STATE OF CALIFORNIA • DEPARTMENT OF TRANSPORTATION
TECHNICAL REPORT DOCUMENTATION PAGE
TR0003 (REV 10/98)

1. REPORT NUMBER
CA24-4015

2. GOVERNMENT ASSOCIATION NUMBER

3. RECIPIENT'S CATALOG NUMBER

4. TITLE AND SUBTITLE
TIRTL Evaluation
Comparing the detection accuracy of The Infra-Red Traffic Logger (TIRTL) to inductive loop detectors at a vehicle count station.

5. REPORT DATE
01/31/2024

6. PERFORMING ORGANIZATION CODE

7. AUTHOR(S)
John Slonaker

8. PERFORMING ORGANIZATION REPORT NO.

9. PERFORMING ORGANIZATION NAME AND ADDRESS
California Department of Transportation
Division of Research Innovation, and System Information
MS-83 1727 30th Street
Sacramento CA 95816

10. WORK UNIT NUMBER

11. CONTRACT OR GRANT NUMBER

12. SPONSORING AGENCY AND ADDRESS
California Department of Transportation
Division of Research Innovation, and System Information
MS-83 1727 30th Street
Sacramento CA 95816

13. TYPE OF REPORT AND PERIOD COVERED
final report
1/15/2023-12/15/2023

14. SPONSORING AGENCY CODE

15. SUPPLEMENTARY NOTES

16. ABSTRACT

The California Department of Transportation (Caltrans) has traditionally used inductive loop detectors for vehicle count stations on freeways, but installation and maintenance of loops requires lane closures that are disruptive to traffic. Caltrans engineers would like the option of using less intrusive types of detection systems in their designs. Caltrans District 4 has equipped two of its vehicle count stations on Interstate 680 with TIRTL (The Infra-Red Traffic Logger) detection systems and wanted to measure their effectiveness relative to preexisting inductive loops. Caltrans' Division of Research, Innovation and System Information (DRISI) worked with District 4 to compare the actuation signals output from the TIRTL to the controller to those of the existing inductive loop detectors. All detection signals from both detection systems were captured with the C1 reader device. The controller was programmed to process only the loop detector data, which it aggregated and transmitted to the District 4 traffic management center as usual. An existing CCTV surveillance camera collocated at the detection site was focused on the adjacent detection zones of the loops and the TIRTL. Video from the camera was recorded concurrently with the detection signals and used to provide "ground truth" in order to determine when either detection system reported vehicles that weren't actually present, i.e., "false positives," or failed to report vehicles, i.e., "false negatives."

17. KEY WORDS
infra-red, inductive loop detector, vehicle count station, traffic signal, C1 Reader, VideoSync, Sensitivity, false positive, false negative

18. DISTRIBUTION STATEMENT
No restrictions

19. SECURITY CLASSIFICATION *(of this report)*
Unclassified

20. NUMBER OF PAGES
21

21. COST OF REPORT CHARGED

DISCLAIMER STATEMENT

Evaluation of the CEOS TIRTL (The Infra-Red Traffic Logger) Traffic Detector for Count Station Applications

Caltrans Division of Research,
Innovation and System Information

John Slonaker, P.E.
Caltrans
Division of Research, Innovation and System Information
January 31, 2024

Acknowledgements

Thanks to all the people listed below who were involved in performing this test. The Division of Research, Innovation and System Information (DRISI) engineers provided C1 Reader support as well as data collection on site; the DRISI student assistants provided VideoSync software support as well as many tedious hours of data analysis; the District 4 engineers provided field support at the count stations; and CEOS and Control Specialists provided remote technical support.

Caltrans, Division of Research, Innovation and Systems Integration (DRISI):

Jerry Kwong, Transportation Engineer, Electrical

Aleksey Zasorin, student assistant, computer science

Julian Hernandez, student assistant, computer science

Hualong Yu, student assistant, electrical and computer engineering

Samah Eltayeb, student assistant, computer science

Caltrans, District 4:

James Kung, Senior Transportation Electrical Engineer

George Young, Transportation Engineer, Electrical

Ahmad Fahimi, Transportation Engineer, Electrical

Madrio Melo, Transportation Engineer, Electrical

Md Huda, Transportation Engineer, Electrical

CEOS:

Kristian Penno, Engineering

Control Specialists:

Terence Tomlin, Regional Sales

Abstract

The California Department of Transportation (Caltrans) has traditionally used inductive loop detectors for vehicle count stations on freeways, but installation and maintenance of loops requires lane closures that are disruptive to traffic. Caltrans engineers would like the option of using less intrusive types of detection systems in their designs. Caltrans District 4 has equipped two of its vehicle count stations on Interstate 680 with TIRTL (The Infra-Red Traffic Logger) detection systems and wanted to measure their effectiveness relative to preexisting inductive loops. Caltrans' Division of Research, Innovation and System Information (DRISI) worked with District 4 to compare the actuation signals output from the TIRTL to the controller to those of the existing inductive loop detectors. All detection signals from both detection systems were captured with the C1 reader device. The controller was programmed to process only the loop detector data, which it aggregated and transmitted to the District 4 traffic management center as usual. An existing CCTV surveillance camera collocated at the detection site was focused on the adjacent detection zones of the loops and the TIRTL. Video from the camera was recorded concurrently with the detection signals and used to provide "ground truth" in order to determine when either detection system reported vehicles that weren't actually present, i.e., "false positives," or failed to report vehicles, i.e., "false negatives."

How the TIRTL Works

A TIRTL installation consists of a transmitter and a receiver, located just above road level on opposite sides of the roadway, and a processor located nearby. The transmitter shines infrared "light" (electromagnetic radiation just below visible wavelengths) at 850 nanometers. One light source is modulated at 90 kHz, while the other is modulated at 110 kHz (shown in the below illustration as red and blue). The light from both sources spreads out in a cone shape, shining on both lenses of the receiver. The light that enters the receiver forms four beams between the two units, two parallel and two oriented diagonally.

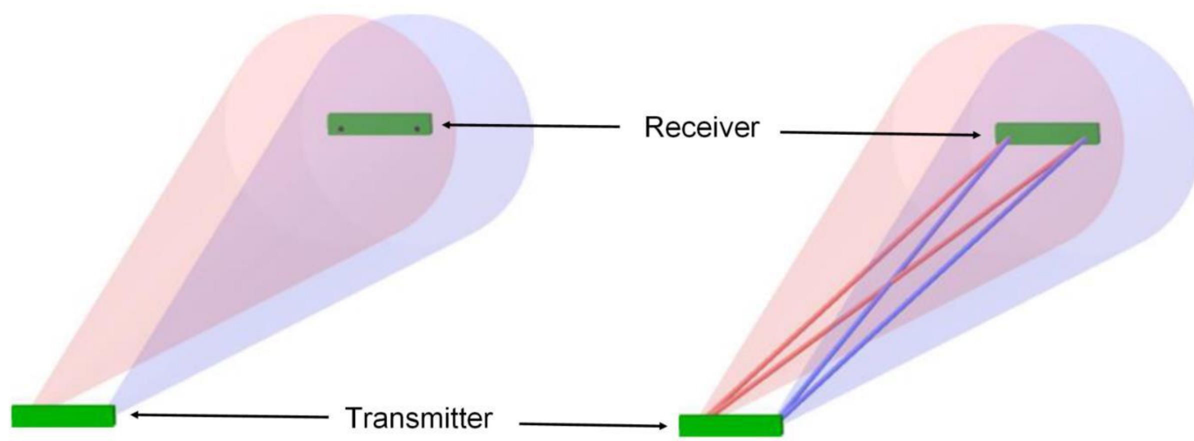

The receiver samples the amplitude of the four beams at 500 kHz and sends the values to the processor. As vehicles pass through the system, each wheel successively blocks each of the four beams. The processor detects each beam being blocked and generates a timestamped beam event; these are called "break" beam events. The same occurs as each beam is unblocked, generating "make" beam events.

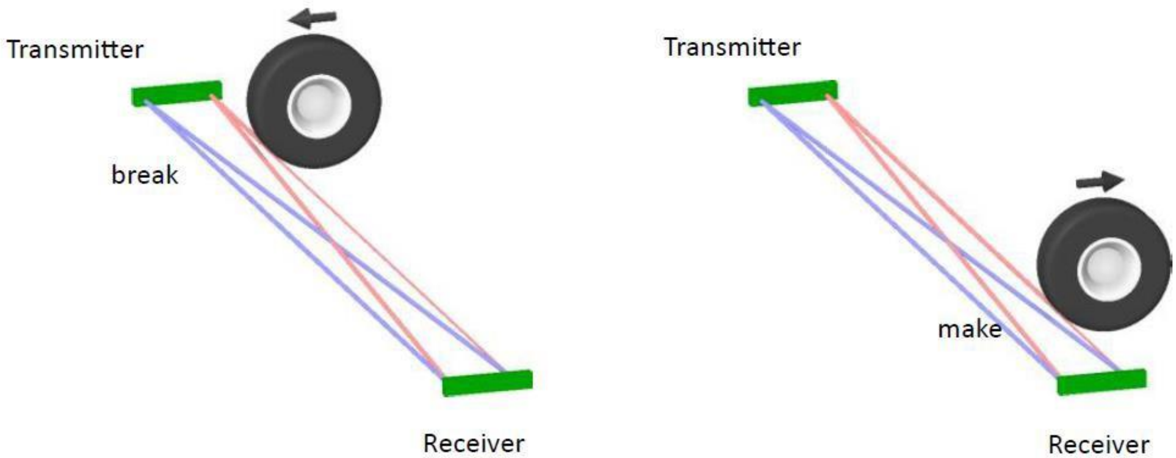

The beams are named from the perspective of the receiver. The parallel beam entering the left lens of the receiver is called Left Parallel (LP), and the parallel beam entering the right lens is called Right Parallel (RP). The diagonal beams are similarly called Left Diagonal (LD) and Right Diagonal (RD).

