier was employed on a non-experimental basis and has been working flawlessly since May 1997. With the additional improvements presented here it can be viewed as one of the most advanced incident detection systems presently deployed in the US. The system consists of 20 CCTV surveillance cameras with pan/tilt/zoom capabilities supplementing the backbone of the project and the automated incident detection/management system. Other technologies that are in place to aid motorists consist of highway advisory radio, variable message signs, and a construction information hot-line.

When the New York State Department of Transportation (NYSDOT) made the decision to revamp the Gowanus Expressway/Prospect Parkway, it was determined that an incident detection system had to be employed to minimize the effects of road construction on traffic flow. The construction on the roadway does not consist of just resurfacing. Major structural work is being performed on the Gowanus Expressway, an elevated roadway, including expanding the width of the road and relocating a two-lane exit ramp. Activities of this nature and magnitude tend to cause major traffic flow disruption. As a result of the automated management system installation, traffic moves freely with minimal slow-down times relating to the clearing of an incident. The time period for constructing this system was roughly about three months, spanning the time period of August 1996 through October 1996. The construction consisted of both field work (installing the cameras, camera equipment, and fiber optic equipment), as well as the construction of the TMC itself. However, system integration was completed in April 1997. At that time the system became fully operational. Though the detection system was initially put into place to monitor the traffic on the roadway during the construction of the Gowanus Expressway, following its effective performance, it was decided to keep it in place when the rehabilitation project is complete, which is scheduled for January 1998.

In addition to the 20 surveillance cameras, 43 Imaging Sensors (cameras meeting machine vision specifications) were mounted on existing luminary poles, covering the entire construction area. Even though the construction area is 8km (five miles) long, a total of 16km (ten miles) of roadway is covered by the 43 Imaging Sensors (IS). The live video is fed to fiber optic transmitters which send the light signals via single-mode fiber optic cable to the TMC. Once the signals reach the TMC, they are transformed back to an electrical signal. The video signals for each of the 43 IS are fed to their

respective input on one of eleven Autoscope processors, called Machine Vision Processors (MVP). At that point the video image is digitized and the processing of the video occurs. Each MVP allows four video inputs for detection and a fifth input for surveillance purposes. The video output of each MVP is terminated to the video input panel of the CCTV matrix switcher. Though the video output of an MVP does feed into the matrix switcher, the TMC operators do not view the image from the IS. For troubleshooting and monitoring purposes, the system was designed so that authorized personnel have the ability to view the video from the image sensors on the bank of monitors by using the matrix switcher. The advantage to this design is that from a central location the video from all of the image sensors may be viewed quickly on a large screen without having to connect and reconnect to the video output of each of the eleven processors. It is possible to tour through each of the 43 image sensors to insure that they, as well as the virtual traffic sensors, are all functioning properly and are aimed at the correct location.

Operation

To provide incident detection, virtual detection zones are overlaid onto the field of view for each of the 43 image sensors via the Supervisor software. These detection zones, or virtual detectors, are then downloaded to the MVP corresponding to the specific IS. As the video images feed into their respective MVP in real time, the processors analyze the video and determine speed, occupancy, volume, stopped vehicles, and other traffic parameters as traffic moves through the detector zones. If the AIDA incident detection software determines that an incident has occurred, an audible alarm is generated, the live video from the corresponding imaging sensor (camera) is displayed, and the TMC operators take the appropriate action to confirm and clear a possible incident.

The TMC operators are an important factor in the process of determining whether an alarm is truly an incident. Six television monitors are utilized to allow the TMC operators to view conditions of the traffic on the roadways. The operators have the ability to pan, tilt, and zoom the 20 cameras that are deployed above the roadway. If an alarm is generated by AIDA, precedence takes place and automatically two CCTV cameras associated with the specific alarm, condition pan, tilt, and zoom to their preset alarm positions. The video images from the associated CCTV cameras are broadcast to two previously determined television monitors, which become dedicated to that alarm condition. At that point the operators make the determination whether the alarm is a true incident and take the appropriate action to clear any possible problems and keep the traffic moving freely.

Following system installation the TMC operators do not miss incidents due to such factors as daydreaming, processing paperwork, or using the restroom. Audible alarms alert the operators to suspected incidents. If an operator steps away from his/her command for a moment and an alarm condition arises, when the operator returns, the suspected incident area will be in view on the monitors and the alarm will not clear until the operator acknowledges that the alarm has been responded to.

