TRB Paper number: 990550

Using Dual Loop Speed Traps to Identify Detector Errors

Benjamin Coifman

Institute of Transportation Studies

University of California

Berkeley, California, 94720

zephyr@eecs.berkeley.edu

http://www.cs.berkeley.edu/~zephyr

voice: (510) 848-5121

Fax: (510) 642-1246

Submitted for presentation and publication at the 1999 Transportation Research Board meeting

REVISED: March 25, 1999

3756 words + (10 figures + 2 extra pages of figures + 2 tables) * 250 = 7256 words

ABSTRACT

Dual loop speed traps have a distinct advantage over single loop detectors because the speed trap detection system is redundant. Each vehicle is observed twice under normal operating conditions, once at each loop. The two observations are normally used to measure velocity, but as this paper demonstrates, the redundancy can also be used to assess the performance of the speed trap and identify detector errors. At free flow velocities, the time each detector is occupied by a vehicle (i.e., the on-time) should be virtually identical, regardless of the vehicle length. Many hardware errors will cause the two on-times to differ. Exploiting this property, we develop a formal methodology for testing speed traps off-line and suggest ways to extend the work to on-line testing. The work is used to evaluate several loop sensor units, revealing problems in two models. A second example shows how the work can be used to detect crosstalk between sensor units.

KEY WORDS

Traffic Surveillance, Freeway Traffic, Inductive Loop Detectors, Error Detection

INTRODUCTION

Inductive loop detectors have been the preeminent vehicle detector for the past several decades and most traffic surveillance applications depend on these detectors. Surprisingly, very little work has been done to develop a method to assess the quality of loop data. Many operating agencies use specialized loop testers to assess the quality of the wiring [1-2], but these tools bypass the controller and loop sensors; thus, they do not analyze the entire detector circuit, nor do they analyze the circuit in operation. To this end, most operating agencies employ simple heuristics such as, "Do the loop sensor indicator lights come on as a vehicle passes?" or simply, "Do the time series 30 second average flow and occupancy seem reasonable to the eye?" Such tests are typically employed when the loops are installed or when the quality of data coming from the detector station is questionable. These heuristics will catch severe errors and help diagnose them, but other problems can easily go unnoticed.

Many practitioners and some researchers [3-5] have worked to formalize the latter heuristic by rephrasing the question, "Are the time series 30 second average flow and occupancy within statistical tolerance?" These strategies are automated and are designed to operate 24 hours a day; thus, they can detect problems before they are noticed by a human operator. These systems often go undocumented in the literature because they are either designed in-house by an operating agency (see [6] for examples) or were developed by a consulting firm using proprietary information. Because these automated systems only use aggregated data, they must accept a large sample variance and potentially miss problems altogether. For example, the systems have to tolerate a variable percentage of long vehicles in the sample population. As the percentage of long vehicles increases, the occupancy/flow ratio should increase simply because a long vehicle *occupies* the detector for more time compared to a shorter vehicle traveling at the same velocity.

Chen and May developed a new approach for verifying detector data using individual vehicle actuations [6]. Their methodology examines the distribution of vehicles' *on-time*, i.e., the time the detector is occupied by a vehicle. Unlike conventional aggregate measures, their approach is sensitive to errors such as "pulse breakups", where a single vehicle registers multiple detections because the sensor output flickers off and back on.

All of these preceding error detection strategies assume the data come from single loop detectors and fail to use all of the information available from dual loop speed traps. It is a simple fact that as a vehicle passes over a speed trap, it yields two measurements, one from the upstream loop and one from the downstream loop. At free flow velocities, it takes a vehicle about 0.25 seconds to traverse the speed trap. Even allowing for hard decelerations, the on-times from the two loops should be virtually identical, regardless of vehicle length. Many hardware and software errors will cause the two on-times to differ. At lower velocities, vehicle acceleration can cause the two on-times to differ even though both loops are functioning properly. Our research exploit these

properties to develop an off-line diagnostic for dual loop speed traps that tests the entire detector circuit, including the controller and sensors. The approach could easily be incorporated into controller software for on-line diagnostics as well.

After presenting the basic error detection methodology, this paper examines two applications: evaluating loop sensor units, and detecting crosstalk between sensors. The former reveals previously unknown problems in two different sensor models, while the latter example provides a useful tool for sensor installation. The final section presents a possible on-line realization of this methodology for continuous error detection.

METHODOLOGY

The beauty of this error detection approach is in its simplicity. It can be summarized in three steps: (1) record a large number of vehicle actuations during free flow traffic, (2) for each vehicle, match actuations between the upstream and downstream loops in the given lane, (3) take the difference between matched upstream and downstream on-times and examine the distribution on a lane-by-lane basis. Assuming the loops are functioning properly, only a small percentage of the differences should be over 1/30 seconds.

For the first step, the controller records vehicle arrival and departure times from each loop, i.e., the sensor output or event data, as shown in Figure 1. For this study, we used 170 type controllers running a program developed by Caltrans District 4. Rather than aggregating the detector data, the program sends the event data to the modem port once every second. A laptop computer collects the data and saves it to disk. The data is later analyzed off-line; however, the analysis is simple enough that it could be easily incorporated in to the controller software.

Next, it is necessary to match events from a given vehicle between the upstream and downstream loops. To this end, the study uses a simple heuristic: match each downstream pulse to the most recent upstream pulse. This heuristic emulates conventional strategies for measuring velocity at speed traps. To illustrate this matching, consider the pulses shown in Figure 2. In part A, pulse 1down is matched to 1up, 2down to 2up, etc.. In part B, pulse 6down is correctly matched to 6up, but, upstream pulse 5up is never considered because it was the first of two successive upstream pulses. Finally, in part C, pulse 7down is correctly matched to 7up, unfortunately, pulse 8down is also matched to 7up because there was no intervening upstream pulse between the two successive downstream pulses. When the speed trap is operating properly, two successive pulses rarely come from one loop without an intervening pulse on the other loop (the loops are typically spaced close enough to ensure that one vehicle will actuate both loops before the next vehicle actuates the upstream loop). As we will soon see, the error detection strategy is sensitive to unmatched pulses and it will detect when this assumption breaks down.

Our research considered a more rigorous approach to matching pulses: identify and eliminate any unmatched pulses, or more specifically, we identify two or more successive pulses from one loop without an intervening pulse on the other loop. To illustrate this process, in Figure 2B, the two successive upstream pulses would cause the algorithm to remove 5up, 6up and 6down (there are cases when pulse 6down could match pulse 5up). While in Figure 2C, the two successive downstream pulses would cause the algorithm to remove 7up, 7down and 8down. This approach will produce better matches, but it is not clear how to treat the unmatched pulses. Hence, the remainder of this paper will use the simple heuristic, matching a downstream pulse with the most recent upstream pulse, which has proven to be sufficient for the analysis.

After matching pulses between loops, the following parameters are calculated for each vehicle: upstream detector on-time (OT_u), downstream detector on-time (OT_d), travel time via the rising edges (TT_r), and travel time via the falling edges (TT_f). Specifically:

$$
\begin{aligned}
OT_u &= t_{FALL\_up} - t_{RISE\_up} \\
OT_d &= t_{FALL\_down} - t_{RISE\_down} \\
TT_r &= t_{RISE\_down} - t_{RISE\_up} \\
TT_f &= t_{FALL\_down} - t_{FALL\_up}
\end{aligned}
\tag{1}
$$

using the four events indicated in Figure 1B: t_{RISE_up}, t_{FALL_up}, t_{RISE_down}, t_{FALL_down}. The controller samples at 60 Hz; so these events are, at best, accurate to 1/60 of a second. Note that an individual vehicle's velocity is simply the loop separation divided by the travel time. For this study, individual vehicle velocity is used to determine if a vehicle is free flowing (velocity > 64 km/h (40 mph)) or if the vehicle should be excluded from further analysis. In this fashion, error detection is applied to all vehicles where OT_u and OT_d should be similar and turned off when vehicle accelerations could cause these parameters to differ significantly.

Finally, taking the difference between OT_u and OT_d for a large number of vehicles during free flow traffic conditions, it is possible to assess the performance of the speed trap. This assessment is presented via the examples in the following sections as we examine four of the five loop sensors from Caltrans' Qualified Products List.

EVALUATING LOOP SENSOR UNITS

This study arose from the need to diagnose the source of poor quality speed trap data. Several applications have been developed using event data from speed traps, e.g., [7]. When attempts

were made to transfer the applications from the large pre-existing Freeway Service Patrol (FSP) database [8] to new data collected using the same tools, the results were unexpected. All of the hardware and software appeared to be identical in both cases, but the new data was unacceptable because OT_u and OT_d differed significantly for free flow vehicles. Investigation revealed that the loop sensor units had been updated and the new revision seems to be faulty. In particular,

Peek GP5 revision E (GP5-E) [the old sensor units].
Peek GP5 revision G (GP5-G) [the new sensor units].

Preliminary work compared data collected using the GP5-G from speed trap stations on State Highway 99 south of Sacramento, CA; I-680 east of Oakland, CA; and State Highway 24 east of Oakland, CA, against the FSP data collected using the GP5-E at 19 speed trap stations on I-880 south of Oakland, CA. Since identifying the apparent source of the problem, the results have been reproduced at four speed trap stations on I-80 in Berkeley, CA. This report will present data from the Berkeley test site, which used 15 GP5-G units and 36 GP5-E units deployed in various configurations at the speed trap stations. Use of the test site allowed for better control because we could switch sensor units while holding all of the other variables constant, such as the physical loops and the controller unit. The work has since grown to include the following "Model 222" sensor units:

Peek GP6 revision C (GP6), 27 units.
Eberle Design Inc. LM222 (EDI), 20 units.
Intersection Development Corporation Model 222 (IDC), 5 units.

The following subsections examine each detector unit in detail. All units were tested for several days at multiple stations. For illustrative purposes the analysis in this report is limited to approximately three hours for each unit from a single speed trap. The given sample period was selected at random and the samples are representative of each unit's performance throughout the study. Although each section presents a single unit, the data were collected with a full rack of similar sensor units: the GP5's, GP6 and EDI sensors were deployed across all 9 lanes at the subject station, while the IDC sensors were deployed across as many lanes as possible. The study uses Caltrans District Four's control cabinet wiring practice: both loops from a given lane are wired to the same sensor unit with the upstream loop connected to channel one and the downstream loop connected to channel two. Unless otherwise noted, sensor frequencies and sensitivity levels were set according to District Four's standard practice. In all cases, the sensors were set to presence mode.

GP5-G Sensors

Figure 3A shows a scatter plot of downstream on-time versus upstream on-time for 2803 matched pulses during free flow traffic using a GP5-G sensor. As one would expect, most of the points fall on the central axis, a line passing through the origin at 45 degrees. Surprisingly though, two side bands are clearly evident, offset by 1/10th of a second from the central band. The scatter plot shows that the problem seems to be independent of on-time, but, the magnitude of the problem is obscured because many observations fall on each point. Plotting the difference between OT_u and OT_d for these same vehicles produces the histogram shown in Figure 3B and in detail in Figure 3C.

Some spreading of the central band is expected. The simplest metric for speed trap assessment is the percentage of on-time differences greater than a threshold value. This study assumes that a difference of ± two controller samples is acceptable, i.e., ±2/60 seconds. These measurements are shown with lighter bars in the histograms, while all other measurements are indicated with darker bars and will be referred to as "bad observations". The threshold value was selected to be liberal, the reader can estimate the results for a different threshold directly from the histograms.

In the case of the GP5-G sample, 17.05% of the observations are bad. Examining the cumulative distribution function for TT_r and TT_f, as shown in Figure 3D, we see the rising edge distribution is much tighter. Because the two distributions come from the same pulses and the rising edge travel times are in the range of what would be expected for free flow velocities (the loops are separated by 6.1 m (20 ft)), the problem is limited to the falling edge. This pattern is typical of the performance of all GP5-G's examined and it appears that the sensors randomly hold on to a pulse for 1/10th of a second after the vehicle leaves the loop's detection region.

GP5-E Sensors

In contrast to the unexpected performance of the later revision, Figure 4A shows a scatter plot of downstream on-time versus upstream on-time for 3067 matched pulses during free flow traffic using a GP5-E sensor. In this case, virtually all of the observations fall on the central axis with a few spurious points due to mismatching. Figure 4B shows the distribution of the difference between on-times. The GP5-E performs quite well, only 0.59% of the observations are bad in this sample, which was typical of all GP5-E's examined.

GP6 Sensors

The GP6 performance is similar to the GP5-E. Figure 5 shows a scatter plot and on-time histogram for 2892 matched pulses using a GP6 sensor. For this particular sample, 0.24% of the observations are bad.

