

Identifying the Onset of Congestion Rapidly with Existing Traffic Detectors

Benjamin Coifman, PhD

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

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

coifman.1@osu.edu
<http://www-ceg.eng.ohio-state.edu/~coifman>
614 292-4282

Accepted for publication in *Transportation Research Part A*

ABSTRACT

From an operations standpoint the most important function of a traffic surveillance system is *determining reliably* whether the facility is free flowing or congested. The second most important function is *responding rapidly* when the facility becomes congested. These functions are complicated by the fact that conventional vehicle detectors are only capable of monitoring discrete points along the roadway while incidents may occur at any location on the facility. The point detectors are typically placed at least one-third of a mile apart and conditions between the detectors must be inferred from the local measurements.

This paper presents a new approach for traffic surveillance that addresses these issues. It uses existing dual loop detector stations to match vehicle measurements between stations and monitor the entire roadway. Rather than expending a considerable effort to detect congested conditions, the research employs a relatively simple strategy to look for free flow traffic. Whenever a unique vehicle passes the downstream station, the algorithm looks to see if a similar vehicle passed the upstream station within a time window that is bounded by feasible travel times. The approach provides vehicle reidentification and travel time measurement on freeways during free flow and through the onset of congestion. If desired, other algorithms can be used with the same detectors to provide similar measurements during congested conditions. The work should prove beneficial for traffic management and traveler information applications, while promising 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, occupancy and velocity) 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 when the facility becomes congested, e.g., when an incident occurs between two detector stations it can take several minutes before measured speeds drop at either of the stations.

The limitation of point data has spurred interest in vehicle reidentification techniques, which match the observations of the same vehicle at successive detector stations, e.g., Kuhne and Immes, 1993, Balke et al. 1995, and Huang and Russell, 1997. All of these earlier works require new detector hardware to extract detailed vehicle *signatures*. Often times, these advanced technologies are developed without consideration for the general goals of traffic surveillance and as a result operating agencies may risk investing in an expensive surveillance system to capture extraneous information. The systems also risk discarding useful information, e.g., in some cases the tools collect link data but are not capable of measuring point data.

From an operations standpoint the most important function of a traffic surveillance system is *determining reliably* whether the facility is free flowing or congested. The second most important function is *responding rapidly* when the facility becomes congested. Other functions, such as quantifying the magnitude of congestion, are desirable but tertiary. Conventional loop detector surveillance satisfies the first and tertiary functions, but the response time to delays between detector stations can be excessive (Lin and Daganzo, 1997). Some of the advanced surveillance technologies promise to satisfy all of these functions, but they have yet to see widespread deployment. In contrast to investing in new detector hardware before evaluating the benefits of vehicle reidentification, this paper proposes a surveillance algorithm using existing dual loop detectors in a new way. The approach will improve the performance of these detectors on the first and second functions. It will also provide a means for operating agencies to assess

the need for the more accurate systems that require new hardware, prior to making a substantial investment.

The algorithm identifies relatively distinct vehicles¹ 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 of *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. In the event a match is found the algorithm can also address the tertiary functions, such as calculating that vehicle's travel time. In other words, the algorithm is capable of reporting free flow travel times or the fact that "traffic is not free flowing". In the latter case, a complementary reidentification method can be used with the same detectors to measure travel time during congestion if the metric proves to be desirable (Coifman and Cassidy, *forthcoming*).

To illustrate some of the potential benefits of the new system and link travel time measurement, consider incident detection using point detectors. Lin and Daganzo (1997) note that two signals emanate from an incident, a backward moving shock wave and a forward moving drop in flow. It is difficult to identify an incident based strictly on changes in point measurements associated with the forward moving wave. Reliable detection of an incident using point measurements can happen only when both of the signals have been received at the detector stations. Although the drop in flow travels at the prevailing traffic velocity, Lin and Daganzo estimated the shock wave velocity to be on the order of 8 mph. Fortunately, the drop in flow reflects the fact that vehicles are being delayed behind the incident and vehicle link travel time will increase after the signal has passed. So rather than waiting for the slow moving shock wave, the new algorithm could be used to quickly and reliably identify the onset of delay corresponding to the drop in flow.

¹ Because vehicle length can be measured at dual loop detectors, long vehicles satisfy this criterion. The reader should note, however, that the algorithm would be equally applicable to any other identifiable vehicle feature, perhaps as measured with a different type of vehicle detector.

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 must fall within a known range of free flow travel times. This concept is illustrated in Figure 1. For the present study, the free flow travel time range is defined as follows:

$$ttR_0 = \left[\frac{\text{distance}}{\max(\nu + 10, 55)}, \frac{\text{distance}}{\max(\nu - 10, 45)} \right] \quad (1)$$

where

ttR_0 = the range of feasible free flow travel times [hours],

ν = local velocity measurement at the downstream detector station [mph],

distance = the known distance between detector stations [mi].

To keep the search window as short as possible, while also being able to accept a wide range of free flow velocities, Equation 1 will select the assumed velocity range to be [45,55] mph if $\nu < 45$ mph and $\nu \pm 10$ mph if $\nu > 55$ mph.²

Chang and Kao (1991) suggest that lane change maneuvers are relatively infrequent during free flow conditions at most locations.³ 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 ttR_0 . Congestion will disrupt this relationship, both because the travel

² Note that these constants were determined empirically, but they have proven robust when applied to many different freeway links.

³ The experimental results in subsequent sections support this observation.

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.

Vehicle 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 dual loop detector, shown in Figure 2A, the controller normally records the four transitions shown in Figure 2B. After accounting for any unmatched transitions, the following parameters are calculated for each vehicle: dual loop traversal time via the rising edges, TT_r , dual loop traversal time via the falling edges, TT_f , total on-time at the first loop, OT_1 , and total on-time at the second loop, OT_2 , as shown in the figure.