The equipment that was used on this project is a clear example of the positive steps being taken to keep the traffic on roadways moving. By using fiber optics as a transmission system, surveillance cameras for monitoring purposes, and the machine vision processors for incident detection and traffic data storage, this system has been proven to be successful.

Further Design Considerations and Performance

The incident detection design consists of parallel systems operating separately and in conjunction with one another. The image processors collect traffic data and provide alarms of suspected incidents. When no incidents are confirmed or are being responded to, the operators use the surveillance cameras to monitor the roadway. It is important to allow the operators to tour through the cameras to view motorists having problems. These two systems work in conjunction with one another during an incident. The video feed from the image sensors is not viewed by the operators for monitoring purposes. Therefore, the operators are using the surveillance to provide the visual information necessary to determine the corrective action needed. Once an alarm is flagged, the surveillance system is automatically directed to position two adjacent surveillance cameras to view the incident location where the alarm was generated. Once the alarm has been responded to, control of all cameras and monitors is given back to the operators. In addition, warnings are being displayed through variable message signs and over the highway advisory system to enable the drivers to avoid the problem area. Warnings of delays are also posted on sites on the Internet.

Prior to AIDA installation, the average time to clear an incident was 1.5 hours. Since the automated incident detection/management system has been put into place, the average time to clear any incident has been substantially reduced. If a vehicle breaks down, the time to aid the motorist from the incident inception to the clearing of the incident is averaging 19 minutes. Similarly, if an accident occurs, the time from inception to clearing is now averaging 31 minutes. To aid in the speed of incident clearing, four tow trucks are stationed at various locations on the roadway.

Since the installation of the system was completed in April 1997, there was not time for long-term tests. However, the performance of the system during the first month of operation was found to be very satisfactory by the TMC operators to the point it was decided to keep the system in place after construction is expected to end. Specifically during the month of May 1997, one hundred incidents occurred within the 16km area where the system was installed. Of these, 81% were correctly detected by AIDA, 19% were missed, and there were no false alarms. Our experience indicates that false alarms have generally discouraged automatic incident detection. The low false alarm rates, of course, came at the expense of incident detection accuracy; however, the relatively small drop in detection accuracy (from over 90% to 81%) was considered more than satisfactory by the system operators.

ADDITIONAL FREEWAY TECHNOLOGY DEPLOYMENT

As mentioned earlier, the enabling machine vision technology described here has now been employed in many ATMS deployment projects worldwide. At the time of the writing of this paper and to the best of the author's knowledge, there have been over 1500 Autoscope installations processing over 5000 cameras for both freeway and intersection applications. In what follows brief descriptions of some representative deployments on freeways are presented to demonstrate the diversity as well as the most common practical uses of the new technology so far.

Atlanta's ATMS Deployment

Atlanta's brand new ATMS, which was designed for the 1996 Olympic Games, monitors traffic flow along I-75 and I-85 for incident management and provides up-to-date traffic information to the traveling public in the greater Atlanta area. The ATMS currently covers 90km (60 miles) of roadways with 316 IS cameras and 57 six-camera video detection units which are used to generate 5000 virtual detection zones. For this ATMS, the Georgia Department of Transportation (GDOT) selected the Autoscope as the most mature wide area video vehicle detection system for a variety of reasons, including minimizing lane closures during installations and any future maintenance activities. In this project the 57 MVPs provide speed, volume, occupancy, level of service, vehicle classification, stopped vehicle detection, and wrong way traveling vehicle detection, as they accumulate traffic statistics. As part of this ATMS project a 32-bit communications ScopeServer was employed, which runs on a Pentium PC with the Windows® NT operating system. A Digi C/CON 16-port concentrator provides separate serial communication channels to different MVP units for multidrop communication. This concentrator is expandable to support a larger number of channels as the Atlanta ATMS grows.