EDI Sensors

Figure 6A shows a scatter plot of on-times for 3007 matched pulses using an EDI sensor. Again, most of the observations fall on the central axis of the scatter plot, but a significant number of pulses seem to be subject to noise. Examining the histogram shown in Figure 6B and in detail in 6C, 5.45% of the pulses are bad for this sample. The source of this noise only becomes evident when looking at the original pulses. Figure 7A shows a detail of three vehicles passing over the speed trap. Pulse breakup is clearly evident for the first vehicle at the downstream loop. This flickering appears to be due to sensor crosstalk, because at all four stations the problem is most severe in the center lane of a given direction and diminishes towards the outer lanes. The data presented in this subsection are typical of EDI sensors in the center lane. Note that these data were collected with the sensitivity set to the lowest level, two steps lower than Caltrans standard practice of using sensitivity level 2 on a scale of 0-7. Finally, the impact from pulse breakup on the velocity measurements can be seen in Figure 7B. Without pulse breakup, all of the measurements would be close to 100 km/h (62 mph).

IDC Sensors

Figure 8 shows a scatter plot and on-time histogram for 4038 matched pulses using an IDC sensor. For this particular sample, 1.26% of the observations are bad.

Comparisons and Comments

Table 1 summarizes sample size and percentage of bad pulses from each sensor unit. The GP5-G and EDI sensors do not appear to be acceptable for the speed trap installations at which they were tested.

Why did these problems go unnoticed until now? Most speed trap controllers and testing procedures aggregate data to 30 seconds or longer, obscuring the problems presented in this section. The problems do not become evident until one examines individual vehicles, which is well beyond the range of typical freeway surveillance applications. In fact, the problem with the

GP5-G is particularly difficult to detect because it occurs randomly on the falling edge and it seems to affect both channels of the sensor unit evenly.

Finally, the reader should note that this analysis can not detect all problems. Any errors that occur at both loops for a given vehicle will go undetected, e.g., if pulse breakup occurs evenly at both loops or if a vehicle is not detected by either loop.

DETECTING CROSSTALK

This section presents a sample application of the error detection strategy: detecting crosstalk between two sensors. Other researchers have examined physical characteristics that cause crosstalk, e.g., [9], but they do not offer a means of detecting the presence or absence of crosstalk. The *Traffic Detector Handbook* [1] provides some advice for correcting crosstalk and localizing which unit is "cross-talking" (see p. 148 of [1]), but again, it does not offer any means of detecting the presence or absence of crosstalk beyond using a specialized loop tester that bypasses the controller and loop sensors. This omission can miss a common source of crosstalk, two sensors set to similar frequencies. On the other hand, error detection strategies based on aggregate measures will miss crosstalk problems unless the condition is severe.

Because each loop of a speed trap should be set to a different operating frequency, even a small amount of crosstalk should affect each loop differently and thus, be detectable via the differencing method presented in this paper. To illustrate the detection process, consider the following example: ten GP5-E sensor units were installed at a loop detector station and event data were collected overnight to evaluate the installation. When examining the on-time histogram for each lane, eight resembled Figure 4B; however, two neighboring lanes showed anomalies. Figure 9A shows a detail view of the histograms for these two lanes. Approximately 6% of the observations from either lane were bad (see Table 2 for the statistics). To address the problem, three sensor units were rotated in the rack (the two problem lanes/slots and one adjacent lane/slot) and then the frequencies for each lane were reset to the lane's values before rotation. In other words, the lane by lane sensor settings remained unchanged, but three physical sensor units changed position in the rack. After the rotation, the histograms from all lanes resembled Figure 4B. Figure 9B shows a detail of the histograms for the subject lanes, now, fewer than 1% of the observations are bad in each lane.

In practice, the lane-by-lane analysis could be incorporated into the controller software and the controller could report each lane's status. When a problem occurs, a technician could simply adjust the sensor frequencies or collect a few hours of event data for off-line analysis.

DEVELOPING AN ON-LINE DIAGNOSTIC

All of the preceding analysis was conducted off-line using event data collected from the speed trap stations. As such, the analysis can only serve to verify the status of the speed traps at the time the data were collected. The approach is simple enough that it could easily be incorporated into controller software for continuous status monitoring and thus, detect temporary loop failures (e.g., crosstalk due to environmental changes) as well as permanent loop failures.

One possible on-line realization of this diagnostic is shown in Figure 10. The most recent actuation from both loops are used to measure travel time over the speed trap and thus, individual vehicle velocity. If the vehicle is free flowing, the difference between on-times is added to a running distribution for the given lane. Otherwise, the difference is not included because acceleration, which can not be measured, becomes a significant factor in the difference. Note that there is no need to store individual vehicle data, just the number of vehicles in each bin of the distribution. The bins can simply be "good" versus "bad", where the division is some reasonable threshold, or they can be the actual time differences. After a fixed number of free flow vehicles an assessment would be made for the lane and the bins would be reset to zero.

CONCLUSIONS

This paper has presented a new approach for detecting data errors at dual loop speed traps. The method is meant to complement existing tests, rather than supplanting them. The examples have shown that this is an important process, both by detecting previously unknown flaws in two sensor models, and by providing a means for detecting crosstalk problems. These errors were only discovered because we examined individual vehicle data from speed traps. Practitioners may want to consider this fact and assess the performance of their speed traps and sensor units using the new tests.

ACKNOWLEDGMENTS

The author would like to thank Sean Coughlin and Joe Palen of Caltrans for their assistance in collecting data for this research.

REFERENCES

[1] Kell, J., Fullerton, I., and Mills, M., *Traffic Detector Handbook, Second Edition*, Federal Highway Administration, Washington, DC, 1990.

[2] Ingram, J., *The Inductive Loop Vehicle Detector: Installation Acceptance Criteria and Maintenance Techniques*, California Department of Transportation, Sacramento, CA, 1976.

[3] Jacobson, L, Nihan, N., and Bender, J., "Detecting Erroneous Loop Detector Data in a Freeway Traffic Management System," *Transportation Research Record 1287*, TRB, Washington, DC, 1990, pp 151-166.

[4] Cleghorn, D., Hall, F., and Garbuio, D., "Improved Data Screening Techniques for Freeway Traffic Management Systems," *Transportation Research Record 1320*, TRB, Washington, DC, 1991, pp 17-31.

[5] Nihan, N., "Aid to Determining Freeway Metering Rates and Detecting Loop Errors", *Journal of Transportation Engineering*, Vol 123, No 6, ASCE, November/December 1997, pp 454-458.

[6] Chen, L., and May, A., "Traffic Detector Errors and Diagnostics" *Transportation Research Record 1132*, TRB, Washington, DC, 1987, pp 82-93.

[7] Coifman, B. "Vehicle Reidentification and Travel Time Measurement in Real-Time on Freeways Using the Existing Loop Detector Infrastructure", *Transportation Research Record no. 1643*, Transportation Research Board, 1998, pp 181-191.

[8] Skabardonis, A., et al, *Freeway Service Patrol Evaluation.* University of California PATH, Berkeley, CA, 1995.

[9] Bhagat, V., and Woods, D., "Loop Detector Crosstalk", *ITE Journal*, Vol 67, No 2, Institute of Transportation Engineers, February 1997, pp 36-37, 47-49.

TABLE 1: Summary of Sensor Unit Statistics.

Sensor Unit	Number of vehicles in histogram	Percentage of bad observations
GP5-G	2803	17.05%
GP5-E	3067	0.59%
GP6	2892	0.24%
EDI	3007	5.45%
IDC	4038	1.26%

TABLE 2: Statistics for the Crosstalk Example.

		Number of vehicles in histogram	Percentage of bad observations
With crosstalk			
	lane 3	2067	5.61%
	lane 4	1705	6.69%
After correcting the problem			
	lane 3	4496	0.53%
	lane 4	3544	0.76%

FIGURE 1: One vehicle passing over a speed trap, illustrating the four critical time measurements.
(A) Time space representation showing the loop detectors and both ends of the vehicle
(B) Detector output, yielding the upstream and downstream, rising and falling edges

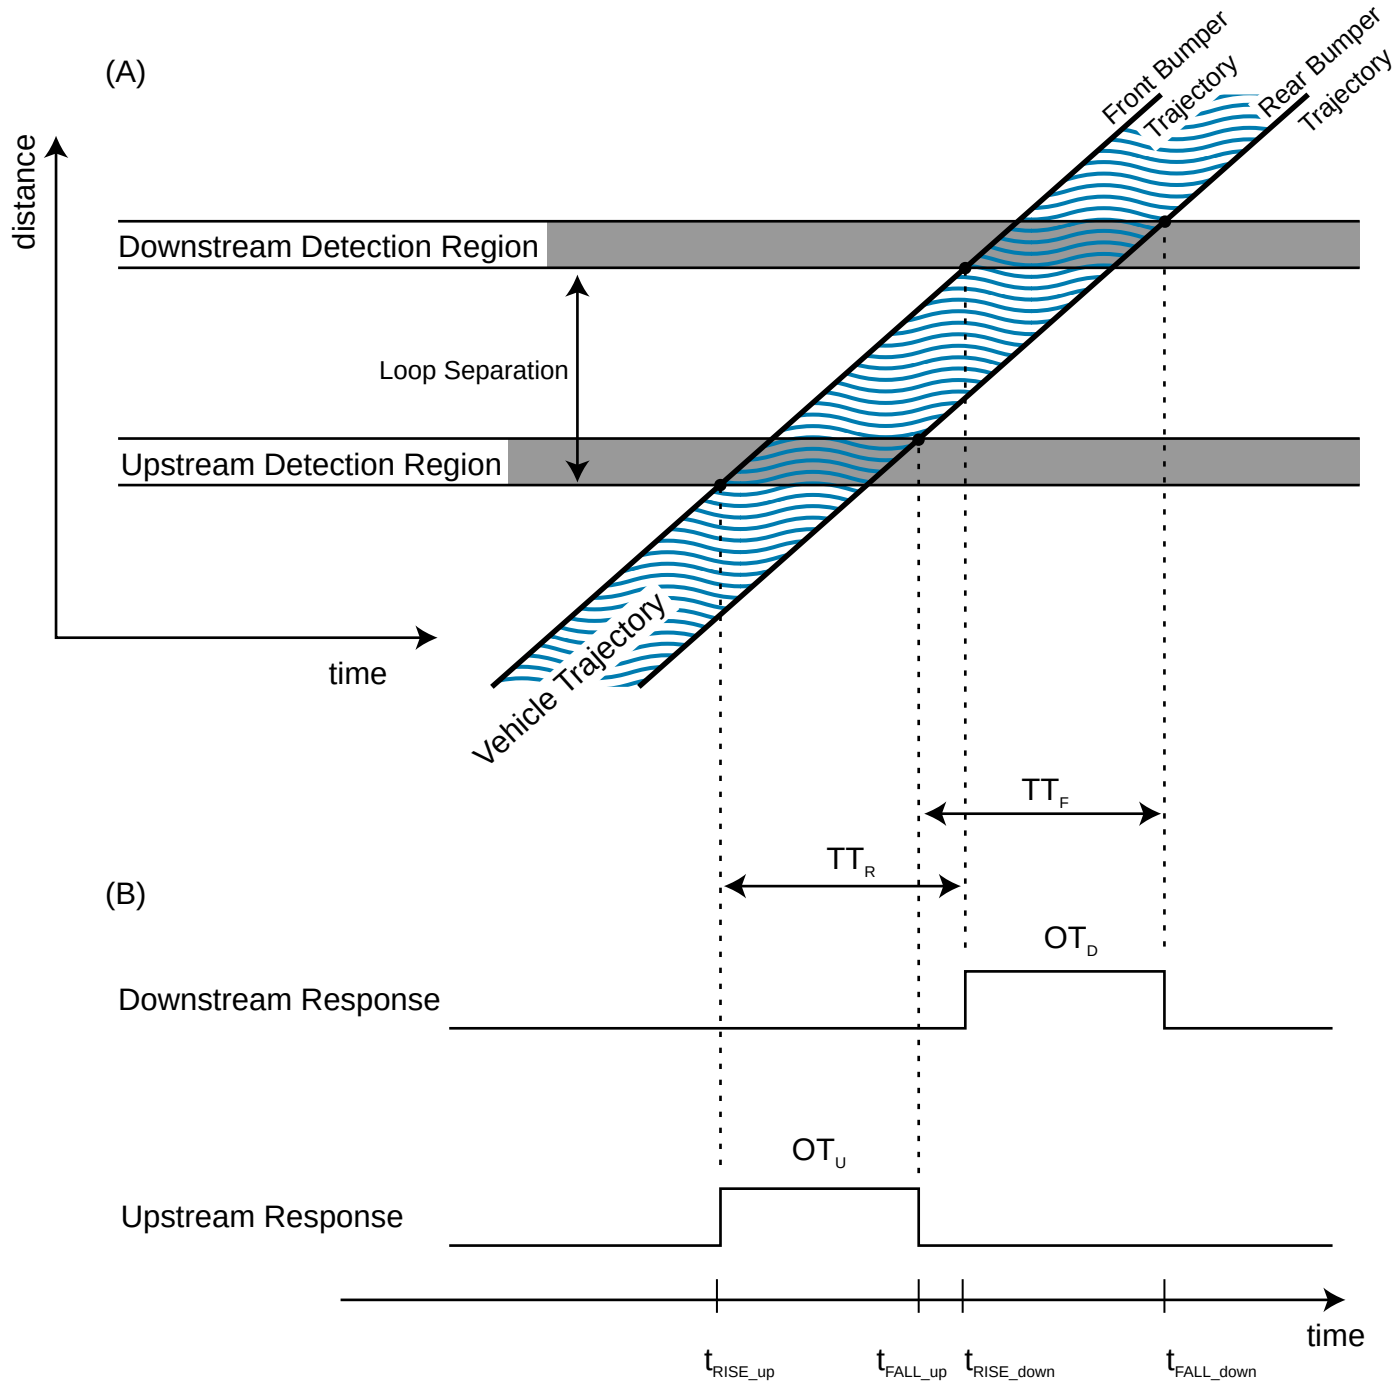

FIGURE 2: Sample loop actuations for three consecutive vehicles, (A) correctly detected at both loops, (B) vehicle 5 is not detected at the downstream loop, (C) vehicle 8 is not detected at the upstream loop.

