Video Vehicle Detector Verification System (V^{2}DVS) Operators Manual and Project Final Report

Presented to Caltrans Division of Research and Innovation

John Slonaker, Caltrans Detector Testbed Manager, Caltrans Irvine Detector Testbed

By Art MacCarley, Ph.D., PE. Principal Investigator
Prof. of Electrical and Computer Engineering
California Polytechnic State University, San Luis Obispo

Task Order 5327, Cal Poly Corporation Project No. 49498

Revision 7, October 24, 2011

Table of Contents

1 The Purpose of this Document

This document serves as the V2DVS system user manual as required by the project contract, and as a comprehensive report on the project, which is provided even though it was exempted as a separate deliverable for the project in the original proposed scope of work.

As a system user manual, this document describes the setup and application of the V2DVS for detector testing in the Caltrans I-405 Detector Testbed. In the context of a final project report, it provides all relevant design and implementation details of the V2DVS.   This includes the hardware and software design details, data acquisition architecture, data fusion methodology, computer vision detection methods, and automated data fusion/reduction and reporting methods. The titles of each of the sections indicate its relative value as a "how-to-do" description for the user manual, or as a "how-does-it-work" or "what-we-did" section primarily intended for project final reporting purposes.

2 Disclaimer

The contents of this report are the sole responsibility of the authors. They do not necessarily reflect the official views or policies of the California Department of Transportation or the State of California. This report does not constitute a standard, specification or regulation.

3 Acknowledgements

This work was conducted at the California Polytechnic State University, San Luis Obispo, California, under contract to the California Department of Transportation Division of Innovation and Research. The Caltrans project monitor is **John Slonaker**, whose guidance and support have been critical to the successful performance of this project.

Except where noted, this report and operator's manual were written by the contract principal investigator, **C. Arthur MacCarley**, Professor of Engineering at the California Polytechnic State University, San Luis Obispo. Although the current project was funded through the Cal Poly Corporation, the original V2DVS was developed and deployed in the Caltrans Detector Testbed by Loragen Corporation of San Luis Obispo, CA by the same Principal Investigator. Computer programming for the most recent versions of all software was performed by research assistant **Erik Sandberg,** and previously **Douglas Brandt**, starting with code from prior V2DVS versions written by Computer Engineering students enrolled in Course ENGR 483-4 and originally developed by Loragen Employees **Neils Nesse and Tabber Watts.** Hardware and software design and architecture, and hardware fabrication tasks were performed by the principal investigator.

4 Motivation and Methods

The accurate detection of the presence, speed and/or length of vehicles on roadways is recognized as critical for effective roadway congestion management and safety. Vehicle presence sensors are commonly used for traffic volume measurement and control of signalized intersections and ramp meters. In addition, vehicle speed and classification from length are important for automated incident detection and the characterization and prediction of traffic demand. Real-time detection is also necessary for actuation of automated driver information systems, for example, the Caltrans Automated Warning System (CAWS) on I-5 in central California. The California Department of Transportation is currently in the process of expanding its detection infrastructure to provide improved traffic information to its Traffic Management Centers (TMCs).

The market has recently seen an increasing number of "out-of-pavement" detection systems that have an advantage over conventional inductive loop detectors because they do not require lane closures for installation and subsequent maintenance. These systems use various non-intrusive sensing methods, and there is evidence from pilot installations, however anecdotal, that these detectors have different detection characteristics and yield different results than inductive loop systems. Known problem situations include ambiguous lane positions during lane-changing, high-clearance vehicles and vehicles with trailers, and possible vehicle occlusion for video-based detectors. Detectors will occasionally fail to detect, falsely detect or multiply detect individual vehicles. Actual non-aggregated (individual vehicle) data is needed on how all types of detectors perform in different situations in order to optimally specify equipment for specific detection requirements and

environments.   There was previously no way to compile this data and understand these anomalies without manually reviewing individual detection event records, a labor-intensive process that is practically impossible for large datasets.

Among the commonly-used vehicle detection mechanisms are fluctuations in a magnetic field (inductive sensors), video image processing, microwave RADAR, optical transmission, laser transmission, reflection or pulse time-of-flight, and acoustic signature discrimination.   Each detection method has advantages and limitations that make them appropriate for some applications, but inappropriate for others. The need to assess the relative accuracy and attributes of each type of sensor against a set of common, objective standards motivated the California Department of Transportation to construct a sensor Testbed on Interstate 405, one of the highest-volume freeways in the state.

With these and other concerns in mind, Caltrans DRI (Division of Research and Innovation) established the ATMS (Advanced Traffic Management Systems) Detector Testbed in southern California on Route 405 near the City of Irvine. An assortment of out-of-pavement detection systems, along with duplex inductive loop detectors, are installed to detect each northbound vehicle passing through the Testbed section of the highway. A central feature of the Testbed is the **Video Vehicle Detector Verification System (V2DVS)** that automates the collection and accuracy assessment of test data from all detectors under test. A maximum of 32 detectors in each of six lanes can be tested concurrently by the V2DVS.

5   How to (and How Not To) Measure the Accuracy of a Vehicle Detector

The accuracy of vehicle presence detectors has been traditionally assessed using aggregated metrics such as cumulative vehicle count over a fixed interval. Field-aggregated data reporting has largely been driven by the computational and communications bandwidth limitations present at the time of implementation, many years ago, of much of the roadway management infrastructure, and has been considered adequate for basic traffic monitoring. But much more information can be gleaned from detector data if individual vehicle records are recorded and processed, both in real time and in post-processing. Field-aggregated metrics such as vehicle count or average occupancy obscure the accuracy of the detector irrecoverably. For example, failures to detect and false detections are self-canceling errors in a total count or volume measurement.

To illustrate further, a recurring situation we have observed in detector installations on California highways is lane connection error: it is not uncommon that the lanes for detectors are mis-ordered due to incorrect wiring of inductive loops in a junction box or on the backplane of the 222 card file. Since adjacent lanes typically have similar traffic characteristics, everything appears to be working perfectly and field-aggregated metrics will report counts, mean speeds and occupancies that appear to be correct, but are actually 100% incorrect since the measurements are from a detector in a different lane. Any detector can be honestly advertised to be 100% accurate because it "does not miss any vehicles" if it continuously generates detections regardless of the actual presence of vehicles. And any detector can be honestly advertised to never false detect vehicles if it is turned off.

Current computational and communications technologies allow the comprehensive recording of individual detection events and the evaluation of detectors based upon fully disaggregated performance metrics. Individual vehicle measurements compared with the ground truth record reveal the true performance of the detector.

Three binary outcomes are possible for any detection event:

1. Correct detection: Vehicle was properly detected
2. Failure to detect: An actual vehicle was not detected by the detector
3. False detection: A detection was reported but no vehicle was actually present

For detectors that report vehicle speed and/or length (for classification), two analog data may be reported for each vehicle that is correctly detected (a subset of the above three outcomes):

4. Vehicle speed (mph or kph)
5. Vehicle length (ft or m)

For presence detection, the properly aggregated complete assessment set must include (percentages relative to the total number of actual vehicles):

1. Total number or percentage of vehicles that were correctly detected (correction detections).
2. Total number or percentage of vehicles that were not detected (failures to detect).
3. Total number or percentage of detections that were not actual vehicles (false detections).

Any statement of performance or specifications not containing all three of these numbers can be potentially misleading.   It is the unfortunate case that most detector vendors do, in fact, provide only field-aggregated metrics such as cumulative count, or only a subset of the metrics above in published product information (January 2009 and August 2011).

While some detector technologies (e.g., Phased-array RADAR and video) are capable of reporting vehicle speeds and/or lengths in a single detection zone, most detectors are primarily designed to provide only vehicle presence detection (e.g., inductive loops, magnetometers, ultrasonic, acoustic, passive and active infrared, laser reflective, some RADAR, some video detectors). When vehicle speed or length measurements are needed, these detectors are arranged as *duplex pairs*, separated by a known distance (typically 20 feet) in each lane. (For identification purposes, single detector installations will be referred to as *simplex* installations.) For detector installations that report vehicle speed and/or length (for classification), the most appropriate aggregated performance measurement metrics are:

4. Average absolute (or RMS) speed measurement error, over a given period of observation of individual vehicle speeds (mph or kph).
5. Average absolute (or RMS) length measurement error, over a given period of observation of individual vehicle lengths (ft or m).

As noted previously, proper testing non-field-aggregated data and metrics require the generation of an accurate ground-truth record accounting for each of these measurements on an individual vehicle basis. Ground truth has traditionally been determined either by the use of a trusted reference detector, or by human observation, by use of slow-motion observation of time-coded recorded video of the traffic. The former approach is limited to the accuracy of the single trusted detector, aware that no detector is completely accurate under all conditions. And for human observations attempting to observe multiple lanes of high-speed traffic, the volume and precision required of the observations can become unwieldy, defying human abilities and accuracy limits.

The V2DVS system generates an accurate ground truth record using an adaptive consensus filter, best described as a weighted consensus of all detector results with the weighting factors determined by the cumulative observations of the past accuracy of each detector. The result is referred to as *inferred ground truth* (IGT), to be explained below.

6   V2DVS – An Automated and Adaptive Detector Assessment System

This document describes the operation, design, installation and maintenance of the Video Vehicle Detector Verification System (V2DVS) developed by the author and research assistants at Cal Poly San Luis Obispo, and installed in the Caltrans I-405 Detector Testbed in Irvine California. The V2DVS system performs a number of functions that greatly reduce the human workload associated with detector performance assessments, especially for long test periods and large datasets: It records an event record, including a digital image acquired at the time of detection for every detector under test. It implements a computer vision system to provide an additional reference image and record of every vehicle while in a standard detection zone (detection zones differ between different detectors). It implements a redundant distributed database system for storage and processing of all data and images. In post processing, it uses the composite detection results from each detector under test in each lane, including the V2DVS video detection record, to create the inferred ground truth (IGT) reference dataset. The IGT typically contains a limited number of "unverifiable" detections – cases that cannot be resolved by the adaptive weighted consensus algorithm – which may be manually resolved by a human operator using provided graphical software tools. Once all unverified cases have been resolved, the IGT may be considered the most accurate possible indication of *actual ground truth* (AGT) possible. The results of each detector are then compared against the verified IGT dataset to assess it's performance with respect to the metrics stated above.

Depending upon how conservative certain parameters are set, the ability of V2DVS to extract IGT for vehicle presence, speed and length is typically 95+% comprehensive, leaving only approximately 5% or fewer of the vehicles to require manual verification. A V2DVS PC application communicates with the server database and provides an easy-to-use graphical software tool to facilitate the final verification of questionable detection cases offline. This includes a graphical tool for rapid human resolution of reported detections for which a clear consensus does not exist. It then compares the results of individual detectors to the ground truth record to generate an overall report of the performance of all detectors under test. It also allows secure Internet-based remote control of all system and video camera functions and settings.

The automated data fusion and IGT generation capability of V2DVS greatly reduces the workload associated with accurate detector testing, especially for large numbers of vehicles or long periods of observation. Presence detection errors requiring human judgment calls are mostly commonly due to ambiguous vehicle lane position, image artifacts, magnetic or inductive cross-lane artifacts, or RADAR reflections, depending on the detection phenomena. Speed or length errors are most commonly due to the intrinsic sensing mechanism used by the detector under test, e.g., ambiguous magnetic or inductive length compared with actual vehicle length.

Figure 6.1 Over-lane video cameras and one of two roadside detector masts.

7 Components, Features and Functions of the V2DVS

The Video Vehicle Detector Verification System is physically comprised of a cluster of data acquisition computers (field machines), one per lane, at each test site, and a central server for archiving and automated processing of data. These are housed in three Caltrans Type 334C field cabinets, numbered one through three left to right, as shown in Figure 7.1. In addition, a workstation (client) program which can run on any MS Windows PC facilitates remote monitoring and control of all field machines, manual ground truth verification, generation of test results interacting with the central server. All field machines are 2-U industrial-rack mount Linux/PC platforms, each interfaced to a video camera located on an overcrossing above an assigned lane. The collected raw data consists of JPEG-compressed images and a database for the time of arrival, speed, and other metrics of every detected vehicle in each lane, and includes a reference record created by the V2DVS system itself based on real-time image analysis. The system supports multiple test sites, with a maximum of eight hardwired detectors and an unlimited number of network or serial signaling detectors for each lane at each site. At maximum traffic capacity, as many as 96,000 records per hour per site are generated.

Figure 6.1 shows the placement of the reference video cameras above each lane, and one of two poles on which wide-area detectors are mounted. Each of the six traffic lanes below each camera are equipped with duplex inductive loops and provision for other interchangeable roadway detectors. The V2DVS field machines are housed in a roadside Caltrans Type 334C cabinet shown in Figure 7.1.

Video cameras are positioned above each lane with a down-looking view, canted out approximately 10 degrees from vertical, and are positioned to view departing traffic. Camera height is 10.7 meters above the roadway for every lane. The field of view for each overcrossing-mounted camera currently extends from immediately below

the overcrossing deck to approximately 25 meters down-road as required to contain most of the zones used by different detectors (exceptions are video-based detection systems which use far-field approaching or departing traffic views). This view can be adjusted remotely, but should be standardized for all lanes. Wider fields of view yield relatively smaller vehicles within the images, which in turn reduce the measurement accuracy for the vehicle position at the time of detection. Images are color, with a configurable JPEG compression ratio. Default compression settings yield images of size 640x240 pixels and between 15 and 30 kbytes.

Each detector under test must *signal* the V2DVS system for each vehicle. Signaling may be *real-time,* in which a record is acquired immediately, or *delay-time*, in which a record is acquired prior to the time of signaling, to accommodate some detectors that do not produce real-time detection signals. The V2DVS computers maintain a circular buffer of the most recent video 300 images, acquired at 60 video fields (images) per second, to facilitate signaling delays of up to five seconds.

All records (image and data) are stored locally on each field machine, and automatically pulled via FTP on a bandwidth-available basis to a central server, which stores data objects in a MySQL database, and compressed images in indexed directories. Local storage capacity allows tests of up to twenty days in duration prior to the need to offload data to the server. Since the field sites are connected to the server via multiple networks, including a slow 802.11b wireless link, periods exceeding 24 hours may be required to push a few hours of traffic data. In the most recent version of the installation, the V2DVS server is located in Cabinet 3, immediately adjacent to Field Machine cabinet, avoiding network bottlenecks in data transfer.

Figure 7.1 V2DVS installed at the Caltrans/Cal Poly Detector Testbed on I-405 near Irvine California. Cabinet 1 contains the detector hardware interfaces including the Type 222 Cardfile. Cabinet 2 contains six V2DVS data acquisition Field Machines, one for each traffic lane. Cabinet 3 contains the V2DVS server.

Post-processing for IGT generation and manual verification uses a PC-based "Workstation" program, which automatically compares all detection times, speeds and lengths reported for each vehicle by the various detectors. This comparison is used to form a weight consensus for each detection event that establishes ground truth based upon the inputs of all the detectors under test, with variable degrees of confidence in the data reported by each detector. An adaptive algorithm continuously adjusts the degrees of confidence associated with each detector based upon agreement with the weighted consensus.

This is an embodiment of a *reputation system* (1), similar to algorithmic learning systems used for information credibility filtering by web search engines, social networking and in military decision-making. For both detector assessment and for image acquisition purposes, the V2DVS uses measured vehicle speeds to compensate for different detection zones used by each detector under test. If a detector reports speed, the speed measurement reported by the detector itself is used to calculate the equivalent pre-trigger time offset. If a detector does not report speed, a trusted source of vehicle speed data may be selected by the user.

7.1   V2DVS Features Summary

The V2DVS provides the following functions and features:

1. The V2DVS system monitors the output of all detectors under test in each lane, and records the exact time of arrival and a jpg image of each vehicle reported by each detector. For the most recent V2DVS version, detection signaling is expected to be hardware-equivalent to the output of a Caltrans Type-222 detector card. Different detection zones can be accommodated for each detector, by user input of the lane offset distance from the baseline position (directly under the downstream side of the Testbed overcrossing) as a settable parameter for the detector.

2. V2DVS also detects vehicles itself via its image processing system, acquiring an image of each vehicle when it is positioned in the baseline detection zone. The computer vision algorithm also measures and records vehicle speed to a precision of 1 mph, although it is not externally reported. Detection accuracy is limited by illumination conditions. The V2DVS video detection ability is treated by the system as just another detector, to aid in the accurate generation of IGT via weighted consensus with the other detectors under test.

3. Detection data, including an image and data entry for the detection event, are stored in a distributed database resident on the field detection computers. During or after data acquisition, the individual vehicle data and images are automatically uploaded to the master database on the V2DVS server, currently located in Cabinet 3 at the Testbed to minimize network delays due to the limited bandwidth of the Testbed's wireless network connection. Records consist of the arrival and departure detection edges at each of two real or synthetic "duplex" detectors. The time of detection, the vehicle speed, and vehicle length are derived from the times recorded for each of these edges in post-processing by the V2DVS workstation (PC) software.

4. Detector tests are initiated, controlled and terminated via either the user interface of any of the field machines in Cabinet 2, or remotely from a network-connected PC running the V2DVS Workstation application. Data accumulated during tests are automatically transferred from the field computers to the central server on a continuous basis. Data is retained locally on the field units until successfully transferred to the central database. A double redundancy transfer protocol is used to insure against data loss. A separate image (14-25K, average 20K) is acquired for each detector for each vehicle. Typical traffic volume is 6,000 vehicles per hour for the six lane segment of I-405 at the Sand Canyon overcrossing of the Testbed. This implies the generation of 120 Mbytes of data per detector per hour per site. A test period of two hours with four detectors in each of six lanes at one site with this traffic volume would yield 1.92 Gbytes of data transferred to the server. If the server is located remotely from the Testbed and connection is via the Internet, the large volume of data coupled with limited bandwidth of the Testbed external network connection will lead to a delay of up to several hours between the completion of a test period to full data availability on the server. For this reason, the server has recently been located permanently in Cabinet 3 at the Testbed, where is is connected to each field machine via a 802.11n local area network.

5. The PC-based Internet-connected workstation application provides remote monitoring and control of each field machine, and the automatic generation of comparative accuracy results for all detectors via interaction with the server database. The client application communicates with the central server, providing images of every car in a time-line-like presentation for all detectors in all lanes. A *ground truth* dataset is created by starting with an automatically-generated set and manually handing all cases that were not assigned by the automatic data reduction process. If the PC is connected to the Testbed LAN directly via a CAT-5 cable, the data transfer between the server and the PC will be nearly instantaneous.

6. Ground truth (actual vehicles, their times of detection, speed and length) are determined by considering the inputs of ALL detectors under test, using confidence factors to weight the relative belief in each detector's reported event. The confidence factor for each detector is constantly adapted over the entire data set based the consistency of its reported detections with the IGT (ground truth) data set. Stated simplistically, highly accurate detectors are progressively assigned greater confidence (belief) factors in the determination of the IGT for each detection event. Failures to detect or false detections reduce the confidence factor for the detector. Correct detections increase the factor. As long as the initial values for the confidence factors do not weight low-accuracy detectors excessively, the ground truth data set will asymptotically optimize, becoming more accurate with time (additional detection events), and generally more accurate as the number of detectors under test is increased.

7. During the automated data reduction process, V2DVS compensates for detection positions outside of the baseline detection zone. This requires vehicle speed information, provided either by the detector itself, or generated by any detector under test selected by the user, e.g., by the duplex inductive loops at the site.

8. Unclear detection cases in the preliminary IGT dataset, in which a vehicle is detected outside of an offset-adjusted admissible arrival position window are left for manual user verification. Following completion of manual ground truth verification, detection accuracy results are calculated relative to the corrected ground truth results. Human correction of the ground truth record is implemented via the V2DVS workstation PC application.

9. Failures to detect as well as false detections are reported for each detector in the Workstation application for support of manual ground truth verification by a user. Each error situation is accompanied by the corresponding image set from all detectors (for failures to detect) or the image acquired for the time of detection reported by the detector (for false detections).

10. Similarly, for detectors capable of reported vehicle speed and length, ground truth values for each of these measurements are generated using a similar adaptive weighted mean of the values reported by each of the detectors for which the vehicle presence was reported.

11. In addition, the server also provides a mechanism (SFTP) to directly download user-selected periods of raw data (JPEG image files and associated text records). This enables detector vendor participants to perform their own analysis on the raw data.

12. The V2DVS workstation application generates several types of comprehensive detector accuracy reports based on comparison of the results of each detector with the verified IGT dataset. These may be viewed from the application, or exported as HTML files for printing or further analysis purposes.

The V2DVS acquires an image and writes a log entry in a database when a detection "event" is reported by any detector under test. We refer to this request as "signaling" the V2DVS system. With multiple detectors under test concurrently, several images and records are acquired for each vehicle detected in each lane. The signaling from each detector may have an offset of up to 5 seconds. V2DVS can perform pre-trigger image recording for detectors with zones outside the field of view of the camera. This capability relies on the speed calculation made by the detector itself or by a user-selected alternative speed detector.

For detection zones outside the field of view of the video camera in a given lane, V2DVS can adjust the time of image acquisition forward or backward in time to align the vehicle with the baseline detection position at the time of image acquisition. This assures that the captured image associated with a detection event shows the actual

vehicle, even if the detection occurs at some distance from the camera's view. V2DVS Field Machine automatically maintains a running queue of 300 video fields (5 seconds), and has the ability to select the appropriate image from this queue to show the vehicle as it was (or will be) passing through the camera field of view.

7.2   V2DVS Components Summary

The Caltrans Video Vehicle Detector Verification System (V2DVS) is comprised of the following components, physically located in a roadside area on the east side or I-405 just north of the north-bound Sand Canyon / Shady Canyon overcrossing.

Hardware:

Six down-looking video cameras with computer-controlled lenses positioned on the Sand Canyon overcrossing above each of the six lanes of northbound lanes of I-405 at the test site.

Six "Field Machines". Each is a 2-U rack-mount Linux PC installed in Type 334C Cabinet 2 at the test site. Each is associated with a single lane and camera, and accepts trigger inputs from all detectors in that lane.

Six detector interface "boxes". Also installed in Cabinet 2 at the test site. Each connects to a maximum of 32 detector signals, conforming to Caltrans Standard Type 222 electrical interface (Section 8, TEES). Each box communicates with its corresponding a field machine via the parallel port. Across all six lanes at the test site, a total of 192 detectors can be tested concurrently.

One "Database Server", sometimes referred to as the "Central Server". Data acquired by the field machines is passed to this central server over a close-coupled LAN. The V2DVS Server is currently located in Cabinet 3 at the test site, but can be located remotely if an Internet Connection with adequate bandwidth to the test site becomes available.

Software:

Running on each Field Machine:

V2DVS – (Linux) detection program. This program starts automatically in the background when the field machine is turned on. A subcomponent of V2DVS running on each field machine is V2View, an image processing program used for local viewing of real-time camera images, and as a supplemental presence detector required when only one detector is under test, since the adaptive ground-truth algorithm requires that at least two detectors be concurrently under test in each lane.

Running on the Central Server

V2Collector – (Linux) starts all processes for the server to start collecting detection data from the field machines. This application starts automatically when the server is booted, so that normally no direct interaction with the server is necessary.

Client applications on PC

V2DVS Workstation – (MS Windows) allows remote control of data collection, activation and deactivation of data collection, configuration of come parameters of the field machines, acquisition and automated data reduction culminating in formatted reports of results. All functions require communications with the server, either locally at the test site by a CAT5 connection to the LAN between the Testbed cabinets, or over the Internet via a wireless Internet connection to the site which is currently unavailable but was previously routed through a UC Irvine VPN.

8 Field Machine Interfaces

8.1 Detector Interface

This interface accepts up to 32 wired inputs from external detectors for signaling detection events. An external box interfaces the PC parallel port to contact closure or active-low digital inputs. The box requires external DC power in the range of 6 to 24 VDC, with 12 VDC preferred. The box recognizes the edges of an external switch. This means the box will detect when any of the 32 input signals changes from VDC to ground and from ground back to VDC. This electrical transition causes the box to generate a pulse on the parallel port interrupt pin. The V2DVS service running on the PC recognizes this pulse, causing an interrupt which immediately captures the current state of all 32 input lines. The state of these lines is compared to a stored value representing the previous state of all the lines. If any line has changed state a detection event is triggered for that line. After processing the detection event to determine the line's corresponding detector number, a log file entry is generated for that detector including the edge detected on the line. Depending on the edge value, an image is also stored at the moment the signal is received. This is discussed further in Section 16.1.

8.2 Video Interface

The V2DVS service continuously monitors the video input from the over-lane camera and processes the images as previously discussed. The physical interface is a Model PXC200AL PCI video frame grabber manufactured by Cyberoptics of Seattle. This provides full-field rate (60 files per second) analysis and acquisition capability to support both the V2DVS detection function, and the acquisition of images for each external detector.

8.3 Lens Control interface

The V2DVS application controls the three functions of the motorized lens mounted on each video camera. Control of the iris by the image processing application is critical to permit reliable detection and standardization of the images acquired for each external detector. The focus and zoom functions are also controllable from the computer, although these are generally set once and not altered after initial setup.

The hardware used to perform this function is a 68HC12 microcontroller board connected to a bipolar driver circuit for driving the three lens motors in either direction. These circuits are all installed inside the PC case, only exposing an RJ45 (CAT5) connector connecting the lens control cable (to the cameras on the bridge deck).

The V2DVS GUI provides both automatic and manual lens control functions, and sends the signal to the microcontroller using the PC serial port. Each serial command generates a 56ms-long pulse for actuation of the lens motors. Under manual control, these pulses are repeated 10 times per second if selection the button is held down. Under automatic iris control, pulses repeat at a maximum rate of once every two second, to provide slow changes to the iris setting that will not interfere with the backgrounding function of the image processing. Approximately 64 pulses are required to move the iris from fully closed to fully open, yielding an end-to-end slew rate of 32 seconds.

Supplemental information on the programming of the 68HC11 microcontroller board with the lens control code is provided in the Appendix.

9 How to Operate the V2DVS

Assuming that all detectors are physically connected in Cabinet 1 and powered, three operations basic are necessary to activate the V2DVS system prior to detector testing:

1. Activate the video cameras on the bridge deck by turning on the power switch located at the bottom of Cabinet 2. Check for the proper field of view if the camera positions have been disturbed since the last use.

2. Start V2DVS by pushing the power switch on each of the six field machines in Cabinet 2. If the cameras have been disturbed or the detection windows have been changed since the last use of the system, it

will be necessary to re-define the detection window for each of the V2DVS field machines via the user interface. Once the detection windows have been defined, this step need not be repeated.

3. Start the V2DVS server in Cabinet 3 by pushing the power button. The server boots up automatically and does not require any user action.

Each task and its precursors are described in further detail below.