The ScopeServer uses the TCP/IP protocol to support platform-independent client communication. This enables, for example, the ATMS software running on Sun Microsystems's Solaris Workstation with the UNIX operating system to communicate with the field Autoscoptes. The ScopeServer IDK for client/server communication is also supported. This IDK was used by GDOT to develop a specially designed Count Station Data Acquisition System (CSDAS). Every 20 seconds the ScopeServer polls each of the image processing units and relays the data to the CSDAS database. The database provides real-time traffic data access to all ATMS applications, such as incident detection and map display. The ScopeServer software is fairly robust and well tested. It can retrieve a data package from the largest cluster or "hub" of all 57 six-camera MVPs in the Atlanta ATMS in merely three seconds, using a communication baud rate of 19,200 bps.

There are three major benefits of this machine vision-based ATMS in Atlanta. First, in the Atlanta ATMS, traffic data flows directly from each field MVP to a PC at the TMC, unlike other systems which use additional field controllers to convert the data from video or loop detectors to transmit to the TMC. Second, data acquisition and incident detection are performed in different computers connected through a client-server network, thereby achieving faster incident response, unlike other systems in which

incident detection must be done concurrently by the same computer which is acquiring the data and thus significantly slowing down the incident response. Finally, machine vision provides a variety of traffic data and a method to archive this massive amount of data for later analysis, which enables the traffic manager to detect traffic trends and problem patterns. Because of its successful deployment, GDOT is currently in the process of expanding the video detection system and the ATMS to an additional 112km (70 miles) and 300 machine vision cameras. This makes the Atlanta installation the largest known freeway deployment of machine vision worldwide.

New Jersey Turnpike

The NJ Turnpike has eleven four-camera MVP units deployed, and nine more are being designed into projects that will start in 1997. The units are deployed near the Newark Airport monitoring six lanes in each direction. In this deployment machine vision is collecting volume, occupancy, and speed. The occupancy data has proven to be compatible with occupancy obtained for inductive loops from neighboring roadways. The data is fed into a freeway management system running an incident detection algorithm based on occupancy (a version of the modified California Algorithm) from other sensors in the system. The machine vision sensors were deployed prior to construction to provide continuous detection and flexibility of detector placement for traffic diversion. Four cameras were deployed at each station (one camera per three lanes) because the agency was not sure of the traffic re-routing requirements during construction. The sensors are interfaced to local M-170 controllers in the field to collect real-time traffic data and communicate to the central traffic management system. The Turnpike Authority is planning to expand the same instrumentation to an additional 134 miles of freeway. Prior to this decision the performance of machine vision compared to radar detection was evaluated. The radar sensors did not pass the evaluation performance tests.

Integrated Corridor Traffic Management

Mn/DOT initiated the Integrated Corridor Traffic Management Project (ICTM) in 1994. Its goal is to instrument a parallel arterial corridor to a major freeway (I-494) with an adaptive intersection control (SCATS) and divert traffic to the arterials during heavy congestion or incidents. The project includes eleven four-camera MVPs deployed on intersections, ramps, and freeways. Five MVPs were acquired for use in the first three phases on a section of arterial roadway under construction. The units are providing stop line detection to enable signal phases during construction and will remain permanently installed at the end of the project to provide strategic and tactical detection for the SCATS. Additional units will be installed in follow-on phases on the freeway for adaptive ramp control, freeway incident detection, and freeway count station detection during roadway resurfacing that will destroy all inductive loops.

The technology is also being used by Mn/DOT at several freeway construction sites for detecting incidents and controlling variable message signs. For this purpose a special

portable video detection system had to be developed which transmits wireless video and data to a central location in the Twin Cities.

Additional Deployments

Another example of an installation that will be using the AIDA system is in Albuquerque, New Mexico. This installation not only uses machine vision-based incident detection at major on and off ramps, but also will control the signalized intersections located at these on and off ramp points. A total of four video detection units, each processing four cameras, will be deployed initially for this project. In this setup, twelve of the sixteen cameras will be used for intersection control, and four cameras will be used for freeway monitoring and incident management. An additional two to three four-camera video detection units will likely be installed in the near future as the project expands.

Two recent examples of machine vision installations which are monitoring freeway status (but not using incident detection), are in Houston, Texas and in Brazil. The installation in Houston is along the I-610 ring road circling the Houston metropolitan area. This installation is being implemented in several phases. The technology described here has already been chosen for two of the phases for a total of 44 MVP units and 127 cameras. A third phase is in the acceptance stage currently. This phase will include an additional 20 MVP units and 73 IS cameras. Machine vision will be used to report vehicle speeds and counts back to a central location via special controllers for traveler information and for general traffic management.