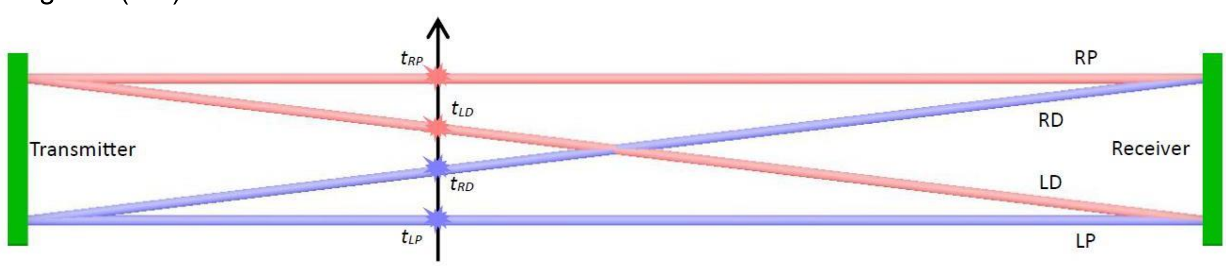

As each wheel traverses the system, the leading edge generates four break events, and the trailing edge generates four make events. Each set of four events is combined to form a single "edge" event.

The distance between the two parallel beams is 153mm. Dividing this distance by the time taken to traverse the parallel beams gives the speed of the edge.

$$speed = \frac{153mm}{t_{RP} - t_{LP}}$$

The direction of travel is found from the order in which the parallel beam events occur. Dividing the time taken to pass from a parallel beam to a diagonal beam by the time taken to traverse the parallel beams gives a value between 0 and 1 representing the location of the wheel between the receiver and transmitter. 0 being at the receiver, increasing to 1 at the transmitter, 0.5 being halfway. Multiplying this value by the total separation of the transmitter and receiver gives the physical distance that the wheel is from the receiver.

$$distance = \frac{t_{LD} - t_{LP}}{t_{RP} - t_{LP}} \times total\_separation$$

This distance can then be applied against the roadway striping pattern to place the wheel in a specific lane.

Each "break edge" and "make edge" has a timestamp, which is used to calculate a speed and a distance from the receiver. The speeds of the two edges are averaged to produce the speed of the axle.

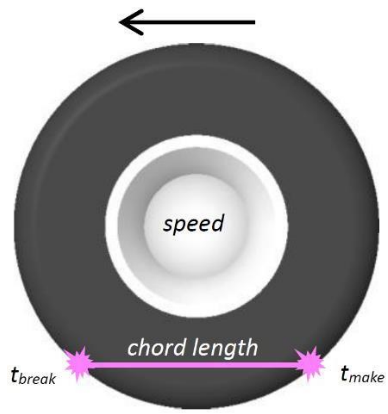

The timestamp, speed and distance from the receiver of the axle events are used to group them together to form vehicles.

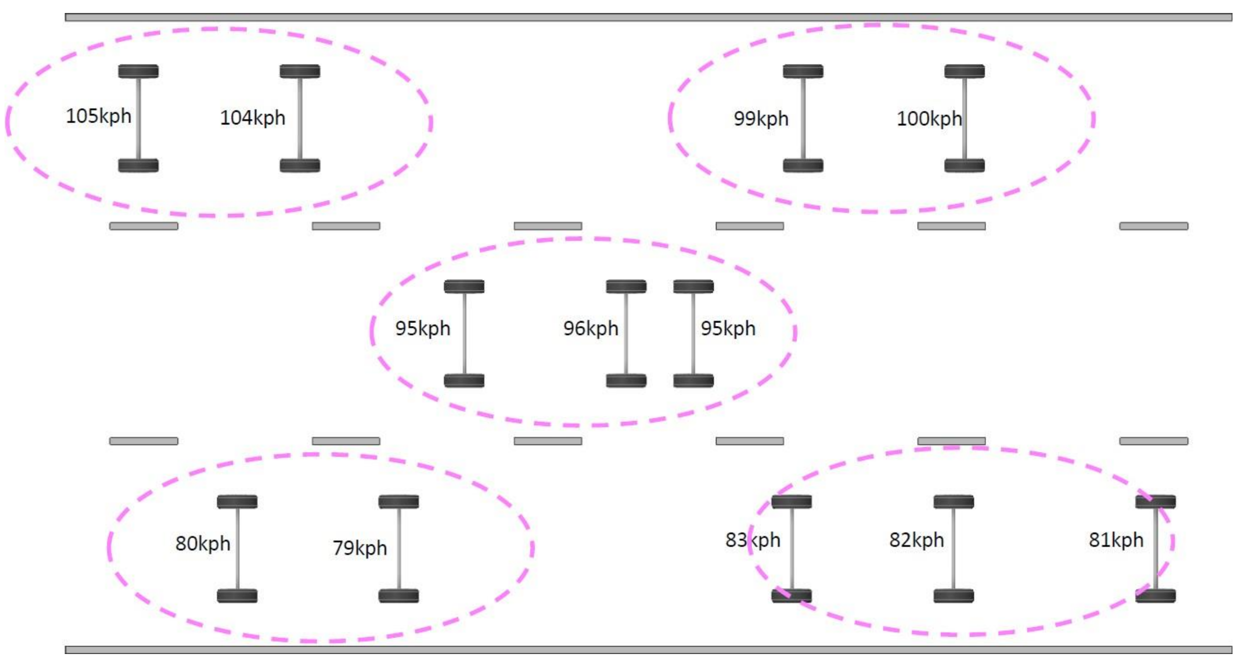

Test Location

The two TIRTL-equipped vehicle count stations in Caltrans District 4 are both located on the southbound (West) side of I-680 in Walnut Creek, one at the North Main Street on ramp, and the other at the South Main Street exit and on ramps. DRISI collected data from both locations, but the dataset analyzed in this report comes from the one at South Main Street.

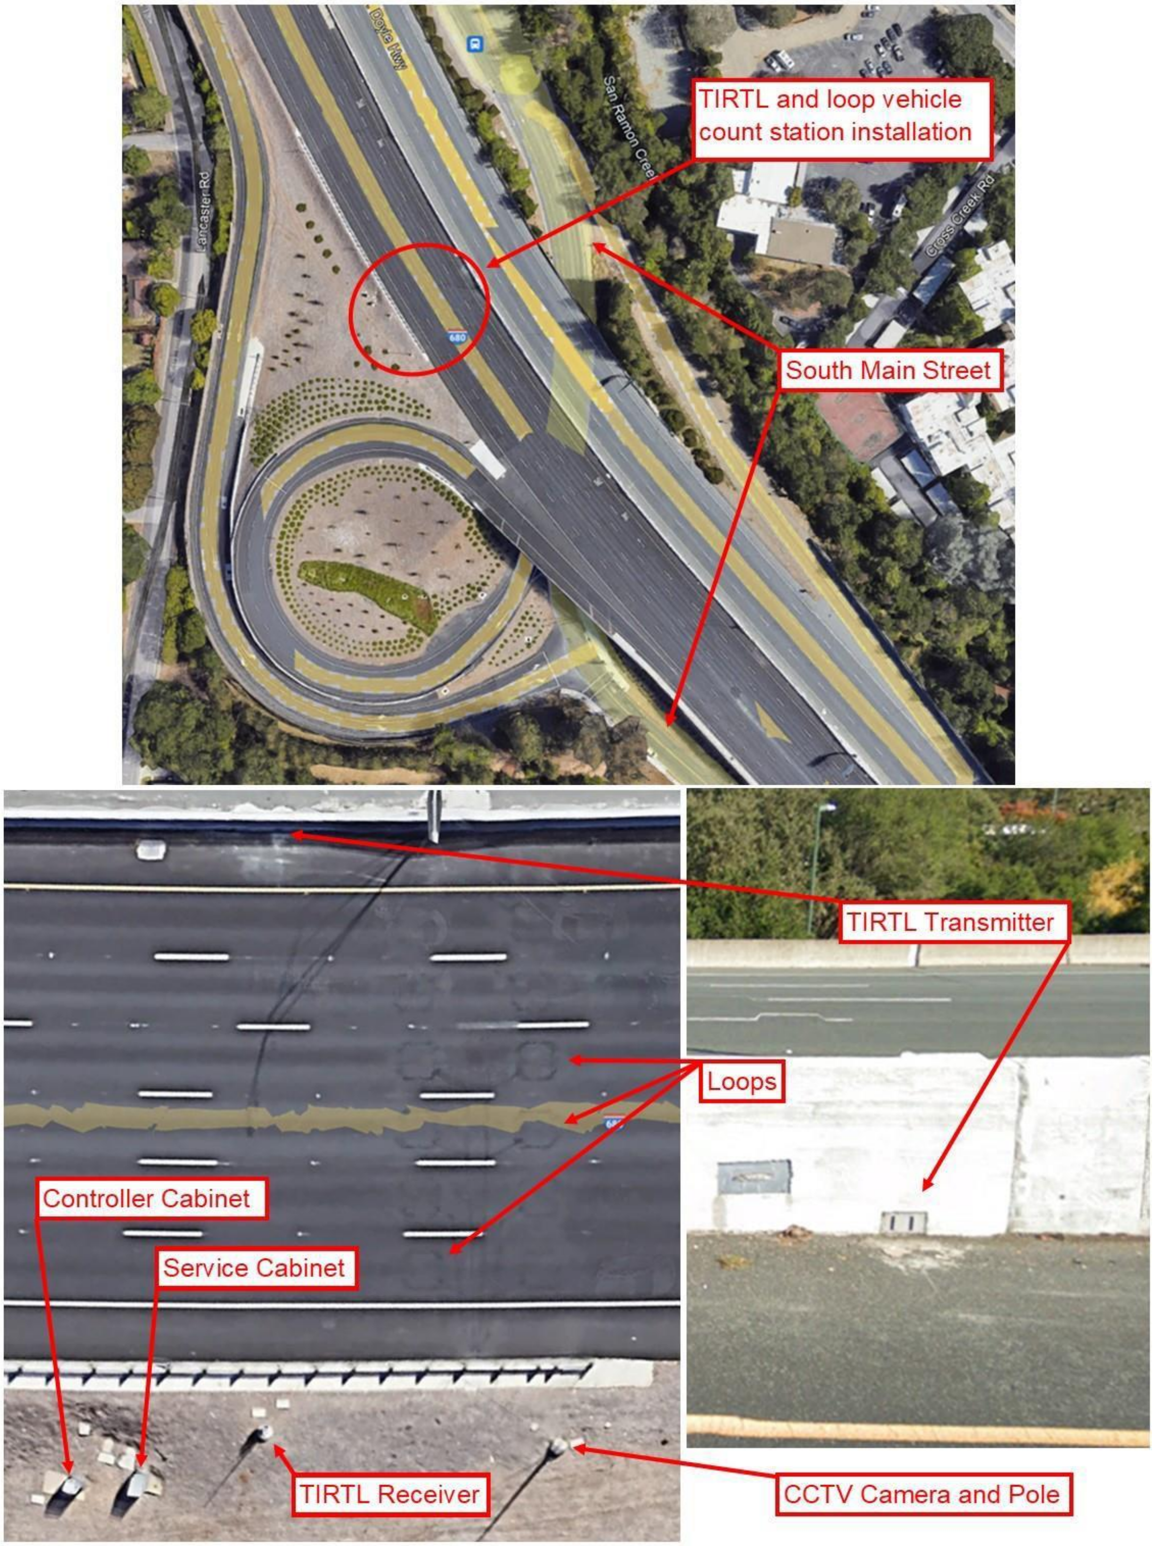

Figure 1 – South Main Street count station site (aerial view)