9.1   Camera Setup

Cameras used by V2DVS for video verification of each detection are positioned above each traffic lane on the side of the overcrossing deck. Cameras are permanently mounted in sealed environmental enclosures, and normally do not need to be adjusted. If, however, the position of the camera has been disturbed, or the settings of remote-controlled zoom or focus have been changed, it may be necessary to reposition or adjust the cameras so the field of view is correct and properly focused. The normal field of view requires that the cameras be tilted approximately 10 degrees from the vertical position, away from the bridge, and that the zoom setting is almost the widest angle position (3.2 mm focal length). This positioning and focal length assure maximum opportunity to acquire images detected by the various detectors under test with widely dispersed detection positions.

The camera angles may be adjusted and the cameras may be serviced safely by rotating the camera mounting arm inward over the sidewalk of the bridge deck, and then rotating them back to their proper position above the traffic lanes. If the camera mounts are rotated, it is necessary to re-align the cameras by observing the camera image and adjusting by hand until the camera is properly positioned. The detection window should also be reset, as will be described below. The position of the camera is important to the proper acquisition of vehicle images – both the images for the detectors under test and for the reference image acquisition by V2DVS itself.

Attached to each camera are three cables: a black coaxial cable for the video connection, an orange round power cable, and gray lens control cable. These are all routed inside the frame tubing of the permanent camera mounts down to a junction box, shown in Figure

Figure 9.1 Camera control electrical box on bridge deck (one of six).

9.1. If the circuit breaker switch in the box is off (in its down position), turn it on by flipping it to the up position. Normally, it should stay in the up (on) position, since power to the boxes is enabled at the controller cabinet rather than in each individual camera electrical box.

If necessary to confirm proper camera positioning, focus, and focal length, an NTSC monitor can be connected to the Coax "T" located in the electrical box at each camera. Power for the monitor may be accessed via the power outlet in the box, although it will be necessary to insert a multiple outlet adapter so that both the power brick and the monitor may receive power. From the monitor, you can verify the following steps and make adjustments to the camera if necessary. To adjust the camera, swing the camera arm back inside the railing, make the needed adjustment, and swing it back out to the correct position over the traffic lane.

Align the camera mounting arms exactly parallel with the direction of travel of the corresponding lane. Make sure the vehicles travel perfectly horizontally across the camera field of view. Verify this using a local monitor or the V2DVS display in the cabinet. If the alignment and position of the camera is correct; the center of the lane should be horizontal and pass down the center of the image. Angled the camera out from the bridge deck just enough that a small sliver of the overcrossing pole structure or other detectors mounted on the V2DVS camera mounting arm are visible on the right side of the image.   Viewed from the bridge sidewalk, the cameras were previously shown in Figure 6.1.

It is important that the focus of the camera is correct. Adjust the focus if the lane lines are not crisp or if details in the road surface are blurred and not distinguishable. Assure that the focal field of view – the camera zoom should be at almost its widest angle position.

Figure 9.2 shows an example of the correct camera field of view, as displayed on the V2DVS program computer screen (same as the video monitor screen). Due to the very wide angle lens, there is significant distortion at the extremities of the view.
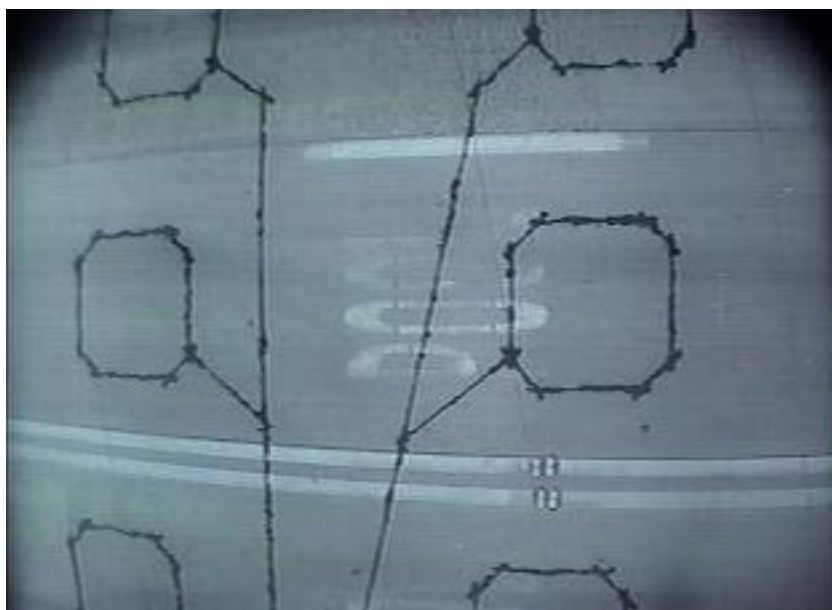

Figure 9.2. Example of correct field of view of V2DVS camera.

9.2   V2DVS Field Hardware Description and Setup

The core of the V2DVS system is the group of six detector data acquisition / image processing computers (referred to generically as "Field Machines") that are housed in Cabinet 2, as shown in Figure 9.4. Each computer handles the detector data acquisition and video image acquisition and processing for all detectors associated with each of the six traffic lanes at the Testbed site. Each cabinet also contains a computer monitor, keyboard, mouse, and an 8-channel (only 6 are used) KVM (Keyboard-Video-Mouse) switch for selection of each computer for setup purposes.

In the most recent version, the V2DVS Server has been permanently located in Cabinet 3, as shown in Figure 9.3. While the server can actually be located anywhere on the Internet, the lack of adequate bandwidth and intermittency of the wireless network link to the site practically prevents this. The communications between the Server and each of the Field Machines as well as the V2DVS Workstation client(s) requires a reliable network connection at DSL-quality bandwidth.

The interfaces between the Field Machines and up to 32 detectors under test in each lane are handled via "detector interface boxes" mounted on the back access door of Cabinet 2, as shown in Figure 9.5. Each accepts 32 inputs, corresponding to 32 detectors in each lane, and communicates with its corresponding Field Machine via the DB-25 parallel port connector. A telephone-type punch-down terminal block located below the

interface boxes handles the connection of the Caltrans Standard 222-compatible outputs from all 192 detectors and provides multiple ground points. The terminal block is labeled and color coded to match the lane and detector number for each connection as shown in Table 9.1.

Table 9.1 Lane Color Code

Lane	
1	Red
2	Violet
3	Blue
4	Yellow
5	White
6	Green

Figure 9.3 V2DVS Server Located in Cabinet 3 at Caltrans Detector Testbed.

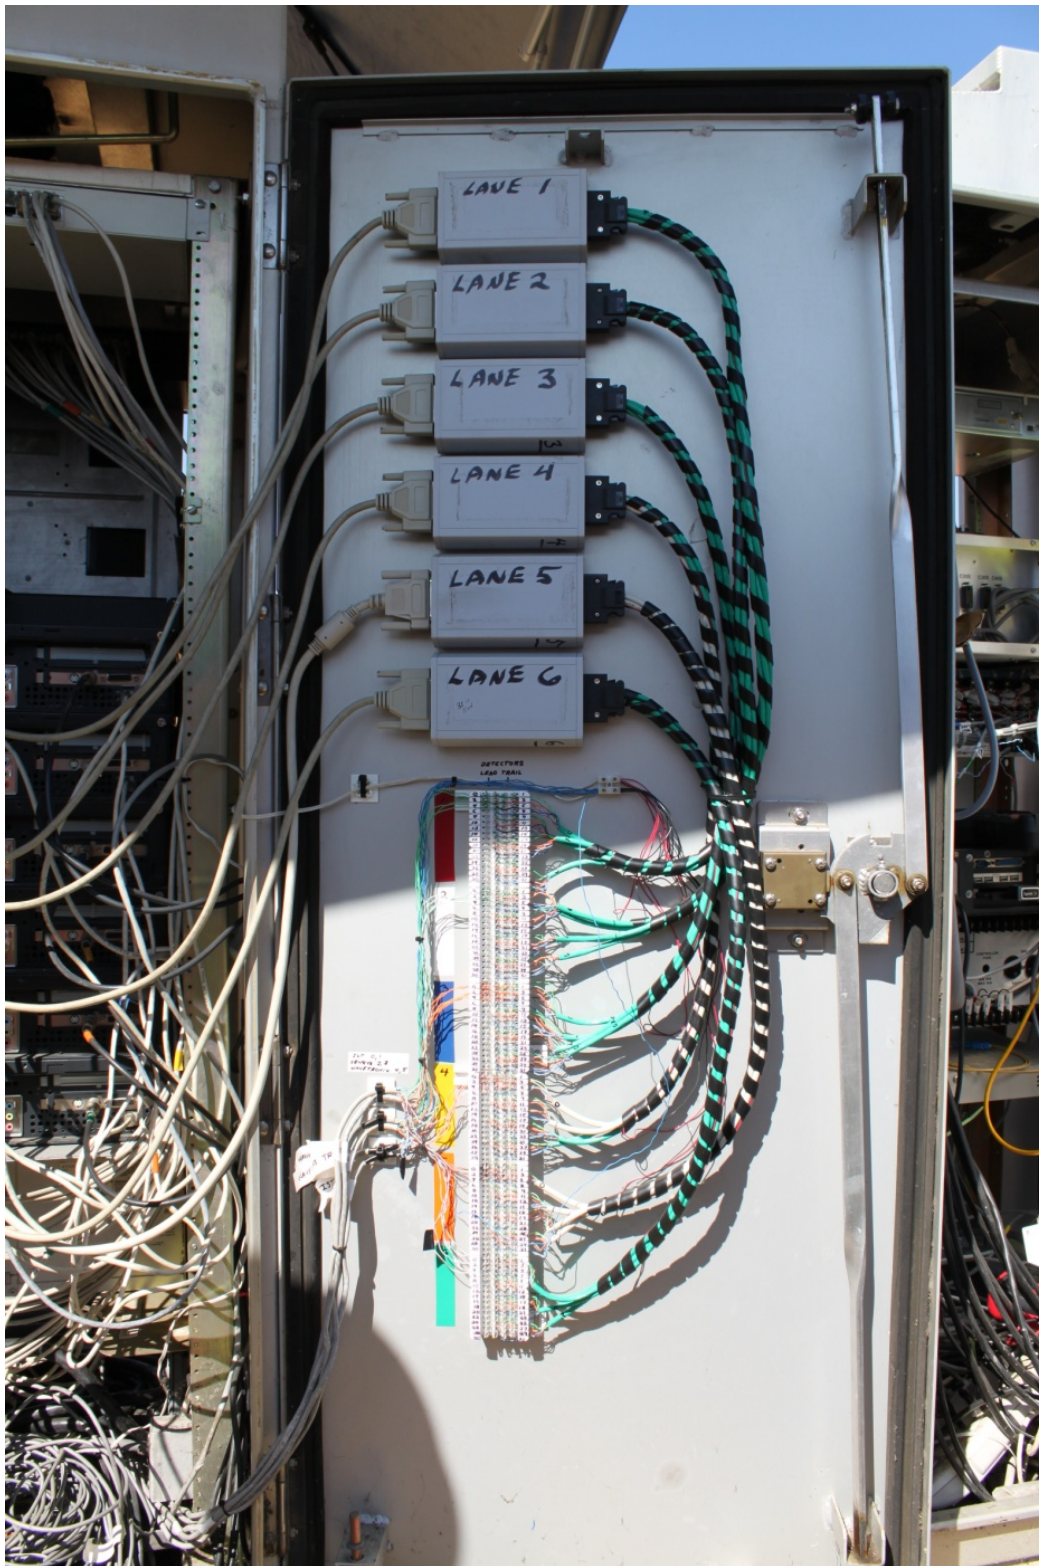**Figure 9.4 V2DVS Field Machines Located in Cabinet 2 at Caltrans Detector Testbed.**

Figure 9.5 V2DVS Detector connection punch-down block and interface boxes on back door of Cabinet 2 at Caltrans Detector Testbed.

Electrical connections from all of the detectors deployed in the Testbed are routed through conduits from the detectors into Cabinet 1. Most are connected to a Type 222 detector cardfile (rack-mount) backplane, shown in Figure 9.6. A multi-conductor twisted-pair cable running between Cabinet 1 and Cabinet 2 connects all of the detectors to the a large telephone-type punch-down board on the inside of the door in Cabinet two, which allows the connection of any detector to any one of 32 input channels for each lane. Each input terminal in the punch-down block is routed to one of the six detector interface boxes, which connect via a multiplexed parallel port cable to each of the six Field Machines. The arrangement is shown in Figure 9.6.
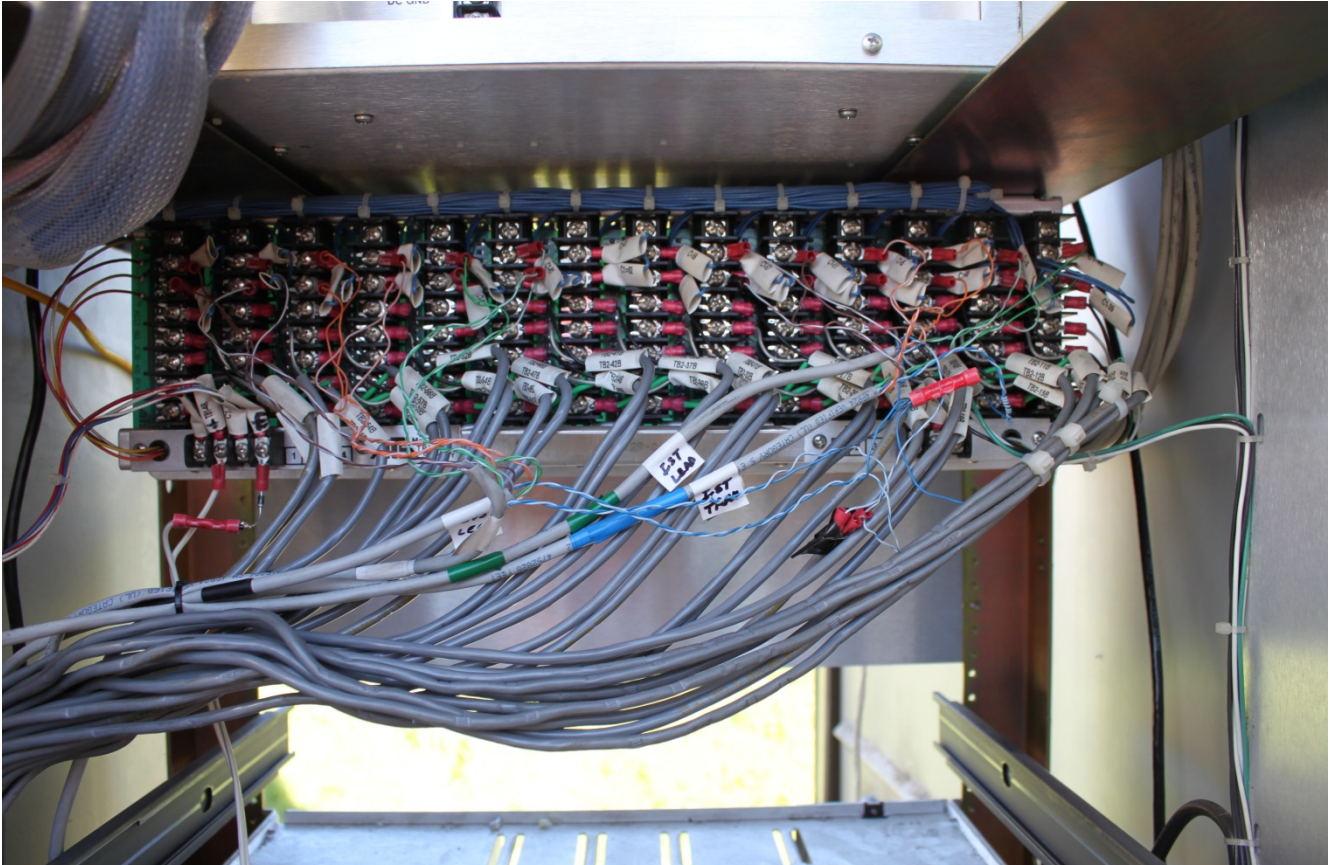

Figure 9.6 Back of Type 222 card file in Cabinet 3, showing connections from multiple detectors routed (to the left and down through a conduit) to the detector interface terminal block of Cabinet 2.

9.3   Operating the V2DVS

9.3.1   System Power

If the Detector Testbed (all three cabinets) is not already powered, turn on the main site power, by switching on circuit breakers 4, 5 and 6 in the main power cabinet. The main power cabinet is the small avocado-green cabinet located approximately 75 feet from the three Type 334 cabinets.

All components of the VDVS are powered through a 3kW UPS (Uninterruptible Power Supply) to protect the equipment from power surges and to assure continuity of test runs in the event of a temporary power outage. The UPS is capable of powering the system at maximum load for approximately 15 minutes. Normally, the UPS should be left on at all times. It should only be turned off if the main power to the cabinets is shut off.

If not already powered on, it is necessary to activate the UPS (Uninterruptible Power Supply) in Cabinet 2 by changing the position of DIP switch 4 on the back of each UPS, as depicted in Figure 9.7. A small precision screwdriver is needed for this. The set of 4 DIP switches is located in the middle area on the back of each UPS.

The DIP switches are very delicate and require gentle action with a small screwdriver only! At the front of the UPS, momentarily press the switch marked "I/O", to power up the outlets of the UPS. The �▢ LED will change from flashing to solid green to verify that it has been activated.
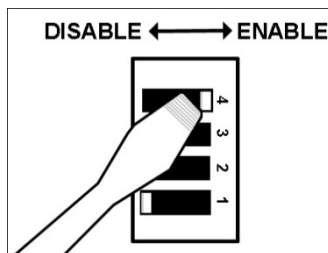
Figure 9.7 Enabling the UPS.

Activate all the V2DVS Field Machines (6 computers in Cabinet 2) and the V2DVS Server (in Cabinet 3) by momentarily pressing the yellow power button on the upper right corner of each computer. If not already on, turn on the power to the video cameras on the bridge deck, via the power switch located at the bottom right side of the back of cabinet 2, as shown in Figure 9.8. This switch can easily be accidentally turned off if one of the many cables in the cabinet brushes against it.

Figure 9.8 Power switch for the bridge deck, which enables the video cameras.

After it is powered on, no further action is required for the V2DVS server. It will boot automatically taking approximately three minutes to complete the process. No configuration is normally required.

For the Field Machines: After each is powered on, the complete Linux boot process will take approximately 90 seconds. Each Field Machine will automatically start V2DVS without any action required by the operator.
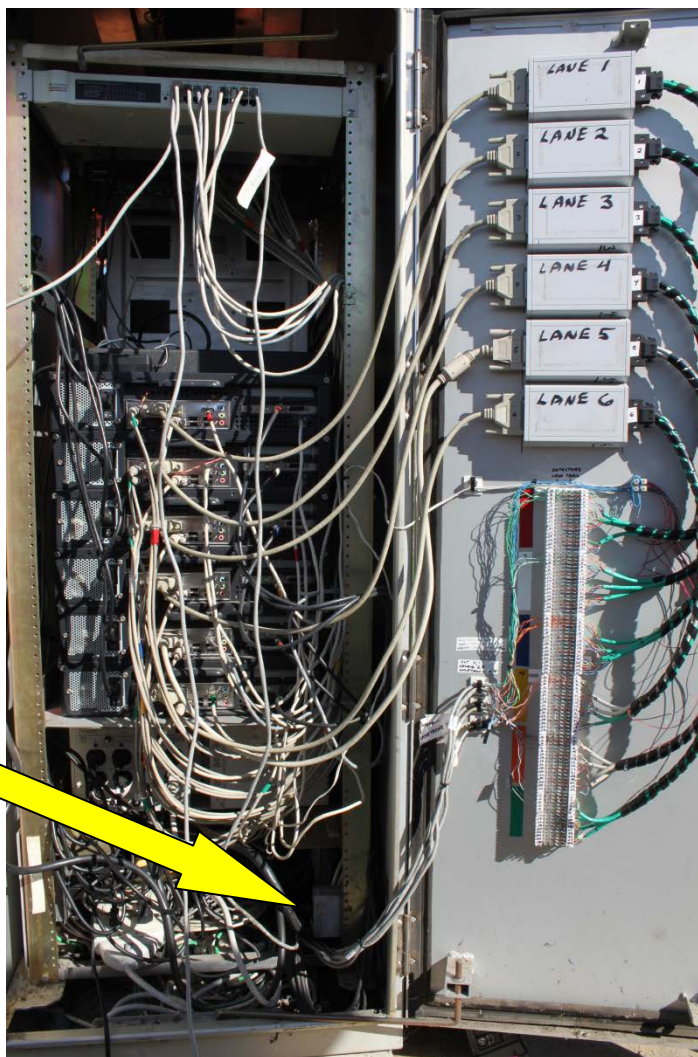
The monitor screen, keyboard and trackball mouse, located in Cabinet, connect to all of the six V2DVS field machines. The particular machine they are connected to can be selected by the KVM (Keyboard, Video and Mouse) switch, indicated in Figure 9.9. The KVM switch will indicate in an LED display the number of the Field Machine, which is the Lane number associated with it (Lane 1 is the HOV lane closest to the meridian, Lane 6 is the "slow lane", in this case the merge lane from the onramp).
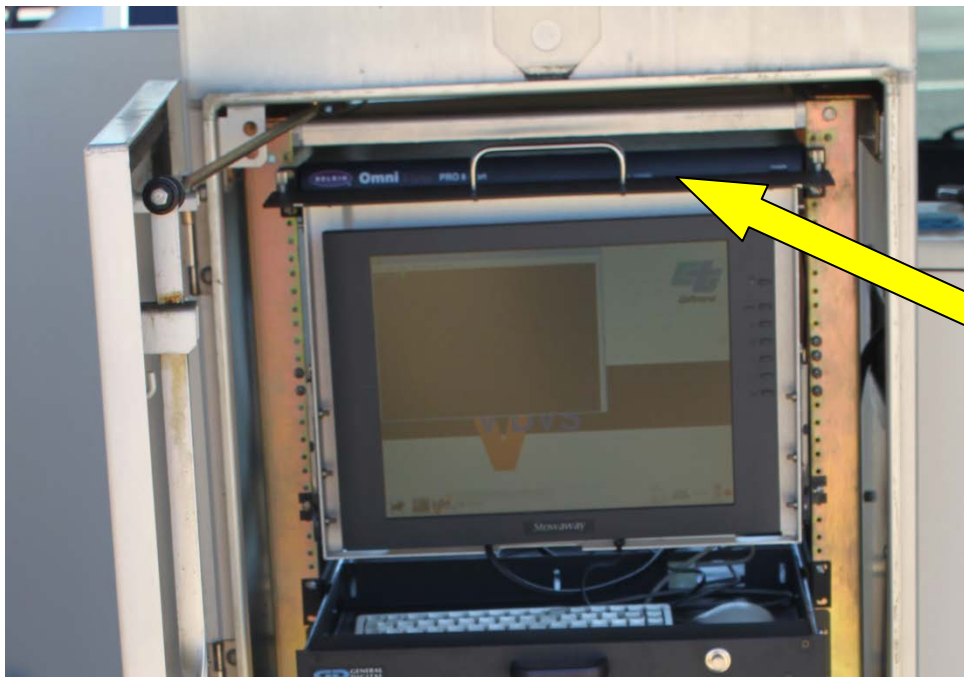
KVM switch for selection of each of six Field Machines to interface with monitor, keyboard and mouse. Note: this unit will be replaced by a newer but similar model prior to the final project demonstration.

Figure 9.9 Location of KVM switch behind flip-up video monitor in Cabinet 2.

9.3.2   Equipment cooling

Equipment cooling units are installed on the sides of Cabinets 2 and 3. If the ambient air temperature is above 70 degrees, or if the system is to be operated with the cabinet doors closed, *be sure that these air coolers are on*. These are powered by 230 VAC, from outlets at the bottom of each of the respective cabinets. The coolers are thermostatically controlled and need not be manually switched. If the cabinet doors are left fully open while operating the Field Machines or Server, the coolers need not be on.

If the V2DVS machines are left on unattended with the cabinet doors closed (locked), the coolers should be left on. Since they are thermostatically controlled, they the coolers will cycle as they are needed to maintain an acceptable operational temperature inside the cabinet. The temperature set point for each cooler is internally set and is not adjustable without removal of the cooler enclosure. There is a temperature readout display on the outside of the cooler, which reads in degrees F. The temperature of Cabinet 2 is the most critical, since it houses the six field machines and a 3kW UPS.

Figure 9.10 V2DVS default window display.

9.3.3   Control and Monitoring the V2DVS

Once selected by the KVM switch, the default monitor display for each of the V2DVS field machines is a default set of the "most-often-used" windows, as shown in Figure 9.10. These windows provide full system control except for some internal operating system functions which are not adjusted during normal operation. Note that all V2DVS Field Machines are fully operational regardless of whether they are selected for viewing via the KVM switch. The monitor can be turned off during data collection without any adverse effects.

The **System Status** window is the main program window from which all other windows may be launched to either set up or monitor V2DVS operation. The V2DVS main program is always running in the background, regardless of the windows displayed on the monitor. In particular, even the **System Status** window does not need to be displayed for V2DVS to be running. If it is closed (by clicking the upper right corner button) it can always be reopened by clicking on the V2DVS logo in the system tray at the bottom left of the screen, shown in Figure 9.11.

**Figure 9.11 Screen icon for activating and
deactivating the Status Window.**

The System Status window displays a number of real-time metrics of interest, including internal operating parameters of the V2DVS hardware such as available memory and the current detection status values such as the state of data collection (ON or OFF). Of particular interest in system diagnostics are the "Logs not collected or deleted" and the "Archives not yet uploaded" numerical displays. The field machines pass data to database on the server via data logs and image archives. These displayed values indicate the number of each that are waiting to be passed to the server. If the server is disconnected or not operating, these values will increase continuously. This will not inhibit data collection since the Field Machines are capable of storing several days of data before transferring it to the server, but it is an indication that there may be a problem with the server or the LAN connecting the field machines to the server.

Data collection can be started and stopped by clicking the **Start Collecting** button at the bottom of the menu. Data collection is started or stopped for all detectors in the given lane. It is not possible to collect data from only one detector, since the V2DVS adaptive consensus filter depends on inputs from multiple detectors to establish an accurate ground truth record.
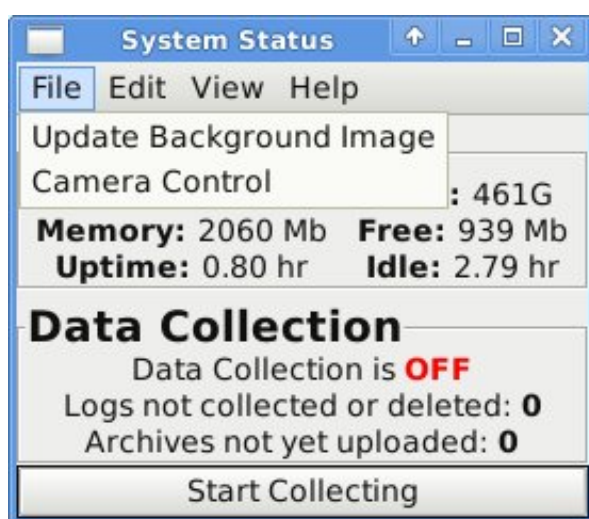
**Figure 9.12 V2DVS status window, showing
File drop-down menu.**

The **System Status** window has four drop-down menus: **File**, **Edit**, **View**, and H**elp.** Each will be described below.

