Vehicle Reidentification and Travel Time Measurement Across Freeway Junctions Using the Existing Detector Infrastructure

Benjamin Coifman
Associate Professor,
Dept. of Civil and Environmental Engineering and Geodetic Science, and
Dept. of Electrical and Computer Engineering,
The Ohio State University,
Hitchcock Hall 470,
2070 Neil Ave.,
Columbus, OH 43210, USA
http://www.ceegs.ohio-state.edu/~coifman
614 292-4282
Coifman.1@osu.edu

Sivaraman Krishnamurthy
Graduate Research Assistant,
Dept. of Civil and Environmental Engineering and Geodetic Science,
The Ohio State University,
Hitchcock Hall 470,
2070 Neil Ave.,
Columbus, OH 43210, USA
krishnamurthy.26@osu.edu

ABSTRACT

Conventional vehicle detectors are capable of monitoring discrete points along the freeway but do not provide information about conditions on the link between detectors. Knowledge of conditions on the link is useful to operating agencies for enabling timely decisions in response to various delay causing events and hence to reduce the resulting congestion of the freeway system. This paper presents an approach that matches vehicle measurements between detector stations to provide information on the conditions over the link between the detectors rather than relying strictly on the point measurements from the detectors. In particular this work reidentifies measurements from distinct vehicles using the existing loop detector infrastructure. Here the distinct vehicles are the long vehicles, but depending on the vehicle population or type of detector used, one might chose to use some other reproducible feature.

This new methodology represents an important advancement over preceding loop based vehicle reidentification, as illustrated herein, it enables vehicle reidentification across a major diverge and a major merge. The examples include a case where the reidentification algorithm responded to delay between two detector stations an hour before the delay was locally observable at either of the stations used for reidentification. While previous loop based reidentification work was limited to dual loop detectors, the present effort also extends the methodology to single loop detectors; thereby making it more widely applicable. Although the research uses loop detector data, the algorithm would be equally applicable to data obtained from many other traffic detectors that provide reproducible vehicle features.

Keywords: freeway traffic monitoring, vehicle reidentification, travel time, loop detectors

1 INTRODUCTION

The most important task of a traffic surveillance system is to reliably determine whether a facility is free flowing or congested. Conventional vehicle detectors are capable of monitoring discrete points along the roadway but do not provide information about conditions on the link between detectors. This heretofore-unavailable information on the link is useful to operating agencies for making timely decisions in response to various delay causing events and hence to reduce the resulting congestion of the freeway system. This paper presents an approach that matches vehicle measurements between detector stations to provide information on the conditions in the link rather than simply extrapolating the point measurements from the detectors to the link.

The basic idea behind the work is to identify 'distinct' vehicles at the downstream detector station, and for each of these vehicles, try to find a matching measurement from the same vehicle in the data from the upstream detector station. Here the distinct vehicles are the long vehicles measured from loop detectors, but depending on the vehicle population or type of detector used, one might chose to use some other reproducible feature. The travel time for traversing the distance between the two detector stations can be found by subtracting the times at which a matched vehicle was observed. After accounting for noisy measurements, if the travel times of matched vehicles fall inside a time window that represents free flow travel time, the link is considered to be in free flow conditions, otherwise the link is considered to be congested. The algorithm can report link travel times or link velocity[1] from free flow conditions down to an average link velocity of 20 mph. As a link becomes congested link velocity will decrease from free flow velocity and usually most of the link will have to be congested before the link velocity drops below 20 mph. So the operating range of the algorithm ensures that typically one or more delayed vehicles will be detected when delays occur and before the link velocity drops below 20 mph.

As has already been demonstrated in Coifman (2003), the starting point for this paper, vehicle reidentification can be used to detect the onset of congestion rapidly over a freeway link using the existing detector infrastructure. The present work holds the same objective, but sets out to develop a more robust algorithm than the earlier effort. The most significant difference between this work and the earlier effort is that the present research does not limit vehicle reidentification to only one lane, which allows vehicle reidentification even when the candidate vehicle changes lanes. This approach allows for reidentification even when many vehicles pass only one of the detector stations without being observed at the other station, i.e., the algorithm is more robust to unmatchable vehicles. While Coifman (2003) presents an algorithm to rapidly detect the onset of congestion, that algorithm does not work once congestion sets in and it ceases reidentifying vehicles until free flow periods return, in effect, the earlier work offered an indicator of free flow or congestion. Separately, we developed an algorithm for vehicle reidentification during congested periods on freeways that seeks out distinct sequences of measured vehicle lengths in the congested traffic stream (Coifman and Cassidy, 2002). This second algorithm assumes that platoons of 5-10 vehicles regularly pass both detector stations in the same lane. Unfortunately conventional dual loop detectors are sampled at 60 Hz, which limits the length measurement resolution. During free flow it becomes impossible to differentiate between most passenger

[1] Link velocity is proportional to the inverse of link travel time.

vehicles. It is only when traffic slows that vehicles will be over the detector long enough to measure their lengths with sufficient precision to differentiate between different passenger vehicles, hence the limitation to congested periods. However, this second algorithm fails in the presence of frequent lane change maneuvers because it assumes that platoons of vehicles regularly pass both detector stations in the subject lane. Taken together, the two separate algorithms can be used to monitor traffic in both free flow and congested conditions (e.g., Coifman et al, 2000), but neither algorithm alone offers comprehensive coverage under both conditions. These earlier vehicle reidentification algorithms (Coifman and Cassidy, 2002; Coifman, 2003) have been running in real time, continuously over six links of the Berkeley Highway Laboratory (BHL) now for over 7 years. The deployment uses the existing detector hardware, and communication channels deployed by Caltrans for conventional traffic monitoring. In addition to providing travel time data, the stations also provide conventional loop detector measurements.

This paper presents a significantly improved algorithm that borrows good ideas from both of our earlier efforts and then goes further. As with Coifman (2003), the present algorithm focuses on vehicles that offer distinct length measurements under all traffic conditions, i.e., the long vehicles. But the recent history is stored in a matrix similar to Coifman and Cassidy (2002). Although employing a similar means of storage, the matrix and the processing in this paper is distinctly different from the earlier work. To accommodate a high frequency of false positives the previous work indexed vehicles by arrival number and searched for sequences in the matrix, as such it could only tolerate a small amount of reordering due to lane change maneuvers or detector errors. The present work eliminates many of the false positives by excluding the common passenger vehicles, i.e., the shorter vehicles, thereby allowing for greater flexibility in the processing. We index potentially matching vehicles by the resulting travel time if a given pair were indeed a match, allowing for much greater reordering among vehicles as they travel between stations before the algorithm is impeded, and by extension, the new algorithm works for larger station spacing compared to the earlier sequence based algorithm. Furthermore, because both of our earlier algorithms rely on sequence information within the given lane to eliminate false positives, they preclude matching across lanes. The present work does not use such lane sequence information and thus, as will be demonstrated herein, it can be applied across multiple lanes.

In short, the present paper combines much of the functionality of our two previous approaches into a single algorithm, extends the performance, and then applies the work to very challenging conditions. Because the new algorithm is much more robust to vehicle reordering and unmatchable vehicles that enter or leave a subject lane between detector stations, this improved algorithm allows for reidentification across major merges or diverges in the freeway segment, where one cannot assume that most vehicles travel along the same lane or even that a given vehicle passes both stations. While earlier research has examined vehicle reidentification across ramps (e.g., Coifman, 2003), to our knowledge, this paper presents the first published attempt by anyone to use loop detector data to reidentify vehicles across a major merge or major diverge. In fact several examples herein examine detector station pairs that span both a merge and a diverge of two interstate freeways. The examples include a case where the reidentification algorithm responded to delay between two detector stations an hour before the delay was locally observable at either of the stations used for reidentification. While the previous work was limited to dual loop detectors, the present effort also extends the methodology to single loop detectors and thereby making it more widely applicable. Although the research uses loop detector data, it

would be equally applicable to data obtained from any other traffic detector that provides a reproducible vehicle feature.

Other emerging vehicle reidentification methods use vehicle signatures from video cameras (e.g., Huang and Russell, 1997; MacCarley, 2001), a much more detailed inductive signature available from new loop detector sensor models[2] (e.g., Kuhne and Immes, 1993; Kwon, 2006), and Automatic Vehicle Identification (AVI) (e.g., Hill and Emmott, 1989, Dion and Rakha, 2006), all of which require new hardware. Other technologies promise travel time measurement by explicitly tracking vehicles and eliminating the need for reidentification altogether, such as cell phone tracking (e.g., Turner, 1995; Yim and Cayford, 2006) or explicitly employing automatic vehicle location (AVL) technology (e.g., Itoh, 1986; Schafer et al, 2002), but here too these approaches require an extensive deployment of new hardware. To date, most of the above systems have only seen limited demonstration or deployment. Meanwhile, some of the systems are starting to be deployed on a large scale, specifically, AVI is becoming commonplace on tollways. If such infrastructure is available it should be considered for improved traffic monitoring, e.g., it is a simple extension to use AVI tags and readers deployed for tolling to also measure link conditions, as is already the case in many cities (e.g., Nishiuchi et al, 2006). Automated license plate recognition (ALPR), a variation of AVI, does away with a separate tag and uses image processing to extract information from the existing license plate[3] (e.g., NYSI&S, 1970; Tindall and Hodgson, 1994; Buisson, 2006). Because AVI assigns a unique identification to each vehicle in the form of a machine-readable license plate or identification tag it eliminates most of the ambiguity and makes vehicle reidentification trivial, though challenges remain (e.g., Robinson and Polak, 2006). Other problems arise with AVI beyond the technical issues, e.g., although AVI data can be used to calculate link conditions the AVI readers are not capable of measuring the equally valuable traffic conditions local to the detector stations. Furthermore, the lack of ambiguity also raises privacy issues that have yet to be fully addressed in the policy arena. In any event, it is not yet clear that the ability to measure link conditions could justify a substantial investment in any of the new detector technologies. In the mean time loop detectors remain the preeminent vehicle detectors, and just like the newer technologies, if such infrastructure is available it should be considered for improved traffic monitoring. Based on experience prototyping related algorithms in real time on the loop detector stations in the BHL, it is estimated that the cost to retrofit loop detector stations could be as low as one percent of the original cost to deploy the detector stations (May et al, 2004). This cost aspect makes this study very relevant as the benefits of this study are two fold: first, it will allow further evaluation of the value of information obtained from vehicle reidentification without major upfront investments, and second, if successful, it may prove to be an inexpensive means of obtaining these data.

Section 2 presents the analysis, starting with the processing used to measure vehicle lengths and a discussion of the distinct vehicles that are easier to identify in the vehicle fleet. Then, the

[2] Such inductive signatures were originally developed for vehicle classification. Although the sensor cards are compatible with existing loops, they require new hardware in the field to process and relay the signatures, and they also need a much wider communication bandwidth compared to the transition data used in the present study.

