Title

Identifying the Onset of Congestion Rapidly with Existing Traffic Detectors

Author
Coifman, Benjamin

Publication Date
1999-11-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Identifying the Onset of Congestion Rapidly with Existing Traffic Detectors

Benjamin Coifman

California PATH Working Paper
UCB-ITS-PWP-99-17

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department Transportation, Federal Highway Administration.

Report for 3010

November 1999

ISSN 1055-1417

CALIFORNIA PARTNERS FOR ADVANCED TRANSIT AND HIGHWAYS

Identifying the Onset of Congestion Rapidly with Existing Traffic Detectors

Benjamin Coifman, PhD

Assistant Professor, Civil and Environmental Engineering and Geodetic Science
Assistant Professor, Electrical Engineering

Ohio State University
470 Hitchcock Hall
2070 Neil Ave
Columbus, OH 43210-1275

coifman.1@osu.edu

614 292-4282

Abstract

From an operations standpoint, the most important task of a traffic surveillance system is determining reliably whether the facility is free flowing or congested. The second most important task is responding rapidly when the facility becomes congested. Other tasks, such as quantifying the magnitude of congestion, are desirable, but tertiary.

To address the first two tasks, this paper presents a new approach for traffic surveillance using existing detectors. Rather than expending a considerable effort to detect congested conditions, the research employs a relatively simple strategy to look for free flow traffic. The work should prove beneficial for traffic management and traveler information applications and it promises to be deployable in the short term.

Keywords: traffic surveillance, loop detectors, travel time measurement, vehicle reidentification, Congested traffic

Introduction

Traditional traffic surveillance strategies use loop detectors to calculate aggregate measures such as flow and occupancy at discrete locations on a freeway. Typically, these point measurements are assumed to be representative of extended links spanning detectors. This assumption is usually not valid during congested conditions, when disturbances propagating through the traffic stream can result in significant changes both temporally and spatially. The limitation of point data has lead to the development of many advanced surveillance technologies [1-4], which promise to match observations of the same vehicle at successive detector stations, so called vehicle reidentification. Often times, these advanced technologies are developed without consideration for the general goals of traffic surveillance, and as a result, an operating agency may risk investing in an expensive surveillance system to capture extraneous information.

From an operations standpoint, the most important task of a traffic surveillance system is determining reliably whether the facility is free flowing or congested. The second most important task is responding rapidly when the facility becomes congested. Other tasks, such as quantifying the magnitude of congestion, are desirable, but tertiary. Conventional loop detector surveillance strategies satisfy the first and tertiary tasks, but the response time to delays between detector stations may be excessive [5]. Some of the advanced surveillance technologies promise to satisfy all of these tasks, but they have yet to see widespread deployment. The cost to develop the new technologies is substantial because they require new detector hardware. In contrast, this paper proposes an interim surveillance algorithm using existing dual loop speed traps in a new way. The approach will improve the performance of these detectors on the first and second tasks.

The algorithm identifies relatively distinct vehicles[1] at the downstream detector station and then for each of these vehicles, it looks for a similar vehicle in the same lane at the upstream station within a time window bounded by reasonable free flow travel times. Thus, if traffic is free flowing over the link between detectors, this approach will usually find a match in the time window. If the freeway is congested, vehicles will be delayed and the true match for a vehicle will not be found in the time window. When a match is found, the algorithm can also address the tertiary tasks, such as calculating that vehicle's travel time. In other words, the algorithm will report free flow travel

[1] Specifically, those vehicles that are longer than the average passenger vehicle, henceforth referred to as "long vehicles".

times or "traffic is not free flowing". In the latter case, other surveillance techniques can be used to measure travel time if the metric proves to be desirable, e.g., [6].

To illustrate the potential benefits of the new system, consider incident detection using point detectors. Lin and Daganzo [5] note that two signals emanate from an incident, a backward moving shock wave and a forward moving drop in flow. Detection of an incident can happen only when both of these signals have been received. Although the drop in flow travels at the prevailing traffic speed, their work estimated the shock wave speed to be on the order of 8 mph. Fortunately, the drop in flow reflects the fact that vehicles are being delayed behind the incident. So rather than waiting for the slow moving shock wave, the new algorithm could quickly identify the onset of delay corresponding to the drop in flow.

Is Traffic Free Flowing?

All vehicles that traverse a link between two detector stations must, by definition, pass both stations. For these vehicles, every downstream observation should have a corresponding upstream observation and the time between these two observations is simply that vehicle's link travel time. These travel times are not known a-priori; however, if the vehicle travels at free flow velocities over the entire link, the travel time is bound by the distance between stations and a reasonable range of free flow velocities. For this study, the travel time range is defined as follows:

$$tt = \left[\frac{\text{distance}}{\max(v+10, 55)}, \quad \frac{\text{distance}}{\max(v-10, 45)} \right] \tag{1}$$

where

tt = the range of feasible free flow travel times [hours],
v = local velocity measurement at the downstream detector station [mph],
distance = the known distance between detector stations [miles].

This constraint is illustrated in Figure 1. Chang and Kao [7] suggest that lane change maneuvers are relatively infrequent during free flow conditions. So a free flow vehicle observed at the downstream station will usually have a corresponding observation at the upstream station in the same lane, in the time window bound by tt. Congestion will disrupt this relationship, both because the travel time will increase beyond the free flow travel time range and because there may be an increase in lane change maneuvers, particularly if one or more lanes are blocked.

Parameter Measurement

For a given downstream vehicle, many upstream vehicles may be observed in the corresponding time range. Effective vehicle length, as defined in this section, is used to differentiate between vehicles. As a vehicle passes over a speed trap, e.g., Figure 2A, the controller normally records four transitions, as shown in Figure 2B. After accounting for any unmatched transitions, the following parameters are calculated for each vehicle: speed trap traversal time via the rising edges, TT_r, speed trap traversal time via the falling edges, TT_f, total upstream detector on-time, OT_u, and total downstream detector on-time, OT_d, where:

$$TT_r = t_{RISE\_down} - t_{RISE\_up} \tag{2}$$

$$TT_f = t_{FALL\_down} - t_{FALL\_up}$$

$$OT_u = t_{FALL\_up} - t_{RISE\_up}$$

$$OT_d = t_{FALL\_down} - t_{RISE\_down}$$

Under free flow conditions, the two traversal times should be approximately equal because any acceleration is negligible during the short period that a vehicle is over the detector, similarly the two on-times should be approximately equal. For this paper, each pair of measurements is reduced to a single value using the harmonic mean,

$$TT = \frac{2}{1/TT_r + 1/TT_f} \tag{3}$$

$$OT = \frac{2}{1/OT_u + 1/OT_d}$$

From Figure 2A, vehicle velocity is simply the loop separation, 20 ft for this study, divided by the traversal time. The effective vehicle length, L, is the velocity multiplied by the on-time,

$$L = \frac{20 \cdot OT}{TT} \ [\mathrm{ft}] \tag{4}$$

The controller samples the loops at 60 Hz, so at best, each parameter in Equation 3 is accurate to ±1/30 seconds. Assuming the times from Equation 2 are expressed in seconds, the length range, LR, is defined as:

$$\begin{aligned} LR &= [\text{minimum length estimate, maximum length estimate}] \\ &= \left[20 \cdot \frac{OT - 1/30}{TT + 1/30},\ 20 \cdot \frac{OT + 1/30}{TT - 1/30} \right] [\text{ft}] \end{aligned} \tag{5}$$

and the measurement uncertainty is defined as the difference between the maximum and minimum length estimates. Finally, to ensure the best measurements possible, any hardware problems such as cross talk between detectors are identified using [8] and corrected.

Length Range

As previously noted, the algorithm compares observations, or length measurements, between detector stations. If the length range for a downstream observation overlaps that of an upstream observation, then the two observations may have come from the same vehicle. Otherwise, the result of the pair-wise comparison can be dismissed as an unlikely match because even allowing for the measurement uncertainty, the two ranges do not intersect. Unfortunately, most observations fall in a small range, which is on the order of the measurement uncertainty during free flow conditions.

For example, Figure 3 shows the distribution of observed vehicle lengths over 24 hours at one detector station. Roughly 80 percent of the observations fall between 16 and 23 feet. During free flow conditions, the measurement uncertainty is on the order of two feet for these short vehicles, making difficult the task of differentiating between them.

In contrast, some length observations are as long as 80 feet. Figure 4 shows a typical example of the measurement uncertainty for vehicles over 23 feet. The large range of feasible lengths and the lower frequency of observations for the long vehicles make it feasible to differentiate between these vehicles even at free flow velocities.

Algorithm Implementation

Using an example to illustrate the algorithm implementation, consider the 1.3 mile freeway segment shown in Figure 5. To eliminate the common vehicles, all downstream vehicles shorter than 23 feet are ignored. Whenever a long vehicle passes the downstream speed trap, the algorithm searches a fixed time earlier, bounded by Equation 1, for any upstream vehicles in the same lane whose length range, as defined by Equation 5, intersects the downstream vehicle's length range. If an intersection is found, the corresponding upstream vehicle is considered a possible match. If more than one intersection is found within the time window, then arbitrarily, the most recent of these upstream observations is considered the possible match. Otherwise, the downstream vehicle does not have a match.

Occasionally, a free flow vehicle will not have a match in the same lane either because the vehicle changed lanes, or because of a misdetection at one of the stations. On the other hand, a delayed vehicle should not have a match, but a false positive may fall within the time window. To eliminate most of these transients, the algorithm takes a moving average of the 10 most recent outcomes (including the current outcome), where a possible match is assigned a value of one and a non-match is assigned a value of zero. The current downstream vehicle is considered free flowing if this moving average is at least 0.5, otherwise, it is considered delayed due to congestion. Using just over 2.5 hours of data from the two speed traps, the circles in Figure 6 show the travel times for all of the long vehicles that were considered free flowing and that had a possible match in the example. These free flowing matches will be referred to as *fast matches*.