9.3.3.1   File Menu

The **File** menu has two options, **Update Background Image** and **Camera Control.**

Update Background Image

If the video feed is interrupted at any time, the background (no-car) reference image used by the image processing system will eventually become non-current. While it will autocorrect over time, it is expedient to click **Update Background** to acquire a new base image. The V2DVS vehicle detector feature will not function correctly without a stable background image, which is expected to change only slowly. Note that detector testing does not require that the V2DVS vehicle detector be operational. Its output is treated like that of any other detector under test for purposes of adaptive ground truth generation. The V2DVS image processing display features are intended only as a convenience to aid in the setup and monitoring of the system during a data collection run.

Camera Control

Clicking **Camera Control** generates the window shown in Figure 9.13. If the user wishes to manually adjust the iris, or the focus or zoom positions of the lens, click the **Manual Lens Control** button. Six buttons below it will now be highlighted for manual control of each lens function. When manual adjustment is complete, be sure to click **Automatic Lens Control** to revert to V2DVS feedback-control of the lens iris during data collection, since the exposure level is critical to the correct function of the video detection algorithm and the quality of the images acquired upon each reported detection.

Figure 9.13 Camera control window, allowing user setup and control of robotic lens on each video camera.

9.3.3.2   Edit Menu

The **Edit** menu has only one option, **Detector Configuration**.
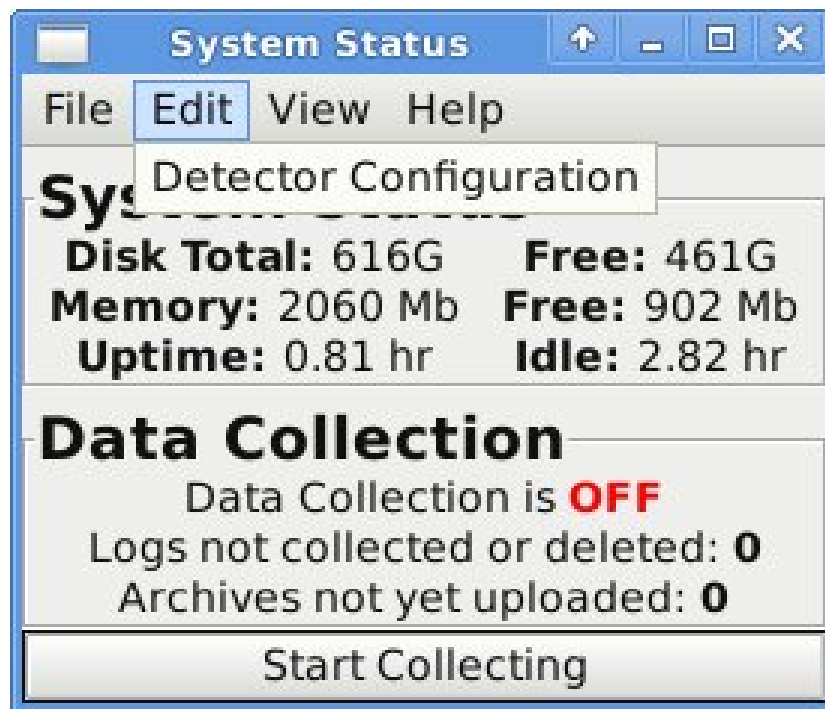
Figure 9.14 Edit menu selected for detector configuration.

Detector Configuration allows user adjustment of certain detector test parameters, in particular, the offset distance of each detection zone relative to the baseline detection position in the lane (immediately below the trailing edge of the overcrossing bridge). Clicking Detector Configuration opens the parameter text file "detectors.conf" for editing. Please note that data collection cannot be in progress if any detector parameter is changed using Detector Configuration, since this could change the fundamental conditions of the test run. If Detector Configuration is clicked while data collection is in progress, an error message will be displayed reminding the user to stop data collection before changing the detection zone offset distances.

This detector configuration file is internally named detectors.conf and is located in /etc/opt/v2dvs/. The file should only be edited only by accessing it by selecting **Detector Configuration** from **Edit** menu tab on the **System Status** window. The format of this file is shown in Figure 9.15 below, and the contents are described in the following three subsections. The values in the file are whitespace-delimited, meaning any amount of whitespace (such as spaces or tabs) between values will be ignored and is used to separate fields. Any text between the '#' character and the end of the line is ignored and may be used for comments.

```
                                    detectors.conf                          ↑  _  ⯐  ✕
File  Edit  Search  Options  Help
              Warning, you are using the root account, you may harm your system.

# Duplex Detectors:
[[duplex]] # Detectors between this line and "[[single]]" are duplex detectors.
# Detector  Lead Terminal  Trail Terminal  Zone Offset (in.)  Latency (ms)
[IST]
  1          0              1               0                  0
[Sensys]
  2          2              3               0                  0
[Wavetronix]
  3          4              5               -200               0

# Single Detectors:
[[single]] # Detectors between this line and "[[trust]]" are single detectors (no speed/length).
# Detector  Terminal  Zone Offset (in.)  Latency (ms)
# [example]
# 4         6         0                  -1000

# Trusted Detectors:
#  In order to determine when to take a picture, a speed must be calculated for
#    each detection. A detector (or set of detectors) must be selected as the
#    "trusted" detector(s). This "trust" is only extended to determining when to
#    take a picture and has NO EFFECT on the final rating of the detectors.
#  You can list multiple detectors here and give them different weights.
#    For example, if you trust both detectors 1 and 2, but you trust detector 1
#    more than detector 2, then you can give detector 1 a higher weight like this:
#       1 2
#       2 1
#  If no detectors are listed, the default is to use the built-in V2DVS detector
#    (detector 0).
[[trust]]
# Detector  Weight
  0          1    # built-in V2DVS detector using video processing on the camera
```

Figure 9.15   Detector configuration information text file format, used to link detector numbers to line numbers (specifying Single and Duplex detectors) and schedule pre-trigger image acquisition for detection zones outside camera field of view.

Relating physical detectors to detector numbers

There are three groups of text information in the detectors.conf file, designated by the headers [[duplex]], [[single]] and [[trust]].   Only the [[duplex]] and [[single]] groups are used to configure detectors for testing. The first is for configuring duplex detector pairs which can provide vehicle speed and length measurements, and the second is for configuring single detectors which provide time-of-arrival (TOA) information only. The V2DVS system contains a precise absolute time base, which is automatically synchronized with one of several national atomic clocks at midnight every night.

The first field in either the [[duplex]] or [[single]] groups is a user-defined detector number (Detector). It must have a value from 1 to 32 inclusive (detector number 0 is reserved for the V2DVS Video Detection System). The next field for single detectors (Terminal) and the next two fields for duplex detectors (Lead Terminal and Trail Terminal) represent the physical terminal block numbers associated with the detector output signals, as described above. The line numbers directly correspond to the physical lines connected from detection devices to the hardware interface box. The allowable numeric range is 0 to 31 inclusive. Note that the physical terminal

numbers are usually NOT the same as the assigned detector numbers. In fact, this is not possible for duplex detector pairs, which are treated as a single detector by V2DVS.

For example, two IST Inductive Loop detectors in each lane are physically connected to terminals 0 (lead) and 1 (trail) on the detector terminal block. But these are collectively treated as a Duplex Detector Number 1 for analysis purposes. Note that only the lead detector (lead loop) is used to report a detection of vehicle presence. The trail detector is used to provide time-of-flight for speed calculation, and a redundant measure of vehicle length based upon time over the loop (single vehicle occupancy).

The next two fields (Zone Offset and Latency) on the line are for adjustment of the detection zone offset for each detector, to allow acquisition of a usable image of the vehicle.

Adjusting Detector Image Acquisition Offsets

Many detectors have detection zones beyond the field of view of the V2DVS cameras, which extends only about 60 feet down-road from the baseline (first loop) position. To enable V2DVS to acquire images for these detectors, the V2DVS stores a continuous queue of 5 seconds (300 fields) of video images *prior* to any reported detection time (TOA). If the offset distance of the detection zone is specified, and the vehicle speed is properly measured, V2DVS can calculate the approximate video field in which the vehicle should be found when it was in the baseline (first loop) position. The image for that detector is then acquired at this previous time, *before* it was actually reported by the detector. Vehicle speeds are calculated by either the duplex detector itself, or by a detector that can be designated as a reliable speed measurement source in on [[trust]] line in the configuration file. NTSC video cameras of the type we use with V2DVS acquire a 'half-raster' image once every 1/60 of a second, or at approximately 16.667 ms sample intervals. The video image sample closest to the calculated TOA in the baseline zone is recorded as the reference image by V2DVS. In cases where a valid speed measurement is not available from the detector itself or the designated trusted detector, the most recently acquired vehicle speed is used. This is a usually the situation for a false detection, not observed by any other detector.

The offset distance, measured in inches, for image acquisition purposes is specific to each lane and detector. Our convention is that for detection zones down the road from the baseline position, this value is *negative*.

The last field (Latency) is an option to compensate for a fixed known time delay between the time detection and the time of reporting by the detector. It is specified in milliseconds. This is used to adjust the time of image acquisition for possible absolute time delays. Reporting delays from actual detection to electronic reporting are compensated by specifying *negative* values in this field, which means that the image will be acquired that amount of time *before* its detection signal is reported to V2DVS. Thus a detector which consistently delays 1 second after detection before sending its detection signal to V2DVS should have "-1000" in this field to compensate for the detector delay. There is no mechanism to compensate for variable reporting time delays.

Special note: the offset and latency compensation methods facilitated by the parameters described above only affect the acquisition of the image for a given reported detection. These are real-time decisions that must be made by the system, since only one image can be retained for each reported detection. They do not affect the adaptive weighted consensus calculation of ground truth from the inputs of all detectors under test, which is performed in post-processing by the V2DVS Workstation program. Offset and latency compensation during this post-processing may be adjusted repeatedly for maximum 'fairness' to the detector when comparing its TOA with that of every other detector.

Regarding the latency of the V2DVS system itself: the maximum possible inaccuracy between the time of the active low signal edge and proper time encoding by the system is 2.0 microseconds (.002 ms). This is an ultimate limitation of the PC-based hardware architecture of the Field Machines.

Trusted Detectors

In order to convert the "Zone Offset" from a physical offset into a time offset, V2DVS needs to know the speed of the vehicle. As mentioned above, the default for a duplex detector pair is to use the speed that it reports itself

for this purpose.   For a simplex detector, or a duplex detector for which speed is known to be unreliable, an alternative source of vehicle speed information may be designated by specifying a "trusted" detector for which speed measurements are known to be more reliable. The number of the designated speed detector is specified in the [[trust]] section of the detector configuration file. Again, this "trust" is only extended to calculating the image acquisition time offset and has NO EFFECT on the accuracy assessment of the detector.

It is possible to list multiple "trusted" speed detectors to provide a mean reported speed using a weighted average. To do this, each additional detector and its associated averaging weight should be listed in the detector configuration file on a separate line. For example, a 60/40 weighted speed average based on the speeds reported by detectors 1 and 2 would be represented with the line:

```
# Detector      Weight
         1       0.6
         2       0.4
```

9.3.3.3   View Menu

The **View Menu** has three options: **View External Triggers, Live Images,** and **Recent Images.**

Figure 9.16 View drop-down menu.

View External Triggers

Clicking **View External Triggers** will bring up a small window that appears to have 32 LED's, one for each contact closure channel, as shown in Figure 9.17. When an external detector signals V2DVS by a contact closure, the corresponding LED will light up for as long as the signal is low, then the LED will turn off. This is a very useful feature for real-time observation of the outputs of the detectors under test.

Figure 9.17. View External Triggers window, showing real-time state of all hardware triggers.

Recent Images

The **Recent Images** submenu allows the display of the most recently acquired image associated with a detection reported by any of the detectors under test. Clicking it will bring up a list of all the detectors under test in that lane. Select any one to open a window, which will freeze an offset-zone-adjusted image acquired from the last detection reported by that detector. The image will be retained until the subsequent detection. This feature is useful explaining why a detector may have falsely detected a vehicle, and it is helpful for adjusting the detector offset parameter (detector configuration, above) to assure the correct adjustment of the acquired image so that the vehicle is captured while it is in the baseline detection zone, which is ideally framed by the cameras.

Live Images

When Live Images is selected, a two-pane view is provided of video images acquired by the V2DVS image processor, shown in Figure 9.18.
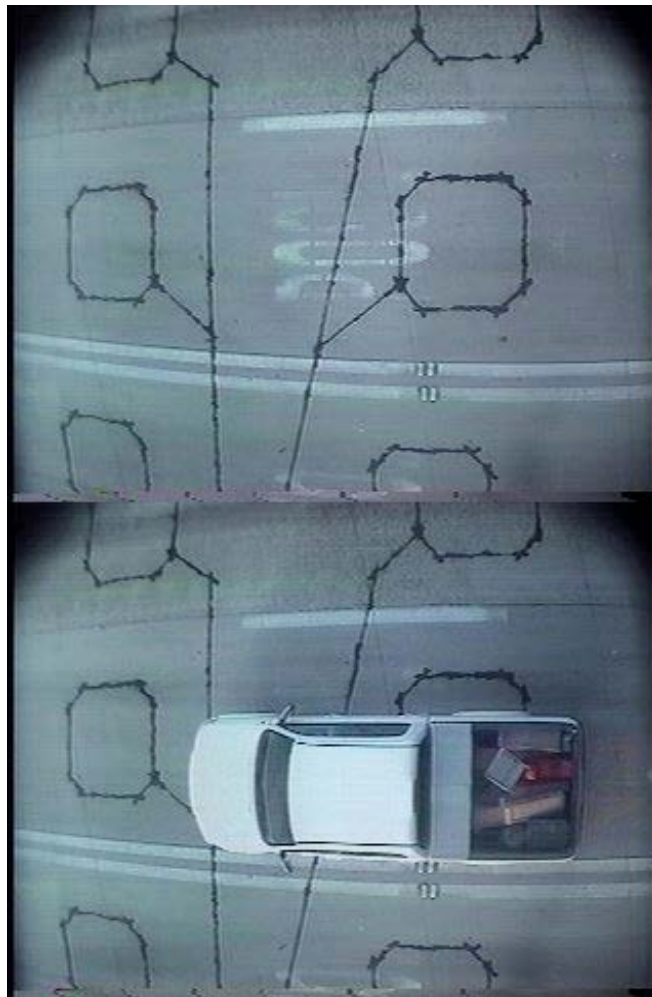
Figure 9.18. Two-pane live image window, showing real-time full motion video in upper pane, and the most recently acquired vehicle image acquired by the V2DVS image processor in lower pane.

Help

Clicking the **Help** tab brings up context searchable version of the user manual (this document, without appendices) in a PDF file reader. The PDF file is internally annotated so that a word search is may be used to locate a particular topic of interest.

10 V2DVS Vehicle Detection Subsystem

The V2DVS system contains a basic video-based vehicle detector, similar to commercial systems for vehicle presence detection. This functionality is provided to aid to system setup and monitoring, and as a supplemental presence detector required when only one detector is under test, since the adaptive ground-truth algorithm requires that at least two detectors be concurrently under test in each lane. As with most commercial video detection systems, it is typically very reliable as a vehicle presence detector under ideal illumination conditions, but degrades significantly in performance under less-than-ideal conditions. Although the image processing algorithm is designed to optimally reject long shadows, sunlight glint, and video image artifacts, its accuracy is poor after sundown, and it shuts off automatically under very low-light conditions or at night.

Note that the V2DVS image processing detector is implemented only to serve as an additional reference detector in cases where the set of detectors under test is inadequate for formation of an adequate consensus-based ground truth conclusion, and for acquisition of a vehicle reference image if none of the detectors under test are close enough to the baseline detection zone to provide a reliable image. It is not an actual project deliverable. It should not be considered as a competing detector to the detectors under test, since the placement of individual cameras above each lane is privileged and is not practical for general detection purposes.

While the V2DVS image processing algorithm internally measures vehicle speed and length, only presence detection is reported, although its speed measurement may be selected for use as a backup source of vehicle speed for detector zone offset purposes.

A large number of internal parameters may be set to optimize the V2DVS vehicle detection subsystem. These are described in detail in Appendix 2.

10.1 V2DVS image processing setup and tuning

The upper pane of the Live Images window always shows the live video feed from the camera at full frame rate. The lower image is selectable to show either the image of the most recently detected vehicle (at the position detected by V2DVS), or the image background (for diagnostic purposes). By default, the most recently detected vehicle is shown. Selection of either option is accomplished by right clicking the mouse while the cursor is above the image panes. A submenu will appear as shown in Figure 10.1 below:
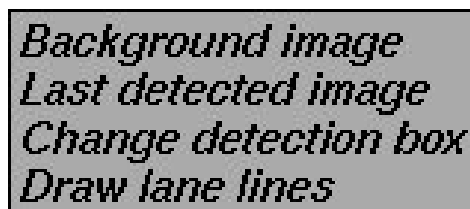

Figure 10.1. Submenu generated by right clicking mouse above the real-time view window.

Each tab on this submenu is selected by a dragging the pointer to the item while the left mouse button is held down, and releasing the button while the pointer is above the tab. Each option is described below.

Background Image

Click this to set bottom pane to show background image of roadway, without cars. This is used primarily for diagnostic purposes.

Change detection box

A detection window (box) must be drawn with a mouse on the live view upper pane to establish the zone of detection for V2DVS. This is required when the system is first set up, and any time that the camera is repositioned. This need only be done once, since V2DVS retains this information even after it is powered off.

Click the *Change detection box tab* to set up the detection window. After clicking on this item, the cursor will become cross-like. Position the cursor over the real-time video pane at the intended upper left position of the *detection box*, click and hold the left mouse button, and stretch the resultant window frame to the right and down, to the lower right corner position of the window. Release the mouse button to set the window. A vertical orange line will then appear through the cursor. This is the *detection position line*, which specifies position in the scene at which the image of each vehicle will be captured. Move the line to somewhere within the left half of the detection box, leaving enough room the right of the line for tracking of the vehicle. The detection window should usually be contained within the lane lines, and extend the equivalent of a physical distance of about 25 feet along the roadway. A typical view of the detection window and detection position line is shown in Figure 10.2 below.

Note: Although it is possible to do so, a detection box should NOT be drawn in the lower (most recent image or background image) window; doing so will cause incorrect or erratic detection since the box will be outside the field of view of the camera.
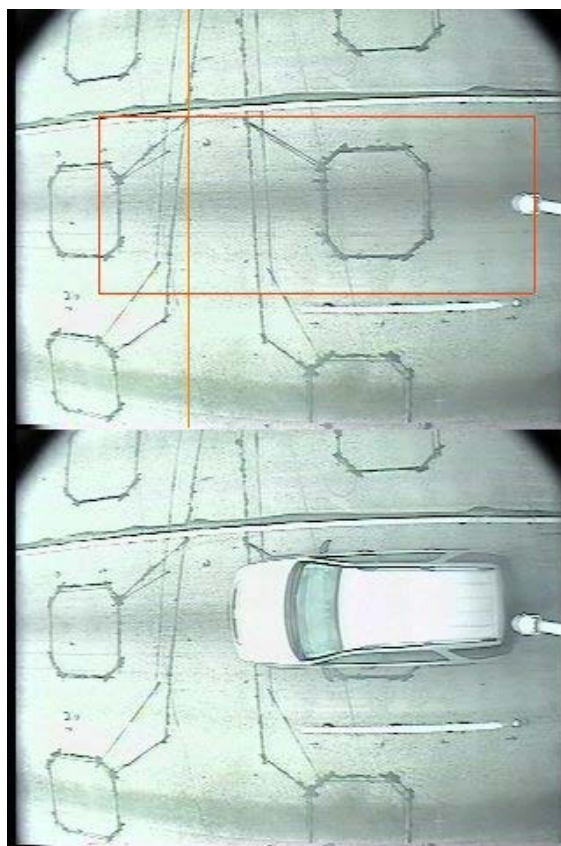

Figure 10.2. Example of correctly positioned detection window (red box) and detection position line (orange vertical line) in the live view pane.

Draw lane lines

It is also necessary, at least once, to establish a dimensional relationship between the camera image coordinates and the actual roadway coordinates, so that distance measurements in the camera image may be translated to actual physical distance measurements. The V2DVS detection subsystem provides a means to do this regardless of the camera zoom or focal length setting, without explicit entry of a conversion factor. This is done by clicking the *Draw Lane Lines* tab, which will cause a sequence of two horizontal orange lines to be displayed which are positioned using the mouse. Position the first line on top of the right (upper) lane line at the midpoint in the image. Left click to set this position. A second horizontal orange line will appear. Position this line using the mouse on top of the left (lower) lane line at the mid-point in the image, as shown in Figure 10.3 below. Left click the mouse to set this line. Upon left clicking, both lines will disappear. This task need only be

performed once, unless the zoom setting (focal length) of the lens is changed. Unless changed again, these settings are retained by V2DVS between sessions, even after powered down.
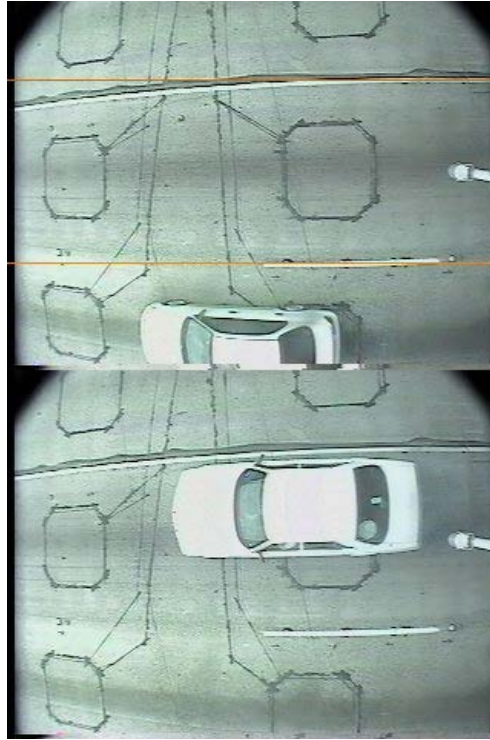

Figure 10.3. Lane lines properly drawn on live camera image.

10.2 V2DVS Image Processing Performance Results

While not a requirement of the system or project, we tested the V2DVS video detection performance under a range of traffic and illumination conditions. Table 10.1 shows the results of a test of the system accuracy relative to a small human-verified (from video tape) ground truth data set:

Table 10.1 Performance of V2DVS vehicle detector subsystem.

Illumination condition	Sample size (Total vehicles)	Correctly detected	Failed to detect	False detected	Velocity measurement
Overhead sun	300	299 (99.7%)	1	0	Excellent
Diffuse moving shadows	300	298 (99.3%)	2	0	Excellent
Crisp moving shadows from adjacent lane	300	295 (98.3%)	5	0	Excellent
Fixed (bridge) shadow in detection zone	200	197 (98.5%)	3	0	Acceptable
Dappled shadows from adjacent lane	200	194 (97.0%)	3	3	Acceptable
Low light (20-30)	300	298 (99.3%)	2	1	Acceptable

Under ideal or near-ideal lighting conditions, V2DVS correctly acquired vehicle images in more than 99% of all cases, and is very good at not falsely detecting vehicles. The ability of V2DVS to reject shadows from adjacent lanes is very good, except when the shadow itself has the texture of an actual vehicle. This can occur when a partial shadow is created by sunlight passing through the windows of the vehicle onto the pavement. Such cases are referred to as "dappled shadows" in the table above, and represent the weakest detection conditions, only 97% accurate with 1.5% false detections. Accuracy degrades under dappled or transient illumination conditions, such as split shadows from the bridge or passing cloud cover.

11 Installation of Detectors to be Tested

As discussed previously, all detectors signal the V2DVS system in real time using hardware triggering compatible with Caltrans Standard Type 222 detector cards. The electrical output of these "loop detector" cards is open collector, active low – with a maximum acceptable 'low' voltage of 0.7 volts. Caltrans Type 170 and 2070 controllers expect these levels for proper interface with loops or other vehicle presence detectors.

All detector outputs are connected to the V2DVS system via the telephone-type punch-down terminal block shown in Figure 11.1. Detector outputs are routed to Cabinet 2 via a conduit from Cabinet 1, which is reserved for use by detector system electronics, including the Type 222 loop detector cards located for the Testbed duplex loops.

As mentioned previously, lanes are color coded to match the color coding of the Field Machines and their cabling:
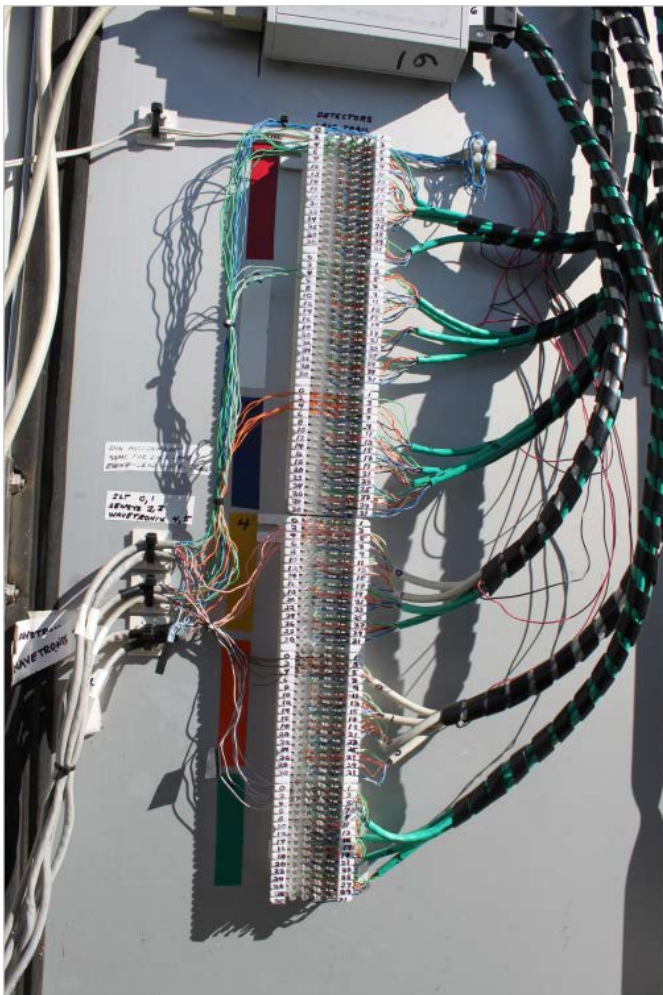

Figure 11.1 Detector interface terminal block located on rear door of Cabinet 2.

Detector outputs enter from the left and are connected to the V2DVS by punching down the 22 AWG wires onto the corresponding terminals using a special punch-down service tool, although a pair of regular pliers works.

Figure 11.2 shows a close-up view of the terminal block on the rear door of Cabinet 2.   Within each color group there are 32 connections, numbered 0 to 31. With four terminal across in each row of the terminal block, the two terminals on the left are a single common connection, and the two terminals on the right are another common connection. Even numbered connections are on the left, while odd numbered connections are on the right. We have established a convention that the lead detector or a duplex pair is always connected to an even terminal, while the trailing detector is connected to an odd-numbered terminal.

