Paper No. 981498

PREPRINT

Title: Vehicle Reidentification and Travel Time Measurement in Real-Time on Freeways Using the Existing Loop Detector Infrastructure

Author: Benjamin Coifman

Transportation Research Board

77th Annual Meeting

January 11-15, 1998

Washington, DC

Vehicle Reidentification and Travel Time Measurement in Real-Time on Freeways Using the Existing Loop Detector Infrastructure

Benjamin Coifman

Institute of Transportation Studies

University of California

Berkeley, California, 94720

zephyr@eecs.berkeley.edu

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

ABSTRACT

This paper presents a new methodology for measuring travel time in real-time on freeways using the existing loop detector infrastructure. The methodology uses vehicle length, a simple signature from paired loop detectors, to extract information about the traffic stream. Other methods of measuring travel time either rely on approximations from aggregate traffic parameters or use expensive, proprietary systems to extract individual vehicle signatures.

The reidentification algorithm is tolerant to noise; instead of finding the ‘best match’ for each vehicle, it finds all possible matches. Using a bottom up approach, the algorithm looks for sequences of vehicles from the possible matches. Even with noisy loop detector data, the sequence detection eliminates most of the possible but incorrect matches while the true matches remain. The algorithm is fast, because it compares each pair of vehicles once and only once. Using loop detector speed traps, the communication bandwidth is small compared to other signature based vehicle reidentification methods. The algorithm can also be generalized to the higher resolution, proprietary detection systems.

The new methodology will be used to examine the applications and benefits of travel time data on real world traffic, without the expensive costs of a field test. Ordinarily, a travel time measurement system would have to be fully deployed before the benefits can be quantified. The algorithm outlined in this paper is designed to work with existing loop detectors, without any hardware modifications.

KEY WORDS:
Travel Time Measurement, Vehicle Reidentification, Traffic Surveillance - Freeways, Loop Detectors, Traffic Data

INTRODUCTION

In recent years, congestion has become a significant problem and it is clear that simply adding more lanes will not solve the problem. Many studies have found trends showing increasing congestion and forecast even worse conditions in the near future [1-2]. These data have led to research on how to use the existing infrastructure more efficiently. This paper presents a new approach to measuring travel time data and could enable or improve several applications such as incident detection or advanced traveler information systems (ATIS).

The potential benefits of incident detection have been known for years [3-5] and countless automated incident detection strategies have been proposed; yet, most jurisdictions continue to rely on labor intensive incident detection strategies (e.g., cellular phones) because of high false alarm rates with point detector based automated algorithms.

A recent report from Caltrans [6] noted that, “rapid changes in link travel time represent perhaps the most robust and deterministic indicator of an incident....” Travel time and related measures such as segment density can be used to quantify conditions between widely spaced detectors, even when the local, point based measurements are not representative of the entire segment. The Caltrans report also stated that, “link travel time ... is perhaps the most important parameter for ATIS functions such as congestion routing.” Many ATIS applications have yet to leave the simulation stage and the loop based system could help evaluate the applications, off-line, before possible deployment.

Several systems have been proposed for measuring travel time directly using vehicle signatures [7-17]. These emerging technologies use specialized hardware to extract the vehicle signatures. In most cases, the vehicle reidentification (VRI) and travel time measurement (TTM) systems have only been installed on a small test site. Full installation is necessary to quantify the benefits of the new systems, precluding cost-benefit analysis before purchasing the hardware. Furthermore, the systems have to be deployed before effective vehicle reidentification algorithms can be developed. Some of the systems use a permanent, unique vehicle identification (AVI) [7-13] that make VRI trivial, but they may compromise personal privacy.

Other systems have been proposed for estimating travel time from aggregate traffic parameters [18-19]. Although these systems appear promising for free flow and lightly congested conditions, they currently do not perform well under heavy congestion.

A third approach is to use cumulative arrivals at successive detector sites to estimate vehicle arrivals [20]. To counter detector drift, these systems use aggregate measurements to recalibrate during free flow conditions. Unfortunately, congestion can last several hours, leading to significant measurement drift between recalibrations.

This paper outlines a new approach: extract vehicle length from existing paired loop detector hardware and use it as a simple signature. A given vehicle length is not unique and may be noisy, but a short sequence of measured lengths is reidentifiable at successive detectors and it is robust to measurement noise.