Verification

Generating ground truth data to verify the algorithm is complicated by the simple fact that vehicle reidentification over extended distances is inherently difficult, both for an automated system and for a human. It is prohibitively time consuming for a human observer to generate exact matches for a large number of vehicles.

Fortunately, it is not necessary to match every vehicle manually. If the algorithm is correctly matching vehicles, it will also yield the true travel times for those vehicles. Although travel time over a freeway link can change dramatically in a short period of time, the travel times for two successive vehicles will be very similar. Thus, a human observer must manually match a sufficient number of vehicles to capture changes in link travel time, but this can be accomplished using a small fraction of the passing vehicles. Manual verification is still a labor-intensive process, but now it becomes feasible to generate ground truth for significant samples. This study used video data, recorded concurrently with the speed trap data, for manual verification and the resulting travel times from the ground truth matches are shown with stars in Figure 6. Finally, Table 1 shows that 7.4 percent of the vehicles were long vehicles in this example and 71 percent of these vehicles were matched.

Detecting the Onset of Congestion

The onset of congestion is characterized by a dramatic increase in link travel times. When this occurs, the true travel times will not fall within the range specified by Equation 1. Notice that the algorithm did not find any *fast matches* after 14.7 hours, which according to the ground truth, corresponds to the onset of congestion.

Although the measured travel times in Figure 6 are useful for traffic surveillance, from this plot, it is impossible to differentiate between congestion and an absence of long vehicles. The true diagnostic power of the method comes from the moving average for the *fast matches*, as shown in Figure 7. The free flow periods are characterized by high average values and congested periods by low values. Unfortunately, there is significant noise in these measurements. During free flow conditions, most of this noise is due to the presence of the two ramps and the long distance between stations. Both of these factors increase the probability that a free flow vehicle will change lanes and thus, not have a match in the same lane. During congested conditions, the long distance between detectors increases the time range in *tt*, and thus, increases the probability of finding a false positive.

To filter out most of the false positives, consider the number of unmatched vehicles preceding each *fast match*, as shown in Figure 8A. Each of the matches preceding the onset of congestion at 14.7 hours have few preceding unmatched vehicles while most of the matches after the onset have many preceding unmatched vehicles. The contrast between the two groups can be increased by taking a moving sum over this data, e.g., Figure 8B shows the results after taking a moving sum of two samples. To eliminate the false positives, all *fast matches* that have more than four unmatched vehicles in Figure 8B are discarded. Figure 9 shows the results after recalculating the moving average over all outcomes. Note that the process has eliminated all of the noise during congestion.

Extending Surveillance into Congestion

Looking closer at the ground truth travel times in Figure 6, there is a transition period between free flow and heavy congestion, characterized by increasing travel times. In an attempt to capture the mildly congested vehicles during the transition, two additional travel time ranges are defined:

$$\text{The range of feasible "medium" travel times [hours]} = tt_m = \left[\frac{\text{distance}}{45}, \quad \frac{\text{distance}}{35}\right] \tag{6}$$

$$\text{The range of feasible "slow" travel times [hours]} = tt_s = \left[\frac{\text{distance}}{35}, \quad \frac{\text{distance}}{28}\right] \tag{7}$$

where the denominators in the bracketed expressions bound the possible link velocities [mph] for each set. As conditions worsen on the link, the true travel times will pass from tt to tt_m, then to

tt_s, and finally exceed all three ranges. The ranges were selected so that collectively they span a larger, continuous range of non-overlapping travel times and each individual range spans approximately the same amount of time between its high and low values.

Repeating the analysis presented in **Algorithm Implementation**, modifying the time window according to Equations 6 and 7, yields *medium matches* and *slow matches*, respectively. Naively treating each group of matches independently, no medium matches are found in the example, but four slow matches are identified, as shown with squares in Figure 10A.

At the onset of congestion, the increasing travel times imply that *medium matches* must be preceded temporally by *fast matches* and *slow matches* must be preceded temporally by *medium matches*. To exploit this phenomenon, modifying the moving average for the *medium matches*, each vehicle is assigned a weight of one if it has a possible *fast match* or possible *medium match*, and zero otherwise. By definition, if a vehicle passes the moving average test for *fast matches*, it must also pass the new test for *medium matches*, thus, a *medium match* will be ignored if the vehicle passes the moving average test for *fast matches*. Similarly modifying the moving average for the *slow matches* such that each vehicle is assigned a weight of one if it has a possible match in any of the three ranges, all *slow matches* that are coincident with *fast* or *medium matches* for the same vehicle are ignored. The results are shown in Figure 10B, where the fast matches have not changed, the medium matches are shown with pentagons and the slow matches are shown with squares.[2] The modified algorithm has found more true matches after the onset of congestion, but as the figure shows, the lower acceptance criteria has allowed a number of false positives to slip through.

Finally, using a hybrid of the two approaches, a *fast match* is only allowed to influence the *medium* or *slow matches* if any of the vehicles in the moving average passed the moving average test for the *fast matches*. Likewise, a *medium match* is only allowed to influence the *slow matches* if any of the vehicles in the moving average passed the moving average test for the *medium matches*. The resulting travel times are shown in Figure 11. The false positives during congestion have been eliminated and compared to Figure 10A, twice as many matches are retained after the onset of congestion.

[2] Note that this analysis does not exploit the filtering discussed in the previous section because some locally common vehicles will cause false positives in more than one time range.

Future research will attempt to optimize the algorithm parameters, e.g., changing the minimum vehicle length, the travel time range(s), and the number of vehicles in the moving average. The use of overlapping travel time ranges could also improve performance.

Conclusions

This paper has developed a new traffic surveillance strategy using existing detectors. Rather than reporting local conditions at the detectors, the strategy identifies when the link between two detector stations becomes congested. This process showed good performance over a 1.3 mile segment with two ramps. Unlike most surveillance strategies that attempt to match vehicle measurements between detector stations, this work is compatible with the existing detector infrastructure. Perhaps more importantly, it is simple enough that it could be implemented on the existing Model 170 controllers, which are based on 20 year old computer technology.

To place the work in context, it is intended to augment, rather than supplant, the local measurements made by conventional surveillance strategies. By combining the existing local data with the new link data, it should be possible to identify transients in either data set and improve performance beyond what would be possible with just one of these data sets.

Acknowledgments

This work was performed as part of the California PATH (Partners for Advanced Highways and Transit) Program of the University of California, in cooperation with the State of California Business, Transportation and Housing Agency, Department of Transportation; and the United States Department of Transportation, Federal Highway Administration.

References

[1] Kühne, R., Immes, S., "Freeway Control Systems for Using Section-Related Traffic Variable Detection" *Pacific Rim TransTech Conference Proc., Vol 1*, 1993, ASCE, pp 56-62

[2] Huang, T., and Russell, S., "Object Identification in a Bayesian Context", *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)* , 1997, Nagoya, Japan. Morgan Kaufmann.

[3] Cui, Y., Huang, Q., "Character Extraction of License Plates from Video" *Proc. 1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE, 1997, pp 502-507.

[4] Balke, K., Ullman, G., McCasland, W., Mountain, C., and Dudek, C., *Benefits of Real-Time Travel Information in Houston, Texas*, Southwest Region University Transportation Center, Texas Transportation Institute, College Station, TX, 1995.

[5] Lin, W., Daganzo, C., "A Simple Detection Scheme for Delay-Inducing Freeway Incidents", *Transportation Research-Part A*, Vol 31A, No 2, 1997, pp 141-155.

[6] Coifman, B., "Vehicle Reidentification and Travel Time Measurement on Congested Freeways", [submitted for publication], 1999. Draft available at: http://.cs.berkeley.edu/~zephyr/BHL

[7] Chang, G., Kao, Y., "An Empirical Investigation of Macroscopic Lane-Changing Characteristics on Uncongested Multilane Freeways", *Transportation Research- Part A*, Vol 25A, No 6, 1991, pp 375-389.

[8] Coifman, B., "Using Dual Loop Speed Traps to Identify Detector Errors", [in publication], Transportation Research Board.

Figure 1, One vehicle traversing an extended link between two speed traps, illustrating the free flow travel time range. (A) The vehicle travels at a free flow velocity and it was observed at the upstream station during the time range; (B) the vehicle traveled slower than the minimum free flow velocity and it passed the upstream station before the start of the time range.