[3] In the case of ALPR, errors remain in sensing the characters and the license plates are not always read correctly, e.g., confusion may arise between the letter "B" and numeral "8". While in the examples this paper uses a physical length measured from loop detectors to differentiate between vehicles, there is no reason why other distance metrics could not be used, e.g., the algorithm could be applied to ALPR data using some distance metric that specifies "B" and "8" are closer to one another than either one is to the letter "T".

section continues with a presentation of the reidentification methodology. Results are presented in Section 3 in detail from several freeway segments, including dual and single loop deployments. Finally, Section 4 closes the paper with the conclusions of the research.

2 ANALYSIS

2.1 Data from Loop Detectors

In conventional practice data from several passing vehicles are aggregated together at a loop detector station to calculate flow, occupancy and so forth. While such aggregation is useful for many applications, the practice also discards information about the individual vehicles. Rather than aggregating vehicle measurements together, we measure individual vehicle lengths. Figure 1A shows a time-space diagram for a vehicle passing over a dual loop detector. The dual loop detector has two detection zones as indicated in the figure. For each passing vehicle the controller measures the dual-loop traversal times using the rising edges TT_r, dual-loop traversal time using the falling edges TT_f, the on-time for the first loop OT_1 and the on-time for the second loop OT_2, as shown in the figure.

For vehicles passing over a dual loop, the length of the vehicle is proportional to the on-time and inversely proportional to the traversal time. The length measurements are subject to resolution constraints of OT_i and TT_i, which in turn depend on the vehicle velocity, the sampling rate of the detectors and separation between the loops. Since the reidentification algorithm explicitly compares measurements of each vehicle between stations, it must accommodate these measurement limitations. To this end, a length range that has sufficient tolerance is developed. Since the algorithm seeks to match the length measurement of a given vehicle detected at the downstream station to the same vehicle detected at the upstream station, the length range aims to increase the probability of intersection between the length ranges of the 'true' upstream and downstream matches while minimizing the number of false positives as much as possible when the measurements do not come from the same vehicle. The following length estimates were used to define the length range.

$$\text{Length estimate \#1: } L1 = S*OT_1/TT_r \quad (1)$$

$$\text{Length estimate \#2: } L2 = S*OT_2/TT_f \quad (2)$$

$$\text{Length range} = [\text{Min}(L1,L2), \text{Max}(L1,L2)] \quad (3)$$

where S is the loop separation in feet.

2.2 Importance of Long Vehicles

In most cases long vehicles constitute a small percentage of the vehicle fleet that travels through an urban freeway. To illustrate this point, Figure 1B-C show the cumulative distribution of measured vehicle lengths observed over 24 hours at one freeway detector station. It can be seen that fewer than 10 percent of the observations range from 23 ft to 80 ft. The long vehicles thus are a 'distinct' category of vehicles in these data.[4] Because the long vehicles encompass a large

[4] Of course exceptions exist where long vehicles make up a large portion of the fleet traveling on urban freeways, but unlike passenger vehicles it remains possible to differentiate between a 40 ft vehicle and a 60 ft vehicle. The present work is constrained by the available data, it is conceivable that some locations may require a more specific vehicle feature to characterize the 'distinct' vehicles (e.g., vehicles between 30 and 50 ft), or even an inversion (e.g., vehicles below 25 ft) before applying this work.

range of feasible lengths and they normally pass with a lower frequency, it makes them easier to differentiate from one another than it is for most vehicles. For a given link, this paper uses vehicles that are longer than a threshold based on the 90th percentile length for all the vehicles passing the downstream loop detector over 24 hours. In all of the analyses that follow, only long vehicles are considered unless otherwise noted.

In each lane at a downstream station, the long vehicles are subsampled out from the rest of the fleet using the following threshold,

$$d_l > D_{threshold} \tag{4}$$

where, d_l is the lower bound of the length range of the vehicle measured at downstream station from Equation 3, and $D_{threshold}$ is the pre-set minimum threshold length of long vehicles. The long vehicles are then numbered consecutively by lane, based on their arrival order.

2.3 Comparison of Length Ranges and Travel Time Representation

Each long vehicle at the downstream station is considered the primary vehicle with a set of candidate vehicles that are feasible matches at the upstream station. This choice of terminology reflects the fact that as soon as a vehicle passes the downstream station we know that if there is a match at the upstream station it has already been observed. The upstream vehicle candidates include all vehicles (no length threshold) and are chosen using the two following rules: first, to ensure positive travel time a candidate must arrive at the upstream station before the arrival of the primary vehicle at the downstream station, and second, the total number of candidates shall not exceed the jam density of link, i.e., the storage capacity, n.

The length range of the primary vehicle is compared with candidate vehicles present within the n most recent upstream station arrivals. The length range from Equation 3 captures measurement uncertainty, so the length range upstream should intersect the length range downstream if the two measurements came from the same vehicle. Of course many other candidate vehicles will have a length range that intersects that of the primary vehicle. Formalizing this concept, denoting the upper and lower bounds from Equation 3 for the primary vehicle as d_h and d_l, and similarly the bounds for each candidate vehicle as u_h and u_l. a primary vehicle candidate pair is considered a possible match if,

$$u_h \geq d_l \text{ and } d_h \geq u_l \tag{5}$$

The possible matches for each primary vehicle are then stored in matrix format indexed by travel time, where the travel time for each matched pair is obtained by subtracting the arrival time of the candidate match at the upstream station from the arrival time of the given primary vehicle at the downstream station. The rows of this Travel Time Matrix (TTM) correspond to the primary vehicle number. The indices of the columns of the matrix represent the possible travel times rounded to the nearest integer second. The TTM is populated with 0's except for the travel times for each of the given primary vehicle's possible matches, which are given a value of 1. To avoid a very large column size and thereby improve computational efficiency, the width of TTM is further constrained to the travel times corresponding to link velocities falling between 2 mph and 90 mph[5]. Figure 2A shows an empirically observed TTM from dual loop detector stations almost

[5] Throughout this description "link velocity" is used to define parameters rather than "link travel time" because the former explicitly accounts for the length of the link.

a mile apart, in this graphical representation the possible matches for the candidate vehicles (the 1's) are shown with solid circles and all other cells (the 0's) are left blank.

2.4 Extracting Information from the Travel Time Matrix

Each row of the TTM typically has many false positives, as evident in Figure 2A, while any given row can have at most, one true match. The false positives are randomly distributed within each row, yielding a low density of possible matches throughout the entire matrix when considering several rows. For the rows in which the true matches exist, the true matches from consecutive vehicles fall in a small range of columns and increase the density of possible matches above the background level of the false positives. This higher density of possible matches in the matrix provides a means to identify the travel time trends and the true matches provided enough vehicles are observed before the link travel time changes. These trends extend from free flow to moderately congested periods but are not clearly discernible during heavily congested periods due to the increased range of the travel times. For example, under these extreme conditions if link velocity drops from 10 mph to 5 mph, the link travel time will double even though it is only a small absolute change in link velocity. Extracting the trends from the TTM reduces the number of possible matches and point to the most probable matches. To this end the algorithm tries to identify the dense areas of the TTM that would capture the trend exhibited in the TTM. As discussed below, the TTM matrix is transformed to the maximum density matrix (MDM), to identify the dense areas. In the first step, to accommodate the fact that successive vehicles may experience slightly different travel times, the possible travel times are extended horizontally to adjacent cells using the following equation,

$$f\left(TTM(K,T)\right) = \begin{cases} 1, & \sum_{H=T-p}^{T+p} TTM(K,H) > 0 \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

where,

$$p = \frac{1}{2}\left(TT_{ffl} - TT_{ffu}\right)$$

and TT_{ffl} and TT_{ffu} are the lower and upper bound on assumed free flow travel times (set to correspond to link velocities of 65 mph and 45 mph, respectively).[6] So this function f() gives a weight of 1 to a cell if any of the adjacent 'p' cells in the same row of TTM had a value of 1. Next, the MDM is generated recursively using the following equation,

$$MDM(K,T) = f\left(TTM(K,T)\right) + \sum_{R=K-1}^{K-q} MDM(R,T) \tag{7}$$

where, K corresponds to primary vehicle number at the downstream station, T corresponds to link travel time yielded by a potential match of vehicle K. The second term of Equation 7 is the sum of the cells in the same column of the previous q rows, so a cell in the MDM gets more weight if preceding vehicles also had the same travel time. The value of 'q' is set to be the number of long vehicles that arrived at the downstream station within the previous 5 minutes with the maximum allowed value of q set at 25. These constraints make sure that vehicles that

[6] Thus, the resulting time window, T-p to T+p, used in the equation corresponds to the range of free flow travel times.

passed the downstream station a long time before the current vehicle do not influence the current extraction of the vehicle's link velocity. Unlike the TTM, which has a value of '1' for cells that correspond to a primary vehicle number-travel time pair, the MDM will exceed 1 for the densest areas because of the summation in Equation 7. Figure 2B shows an example of the MDM superimposed on the original TTM from Figure 2A.

The MDM is then subsampled to pull out a new matrix, MDM_{20}. First, the column corresponding to the maximum value in each row of the MDM is found. These cells are given a value of 1 in MDM_{20} while all other cells in that row of MDM_{20} are assigned a value of zero. MDM_{20} is then truncated so that all columns with a travel times corresponding to link velocity slower than 20 mph are unilaterally set to zero to exclude these slower link velocities where small changes in velocity translate to large changes in travel time.[7] When the link velocity is above 20 mph, the densest cells of each row in the TTM should normally fall within the non-zero range of MDM_{20}. The travel times corresponding to these cells are henceforth called the most probable travel times (MPTT). The final non-unique set of most probable matches is obtained by finding the vehicles that are responsible for these travel times. Since one is not guaranteed an intersection between the densest areas of MDM and TTM, for every primary vehicle the MPTT are compared with the travel times resulting from the initial set of primary vehicle-candidate pairs. If there is a single intersection, that pair is then considered to be a most probable match and stored. It is possible that sometimes there might not be an intersection, in which case any pair whose absolute travel time difference with the MPTT for that primary vehicle is not greater than 1/8th of the travel time difference between upper bound and lower bound of assumed free flow velocity range[8] is considered a most probable match and stored, i.e., if

$$\left|MPTT - TT_x\right| \le \frac{1}{8}\left(TT_{ffl} - TT_{ffu}\right) \tag{8}$$

where TT_{ffl} and TT_{ffu} are defined above, and TT_x is a candidate travel time. When Equation 8 is true, the pair of vehicle measurements corresponding to TT_x is then retained as a MPTT for that primary vehicle. This method for looking for matches increases the number of vehicles found by including those matches from vehicles that might be traveling slightly slower or faster than the densest areas of the TTM.

2.5 Unique Matches

In the final set of MPTT, it is possible for a primary vehicle to have more than one match and a candidate vehicle to be matched with more than one primary vehicle. To address the former issue, the matched vehicle pairs are subjected to a two-step filtering process that results in at most one unique match for every primary vehicle. The first step is to remove any potential match that has an improbably fast link velocity based on the posted speed limit, e.g., the filtering threshold is 80 mph for a link whose posted speed limit is 65 mph in order to accommodate drivers that travel faster than the posted speed limit. In the second step, for every primary vehicle with more than one candidate vehicle match, the match that yields the median of the possible link velocities from the set of matches is accepted arbitrarily if there is an odd number matches, or the one immediately below the median velocity is chosen if there is an even number of

[7] Thus, when the travel times correspond to link velocities below 20 mph, those rows of MDM_{20} will only contain zeros.