As shown in Figure 11.2, detector terminals 0 through 5 for each of the six lanes are currently assigned as follows:

Terminals 0 and 1: IST inductive loop detectors
Terminals 2 and 3: Sensys wireless magnetometer detectors
Terminals 4 and 5: Wavetronix phased-array RADAR detectors.

Additional detectors should be added following this convention. Note that the V2DVS image processor is treated as a detector for analysis and display purposes, but does not use input terminals on this block.
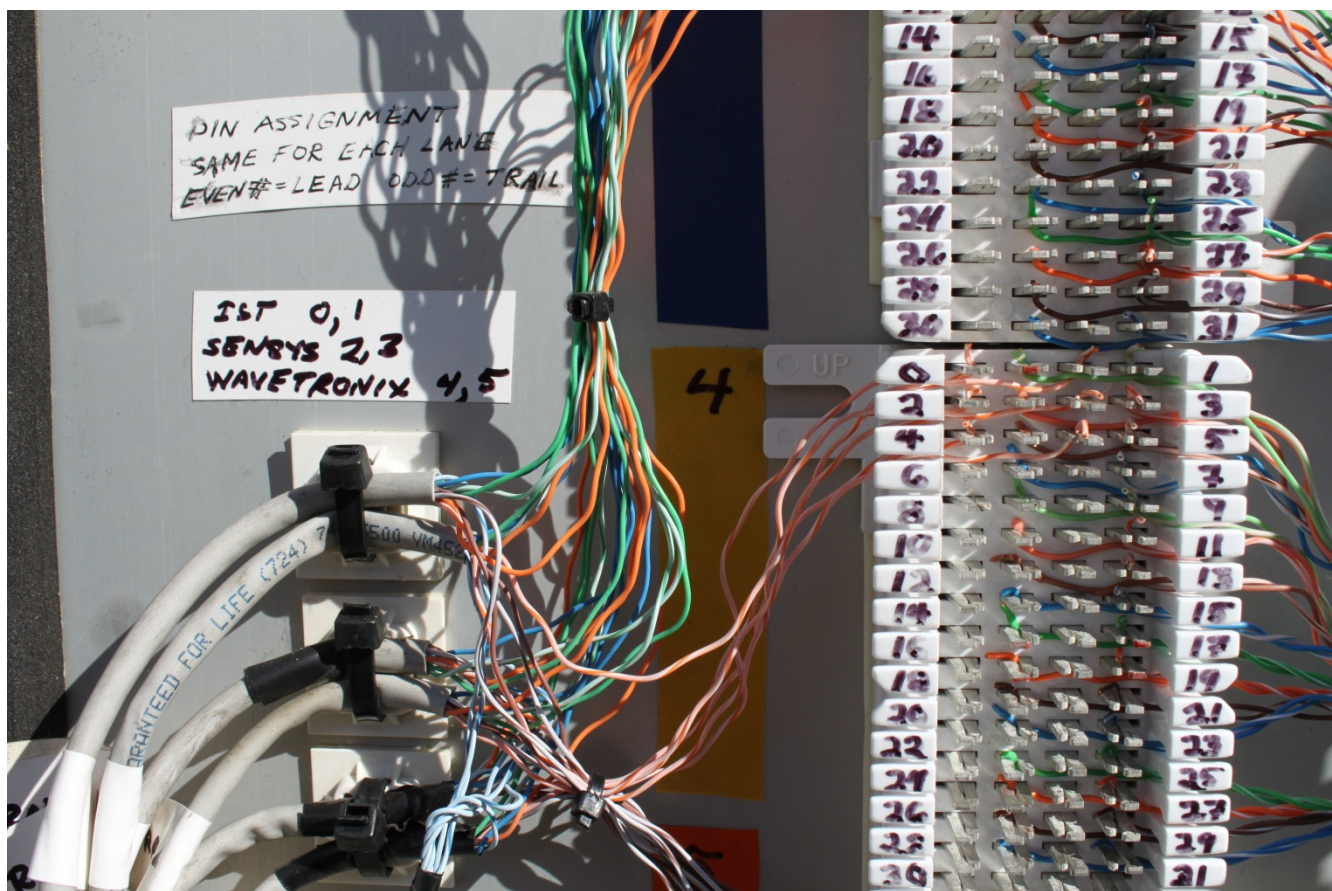
Figure 11.2 Close-up view of the detector connection terminal block.

Power for the detectors is available in Cabinet 3: 110VAC, 24 VAC, 12 VDC and 5 VDC.

Type 222 specifications call for a "switch closure" type of detector output, with both output terminals floating with respect to either the power supply or device ground. The common detector low-side switch contact (or *emitter contact* as referenced by some Type 222 loop detector manufacturers, is referred to in TEES (2009) as the "C1 Harness DC Ground" (2) since it is the ground connection for the Type 170 controller via the large "C1"

Amphenal connector. If the 170 Controller is not connected, this terminal is not connected to any ground point in the cabinet. Since the testbed installation does not include a 170 controller, we must provide this pseudo-ground connection. For practical purposes, we will simply refer to it as the *detector ground*. It is located on Terminal 4 of Terminal Block T-5 at the bottom left of the back of the Type 222 cardfile in Cabinet 3, as shown in

Figure 11.3. This common connection point is also found on a terminal block located above and to the right of the detector terminal block in Cabinet two. Note that the detector ground is not the same as the Equipment Ground or "DC Ground" provided by the main power supply, and these should not be connected together.

Appendix 1 provides the full TEES specifications, and shows the 222 backplane wiring diagram.

We found that many Type 222 loop detector cards and 222 emulation cards do not faithfully implement both switch contacts as floating connections – they implement the more typical "common-emitter" active-low connection, with the emitter of the final drive transistor connected to the DC ground of the card. Further, the final drive is typically done using a Darlington connection of two NPN transistors, with the emitter of the first stage connected to the base of the second stage. Of significance here is that the minimum "closed switch" voltage drop of such a connection is 0.8-0.9 Volts. Our detector interface box were designed assuming that a true "switch closure" type of output would be provided, as specified in TEES, with a minimum voltage drop when active low (closed) of less than 0.7 volts.

To remedy this problem for all Type 222 cards or interfaces, we take advantage of the lack of the Type 170 controller in the cabinet, and establish the detector ground at approximately 1.2 volts below the "DC Ground" level found on terminal 2 of the same T-5 terminal block. This was accomplished by using an external 3.0 Volt regulated DC power supply with the positive supply side connected to Terminal 2 (DC equipment ground) and the negative side connected to Terminal 4 (detector ground) through two silicon diodes in series**.** By depressing the detector ground level to -1.6 volts (3.0 – 1.4), even the weakest detector output will appear to the detector interface box as providing a solid active-low switch closure voltage at or below zero volts, which will be reliably read as a presence detection trigger signal. The 3.0 Volt external power supply is a power brick that is plugged into a power strip in Cabinet 1. *This power brick must be on for the V2DVS to operate correctly.*
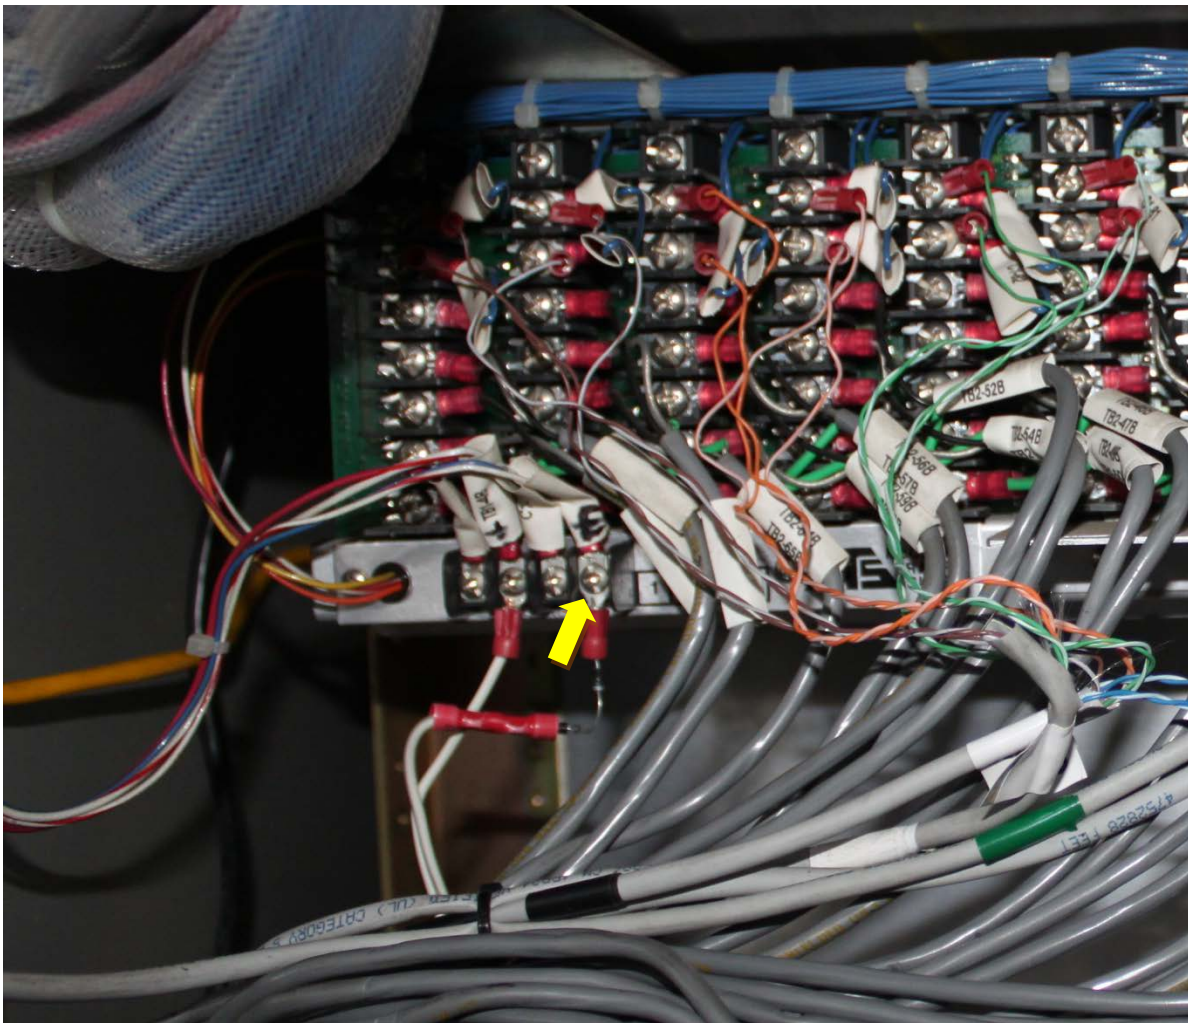
Figure 11.3 Location of common detector ground Terminal 4 on rear of 222 cardfile in Cabinet 3.

11.1 Internal Utilities

While not normally ever needed by a user, some additional utility programs are available to assist in troubleshooting or setting up the V2DVS on the Field Machines.

A command line window may be brought up outside of the V2DVS program by clicking on the "terminal" icon in the icon tray at the bottom left of the main screen.   All standard Linux/Unix utilities are available, and in addition, some specific utility programs to help quickly identify and set configuration parameters used by V2DVS.

One internal utility of particular value is rt_config. This command takes as its first parameters one of three arguments: list, get or set.

list – Displays the current values of each real-time parameter used by V2DVS for detection.

get – Requires one argument, the name of the single real time parameter that you wish to display.

set - Requires two arguments, the first the name of the real time parameter you wish to set, and the second the value you wish to set it to.

Cautionary note: Entering an inappropriate value in rt_config can potentially crash the program, and is recommended for diagnostic purposes only. All normally-needed configuration parameters are editable via the V2DVS window interfaces.

```
Usage:
      rt_config list          -List all configuration values
      rt_config get <param>   -List 'param' configuration value
      rt_config set <param>   -Set 'param' configuration value
Examples:
      rt_config get c_collect
      rt_config set c_collect 1
```

Figure 11.4   rt_config help (displayed when rt_config is run with no parameters)

12 The V2DVS Server

The V2DVS central server computer is connected to the field machines either directly via a local area network cable, or over the Internet from a remote location. At the present time, it is connected via a CAT5 network cable to a network switch in Cabinet 2, through a conduit between Cabinets 2 and 3. In normal operation, the user simply turns on the server and does nothing further. The server is designed for "plug-and-play" operation. Upon activation, the server will display the screen image of Figure 12.1

The central server collects data from the field machines. It acquires all vehicle record data and images, storing the record data in a MySQL database and the files in a referenced folder hierarchy.

The central server is physically a 4-U rackmount computer with a total of 1.2 Terabytes of data storage capacity. It communicates with all field machines via a CAT-5 (100-Base T) Ethernet connection, and is designed to function automatically unattended on a 24/7 basis.

Once the server is powered up, it automatically launches the database and communications services. A simple user interface is generated to show the status of the server and data collection activities. Whenever the server is running it will automatically start attempting to collect data from field machines, and start hosting Workstation (PC client) data requests from remote V2DVS Workstation sessions. The application may be exited by clicking the ON/OFF symbol at the right bottom of the icon tray on this screen. Select "shutdown" to turn off the server.
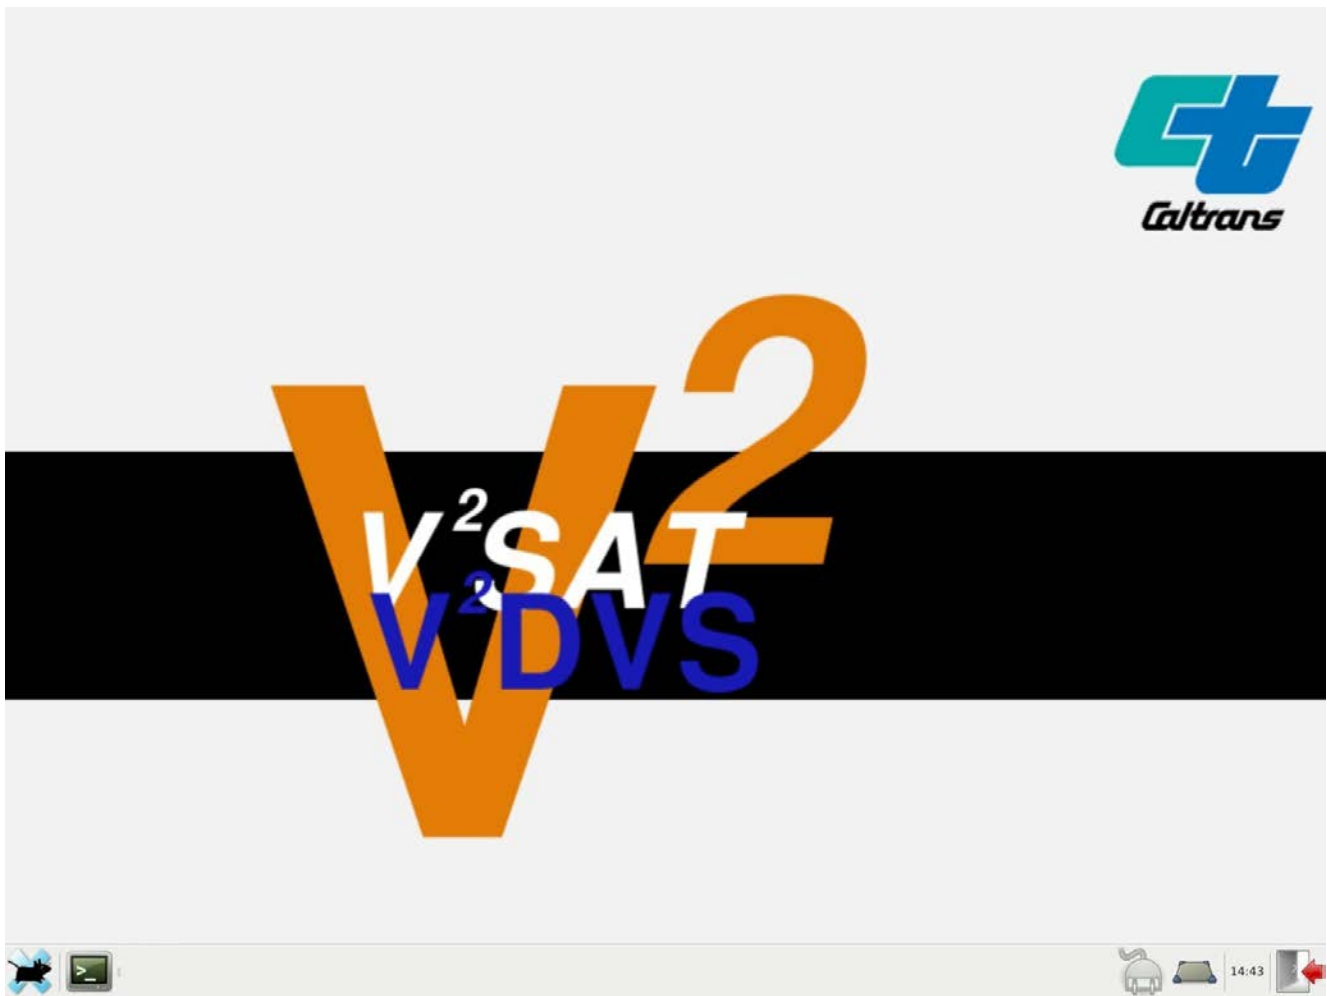
Figure 12.1. Opening screen image of V2DVS Server.

To monitor the operation of the server, a command-line window can be activated by clicking the terminal icon at the bottom left of the screen. This provides access to several command-line utilities related to the V2DVS server as well as all Linux operating system functions. Under normal conditions, the user need not use any of these capabilities.

The V2DVS is designed to be able to collect detector data even when the server is off or disconnected from the network. The field machines can each accumulate up to 600 Gbyte (several days) of image and log data that will be uploaded to the server whenever it is accessible. Once the server has established a network connection to the field machines, transfer of data starts automatically. In view of the huge size (in Mbytes) of the image data generated for typical traffic by each field machine, it may take a long period of time to retroactively transfer field data to the server over a slow network connection. A slow network connection between the field machines and the server does not prevent the transfer of data, but it can easily slow it to the point that the data flow rate over the network is less than the data accumulation rate of all the field machines operating at once. It is for this reason that the server is currently located in Cabinet 3, immediately adjacent to the Field Machines in Cabinet 2, on the same local area network. This implementation incurs no possible network-related delays.

13 V2DVS Workstation (PC client application)

An MS Windows (XP, 7) program is provided to initiate and stop data collection, control the video cameras remotely, and to process and verify all detector data collected by the field machines. It interacts with the server database, which contains all data and vehicle images acquired by the field machines during data collection periods. The program, referred to as The V2Workstation, is packaged for installation as a standard windows application with installation wizard.

13.1 System Settings

The PC monitor must support a minimum resolution of 1024x768 in order to view the complete windows provided by this program. Other than this requirement, any network-capable Windows PC or laptop computer should be compatible.

13.2 Program Installation and Setup

The installation package will appear as the following icon:

V2DVS_Workstation_v1.4.exe

It can be found on the CD accompanying this manual. The file can be placed anywhere on the computer, such as the windows desktop. To install the V2DVS Workstation Program, double-click the program icon to start the installation. Follow the prompts of the installation wizard. It is not necessary to uninstall any previous versions that may already exist on the computer before installing the latest version.

13.3 Program Usage

After installation, a program icon will be added (if this option was selected during installation) to the desktop:

V2DVS Workstation.lnk

Double click this desktop icon to launch the V2DVS Workstation Program.   Your PC must have a network connection with the server so that the program can communicate with the database on the V2DVS Server and optionally, a network connection with all field machines to use the data collection start/stop capabilities and remote camera controls. Normally, this is accomplished at the test site by a CAT5 connection between the PC and one of the open CAT5 connectors on the 16-port D-Link switch in Cabinet 1. It is also possible to communicate with the Server and Field machines over the Internet if a broadband network connection is provided to the Testbed. The local IP address of the server is 192.168.40.50. The Field Machine IP addresses are 192.168.40.51 through 56. These local IP address are not user-changeable. For bidirectional communication with the V2DVS Server and Field Machines, the PC's network card must be set to a fixed IP address in subdomain 192.168.40 between .00 and .255, but outside the range .50 through .56. **We suggest using IP address 192.168.40.57**.

Upon launch, V2DVS Workstation will display the Startup Screen shown in Figure 13.1
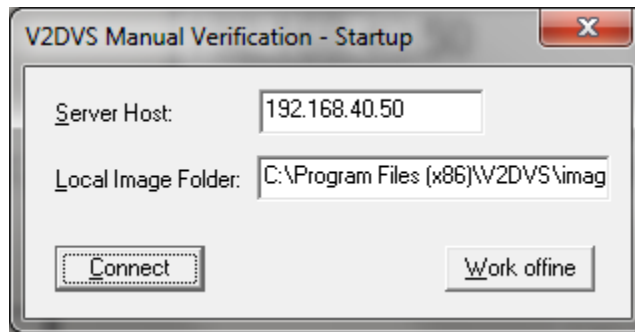
Figure 13.1 V2DVS Workstation Startup Screen – Usually just click "Connect".

At this point, the user has two choices. If data and images have been previously downloaded onto the local PC, this data may be viewed and analyzed by clicking **Work Offline,** which will display the blank canvas screen of Figure 13.2.
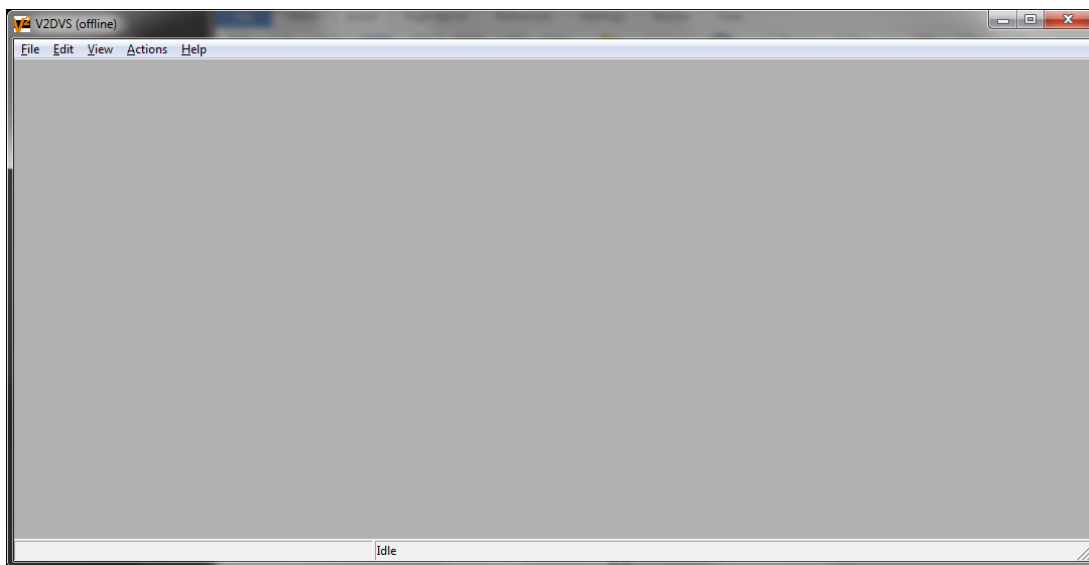
Figure 13.2 V2DVS Workstation Main Window, with Welcome Menu.

If **Connect** is clicked and the V2DVS Server is online, the Welcome box shown in Figure 13.3 will be displayed.

Figure 13.3 V2DVS Welcome Menu, for selection of task to be performed.

The three options to select within the Welcome box are **V2DVS Status and Control, Verify New Detector Data, and Continue Previous Verification.** All of these functions are also available on the main menu bar.

The main V2DVS menu bar has five options as shown in Figure 13.4, each corresponding to a group of features. If you have chosen to work offline, only a subset of these features will be available.

Figure 13.4 V2DVS main menu bar.
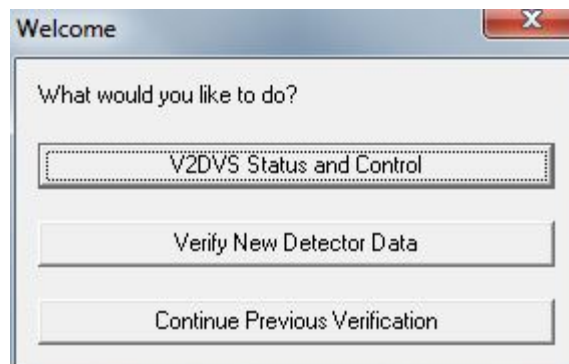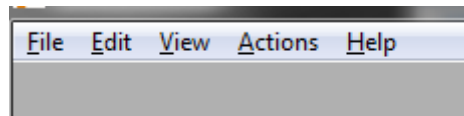
File
> Under **File** on the main window are the options:
> **Save Session:**   save the current session for later
> **Save Session As:**   save the current session for later specifying what the file will be
>   called and where it will be located.
> **Exit:**   exit the program

Edit
> Under **Edit** on the main window are the options:
> **Lane Manager:**   opens the lane manager window for sessions that are being/have been
>   generated
> **Detection Offsets:**   allows the user to view/adjust (local only) the detector offsets

View
> Under **View** on the main window there are the options:
> **Results:**   displays the results for the current report
> **Toolbar:**   various other options are displayed
> **Freeze Frame:**   captures the current view of the detection verification window

Actions
> Under the **Actions** on the main window are the options:
> **V2DVS Status and Control:**   shows the status of the V2DVS, and start or stop data collection
> **Verify New Detector Data:**   start a manual verification session with new detector data
> **Continue Previous Verification:**   open a previously saved session and continue manual verification
> **Data Deletion:**   remove data from the database and the user's machine
> **Automatic Correlation:**   displays the correlation parameters and allows the user to run   an
>   automatic correlation for the current report.

Help
> Under **Help** on the main window there are the options:
> **About:**   displays information about the current program
> **Users Manual:** displays the user's manual

Most actions are self-explanatory once selected.

If **Actions** -> **V2DVS Status and Control** is selected, the window of Figure 13.5 is displayed showing the current operation conditions of all the Field machines and allowing the user to start and stop data collection in any or all lanes.

V2DVS Status

Lane	Status						Data Collection					Control Camera
	Disk		Ram		Time		Running	Logs	Images	Last Detection	Start / Stop	
	Size	Free	Total (K)	Free (K)	Uptime (Hr)	Idle (Hr)						
1	616G	584G	2060496	50576	0.95	3.26	ON	0	1	10/24/11 03:40:29 PM	STOP	Lane 1
2	613G	582G	1020136	62140	0.06	0.22	ON	0	6	10/24/11 03:40:00 PM	STOP (START ALL)	Lane 2
3	616G	584G	2060496	908420	2.84	9.89	ON	0	10	10/24/11 03:43:00 PM	STOP	Lane 3
4	616G	585G	2060496	834168	0.26	0.92	ON	0	2	10/24/11 03:42:12 PM	STOP	Lane 4
5	616G	584G	2060496	273856	4.67	16.55	ON	0	1	10/24/11 03:41:02 PM	STOP (STOP ALL)	Lane 5
6	616G	585G	2060496	715832	0.36	1.28	ON	0	12	10/24/11 03:41:03 PM	STOP	Lane 6

Quick Refresh   Full Refresh   Close

Figure 13.5 V2DVS Status and Control Window.