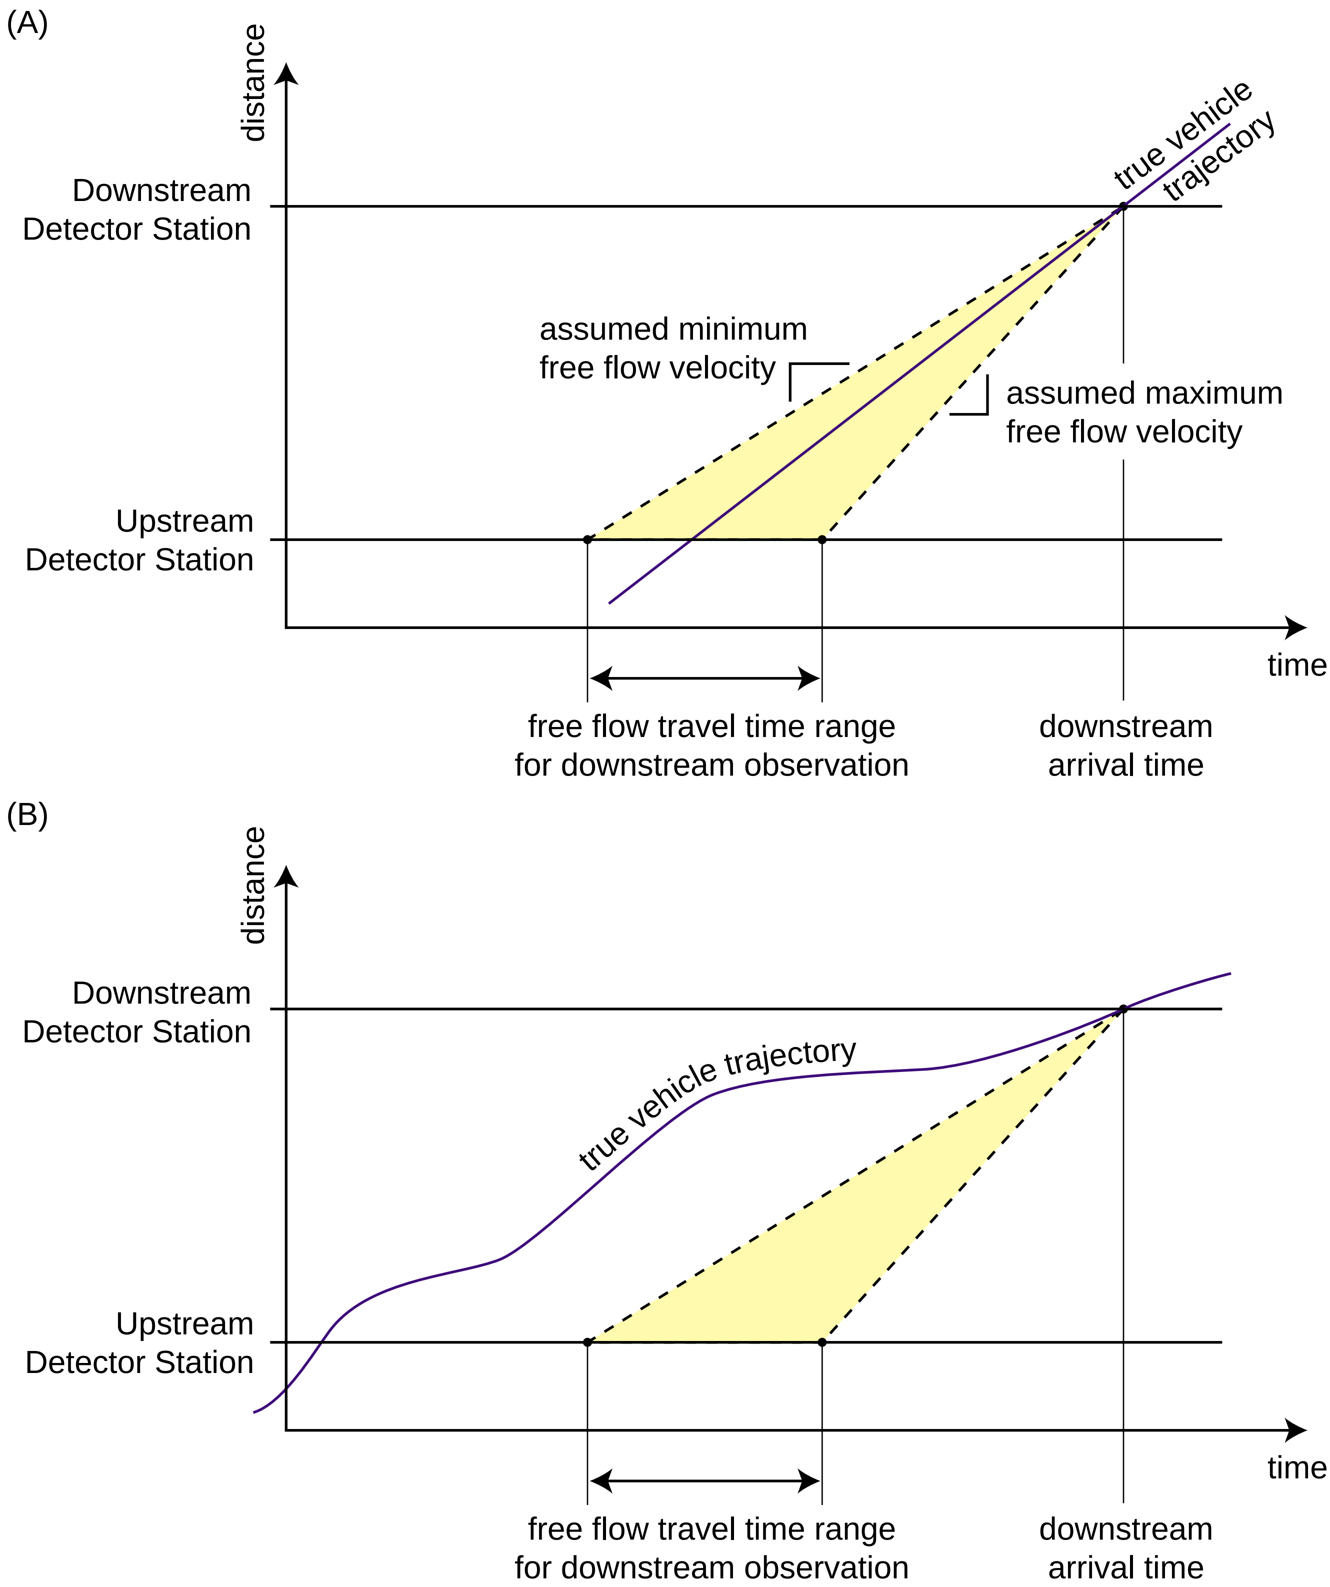

Figure 2, One vehicle passing over a speed trap, (A) detection zones and the vehicle trajectory as shown in the time space plane. The height of the vehicle's trajectory reflects the non-zero vehicle length. (B) The associated detector transitions, including the upstream and downstream, rising and falling edges at the indicated times.