The approach is similar to Pfannerstill, Kühne, et al [15-17], who simplified the problem of VRI by grouping vehicles into classes and then matched sequences of vehicle classifications rather than individual vehicles. In this fashion, a detector with limited resolution (e.g., only capable of identifying a small number of classes) can reliably reidentify vehicles via sequence matching. Pfannerstill's and Kühne's systems use existing loops in conjunction with proprietary hardware and software to extract a magnetic vehicle signature and assign a vehicle-class.

Instead of grouping vehicles into a finite number of vehicle classes, we demonstrate automated vehicle reidentification via sequence matching of vehicle lengths measured from the standard Caltrans speed trap. Preliminary results show that, unlike most loop based travel time measurement strategies, the system performs as well under congestion as it does under free flow conditions. Unlike Pfannerstill and Kühne, our system uses off-the-shelf 170 controllers and controller software already developed by PATH and Caltrans for the I-880 FSP study [21]

The reidentification rate based on speed trap length measurements is not as high as the emerging signature extraction technologies; but, because it can be implemented using the existing detection hardware, the benefits of travel time measurement can be quantified before a jurisdiction commits to purchasing a TTM system. In fact, the system can be implemented on the pre-existing FSP database which includes 20 detector stations, over a 11.5 km (7.2 mi) length in two directions for 50 days, with comprehensive incident data and approximately 10,000 link travel time measurements from probe vehicles [21].

Although this section presents competing technologies for measuring travel time, it is not intended to give the reader the impression that any one of the technologies is better than the others under all conditions. In fact, a hybrid between two or more systems will likely yield better performance than any one of the systems operating independently.

The remainder of this paper will outline the results of a pilot study to develop an automated VRI/TTM algorithm. First the basic theory and the data set are presented. Then, two VRI algorithms are explained *via* example. The first algorithm attempts to reidentify all of the vehicles, while the second algorithm only attempts to reidentify long vehicles and is applicable even when there are significant lane changes between detector stations. The paper closes with a brief discussion on field implementation.

VEHICLE REIDENTIFICATION USING SIMPLE SIGNATURES

This section illustrates the principles of extracting identifiable features at widely spaced detector stations. Because travel time is simply the difference between arrival times at two locations, the emphasis is on VRI rather than the final step of TTM.

The study uses 60 Hz event data from two paired loop detector speed traps in the same lane, 1.55 km (5100 ft) apart, as shown in Figure 1. Two hours of data from March 10, 1993 [21] are used to illustrate the VRI process. Throughout this paper, time is expressed in seconds with zero corresponding to 7:44 AM.

At each of the speed traps, upstream and downstream loop pulses were matched using an automated procedure [22] and any unmatched pulses were discarded (less than 0.5 percent of all pulses). Vehicle lengths were calculated from the matched pulses using the individual vehicle occupancy and travel times over the speed trap. Thus, the system measures effective vehicle length rather than the physical vehicle length and the implications for field implementation are addressed at the end of the paper.

Within each speed trap, the two loop detectors are 6.1 m (20 ft) apart and sample at 60 Hz. As a result, the vehicle length resolution ranges from 15 cm at 32 km/h to 60 cm at 125 km/h (0.5 ft at 20 mph to 2 ft at 80 mph). In addition to the resolution constraint, measurements are subject to external noise.

The VRI algorithm finds all *possible* matches for a given vehicle and then searches for sequences of vehicles within the possible matches. The algorithm is fast, because it compares each pair of vehicles once and only once. Sequence matching makes the VRI algorithm robust to noise and resolution constraints on individual measurements because the sequence information from a platoon contains significantly more information than a single vehicle measurement. In addition, the VRI algorithm exploits the fact that some vehicles are distinct from the general population (e.g., long trucks) and the fact that vehicle length resolution improves as velocity decreases.

Using loop detector speed traps, the communication bandwidth is small compared to other signature based vehicle reidentification methods. The VRI algorithm can also be generalized to the higher resolution, proprietary signature extraction systems.