[8] The value of 1/8 was arrived at empirically for reducing the number of false positives that resulted from this approach.

matches. In either case, all the other entries for that primary vehicle are removed. Figure 2C shows the MPTT superimposed on the original TTM.

Because a candidate vehicle could later be matched to another primary vehicle, the entire set of primary vehicles for a candidate vehicle is not known in real time. As such, this paper does not test for duplicate matches for a given candidate. For applications that do not need real time conditions, the filtering process could easily be repeated for the candidate vehicles to ensure that no candidate is matched to more than one primary vehicle, but this filtering could only occur after the last feasible candidate is observed. Obviously, more sophisticated means of establishing unique matches are possible as well, e.g., that account for more than one primary vehicle at a time.

2.6 Calibration

As already noted, many of the parameters were established empirically. This study examined only two corridors because the necessary transition data, e.g., the TT's and OT's as shown in Figure 1A, are not normally transmitted outside of the controller cabinets. The transition data are readily available at most loop detector stations but these data are usually aggregated and then discarded. Although not shown, the two subject corridors exhibited stable performance to small changes in the parameters. Undoubtedly the vehicle fleet, driver population, or geometry of the roadway will likely impact the sensitivity of the algorithm to the empirically established parameters. So the parameters presented herein should be considered a starting point. An application of this work to a new corridor may require retuning the parameters, but this tuning should not be difficult, e.g., for Equation 8, one would collect a few days of data and then regenerate Figure 2C with different scale factors.

2.7 Multi Lane Vehicle Reidentification

The primary difference between the single lane case described above and the multi lane case is that in the multi lane case the first step is to combine the candidate vehicles from all relevant lanes and sorting the candidates sequentially by arrival time.[9] The vehicle lengths and length ranges are calculated as explained in the single lane case. It should be noted that the number of candidate vehicles upstream to be used for comparison with each primary vehicle downstream would be higher in the multi lane case because the storage density of the lanes that have been combined should be taken into consideration.

The travel time representation and extraction of information are the same as the single lane case. Unlike the single lane case, the multi lane case provides matches for long vehicles that changed lanes. At present the analysis remains limited to a single downstream lane because adjacent lanes can exhibit significant differences in travel time. If one uses multiple lanes at both stations the MDM could have two or more distinct dense areas in a given row that reflect true travel times, with small fluctuations in density causing the MPTT to jump between the dense areas over several successive rows. An obvious extension of this method would be to accommodate more than one MPTT, one for each downstream lane. Presently, however, multiple downstream lanes are addressed simply by applying the algorithm independently to each downstream lane and any

[9] One can easily store the lane in which each candidate passes as well as the order in which candidates from a given lane pass if desired for subsequent analysis, e.g., tallying the number of reidentified vehicles that changed lanes.

candidate vehicle that is chosen for more than one downstream lane can then be addressed, e.g., assigned to a single lane.

2.8 Extending to Single Loop Detectors

Unlike dual loop detectors, a single loop detector has a single detection zone, e.g., only the first detection zone in Figure 1A and thus can only report the first detector response in the figure. In this configuration it remains possible to measure vehicle on-time, but since there is no observed traversal time it becomes impossible to measure velocity. Velocity can only be estimated from single loop detectors and the conventional estimate is very noisy because the assumption of a fixed average vehicle length, L, that may or may not be representative of a given sample in the following equation,

$$\hat{v}_{old} = \frac{L \cdot flow}{occupancy} = \frac{L}{mean(on-time)} \tag{9}$$

The mean(on-time) for a given velocity will fluctuate relative to the assumed L as the percentage of long vehicles changes, and thus, so will the velocity estimate from Equation 9. Earlier work by our group has shown that processing the data from a single loop detector in unconventional ways can make the performance of single loop detector based velocity estimation approach the velocity measurement accuracy from a dual loop detector. A simple but important revision to Equation 9 to calculate the median rather than the mean,

$$\hat{v}_{new} = \frac{L}{median(on-time)} \tag{10}$$

greatly reduces the sensitivity to fluctuations in the composition of the vehicle fleet relative to the assumed L (Coifman et al, 2003). The estimated vehicle lengths, $\hat{L}$, are then obtained by multiplying the on-time of each vehicle by the estimated velocity from Equation 10.

$$\hat{L} = on-time \cdot \hat{v}_{new} \tag{11}$$

Coifman et al (2003) showed that when the estimated velocity of the sample was greater than 20 mph the average absolute error in estimated length was less than six percent for the vehicle set considered. The paper notes that the error in the length estimates is expected to increase in periods with heavy truck flows. Fortunately, heavy truck flows are not common on urban freeways during most periods when vehicle reidentification would be of greatest interest and hence the above formula for estimation of length from single loop detectors was used. If truck flows are of concern one could use more sophisticated velocity estimation algorithms, e.g., Neelisetty and Coifman (2004). In any event, since there is only one length estimate per vehicle, the length range from Equation 3 is replaced by the following:

$$\text{Length range} = [0.8^* \hat{L},\ 1.2^* \hat{L}] \tag{12}$$

The values 0.8 and 1.2 are used to give the lengths a reasonable tolerance and thereby increase the accuracy and percentage of vehicle reidentification. The 20 percent threshold was arrived at by trial and error. For an estimated length of 20 feet, the above definition would give a tolerance of 8 feet, while an estimated length of 80 feet it would yield a tolerance of 32 feet. In all of the subsequent analysis, Equation 3 was used for the dual loop length range and Equation 12 was used for single loops. The rest of the processing was identical under the two conditions.

3 RESULTS

Several test sites are examined below and in each case between 10 and 40 days of data were studied at a given site. While the traffic patterns changed from day to day, the relationships between the link velocity and local velocities in these examples are typical those observed at the given location.

3.1 Dual Loop Detectors

The multi lane vehicle reidentification algorithm was implemented over different sections of the Berkeley Highway Laboratory (BHL) (Coifman et al, 2000). Figure 3 shows a schematic of a 0.91 mi section of westbound I-80 between station 2 and station 5. All lanes have dual loop detectors and there are two intervening stations that were not used in the vehicle reidentification. Figure 4 shows an example of the results of the vehicle reidentification algorithm on the section using data from July 15, 2003. Vehicles from lane 3 of downstream station 5 were matched with vehicles from lanes 2, 3 and 4 of upstream station 2. The figure shows the resulting link velocities for the 1,277 matched vehicles superimposed on the time series local velocity measurements at each of the detector stations. Note that there are few matches between 17:00 and 19:00 because the subsampling of MDM_{20} will exclude true link velocities below 20 mph and during this time window both ends of the link are below the threshold velocity.

The BHL includes several ramps, but it does not include any major freeway merges or diverges. Moving to the I-71 freeway corridor in Columbus, OH, Figure 5 shows a schematic of a 0.97 mi section of southbound I-71 between stations 108 and 103. As before, lanes are numbered from left to right relative to the flow of traffic, starting with lane 1 on the left. The section has two merges and two diverges including the I-71/I-70 interchange. Both stations 108 and 103 in the southbound direction have dual loop detectors. Intervening station 106 shown in the schematic has single loop detectors and was not used in the reidentification. Figure 6 shows an example of the results of the vehicle reidentification algorithm on the section using data from May 29, 2003. Vehicles in lane 2 of downstream station 103 were matched with vehicles from lanes 2 and 3 of upstream station 108. The link velocities of the 1,548 matches obtained through the vehicle reidentification algorithm over 21 hrs are shown in the figure, superimposed on the time series local velocity measurements at each of the detector stations. In fact Figure 2 shows a portion of the TTM, MDM and MPTT from these stations on this day. Figure 6 also shows the measured travel times from two probe vehicle trips through the corridor, as measured by an on-board global positioning system (GPS) receiver. During the afternoon peak period the tail of the queue fell between the two stations used for reidentification, as evidenced by the fact that the upstream station had free flow conditions during most of this period while the downstream station was congested.

If one considers the evolution of the traffic state (in particular velocity) over time and space, the local measurements provide independent slices through the time-space plane at the respective fixed locations. First looking strictly at these local time series from the detector stations, at the end of the day the evolution of the traffic state is evident. Whether free flow or congested, most of the time the local velocity time series is fairly stable, but there are a few major changes in velocity and these are the features that an operating agency wants to detect. For example, an incident or recurring bottleneck between the detector stations can cause link velocity to drop as a queue forms. If the delays are significant, the queue will eventually reach the upstream station though it may be some time between when a queue forms and when it is first observable locally

at a detector station. Any queue longer than the distance between detectors will eventually be seen locally at one or more stations. Queue growth from a bottleneck downstream of the link is evident in the afternoon peak of both Figure 4 and 6, with the downstream station becoming congested followed the upstream station, similarly, queue decay is also evident as demand wanes and the queue discharges with the upstream station first clearing then the downstream station. But this big picture is not available from the local stations in real time; it is obscured both by normal traffic fluctuations and detector errors, and for bottlenecks between stations the delay between when a queue forms and when it is first observable locally at a detector station.

The link velocities from the matched vehicles provide a separate set of slices through time-space of the evolving traffic state. Unlike the fixed detectors that only average over time, vehicle reidentification also averages over space, along the trajectory of the matched vehicle. While the complete picture from the local velocity shown in Figures 4 and 6 is only available after the fact, it provides an independent validation for the over 1,000 reidentified vehicles in each of the figures. Figures 4 and 6 both have three periods where traffic is free flowing for several hours while the morning and evening peaks are congested. Except for a small number of stray measurements evident in Figures 4 at 18:00 and 6 just before 24:00, the link velocity from the matched vehicles follow the general trend evident in the local velocity at each of the detector stations. During the off peak periods the reidentification algorithm found free flow link velocities, consistent with the trends revealed in the local conditions. As one reviews these figures it is important to remember that the link velocity from the matched vehicles is measured independent of the local velocities at the detector stations. The reidentification algorithm searches the same time window, regardless of traffic conditions. So even when traffic is free flowing and the true matches should fall on the left-hand side of the TTM, the algorithm still searches the entire TTM. Moving to the congested periods, upon closer inspection of Figure 4, one can see the impact of a downstream queue sweeping through the link around 14:00 and receding around 19:00. During these transitions the matched vehicle velocities pull apart from the local velocity measured at either of the stations used for reidentification, falling in between the two local velocities and reflecting the fact that a queue covers only part of the link. In this case, during these transition periods the link velocity is much closer to the local velocity at an intervening station not used in reidentification. Figure 7A repeats this comparison against the intervening station, omitting for clarity the local velocity from the two stations used in reidentification. In other words, the reidentification provides a better representation of conditions on the link than is available locally from the detector stations used for reidentification. Likewise, Figure 7B provides a similar comparison against an intervening station for the data from Figure 6.[10] In this case, during the afternoon peak period the intervening station's velocities often fall below the MDM_{20} threshold velocity.