The Brazilian installation uses two video detection units and eight IS cameras to monitor bi-directional traffic on one stretch of freeway. The MVP units will report real-time data via a serial link back to a central location which will then disseminate the information and provide travelers with updated reports about traffic conditions via the Internet.

Overseas Deployments

As the technology is being accepted and deployed overseas, large-scale deployment is being experienced there as well. The largest known so far is a multiphase freeway traffic management project in Korea. The first phase of the project will cover 18km of the eight-lane Olympic Road highway from the airport to downtown Seoul.

The management system is designed to monitor the freeway traffic as well as manage and control day-to-day traffic. Machine vision will provide average vehicle speed, volume, and occupancy data, while AIDA will detect incidents on the freeway. The data and live video signals will be transmitted to the City's traffic management center via fiber optic lines. Special-purpose custom application software will poll each MVP every 30 seconds through the ScopeServer communication server. Processed data will then provide estimated travel time, and Changeable Message Signs (CMS) will display this information to motorists. The initial phase of the project included 34 cameras and seventeen two-camera Autoscoptes. Construction commenced in 1996, and the machine vision equipment was installed in May of 1997.

CONCLUDING REMARKS

After many years of experimentation and several generations of video sensors for wide area vehicle detection, it appears that the Minnesota technology is now mature and cost-effective for truly advanced traffic management applications. This, of course, did not occur without relentless long-term field experimentation and countless tests, comparisons and feedback from government agencies, as well as the willingness of the engineers responsible for the many projects incorporating the technology to take risks. To be sure, many difficulties were encountered, ranging from poor camera quality and placement to deficiencies in communications and the machine vision software. In the end, initial failures were turned into successes because of the willingness of the engineers to work with the developers of the technology, the spirit of pioneering new technology and, perhaps most importantly, the persistence and determination to succeed.

In the Gowanus Expressway project the decision to keep the system after construction is complete demonstrates the confidence of the responsible authorities in the system's reliability and performance, as well as its user-friendliness. In Atlanta, bringing together many different technologies, including machine vision, meant overcoming integration challenges of both infrastructure and software. Despite the initial problems encountered, the satisfaction of the GDOT and the benefits for Atlanta and the surrounding communities far exceeded expectations. It is for this reason that the ATMS system, including the machine vision part, is currently more than doubling its Centennial Olympics' size.

Hopefully, this paper presented at least a glimpse of the potential of successful machine vision deployment in innovative ATMS applications. As this and other technologies improve and take hold, traffic management systems such as the ones described here will become more common. The lessons learned in such projects should not only lead to more effective utilization of machine vision technology, but also to greater technology accessibility as well to lower costs.

ACKNOWLEDGEMENTS

Financial support for the basic research was provided by the Minnesota Department of Transportation, the Federal Highway Administration, and the Center for Transportation Studies at the University of Minnesota.

REFERENCES

- (1) Culver, Marcus. Video Detection: The Atlanta Experience. Traffic Technology International, Jan. 1997, pp. 40-43.
- (2) Vinger, Stephanie. Cameras for the Olympic Road. Traffic technology International, July 1997, pp. 113-114.

- (3) Michalopoulos, P.G. Vehicle Detection Through Video Image Processing: The AUTOSCOPE System. IEEE Transactions on Vehicular Technology, Vol. 40. 40, No. 1, 1991, pp. 21-29.
- (4) Michalopoulos, P.G., Jacobson, R.D., Anderson, C.A., and DeBruycker, T.B., Automatic Incident Detection Through Video Image Processing. Traffic Engineering & Control, Feb. 1993, pp. 66-75.