Each video camera, one per lane, is equipped with a remote control wide-angle zoom lens, which can be controlled from either the V2DVS field machine program, or from the V2DVS Workstation program running on a network-connected PC. The user interface from either unit is the same.

Remote camera control can be selected from the V2DVS Status window. The field machine must be connected in order to communicate with the camera, which requires clicking the corresponding "Lane" under the "Control Camera" heading. The camera control window is then launched, as shown in Figure 13.6.

Camera Control - Lane 1

Camera Controls

Change to Manual Lens Control

Iris: Open / Close
Zoom: In / Out
Focus: Near / Far

Live   Last Detection   Background Image

Refresh   Close

Figure 13.6 Remote Camera Control Window.

From this window, the user may control the aperture, focus and zoom (focal length) of each down-looking video camera on the overcrossing. At the bottom of the camera control window are three images, from left to right: the most recent video image from the camera at the time the window was activated or the refresh button was clicked, the image of the most recently detected vehicle (detected by the V2DVS image processor), and the current background image (for diagnostic purposes). Each image can be reloaded by clicking on the refresh button located at the bottom left corner of the window.

There are two modes of camera control: automatic and manual. These modes can be toggled by clicking on the long top button. If *automatic mode* is selected, the V2DVS system controls the camera aperture (iris) to assure optimal video image quality, required by the image processor. The focus and zoom remain at the last values that are set manually. If *manual mode* is selected, the user has the option to adjust the camera aperture manually, in addition to the zoom and focus. During manual operation, the camera lens functions are controlled using the six buttons at the top of the window, allowing the user to zoom in or out, refocus, or manually adjust the iris of the lens. During normal operation, it is important that Automatic camera control is selected, to allow V2DVS to control the aperture as the light level changes.

If the user is communicating with the field machine via a remote Internet connection, there may be considerable latency for some of the camera control actions.

13.3.1 Starting or Stop Data Collection

Also on the Status and Control Window are a series of commands to start or stop data collection, and monitor the data collection process. A close-up view of the applicable part of this window is shown in Figure 13.7. The START ALL and STOP ALL buttons activate data collection concurrently for all detectors in all six lanes. The amount of data waiting to be transferred to the server by each field machine can be monitored by observing the number of vehicle detection log entries and images that remain to be transferred, shown in columns 2 and 3.

Data Collection					
Running	Logs	Images	Last Detection	Start / Stop	
ON	642	652	10/22/11 02:23:28 PM	STOP	START ALL
—	—	—	Not Connected	START	
—	—	—	Not Connected	START	
—	—	—	Not Connected	START	STOP ALL
—	—	—	Not Connected	START	
—	—	—	Not Connected	START	

Figure 13.7 Control and monitoring of data collection.

13.3.2 View and Verify Detector Data

A *session* is a collection of detection data over all available lanes for a defined time interval. To generate a session, click on the menu option **Actions->Verify New Detector Data**. This will open the **Verify New Detector Data** window shown in Figure 13.8. The calendar will indicate dates on which data is available. The user may jump to any specific month or year by clicking on the month name or year at the top of the calendar Select a day on the calendar, and click START DATE or END DATE to automatically transfer this date to the corresponding selection boxes. The lower area of the window will then display the time periods within the

selected range of dates during which data is available for each lane. The start and stop times for the selected date range must be manually entered into the "Starting Time" and "Ending Time" selection boxes. However, a selection "spinner" for at the right of each box expedites the time specification process. After the date and time ranges have been fully specified, click **Next** to proceed to the **Lane Manager** in which the lane or lanes to be viewed or verified are selected, as shown in Figure 13.9.

Figure 13.8 Window for specification of detector data to incorporate into a session.
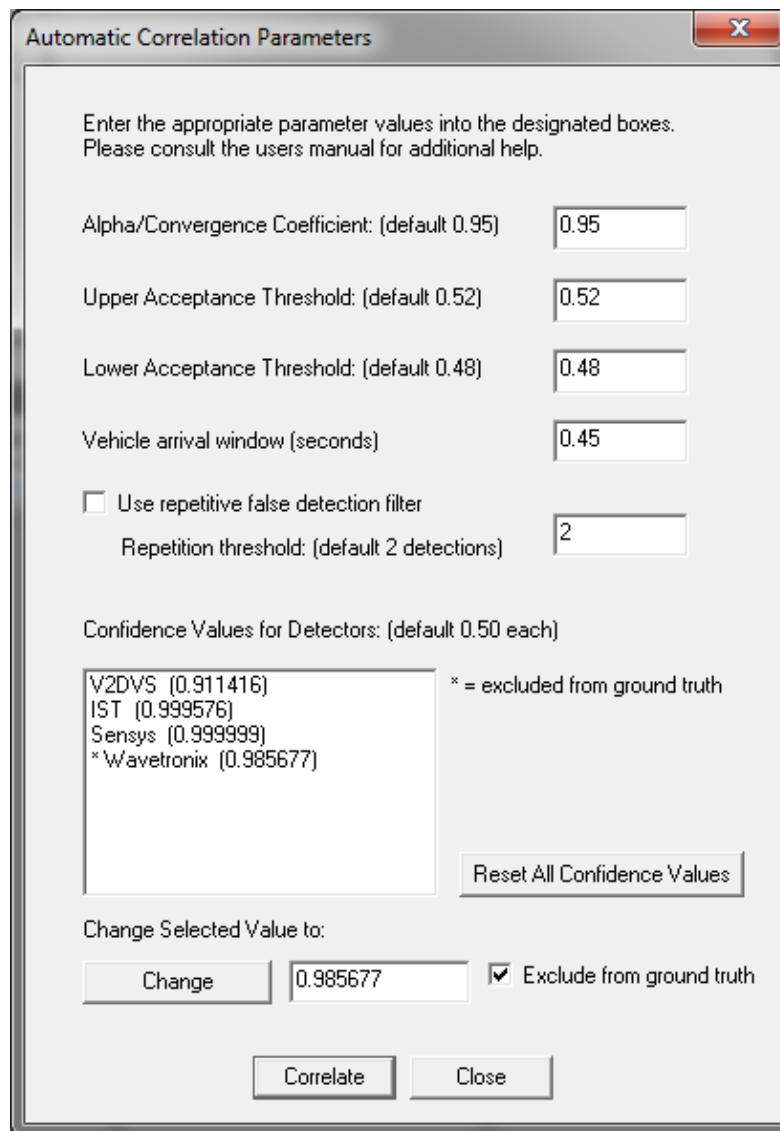
Figure 13.9 Lane manager window, for selection of lanes for which data is available during the session.

13.3.3 Automatic Ground Truth Generation

As will be discussed in detail later, the V2DVS Workstation application determines the accuracy of each detector under test by comparison with a ground truth dataset for presence, speed and length. This reference dataset is generated using a weighted consensus of all detectors under test including the V2DVS image processing detector. A "learning" algorithm continuously adjusts confidence factors for each detector used in the formation of each ground truth decision to assure the most accurate possible ground truth dataset. It is important to note that the ground truth dataset is generated for each session in post processing by the V2DVS Workstation application; it is not generated by the V2DVS field machine in real time.

To run the adaptive ground truth generation algorithm, click on the **Automatic Correlation** tab under the **Actions** menu item on the main window. The window of Figure 13.10 will appear, from which the automatic correlation parameters are displayed and may be changed for the given session. The Alpha/Convergence Coefficient, and Upper and Lower Acceptance Thresholds will be discussed later in the context of the adaptive consensus algorithm. The "Vehicle Arrival Window" is the time aperture (seconds) during which offset-correct detections generated by different detectors are admitted as representing the same vehicle. Typical values between 0.4 and 0.5 seconds have been found to be optimal for the best trade-off between acceptance of the detection and rejection of proximate false detections.

Figure 13.10   Submenu window for displaying and optionally changing ground truth generation parameters.

The optional feature "Use Repetitive False Detection Filter" may help to discriminate cases of detectors which are prone to multiple detections of the same vehicle. Individual detectors may optionally be removed from usage by the adaptive ground truth algorithm in its formation of the ground truth dataset. This feature would normally be used if a particular detector was not operational during some significant period of the session, since its perceived continuous failures to detect during these drop-out periods could otherwise skew the accuracy of the automated correlation algorithm, at least until the confidence factor for that detector is eventually nulled to zero due to its persistent inaccuracy.

Click **Correlate** to initiate the process of automatic correlation of the reported detections from each detector to identify actual vehicles, and discriminate false detections to form the automated ground truth dataset. All detections are then compared with this dataset and identified as correct detections, false detections or unverified due to marginal correlation.

At the completion of this process, a graphical data display similar to that of Figure 13.11 will appear, which color-codes the correlated detection (actual vehicles) in dark blue, false detections in red, and unverified detections in yellow. A small window summarizing the results of the automatic correlation process will also appear, as shown in Figure 13.12.
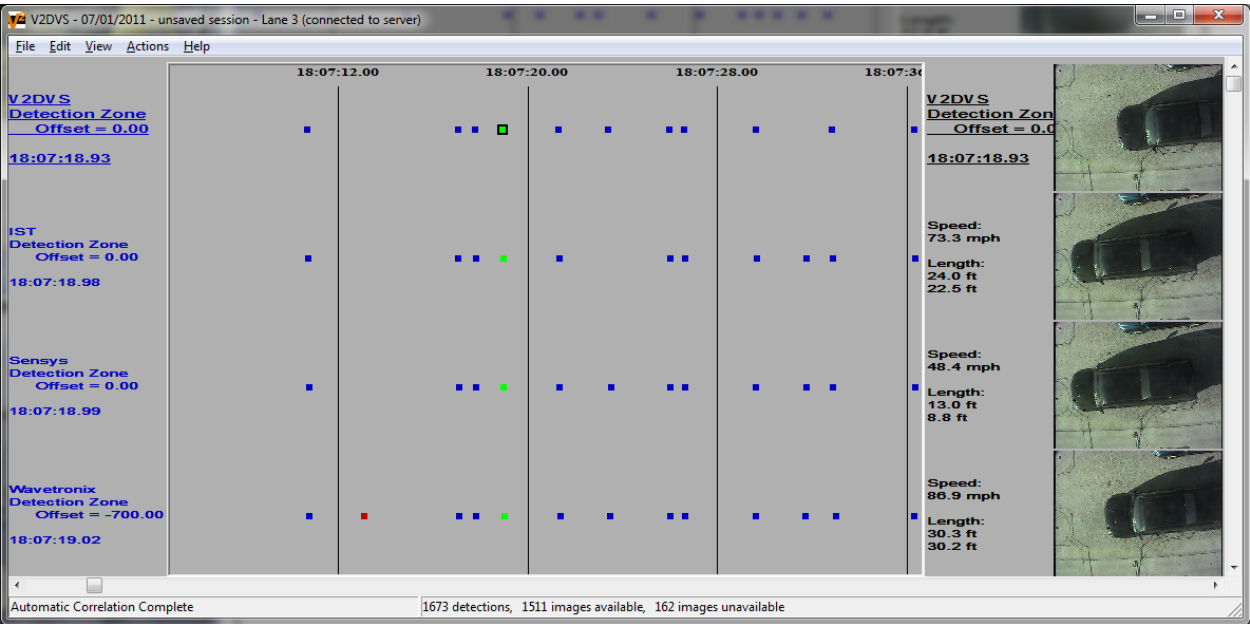
Figure 13.11 Session Window.
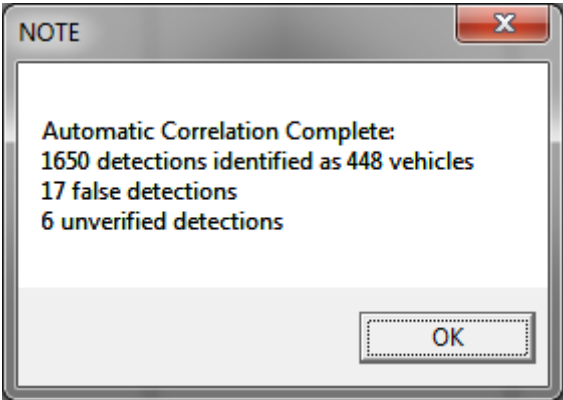
Figure 13.12   Summary of results of automatic correlation.

13.3.4 Manual verification for finalization of the ground truth dataset

Although the automatic correlation process is typically 98+% complete in its generation of accurate ground truth from the inputs of all detectors under test, some events may remain "unverified", marked as yellow dots. Situations in which the V2DVS adaptive consensus algorithm cannot adequately classify a detection event as correct or incorrect must be "manually verified", that is, resolved by a user graphically by referencing the images acquired for each detector for the given vehicle or false detection. This final correction process is facilitated by the Session Window shown in Figure 13.11. When a detection dot is selected, the corresponding image acquired at the time of the detection is shown on the right of the window, and may be enlarged by clicking on it.

The **Toolbar** menu shown in Figure 13.13 appears whenever the Session Window is displayed. This provides an assortment of features for expediently inspecting the dataset, correcting and possible errors in the automatic correlation results, and verifying the "unverified" cases.
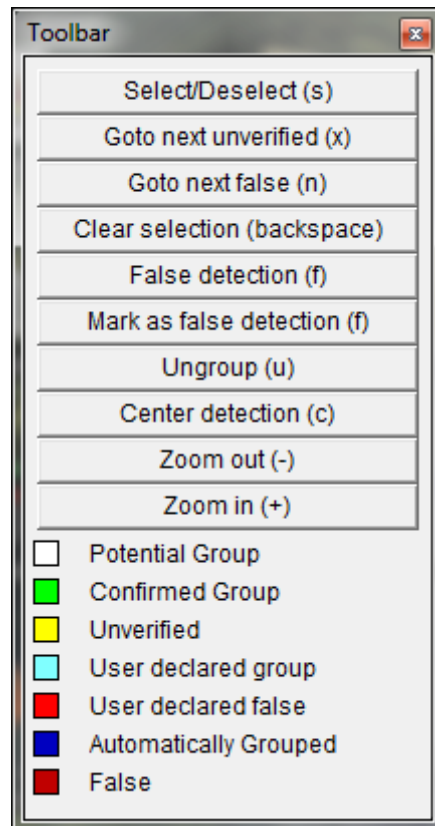
Figure 13.13 Menu for selection of appropriate action for any selected vehicle "dot" or "group".

The user can traverse the detections using either the mouse or the arrow keys on the keyboard. The selected dot (detection) is shown highlighted by a black square surrounding the detection dot. Once selected, verified detections becomes green dots. Multiple dots may be highlighted to indicate all detections that the user believes correspond to the same actual vehicle. The user can confirm the group (or identify false detections), by clicking on the appropriate button on the Toolbar menu of Figure 13.13, or by a using the corresponding shortcut keystroke. The possible actions are:

Select/Deselect (s): Selects an individual detection or member of a detection group corresponding to the same vehicle. Selected detections become highlighted as white dots.

Goto unverified (x): Jumps forward to the next unverified detection (yellow dot) on the timeline.

Goto next false (n): Jumps forward to the next false detection (red dot) on the timeline.

Clear selection (backspace): Clears all selected detections.

Group (g): Selected dots are identified as a group, which corresponds to the detection of the same vehicle. Once the group is formed in this way (manually), the dots in the group change color to teal. Note that if a selected group of dots (detections) are grouped while some were initially members of another group (vehicle), the remaining detections of the other group will remain in their own group.

False detection (f): Toggles a highlighted or selected unverified detection as false. False detections are represented with bright red dots.

Ungroup detection (u): Toggles a highlighted or selected verified detection as *unverified*, and the dot turns yellow to indicate this.

Unverified detection (f): Toggles a highlighted or selected false detection as *unverified*.

Center detection (c): Used to center the currently-selected dot (detection) in the middle of the time line, as a convenient alternative to scrolling forward or backwards in time.

Zoom in (+): Zooms in and centers a highlighted detection. Increases visible time scale of the detection.

Zoom out (-): Zooms out and centers a highlighted detection. Decreases visible time scale of the detection.

Use scroll bar slider to move right (forward) or left (backward) through the timeline. The program also supports the use of a mouse wheel and "click and drag" to perform this function. The page up and page down keys move the timeline one quarter of the current window width. The arrow keys can also be used to move the detection cursor. Moving the detection cursor to a detection on either edge of the visible timeline causes the timeline to automatically scroll.   The user can also use the **x** key on the keyboard to proceed directly to the next unverified detection on the timeline.

Displayed in Figure 13.14 is the **Detector offsets** window, which allows for accurate treatment of detection times despite different detection zones in the same lane. Note that negative values represent detection zones that are located past (down the road from) the baseline detectin zone, which is immediately below the overcrossing bridge deck. If the detector under test reports bad/unreliable speeds, or an incorrect offset distance is set, then the offset-adjusted positions of some vehicles may appear incorrectly ordered in the manual verification display. When viewing or performing manual verification, scrolling right or left through the vehicle detections may appear to jump over an affected vehicle. This problem may be corrected by the user by changing the detector offset from this window.

For detection zone offset time calcuation purposes, the V2DVS Workstation application requires a reliable source of the vehicle speed. By default, this is the detector itself assuming that it reports speeds by generating a real or synthetic duplex detector output (two edges for each vehicle). If detectors don't report speeds, the user may select the particular detector or detectors to use as an alternative source of speed for offset adjustment purposes, via the Detecter Offsets window shown below. The asterisk(*) next to the detector name in the window denotes the alternative speed source to be used.
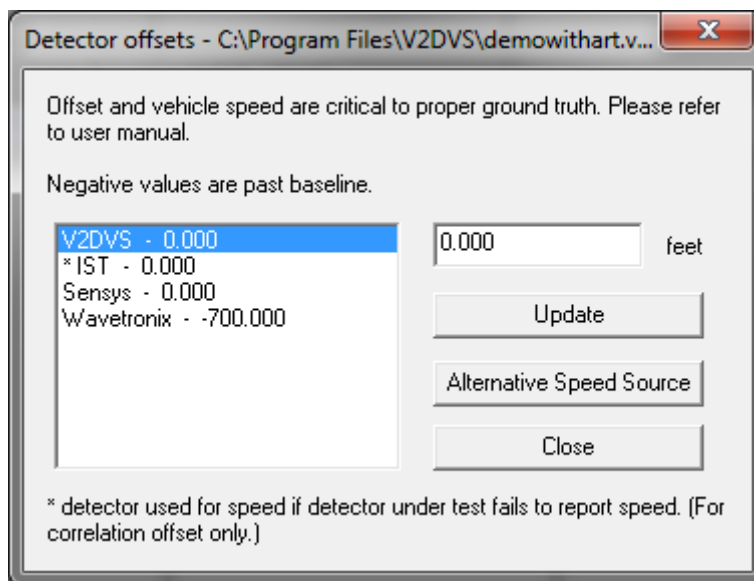
Figure 13.14 Window in which the position (offset from the baseline position) of the detector's detection zone may be specified, in feet. Negative values are down-road of the baseline zone.

After all detections during the session in the selected lane have been processed, the automated correlation and manual verification procedures may then be repeated for another lane for which data is available by re-selection of the Lane Manager window, activated by selection of **Lane Manager** under the **Edit** tab of the main window. The process is repeated until all lanes have been automatically correlated and manually verified, resulting in a complete ground truth data set, and individual-event accuracy records for all detectors in all lanes for the session.

While viewing and/or manually verifying detections, a user may select **Edit->Freeze Frame** to capture and freeze a copy of the current Manual Verification window, as shown in Figure 13.15. Any number of window images may be frozen for future reference, although the computer screen may become cluttered if more that 2 or 3 are captured. A Freeze Frame window is closed by clicking on the **x** in the upper right corner.
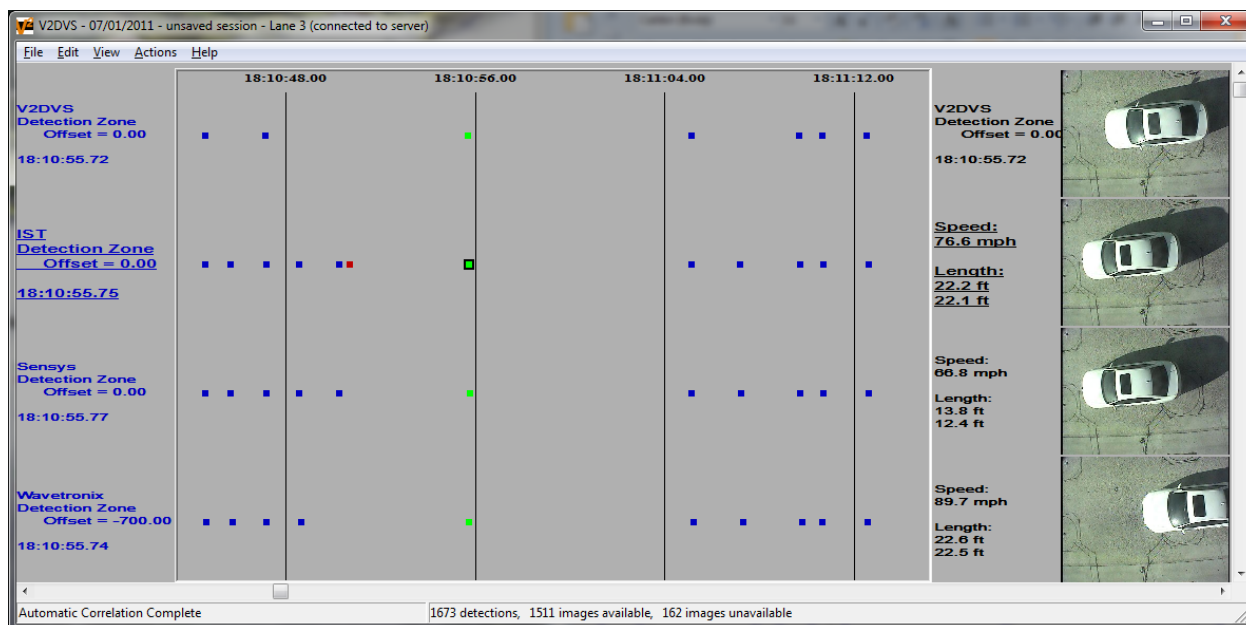
Figure 13.15 Sample frozen image of manual verification window, used for comparisons over spans of time that exceed the timeline width of the window.

13.3.5 Naming and Saving a Session

The session and all verification actions may be saved at any time by selection of either **Save Session As** or **Save Session** under the **File** tab of the main window. Saved sessions may be incompletely verified, and may be restored to permit completion of verification or viewing at a later time. The session is stored as a file with the suffix "**.vrp"** in a user-selected folder, as shown in Figure 13.16.

This file can be opened by the manual verification client at a later point in time. A connection to the server is **not** required to open a previously saved session file.
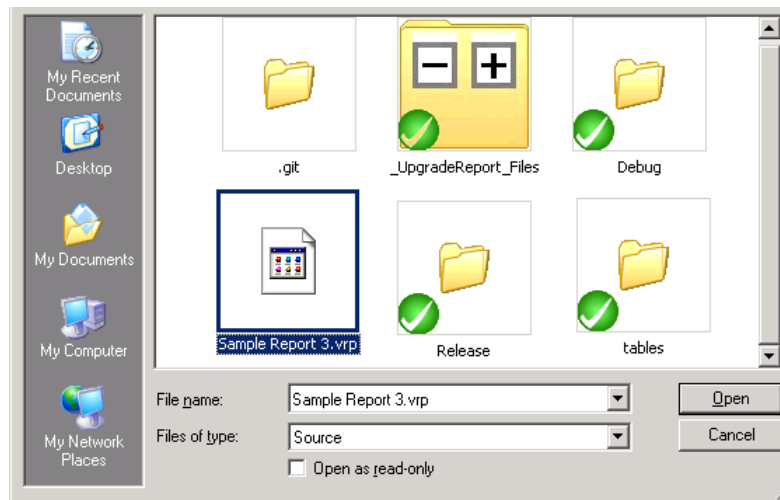
Figure 13.16 Selection window for storing a session as a file for later use.

The previously named "unsaved session" window will now be titled accordingly on the top title bar, as shown in Figure 13.17.
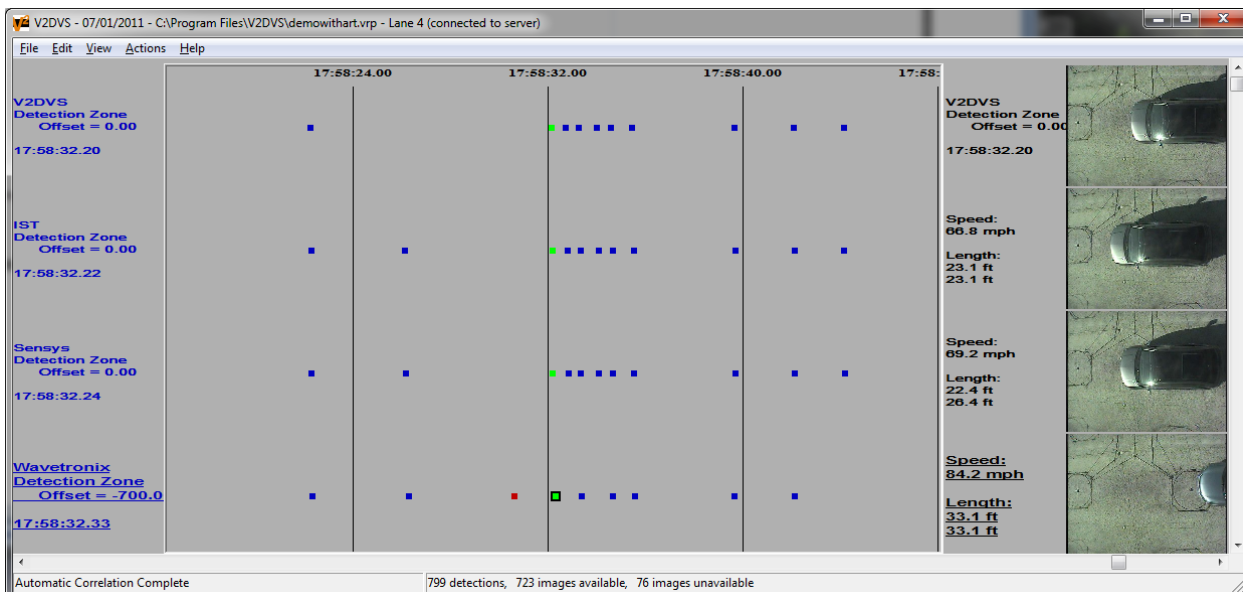
Figure 13.17 Named session, after "Save Session As" is used to designate a session name and store the session data.