Consistency with local conditions at the detector stations during stable traffic is a good feature, but the real benefit of the reidentification algorithm comes during unstable traffic by detecting delays before a queue reaches a detector station. Figure 8 shows a schematic of a 0.95 mi section of northbound I-71 between stations 105 and 109. The section crosses the I-71/I-70 interchange in the opposite direction from the previous example and includes a second major diverge on the

[10] Note that the estimated velocities from the intervening single loop detector station are noisy before 5:00, showing a few seemingly unexpected slow samples. This fact arose due to normal early morning low flow conditions and can be addressed with an occupancy filter (Jain and Coifman, 2005)

north end, I-71/I-670. Lanes are still numbered from left to right relative to the flow of traffic. Both stations 105 and 109 in the northbound direction have dual loop detectors. Intervening stations 106 and 107 shown in the schematic have single loop detectors and were not used in the reidentification. Figure 9 shows an example of the results of the vehicle reidentification algorithm on the section using data from February 28, 2002. Vehicles in lane 2 of downstream station 109 were matched with vehicles from lanes 1 and 2 of upstream station 105.[11] In this case there were 1,809 matches and four probe vehicle travel time measurements through the corridor were recorded throughout the day. The most significant feature of this plot is the fact that the vehicle reidentification responded to delays between 7:00 and 9:00 that were not observable at either station used in the reidentification (105 or 109). This delay is evident at the intermediate stations, e.g., station 107 shown in the plot, and one of the probe vehicle runs. As with the two previous examples, Figure 7C repeats this comparison against the intervening station, omitting for clarity the local velocity from the two stations used in reidentification. In the afternoon the stations at either end of the link eventually become congested, but here the vehicle reidentification algorithm responded almost an hour before the delay was observable at either of the stations (once more the delay is evident at the intermediate station not used in reidentification). So the reidentification can provide better fidelity in the data from existing detectors or provide a comparable fidelity with a lower density of detector stations. Table 1 presents the reidentification rates from these three examples.

3.2 Single Loop Detectors

The previous examples used the dual loop data, but not all loop detector equipped freeways have dual loop detectors. If anything, single loop detectors are more common than dual loop detectors. To emulate facilities equipped with single loop detectors, this section limits the input to a single loop from each of the dual loop detectors and this section includes a single loop detector station that was not previously used in the reidentification. Using this single loop data Figure 10 shows the results of reidentification between stations 105 and 109 on the same day and section shown with dual loops in Figures 8-9. So this new figure shows the resulting travel time when matching observations from the single loop detector data in lane 2 of station 109 with single loop detector data in lanes 1 and 2 of station 105. It can be seen that the algorithm works quite well using the single loop detector data, although the number of matches has decreased slightly to 1,284. Again, the reidentification algorithm found the delay between the stations long before it was locally evident at either end of the link. One difference however was that a higher threshold of 25 ft was used for defining long vehicles. Note that in this figure the estimated velocities from the single loop detectors are somewhat noisy before 5:00, this fact arose due to normal early morning low flow conditions and can be addressed with an occupancy filter (Jain and Coifman, 2005).

The third row of Table 1 includes the reidentification rates for this link when emulating single loop detectors. The reidentification percentage from the single loop detector data is less than the dual loop detector data for the 0.95 mi link (I-71 Northbound 105 to 109). The final row of the table shows the reidentification rates when the vehicle reidentification algorithm is applied to a sub-segment of the same link on the same day. The algorithm was applied to the single loop detector data from lane 2 between stations 106 and station 107 in the I-71 northbound corridor, shown in Figure 8. The distance between the two stations is about 0.39 mi. Performance was

[11] Note from the schematic in Figure 8 that lane 2 at the downstream station 109 originates from I-70 westbound, and does not pass the upstream station 105. So any matched vehicle will have had to make at least one lane change maneuver.

similar to Figure 10 and for brevity, not shown. The percentage of reidentified vehicles from the single loop detectors increases when the link distance is smaller, as evident in the table. This improvement arises for three reasons: first, the shorter link distance reduces the dispersion between successive true travel times in the TTM. Second, the shorter link narrows the search time windows in Equations 6 and 8, reducing the impact of false positives. Third, as evident in Figure 8, the shorter link excludes two major diverges that are present in the longer link. So a greater percentage of the long vehicles pass both stations without diverting. Investigating the specific impacts of various features such as these, as well as traffic conditions and other factors, is left to future work.

4 CONCLUSIONS

Conventional vehicle detectors are capable of monitoring discrete points along the roadway but do not provide information about conditions on the link between detectors. This paper developed a robust vehicle reidentification algorithm to quantify conditions between detector stations. The distinct vehicles were key to this reidentification, namely the long vehicles in this case.

The reidentification results were used to find link travel times until link velocity dropped below 20 mph. About 40 percent of long vehicles were consistently reidentified over links on the order of 1 mi long using dual loop detector data, and a slightly lower reidentification rate when using single loop detector data. The travel time measurements obtained from the reidentified vehicles will be useful for studying the traffic conditions on the freeways. This work advances our earlier efforts significantly. First, the ability to match vehicles between single loop detector stations opens up additional freeways for analysis. But the most important contribution of the new algorithm is the fact that the present research does not limit vehicle reidentification to only one lane, which allows vehicle reidentification even when the candidate vehicle changes lanes. The improved detection allows for reidentification across major merges or diverges in the freeway segment, where one cannot assume that most vehicles travel along the same lane. Two out of the three links examined included a major merge and at least one major diverge of two interstate freeways. While earlier research has examined loop detector based vehicle reidentification across ramps, to our knowledge, this paper presents the first published example using loop detector data to reidentify vehicles across a major merge or diverge. The examples include a case where the reidentification algorithm responded to delay between two detector stations an hour before the delay was locally observable at either of the stations used for reidentification and another case where the reidentification algorithm found a delay observable an intermediate station, but the delay was never observable locally at the stations used for reidentification.

The new information from vehicle reidentification promises to be useful to operating agencies for making timely decisions in response to various delay causing events and hence to reduce the resulting congestion on the freeway system. Agencies should seek to extract as much information as possible from the existing surveillance infrastructure. Because many cities already have an extensive network of loop detectors compatible with this algorithm, the cost to deploy this work to measure link conditions will often be far lower than most of the alternatives for measuring link conditions. Although the individual vehicle data used for reidentification typically exist in the traffic controllers, the research has been limited by the fact that these data are usually aggregated at the controller and are not available for research, so the algorithm is potentially subject to further calibration as enumerated herein. Finally, although the research uses loop detector data, the algorithm would be equally applicable to data obtained from many other traffic detectors that

provide reproducible vehicle features such as vehicle length, height, class, color, or even partial characters read from the license plate.

5 ACKNOWLEDGEMENTS

The authors would like to acknowledge the valuable feedback from the anonymous reviewers, their constructive input has helped make this a better paper.

This material is based upon work supported in part by 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. 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.

6 REFERENCES

Buisson, C. (2006). Simple Traffic Model for a Simple Problem: Sizing Travel Time Measurement Devices, *Transportation Research Record 1965*, TRB, pp 210-218.

Coifman, B., Lyddy, D., 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., Cassidy, M. (2002). Vehicle Reidentification and Travel Time Measurement on Congested Freeways, *Transportation Research: Part A*, 36(10), pp. 899-917.

Coifman, B. (2003). Identifying the Onset of Congestion Rapidly with existing Traffic Detectors, *Transportation Research Part A*, 37(3), pp 277-291

Coifman, B., Dhoorjaty, S., Lee, Z., (2003). Estimating Median Velocity instead of Mean Velocity at Single Loop Detectors, *Transportation Research Part C*, 11(3-4), pp 211-222

Davies, P, Hill, C., Emmott, N., (1989). Automatic Vehicle Identification to Support Driver Information Systems, *Proc. of IEEE Vehicle Navigation and Information Systems Conference*, 11-13 Sep., pp A31-35.

Dion, F., Rakha, H.. (2006). Estimating Dynamic Roadway Travel Times Using Automatic Vehicle Identification Data for Low Sampling Rates, *Transportation Research. Part B*, 40(9), pp 745-766.

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

Itoh, T., (1986). Navigation System Using GPS for Vehicles, *SAE transactions*, pp 5.236-5.248.

Jain, M., Coifman, B., (2005). Improved Speed Estimates from Freeway Traffic Detectors, *ASCE Journal of Transportation Engineering*, Vol 131, No 7, pp483-495.

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

Kwon, T. M., (2006). Blind Deconvolution Processing of Loop Inductance Signals for Vehicle Reidentification, *Proc. of the 85th Annual Meeting of the Transportation Research Board*.

MacCarley, A.C., (2001). *Video-based Vehicle Signature Analysis and Tracking System Phase 2: Algorithm Development and Preliminary Testing*, California PATH Working Paper, UCB-ITS-PWP-2001-10

May, A., Coifman, B., Cayford, R., Merritt, G., (2004). *Automatic Diagnostics of Loop Detectors and the Data Collection System in the Berkeley Highway Lab*, California PATH Research Report, UCB-ITS-PRR-2004-13.

Neelisetty, S., Coifman, B., (2004). Improved Single Loop Velocity Estimation in the Presence of Heavy Truck Traffic, *Proc. of the 83rd Annual Meeting of the Transportation Research Board*.

Nishiuchi, H., Nakamura, K., Bajwa, S., Chung, E., Kuwahara, M., (2006). Evaluation of Travel Time and OD Variation on the Tokyo Metropolitan Expressway Using ETC Data, *Research into Practice: 22nd ARRB Conference Proceedings information*, Australian Road Research Board.

NYSI&IS, (1970). *Automatic License Plate Scanning (ALPS) System-Final Report*, New York State Identification & Intelligence System, Albany, NY.

Robinson, S., Polak, J., (2006). Overtaking Rule Method for the Cleaning of Matched License-Plate Data, *Journal of Transportation Engineering*, Vol. 132 No. 8, ASCE, pp 609-617.

Schafer, R., Thiessenhusen, K., Wagner, P., (2002). A Traffic Information System by Means of Real-Time Floating-Car Data, *Proc. 9th World Congress on Intelligent Transport Systems information*, ITS America.

Tindall, D., Hodgson, R. (1994). Evaluation and Trials of an Automatic License Plate Recognition System Employing Neural Network Techniques, *Moving Toward Deployment. Proc. of the IVHS America Annual Meeting.* Vol 1, pp. 329-334.

Turner, S, (1995). Advanced Techniques for Travel Time Data Collection, *Proc. of the 6th Vehicle Navigation and Information Systems Conference,* IEEE, pp 40-47.

Yim, Y, Cayford, R. (2006). Field Operational Test Using Anonymous Cell Phone Tracking for Generating Traffic Information, *Proc. of the 85th Annual Meeting of the Transportation Research Board*.