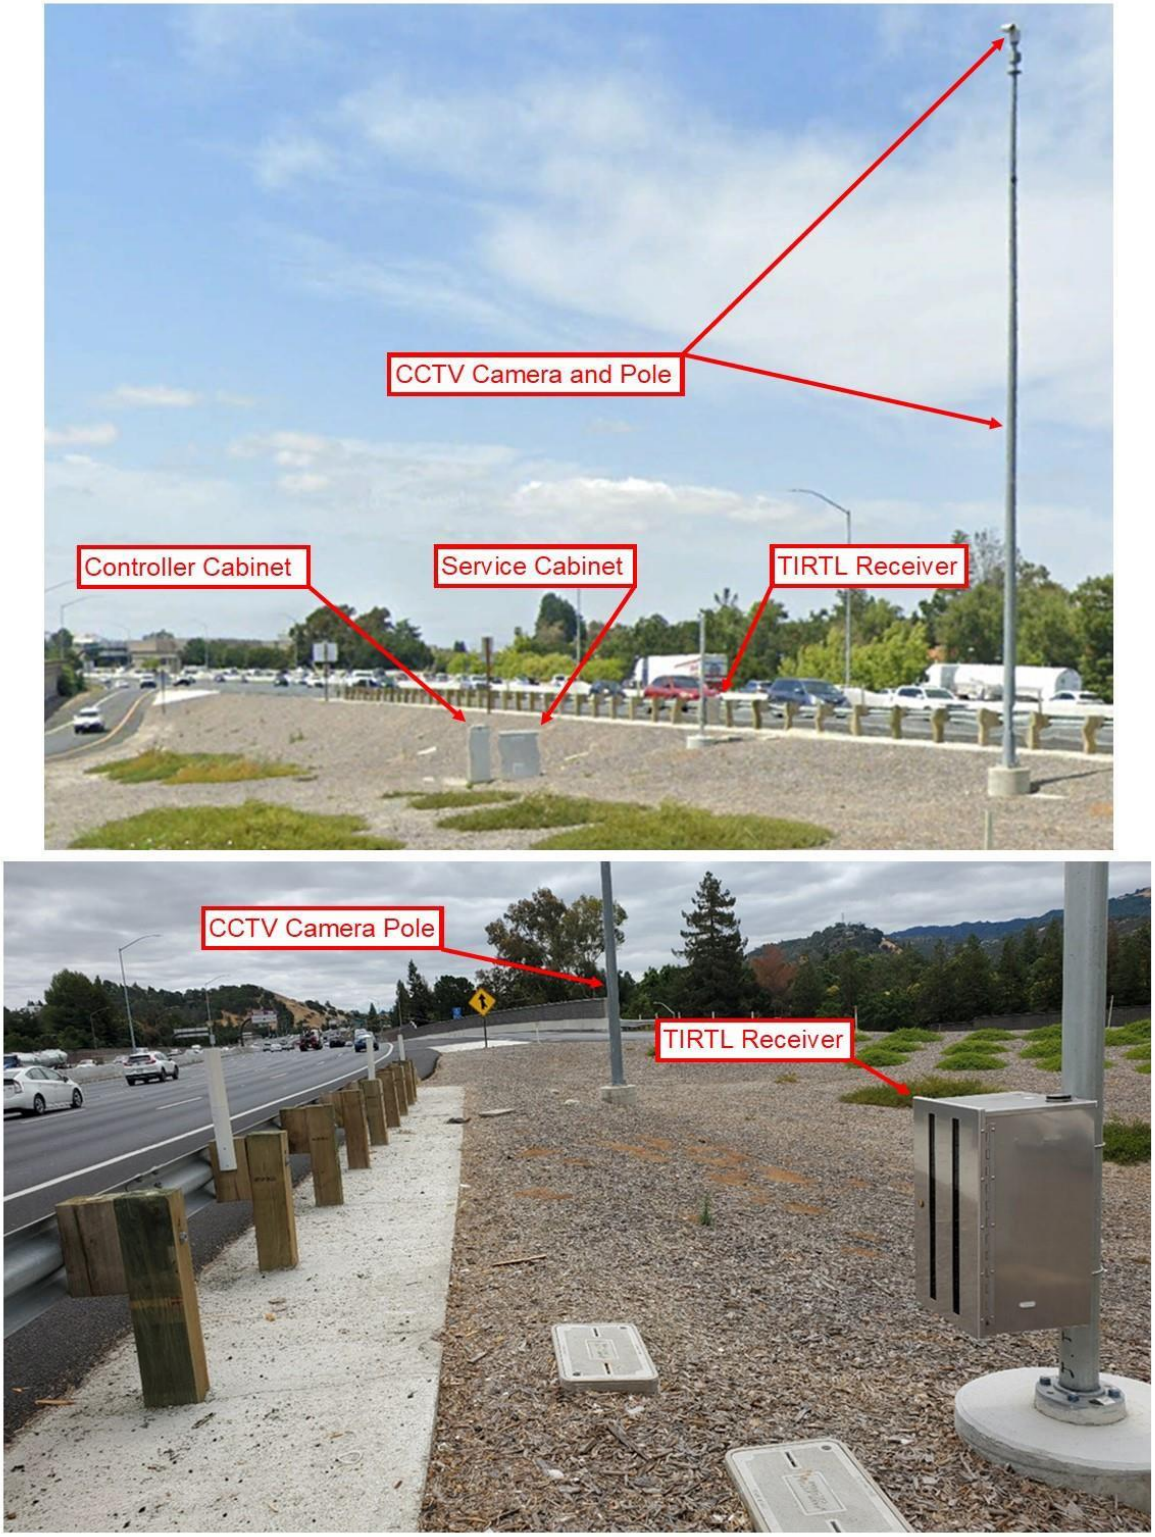

Figure 2 – South Main Street count station site (side view)

Data Collection

The TIRTL detector connects to a separate loop emulation system called the HARE that provides two outputs per lane, one to emulate an upstream loop and the other to emulate a downstream loop. Caltrans District 4 staff wired the detection outputs of the HARE to the controller inputs on slots one through six of the controller cabinet input file. The output corresponding to the upstream loop in Lane 1 was connected to slot one, upper channel, and the output corresponding to the downstream loop in Lane 1 was connected to slot one, lower channel. They continued this pattern for all six lanes. Figure 3 shows the color-coded connections, with white, red, orange, green, blue, and black wire pairs corresponding to lanes one through six respectively. The detector lead in cables from the actual loops, located just downstream of the TIRTL detection zone, were connected to the loop detector card inputs on slots seven through twelve of the controller cabinet input file. The pattern was the same as the TIRTL/HARE connections, with the upstream loop in Lane 1 connected to the upper channel of slot 7, etc. Figure 4 shows loop detector cards plugged into slots seven through twelve from the front of the input file. Slots one through six are empty because the HARE detection outputs work the same way as the loop detector card outputs, pulling the voltage supplied by the controller input pins down to zero to signal a detection. District 4 programmed the controller to aggregate and transmit data only from the actual loops to their traffic management center.

Caltrans DRISI installed a coaxial cable splitter between the analog CCTV camera and the digital video encoder in the controller cabinet. One split video signal was fed to a Flir MPX digital video recorder, and the other was fed to the encoder as before. DRISI then positioned the camera so the loops and the TIRTL detection zone filled its field of view. The remote pan/tilt/zoom control was disabled so the camera wouldn't be moved until the data collection had been completed. Finally, they installed the C1 Reader, a signal sampling and data acquisition device developed by DRISI, between the input file and the controller. Figure 5 shows the C1 Reader connected to the controller and cabinet.

The C1 Reader samples all logic signals on the 104-pin C1 connector cable between the controller and the cabinet, as shown in Figure 6. This includes the detection input file, into which the inductive loop detector cards were plugged and to which the TIRTL/HARE contact closure outputs were connected. The C1 logic signals from both detection systems were recorded concurrently with the video from the CCTV camera. Approximately one week of C1 data and video were collected from August 9th to August 16th, 2023.