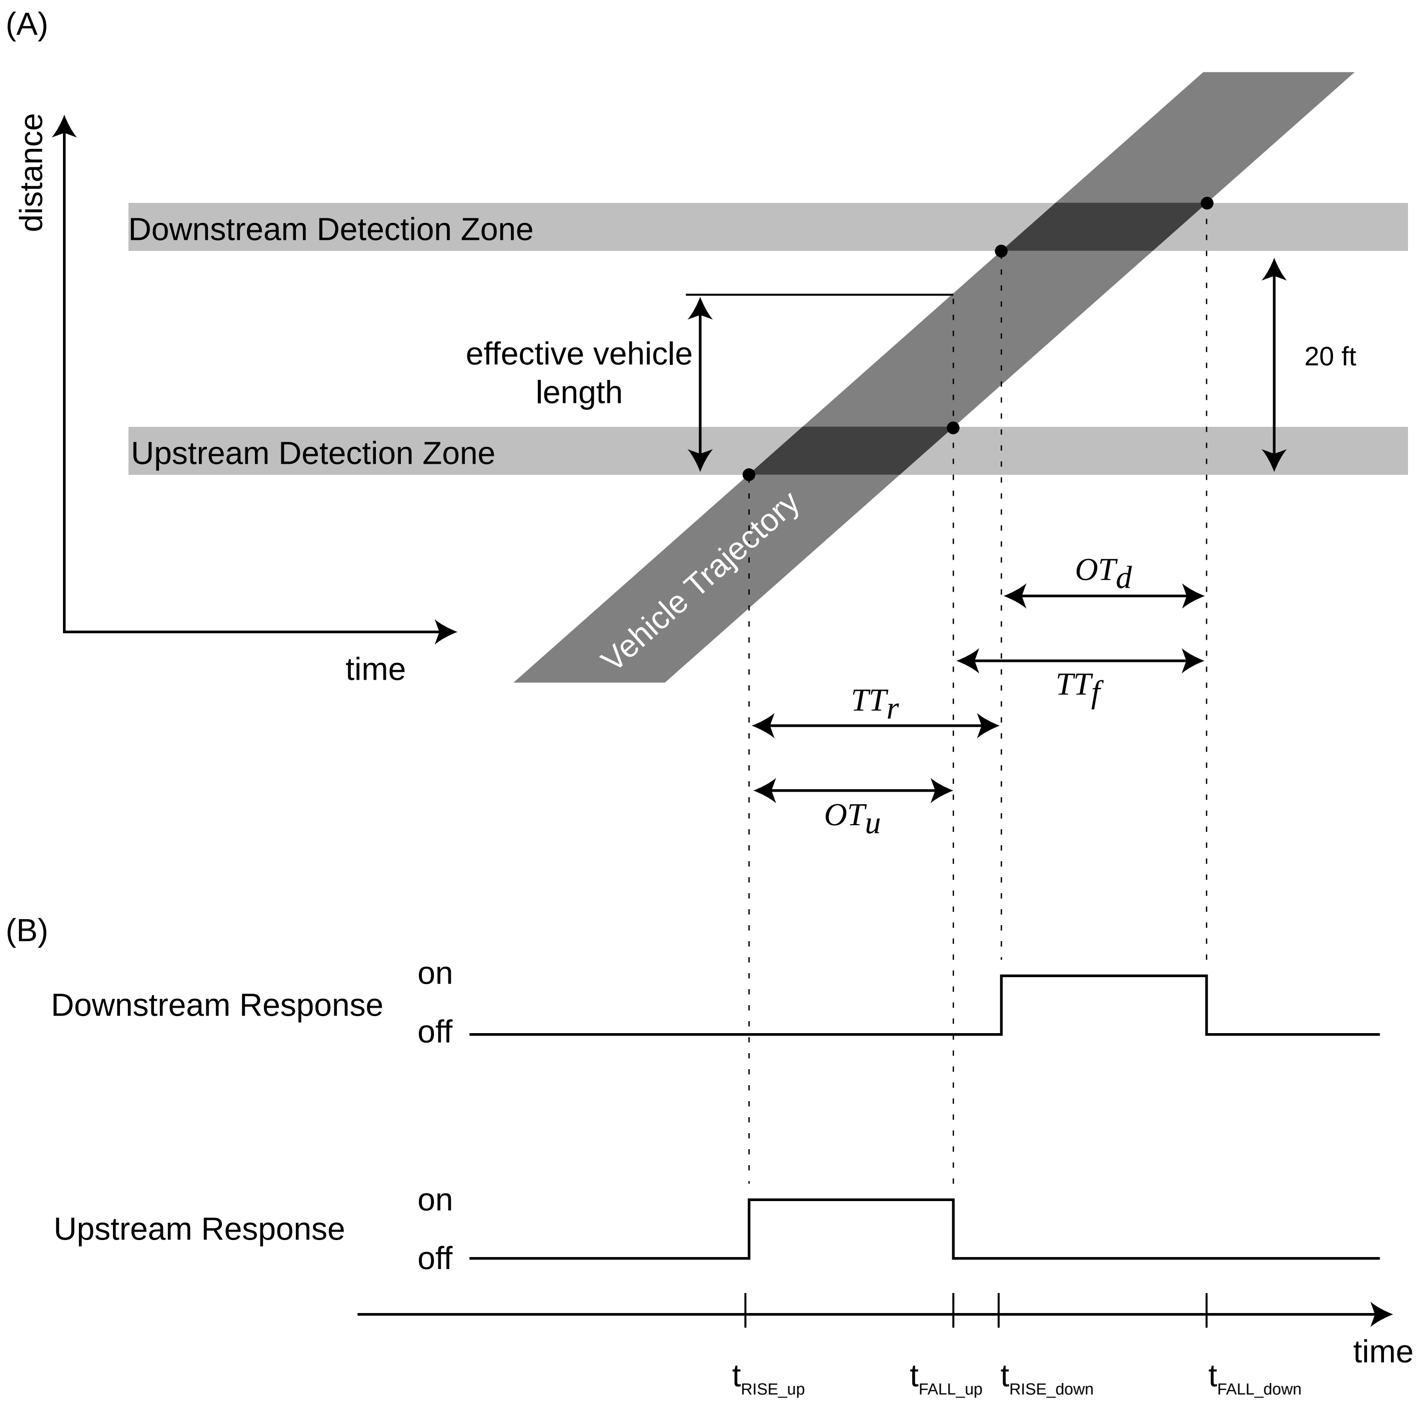

Figure 3, (A) Cumulative distribution of individual vehicle lengths over 24 hours at one detector station. (B) detail of part A.

(A)

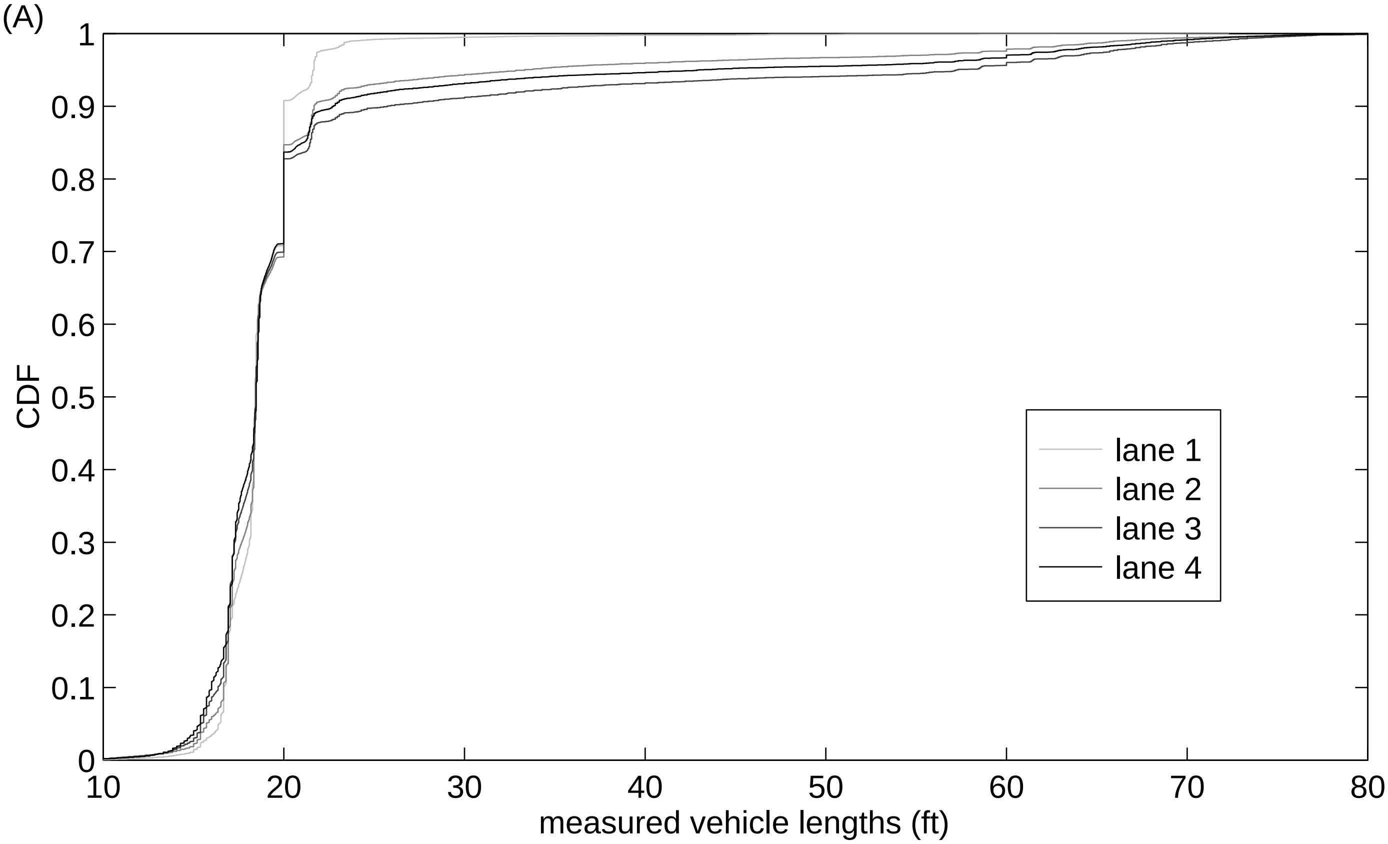

(B)

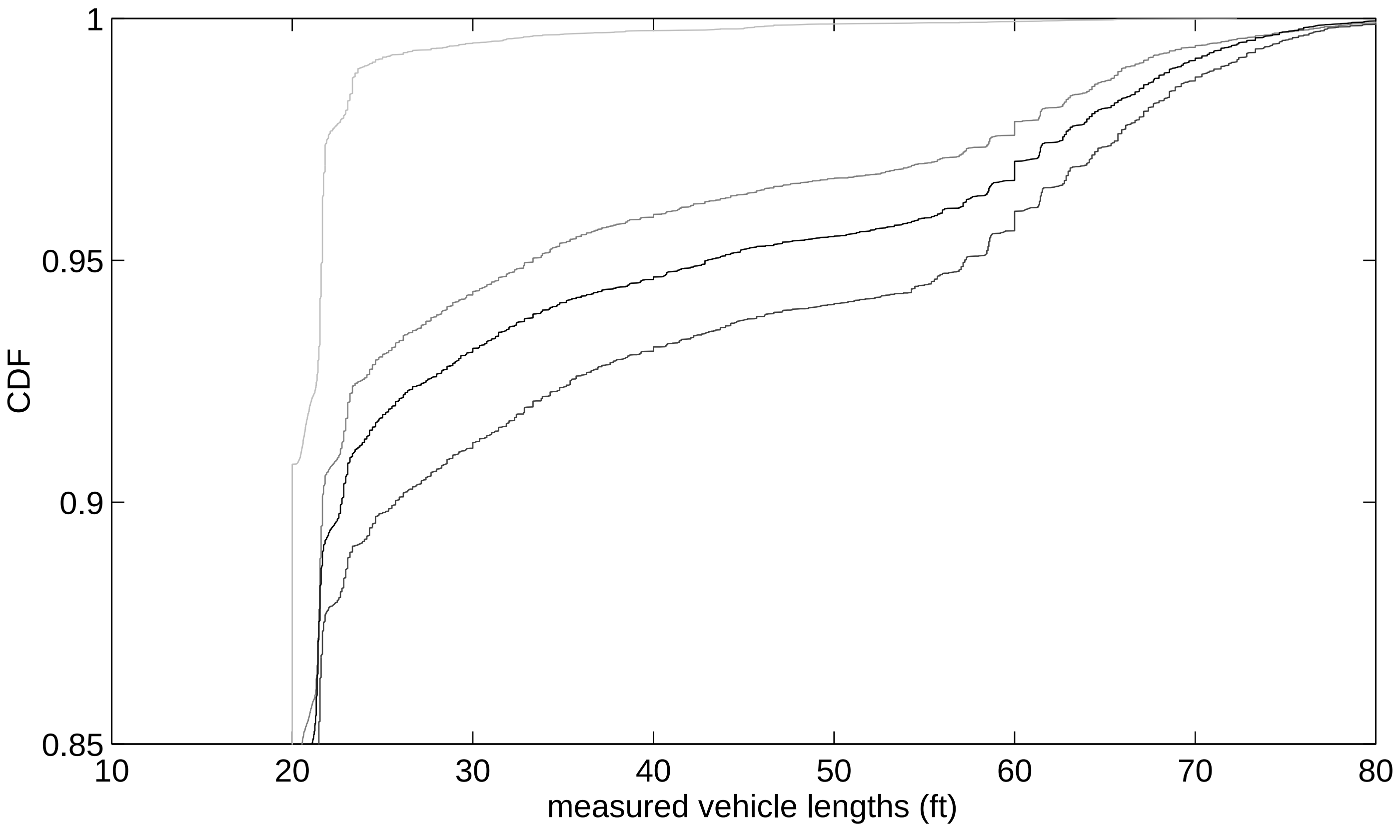

Figure 4, Measurement uncertainty versus measured length for long vehicles at one speed trap over a three hour period. The observations are grouped into three categories based on speed, note how the measurement uncertainty increases with velocity.