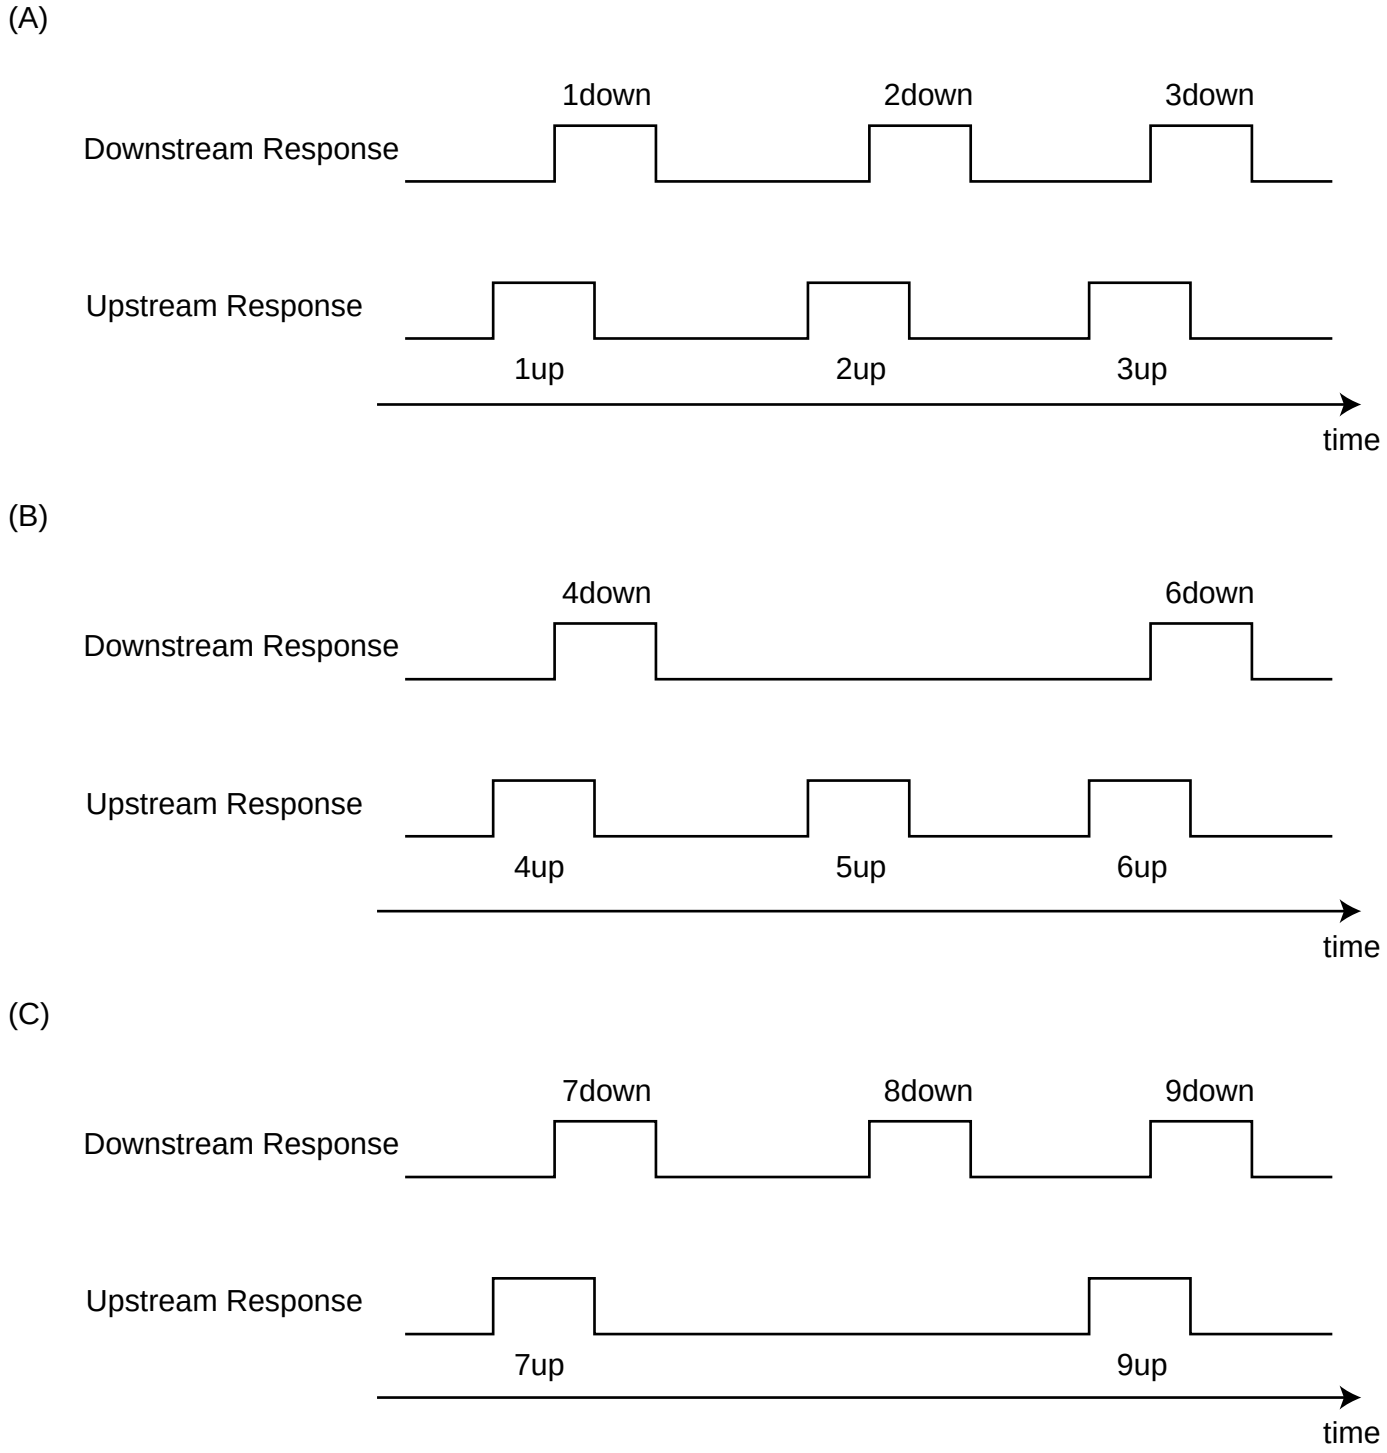

FIGURE 3: (A) Scatter plot of on-time for 2803 matched pulses using a GP5-G sensor (B) Distribution of OT_u - OT_d for the same pulses.