Figure 3 – TIRTL/HARE detection outputs

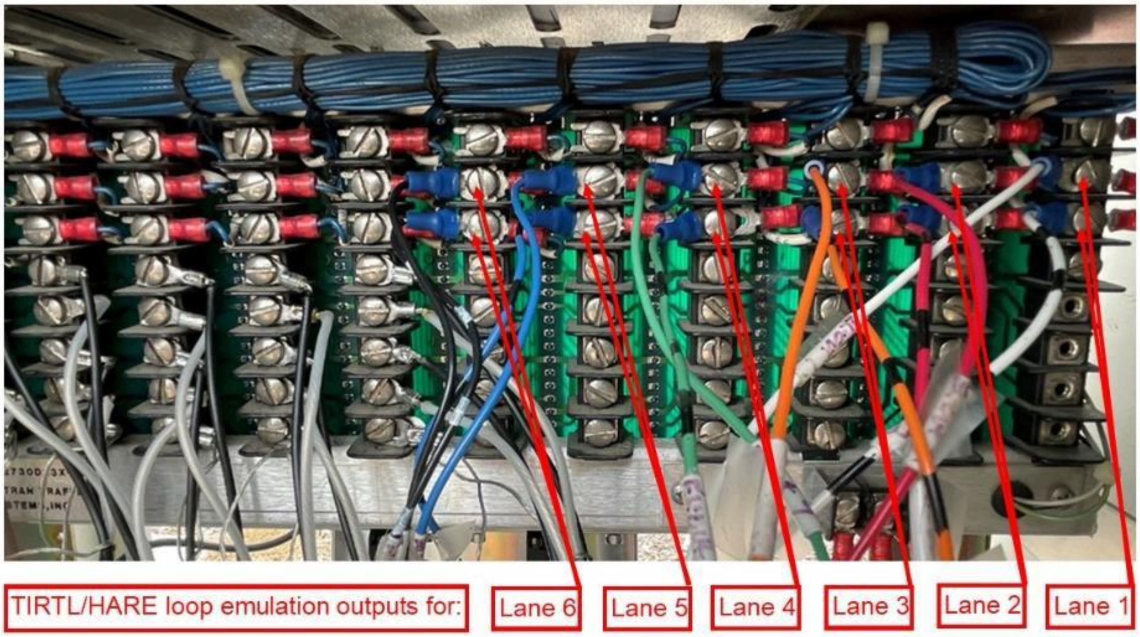

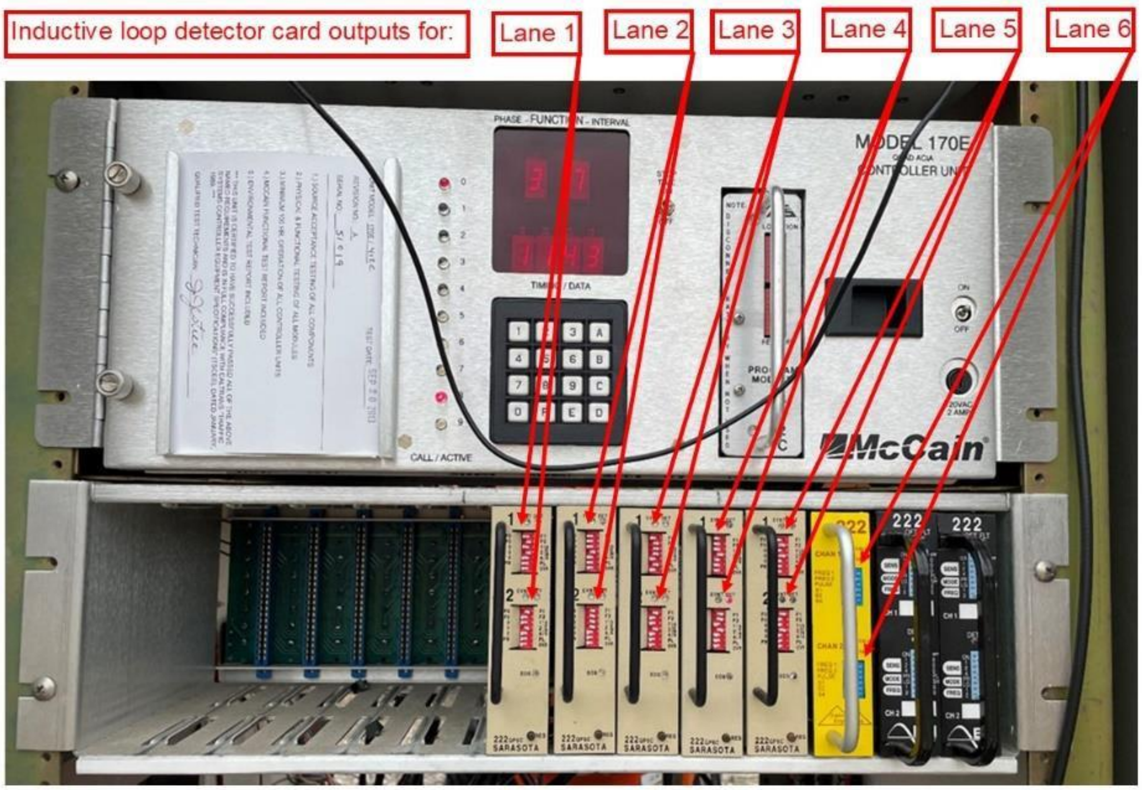

Figure 4 – inductive loop detector card outputs

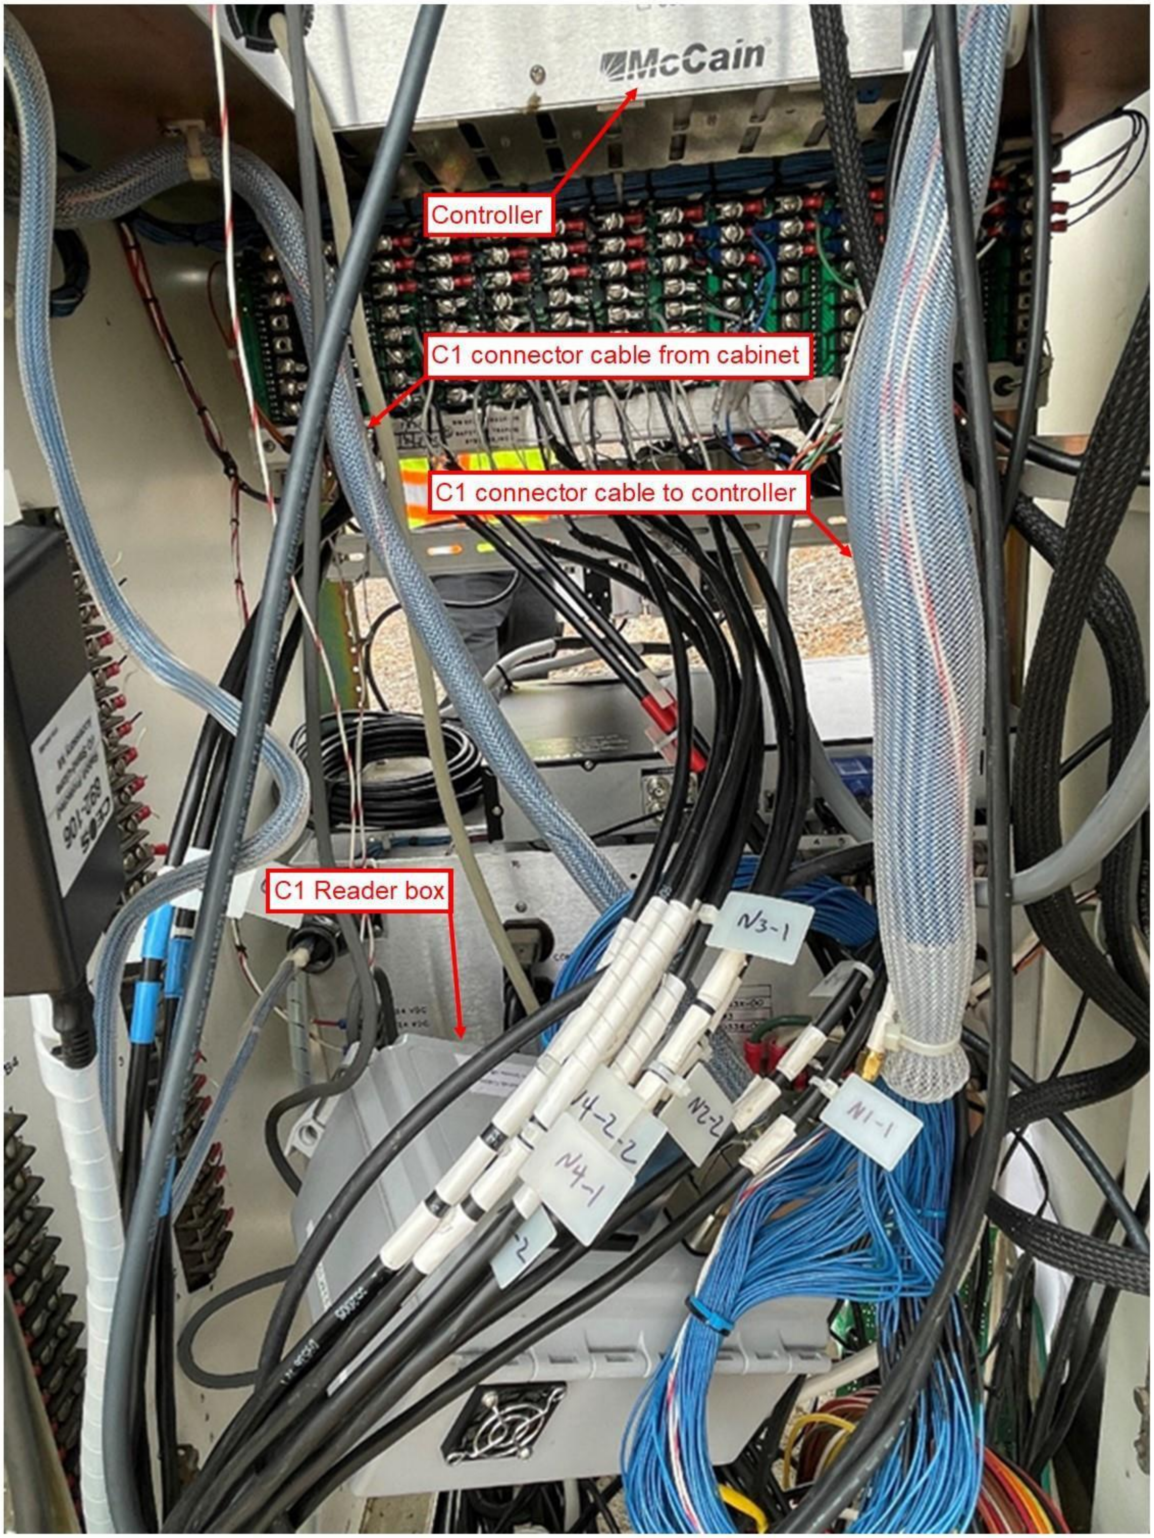

Figure 5 – C1 Reader connected between the cabinet and the controller

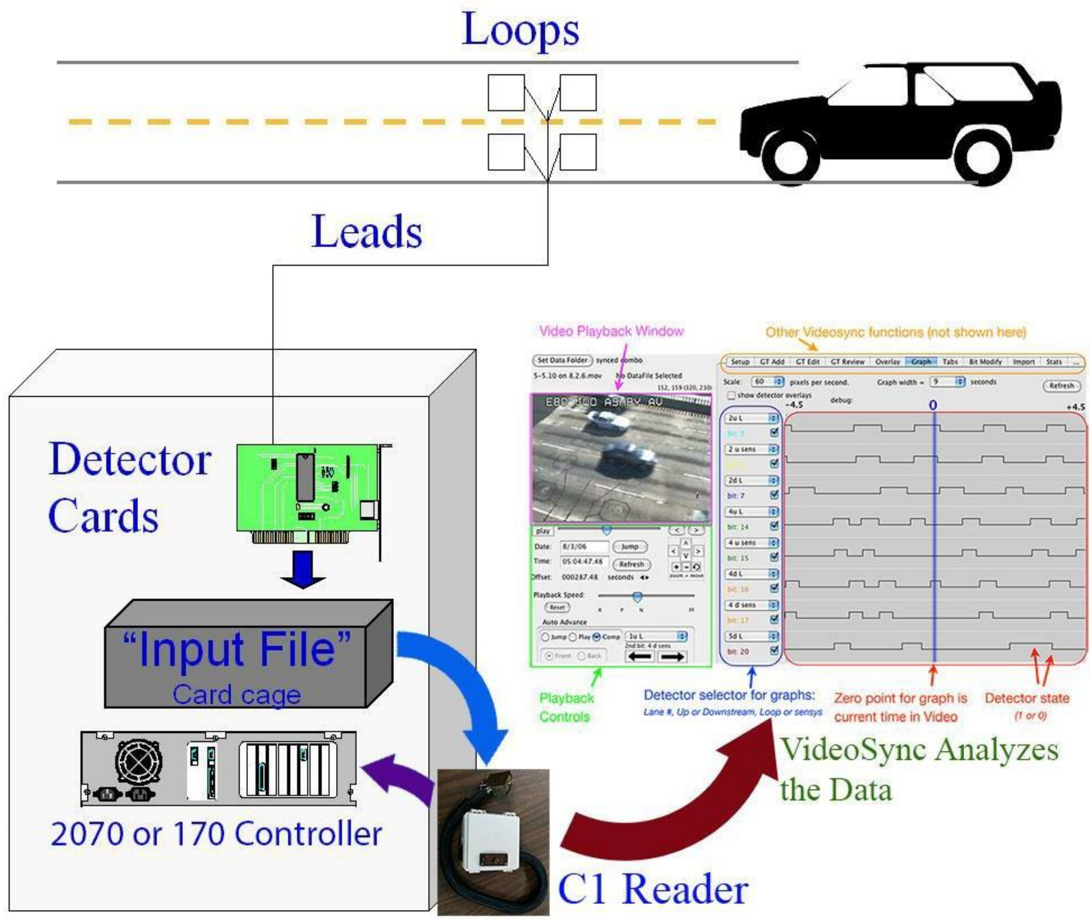

Figure 6 – C1 Reader connected into an inductive loop count station

Data Analysis

After all the C1 detection data and video had been captured, DRISI processed it for analysis. It would have been labor-prohibitive to look at all the data and video, so DRISI chose the hour that had the most congestion: Wednesday, August 9th, from 5 to 6 PM. DRISI thought that this period would be the most challenging for the TIRTL, since the more congestion, the more chance for vehicles being occluded by others in front of, or behind them. Also, the lower the speed of the vehicles, and the more fluctuations in velocity, the harder it would be to group detected axils into individual vehicles correctly. DRISI had already collected about two hours of preliminary data from the North Main Street count station under free flow conditions, and the results from the TIRTL had been very promising, but they wanted to make sure that the TIRTL was capable of accurate detection even in congested conditions. The analysis of the congested hour at the South Main Street count station included a total of 11,196 individual vehicles.

As shown in Figure 7, VideoSync displays "ground-truth" video alongside a graphical representation of the detection logic signals on user selected C1 connector pins. VideoSync includes a pattern recognition algorithm that looks at the spacing of vehicle platoons and matches the vehicles to corresponding detection signals with like spacing. In some cases, this can be used to automatically apply a time offset to synchronize the video with the detection logic signals. In other cases, depending on offset length and video quality, synchronization needs to be done manually. In this test, DRISI was able to use automatic offset detection. Once the detection logic signals and video are synchronized, false detections, i.e., false positives, missed detections, i.e., false negatives, dropped calls, detector contact bounce and other erroneous reported detections are readily visible. The operator then looks at each event where there is disagreement among the detection logic signals from the detection systems (loops and TIRTL in this case) and the video and classifies the events as false positives or false negatives for one or more detectors. VideoSync includes tools that use these data sets, once compiled and analyzed, to generate statistics on the accuracy of any vehicle detector under test.

Figure 7 – Video and graphical display of corresponding detection signals in VideoSync

For this test, the degree of accuracy of a vehicle detector is defined as its “Sensitivity,” where

$$\text{Sensitivity} = \frac{\text{True Positive}}{\text{False Positive} + \text{False Negative} + \text{True Positive}} \times 100.0$$

According to this definition, the detector is penalized equally and cumulatively for each false positive and false negative. The more of either, the lower the Sensitivity. If there are none of either, the Sensitivity equals 100%.

Results