The session file may be retrieved by selection of **Continue Previous Verification** under the **Actions** tab of the main window. A standard Windows file selection window will appear, as shown in Figure 13.18.
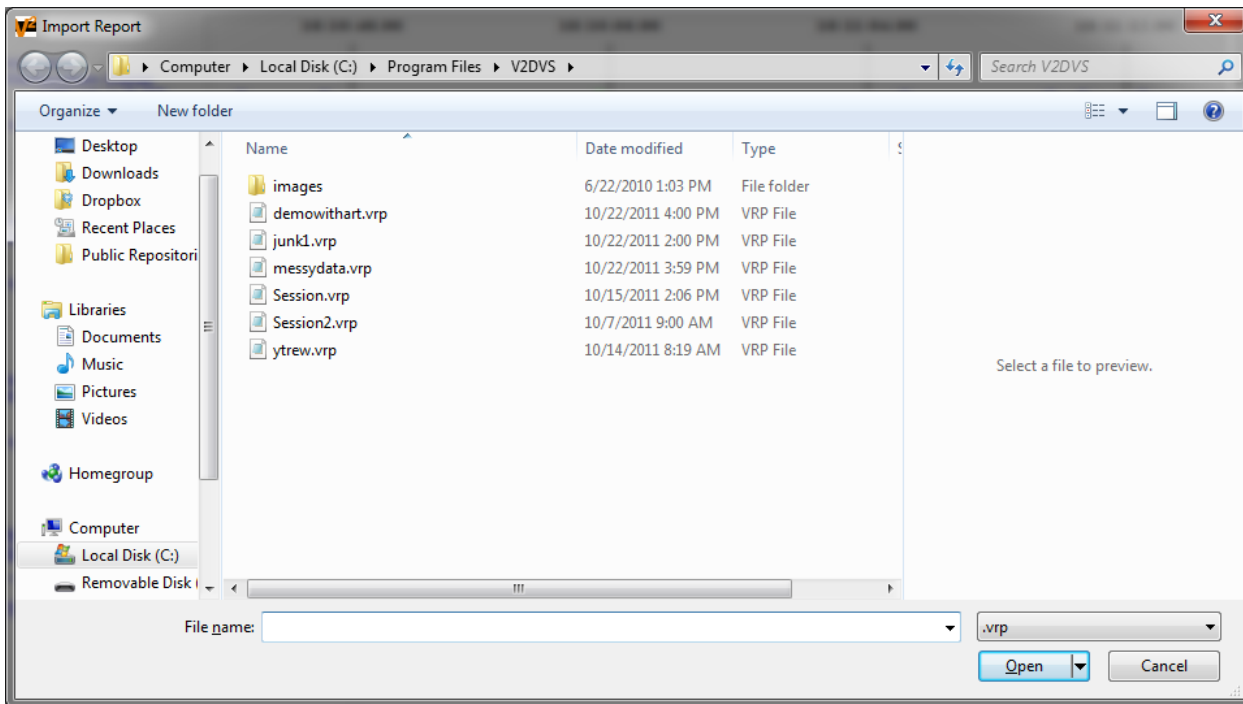
Figure 13.18 Retrieval of previous session report for continued verification work or viewing.

13.3.6 Reporting of Detector Test Results

Finally, the numeric accuracy results for each detector relative to the ground truth data set may be reported. Any number of lanes may be included in the report, from a single lane to all lanes. When multiple lanes are selected, the arithmetic means of the metrics for each lane are displayed. The lanes to include in the Results Report may be selected by clicking on the **Results** button under the **View** pull-down menu of the main window, as shown in Figure 13.19.
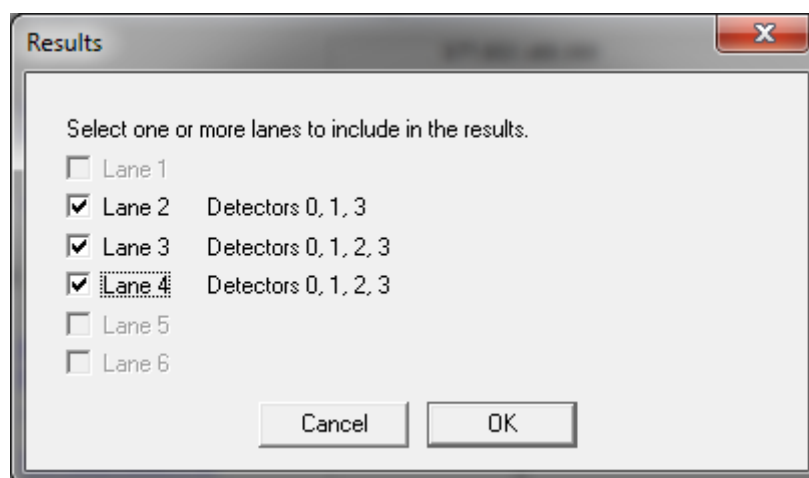
Figure 13.19 Menu for selection of lane(s) to include in the results report.

Once you have selected the lane(s) to include in the report, click **OK** and the results report will be generated and displayed. Figure 13.20 shows a report for a single lane.

Figure 13.20 Results report for a single lanes (in this case, lane 6).

A composite report for lanes 2, 3 and 4 is shown in Figure 13.21.

Figure 13.21 Multi-lane results report (lanes 2,3 and 4).

Either report can be exported and saved as an html file by clicking the **Export as HTML** button at the bottom left corner of the report. The HTML format preserves the report formatting when imported into an MS Word document or displayed on a web page.

13.3.7   Data Deletion

In normal operation, the server continues to accumulate all data and vehicle images from all field machines indefinitely, retiring the oldest data first in the case of exhaustion of the disk space. The server contains 1.2 TBytes of hard disk capacity, which is sufficient for several months of data and image accumulation. However, as a maintenance item, it is advisable to periodically remove old accumulated data and image files from the server. During the process of data reduction and verification, all these data files are copied to one or more client machines. Therefore, they are no longer required on the server. The server response speed will also be increased by deleting unnecessary old records, which reduces the size of the SQL database.

The data on the server (both the images and data stored in the database) along with the images (only) stored on the windows machine running the V2DVS Workstation application can be removed by selecting **Data Deletion** under the **Actions** tab on the main menu bar. Clicking on this tab will open the **Data Deletion Window shown in** Figure 13.22**,** which is very similar to the session data selection window. Data is deleted by selection of a range of dates and times, as with the original session data selection. Data selection does not remove stored sessions (.vrp files) or exported HTML Results Reports, although it does remove the stored images associated with previously stored sessions.

Data and images may also be deleted directly on the server. The procedure is provided in Appendix 2.
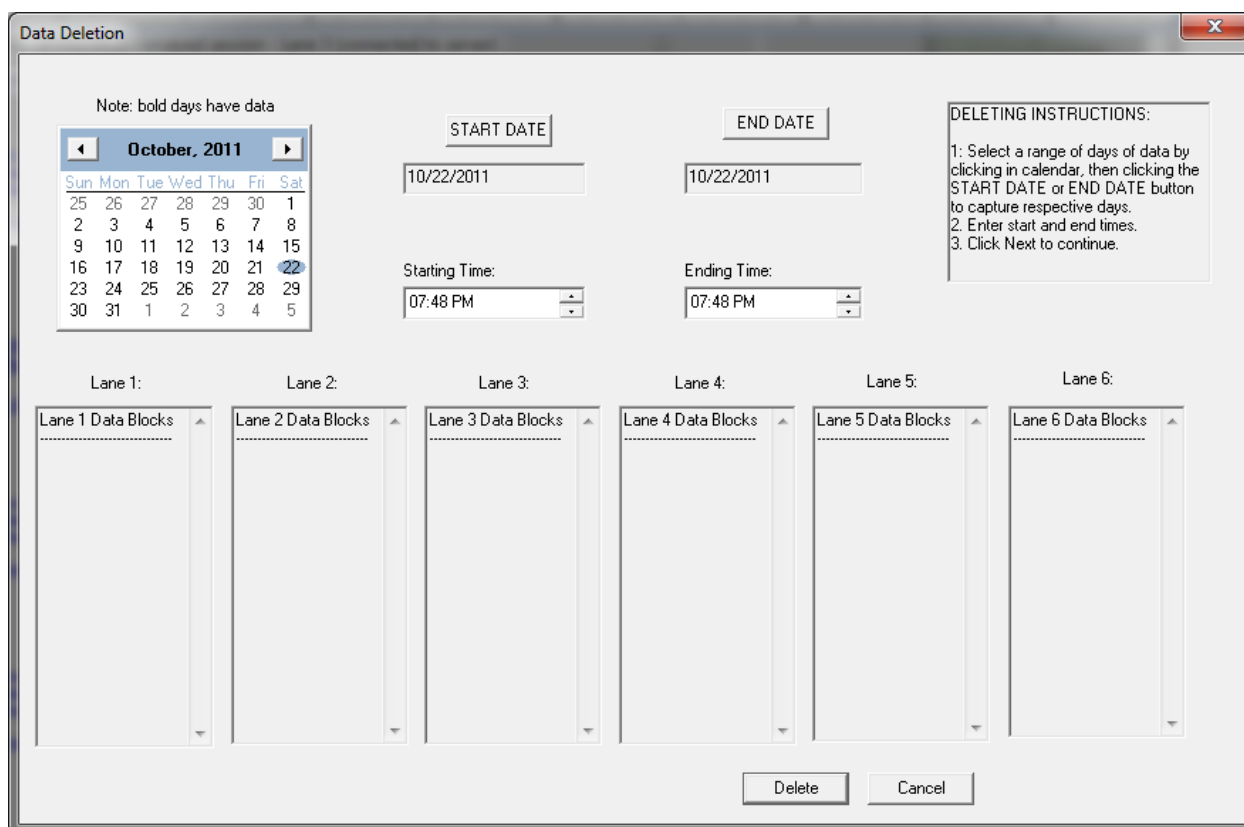
Figure 13.22 Data deletion window, for selection of range of times in the database to be deleted permanently.

13.3.8 Ending the V2DVS Workstation Program

To terminate the V2DVS Workstation program, click **Exit** on the **File** menu, or click the **X** in the upper right corner of the main program window. The program will display a warning message to avoid accidental program termination, as shown in Figure 13.23.
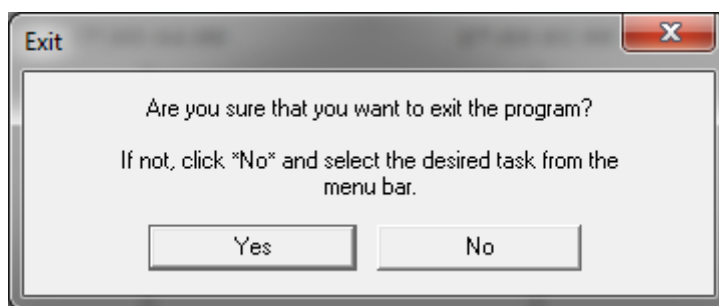
Figure 13.23 Warning message upon program termination.

13.3.9 Uninstallation of the V2DVS Workstation Program

The V2DVS Workstation program may be uninstalled by execution of the included uninstall program accessed through the **Start menu: All Programs->V2DVS Workstation->Uninstall**.   V2DVS does not appear as a program under Windows, so it cannot be uninstalled by the usual Add/Remove software feature on the Control Panel of MS Windows. The uninstallation process will remove all components of the program but will leave all data, images, stored sessions and reports in the V2DVS folder.

14 V2DVS Automated Detector Testing Methodology

14.1 How V2DVS generates ground truth

In traffic monitoring, the notion of "ground truth" has been used loosely to describe estimates of various metrics for traffic characterization. These span the range from simple traffic count and average occupancy, requiring only a vehicle presence sensor such as an inductive loop detector, to metrics derived from comprehensive non-aggregated records of the time of arrival (TOA), speed and length of every vehicle at a fixed location on a highway by detectors with these capabilities. All traffic sensors have error characteristics. This ultimately limits the utility of a single trusted detector as ground truth estimator. This has been recognized formal detector testing standards such as ASTM E 2532 – 06 (3), which uses consider a trusted detector as only an "Accepted Reference Value" to distinguish the data it generates from actual ground truth. Human observation is similarly limited for large traffic volumes and in multi-lane situations, due to both human observation errors and the 16.67 ms NTSC video (field) sample period of video records used to facilitate the analysis. Faced with the challenge of testing and fairly evaluating commercial traffic detectors, it was imperative that we use as a reference a ground truth estimate that was as accurate as possible, and without bias in favor of or against any particular sensing technology.

Reputation theory is a relatively new but well-established field of system science related to optimum decision making based upon information from a collection of sources of information of unknown reliability. The theoretical objective is to establish the "reputation" or accuracy of each information source. As applied to remote sensing, it is known that sensing based upon an adaptive consensus of multiple independent but error-prone observations may be utilized to achieve the most reliable possible estimate of the measurement sought, for example (4). In the V2DVS, we have applied this approach to the determination of traffic ground truth based upon data provided by the complete collection of detectors under test, each with unknown accuracy. In the process, we establish the relative reliability of the information provided by each of the detectors.

In the first version of the V2DVS (Gen-I, presence detector testing only), the ground truth consensus process was based upon fixed "reliability coefficients" associated with each detector under test, with detectors known to be more accurate affecting a larger influence on the consensus results. In further work, an adaptive method has been developed which automatically adjusts these coefficients based upon an evolving metric of reliability associated with each detector, known as a confidence factor or "CF". Reliability, in this case, is continuously assessed by the algorithm based upon agreement with the weighted consensus.

14.2 Vehicle presence detection

During post-processing of acquired detector data for a given period, detection events are considered successively as a forward time sequence. The k^{th} detection event may be initiated by any of the $n = 1,…,N$ detectors under test. The event will correspond either to an actual vehicle or a false detection. For each detection event, the algorithm attempts to correlate a group of temporally-proximate detections by the N detectors under test, with the belief that these represent the same vehicle.

Different detectors have different zones of detection in the same lane, some offset by as much as 300 ft (92 m) down-road from a designated baseline position defined by the first-encountered detection zone. The arrival times of the (assumed) vehicle reported in each of the successive detection zones are adjusted negatively to align them in time as much as possible, based upon knowledge of the each detection zone offset distance (fixed) and the vehicle speed (variable).

$$t_{k,n} = \frac{x_n}{v_{k,n}}$$

where: $t_{k,n}$ = the advance offset time for detection zone *n* during detection event *k*.

x_n = the known offset distance of the detection zone from the baseline.

$v_{k,n}$ = the velocity of vehicle *k*, assumed to be invariant between the multiple detection zones.

n = detector index, $i = 1,2,….N$.

k = detection event (possible vehicle) index. $k = \{1,2,….K\}$

At least one detector must be capable of measuring speed during the event. If the detector associated with a given zone is capable of measuring speed, this speed value is used for the zone offset time correction. Alternatively, the V2DVS reverts to one of several alternative speed measurements sources including duplex inductive loop detectors, and the V2DVS computer vision system itself. If the initiating detection for event *k* is a false detection, no speed will be available from other detectors. For these cases, the speed of the most proximate confirmed vehicle is used.

14.3 Consensus Formation

After appropriate time corrections for each detection, the event is concluded to be an actual vehicle if a sufficient number of detectors, each with a weighted degree of confidence, report a detection within an admissible time aperture, typically 0.5 seconds in duration, starting with the time of initiation of the detection event.

$d_n(k) = \{0,1\}$ is the binary result for detector *n* during any detection event k. 0 means that detector *n* did not report a detection in the time window for this event. 1 means that detector *n* did report a detection in the time window.

$g(k) = $ analog conclusion for k^{th} detection event (possible vehicle).

$G(k) = \{0,1\}$ binary (incorrect or correct) conclusion for k^{th} detection event.

The consensus conclusion $G(k)$ for event k is found by thresholding the weighted numeric average

$$g(k) = \left(\frac{\sum_{n=1}^{N} a_{n,k} d_{n,k}}{\sum_{n=1}^{N} a_{n,k}} \right), \quad \text{where } a_n(k) = \text{confidence factor for detector } n \text{ during detection event } k.$$

$$0 < g(k) < 1 \text{ and } 0 < a(k) < 1.$$

$$G(k) = \begin{cases} 0 & g(k) < \gamma_{lower} \\ \beta_k & \gamma_{lower} \leq g(k) < \gamma_{upper} \\ 1 & g(k) \geq \gamma_{upper} \end{cases} \quad \text{where typically } \gamma_{lower} = 0.48, \gamma_{upper} = 0.52$$

0 means that the consensus for event *k* was that no vehicle was actually present.
1 means that the consensus for event *k* was that a vehicle was actually present

β_k is a binary {0,1} variable assigned when no consensus is formed for event *k*, which remains indeterminate until the event is confirmed or rejected by a human observer using the acquired detection images via the V2DVS Workstation program.

γ_{lower} and γ_{upper} can be thought of as defining acceptable thresholds of belief or disbelief to establish a binary (vehicle or not) consensus for event *k*.

If $G(k) = 1$, a similar weighted consensus is used to determine the estimated TOA $t(k)$ for the confirmed vehicle.

$$t(k) = \left(\frac{\sum_{n=1}^{N} a_{n,k} t_n(k)}{\sum_{n=1}^{N} a_{n,k}} \right)$$

where $t_n(k)$ is the offset-adjusted TOA reported by detector *n* during event *k*. $t(k)$ is undefined if $G(k) = 0$ for the event.

The relative presence detection accuracy of each detector under test is determined from the set of values of $c_n(k)$ and $G(k)$, $k = 1,...K$, for a given data run:

$$\text{Number of correct detections by detector } n = \sum_{k=1}^{K} d_n(k) \bullet G(k)$$

$$\text{Number of failures to detection by detector } n = \sum_{k=1}^{K} \overline{d_n(k)} \bullet G(k)$$

$$\text{Number of false detections by detector } n = \sum_{k=1}^{K} d_n(k) \bullet \overline{G(k)}$$

Where $\bullet$ is the logical AND function and the overbar indicates the logical inverse. Results are usually reported in normalized form (as a percentage) by division by the total number of confirmed vehicles $\sum_{k=1}^{K} G(k)$.

These values are adjusted under special conditions. Multiple detections by the same detector within the same non-overlapping aperture are considered mutually exclusive, with only one detection being correct and the others false detections. Proximity to the consensus-derived center time is used to discriminate the single valid detection among the multiple instances. A similar resolution method is used when the admissibility apertures for successive events overlap, possible in the case of very-closely-following vehicles. The occurrence of this scenario is minimized by the use of a relatively small admissibility time aperture, typically 500 ms or less.

14.4 Vehicle Speed and Length Ground Truth Estimation

New to the Generation II version of V2DVS is the determination of the accuracy of the speed and length measurements of detectors which are capable of generating this information.

For events k resulting in $G(k) = 1$ (a positive vehicle presence conclusion), the vehicle speed and length reported by detectors in agreement with the presence conclusion are used to establish analog estimates of these ground truth values:

$$s(k) = \left(\frac{\sum_{n=1}^{N} a_{n,k} s_n(k)}{\sum_{n=1}^{N} a_{n,k}} \right) \quad \text{where } s_n(k) = \text{vehicle speed reported by detector n for event } k \, ..$$

$$l(k) = \left(\frac{\sum_{n=1}^{N} a_{n,k} l_n(k)}{\sum_{n=1}^{N} a_{n,k}} \right) \quad \text{where } l_n(k) = \text{vehicle length reported by detector n for event } k \, .$$

$s(k)$ and $l(k)$ are undefined for values of k for which $G(k) = 0$, indicating a non-existent vehicle.

In the current system embodiment, the same confidence factors used for presence estimation are used for speed and length estimation above, although we recognize that separately-adapted confidence factors for each of these measurements could possibly yield superior accuracy, since the accuracy of the speed or length measurement ability of a detector is likely independent of its ability to determine vehicle presence.

As with presence, the speed and length measurements of the individual detectors under test are compared with the estimated ground truth record of speed and length, and the cumulative difference over all vehicles is reported as the average absolute error or mean square error.

14.5 Adaptation of Confidence Factors for Weighted Consensus

Adaptive estimation based upon multiple independent observers of questionable accuracy has been utilized in many applications to achieve the most reliable possible estimate of the measurement sought, for example (5,6,7,8). The initial ground truth estimation process used in the Test Bed was based upon fixed confidence factors in the formation of the weighted average. These values were determined intuitively, which allowed a possible inadvertent bias in the ground truth estimation. Improving upon this, confidence factor adaptation based upon agreement or disagreement with the consensus is currently being used.

$c_n(k) = d_n(k) \otimes G(k)$ is a binary $\{0,1\}$ "agreement variable" for detector n and event k. $\otimes$ is the logical XNOR function. $c_n(k) = 1$ only if detector n agrees with the consensus conclusion for event k.

A large number of learning adaptation methods are possible, but a simple first order autoregressive filter has been found to be adequate for the present ground truth estimation purpose, gradually adjusting the confidence factor $a_n(k)$ for detector n based upon successive agreement or disagreement with the consensus $G(k)$ for each detection event $k = 1,\dots K$.

$$a_n(k+1) = ba_n(k)(b) + (1-b)c_n(k)$$ where b is the filter eigenvalue, typically = 0.95 .

For each event k, agreement of detector n with the weighted consensus slightly increases a_n toward unity, while disagreement slightly decreases it toward zero. For consistent traffic conditions and uniform distributions of detector errors, the final value $a_n(k = K)$ converges to the cumulative ratio of the number of agreements with the consensus to the total number of detection events K, but is subject to a small sample-to-sample noise component.

$$\lim_{k \to K} a_n(k) = \frac{\sum_{k=1}^{K} c_n(k)}{K}$$

If every detection by detector n agrees with the consensus, $a_n(k)$ asymptotically approaches unity, while if the detector consistently disagrees with the consensus, $a_n(k)$ asymptotically approaches zero. Consequently, the contribution to the consensus is gradually diminished for consistently inaccurate detectors, and enhanced for consistently accurate detectors.

As with most iterative optimization strategies, the values of the confidence factors are susceptible to seeking turning points which could represent either maximum accuracy or maximum inaccuracy, dependent upon the initial values of the confidence factors relative to the accuracy of the corresponding detectors. Our relatively simple adaptive approach assumes that the detectors assigned higher initial confidence factors will be at least initially correct most of the time. Failing this assumption admits the possibility that the initial consensus decision is incorrect, and successive decisions will perpetuate the incorrect conclusions regarding the presence of non-presence of an actual vehicle. The convergence behavior of the algorithm is therefore potentially unstable. A reasonable set of initial values for the detector confidence factors based upon a priori general knowledge of relative detector accuracy is therefore important to assure optimization of accuracy rather than optimization of inaccuracy. In all tests reported herein, all detectors start with equal confidence factors $a_n(0) = 0.5$.

Optimization of maximum inaccuracy can occur only if the majority of the detectors under test have error rates exceeding 50%, and they are highly weighted initially, an avoidable condition.

Convergence rate is a related primarily to the filter eigenvalue b, with larger values (approaching unity) implying longer filter time constants which lead to longer convergence times. The suggested value (b = 0.95) is not ideal for rapid convergence, but has been found an acceptable tradeoff between noise and convergence rate. In actual practice, an acceptable level of convergence is usually achieved after between 50 and 100 detection events.

Since the optimization process is performed entirely during post-processing of existing datasets, it is possible to do a first "pass" on the dataset to optimize the detector confidence factors, then a second pass to generate the actual ground truth record. Successive iterations tend to generate better results, as the confidence factors begin as stable values, and are then further refined. For longer datasets during which traffic or detection

conditions (such a illumination for video detectors) gradually change, it is intuitively preferable that the confidence factors be allowed to continuously change over the course of the data run. For this reason, we elected to conduct all tests reported herein as single-pass post-processing runs, accepting the increased probability of errors in ground truth estimation early in the data run.

It should be noted that the final value of the confidence factor $a_n(k = K)$ for each detector is a normalized indicator of detector accuracy, with higher values indicating a more accurate/reliable detector.

15 An Assessment of V2DVS ground truth estimation accuracy using real and synthetic data

The performance of the ground truth estimation algorithm may be assessed using several metrics, two of particular relevance for large-scale detector testing:

1. The ability of the algorithm to classify as many detection events as possible, leaving a minimum number for human verification.
2. The accuracy of the resultant estimated ground truth.

While metric 1 can be observed readily using actual traffic data, metric 2 cannot. This is because any real traffic dataset must be generated by some detection mechanism, and cannot be considered ground truth, but only a set of *accepted reference values*. For an absolute accuracy assessment over large datasets, it is necessary to synthetically generate a ground truth dataset with a priori known characteristics, and the generate multiple "imperfect" variations of this set representative of the outputs of typical detectors we have encountered in field testing, each with specified error rates. The synthetic ground truth generator mimics actual car following and platoon behavior under various levels of service by taking as input parameters the minimum vehicle length, maximum vehicle length, minimum speed, maximum speed, tailgate probability, safe-following probability, long gap probability, minimum tailgate following distance, maximum tailgate following distance, minimum safe following distance, maximum safe following distance, minimum long gap following distance, maximum long gap following distance, and either the total number of vehicles or the start/stop data collection times for the dataset.

Synthetically-generated detector data starts with a synthetic ground truth data set, and introduces errors of the following types:

1. Elimination of a specified percentage of vehicles (failures to detect)
2. Introduction of a specified number of additional vehicles (false detections)
3. Variation of vehicle arrival times over a specified range and statistical distribution
4. Variation of vehicle speed over a specified range and statistical distribution
5. Variation of vehicle length with a specified offset, range and statistical distribution

Type 2 errors (false detections) are injected into the dataset at uniformly random positions between existing vehicles. Type 3 errors (speed variation) consider the gap between successive vehicles in allowing speed variations, since proximate vehicles travel at nearly identical speeds.

Four ground truth data sets, each containing 1000 vehicles were used as the root for the generation of 16 different synthetic detector data sets, with characteristics shown in Table 15.1.

Table 15.1 Synthetic Detector Data Characteristics.

Synthetic Detector No.	Failures to detect (%)	False detections (%)	Root ground truth set
1	1	1	1
2	1	10	1
3	10	1	1
4	5	5	1
5	10	10	2
6	20	5	1
7	5	20	1
8	1	20	1
9	20	1	1
10	0	50	3
11	50	0	3
12	40	10	3
13	10	40	3
14	40	20	3
15	20	40	3
16**	100	100	4

* Uniform statistical distributions were used for variations over the specified range.
** Synthetic detector No. 16 was actual Testbed data from a duplex inductive loop detector located in a lane other than the one observed by all other detectors. This uncorrelated data mimics a common wiring error for field-installed inductive loop detectors.

As a reality check for the synthetic data, we examined and present in Table 15.2 the observed behavior of four actual detectors under test over a 30 minute period, approximately 1000 vehicles, under typical traffic conditions.

Table 15.2 Error Characteristics for Four Detection Technologies, Lane 3, Mid-day, Full Sun

Detector Technology	Average Absolute TOA error (ms)	% failures to detect	% false detections
Video processor	37.3	6%	1%
Inductive loop	11.8	<1%	1%
Magnetometer	24.5	1%	<1%
RADAR	45.6	5%	2%

Note that the ground truth reference dataset used for actual detector data is the estimated ground truth generated by the adaptive consensus algorithm, not an absolute ground truth record as used in the synthetic detector tests. It is therefore not possible to assess the accuracy of the ground truth estimation when using real data, only the accuracy of the detectors under test, assuming the ground truth estimate to be reliable. The purpose of this sample is to suggest that most detectors exhibit fail and false detection error rates well within the bounds of the majority of the synthetic detector data used for the testing the consensus-based ground truth estimation algorithm, as specified in Table 15.1.