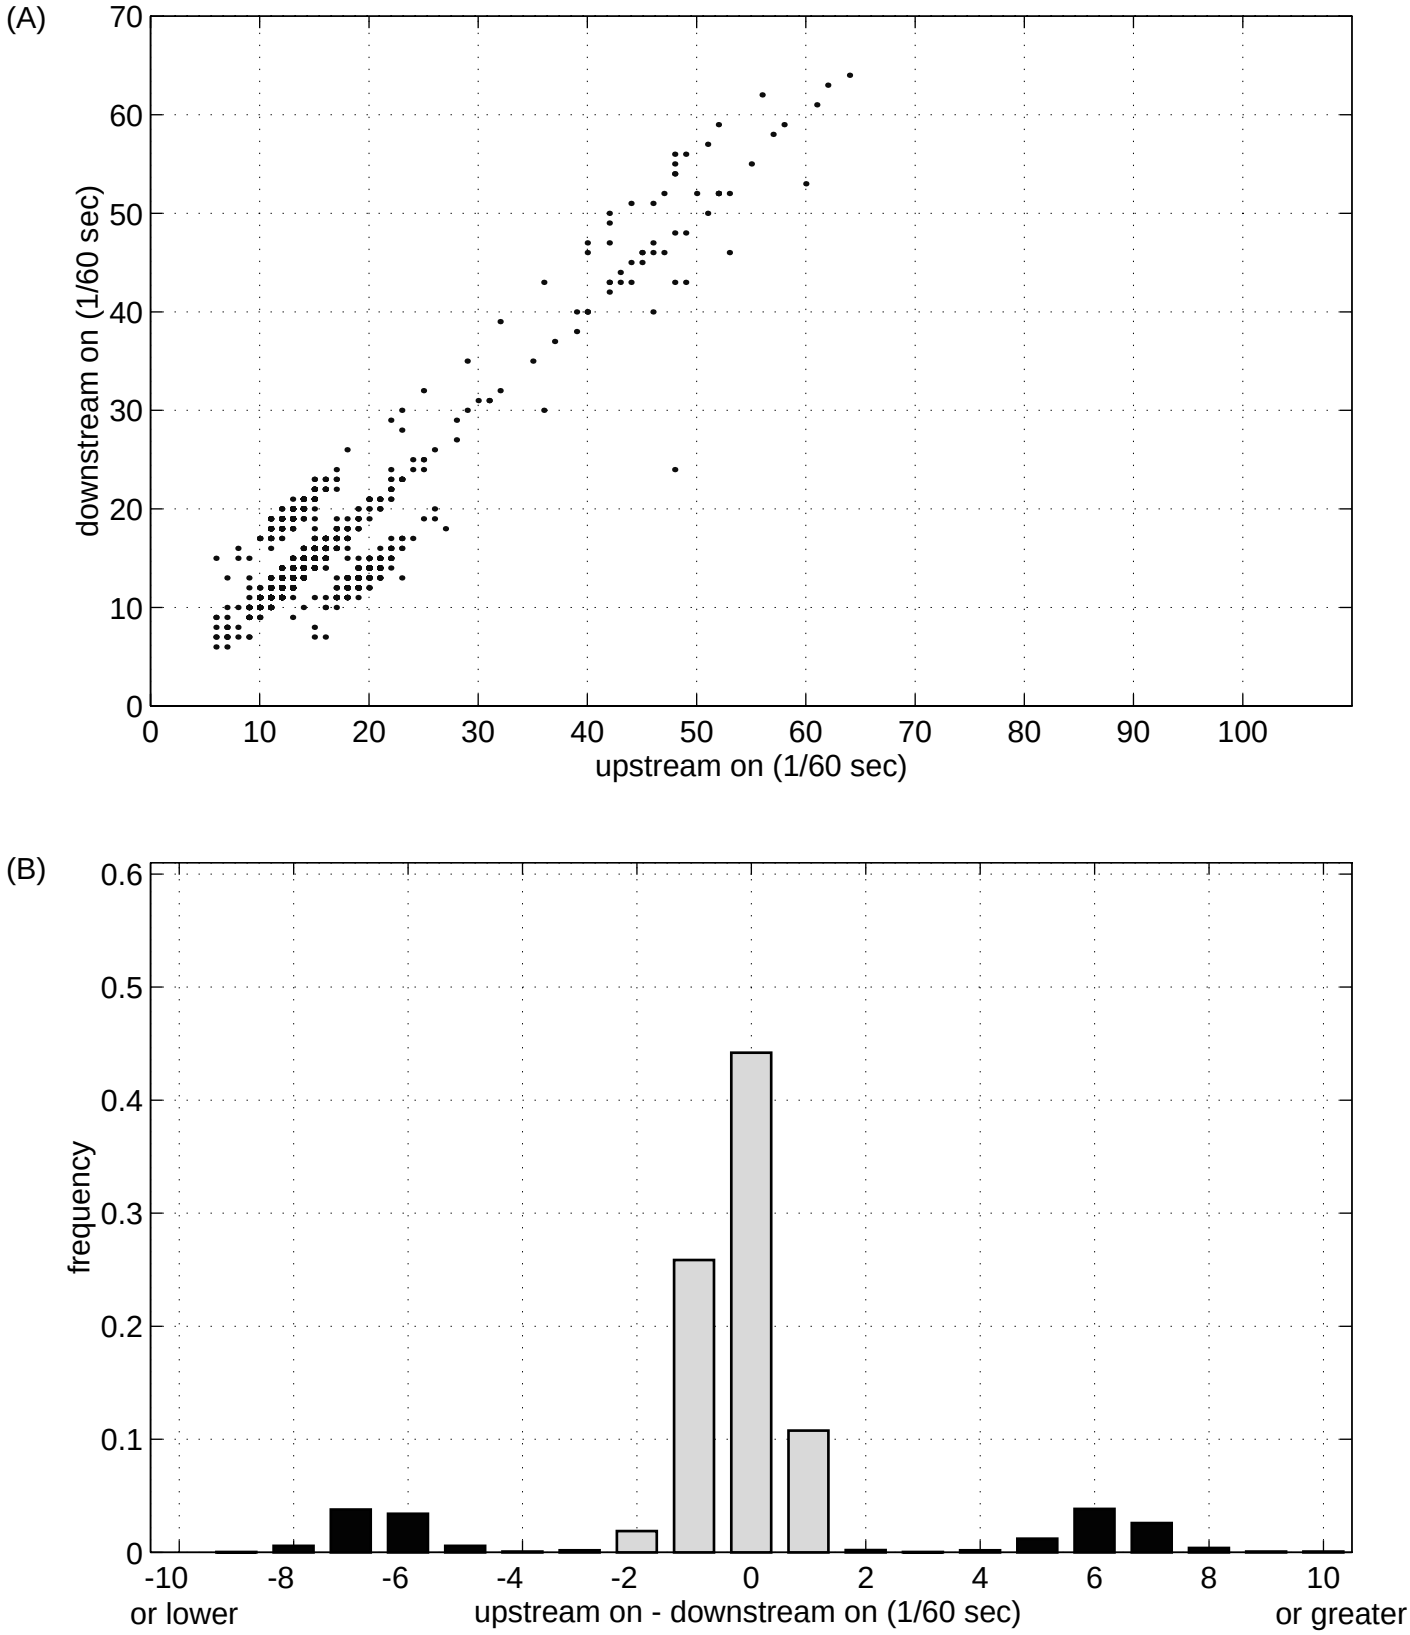

FIGURE 3: (C) detail of part B (D) CDF of travel time between loops (TT_r and TT_f) for the same pulses.

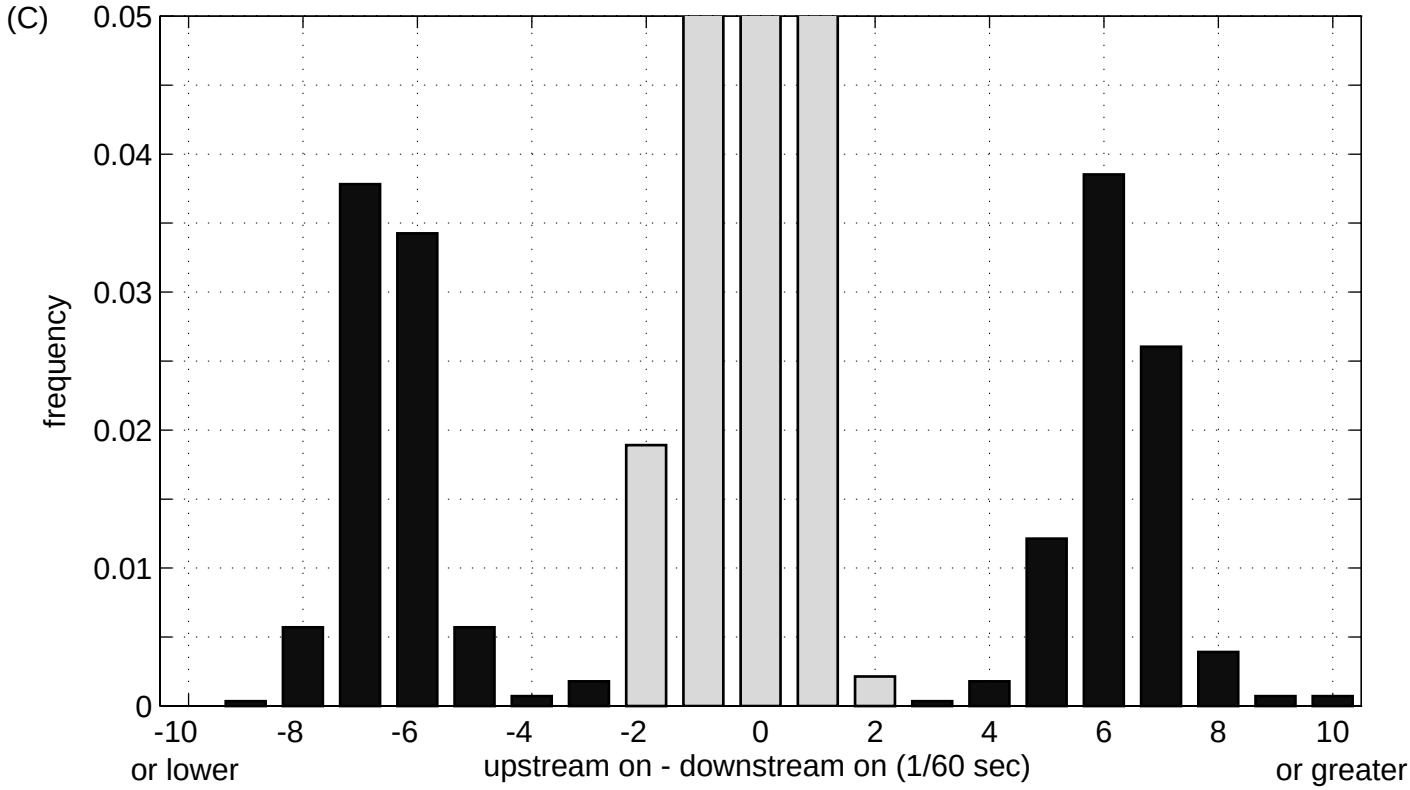

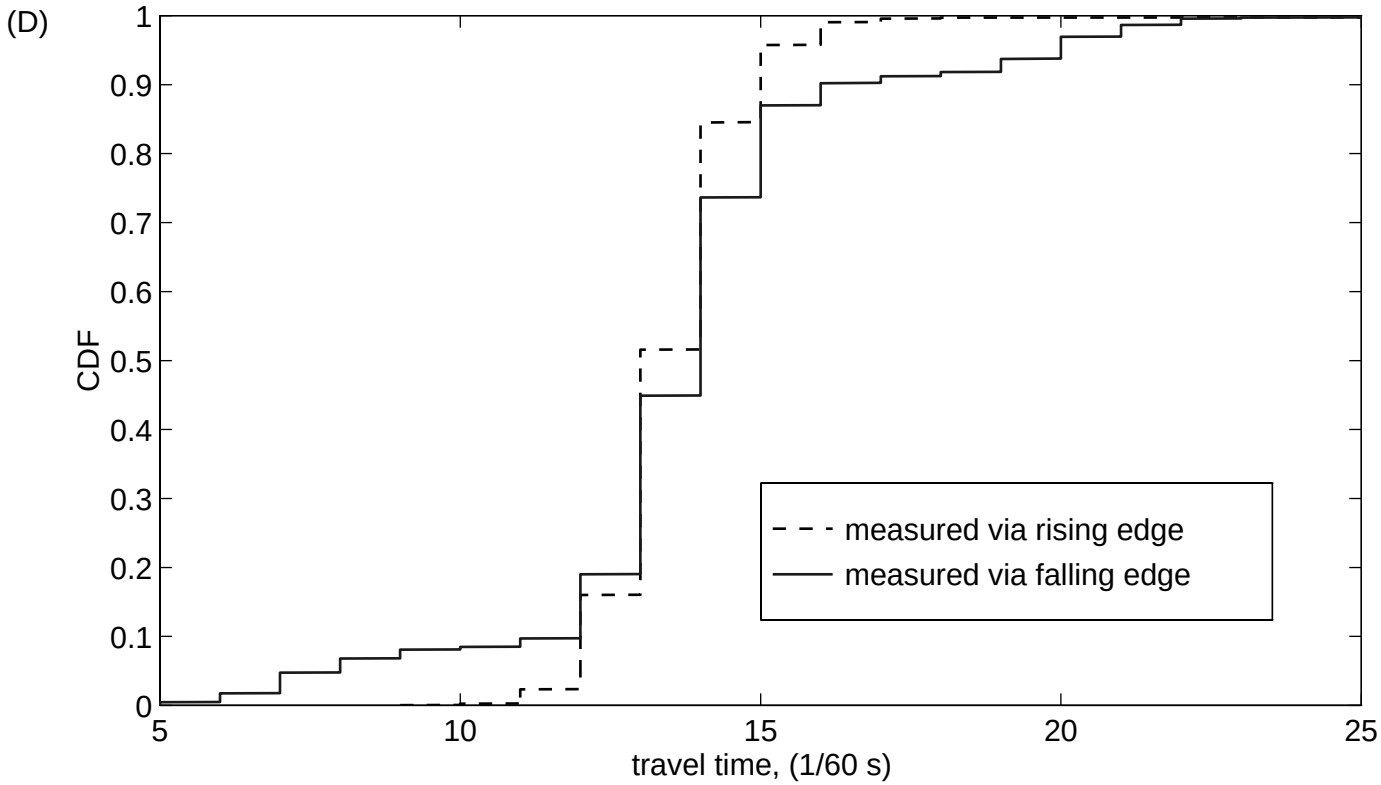

FIGURE 4: (A) Scatter plot of on-time for 3067 matched pulses using a GP5-E sensor (B) Distribution of OT_u - OT_d for the same pulses.