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,

$$TT = \frac{2}{1/TT_r + 1/TT_f} \quad (2)$$

$$OT = \frac{2}{1/OT_1 + 1/OT_2}$$

Equation 2 is used to reduce the impact of occasional measurement errors, but the method also works using just one of the traversal times and one of the on-times for each vehicle. From Figure

2A, vehicle velocity is simply the loop separation, 20 ft for this study, divided by the traversal time. The effective vehicle length⁴, L , is the velocity multiplied by the on-time,

$$L = \frac{20 \cdot OT}{TT} \quad [\text{ft}] \quad (3)$$

The controllers used for this study sample the loops at 60 Hz, so at best, each parameter in Equation 2 is accurate to 1/60 seconds. Assuming the times from Figure 2 are expressed in seconds, the length range, LR , is defined as:

$$LR = [\text{minimum length estimate}, \text{maximum length estimate}] \\ = \left[20 \cdot \frac{OT - 1/30}{TT + 1/30}, 20 \cdot \frac{OT + 1/30}{TT - 1/30} \right] \quad [\text{ft}] \quad (4)$$

and the measurement uncertainty is defined as the difference between the maximum and minimum length estimates. Of course the occasional measurement error, such as a vehicle changing lanes over the detector, will result in an erroneous LR for that vehicle. As will be shown in a subsequent section, the methodology was specifically designed to accommodate these errors. Finally, to ensure the best measurements possible, any hardware problems such as cross talk between detectors are identified using the methodology presented in Coifman (1999) and corrected.

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.

⁴ The "effective vehicle length" is simply the length measured by the detector. As shown in Figure 2, it includes the length of the detection zone. Provided care is taken to ensure that all detectors are set to the same sensitivity, this additional length does not impact the methodology.

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 ft and 23 ft. During free flow conditions, the measurement uncertainty is on the order of 2 ft for these short vehicles, making difficult the task of differentiating between them. In contrast, some length observations are as long as 80 ft. The large range of feasible lengths and the lower frequency of observations for the long vehicles make it possible to differentiate between them even at free flow velocities.

ALGORITHM IMPLEMENTATION

Using an example to illustrate the algorithm implementation, consider the 1.3 mi freeway segment from the Berkeley Highway Laboratory (Coifman et al., 2000) shown in Figure 4. To eliminate the common vehicles, all downstream vehicles shorter than 23 ft are ignored. Whenever a long vehicle passes the downstream dual loop detector, 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 4, 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 later upstream arrival is considered the possible match. Otherwise, the downstream vehicle does not have an identifiable free flow match in the lane.

Obviously, a free flow vehicle will not have a match in the same lane if the vehicle changed lanes, entered the freeway between detectors, 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. Figure 5A shows this moving average for just over 2.5 hours of data from the two dual loop detectors. Applying a threshold to these

data, Figure 6 shows the travel times for all of the long vehicles that had a possible match and a moving average over 0.5. These free flowing matches will be referred to as *fast matches*.

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 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. To this end, the study used video data recorded concurrently with the dual loop detector data and an observer matched all visually distinct vehicles that passed both detectors in the lane during the study period. The resulting travel times 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 the long vehicles were matched. At first glance, this low rate of detection may not seem very useful. But one must keep in mind that traffic is free flowing when the matches are found, so travel times are relatively constant. Furthermore, the goal of the algorithm is to detect when travel times start to increase, which is the subject of the next section.

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 corresponds to the time when the ground truth travel times start increasing due to a queue overrunning the downstream station. Although the measured travel times in Figure 6 are useful for traffic surveillance, the true diagnostic power of the method comes from the moving average in Figure 5A. 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, it will not have a match in the same lane. On the other hand, the long distance between detectors increases the time range in ttR_0 and thus, increases the probability of finding a false positive during congested conditions.

To filter out most of the false positives, consider the number of unmatched vehicles preceding each *fast match*, as shown in Figure 7A. 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 7B 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 7B are discarded. Figure 5B shows the results after recalculating the moving average over all outcomes. Note that in this example, the process has eliminated all of the noise during congestion.

EXTENDING SURVEILLANCE INTO CONGESTED CONDITIONS

Looking closer at the ground truth travel times in Figure 6, there is a transition as the queue first overruns the downstream detector and eventually covers the entire link. This transition is characterized by increasing travel times. In an attempt to capture the increasing delays during the transition, four additional travel time ranges are defined:

$$\begin{aligned}
 ttR_1 &= \left[\frac{\text{distance}}{50}, \frac{\text{distance}}{40} \right]; & ttR_2 &= \left[\frac{\text{distance}}{45}, \frac{\text{distance}}{35} \right]; \\
 ttR_3 &= \left[\frac{\text{distance}}{40}, \frac{\text{distance}}{33} \right]; & ttR_4 &= \left[\frac{\text{distance}}{35}, \frac{\text{distance}}{28} \right]
 \end{aligned} \tag{5}$$

where the denominators in the bracketed expressions bound the possible link velocities [mph] for each set. Notice how all five ranges overlap.

As a queue grows across the link, the true travel times will pass from ttR_0 to ttR_1 , and so on through ttR_4 , until finally the travel time exceeds all five ranges. Repeating the analysis presented in the previous section (i.e., Figure 5B) for ttR_1 and ttR_2 yields the dashed lines in Figure 8A. Notice how the curve for ttR_1 starts increasing before ttR_0 drops due to congestion, and similarly the curve for ttR_2 relative to ttR_1 . This plot also shows a few false positives, such as the rise in ttR_2 at 14.85 hrs. Once the link becomes sufficiently congested, the history implicit in the moving average is lost and the method can not differentiate between false positives and true matches. To reduce the influence of false positives in the longer travel time ranges, the algorithm will only accept a sequence of non-zero values if the next faster ttR moving average is non-zero at the start of the sequence. Thereby allowing the algorithm to follow increasing travel times, but blocking any long travel times once the travel time has exceeded ttR_4 . This process is evident in the right hand side of Figure 8B, where the false positives have been discarded. Finally, the algorithm selects the range with the highest average value at the given instant, as shown in Figure 8C. Figure 9 shows the resulting travel times after extending the method to ttR_3 and ttR_4 . The algorithm does not find any matches once the queue has covered the entire link and the ground truth travel times level off.