The overall Sensitivity of the inductive loops was slightly higher than that of the TIRTL. The difference was only 0.37%. This difference is much lower than that of microwave radar and machine vision detectors that DRISI has tested in the past, which typically have Sensitivity values from about 3% to 7% lower than loops. The only other detection system that DRISI has tested with Sensitivity values as close to that of loops is the wireless magnetometers from Sensys Networks. Table 1 shows a breakdown of the Sensitivity of each detector by the twelve detection zones, e.g., lane 1 upstream, etc. The Sensitivity of the loops is slightly higher for nine zones, but the Sensitivity of the TIRTL is higher for three zones.

Table 1 - Detector Sensitivity

Detection Zone	Loops	TIRTL	Difference
All lanes combined	99.54%	99.17%	0.37%
Lane 1 upstream	96.03%	99.34%	-3.31%
Lane 1 downstream	99.83%	99.34%	0.49%
Lane 2 upstream	99.94%	99.49%	0.45%
Lane 2 downstream	98.80%	99.55%	-0.75%
Lane 3 upstream	100.00%	99.20%	0.70%
Lane 3 downstream	99.92%	99.20%	0.72%
Lane 4 upstream	100.00%	98.74%	1.26%
Lane 4 downstream	100.00%	98.74%	1.26%
Lane 5 upstream	100.00%	98.52%	1.48%
Lane 5 downstream	100.00%	98.52%	1.48%
Lane 6 upstream	99.43%	99.61%	-0.18%
Lane 6 downstream	99.62%	99.61%	0.01%

As shown in Table 2, the inductive loop detectors produced 48 false positives but only 4 false negatives, whereas the TIRTL detection system produced only 8 false positives but 85 false negatives. Consequently, the loops overcounted the actual number of vehicles by a factor of 0.393%, while the TIRTL undercounted the vehicles by a factor of -0.741%, as shown in Table 3. If this tendency of loops overcounting and the TIRTL undercounting were shown to be a systematic bias, it might be appropriate to introduce a factor corresponding to a particular detection device into the controller software when reporting aggregated counts to the traffic management center. However, more data from different locations would be needed to determine if this effect is local to this installation or if it is common across inductive loop detectors and TIRTL detection systems in general.

As shown in Table 2, most of the false positives reported by the loops were from either the upper channel of slot 1 ("Lane 1 upstream") or the lower channel of slot 2 (Lane 2 downstream"). A significant increase in performance of this vehicle count station might be realized by replacing or adjusting the loop detector cards in those two slots, however, the problem might be in the physical loop wires or detector lead in cables. Figure 8 shows an instance of a false positive in Lane 1 where the output "bounced" three times, however, since this happened over a short period of time, i.e., about the same amount of time as a normal detection, it was only penalized for one false positive in the results. Figure 9 shows a more typical instance of a false positive in Lane 1 where the output was a single pulse.

Figure 10 shows a false positive from both loop detectors in slot 6 ("Lane 6 upstream" and "Lane 6 downstream"). The TIRTL reports this black van in Lane 5, where it appears to be from the video, but not in Lane 6. However, the van can be seen on top of, and perhaps a few inches on the Lane 6 side of, the demarcation striping between Lanes 5 and 6. The loop detectors in slot 6 appear to have been triggered by the van's proximity in this instance of "crosstalk" from Lane 5. The TIRTL, however, was able to report the van as only in Lane 5.