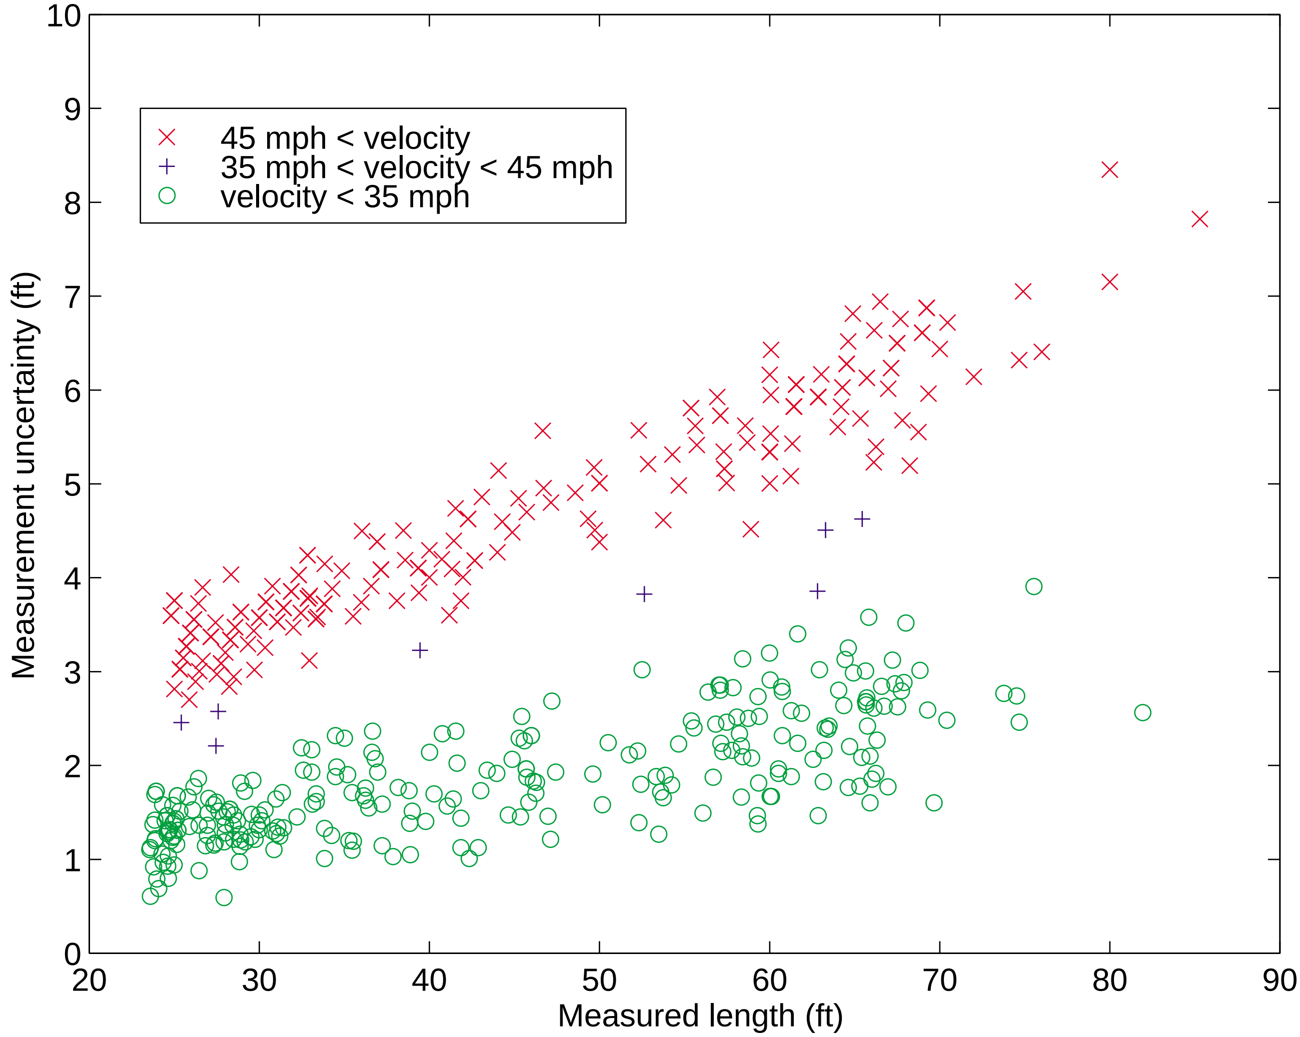

Figure 5, The segment of Interstate-80 in Berkeley, California used to illustrate and verify the Surveillance Algorithm.

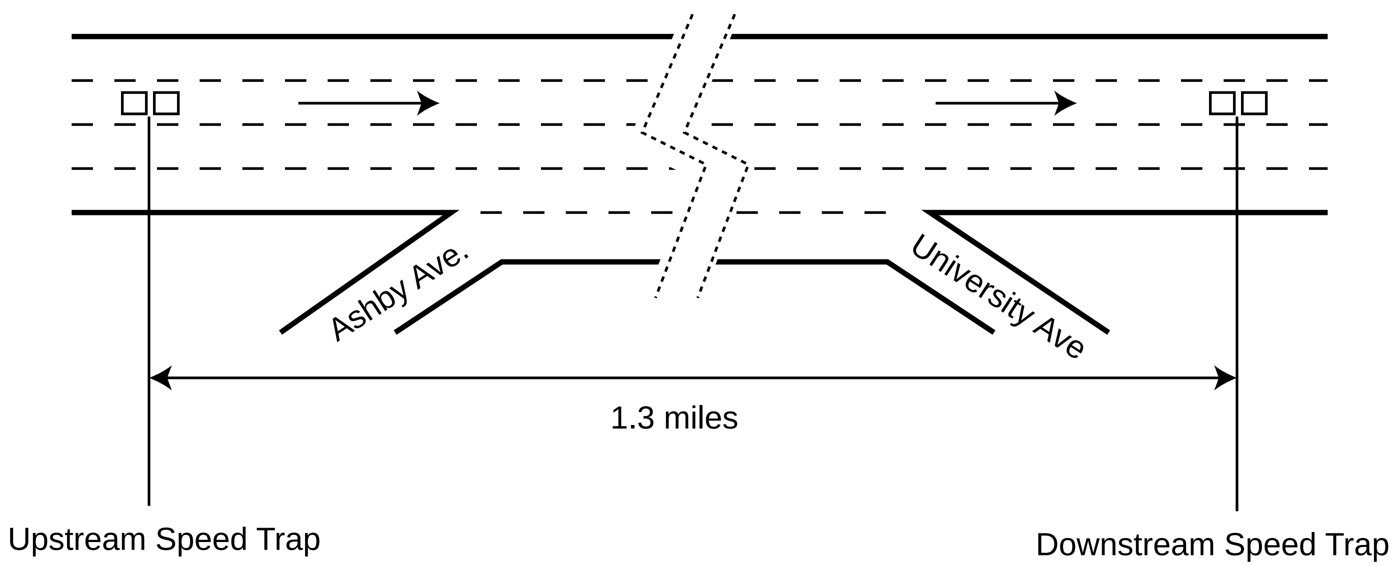

Figure 6, Measured travel times for matched vehicles ("fast matches") and ground truth travel times. As designed, the algorithm did not find any matches after the true travel times increased due to congestion.

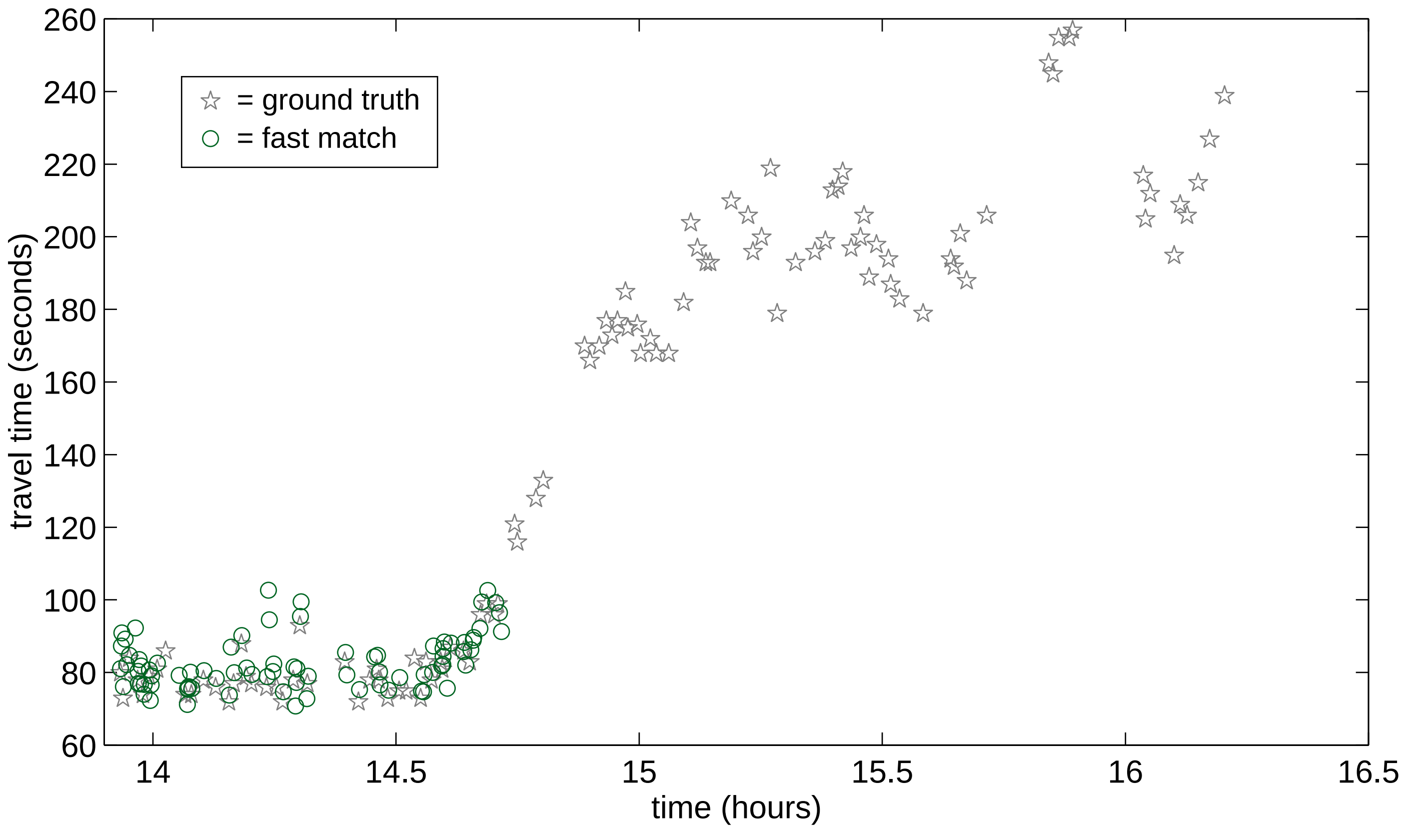

Figure 7, Each pair-wise test is assigned a value of 1 if a match is found and 0 otherwise. This plot shows the moving average of 10 sequential outcomes. There is a noticeable drop at 14.7 hours due to the onset of congestion.