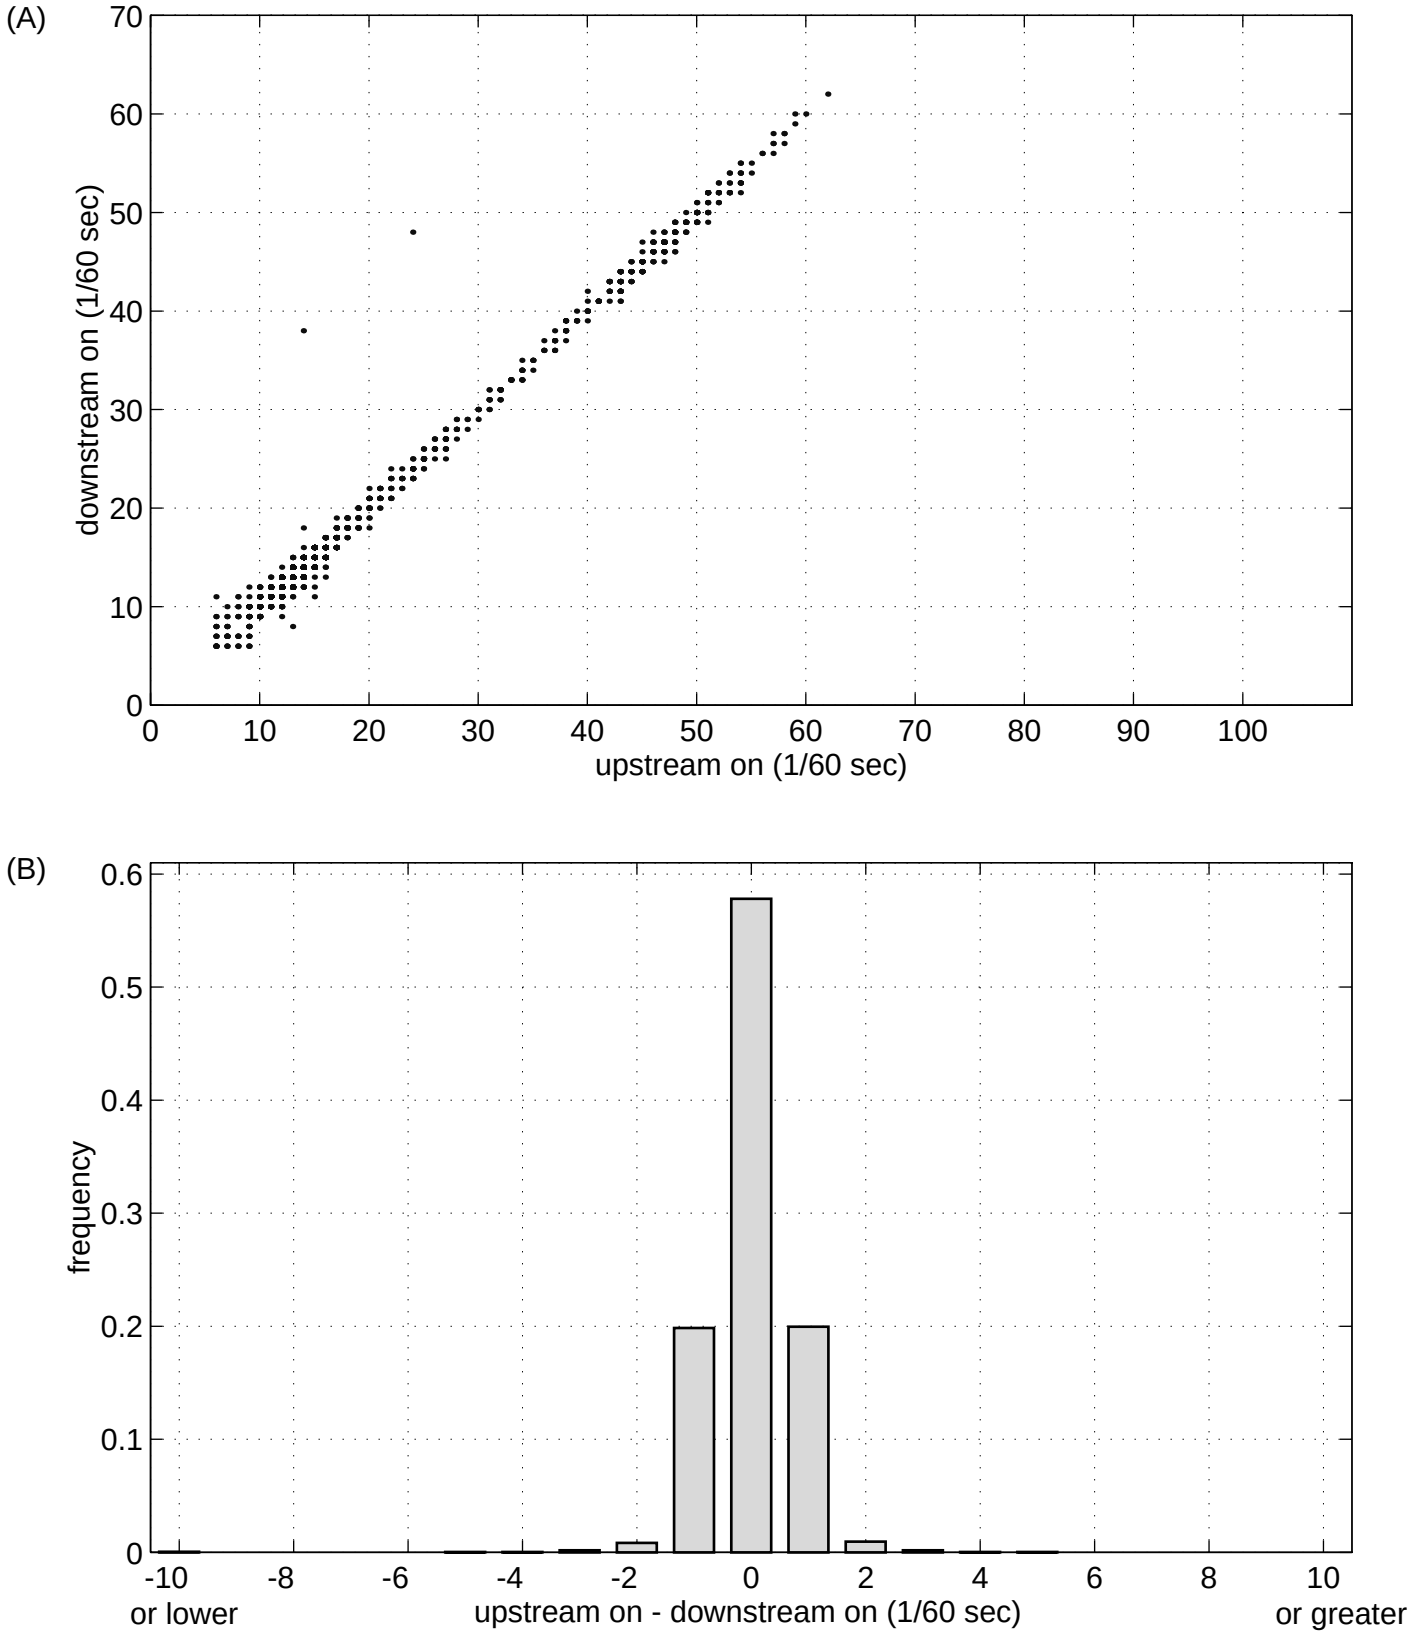

FIGURE 5: (A) Scatter plot of on-time for 2892 matched pulses using a GP6 sensor (B) Distribution of OT_u - OT_d for the same pulses.

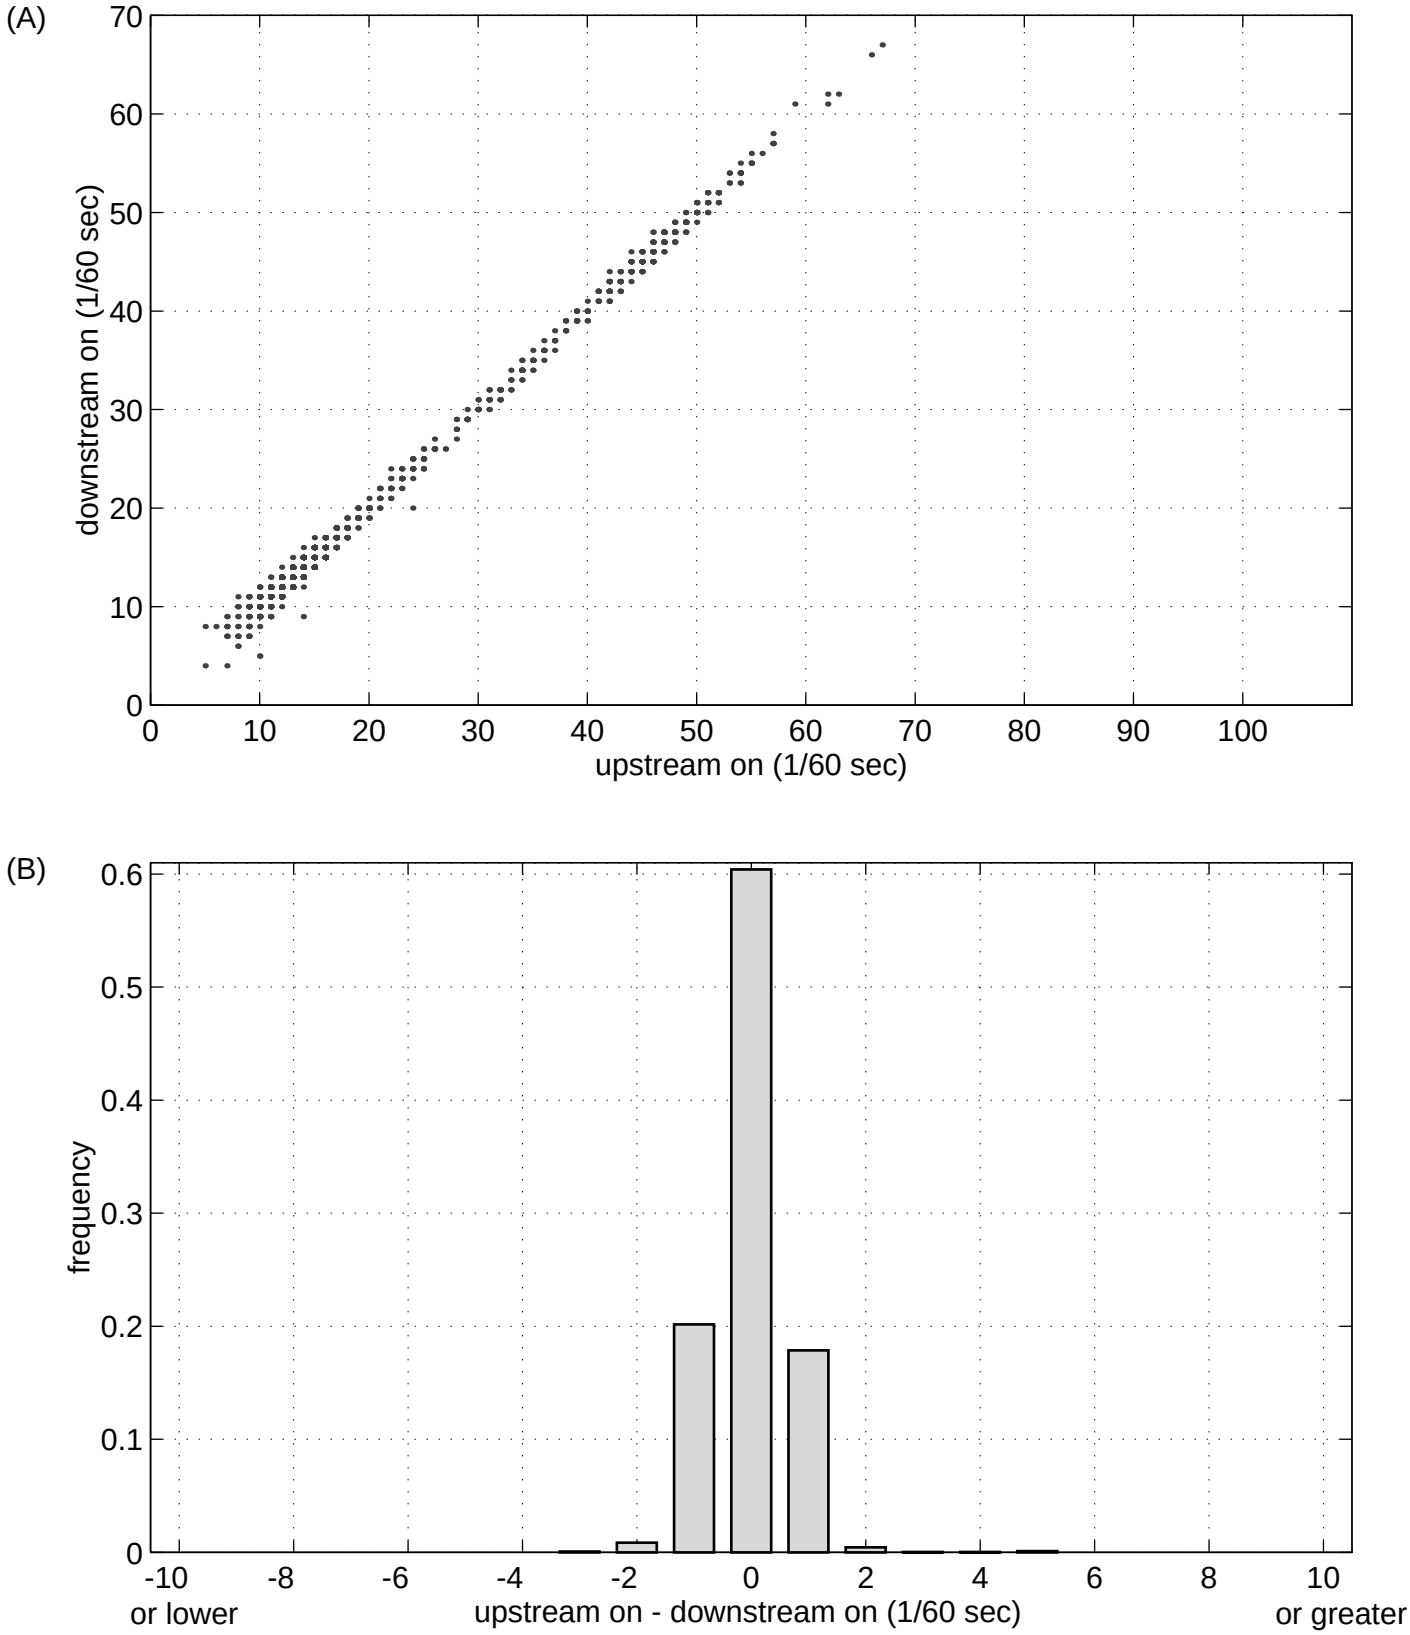

FIGURE 6: (A) Scatter plot of on-time for 3007 matched pulses using an EDI sensor (B) Distribution of OT_u - OT_d for the same pulses.