Table 2 – False Positives and Negatives

Detection Zone	False Positives		False Negatives	
	Loops	TIRTL	Loops	TIRTL
All lanes combined	48	8	4	85
Lane 1 upstream	24	0	1	4
Lane 1 downstream	0	0	1	4
Lane 2 upstream	0	0	1	8
Lane 2 downstream	18	0	1	7
Lane 3 upstream	0	1	0	9
Lane 3 downstream	1	1	0	9
Lane 4 upstream	0	2	0	10
Lane 4 downstream	0	2	0	10
Lane 5 upstream	0	1	0	10
Lane 5 downstream	0	1	0	10
Lane 6 upstream	3	0	0	2
Lane 6 downstream	2	0	0	2

Table 3 – Difference in Vehicles Counted

Detection Zone	True Number of Vehicles	Vehicles Counted by Loops	Percent Difference	Vehicles Counted by TIRTL	Percent Difference
All lanes combined	11,196	11,240	+0.393	11,119	-0.688
Lane 1 upstream	606	629	+3.795	602	-0.660
Lane 1 downstream	606	605	-0.165	602	-0.660
Lane 2 upstream	1,561	1,560	-0.064	1,553	-0.512
Lane 2 downstream	1,561	1,578	+1.089	1,554	-0.448
Lane 3 upstream	1,252	1,252	0.00	1,244	-0.639
Lane 3 downstream	1,252	1,252	0.00	1,244	-0.639
Lane 4 upstream	919	919	0.00	911	-0.871
Lane 4 downstream	919	919	0.00	911	-0.871
Lane 5 upstream	741	741	0.00	732	-1.215
Lane 5 downstream	741	741	0.00	732	-1.215
Lane 6 upstream	519	522	+0.578	517	-0.385
Lane 6 downstream	519	521	+0.385	517	-0.385

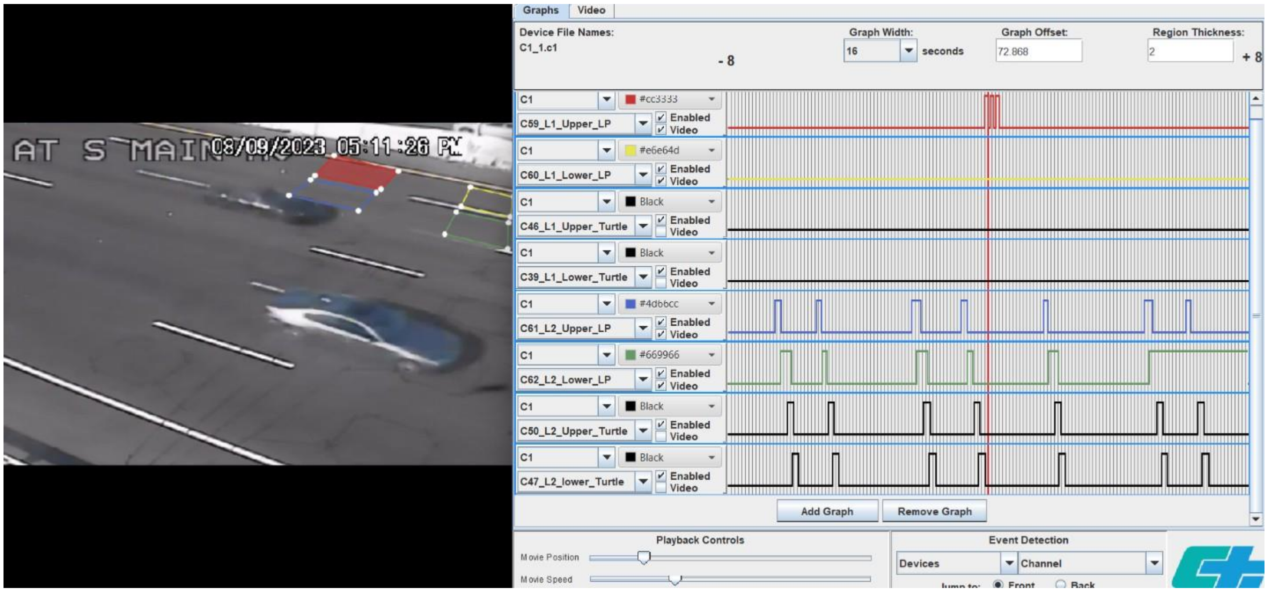

Figure 8 – False positive from Lane 1 upstream loop with output "bounce."

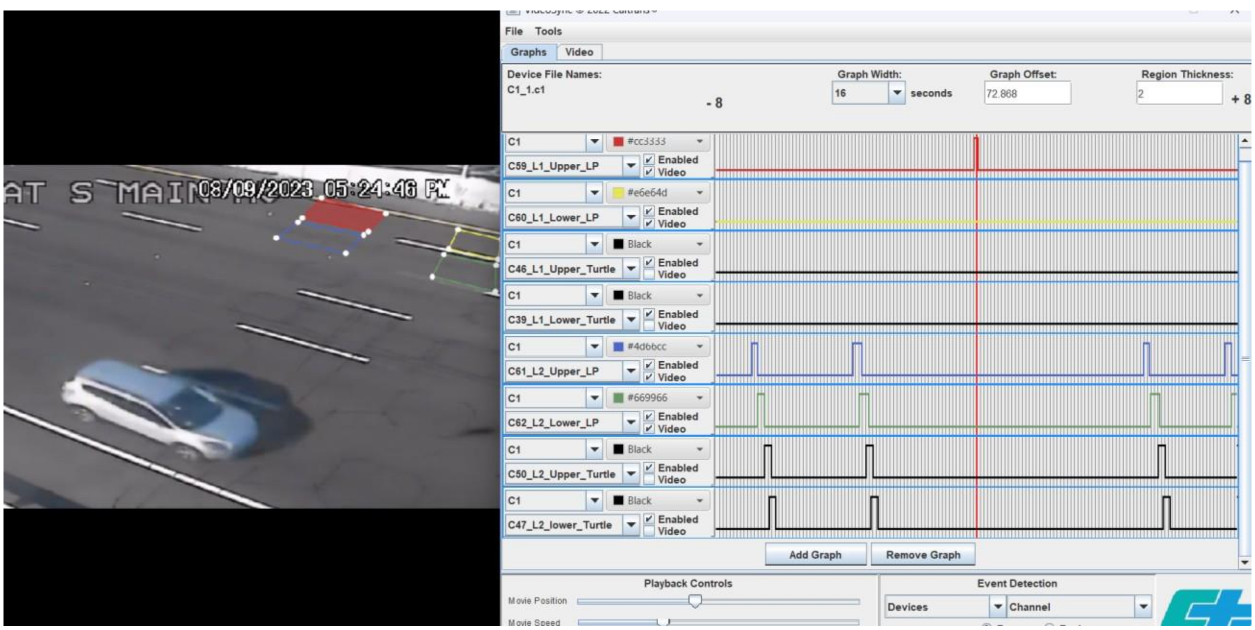

Figure 9 – False positive from Lane 1 upstream loop with typical single pulse output.

Table 2 shows that the loops only reported a total of four false negatives. These happen to be distributed as one each for the upper and lower channels of Lanes 1 and 2, however, this doesn't necessarily mean that both loops missed the same vehicle. For example, Figure 11 shows the upper channel of Lane 1 (in purple) about to detect the white car, but the lower channel of Lane 1 (in orange) never detects it.

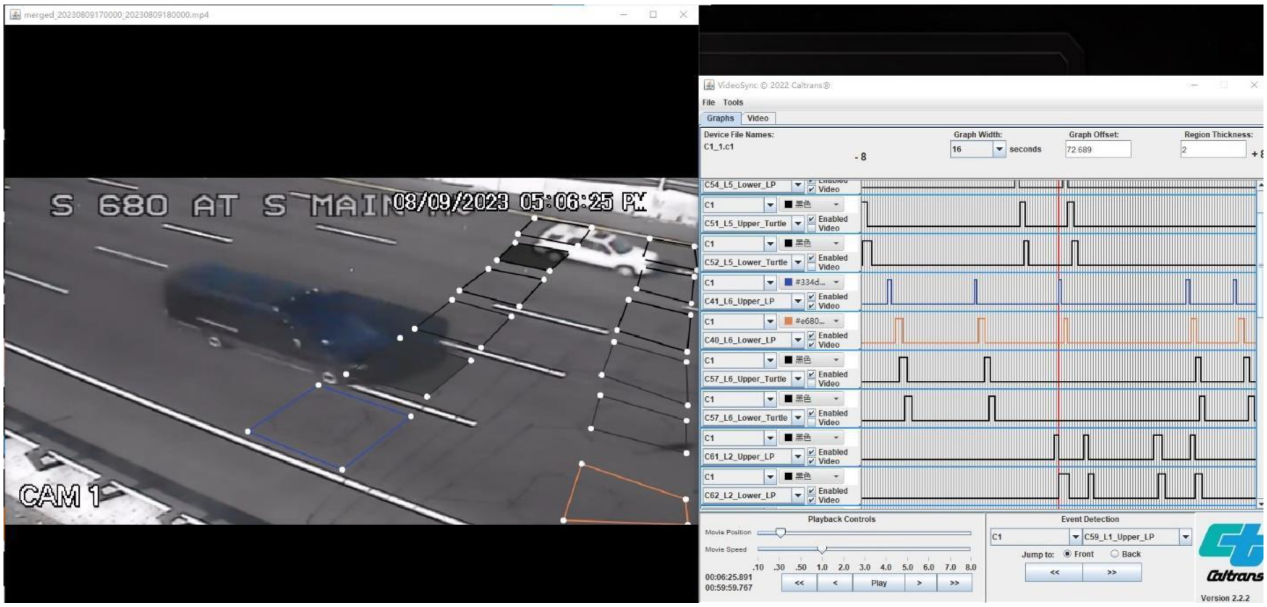

Figure 10 – False positive from the Lane 6 loops due to “crosstalk” from the encroaching van