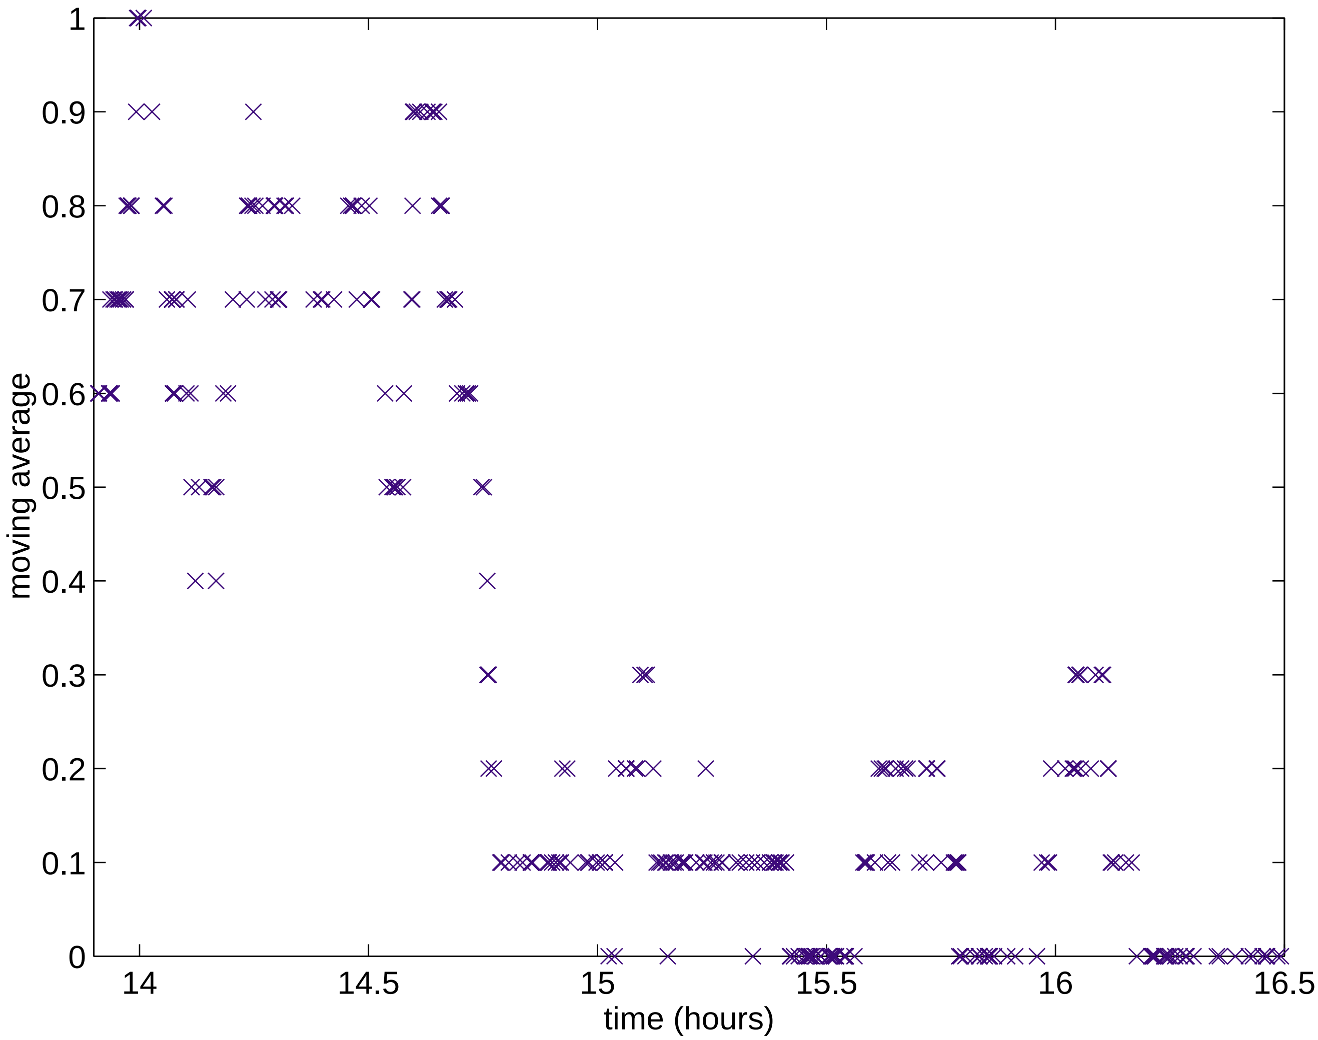

Figure 8, (A) The number of unmatched vehicles preceding each fast match. The free flow matches typically have few unmatched vehicles compared to the congested matches after 14.7. (B) Moving sum of two measurements from part (A).

(A)

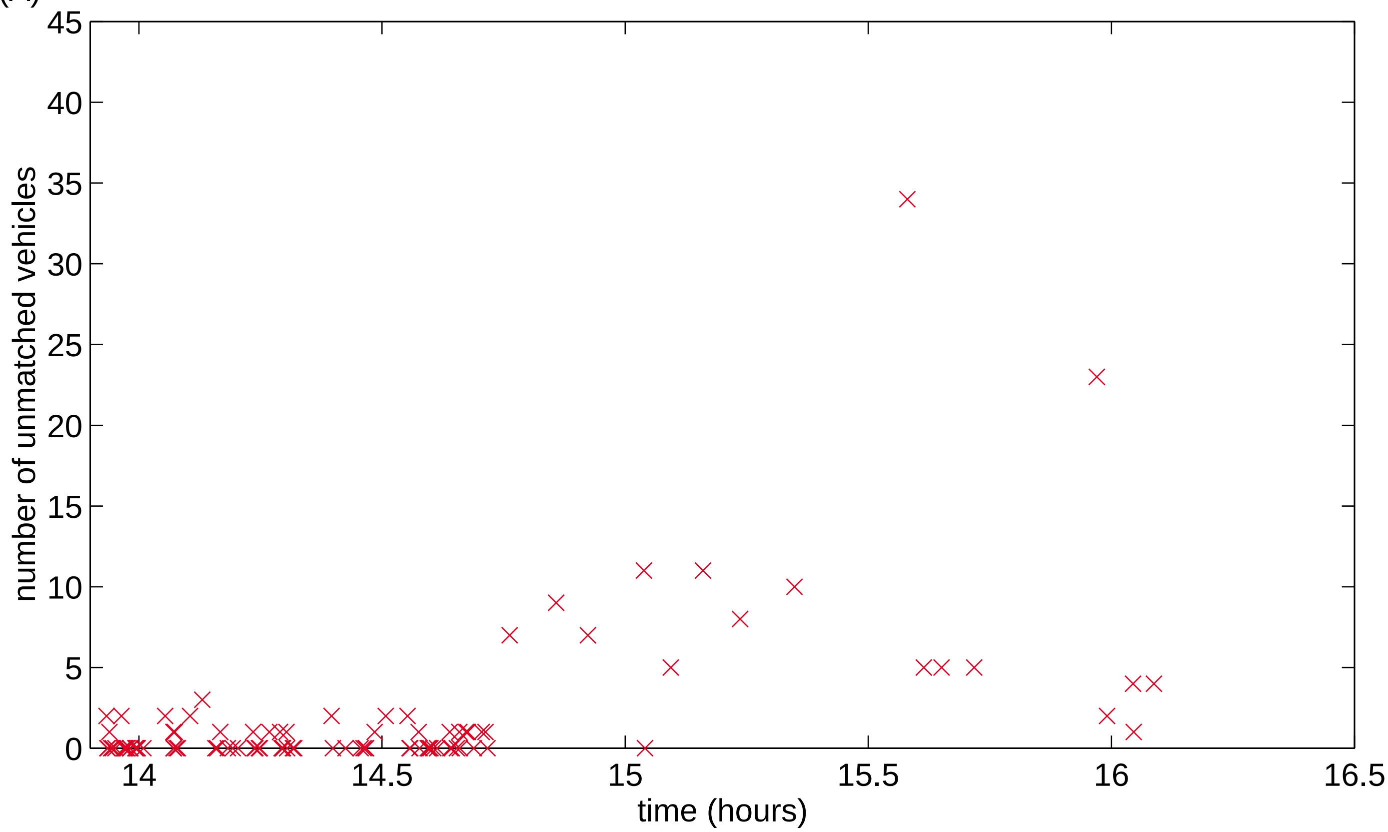

(B)