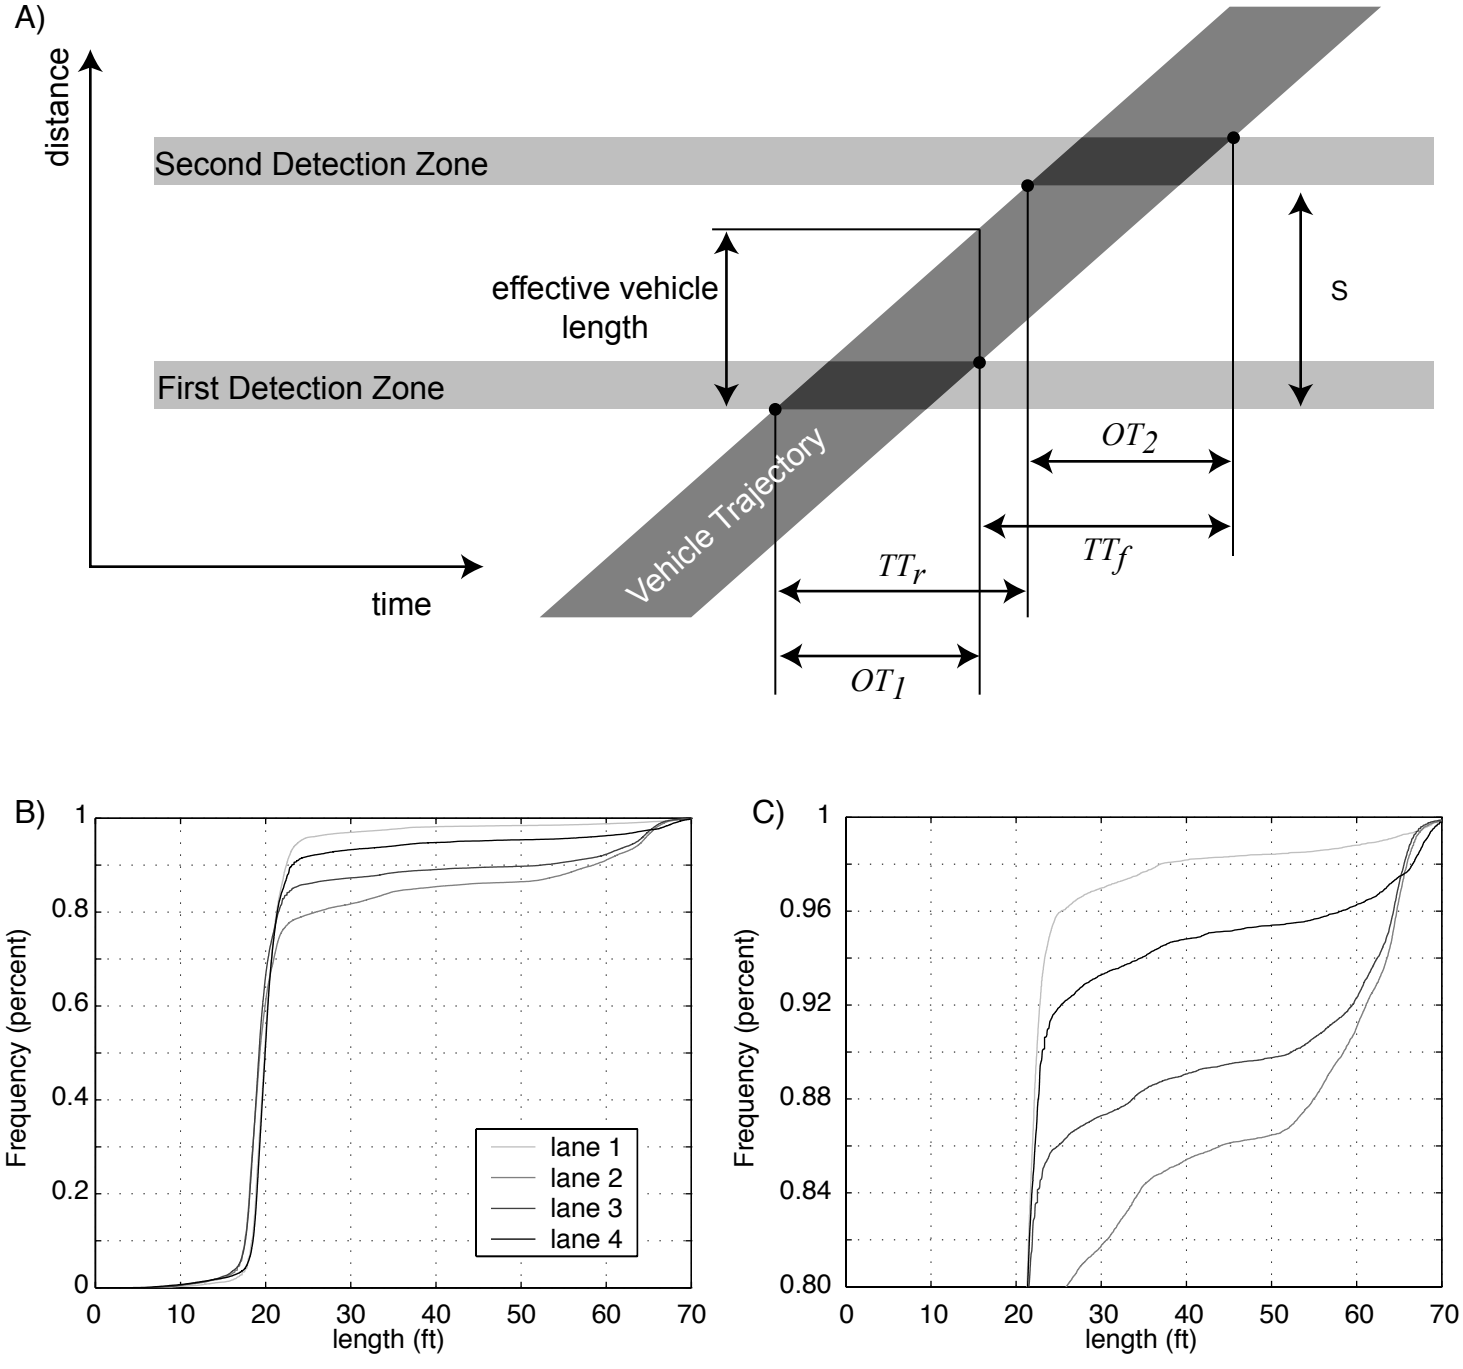

Figure 1, (A) One vehicle passing over a dual-loop-detector showing the two detection zones and the vehicle trajectory in the time space plane and the resulting measurements from the detectors, (B) Cumulative Distribution of lengths over 24 hours from one freeway detector station on I71, (C) Detail of Part B

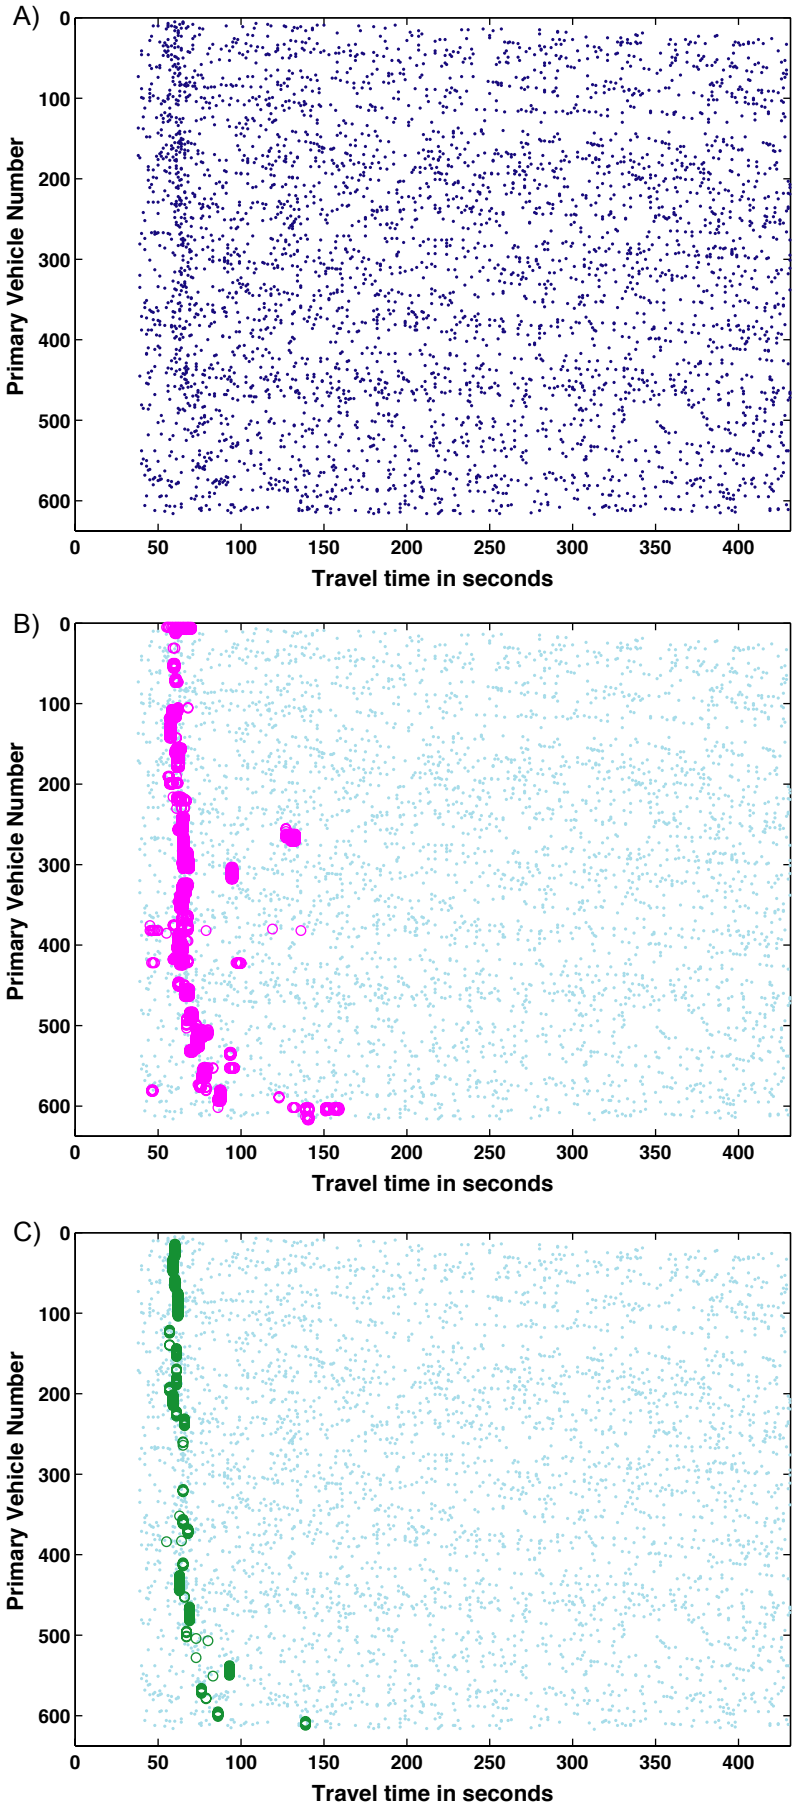

Figure 2, (A) Original TTM for a sample data set from two stations almost one mile apart, (B) MDM superimposed on the TTM from A, (C) MPTT after finding the unique matches superimposed on the same TTM.