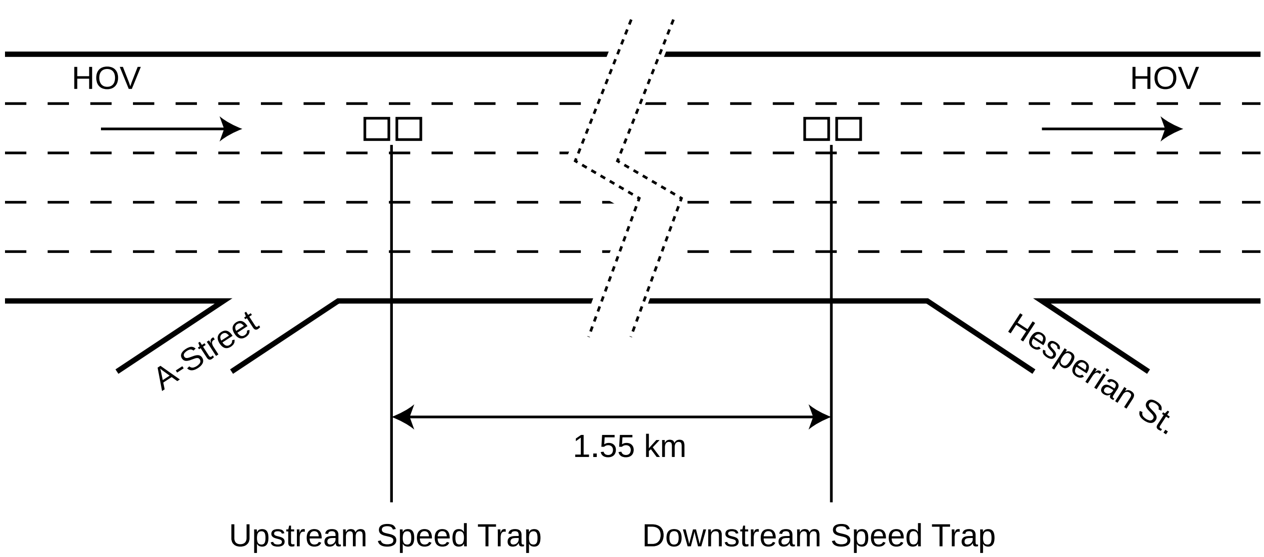

FIGURE 1: Region of pilot study on I-880, south of Oakland, California

Vehicle Length Data From Two Consecutive Detector Stations, Manual VRI

For this pilot study, the VRI algorithms are tested against manually generated reidentifications. Later analysis will use probe vehicle travel times to verify the VRI/TTM.

To illustrate the manual vehicle reidentification process, Figure 2 shows just over two minutes of time series vehicle length data extracted at the two speed traps. Note that upstream and downstream series are observed at different times. The data come from congested conditions; velocities were between 8 km/h and 32 km/h (5 mph and 20 mph) at both stations during this sample.

Indexing vehicles by arrival number rather than time, Figure 3A shows the two vehicle length sequences superimposed on the same plot while Figure 3B shows the corresponding velocities for reference. Note that four breaks were manually inserted in the upstream sequence to improve the match and presumably, to represent vehicles that entered the lane between the two speed traps. The breaks were inserted strictly on the basis of improving the match between the sequences.

The sequence of vehicle lengths at the upstream and downstream speed trap are very similar. The difference between upstream and downstream length is less than 15 cm (0.5 ft) for 75 percent of the vehicles in the sequence. The strong similarity between the two sequences, in conjunction with the correlation of the two long vehicles (points A & B), suggest that VRI based on sequential vehicle length is feasible.

The VRI algorithms attempt to automate this procedure. They look for short sequences with strong similarity and thus, the algorithms are tolerant to lane changes between the sequences. Furthermore, the algorithms exploit any distinct vehicles such as the long vehicles in the previous example.

Automated Vehicle Reidentification Using a Fixed Window: 'The Basic VRI Algorithm'

The basic VRI algorithm attempts to match each vehicle's length measurement at the downstream station with its corresponding upstream measurement. The algorithm starts by comparing individual length measurements between the two stations using a resolution test, as described below. If the difference between the upstream and the downstream measurements exceed the measurement resolution then the observations probably did <u>not</u> come from the same vehicle. The pair of vehicles can then be marked as an *unlikely* match. Otherwise, the pair of measurements can not be eliminated by this test and the pair is marked as a *possible match*.