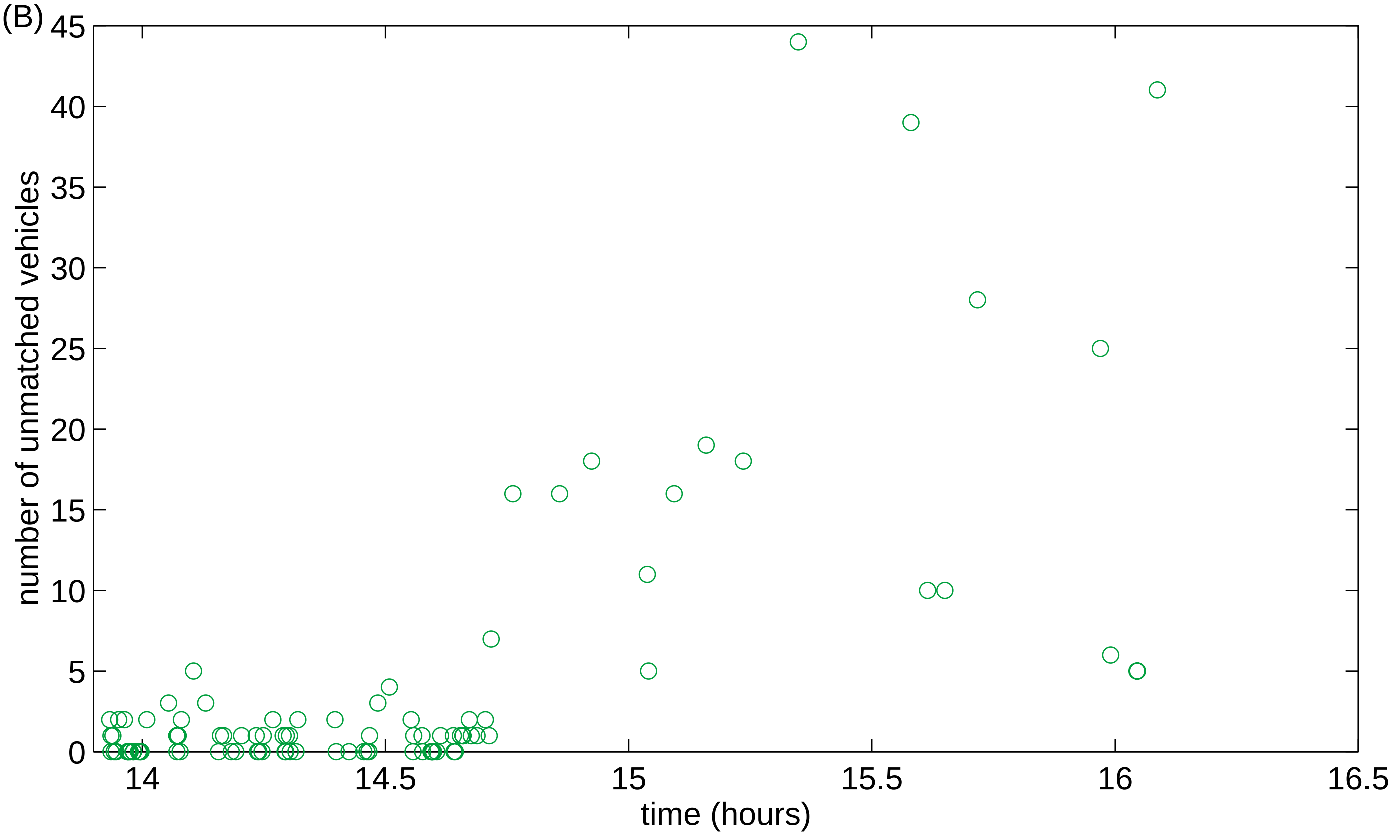

Figure 9, Discarding all matches preceded by a large number of unmatched vehicles before calculating the moving average, this plot yields a better contrast between free flow and congested conditions.

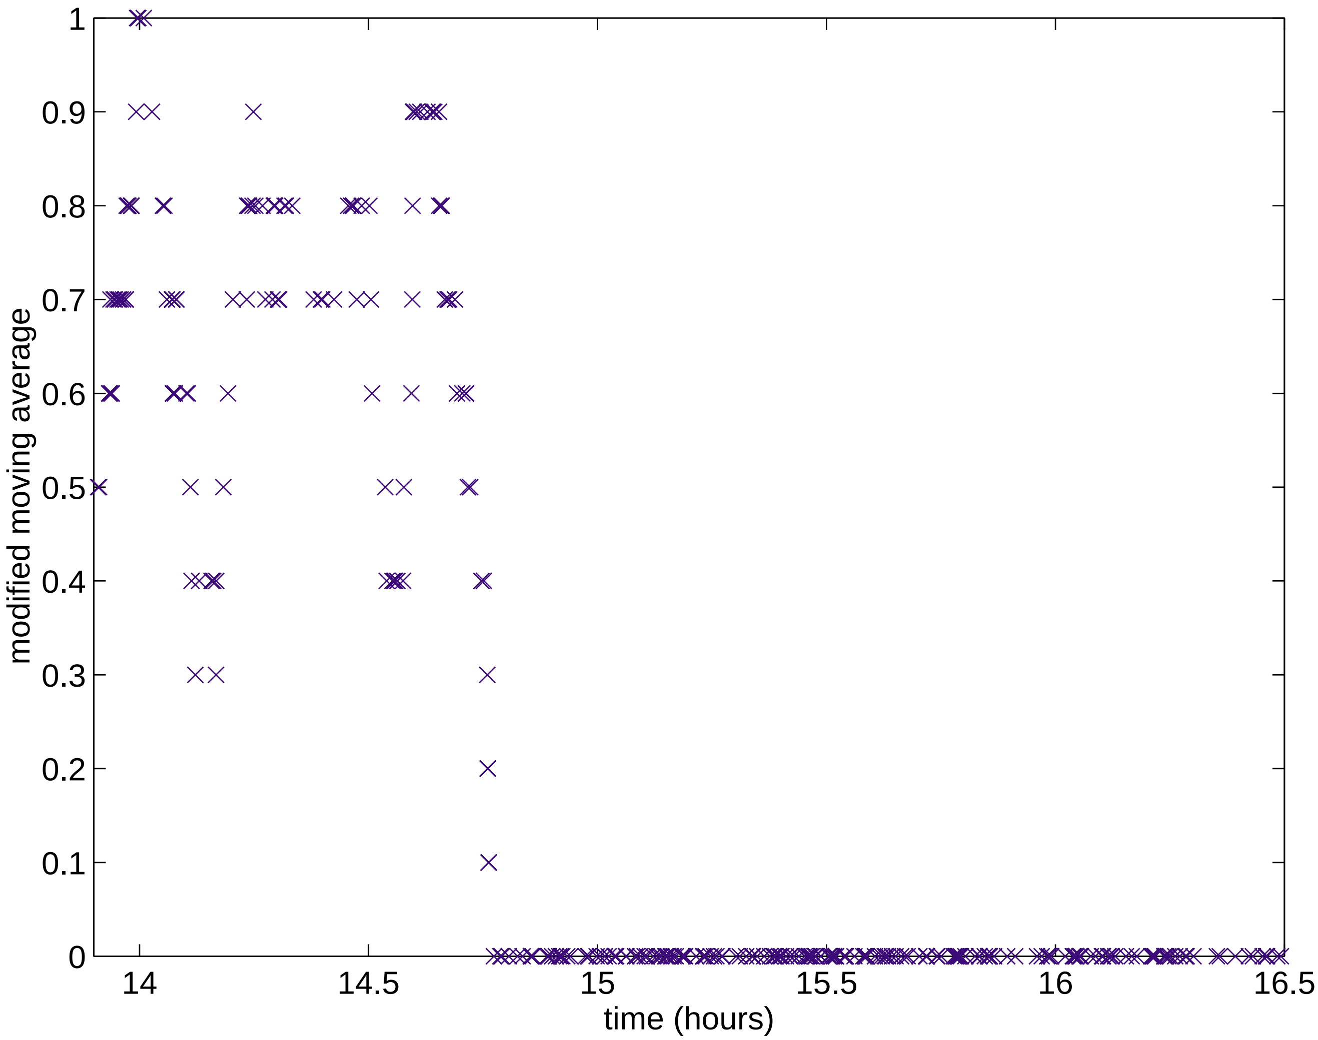

Figure 10 The resulting travel times after extending the matching to two lightly congested conditions (A) assuming each population is independent, (B) including any faster matches in the moving average for a slower group.

(A)