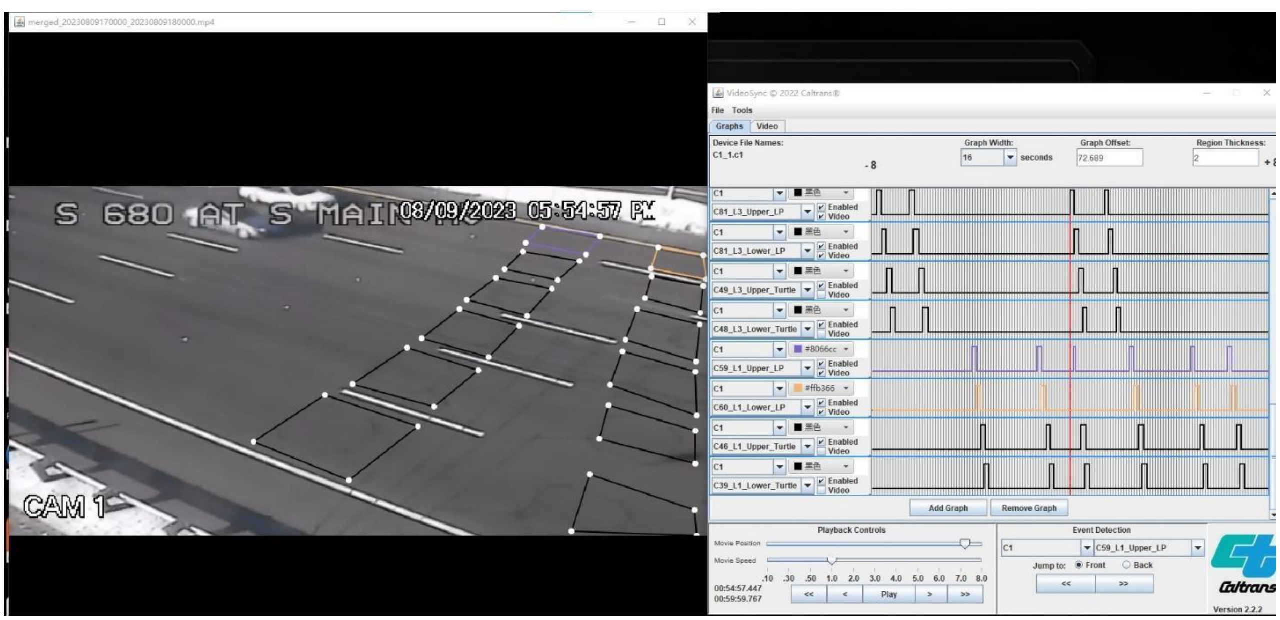

Figure 11 – False negative from the Lane 1 downstream loop

Table 2 shows the TIRTL as having reported a total of 85 false negatives, i.e., its upper and lower channels missed 85 detections. This happened in all lanes, as shown in Figures 12 through 17.

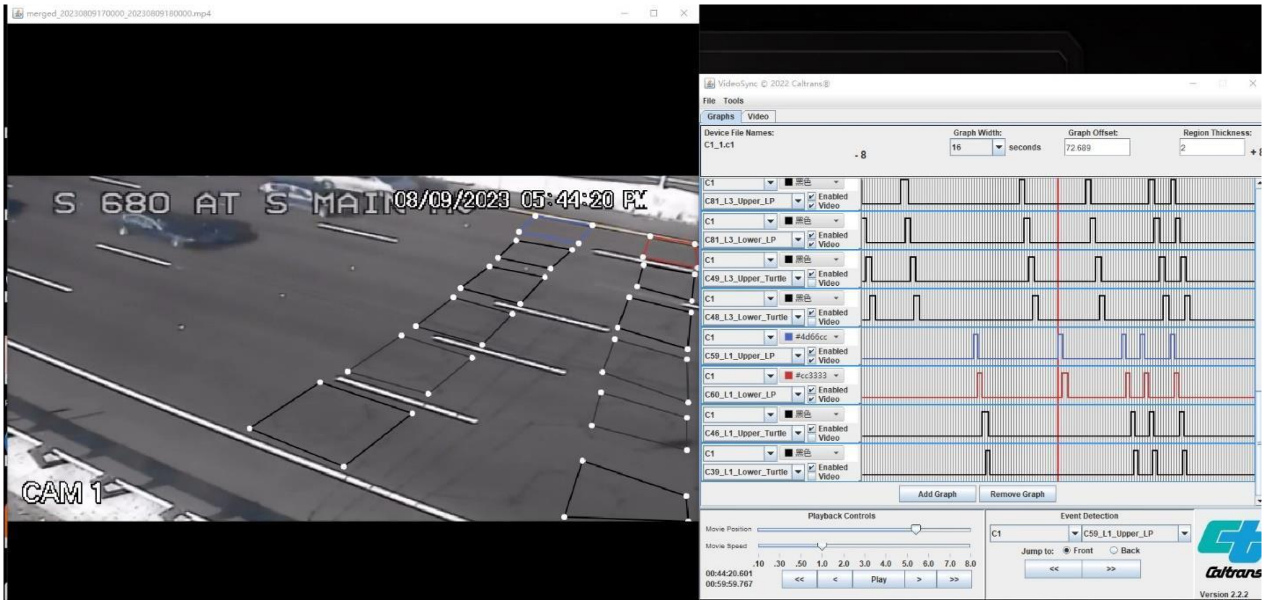

Figure 12 – False negative from the Lane 1 TIRTL (bottom two rows)

Figure 13 – False negative from the Lane 2 TIRTL (underneath yellow row)

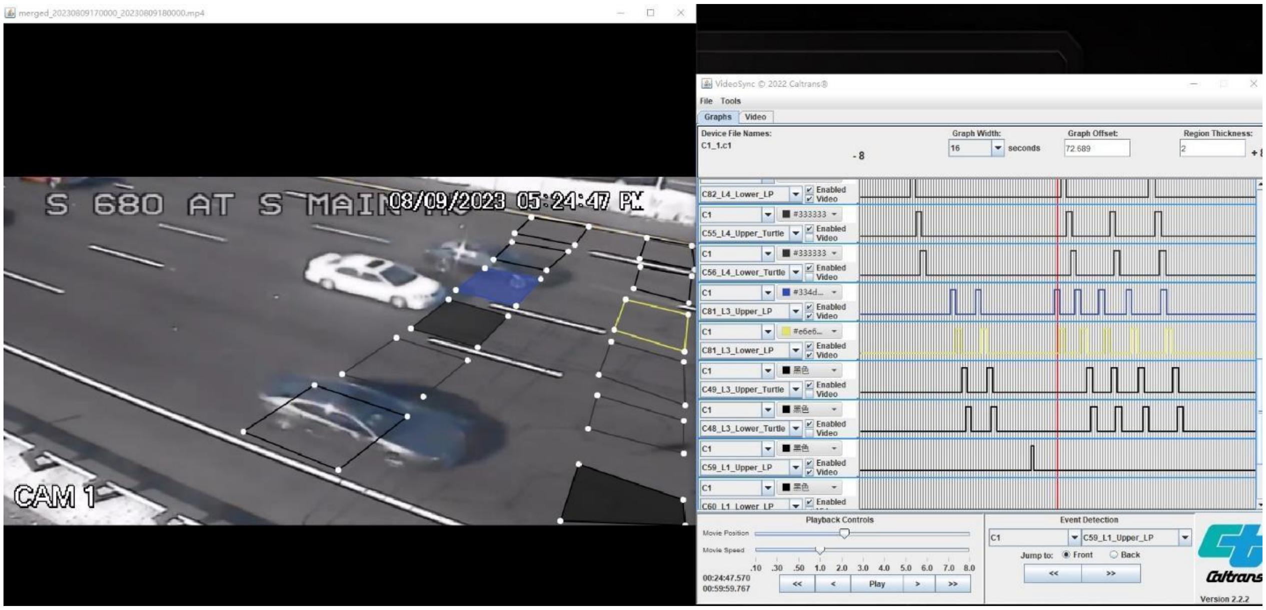

Figure 14 – False negative from the Lane 3 TIRTL (underneath yellow row)

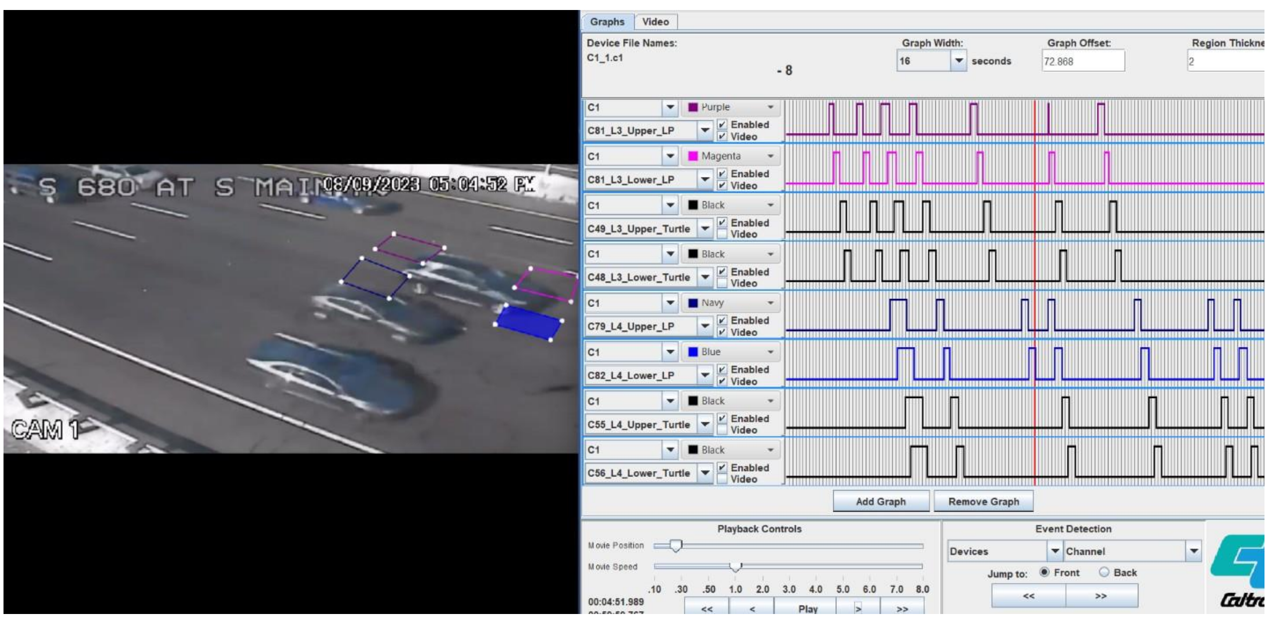

Figure 15 – False negative from the Lane 4 TIRTL (bottom two rows)