RESPONSE TIME AND ALGORITHM EXTENSIONS

After the onset of congestion in the preceding example, it took 3.5 minutes for the ttR_0 moving average to go to zero (Figure 5B), clearly indicating that the link travel time had exceeded the time range. Combining the information from multiple time ranges yields a more responsive indicator. The moving average for ttR_1 exceeds the respective measurement for ttR_0 just 15 seconds after the last free flow vehicle was observed (Figure 8B), rapidly indicating that the travel time has exceeded ttR_0 . In contrast, using the empirically measured end-of-queue velocity reported in Lin and Daganzo (1997), it would have taken approximately 9 minutes for the queue to span the

entire link. Examination of the local detector data revealed that it actually took the queue 10.5 minutes to span the link in this case. After the queue reaches the upstream detector, following conventional practice, it would then take several additional sample periods⁵ to differentiate between a transient event and an actual drop in velocity. Of course the response time of the present algorithm is a function of truck frequency. In this case the long vehicles made up almost 10 percent of the population. If another facility has a lower truck frequency, one could change the parameters to give the long vehicles greater weight, e.g., using a smaller moving average in Figure 5, since these vehicles would be more unique.

Although the example considered recurring congestion, it is reasonable to assume that vehicle movement within a queue should be similar regardless of whether the delay is recurring or due to an incident. Any deviations due to an incident, such as lane change maneuvers out of a blocked lane, would only serve to improve the response time of the algorithm. In fact, an incident during low flow conditions may cause an increase in lane change maneuvers without any measurable changes in local velocities or even the link travel time.

Although no effort was made in this algorithm to detect lane changes, an obvious extension would be to look across all upstream lanes rather than a single lane. Nonetheless, as demonstrated in this paper, the algorithm works using data from a single lane. One could also treat each lane as an independent process and integrate the results from all lanes for a more robust indicator. For example, if several adjacent lanes simultaneously showed a drop in the ttR_0 moving average, one could respond before any of them reached zero. On the other hand, if lane changes are frequent under free flow conditions, such as in a weaving section, one could map the normal pattern of movements. Then, for a given downstream observation, explicitly look in a different upstream lane. Future research will address these points and attempt to optimize the algorithm

⁵ The typical sampling period is on the order of 30 seconds in conventional practice.

parameters, e.g., changing the minimum vehicle length, the travel time range(s), and the number of vehicles in the moving average.

As it stands, the algorithm has been running in real-time for several years, across seven detector stations in the Berkeley Highway Laboratory (Coifman et al., 2000). Although only a few hours of ground truth data have been generated, this long testing period has provided ample evidence of the algorithm's performance. Figure 10 shows an entire day's worth of travel time data from five lanes in one direction on a one-third mile link. The algorithm's measurements are shown with dark points. The algorithm found the fewest matches in lane one⁶ and the most in lane three, 542 and 2,100 matches, respectively. Throughout the entire day the algorithm gave consistent results across all five lanes. For reference, this example uses a second algorithm to match vehicles during congested conditions, as shown with light points (Coifman and Cassidy, *forthcoming*).⁷ It is worth noting that both algorithms were run throughout the entire day; each one is self-selecting, only matching vehicles when it can.

Obviously, it would be desirable to test the methodology at many locations. But the research has been constrained by the simple fact that most operating agencies do not collect data on individual vehicles. Conventional practice aggregates the measurements over fixed sample periods in the field.

CONCLUSIONS

This paper has developed a new traffic surveillance strategy using existing detectors. Rather than reporting local conditions at the detectors, the strategy identifies periods when the link between two detector stations becomes congested. Whenever a unique vehicle passes the downstream

⁶ Due to a truck restriction, only 2 percent of all vehicles in this lane were long enough to be considered by the algorithm.

⁷ The second algorithm is more complicated and requires slower traffic to reduce the length measurement uncertainty in the dual loop detector data. Also note that the total number of matches reported above does not include any matches from the second algorithm.

station, the algorithm looks to see if a similar vehicle passed the upstream station within a time window that is bounded by feasible travel times. This process showed good performance over a 1.3 mi 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 has been implemented using the existing Model 170 controllers, which are based on 20 year old computer technology. The key to this implementation is using the controllers to relay the data to a central server, as discussed in Coifman et al. (2000). For this limited deployment, the additional cost per station was roughly one percent of the cost for a new detector station. Presumably these costs would be even lower for a large scale deployment.

Although this paper used dual loop detectors to demonstrate the algorithm, the methodology would be equally applicable to any other identifiable vehicle feature, perhaps as measured with a different type of vehicle detector. In fact Coifman and Banerjee (2002) has demonstrated that the algorithm can also be deployed with single loop detectors.

To place the work in context, the algorithm does have a lower reidentification rate than the other methods that require new hardware, but perhaps the higher rate is not necessary. One could view the algorithm as a low cost means to investigate the benefits of vehicle reidentification and travel time data before investing in a new surveillance system. In any event, the algorithm is intended to augment, rather than supplant, conventional point detector measurements. By combining point detector 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.

The Contents of this report reflect the views of the author who is responsible for the facts and accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the State of California. This report does not constitute a standard, specification or regulation.

REFERENCES

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

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

Coifman, B. (1999) "Using Dual Loop Speed Traps to Identify Detector Errors", *Transportation Research Record no. 1683*, Transportation Research Board, pp 47-58.

Coifman, B., Lyddy, D., and Skabardonis, A. (2000) "The Berkeley Highway Laboratory-Building on the I-880 Field Experiment", *Proc. IEEE ITS Council Annual Meeting*, pp 5-10.

Coifman, B., Banerjee, B. (2002) "Vehicle Reidentification and Travel Time Measurement on Freeways Using Single Loop Detectors- From Free Flow Through the Onset of Congestion", *Proc. of the 6th International Applications of Advanced Technologies in Transportation Engineering*, ASCE, Cambridge, MA.

Coifman, B., Cassidy, M. (*forthcoming*) "Vehicle Reidentification and Travel Time Measurement on Congested Freeways", *Transportation Research- Part A*. Draft available at: <http://www.ceegs.ohio-state.edu/~coifman/documents>

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

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

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

FIGURE CAPTIONS

- Figure 1, One vehicle traversing an extended link between two detector stations, 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 dual-loop-detector, (A) the two 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 turn-on and turn-off transitions at each detector.
- Figure 3, (A) Cumulative distribution of individual vehicle lengths over 24 hours at one detector station. (B) Detail of part A.
- Figure 4, The segment of Interstate-80 in Berkeley, California used to illustrate and verify the Surveillance Algorithm.
- Figure 5, (A) 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. (B) 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 6, Measured travel times for matched vehicles ("fast matches") and ground truth travel times. Note that the algorithm did not find any matches after the true travel times increased due to congestion.