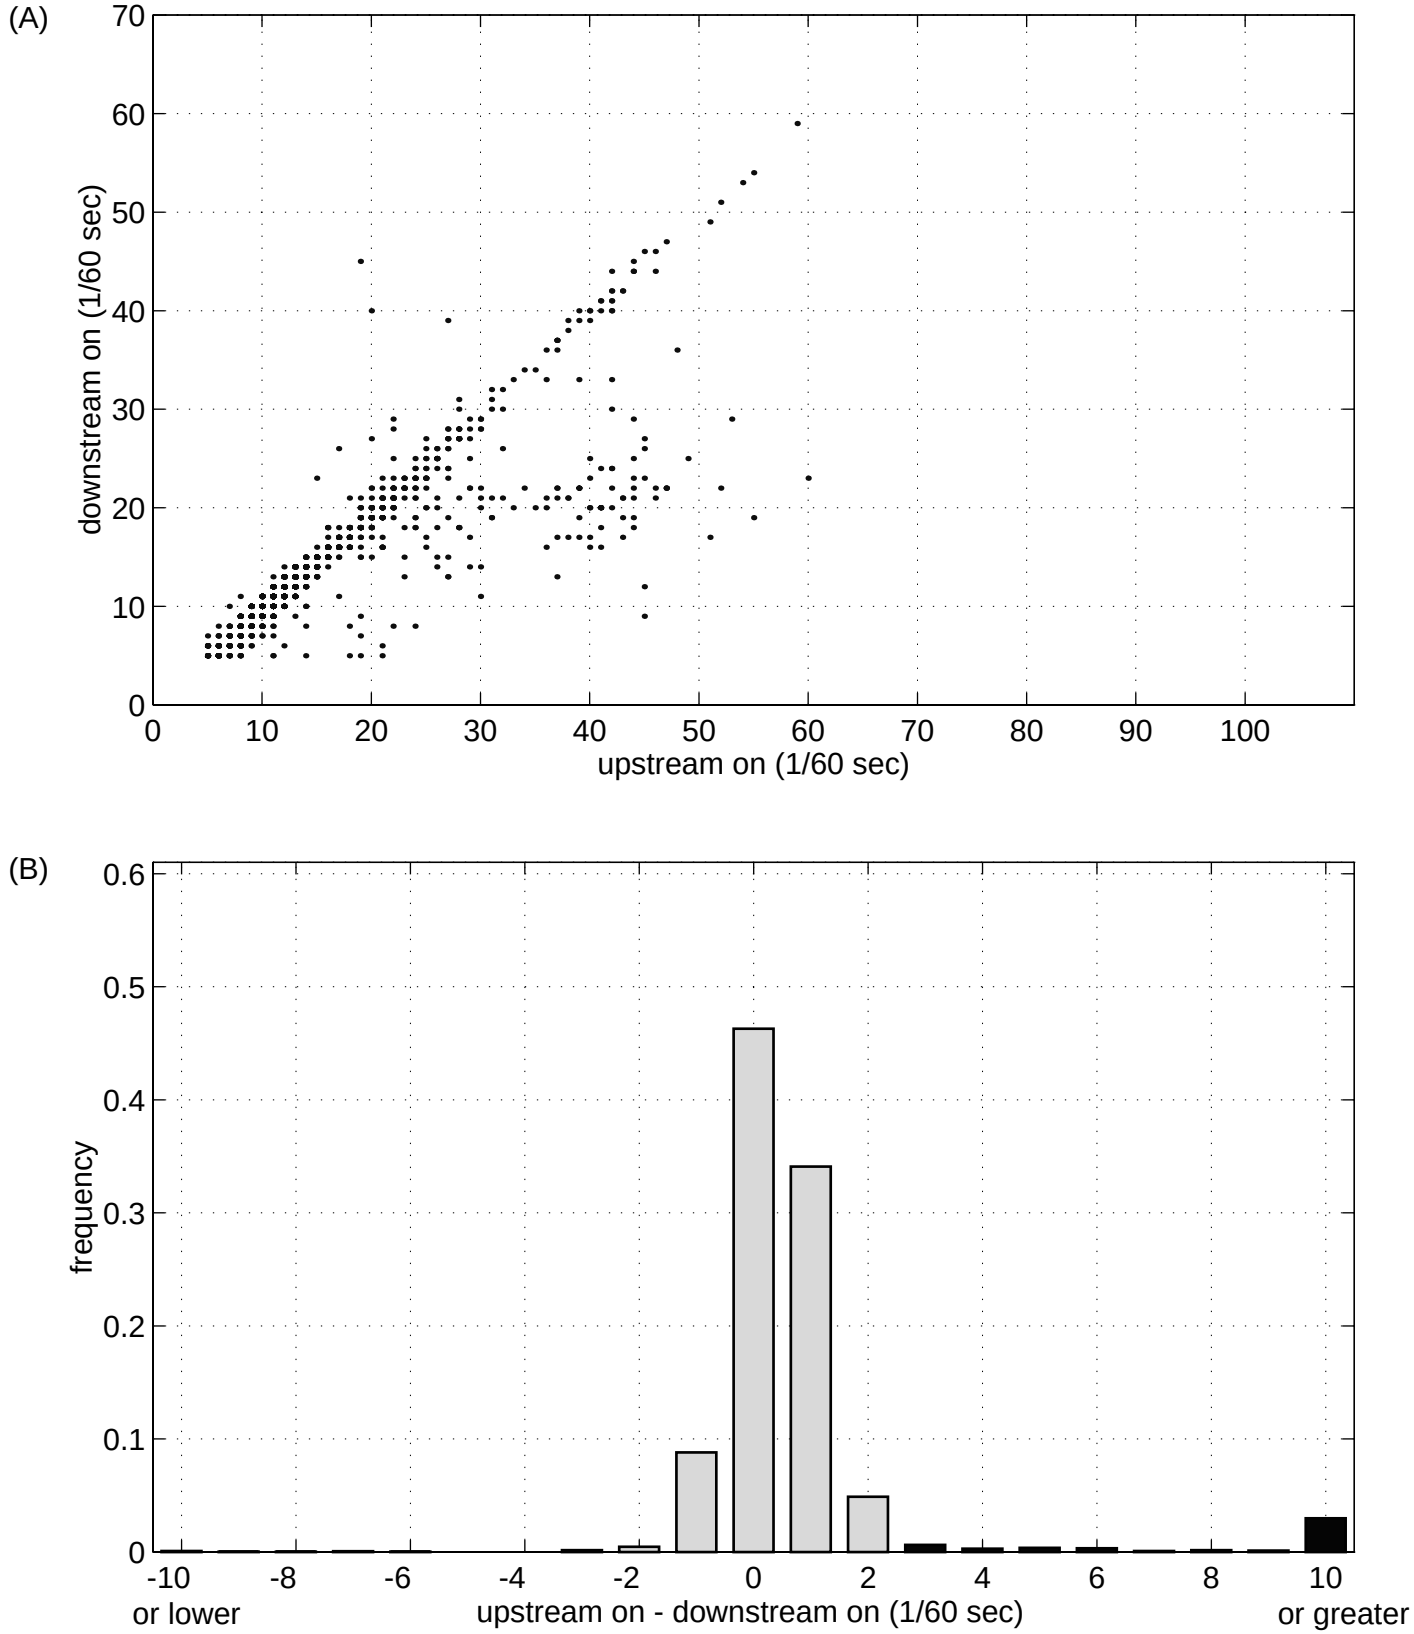

FIGURE 6: (C) detail of part B

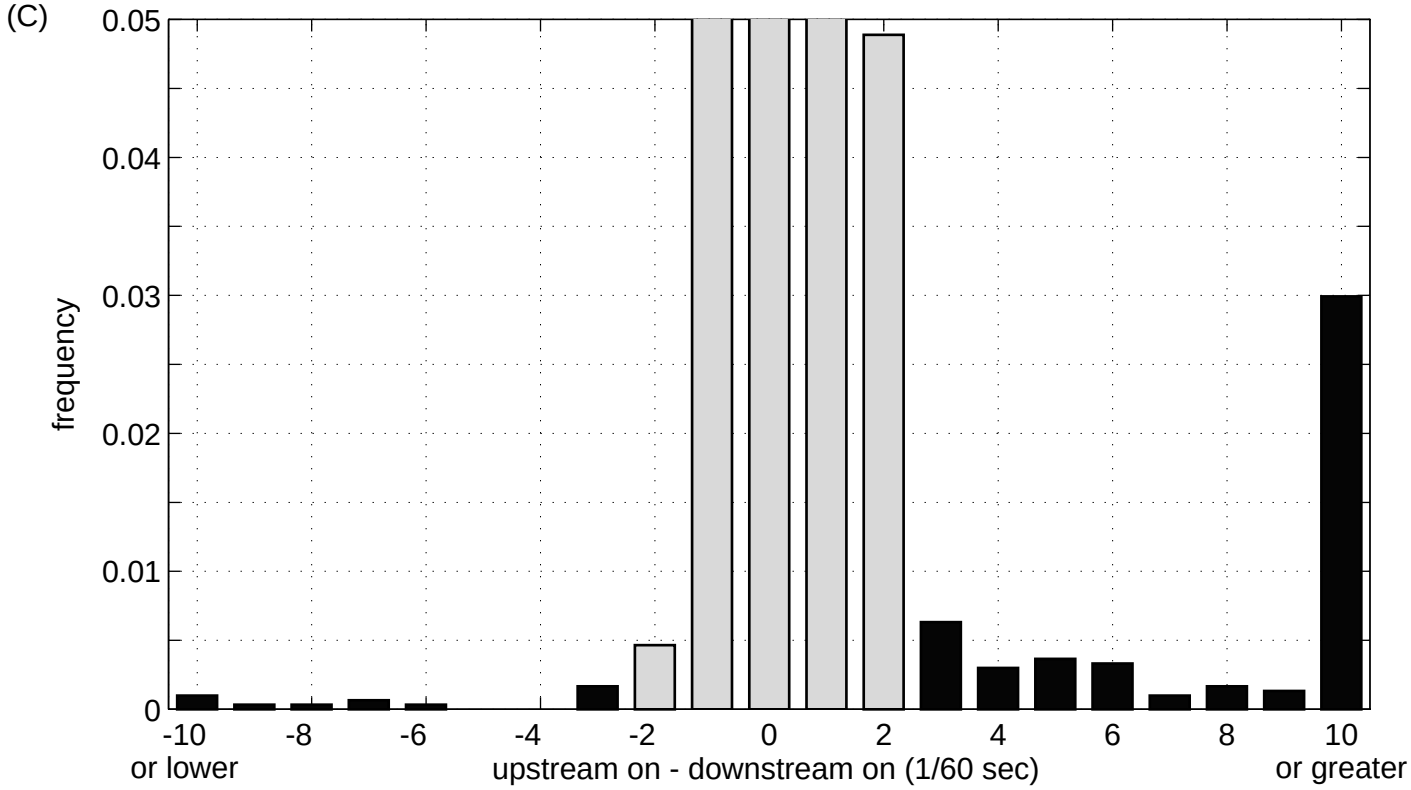

FIGURE 7: (A) Example of pulse breakup from an EDI sensor (B) Individual vehicle velocities for the pulses shown in Figure 6.