The test methodology involved using collections of synthetic detectors "under test" with various degrees of accuracy as inputs to the adaptive consensus algorithms for presence, speed and length, and then comparing the ground truth estimates produced by the algorithms with the original ground truth dataset from which the synthetic detector data were derived. Table 15.3 presents the results of nine test using various combinations of the synthetic detectors, and three TOA variations applied to each detector dataset. Each consisted of 1000 vehicles. Speed errors introduced into all detector datasets were uniformly distributed over a ± 16 mph (± 25.6 kph) range, subject to car following limitations. For all tests, the admissible time window for correlation of the offset-adjusted detection times was 500 ms. Length errors introduced into all detector datasets were uniformly distributed over a skewed ± 0.83 ft (± 0.25 m) range, but since these have no effect on the weighted consensus

algorithms for vehicle presence, length-related results are not reported. All results were reported *prior to* any human resolution of unverified detection events (manual verification).

Table 15.4 presents the results of an examination of the possible improvement in the accuracy of the ground truth estimate if manual verification of the unverified detections is performed following generation of the consensus-estimated ground truth. Table 15.5 focuses on the corresponding accuracy results for each of the synthetic detectors, that is, the ability of the V2DVS to correctly identify the number of correct detections, failures to detect and false detections for each of the detectors under test.

Table 15.3 Performance of Adaptive Consensus Algorithm for Synthetic Detector Data

Test No.	Detector No. or Estimated GT	TOA variation range (+/- ms)	Correctly identified vehicles	Incorrectly Identified Vehicles	Failures to identify a vehicle	Est mean vehicle speed (mph)	Error in est mean speed (mph)	Final value of confidence factor	Detections marked by V2DVS as unverified
1	Est GT		1000	0	0	66.7	0.1		
	1	100						0.96	0
	1	100						1.00	0
	1	100						0.99	0
	1	100						0.95	0
	1	100						0.95	0
2	Est GT		999	0	1	66.8	0.0		
	4	100						0.91	0
	4	100						0.96	0
	4	100						0.93	0
	4	100						0.92	0
	4	100						0.90	0
3	Est GT		1000	0	0	66.6	0.2		
	1	50						0.97	0
	1	50						1.00	0
	1	50						0.99	0
	1	50						0.96	0
	1	50						0.97	0
4	Est GT		996	0	4	66.7	0.1		
	4	50						0.89	0
	4	50						0.91	0
	4	50						0.97	0
	4	50						0.84	0
	4	50						0.91	0
5	Est GT		1000	0	0	66.8	0.0		
	1	25						0.99	0
	1	25						0.94	0
	1	25						0.98	0
	1	25						1.00	0
	1	25						0.99	0
6	Est GT		999	0	1	66.8	0.0		
	4	25						0.97	0
	4	25						0.98	0
	4	25						0.98	0
	4	25						0.86	0
	4	25						0.91	0
7b	Est GT		992	0	8	66.6	0.2		
	1	100						0.88	4
	2	100						0.94	4
	3	100						0.92	3
	10	100						0.72	2
	14	100						0.53	1
8b	Est GT		994	0	6	67.0	-0.2		
	1	100						0.85	5
	2	100						0.94	5
	3	100						0.92	2
	11	100						0.61	2
	13	100						0.74	4
9b	Est GT		986	0	14	66.7	0.1		
	1	100						0.90	8
	2	100						0.83	10
	3	100						0.89	3
	4	100						0.96	9
	16	100						0.13	7

Table 15.4 Relative Improvement of Ground Truth Estimate Following Manual Verification of Detection Events Found to be Indeterminate by the Consensus Algorithm ($0.48 < \beta < 0.52$)

Test No.	Detector	Before or after resolution of unverified detections	Correctly identified vehicles	Failures to identify a vehicle	Final value of confidence factor	Detections marked by V2DVS as unverified
7b	Est GT	before	992	8		
	1				0.88	4
	2				0.94	4
	3				0.92	3
	10				0.72	2
	14				0.53	1
7a	Est GT	after	997	3		
	1				0.88	0
	2				0.94	0
	3				0.92	0
	10				0.72	0
	14				0.53	0
8b	Est GT	before	994	6		
	1				0.85	5
	2				0.94	5
	3				0.92	2
	11				0.61	2
	13				0.74	4
8a	Est GT	after	999	1		
	1				0.85	0
	2				0.94	0
	3				0.90	0
	11				0.59	0
	13				0.72	0
9b	Est GT	before	986	14		
	1				0.90	8
	2				0.83	10
	3				0.89	3
	4				0.96	9
	16				0.13	7
9a	Est GT	after	996	4		
	1				0.90	0
	2				0.83	0
	3				0.89	0
	4				0.96	0
	16				0.13	0

Table 15.5 Assessment of the Ability of V2DVS to Correctly Identify the Number of Correct Detections, Failures to detect, False detections and Speed Measurements of Each Detector Under Test

Test No.	Detector	Final value of confidence factor	Cars in detector data	Correct detections as reported by V2DVS	V2DVS errors in correct detection results	Missing cars in detector data	Failures to detect reported by V2DVS	V2DVS errors in failure-to-detect results	False detections in detector data	False detections reported by V2DVS	V2DVS errors in false detection results	Mean Av Abs Speed Error in detector data (mph)	Mean Av Abs Speed Error reported by V2DVS (mph)	V2DVS error in detector speed results (mph)
1	1	0.96	986	986	0	14	14	0	11	11	0	8.37	7.32	1.05
	1	1.00	990	990	0	10	10	0	9	9	0	8.01	6.95	1.06
	1	0.99	991	991	0	9	9	0	11	11	0	8.08	7.01	1.07
	1	0.95	993	993	0	7	7	0	12	12	0	8.08	6.99	1.09
	1	0.95	992	992	0	8	8	0	10	10	0	8.06	7.21	0.85
2	4	0.91	949	948	1	51	51	0	44	45	-1	8.18	7.15	1.03
	4	0.96	954	949	5	46	50	-4	55	60	-5	8.05	7.04	1.01
	4	0.93	957	957	0	43	42	1	56	56	0	8.31	7.43	0.88
	4	0.92	955	954	1	45	45	0	48	49	-1	8.18	6.95	1.23
	4	0.90	953	952	1	47	47	0	56	57	-1	8.04	7.06	0.98
3	1	0.97	989	988	1	11	12	-1	8	9	-1	8.15	7.17	0.98
	1	1.00	987	987	0	13	13	0	14	14	0	7.97	6.83	1.14
	1	0.99	995	994	1	5	6	-1	6	7	-1	8.12	7.00	1.12
	1	0.96	992	992	0	8	8	0	16	16	0	7.88	6.98	0.90
	1	0.97	990	990	0	10	10	0	8	8	0	8.06	6.89	1.17
4	4	0.89	958	955	3	42	41	1	43	46	-3	7.70	6.65	1.05
	4	0.91	937	936	1	63	60	3	50	51	-1	8.22	6.80	1.42
	4	0.97	954	951	3	46	45	1	47	50	-3	8.13	7.13	1.00
	4	0.84	948	949	-1	52	47	5	48	47	1	7.92	6.95	0.97
	4	0.91	944	943	1	56	53	3	54	55	-1	8.31	7.00	1.31
5	1	0.99	992	992	0	8	8	0	11	11	0	8.21	7.22	0.99
	1	0.94	991	991	0	9	9	0	11	11	0	8.03	7.08	0.95
	1	0.98	991	991	0	9	9	0	10	10	0	7.96	6.92	1.04
	1	1.00	991	991	0	9	9	0	9	9	0	8.29	7.01	1.28
	1	0.99	994	994	0	6	6	0	8	8	0	7.98	7.00	0.98
6	4	0.97	945	945	0	55	54	1	49	49	0	7.81	6.87	0.94
	4	0.98	951	952	-1	49	47	2	42	41	1	7.96	6.79	1.17
	4	0.98	940	941	-1	60	58	2	55	54	1	8.09	6.79	1.30
	4	0.86	957	958	-1	43	41	2	55	54	1	8.02	6.87	1.15
	4	0.91	945	945	0	55	54	1	42	42	0	8.08	6.81	1.27
7	1	0.88	988	975	13	12	17	-5	7	16	-9	8.15	6.87	1.28
	2	0.94	987	972	15	13	20	-7	97	108	-11	8.17	7.19	0.98
	3	0.92	905	895	10	95	97	-2	11	18	-7	7.87	6.47	1.40
	10	0.72	1000	987	13	0	5	-5	498	509	-11	8.41	7.28	1.13
	14	0.53	560	576	-16	440	416	24	201	184	17	8.11	7.27	0.84
8	1	0.85	988	978	10	12	16	-4	15	20	-5	7.93	6.77	1.16
	2	0.94	988	978	10	12	16	-4	94	99	-5	8.01	6.83	1.18
	3	0.92	909	901	8	91	93	-2	14	20	-6	8.05	6.83	1.22
	11	0.61	497	491	6	503	503	0	0	4	-4	8.33	7.36	0.97
	13	0.74	915	916	-1	85	78	7	371	366	5	8.09	6.90	1.19
9	1	0.90	991	972	19	9	14	-5	6	17	-11	8.15	6.58	1.57
	2	0.83	991	964	27	9	22	-13	99	116	-17	7.77	6.42	1.35
	3	0.89	917	906	11	83	80	3	11	19	-8	7.87	6.59	1.28
	4	0.96	936	921	15	64	65	-1	38	44	-6	8.21	6.89	1.32
	16	0.13	0	237	-237	1000	749	251	1000	979	21	0.00	15.12	-15.12

In preliminary testing of the Gen I V2DVS using eight detectors under test and 1600 actual vehicles, approximately 98.5% of detection events reported by the detectors under test were identified as correct, false, or failures-to-detect by the consensus algorithm, leaving only 1.5% unverified requiring manual verification. This represented an improvement over the use of fixed (assumed) confidence factors for each detector, which achieved 97.0% event identification.

Using synthetic absolute ground truth and detector data, the accuracy of the ground truth estimate is examined. Among other observations, the nine tests represented in Table 15.3 reveal the effects of the quality of the detectors under test on the quality of the consensus-based ground truth estimate. Note that the sum of the "Correctly identified vehicles" and "Failures to identify a vehicle" are always equal to the total number of actual vehicles in absolute ground truth dataset. For Tests 1, 3 and 5, each based on five detectors all with 1% fail and 1% false rates, the ground truth estimate is identical to the absolute ground truth, with respect to all metrics: all 1000 vehicle were correctly identified in each test. The mean speed estimation errors in each of these tests were very low, with Test 2 showing the largest error at 0.2 mph. At the other extreme are tests 7, 8 and 9, which contained detectors 10, 13, 14 and 16, each with errors rates that equal or exceed their correct detection rates.

However, even in these cases the number of correctly identified vehicles remained high: 992, 994 and 986 respectively. Test 9 included detector 16 for which 100% of the vehicles were not detected, and 100% of the reported detections were false. The apparent robust performance of the consensus algorithm in these cases can be explained by the fact that in each test, detectors 1,2 and 3 were included, each with low errors rates (see Table 15.1). Intuitively, as long as a simple majority of the detectors under test have error rates (failures to detect in particular), that do not exceed the number of correct detections, the estimated ground truth dataset will be very accurate (98+% in all tests), since the confidence factors for the minority "bad" detectors will eventually be reduced to a point where their influence on the consensus is minimal. In Test 9, the final value of the confidence factor for detector 16 was 0.13.

With the admissible time window equal to 500 ms in all tests, the algorithm is conservative in identification of vehicles. This is apparent from observation of the zero number of incorrectly identified vehicles in the estimated ground truth datasets for all tests.

Random variations in the time of arrival reported by each detector have little effect on the accuracy of the estimated ground truth dataset, as long as the variation ranges remain less than the duration of the admissible time window. The ± 100 ms (200 ms total) variation range used for the majority of the tests was well within the 500 ms window.

Despite the large introduced speed errors (± 16 mph, ± 25.6 kph) in all detector data sets, the average absolute speed error in the estimated ground truth remained very low, never exceeding 0.2 mph (0.32 kph). This is understandable since the random speed errors for each detector were independent and symmetric about the mean.

The number of "failures to identify" in the estimated ground truth can be partially reduced in the optional manual verification process, in which a human resolves the unverified detections for each detector by inspecting the image acquired at the time of detection using the graphical tools provided in the V2DVS workstation program. This is illustrated in Table 15.4, in which the least-accurate ground truth estimates from Tests 7, 8 and 9 are displayed again as "before" or "after" manual verification. The number of unverified events for each detector are shown in the rightmost column, "Detections marked by V2DVS as unverified". For Test 9, human resolution of the designated unverified events results in an improvement in correctly identified vehicles from 986 to 996 out of the 1000 total.

After generation of the ground truth estimate, the detectors under test are compared with this set with respect to several metrics, shown in Table 15.5. The columns are labeled to be self-explanatory. Of ultimate concern are the errors of the V2DVS system in its assessment of detectors, indicated by columns headed "V2DVS errors in…" correction detection results, failure to detect results, false detection results, speed results. V2DVS errors in assessment results are very small for collections of reasonably accurate detectors under test, but error rate increase dramatically in the results for very poor detectors. The extreme case of Detector 16 in Test 9, for which every detection was incorrect by design, yielded the highest rate of errors in the corresponding V2DVS test results. This is due to both to increased errors in the estimated ground truth, as well as the extreme characteristics of the completely uncorrelated detections from this device. It is clear that V2DVS errors in

assessing a detector increase with the error rate of the detector itself for all metrics. Note that this leads to a *less severe* conclusion regarding the performance of the detector, leaving many errors unaccounted for.

As previously mentioned, the final value of the confidence factor $a_n(k = K)$ for each detector may also serve as a normalized indicator of its presence detection reliability.

Does the consensus filter yield a more accurate ground truth estimate than a single highly accurate trusted detector? Examining Table 15.3 aware of the designed error rate of each detector as specified in Table 15.1, we first consider Tests 1-6, consisting entirely of sets of detector with either 1% or 5% error rates. In all these tests, the algorithm correctly identified 99.9+% of the actual vehicles, and had no false identifications. Detector 1 data, by design, had an average correct detection rate of 99% and false detection rate of 1%. Detector 4 data, by design, had an average correct detection rate of 95% and false detection rate of 5%. For Tests 7-9, each containing at least one detector of very poor accuracy, the algorithm identified between 99.2% and 98.6% of the actual vehicles correctly, and again had no false identifications. It can be concluded that while the accuracy of the ground truth estimate does degrade somewhat as the accuracy of the detectors under test are reduced, the consensus filter almost always generates more reliable ground truth data than the best of the detectors under test. This is true even if some of the detectors are extremely inaccurate, as long as they do not constitute the majority of the test set.

16 Detector Signaling Details

16.1 How detectors communicate with the V2DVS Field Machines

V2DVS will accept up to thirty-two binary inputs *per lane* for real-time hardware signaling (total of 192 binary inputs across six lanes). Each binary input corresponds to a specific lane and a specific line input. Each line could correspond to a different detection zone within that lane, since a different detection position may be specified for each detector in its setup configuration. Pairs of binary inputs can be used to derive speed and length of a vehicle using the same detector type on equally space detection zones. Signaling is active low edge detection. The hardware interface is open collector ground-referenced, which is compatible with both 0-5V TTL and 0-24V typical of Type 222 detector card outputs. As previously discussed, the "detectors.conf" file on each V2DVS Field machine allows for configuration of detectors to particular terminals of the hardware interface, as previously discussed.

Each V2DVS computer uses its parallel port to receive coded data from the detector interface boxes representing the electrical state of the output of each detector. The core of the interface box is a Complex Programmable Logic Device (CPLD) that processes 32 inputs into an 8-bit binary byte, and provides a 2 ms interrupt to signal the field machine of the presence of a new detection "edge" or state change on any input.

A low-level driver program running on the V2DVS Field Machine services the interrupt by addressing the thirty two binary inputs eight lines at a time. The interface boxes are the gray boxes shown previously in Figure 9.5, attached to the rear door of Cabinet 2. The interface CPLD logical block diagram is shown in Figure 16.1. The schematics for the detector interface box circuits are shown in Figure 16.2 and Figure 16.3 . All inputs go through a de-bouncing and high-voltage protection circuit before reaching the CPLD, for protection and high frequency noise immunity.

The VHDL-coded logic in the CPLD works as follows: Data from any of the thirty-two detector inputs to the box is latched after going through a de-bounce circuit. If the data coming out of the de-bounce circuit is not equal to the latched data, then a change of state has occurred on one of the thirty-two inputs and an interrupt is generated. The PC acknowledges the interrupt by reading in the byte representing the current state of all the detectors via the parallel port. If multiple detector lines change concurrently (within a maximum of 2 microseconds of each other), a FIFO of state bytes is maintained and streamed to the Field Machine in line-order succession by the field machine. No data is lost, and maximum detection time error of 64 microseconds is possible (if theoretically all detectors triggered at exactly the same time). The maximum *relative* event time inaccuracy for measurements made by the same field machine is less than 10 microseconds, which is well below the detection data time stamp resolution of 10,000 microseconds (10 ms). The absolute time inaccuracy can be much larger if the computers are not synchronized.

Time stamping of all data is done by the Linux field machine, which, if the feature is enabled, can periodically synchronize its system time clock with that of the V2DVS server. The server has the capability to synchronize over the Internet with the National Bureau of Standards Timer Server, although lacking an external Internet connection to the TestBed, this feature is not currently activated.
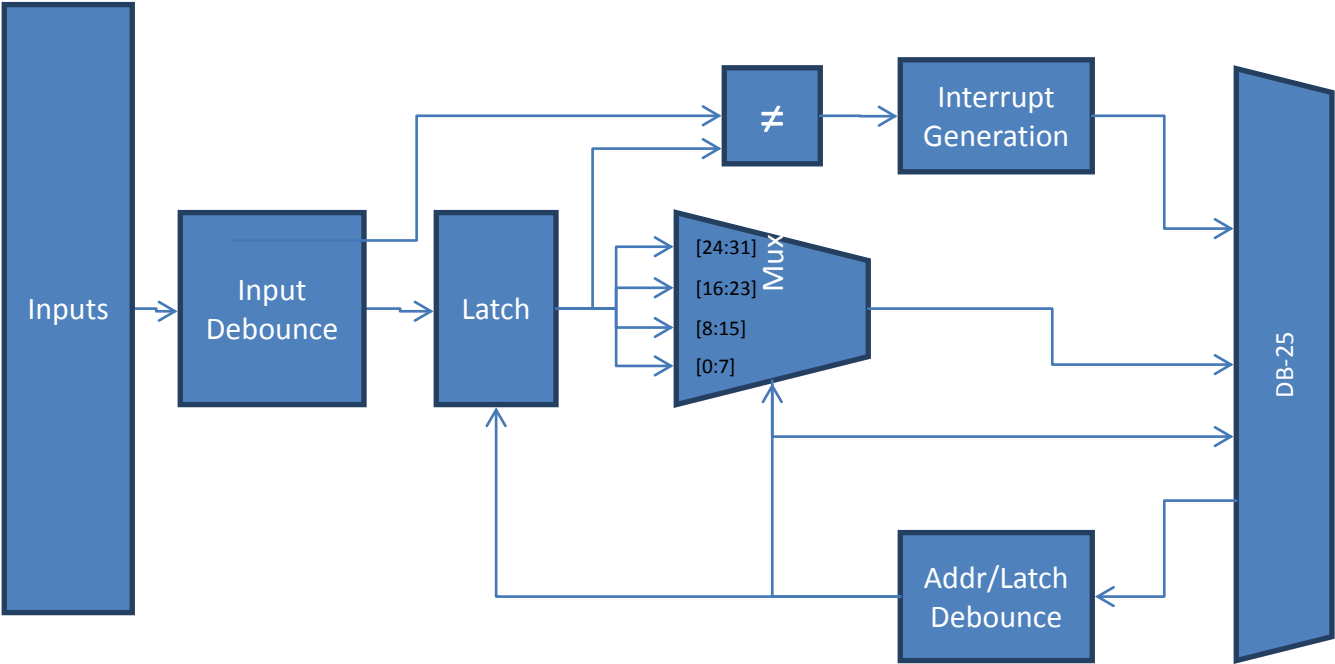
Figure 16.1 Block diagram for CPLD-based detector interface logic.
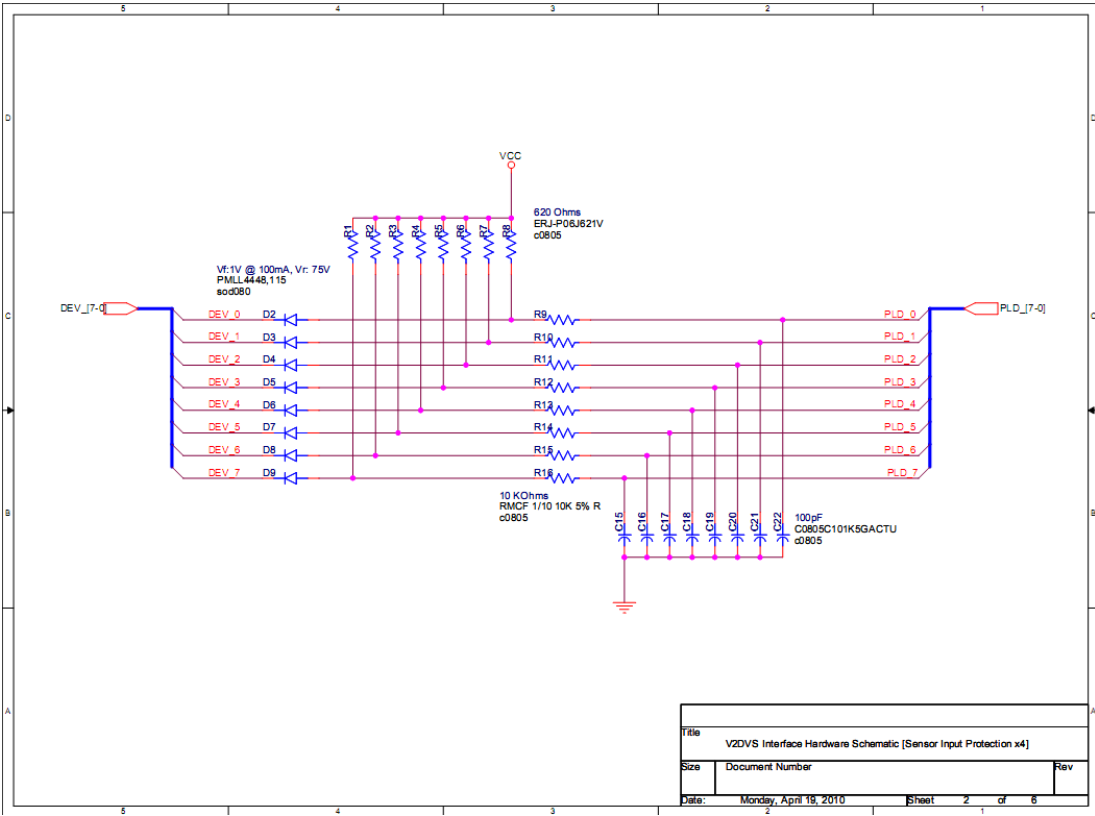
Figure 16.2   High-voltage input protection and de-bounce circuit schematic.

Figure 16.3 Schematic for contact-closure interface box (one for each V2DVS field machine).

Figure 16.4 Photo of PCB inside each detector interface box.

16.2 Detector Hardware Connections

Each detector interface box receives inputs from four sets of standard 8-conductor Cat-5 cables, for a total of 32 inputs. These cables accommodate up to 32 input individual (simplex) detectors, or 16 duplex (pairs of) detectors. Each cable is labeled with both a cable number and a matching number of stripes, as indexed in Table 16.1 and shown in Figure 16.5.

Table 16.1   Cable lead assignment between detector interface box and detector terminal block.

Line #	Ethernet Cable	Wire Color
1	0	Green
2	0	Green/White
3	0	Orange
4	0	Orange/White
5	0	Brown
6	0	Brown/White
7	0	Blue
8	0	Blue/White
9	1	Green
10	1	Green/White
11	1	Orange
12	1	Orange/White
13	1	Brown
14	1	Brown/White
15	1	Blue
16	1	Blue/White
17	2	Green
18	2	Green/White
19	2	Orange
20	2	Orange/White
21	2	Brown
22	2	Brown/White
23	2	Blue
24	2	Blue/White
25	3	Green
26	3	Green/White
27	3	Orange
28	3	Orange/White
29	3	Brown
30	3	Brown/White
31	3	Blue
32	3	Blue/White

Figure 16.5   Four cables with a total of 32 individual wires provide the detector connections to each interface box.

17 Robotic Lens Control Hardware and Firmware

The robotic lenses for each video camera are controlled by a modified 68HC11EVB microcontroller board housed inside each Field Machine case. These boards include power transistors to provide positive or negative control voltages for each of the primary lens control actions: lens focus, lens zoom, and lens aperture (iris). The boards accept as their inputs serial data commands from the serial port of the Field Machine, and translate specific ASCII characters into control actions which power the lens control inputs via a long 8-conductor cable connecting each camera/lens on the bridge deck to its corresponding Field Machine. The source code running on each microcontroller is included in the accompanying CD delivered at the end of the project.
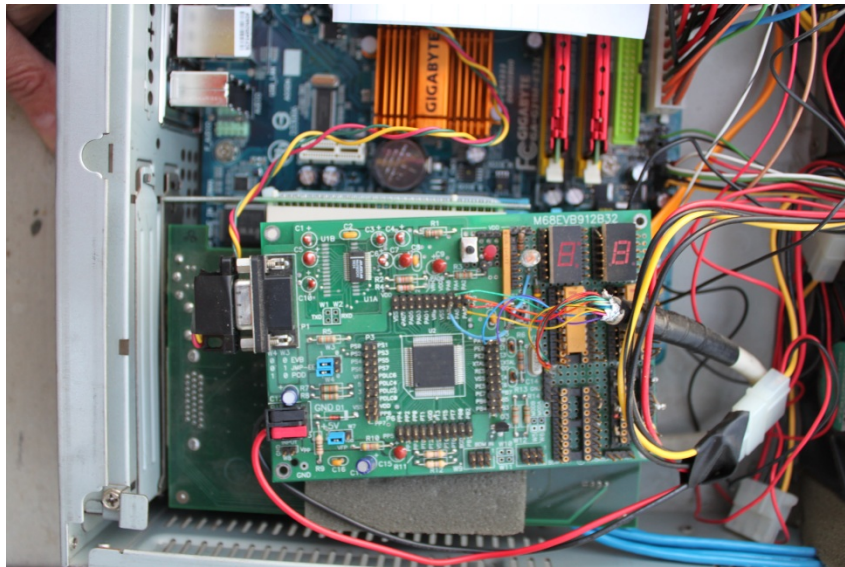
**Figure 17.1 68HC11-based microcontroller lens control board,
installed inside the case of each Field Machine.**

18 Known Software Issues

The windows client system may occasionally disconnect from the server if a large time period passes between uses of the system. The only option if this occurs is to save the session and restart the program. This issue is related to excessive network latency causing the timeout of network interface elements in the C++ *libssh* library.