- Figure 7, (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 hours. (B) Moving sum of two measurements from part (A).
- Figure 8, (A) Repeating the analysis shown Figure 5 for the first three time ranges. (B) Only accept non-zero values for the new time ranges if the immediately faster range was non-zero at the start of the period. (C) Select the time range with the highest moving average at the given instant.
- Figure 9, Measured travel times for matched vehicles from the five travel time ranges, as indicated, and the ground truth travel times.
- Figure 10, Measured travel time across five lanes over one-third mile for an entire day using the algorithm in this paper (dark points) and a complementary algorithm during congestion (light points).

TABLE CAPTIONS

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

Figure 1, One vehicle traversing an extended link between two detector stations, 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.

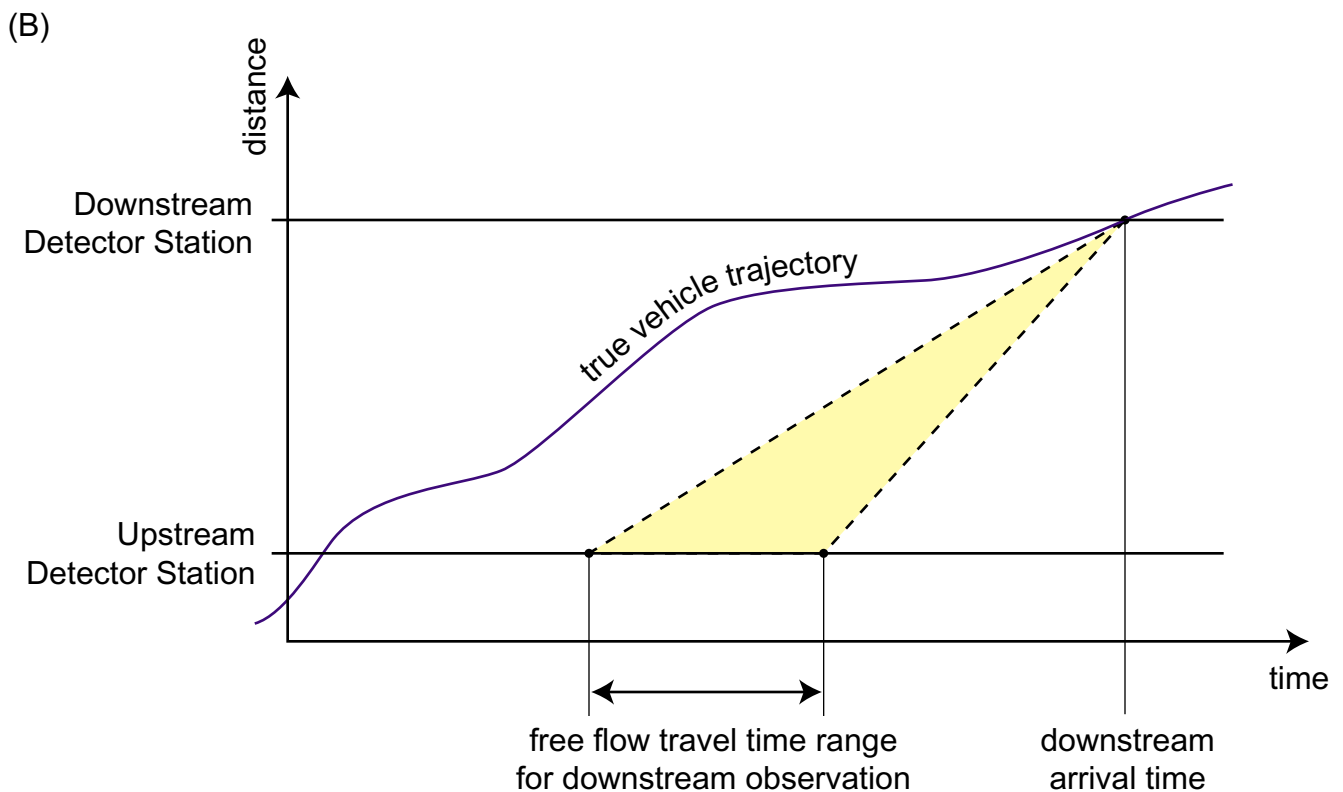
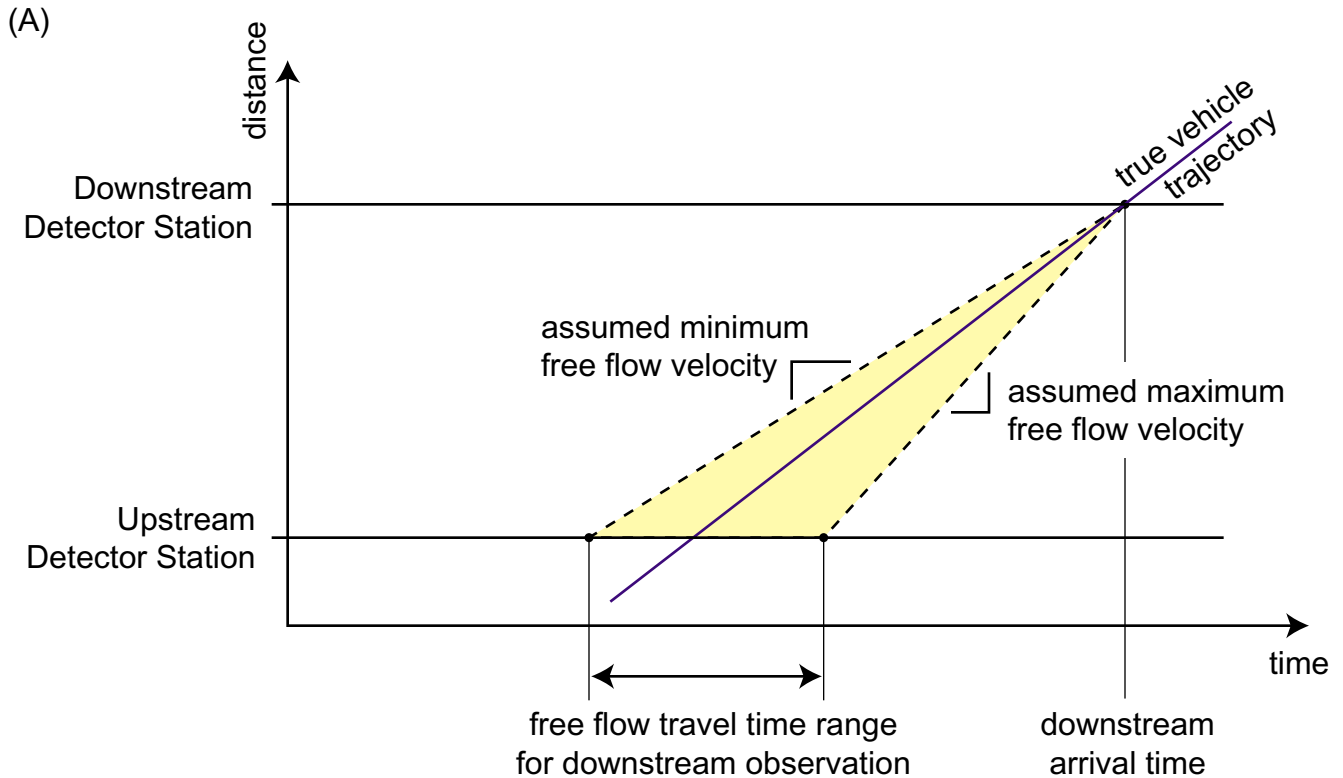