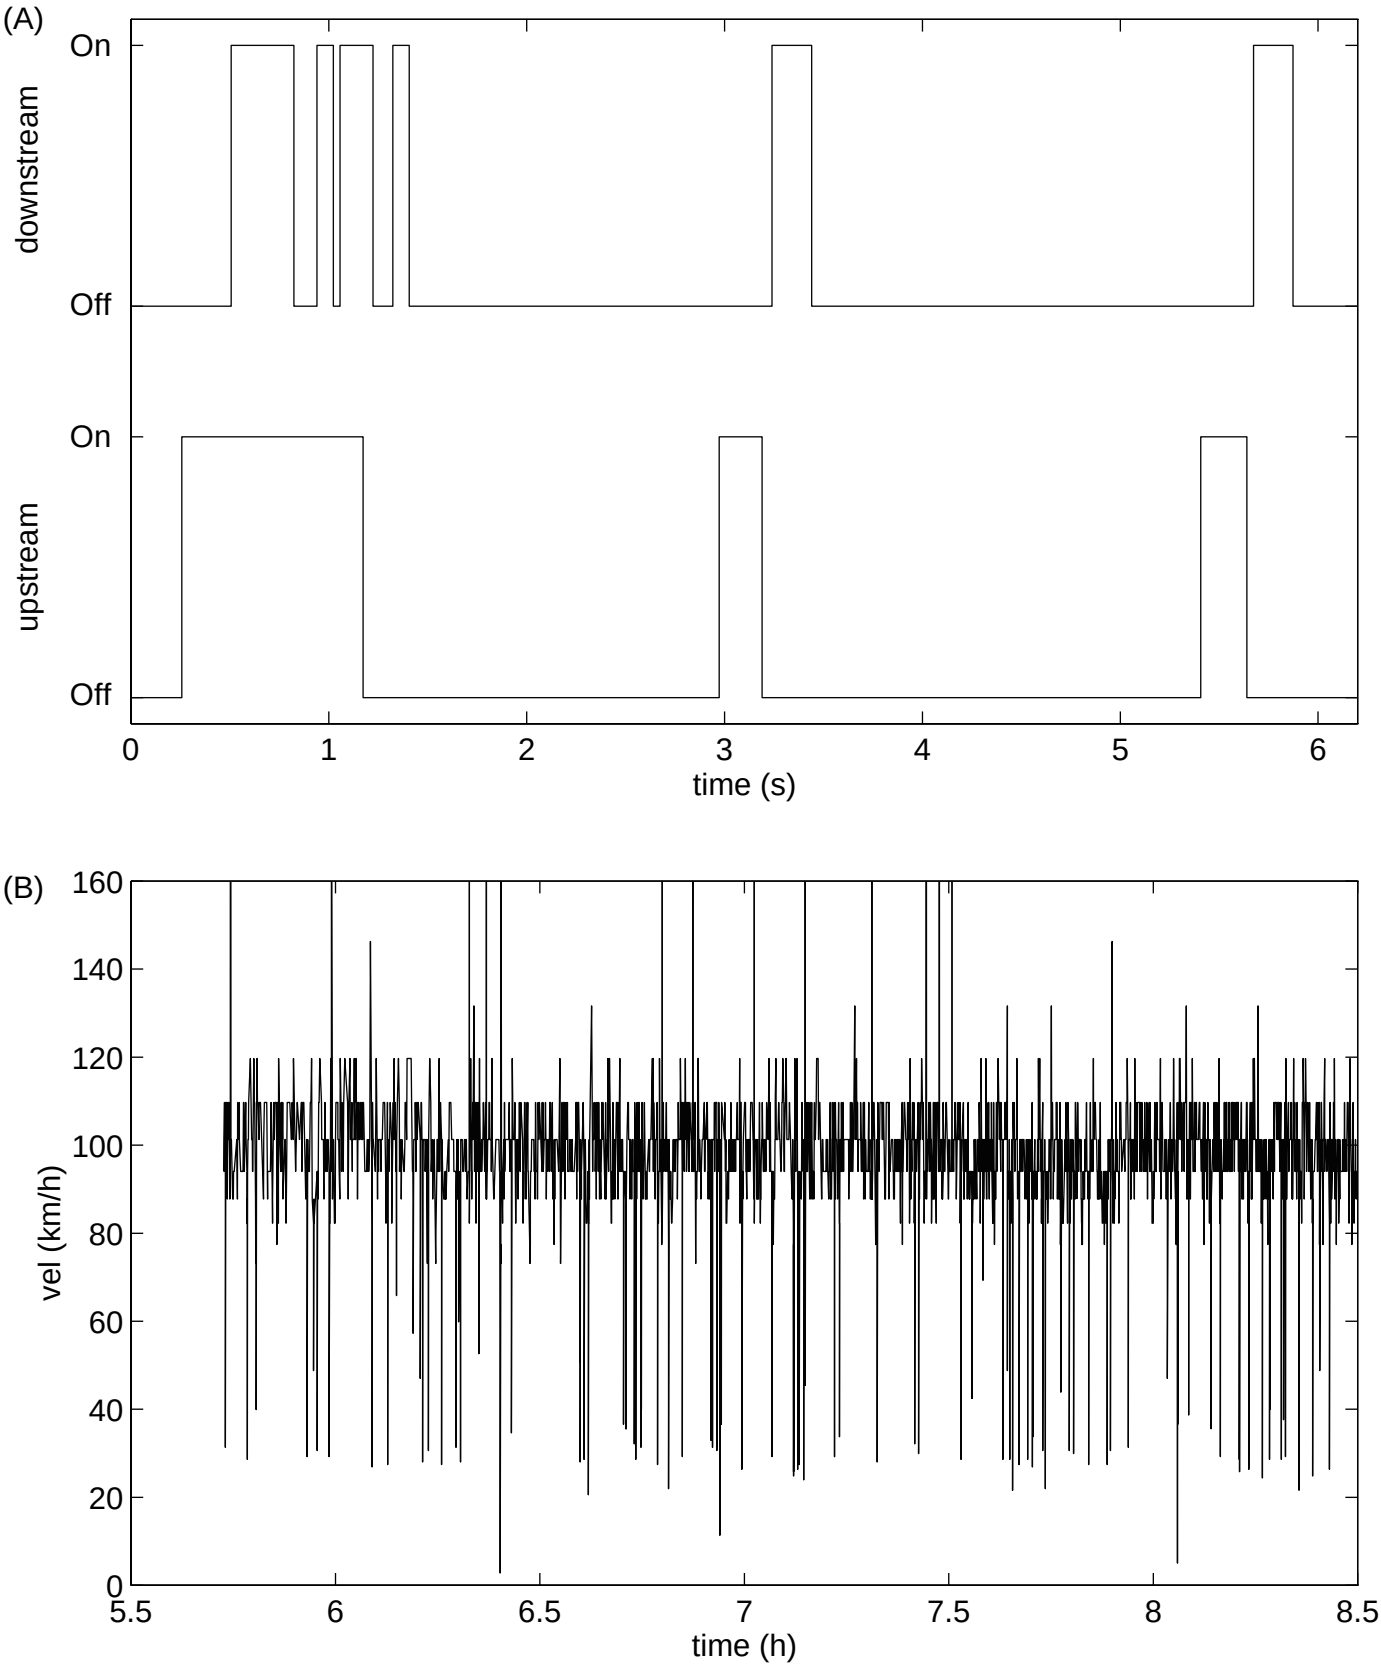

FIGURE 8: (A) Scatter plot of on-time for 4038 matched pulses using an IDC sensor (B) Distribution of OT_u - OT_d for the same pulses.

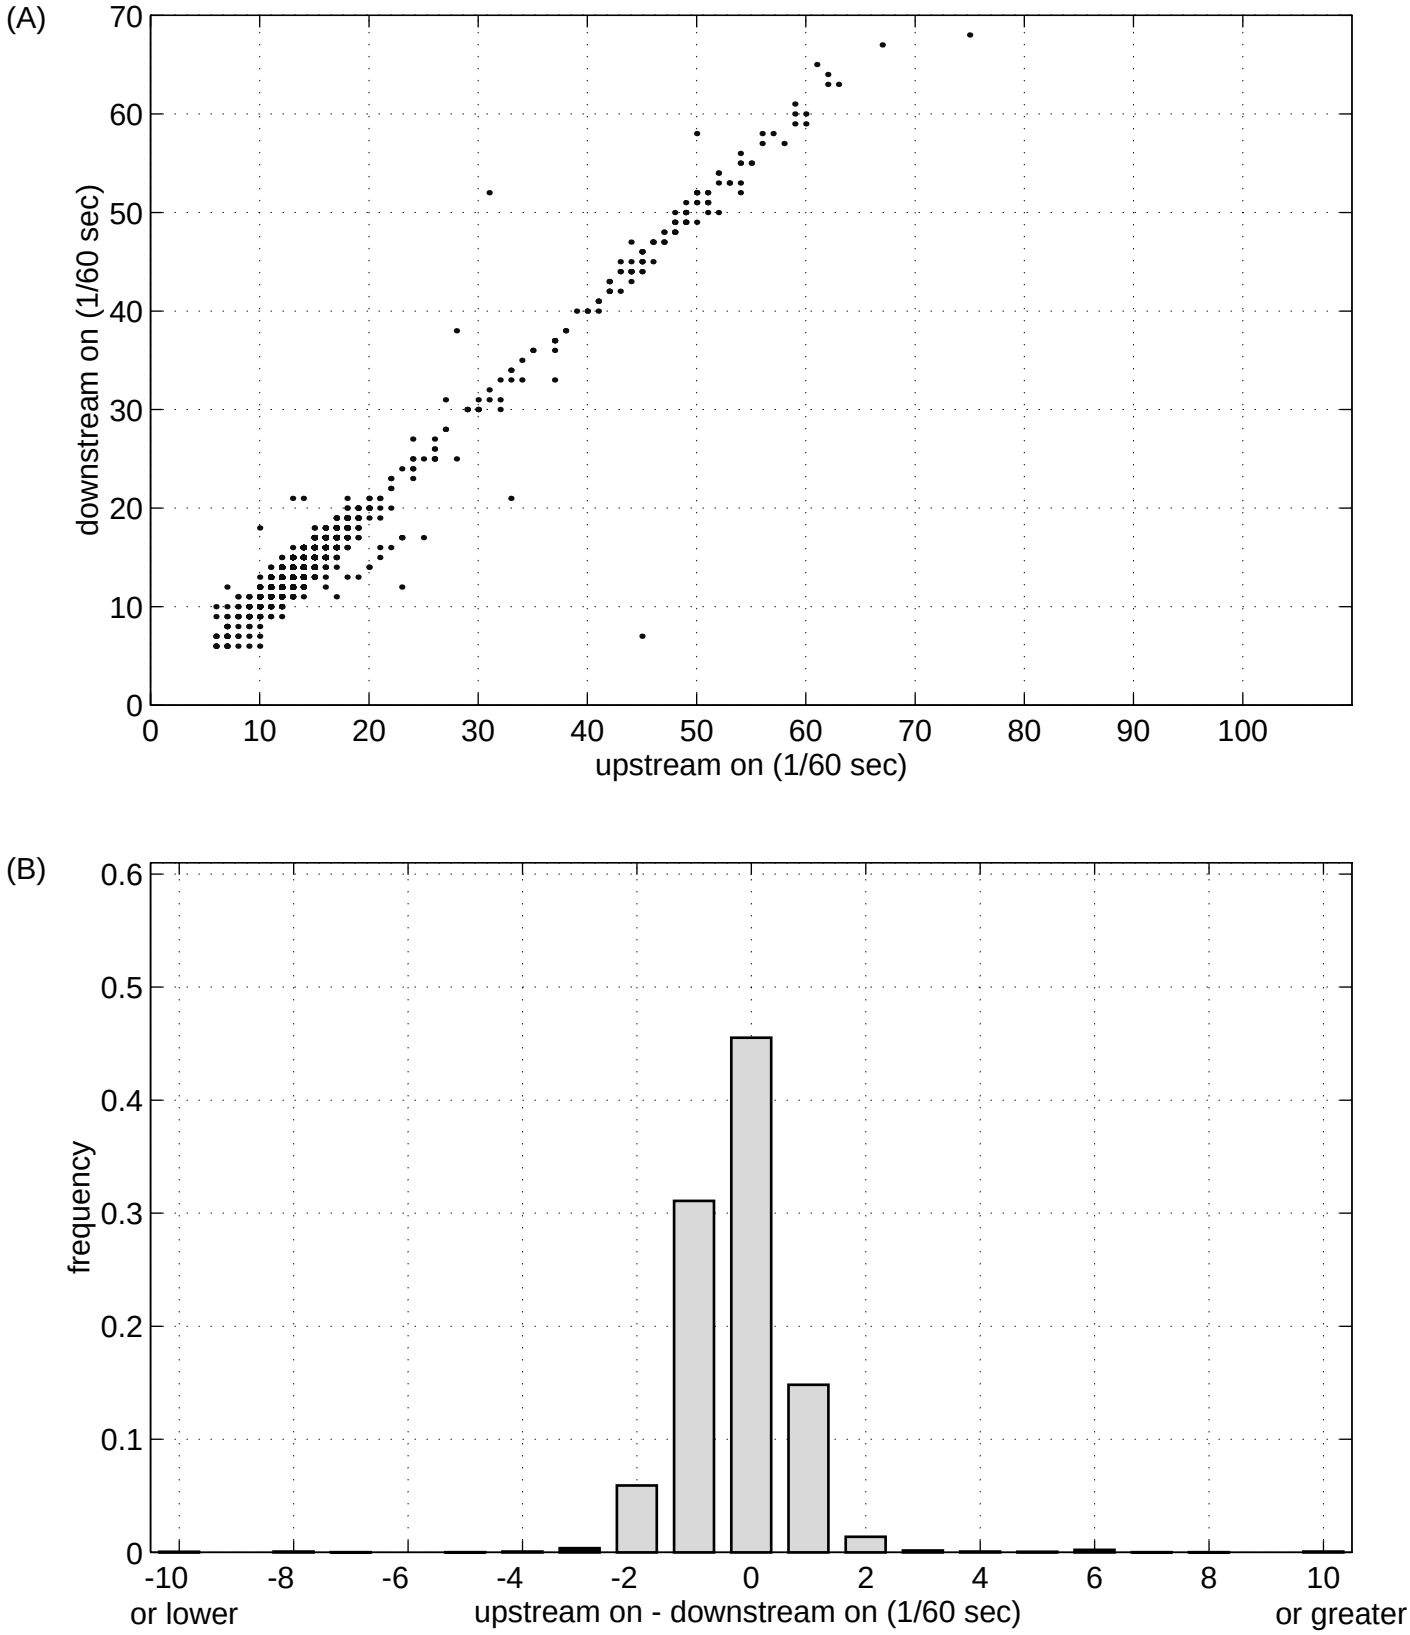

FIGURE 9: (A) Distribution of OT_u - OT_d for two lanes in the presence of crosstalk (B) Distribution of OT_u - OT_d for the same lanes after correcting the problem.