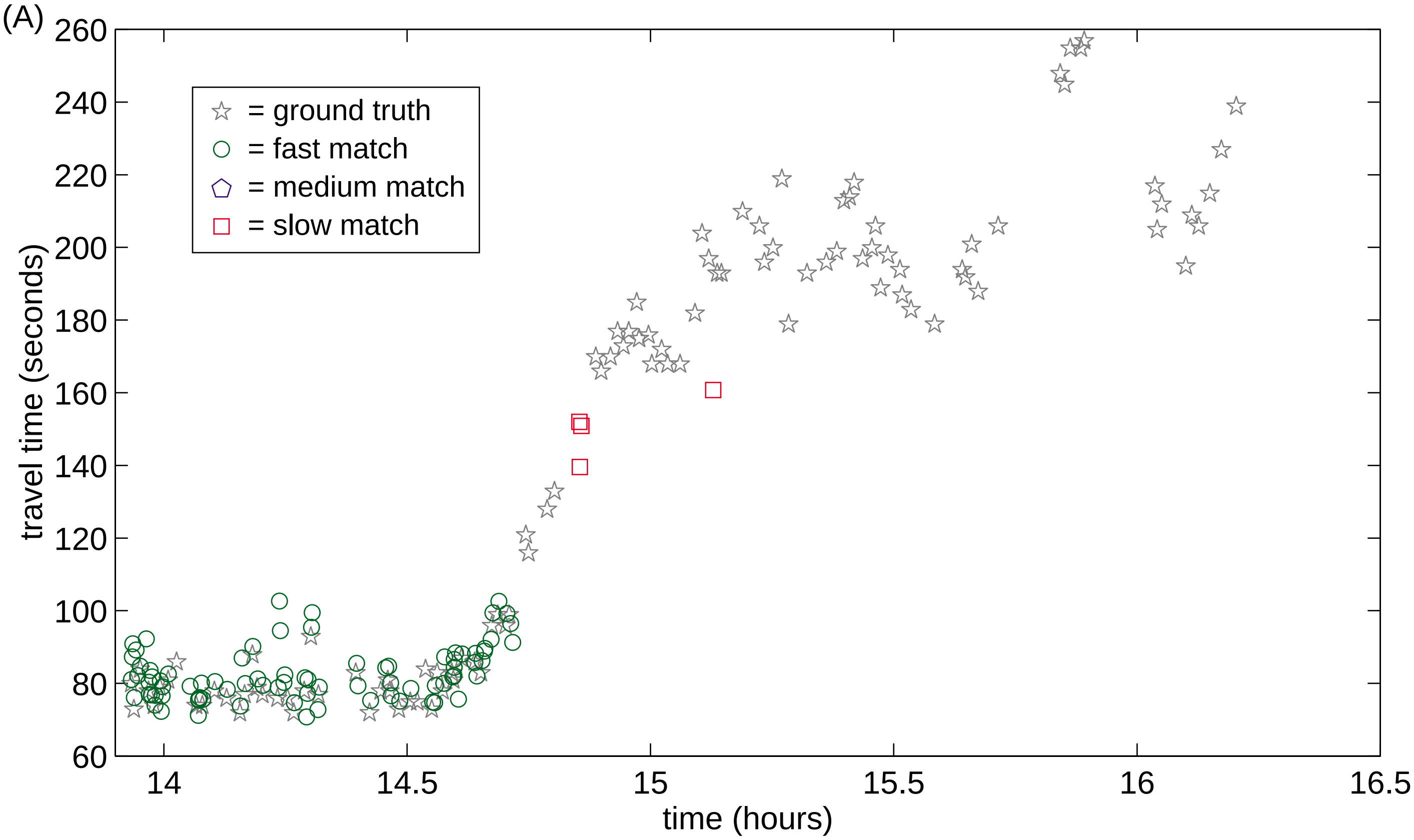

(B)

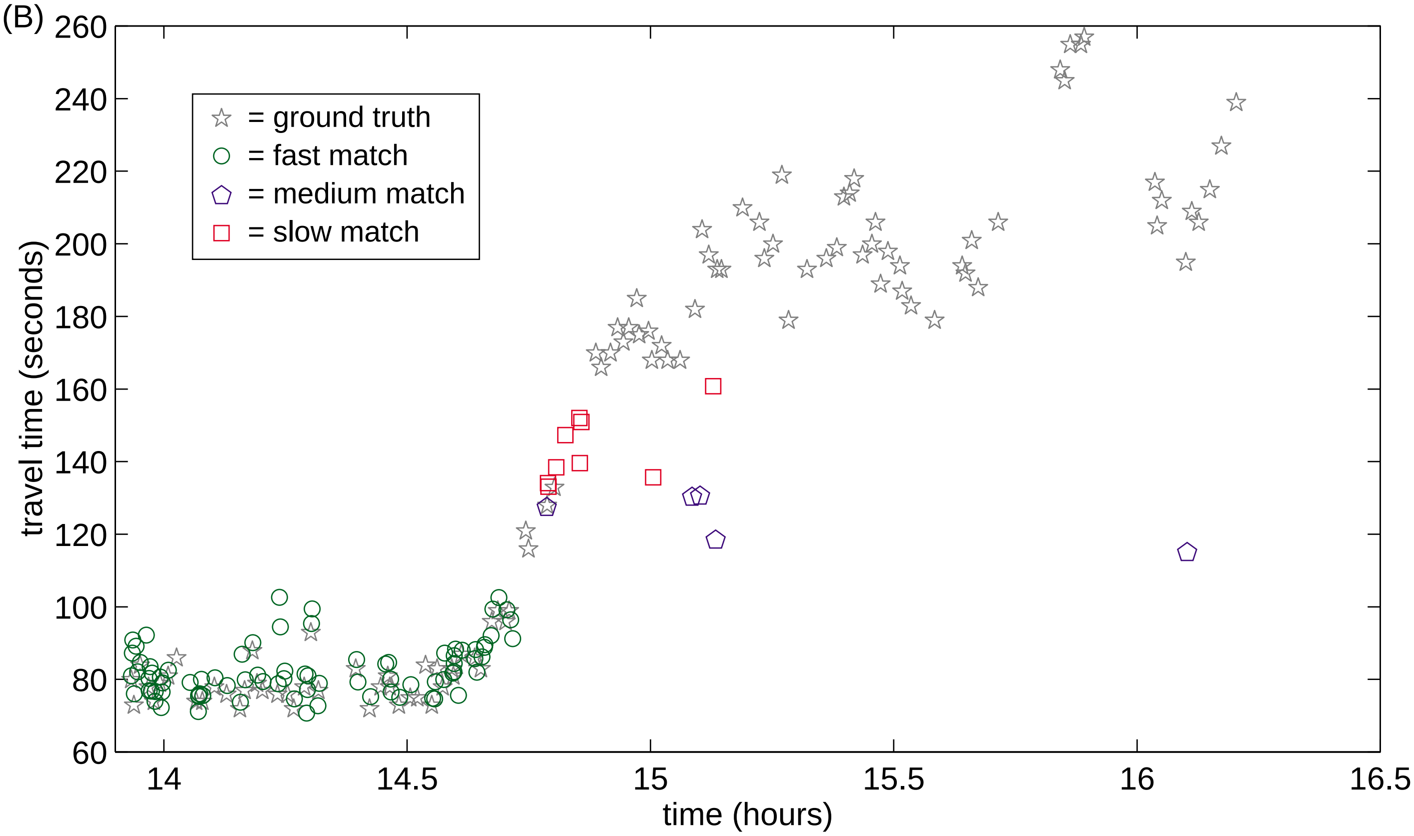

Figure 11 This figure shows the resulting travel times from the earlier matches after limiting the influence of faster matches to those vehicles that were accepted at the given level.

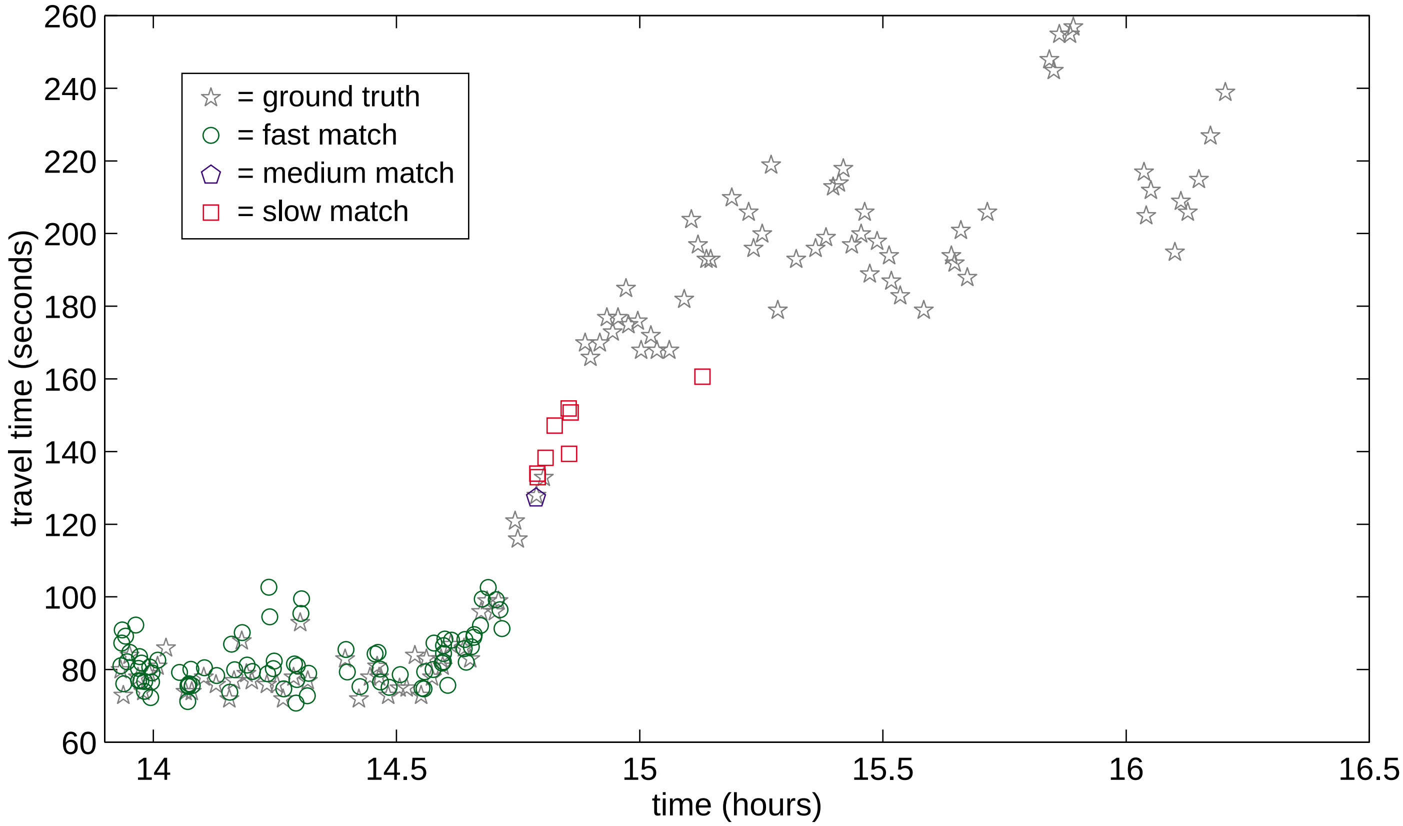

Table 1 The number of vehicles in various subgroups for the example.

Subgroup	Size
total number of vehicles in sample	4344
total number of long vehicles in the sample	320
number of ground truth matches	106
number of long vehicles before onset of congestion	115
number of fast matches	82