Figure 2, One vehicle passing over a dual-loop-detector, (A) the two detection zones and the vehicle trajectory as shown in the time space plane. The height of the vehicle's trajectory reflects the non-zero vehicle length. (B) The associated turn-on and turn-off transitions at each detector.




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


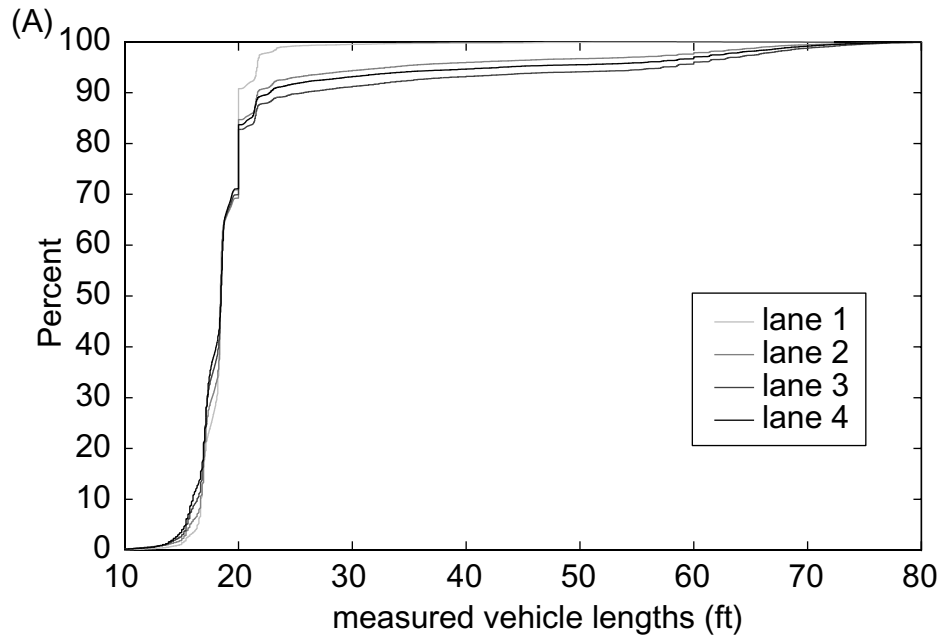


Figure 4, The segment of Interstate-80 in Berkeley, California used to illustrate and verify the Surveillance Algorithm.




Figure 5, (A) 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. (B) 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.