The algorithm applies the resolution test to each pair of upstream and downstream measurements from some specified group of vehicles. In this study, the subject group of vehicles was limited to a fixed set of upstream and downstream measurements, e.g., the two sequences shown in Figure 3. The results of these resolution tests can be summarized in a *vehicle match matrix*. The matrix is indexed by arrival number at each station (upstream and downstream) and each element of the matrix is the outcome of a single pair-wise resolution test. Figure 4A-C shows an example of the notation used in the *vehicle match matrix.*

The data from Figure 3 yield the *vehicle match matrix* shown in Figure 5. The horizontal axis is indexed by upstream arrival number and the vertical axis is indexed by downstream arrival number. In

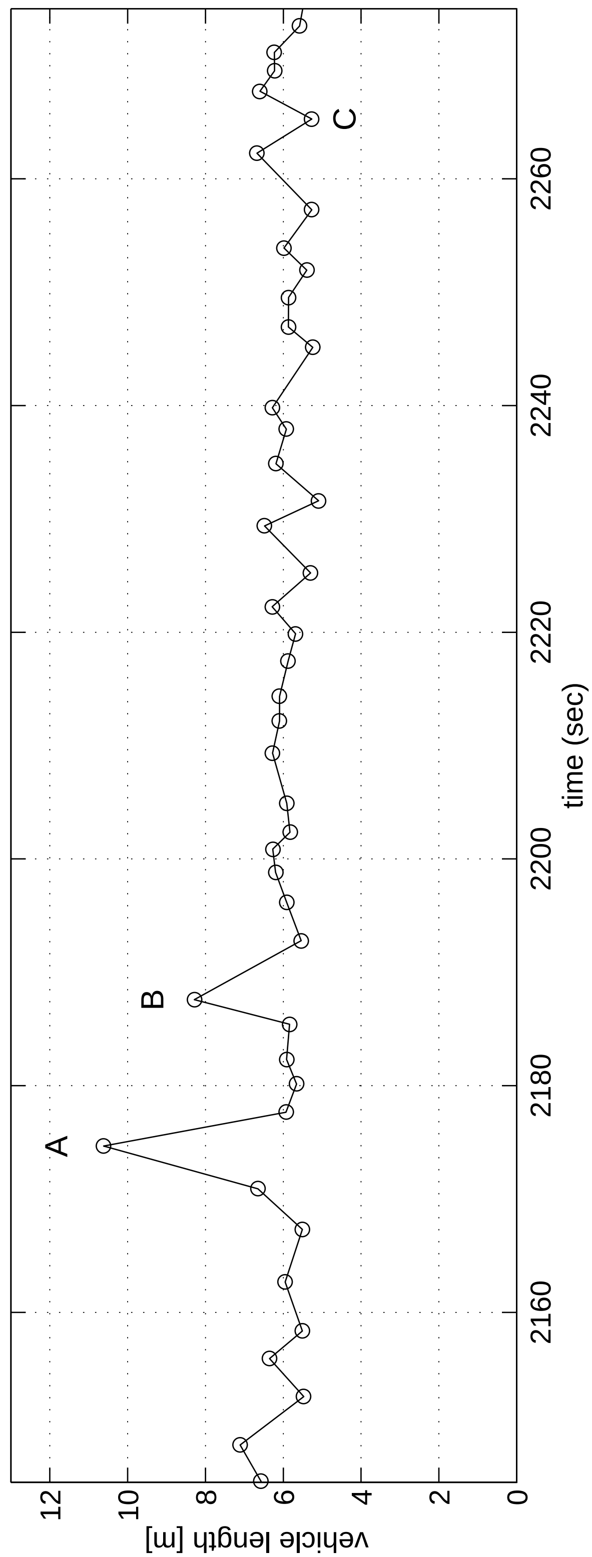

FIGURE 2A: Detail of the upstream vehicle length time series

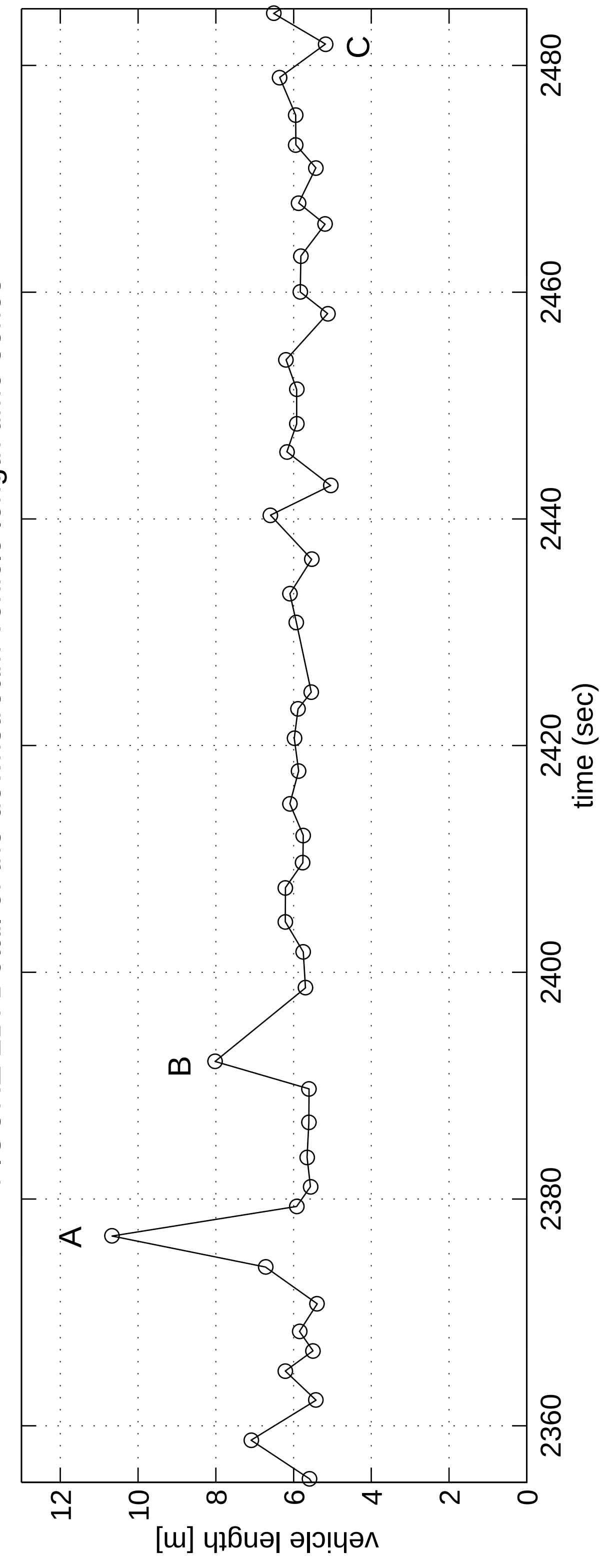

FIGURE 2B: Detail of the downstream vehicle length time series

FIGURE 3A: Superposition of the vehicle lengths from Figure 2

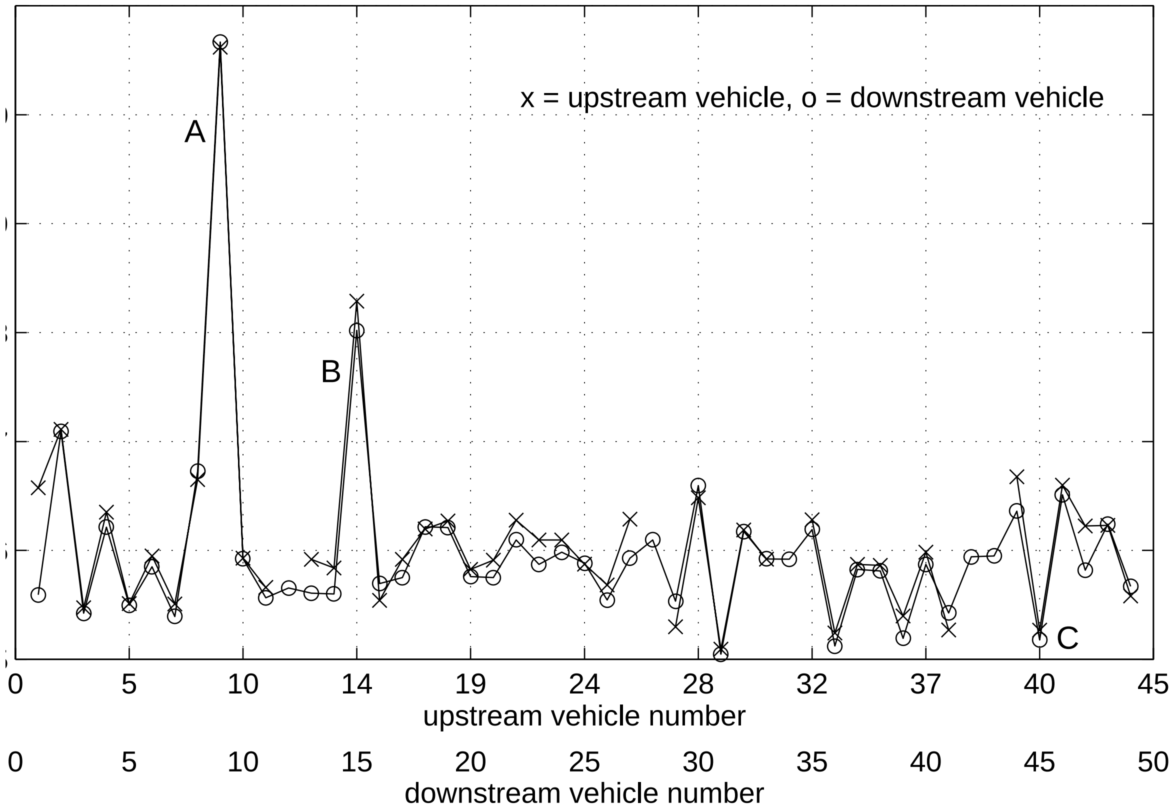

FIGURE 3B: The corresponding measured vehicle velocities

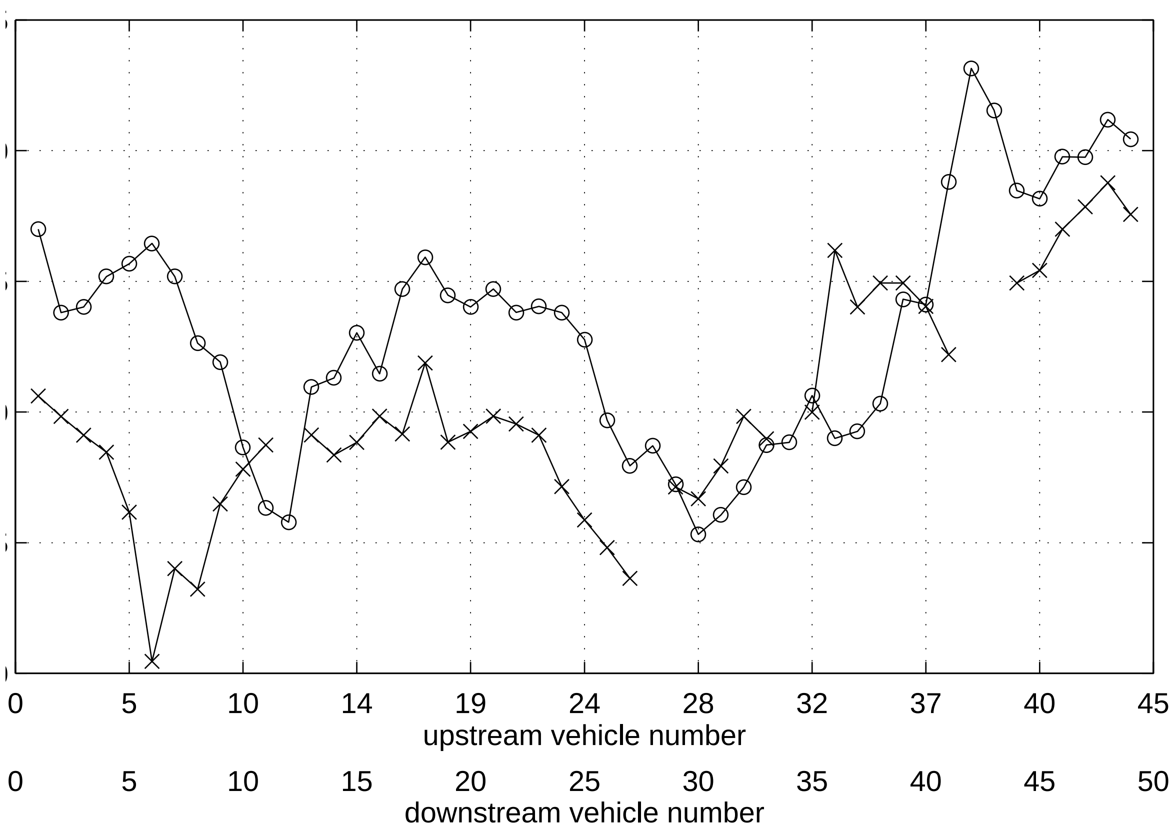

FIGURE 4: Part 1, simple example of notation: (A) measured vehicle lengths, (B) possible correlations with a length measurement tolerance of 1 unit, (C) resulting vehicle match matrix. Part 2, an example illustrating