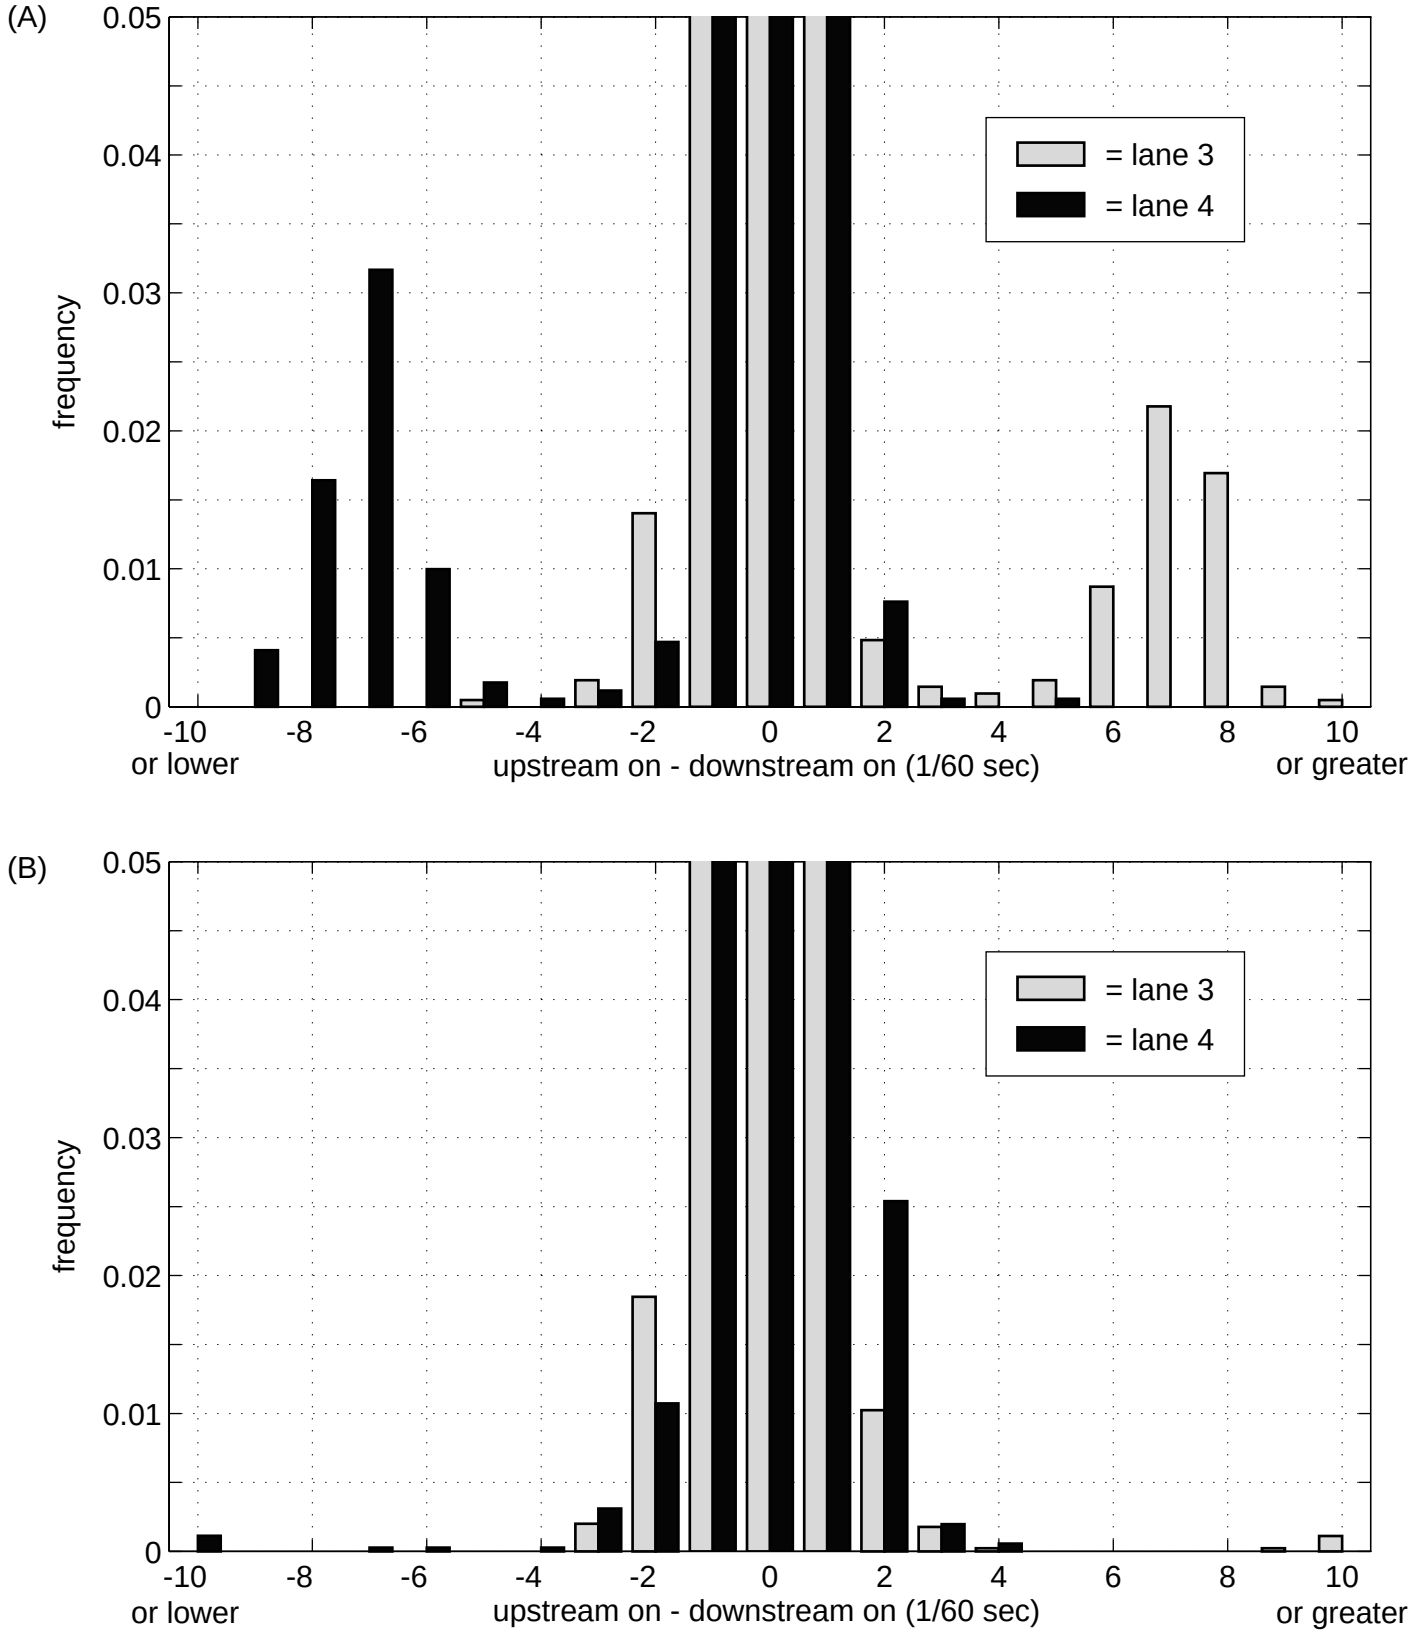

FIGURE 10: Flow chart showing one possible on-line realization of the error detection methodology. The assessment, shown in black, would be conducted in parallel with traditional speed trap measurements, shown in the gray box.

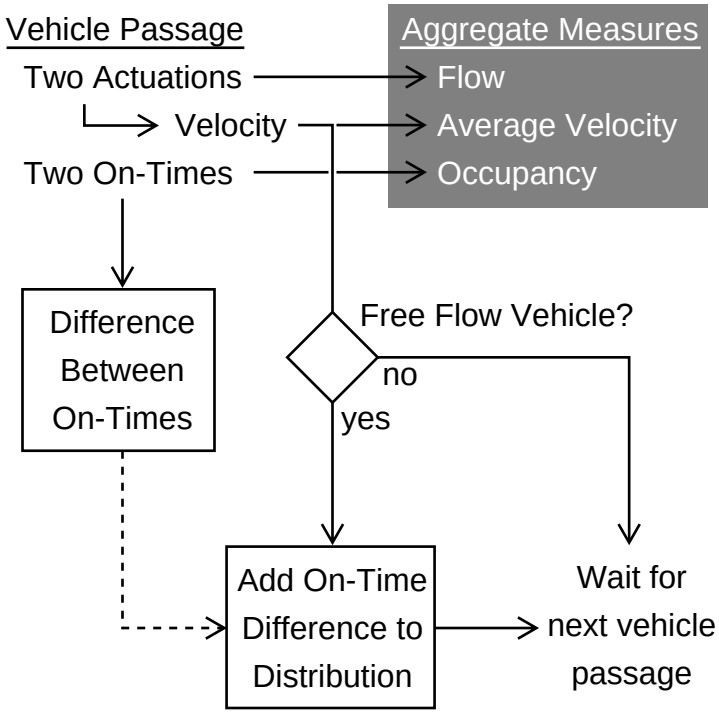

Table titles and figure captions:

TABLE 1: Summary of Sensor Unit Statistics.

TABLE 2: Statistics for the Crosstalk Example.

FIGURE 1: One vehicle passing over a speed trap, illustrating the four critical time measurements. (A) Time space representation showing the loop detectors and both ends of the vehicle (B) Detector output, yielding the upstream and downstream, rising and falling edges.

FIGURE 2: Sample loop actuations for three consecutive vehicles, (A) correctly detected at both loops, (B) vehicle 5 is not detected at the downstream loop, (C) vehicle 8 is not detected at the upstream loop.

FIGURE 3: (A) Scatter plot of on-time for 2803 matched pulses using a GP5-G sensor. (B) Distribution of OT_u - OT_d for the same pulses. (C) detail of part B. (D) CDF of travel time between loops (TT_r and TT_f) for the same pulses.

FIGURE 4: (A) Scatter plot of on-time for 3067 matched pulses using a GP5-E sensor. (B) Distribution of OT_u - OT_d for the same pulses.

FIGURE 5: (A) Scatter plot of on-time for 2892 matched pulses using a GP6 sensor. (B) Distribution of OT_u - OT_d for the same pulses.

FIGURE 6: (A) Scatter plot of on-time for 3007 matched pulses using an EDI sensor. (B) Distribution of OT_u - OT_d for the same pulses. (C) detail of part B.

FIGURE 7: (A) Example of pulse breakup from an EDI sensor. (B) Individual vehicle velocities for the pulses shown in Figure 6.

FIGURE 8: (A) Scatter plot of on-time for 4038 matched pulses using an IDC sensor. (B) Distribution of OT_u - OT_d for the same pulses.

FIGURE 9: (A) Distribution of OT_u - OT_d for two lanes in the presence of crosstalk. (B) Distribution of OT_u - OT_d for the same lanes after correcting the problem.

FIGURE 10: Flow chart showing one possible on-line realization of the error detection methodology. The assessment, shown in black, would be conducted in parallel with traditional speed trap measurements, shown in the gray box.