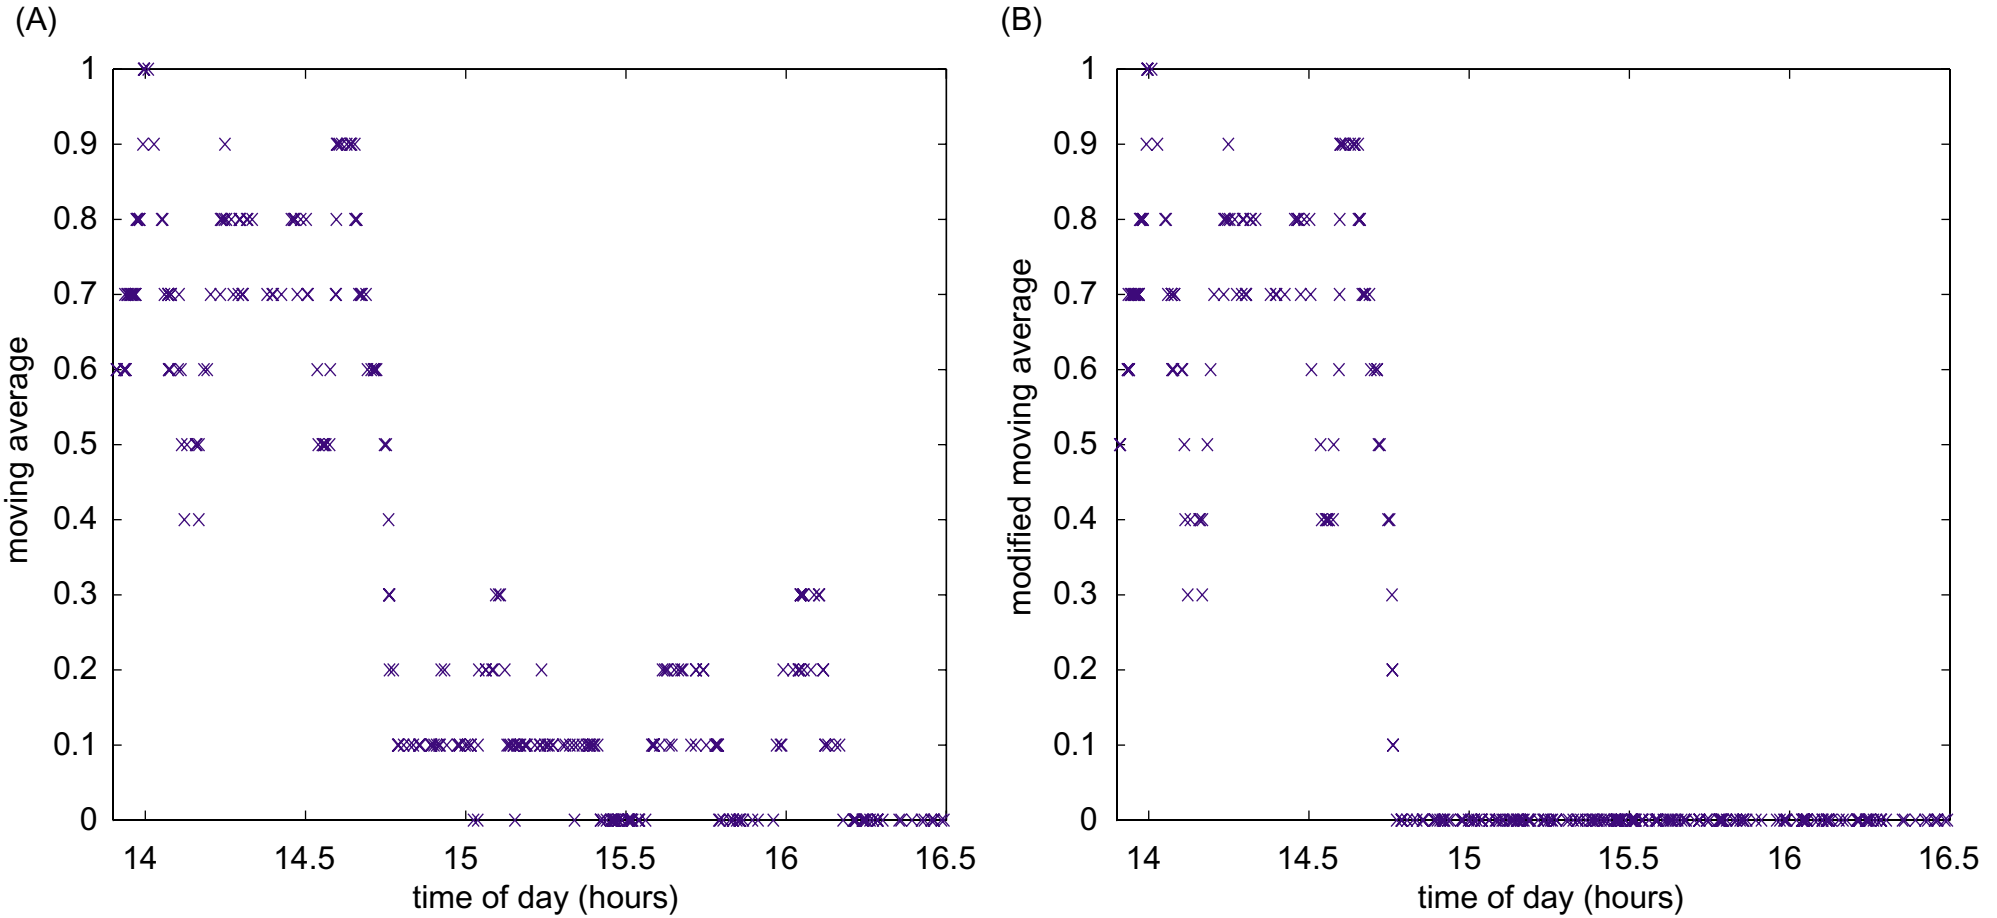


Figure 6, Measured travel times for matched vehicles ("fast matches") and ground truth travel times. Note that the algorithm did not find any matches after the true travel times increased due to congestion.