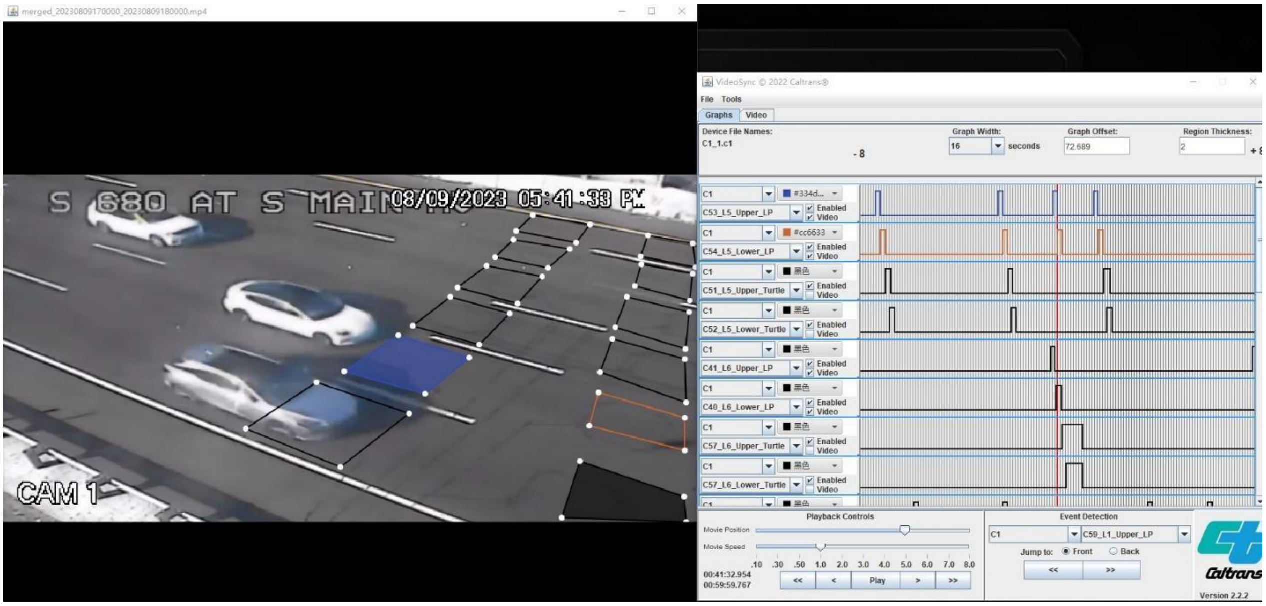

Figure 16 – False negative from the Lane 5 TIRTL (rows 3 and 4)

Figure 17 – False negative from the Lane 6 TIRTL (bottom two rows)

In all the examples in Figures 12 through 17, the TIRTL outputs two detection pulses per lane in order to emulate an upstream and a downstream loop, even though the internal electronics only generate one detection per vehicle. Therefore, the 85 false negatives shown in Table 2 correspond to significantly fewer than 85 missed vehicles. One could argue that the TIRTL is effectively penalized twice for each false positive and false negative because it outputs two pulses per detection, whereas the loops typically falsely detect or miss on either an upstream or a downstream channel and thus only output a single pulse per false positive or false negative. One might also assume that the TIRTL would always output two pulses for a true positive or a false positive and zero pulses for a false negative. However, as shown in Figure 18, this is not always the case. The figure shows a true positive as detected properly by the loops and shown in the video (the light-colored car), but the TIRTL outputs a false negative on its virtual upstream channel and a true positive on its virtual downstream channel. This explains why Table 2 shows an odd number of total false negatives from the TIRTL and why it shows a different number of false negatives from the virtual upstream and downstream channels of the TIRTL for Lane 2. Given the demonstrated possibility of the TIRTL to output a single pulse corresponding to a false negative, it's unclear exactly how to compare its error rate to that of loops. In this report, the collected data is tabulated, and the reader can draw his or her own conclusion.

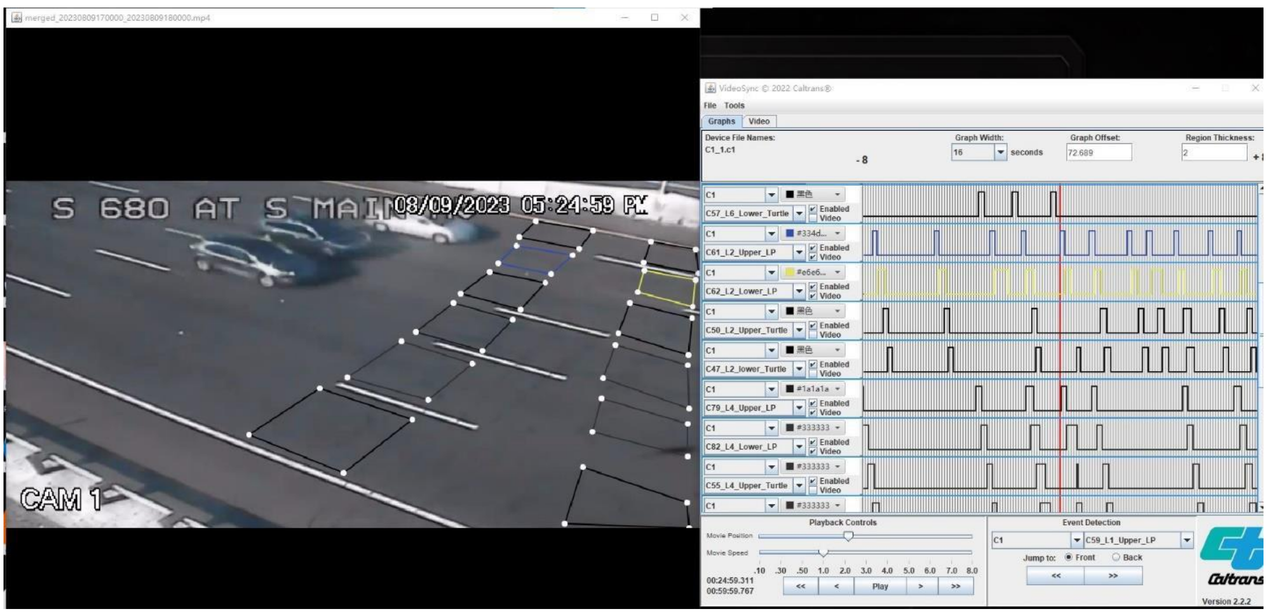

Figure 18 – Both a false negative and a true positive from the Lane 2 TIRTL (rows 4 and 5)

Table 2 shows that the TIRTL only reported a total of eight false positives. One of these is shown in Figure 19. A silver truck is shown in the video in Lane 3, and it is detected properly by both the loops and the TIRTL. The video shows nothing in Lane 4, and indeed, the loops do not report anything, however the TIRTL outputs virtual upstream and downstream pulses.

Figure 19 – False positive from the Lane 4 TIRTL (bottom two rows)

Conclusion

Based on the sensitivity measurements in this test, the TIRTL detection system was very close in accuracy to the inductive loops. However, the loops were still a little more accurate in general. The TIRTL system still offers some advantages over loops. Most notably is the possibility of installation and maintenance without having to close all monitored lanes. Another is the ability to reconfigure lane positions on a roadway, e.g., during construction, without affecting detection operation. The TIRTL is one of a few alternative detection systems available to Caltrans design engineers, and the tradeoffs between these systems should be considered when planning detection installations not based on inductive loops, which DRISI has consistently shown to be the most accurate detection system available.

TIRTL and Sensys are much closer to inductive loops in accuracy (as measured by "Sensitivity") than the microwave radar or machine vision systems that DRISI has also tested. However, microwave radar and machine vision systems can typically be mounted on roadside poles or other existing overhead structures, which allows access for installation and maintenance without having to close lanes. Sensys magnetometers, on the other hand, require lane closures for all lanes in which they are installed, although the procedure is typically significantly faster per lane than that of loop installation. The TIRTL systems on I-680 in Walnut Creek at least required a number one lane closure to install the transmitter in the median K-rail. Effectively, there is a tradeoff between the size and distance to the roadway of detection instrumentation and the accuracy of the corresponding detection systems. In some cases, it may be appropriate to choose ease of installation and maintenance at the expense of accuracy. For example, vehicle count stations ultimately report average speeds and aggregated counts per time-period to users in the traffic management center (TMC). As such, a few missed or double-counted vehicles from the detection system wouldn't significantly degrade their overall functionality. In fact, if a detection system happened to report the same number of missed and double-counted vehicles to the controller in a certain aggregation period, the controller's output to the TMC would be the same as if the detector had been 100% accurate. In any case, a detector that output both double counts, i.e., false positives, and missed vehicles, i.e., false negatives, would result in an aggregated count closer to the truth than a detector that only output one or the other, assuming a comparable number of erroneous outputs by both detectors. In other cases, it may be appropriate to choose accuracy at the expense of ease of installation and maintenance. For example, detection for traffic signals needs to be more accurate and reliable that detection for count stations because even a single missed detection, i.e., false negative, could result in a skipped phase and consequently, a waiting vehicle not being served. This would be more likely in very low traffic conditions, e.g., late at night with only one vehicle present to actuate a phase. Likewise, a single false positive could result in a phase being served or extended for no reason and consequently, a delay in service for vehicles waiting for other phases. Therefore, the suitability of a given type of vehicle detection system to a particular application should be determined by the design engineer on a case-by-case basis.

The scope of this test was limited to the use of the TIRTL system for vehicle count stations, for which detection accuracy and reliability are less critical than for other applications such as traffic signals and ramp meters. For this test, DRISI analyzed data collected under heavy traffic conditions, during the "PM peak," from 5 to 6 PM. DRISI also collected the data during August, when ambient reflected heat would be likeliest to present a challenge to the optics of the system. Despite these conditions, the TIRTL system performed impressively, with "Sensitivity" numbers very close to those of the inductive loops. In some Caltrans districts, count stations are configured to monitor only a single row of loops. Had the TIRTL been installed at this type of count station, it could have been configured to output only a single pulse per vehicle and, according to the methodology of this test, would only have been penalized for a single false negative or false positive per erroneous vehicular detection. In this case, its error rate would compare even better to that of a single row of loops than it did for the double row of loops in this test. It is therefore the opinion of the author that the TIRTL detection system is indeed fit for the purpose of replacing inductive loops for vehicle detection at Caltrans vehicle count stations.