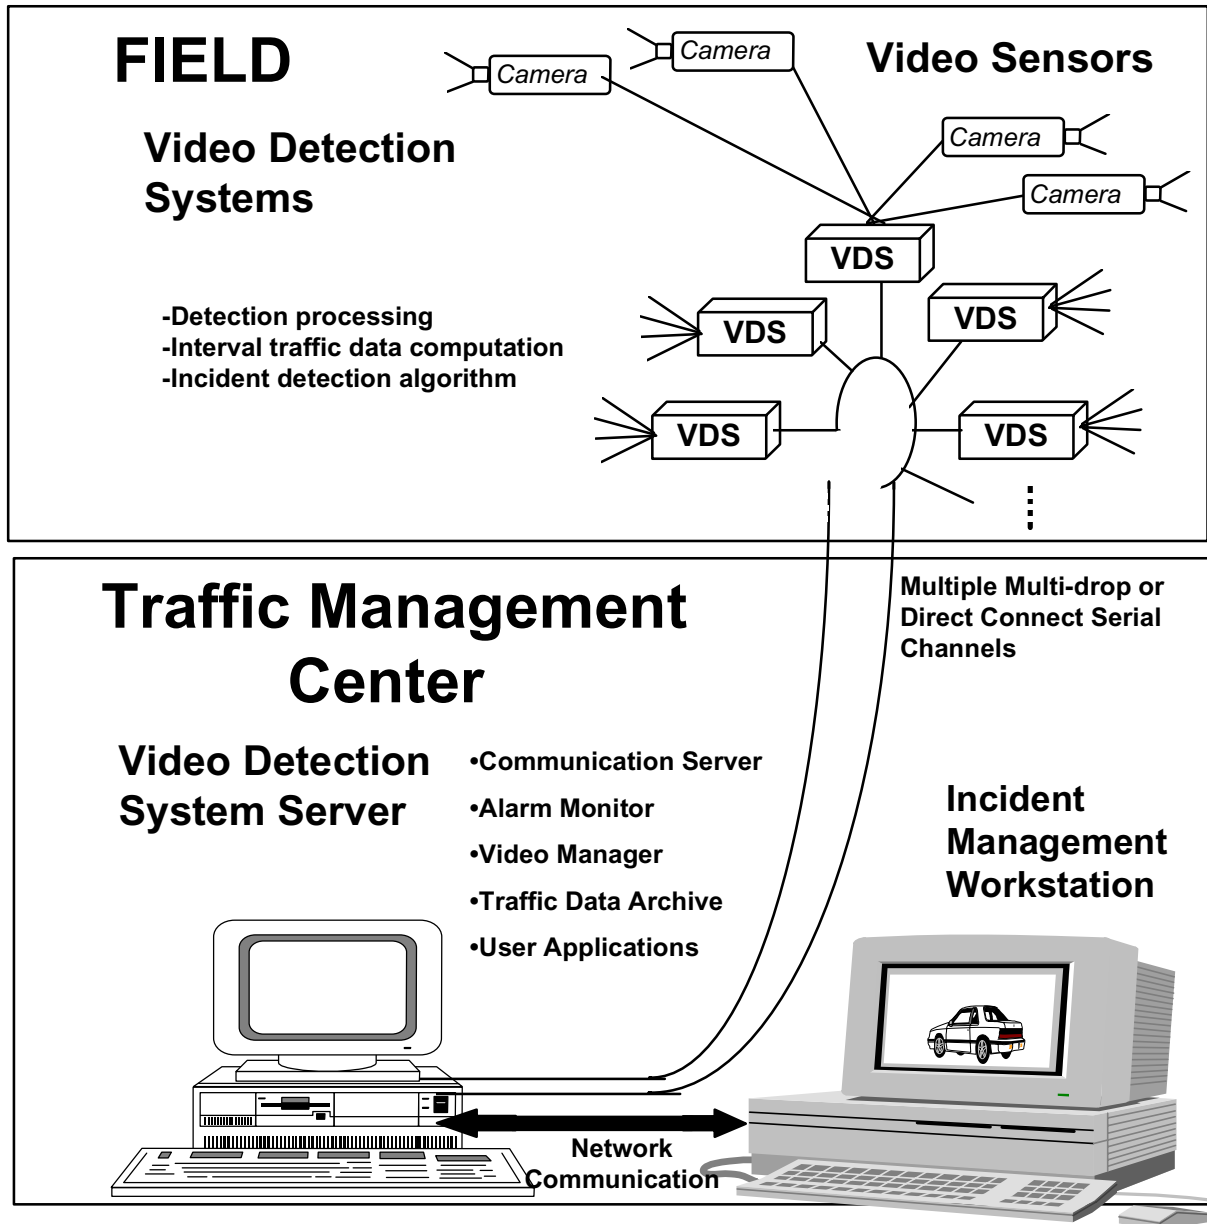


Figure 1. AIDA-Based Video Traffic Management Station Configuration

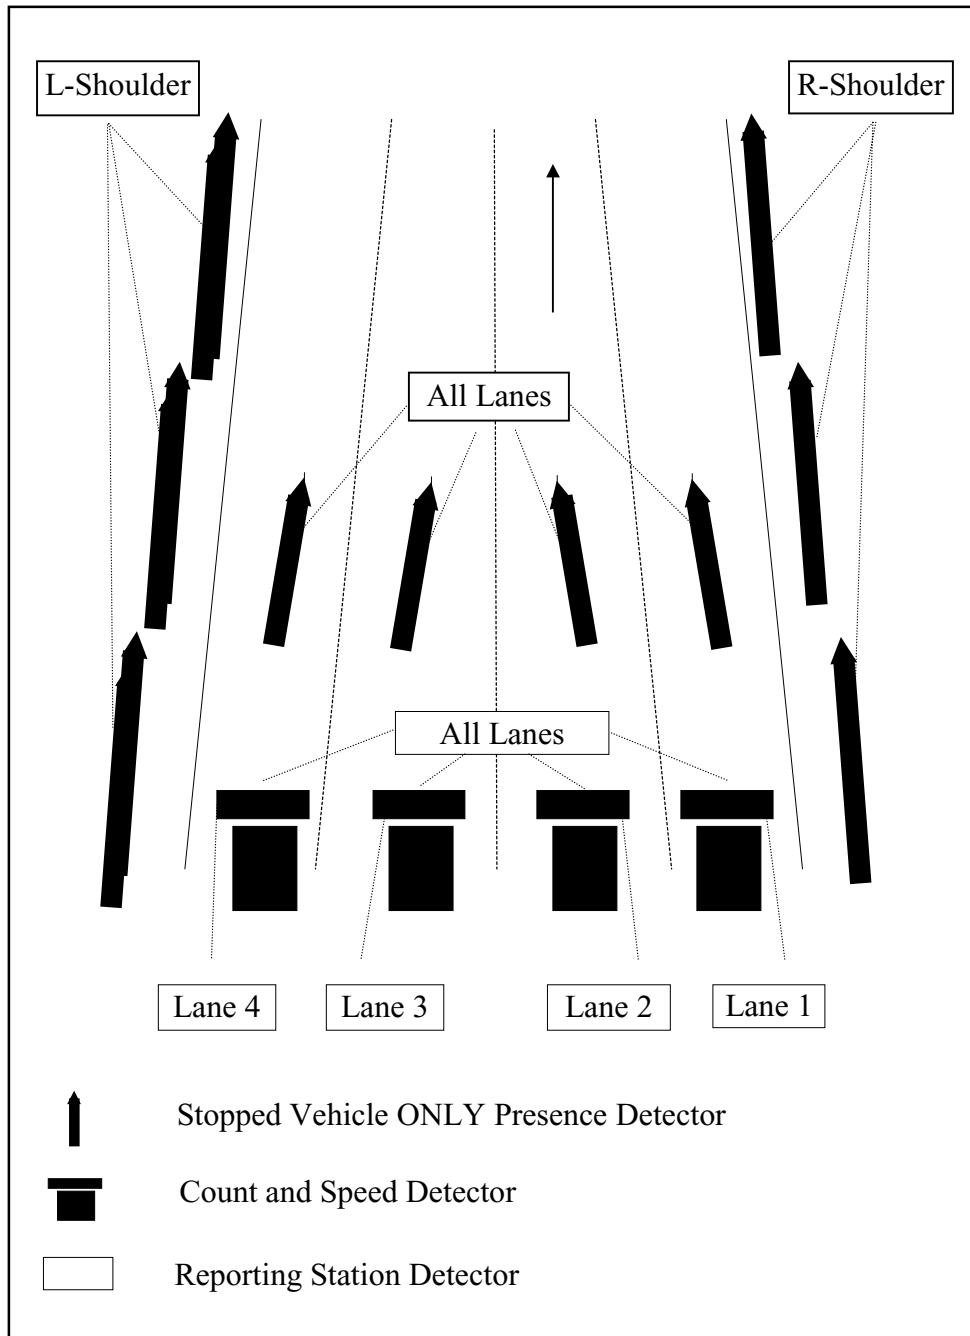