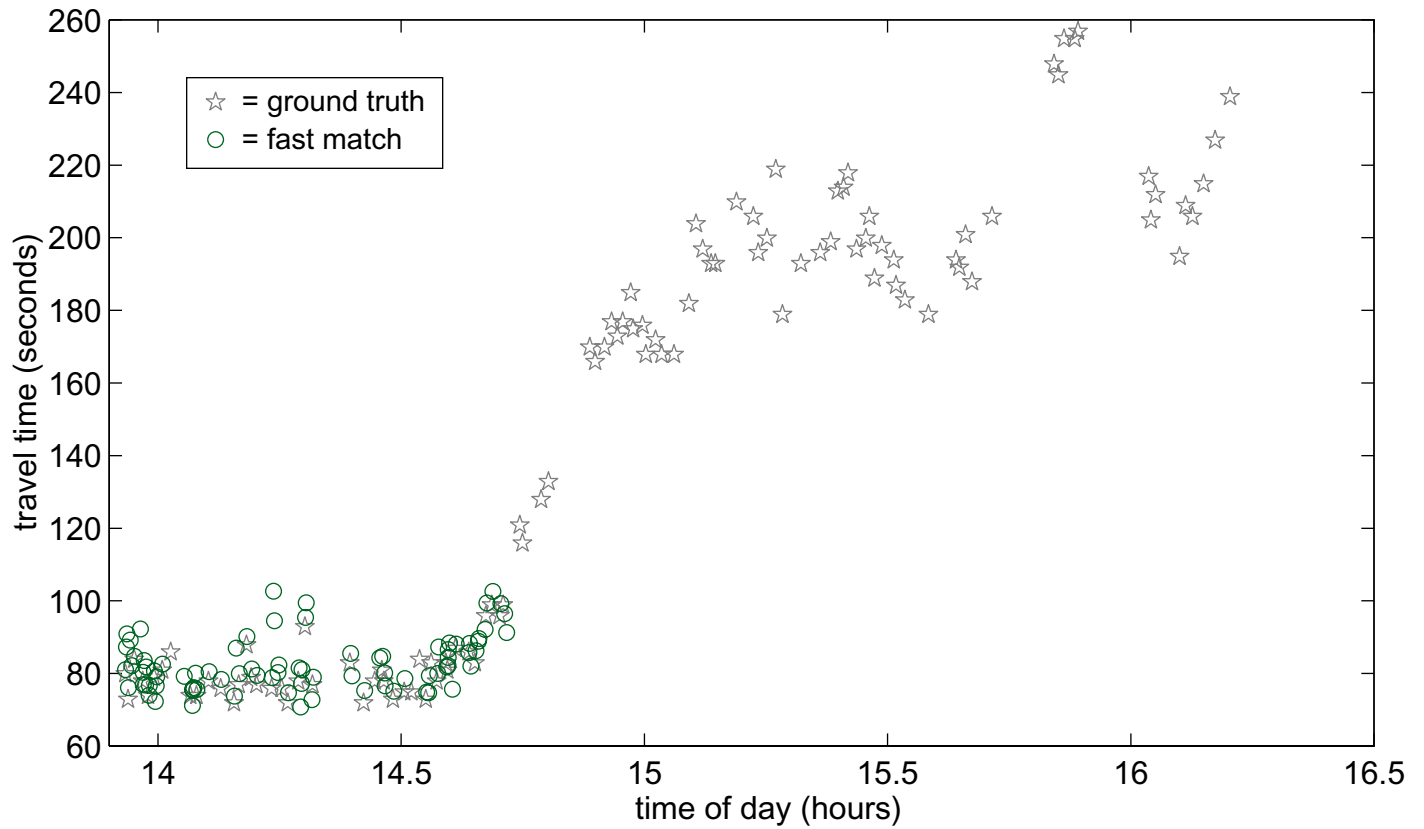


Figure 7, (A) The number of unmatched vehicles preceding each fast match. The free flow matches typically have few unmatched vehicles compared to the congested matches after 14.7 hours. (B) Moving sum of two measurements from part (A).





Figure 8, (A) Repeating the analysis shown Figure 5 for the first three time ranges. (B) Only accept non-zero values for the new time ranges if the immediately faster range was non-zero at the start of the period. (C) Select the time range with the highest moving average at the given instant.




Figure 9, Measured travel times for matched vehicles from the five travel time ranges, as indicated, and the ground truth travel times.