the transition from (D) a vehicle match matrix to (E) the Sequence Matrix. Each non-zero element in the Sequence Matrix indicates the total number of Possible Matches in the sequence up to and including the given matrix element.

(A)

downstream:

vehicle #	length
1	20
2	30
3	16

upstream:

vehicle #	length
1	20
2	31
3	21
4	15

(B)

possible correlations (d,u)
(1,1)
(1,3)
(2,2)
(3,4)

(C)

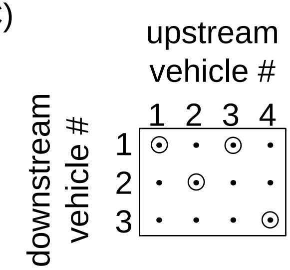

"o"= length within tolerance, thus, a possible match

(D) Vehicle Match Matrix

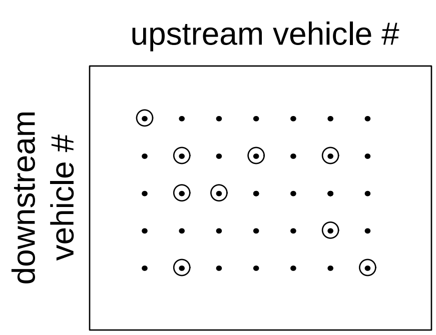

"o"= possible match

(E) Sequence matrix

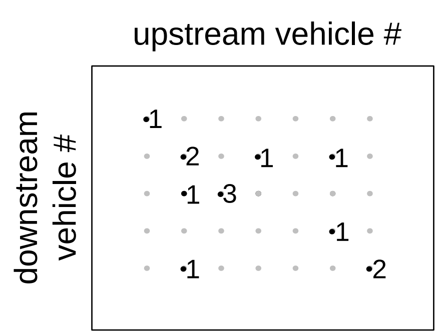

this figure, "O" indicates a *possible match* because the two length measurements are within the length resolution, while all other elements are left empty to indicate that a match is unlikely between the pair of vehicles. The solid line in this figure indicates the manually generated reidentifications.

Even without knowing the ground truth matches, many false positives are clearly evident in Figure 5. Each vehicle can only have, at most, one true match, yet most rows and columns have far more than one *possible match* for the given vehicle. Assuming that any two successive length measurements at a speed trap are independent of each other, the false positives are manifest as random noise in the *vehicle match matrix*. If the probability of a false positive occurring is less than 0.5, a false positive