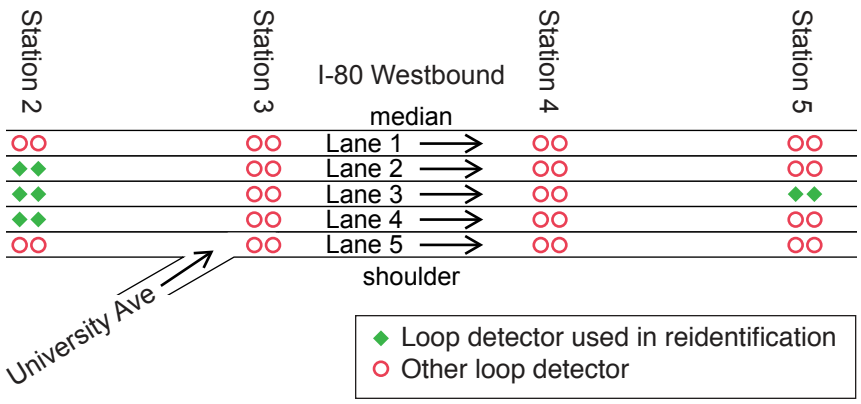

Figure 3,□ Schematic of the section between station 2 and station 5 in the westbound I-80 corridor (not to scale).

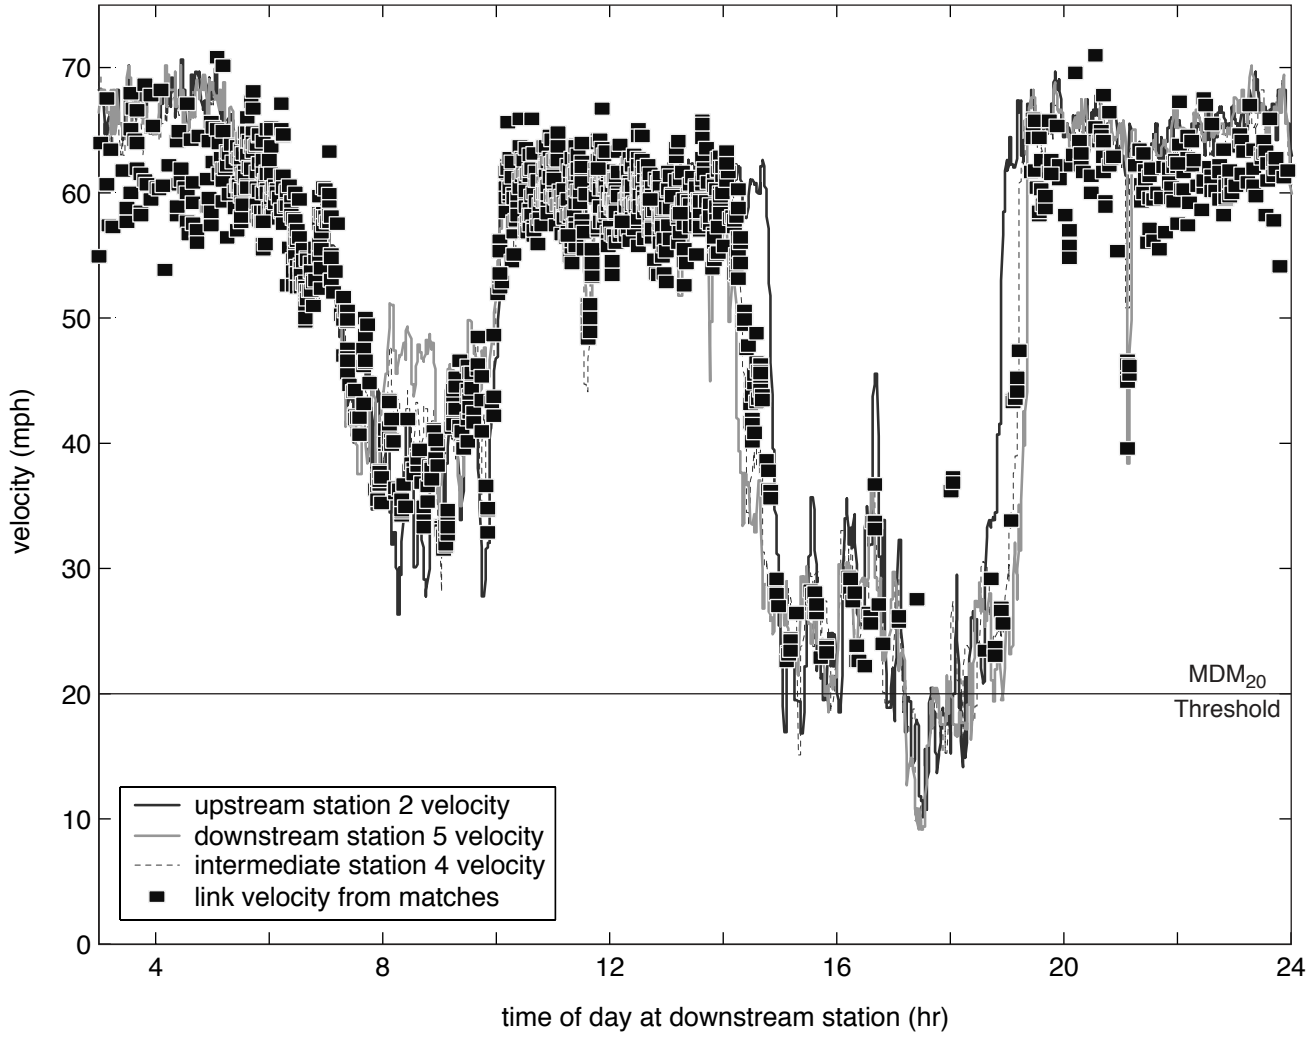

Figure 4,□ Results of vehicle reidentification between lane 3 at station 5 and lane 2, 3 and 4 at station 2 on the westbound I-80 corridor.

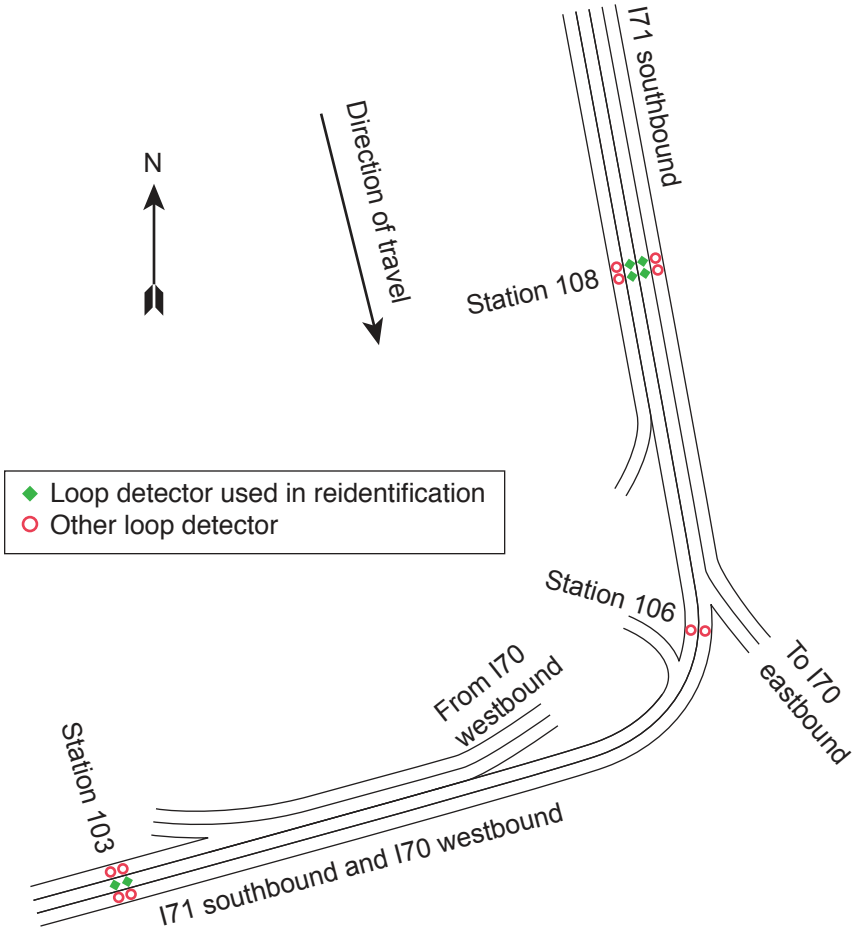

Figure 5,□ Schematic of the section between station 108 and station 103 in the southbound I-71 corridor (not to scale).