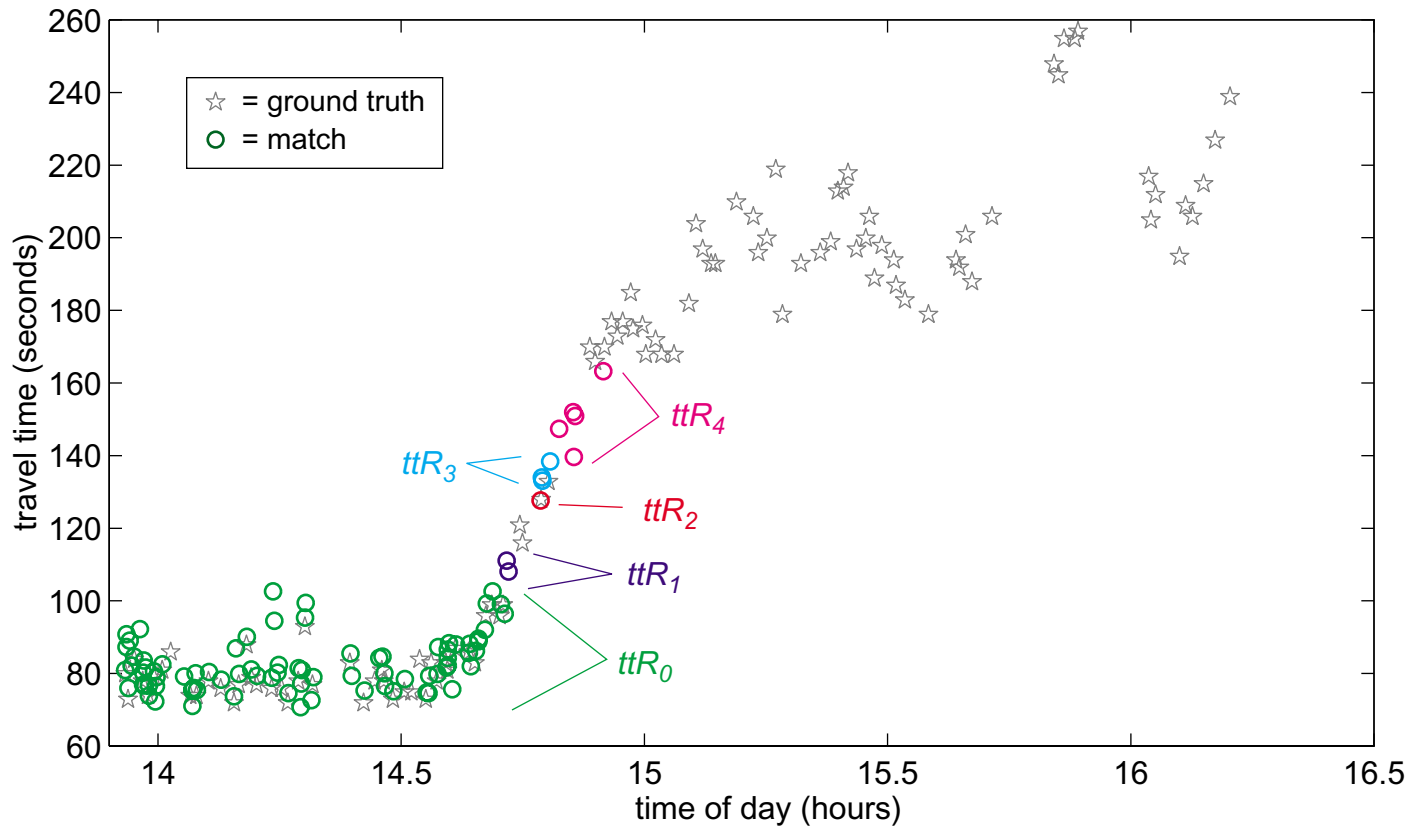
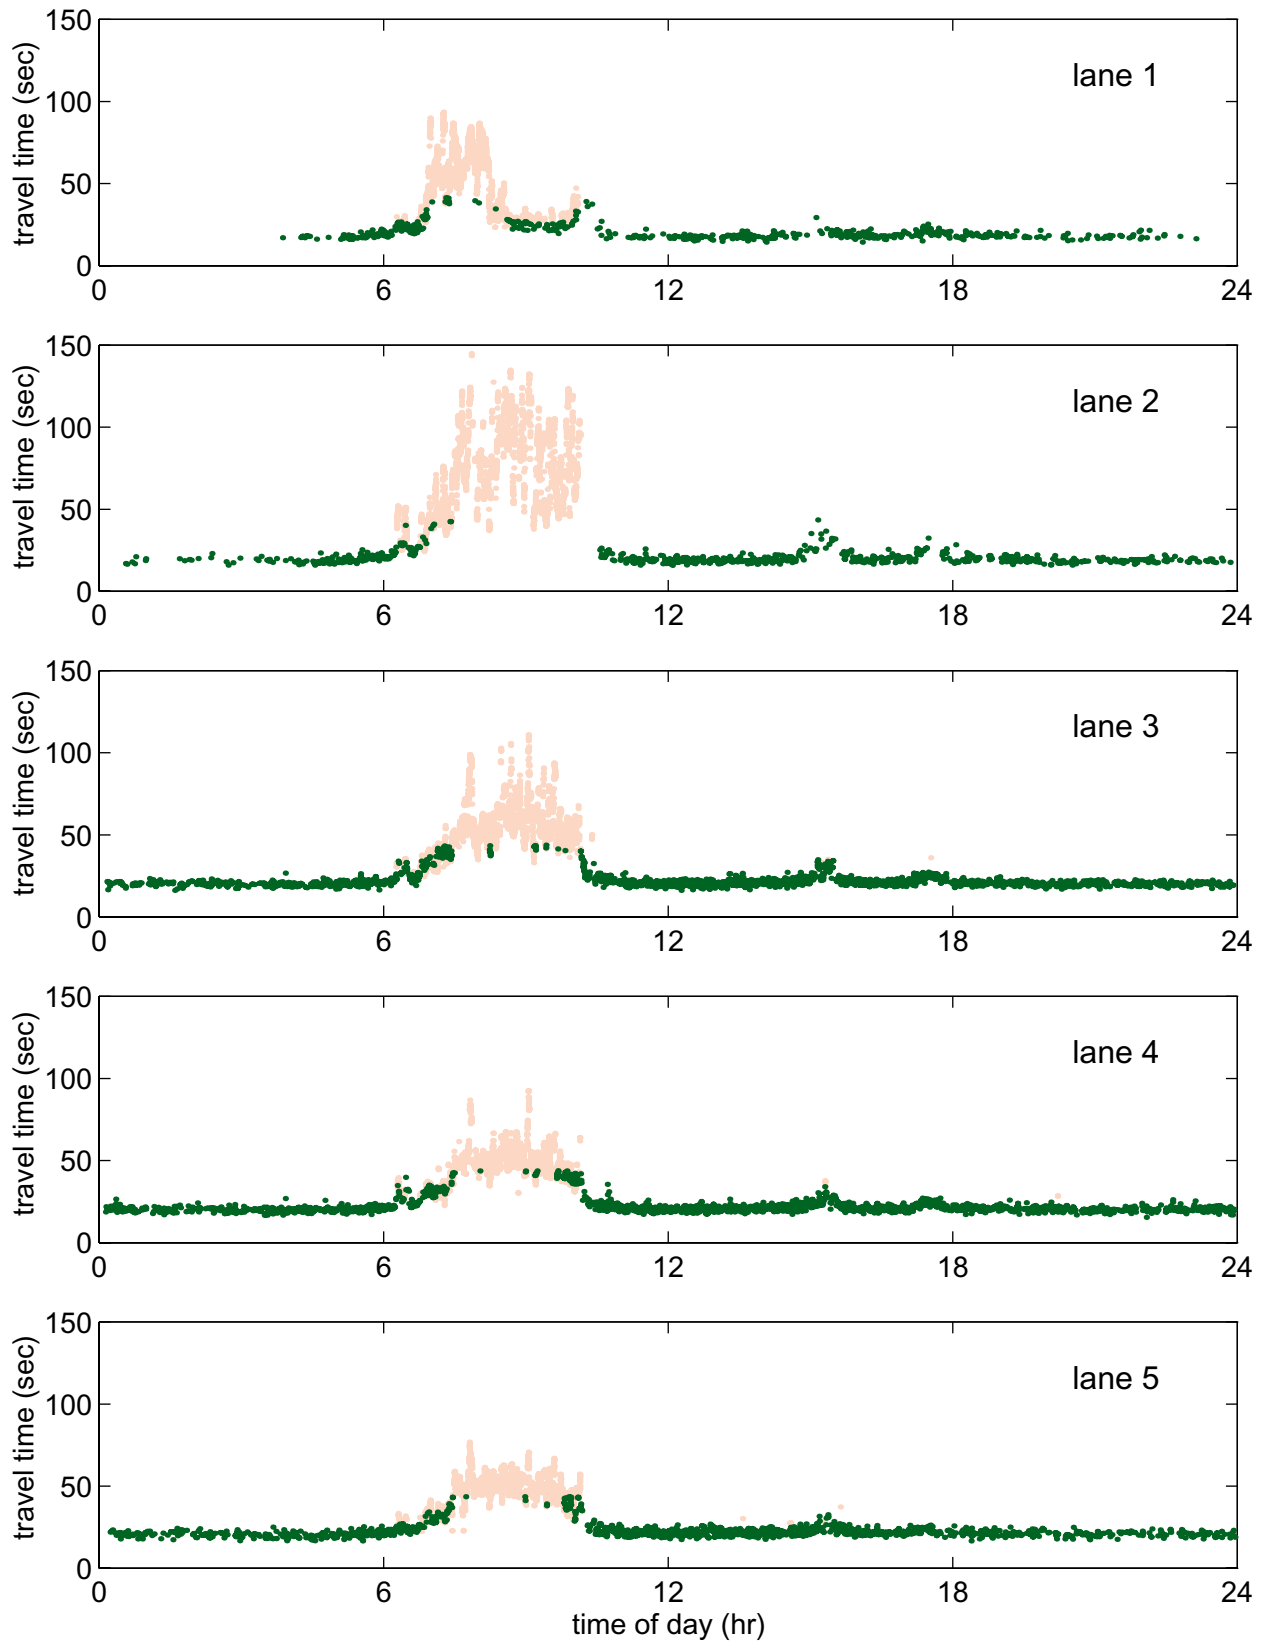


Figure 10, Measured travel time across five lanes over one-third mile for an entire day using the algorithm in this paper (dark points) and a complementary algorithm during congestion (light points).



Coifman, B.

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

| <u>Subgroup</u> | <u>Size</u> |
|--|-------------|
| total number of vehicles in the 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 |