Figure 2. AIDA, Incident Detection Detector Configuration

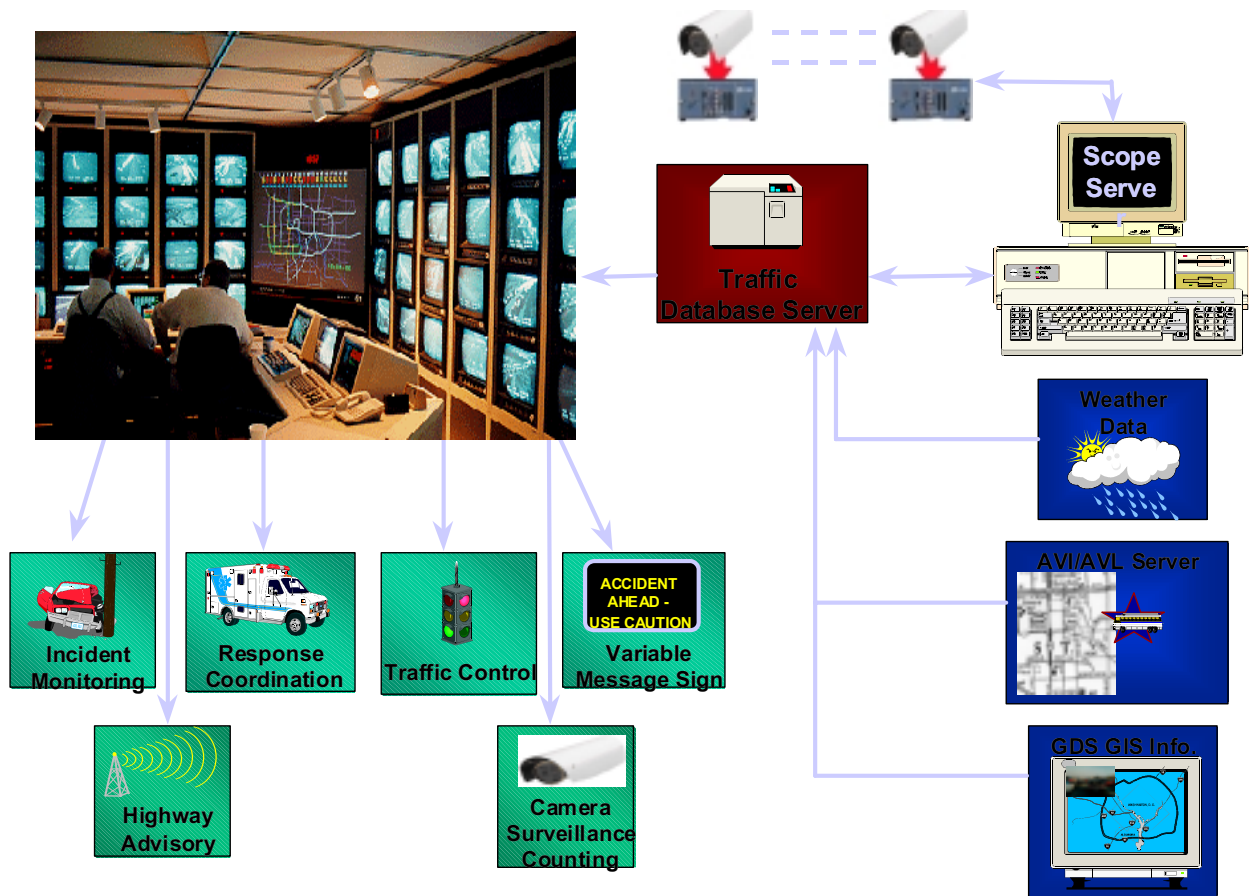


Figure 3. Integration of ScopeServer in a TMC Environment

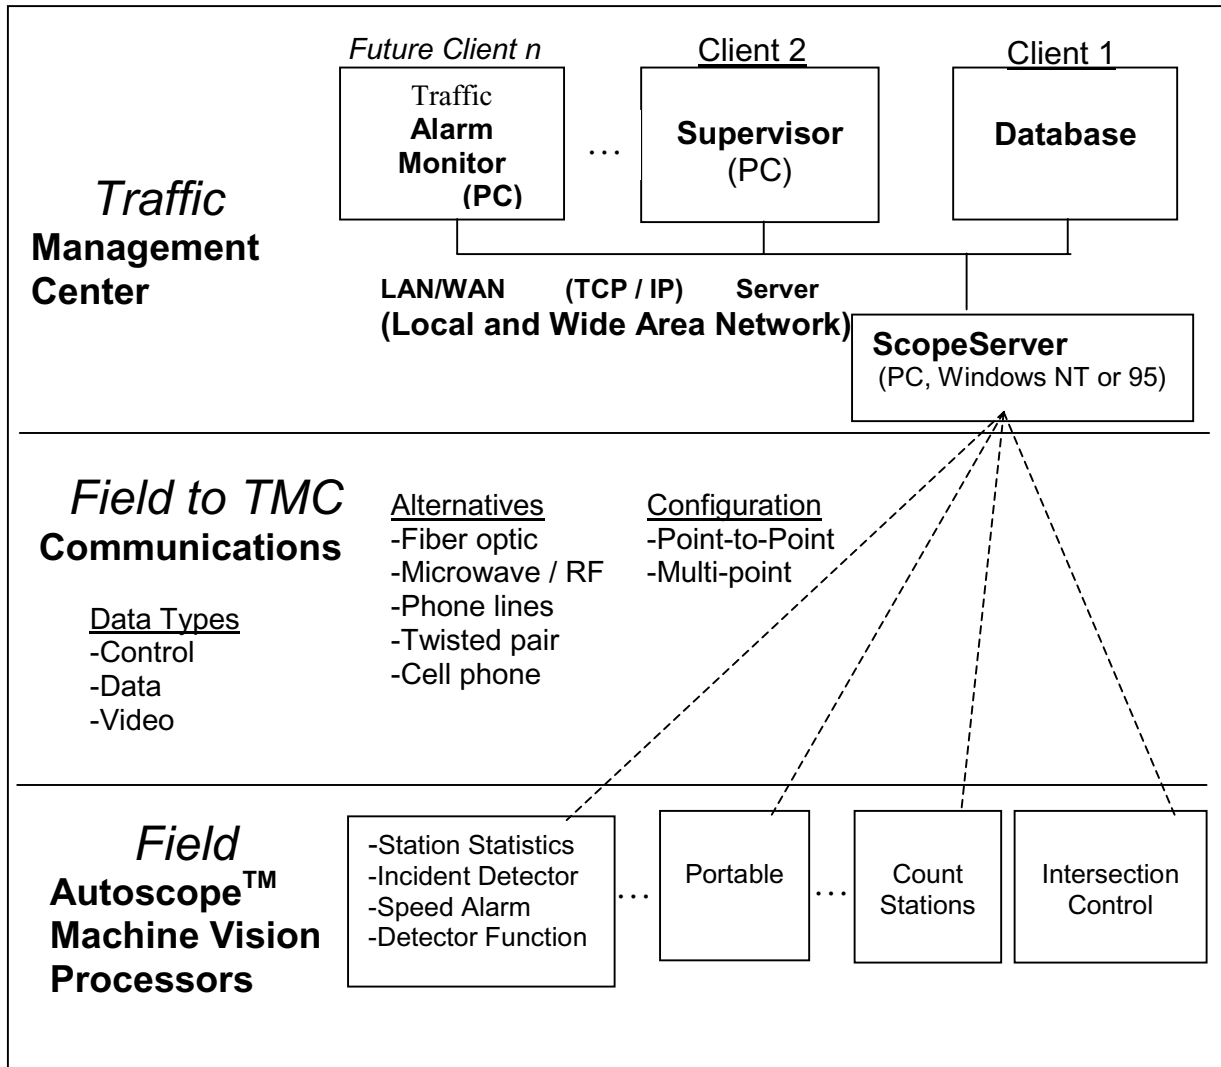


Figure 4. Client-Server System Architecture for Advance Traffic Management

Figures:

Figure 1. AIDA-Based Video Traffic Management Station Configuration

Figure 2. AIDA, Incident Detection Detector Configuration

Figure 3. Integration of ScopeServer in a TMC Environment

Figure 4. Client-Server System Architecture for Advanced Traffic Management