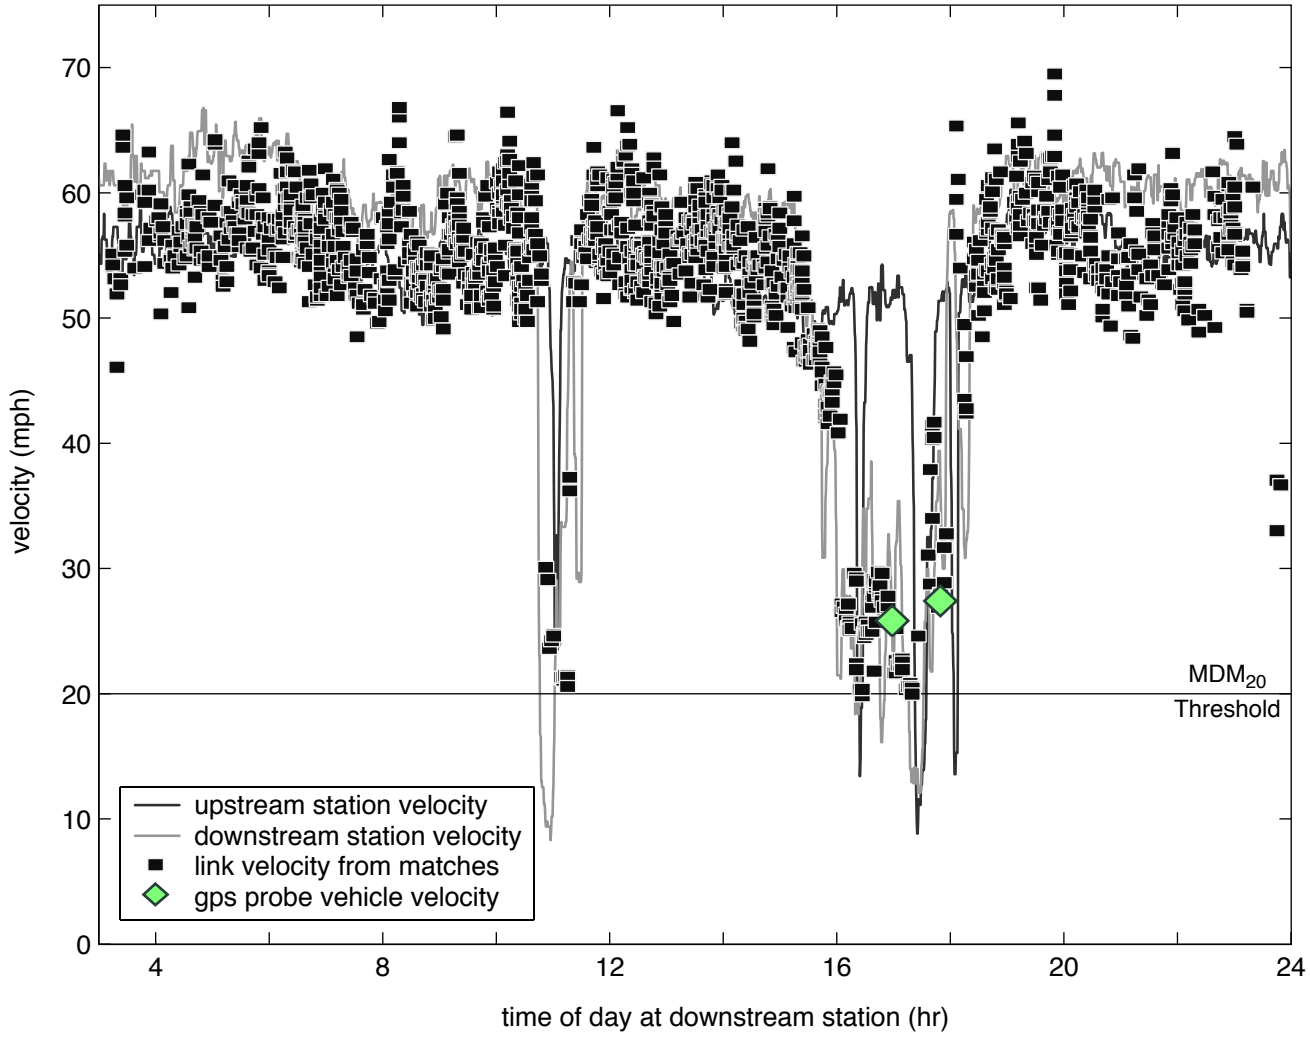

Figure 6, Results of vehicle reidentification between lane 2 at station 103 and lane 2 and 3 at station 108 on the southbound I-71 corridor.

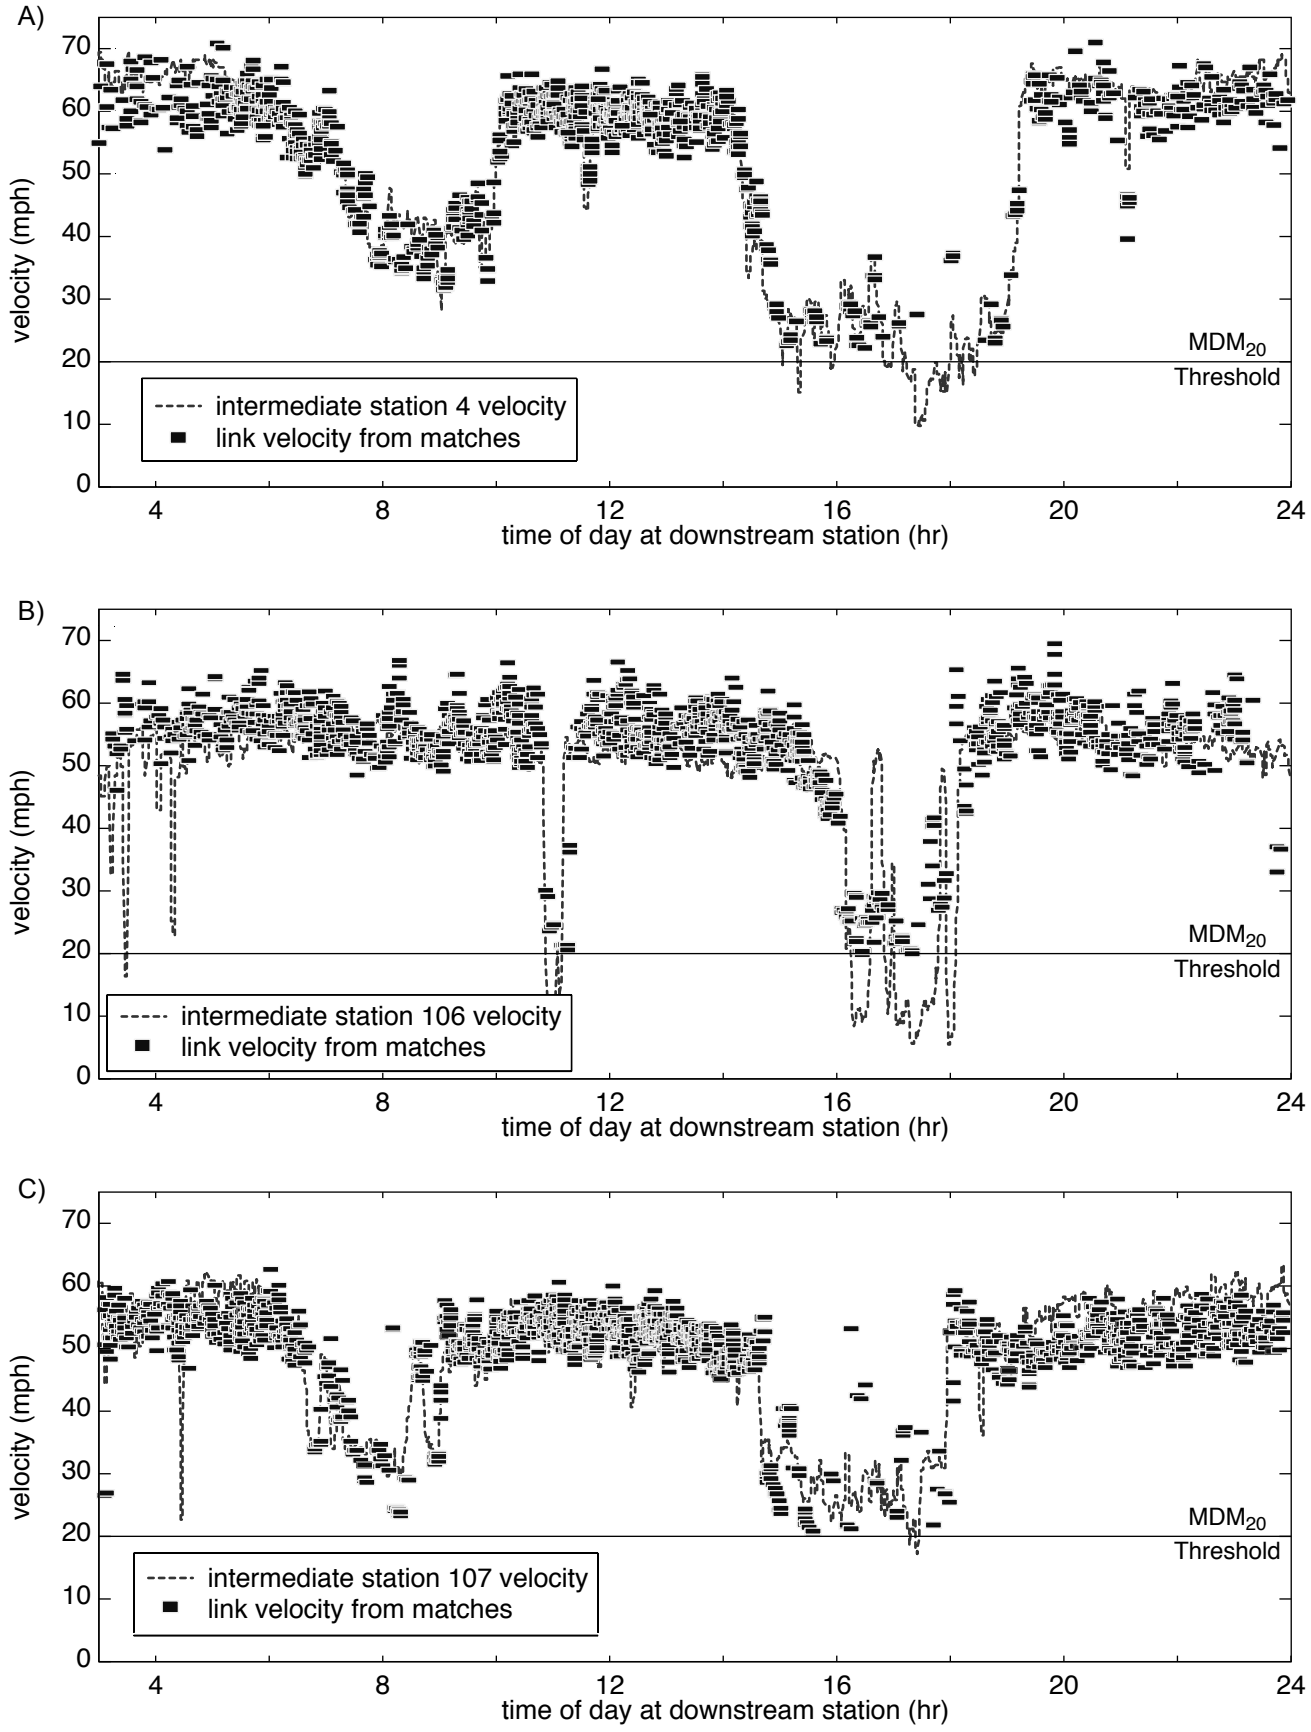

Figure 7, Results of vehicle reidentification compared against an intermediate station not used in the matching (A) between station 5 and station 2 westbound I-80, (B) between station 103 and station 108 southbound I-71, and (C) between station 109 and station 105 northbound I-71.

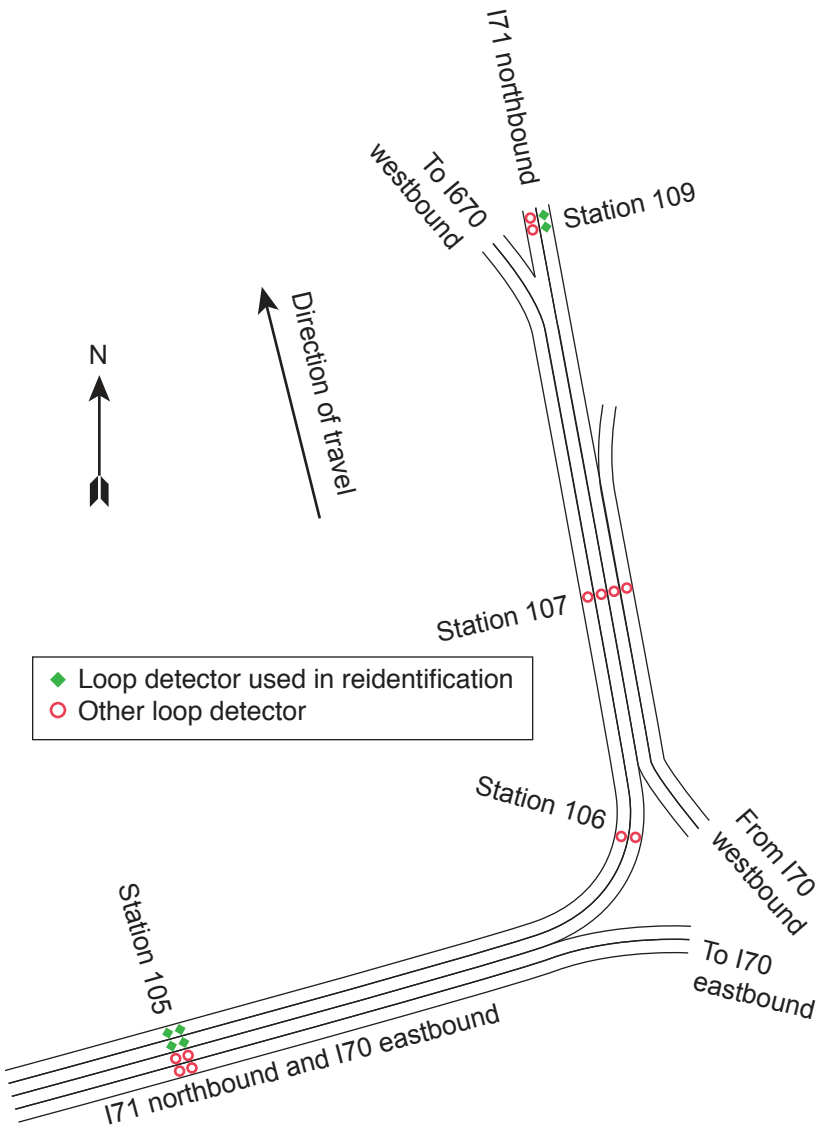

Figure 8,□ Schematic of the section between station 105 and station 109 in the northbound I-71 corridor (not to scale).