The Workstation Status and Control window and camera control windows remain on top of other V2DVS windows, but may be moved aside with the mouse. This was an early design decision that in actual use appears to be more of an annoyance.

If the database becomes corrupted, the Workstation client software may crash while attempting to process the downloaded data. This should never occur under normal operation, but has been observed with artificial data introduced into the database for testing purposes.

The transfer of log files or image archives from the field machine to the server has occasionally (rarely) fails to occur. The cause is not yet understood and remains under investigation.

It is advisable to not run the V2DVS Workstation application at the same time as other memory-intensive Windows programs. It is possible that the memory demands of the V2DVS program may eventually cause the crash of the other windows programs. Only one case has been observed: multiple instances of MS Word 2010 running concurrently with V2DVS Workstation. Word crashed.

Detection of vehicle speed requires signals from both detectors in a duplex pair, or synthesized duplex pair. It is possible that a failure to detect by the second detector in a duplex pair followed by a failure to detect by the first detector in a duplex pair will be interpreted as an extremely low vehicle speed. This is a highly unlikely event which is fundamental to the use of duplex detectors for measurement of vehicle speed.

19 V2DVS Technical Support Contacts

Art MacCarley, V2DVS System Designer and Project PI. amaccarl@calpoly.edu , 805 756-2317.

Lawrence Klein, Testbed Detector Consultant. larry@laklein.com

John Slonaker, Caltrans Detector Testbed Manager. John_slonaker@dot.ca.gov

20 Cited References

1. Roger Dingledine, Michael J Freedman, David Molnar, David Parkes, Paul Syverson, *Reputation,* April 2003, http://freehaven.net/~arma/jean.html

2. Caltrans Transportation Electrical Equipment Specifications (TEES), 2009, A6-16 and A3-5.

3 American Society for Testing and Materials (ASTM) E 2532 – 09*, Standard Test Methods for Evaluating Performance of Highway Traffic Monitoring Devices.* ASTM International, 100 Barr Harbor Drive, P.O. Box C700, West Conshohocken, Pennsylvania, United States, 2009.

5 Ren, W., H. Chao, Y. Bourgeous, N. Sorensen, Y. Chen, Experimental implementation and validation of consensus algorithms on a mobile actuator and sensor network platform. *IEEE International Conference on Systems, Man and Cybernetics,* Montreal, QC, Canada, 2007.

6 Olfati-Saber, R. Shamma, J.S., *Consensus Filters for Sensor Networks and Distributed Sensor Fusion,* 44th IEEE Conference on Decision and Control, and 2005 European Control Conference CDC-ECC '05, Dec. 12, 2005, pp. 6698 – 6703.

7 Reza Olfati-Saber and Jeff S. Shamma, *Consensus filters for sensor networks and distributed sensor fusion,*Proc. 2005 American Control Conference, pp. 2371– 2378, 2005.

8 Franco, E.; Olfati-Saber, R.; Parisini, T.; Polycarpou, M.M. *Distributed Fault Diagnosis using Sensor Networks and Consensus-based Filters*, 45th IEEE Conference on Decision and Control, San Diego, CA, pp. 386-391, Institute of Electrical and Electronic Engineers, http://dx.doi.org/10.1109/CDC.2006.376797 Dec. 2006.

21 Appendix

Appendix 1. Input File Terminal Assignment Detail and Wiring Diagram. Caltrans TEES 2009.

From P. 85. 2009 TEES, Section 6.4, regarding the input file specifications:

6.4.4 Input File

6.4.4.1 Depth
The file shall have a maximum depth of 8.5 in and shall intermit with and support 14 two-channel detector sensor or isolator units.

6.4.4.2 Connectors
The file shall provide a PCB 22/44S connector centered vertically for each twochannel detector. The associated number and letter side connectors shall be shorted internally. Pins D, E, F, J, K, L and W shall be brought out to a 8 position terminal block on the back of the file. The output emitters shall be common grounded with the ground terminating at TB 15, Position 4. Position 8 of the terminal block is assigned to Equipment Ground and is used to terminate lead in shields.

We need pins H and X brought out since these are the emitter pins:

From Operating Manual, Sarasota Model 222/224 Gp7

CONNECTIONS

Connection is made via a double-sided card edge connector 0.156" spacing mating with Cinch Jones 50-44A-30M connector designed for keys between contacts B and C, E and F plus M and N. Lettered pins are on the bottom of the PCB. The numbered pins are on the top or the component side of the PCB. The 222 Series functions are listed in brackets [ ].

Table 4 – Pin outputs of the Model 222/224 Detection cards

Pin Function Pin Function

A DC common (-) 1 Spare
B DC 24V (+) 2 Spare
C Reset 3 Spare
D Ch. 1 loop input 4 Spare
E Ch. 1 loop input 5 Spare
F Ch. 1 output collector (+) 6 Spare
H Ch. 1 output emitter (-) 7 Ch. 1 status output
J Ch. 2 loop input 8 Spare
K Ch. 2 loop input 9 Spare
L Chassis ground 10 Spare
M Not used 11 Spare
N Not used 12 Spare
P Ch. 3 loop input [Spare] 13 Spare
R Ch. 3 loop input [Spare] 14 Spare
S Ch. 3 output collector (+) [Spare] 15 Spare
T Ch. 3 output emitter (-) [Spare] 16 Ch. 3 status output
U Ch. 4 loop input [Spare] 17 Spare
V Ch. 4 loop input [Spare] 18 Spare
W Ch. 2 output collector (+) 19 Spare
X Ch. 2 output emitter (-) 20 Ch. 2 status output
Y Ch. 1 output collector (+) [Spare] 21 Spare

From TEES, the wiring diagram for input file backplane is shown below. Note that the emitter pins H and X are not brought out on terminals, but are connected on the backplane itself to "CONTROLLER DC GROUND". This explains why the 170 controller must be connected for a signal to appear on the backplane terminals.

INPUT FILE DETAIL

17.50 MAX (OUTSIDE DIMENSIONS)

0.50±0.20

CENTER TO CENTER TERMINAL DIMENSIONS FOR
TERMINAL BLOCKS 1-14 SHALL BE 1.50±0.02

MOUNTING LOCATION IF T15
MOUNTING LOCATION OF T16

TERM 1 → 4
TERM 1 → 3

REAR VIEW

INPUT FILE TERMINAL ASSIGNMENT DETAIL

T1 - 14

TERM	PIN - FUNCTION
1	SP - SPARE
2	F - CHANNEL 1 OUTPUT
3	V - CHANNEL 2 OUTPUT
4	R - CHANNEL 1 INPUT
5	E - CHANNEL 1 INPUT
6	J - CHANNEL 2 INPUT
7	E - CHANNEL 2 INPUT
8	L - EQUIP. GROUND

T15

TERM	FUNCTION
1	+24 VDC
2	DC GROUND
3	DETECTOR RESET
4	C1 HARNESS DC GROUND

T16

TERM	FUNCTION
1	AC+ CLEAN
2	AC- CLEAN
3	EQUIP. GROUND

INPUT FILE WIRING DIAGRAM

DC GROUND
+24 VDC
DETECTOR RESET
OUTPUT
OUTPUT
EXCITATION
EXCITATION
CHASSIS GROUND
AC-
AC+

CONTROLLER DC GROUND
EX GND
AC-
AC+

OUTPUT
OUTPUT
EXCITATION
EXCITATION
OUTPUT
OUTPUT
OUTPUT
OUTPUT

CONTROLLER DC GROUND

SLOT 1   SLOT 2   SLOT 3 → SLOT 14

TITLE: INPUT/OUTPUT FILE DETAILS SHEET 1 OF 5

NO SCALE

A6-16

TEES 2009

NOTE:
All dimensions shown are in inches.

Appendix 2. **Video Image Processing Internal Parameters**

With the exception of the lane_inches, the user normally will not access these settings. The lane width need only be set once for a location, must be set only if a camera is repositioned or replaced, since this measurement serves as the calibration between the scene coordinates and the image coordinates.

You can set each of the following parameters from a command line window on a field machine by typing:
/opt/v2dvs/bin/`rt_config set <parameter> <value>`

<u>Display Options</u>

Menu items in this category are for diagnostic purposes only.

stripe – Set to 1 to make the virtual detection "stripes" visible overlaid upon the real-time camera view in the upper pane of the V2View window. These stripes retract with the approach of the front end of a vehicle, and change colors dependent on whether V2DVS determines the optical flow front to be due to a true vehicle leading edge (red stripes) prior to the capture zone, or shadow (yellow stripes). Red stripes turn green after the vehicle successfully enters the capture (left) zone, indicating a successfully acquired vehicle. Set to 0 to disable.

car_pos – Set to 1 to make visible a green vertical line which, on a frame-by-frame basis, tracks the progress of the front end of each vehicle from right to left across the detection box. Set to 0 to disable.

blockify_fg - Set to 1 to reveal the results of an internal video processing step in which the contents of the detection window are "blurred" by annealing into larger squares each based on a field of 100 (10x10) native pixels. Each pixelized block is itself 5x5 pixels in size. Set to 0 to disable.

<u>Stripe Parameters</u>

Detection and tracking of the front edge of a vehicle requires that at least one of the following three tests be satisfied on a field-by-field basis:
For edges that are brighter than the background, the intensity change must exceed the **Bright threshold (% deviation)**, stated as a percentage.

For edges that are darker than the background, the intensity change must exceed (be darker than) the **Dark threshold** absolute minimum level, in the range of 0 to 255, AND the foreground in the vicinity of the edge (at least three pixels prior to the edge) must be brighter than the **Dark threshold** setting by an amount equal to the Bright threshold percentage. Three pixels immediately to the right of the detected edge must differ by **Color threshold (milliradians)** in their color angle (in the RGB space) compared with the background color angle beneath these pixels AND all three pixels must be greater than the **Darkness threshold**.

User-settable parameters affecting basic detection are:

shadow_reject –The detection algorithm uses several methods to reject false detections due to the presence of long shadows from vehicles in adjacent lanes. Inevitably, this involves some trade-offs which could compromise the proper detection of some vehicles (especially ones that look like shadows) under poor lighting conditions. If set to 0, the shadow rejection features will be disabled. Under conditions of no shadows but otherwise very poor lighting, detection accuracy may be enhanced. Under normal conditions, this should always be set to 1.

num_stripes – The primary detection method used by V2DVS to track, detect and measure the speed of each vehicle is based on the propagation of a set of stripes from left to right through the detection box. As vehicles progress from right to left, these stripes retract at the perceived leading edge of the vehicle. This parameter box reports the number of stripes used in the current detection window. This number is automatically determined when the detection box is created. It is possible to override this number directly for diagnostic purposes only, but this is not advised for normal operation.

stripe_separation – This parameter is the number of pixels measured vertically between each detection stripe. This is normally set to an optimal overall value and not adjusted by the user. Decreasing this number increases the number of detection lines contacting a given vehicle, which improves the probability of detecting very dark cars with few distinguishing features, but increases the processor load.

stripe_start – A stripe group is a collection of detection stripes which terminate at horizontal positions (indicative of the front end of a vehicle) within a range of positions less than the Maximum Group Separation, defined below. The value set in the Minimum Stripe Group parameter box is the minimum number of such stripes required to define a valid front-end position for the vehicle.

stripe_end – This parameter defines the maximum spread of horizontal positions at which stripes may terminate, in order for them to meet the requirements of a valid stripe group. The units of this distance parameter are pixels.

ratio_t – The shadow rejection feature of the detection algorithm requires that a true shadow extends for some minimum horizontal depth. If the potential "shadow" area does not exceed this depth, it is not determined to be a shadow, but is detected as a regular vehicle feature. The Ratio Test Depth parameter defines this horizontal depth in units of feet, based on the current pixels/feet calibration.

A word on Speed measurement methods and criteria:

Once a valid set of samples of the vehicle front end are detected from the array of detection stripes, these are processed to determine if the set represents a valid vehicle (as opposed to a shadow or video artifact) and the speed of the vehicle is measured from this dataset. A hierarchical set of three methods is used for the measurement of the vehicle speed. If the criteria for one method are not satisfied, the next method is used. If the criteria for all three methods are not satisfied, no speed is measured and the potential vehicle is rejected. The three methods for speed calculation, in order of application, are:

Linear least curve fit (code F) – The time-position ordered pairs determined along all stripes are submitted to a least squares linear curve fitting algorithm, and resulting slope (rate) coefficient is reported as the mean speed of the vehicle through the detection window.

Simple Time of Flight (code P) – A qualified set of position samples is used to establish the maximum interval over which the vehicle progression is considered "smooth", e.g., with a consistent progression between successive video fields. The endpoints of this qualified interval are then used to determine a distance, which is divided by the elapsed number of video fields, each of know duration, to calculate the vehicle speed. This method is straightforward but conservative due to its requirement for a sufficiently large smooth set of position samples.

Clustering Algorithm (code C) – Many speed measurements are made along each individual stripe, based on every pairing of valid position samples separated by three successive video fields. A time-of-flight calculation is performed based on each of these local speed calculations, and these numbers are classified into bins, each within the error bounds defined by the block size (the position error). These speed measurement bins are allowed to overlap in a histogram. After the passage of vehicle, the speed corresponding to the peak point in the histogram is declared to be the speed of the vehicle. This approach essentially yields the local speed measurement which is measured most often during the passage of the vehicle. It is the most liberal of the speed measurement methods since it only requires some minimum number of valid pairs of position samples.

Parameters associated with each method are described below:

<u>Curve fit speed calc.</u>

Vehicle progress from right to left through the detection window is sampled as the retraction of each stripe. During every video field, each stripe that terminates on the front end of the vehicle generates an ordered time-position pair each

fit_min_seq – Position samples along each stripe are admitted to the curve fit dataset for a given vehicle only if a minimum number of valid position samples have been detected. This criteria is applied to each individual stripe as a requirement for utilization of its position samples. This parameter establishes this minimum quantity.

fit_min_stripe – Minimum number of stripes in each video field that must satisfy the minimum stripe sample requirement (above) in order for a valid vehicle position to be detected and tracked.

fit_min_samp – Minimum number of total valid vehicle position samples, accumulated over all stripes, required in order for the curve fit algorithm to be run. Failing this test causes this detection method to fail.

Clustering speed calc.

vote_min_votes – The minimum number of position sample pairings, obtained from any of the stripes, required for a valid detection to be declared by this algorithm and for the probabilistic speed analysis to be performed.

vote_t – The clustering algorithm groups pairs of position samples a fixed number of fields apart (currently fixed at three fields for each pairing). The *minimum voting inverse ratio* is the divisor of a fraction which establishes the minimum number out of the total number of valid position pairs required for a valid clustering-based speed measurement to be performed.

Noise Parameters

bright_t – Minimum noise level, on the 0 to 255 intensity scale, of two adjacent 10x10 pixel blocks which is tolerated prior to acceptance as a valid optical flow field edge. The default value is 4.

bkgnd_nudge – Overall tolerated video intensity error amount, considering both baseline noise and the image brightness intensity threshold, on a scale of 0 to 255, prior to acceptance of a valid optical flow field edge. This may possibly be increased in situations where excessive errors occur due to abrupt lighting changes. The default value is 180.

int_cutoff – The overall foreground image brightness is divided by this quantity and added to the baseline noise parameter to establish the dynamic noise threshold, which adapts to variations in illumination level.

Zone detection parameters

v_box – Set to 1 to activate an alternative method of detection based upon bulk intensity change in two sequential detection boxes positioned with the mouse on the roadway. This is the detection method used by almost all commercial video vehicle presence detection systems. It is inferior to the tracking methods implemented as the default by V2DVS, especially with regard to shadow-rejection ability.
box_ratio_h_t – The fraction bulk intensity change in each of the detection boxes in order for vehicle presence to be recognized. Presence is considered "on" when this threshold is exceeded.

box_ratio_l_t - The fractional bulk intensity change in each of the detection boxes in order for vehicle presence to be non-recognized. Presence is considered "off" when this threshold is exceeded. This setting, in combination with the *box_ratio_h_t* , implements a form of hysteresis which reduces false detections.

box_h_min – The absolute intensity threshold that must be exceeded in a detection box in order for vehicle presence to be recognized. Presence is considered "on" when this threshold is exceeded.

box_l_min - The absolute intensity threshold that must be achieved in a decreasing sense in a detection box in order for vehicle presence to be non-recognized. Presence is considered "off" when this threshold is exceeded.

Roadway geometry

lane_inches – The separation distance in inches between lane lines in the scene. The parameter must be entered at least once for speed measurement to be performed correctly. The lane width serves as the calibration between the scene coordinates and the image coordinates. It need not be reset again, even if the

camera position or field of view is changed, since it is characteristic of the roadway itself. The default value is 132 inches (11 ft.).

Scene parameters

bright_t – The set point, on an intensity scale of 0 to 255, for the automatic iris control. Default value is 100.

int_cutoff – The minimum average scene intensity at which vehicle detection by V2DVS is performed. The default value is 20 on a scale of 0 to 255, with 0 corresponding to zero light and 255 corresponding to a completely white image.

Appendix 3. Manual Removal of Detection Data from the V2DVS Database

In normal operation, the server continues to accumulate all data and vehicle images drawn from all field machines, each monitoring ten or more detectors. The server contains 600 Gbytes of hard disk capacity, which is sufficient for months of data accumulation. However, as a maintenance item, it is necessary to periodically remove the accumulated data and image files from the server to avoid exceeding the capacity of the disks. During the process of data reduction and verification, all these data files are copied to one or more client machines. Therefore, they are no longer required on the server.

Preferred method: Please see Section 10.3 for data removal via the user interface. The data on the server (both the images and data stored in the database) along with the images stored on the windows machine running the client can be removed using the **Data Deletion** tab under **Actions** in the main window (as described above).

Manual method: If the preferred method Section 10.3 (using the V2DVS Workstation Application) is no not possible, or if data removal must be done exclusively from the Server itself, the procedure for manually removing image files is as follows.

1. Log into the server on console or remotely via SSH as root. The root password is in the possession of the Caltrans Detector Testbed manager, John Slonaker. After logging in, use a Linux command window for the following steps.

2. To remove all image files type: `"rm –rf /home/dfserv/store/*"`

This will remove ALL past data and images. It is also possible to selectively remove data accumulated before a particular date or for a particular detector or field machine, by using the path

`/home/dfserv/store/<year, 4 digits>_<month number>_<day number>/<hour, 24-hour format>/<lane number>/<detector number>`

month, day, and hour numbers are NOT zero-padded, i.e., no need to put a leading zero in front of the number 2 (ie. 02).

To flush the data from the database on the server, type at the Linux command prompt:

`mysql --user=tabber --password=loragen805 central`

You will get a mysql command prompt.

Note: the root username and password for mysql is: user=root password=loragen805

Type the following SQL commands:

`delete from error_log where time_sec between UNIX_TIMESTAMP('yyyy-mm-dd 24HH:MI:SS') and UNIX_TIMESTAMP('yyyy-mm-dd 24HH:MI:SS');`
(records associated with a report)

```
delete from field_log where stamp between 'yyyy-mm-dd 24HH:MI:SS' and 'yyyy-mm-dd
24HH:MI:SS';
```
(records associated with all detections)

3. To remove all images stored on the windows machine running the V2DVS client, simply delete the desired folder(s) from the "images" directory in the V2DVS installation location. These folders are likely to be stored in "C:\Program Files\V2DVS\images..." organized by date in the following format: `\<year, 4 digits>_<month number>_<day number>\<hour, 24-hour format>\<lane number>\<detector number>`.

Appendix 4. V2DVS Field Machine or Server Upgrade

In order to upgrade the software on the field machine or servers, the most recent code must be in the git repository, make sure this is the case before attempting the following instructions.

After backing up if desired, run the following commands within a terminal connected on the local area network:

bash$: su
[note: root password is bbroygbVgw]

bash$: cd /root/src/v2dvs_linux

bash$: git pull origin master
[pw is the same as root password above, after entering, source will update]

If you are on a field machine:
bash$: make dfclient_make
bash$: make dfclient_install

If you are on the server:
bash$: make dfserv_make
bash$: make dfserv_install

bash$: shutdown -r now

The upgraded V2DVS should run when the machine reboots.

Appendix 5. Source Code

A CD is attached containing source code and executable program files for all software developed under the V2DVS Gen II contract. A descriptive CONTENTS.TXT text file on the CD provides an index to the CD contents.

Appendix 6. Replacing a V2DVS Field Machine with the "hot swap" spare Field Machine

Some configuration is necessary before using the hot swap field machine. Take note of the number of the lane that the new machine is being used for, it will be utilized for configuring the networking.

Physically install the new machine, boot it up, and open a terminal window:

bash$: vim /etc/rc.d/rc.inet1.conf

[Once the file has opened, find the line IPADDR[0] underneath the comment: "Config information for eth0:" and set the last octet in the IP to 5X, where X is the lane number. Write the file and quit vim (:wq).
e.g. lane 5 would look like IPADDR[0]="192.168.40.55"]

bash$: /etc/rc.d/rc.inet1 restart

bash$: vim /etc/HOSTNAME
[Edit the line at the top to say dfclientLX where X is the lane number, write the file and quit vim]

bash$: vim /etc/hosts
[Change the 2nd uncommented line to say dfclientLX. It should now look like:
127.0.0.1   darkstar.example.net dfclientLX
write the file and quit vim]
Change the Lane # and Machine # to the machine being swapped out:
bash$: su
[note: root password is bbroygbVgw]

bash$: cd /root/src/v2dvs_linux

bash$: vim app/config.h
[open the config.h file and change the #define for LANE and MACHINE_NO to the lane being swapped out]

Make any configuration changes to detectors.conf (setting up external detectors and offsets). See Section 7.3.

bash$: shutdown -r now

Appendix 7. **Installation or Upgrade of Field Machine Software on Site**

After physically installing a field machine, certain configuration settings need to be setup and tested. The following steps should be followed:

Replace complete section with installation from flash drive via script.

Verify hostname and IP address:

   bash$: hostname
   [This should print out the current hostname of the computer. If different from what it should be (based on the lane it is for) it can be changed using the following commands:
      bash$: vim /etc/rc.d/rc.inet1.conf
      [Once the file has opened, find the line IPADDR[0] underneath the comment: "Config information for eth0:" and set the last octet in the IP to 5X, where X is the lane number. Write the file and quit vim (:wq).
      e.g. lane 5 would look like IPADDR[0]="192.168.40.55"]

      bash$: /etc/rc.d/rc.inet1 restart
   ]

   bash$: ifconfig
   [This command must be run as root. If the IP is different from what it should be (based on the lane it is for) it can be changed using the following commands:
      bash$: vim /etc/HOSTNAME
      [Edit the line at the top to say dfclientLX where X is the lane number, write the file and quit vim]

      bash$: vim /etc/hosts
      [Change the 2nd uncommented line to say dfclientLX. It should now look like:
      127.0.0.1   darkstar.example.net dfclientLX

       write the file and quit vim]
  ]

Verify the source code is set to the correct Lane and Machine number:

    bash$: su
    [note: root password is bbroygbVgw]

    bash$: cd /root/src/v2dvs_linux

    bash$: vim app/config.h
    [Open the config.h file and change the #define for LANE and MACHINE_NO to the lane this machine is for.]

Add/remove/change detectors:

    bash$: vim /etc/opt/v2dvs/detectors.conf
    [Open configuration file and make any necessary additions/changes to the file. Details about the file are in Section 7.3]

Reboot the system:

    bash$: shutdown -r now

Verify the Field Machine is getting detections:

    bash$: tail -f /home/dfclient/log/main.log
    [This will show you all current activity of v2dvs detections]

    Using the v2dvs gui, open external triggers from the menu and verify the lines that are connected are being triggered by looking at the visual LEDs blinking. Make sure these are the lines that are expected to be triggered and that the detectors.conf file is setup correctly to pair duplex detectors using the correct lines.

Verify the Server is receiving detections:

    On the server, the following commands can be run to verify the server is receiving data from each field machine:

    bash$: tail -f /home/dfserv/log/unsent_log_check.log
    [The output of the above command will show when the server has successfully retrieved log information and images from a connected field machine, as well as a timeout for which field machines it was not able to connect to]

    To access the mysql database to view detection information run the following commands as well:

    bash$: mysql --user=tabber --password=loragen805 central
    [This will log into the database. The following commands show a new prompt representing the database]

    mysql> select * from field_log order by id desc limit 20;
    [The above command will show the 20 most recent items entered into the database. The parameters may be changed to access specific data based on what the user wants or what is being tested. This requires knowledge of mysql queries and can be looked up in the man page for mysql]

Appendix 8. Reprogramming and Installing the 68HC11 Lens Control Boards

EPROM Programming Process:

NOTE: If communication cannot be made with the board, success has been obtained by removing all connector wires on the board (i.e., bringing it back to stock). After the programming process, re-solder the wires as described in the next section.

- Connect the wires 12V to V_{PP} and 0V to ground   (Item W8 on the board)
- Connect the wires 5V to the connector closer to the 12V V_{PP} and 0V to the other (Item P5 on the board)
- JMP - EE should be 0 (for both of them)
- the other jumper (Item W7 on the board) should be on VPP (**not** VDD).
- Lastly the BDM cable should be connected to BDM IN (Item W9)
    - Note the orientation of connection matters. The bottom side of the connector (the side with the wire) should be closer to the FPGA chip.
- Plug the USB side of the PEMicro into the computer, the blue light should light up on the PEMicro
- Power the board (both 5V and 12 V simultaneously).
    - The yellow light on the PEMicro should light up. If it does not light up, it is likely that the wire connecting to the board is orientated wrong. Start over and flip the connector the other way (on the W9 Item on the board).
- OPEN the "NoICE for 68HC12" program communication
- Click FILE, then LOAD
    - Select the completely built program code (in the executable folder of the IAR code)
- Click Ok/DONE and the board will burn.
- If no errors result, the board is ready for install.......Continue to the next section.

Installing the Lens Control Boards in the Field Machines

- With all of the main parts soldered to the board.
- Move the jumper (Item W7 on the board) to the VDD position (**not** VPP)
- From the Blue box wires (aka the RS-232) solder the wires as follows:
    - Purple wire connects to PA0 (PA zero)
    - Blue wire connects to PA1
    - Orange wire connects to PA2
    - Red wire connects to PA3
    - Green wire connects to PA4
    - Yellow wire connects to PA5
    - The following striped wires are connected together... yellow/black, red/black, green/black, and orange/black.
    - All other wires are not connected to the board and appear to be individually wrapped and then wrapped to the main connecting wire so that they are out of the way.
- Note: there is/was a resistor (colors green, brown, yellow) and capacitor (1 uF - blue with polarity + on one side and the number 35) setup but the result doesn't seem to provide any result.