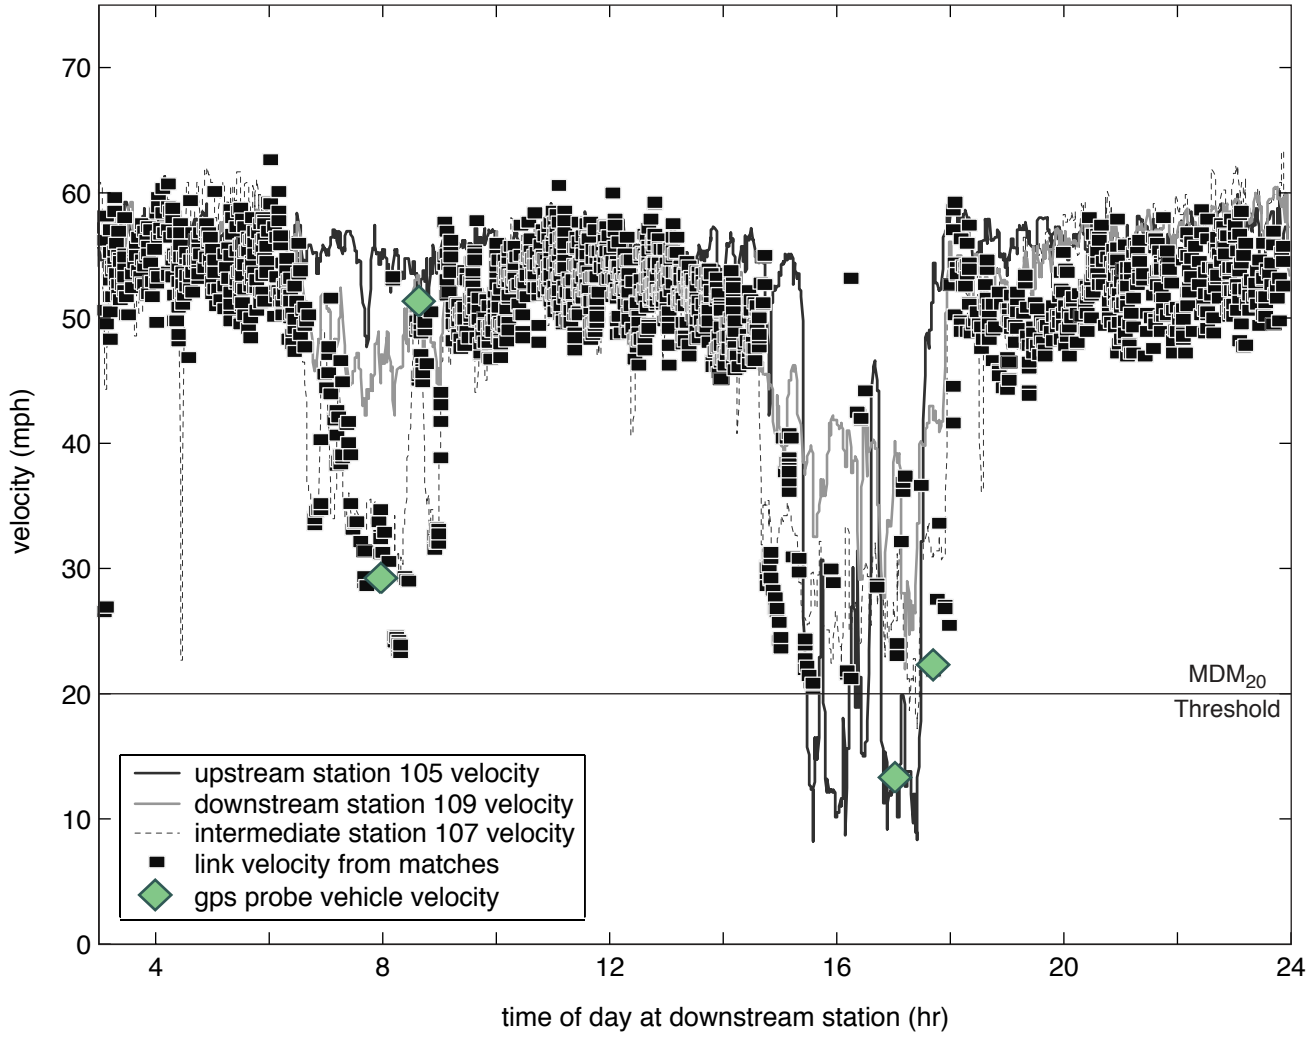

Figure 9, Results of vehicle reidentification between lane 2 at station 109 and lane 1 and 2 at station 105 on the northbound I-71 corridor using dual loop data.

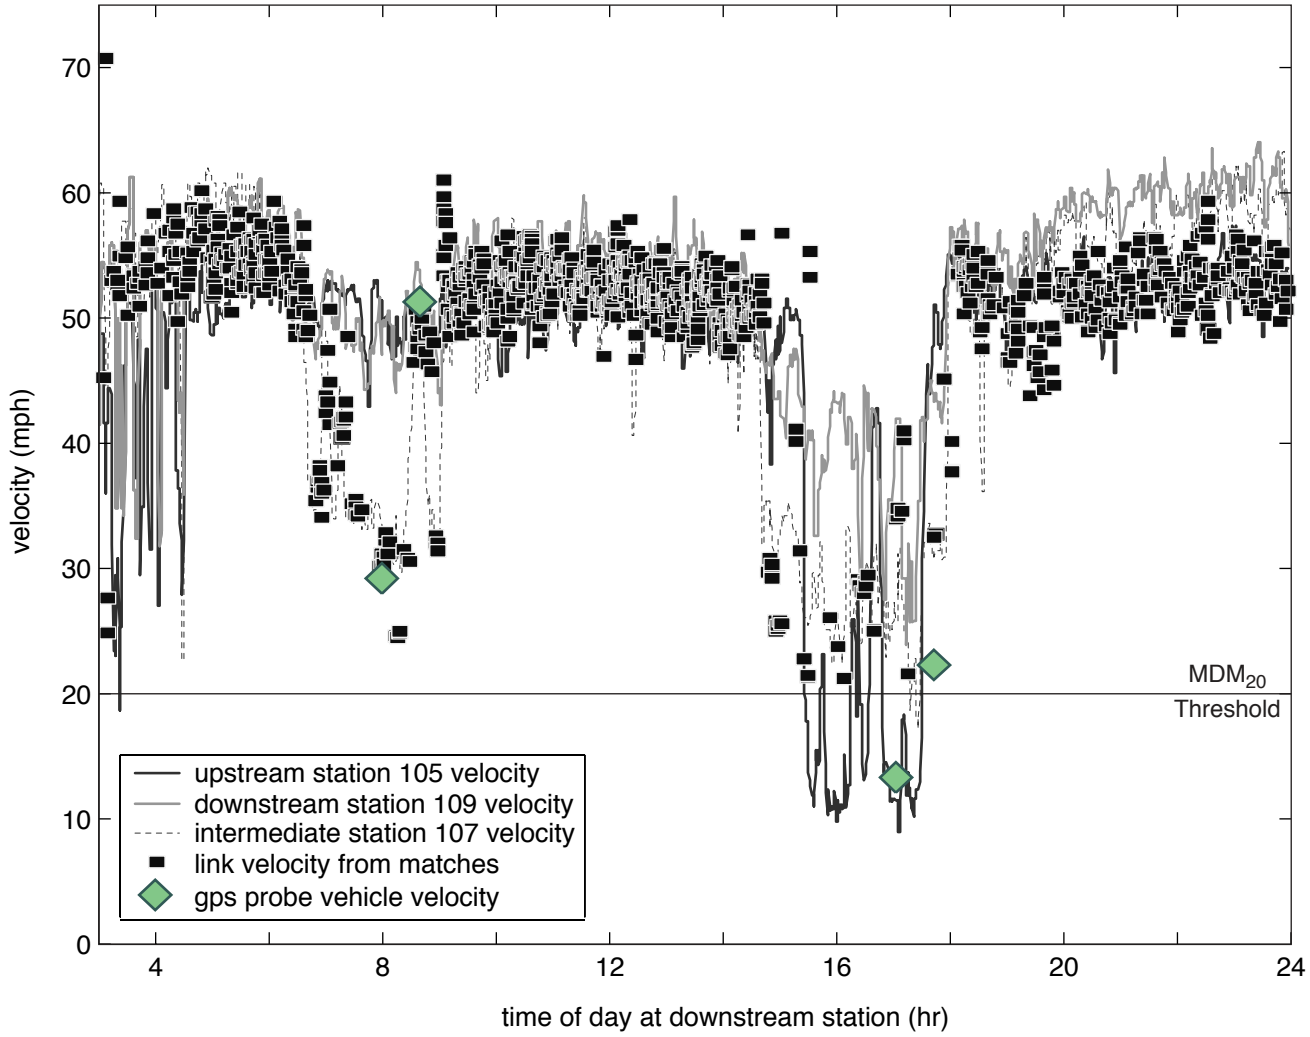

Figure 10, Results of vehicle reidentification between lane 2 at station 109 and lane 1 and 2 at station 105 on the northbound I-71 corridor using single loop data.

Table 1, Reidentification results between 0300-2400 hrs

Case	Facility	Date	Upstream	Downstream	Distance (miles)	Total No vehicles downstream	Total No of long vehicles downstream	Total No of vehicles reidentified		% of long vehicles reidentified	
								dual loop	single loop	dual loop	single loop
1	I-80 westbound	July 15, 2003	station 2, lanes (2,3,4)	station 5, lane (3)	0.91	26,095	3,533	1,277		36%	
2	I-71 southbound	May 29, 2003	station 108, lanes (2,3)	station 103, lane (2)	0.97	28,668	4,254	1,548		36%	
3	I-71 northbound	February 22, 2002	station 105, lanes (1,2)	station 109, lane (2)	0.95	33,960	4,430	1,809	1,284	41%	30%
4	I-71 northbound	February 22, 2002	station 106, lane (2)	station 107, lane (2)	0.39	29,060	4,371		1,579		36%