FIGURE 5: Vehicle match matrix

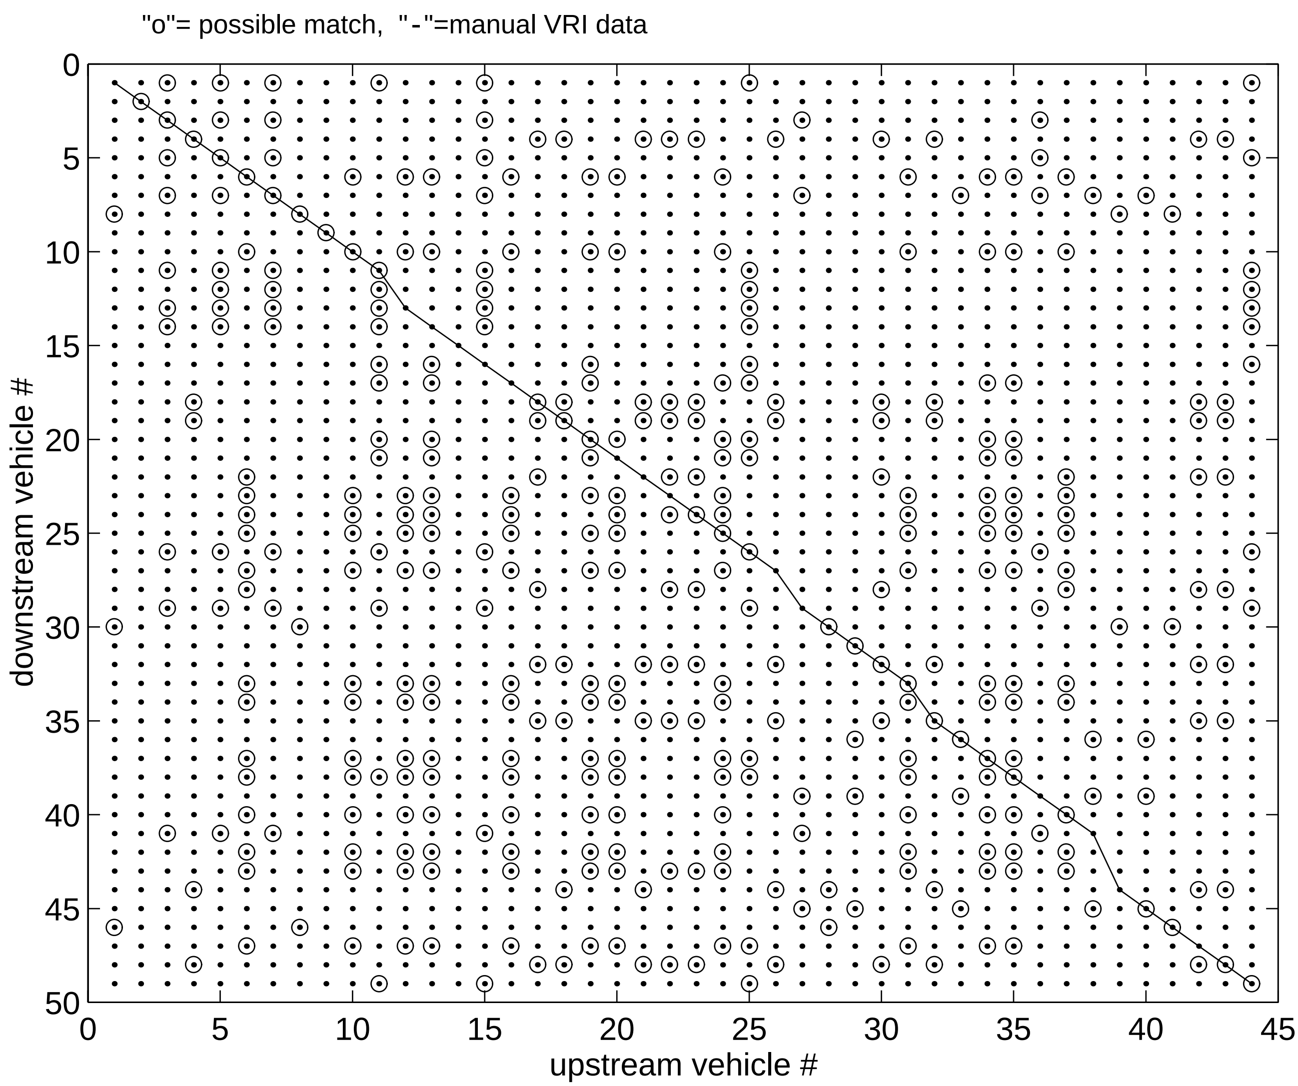

should usually be preceded (moving up one row and shifting left one column in the matrix) by an *unlikely* element. Whereas, if vehicles maintained their order between the two stations and the probability of a false negative is less than 0.5, a true match should usually be preceded by a *possible match* element. Relaxing the order constraint somewhat by assuming that vehicles usually maintain their order between stations, the true (but unknown) matches should manifest themselves as sequences (diagonal lines at -45 degrees) of *possible matches* in the *vehicle match matrix*. In other words, false positives will typically form short sequences while the true matches will usually form longer sequences in the *vehicle match matrix*. To exploit this property, the algorithm looks for sequences of *potential matches* in the *vehicle match matrix* and tallies how many sequential vehicles matched at both stations. These totals are stored in the *sequence matrix;* each integer is the cumulative number of *potential matches* in a sequence up to the given element. Figure 4D-E shows a simple example of the conversion to the *sequence matrix*. The *sequence matrix* for the on-going example is shown in Figure 6, where elements of length one have been omitted for clarity.

FIGURE 6: Sequence matrix, indicating the sequential number of possible matches

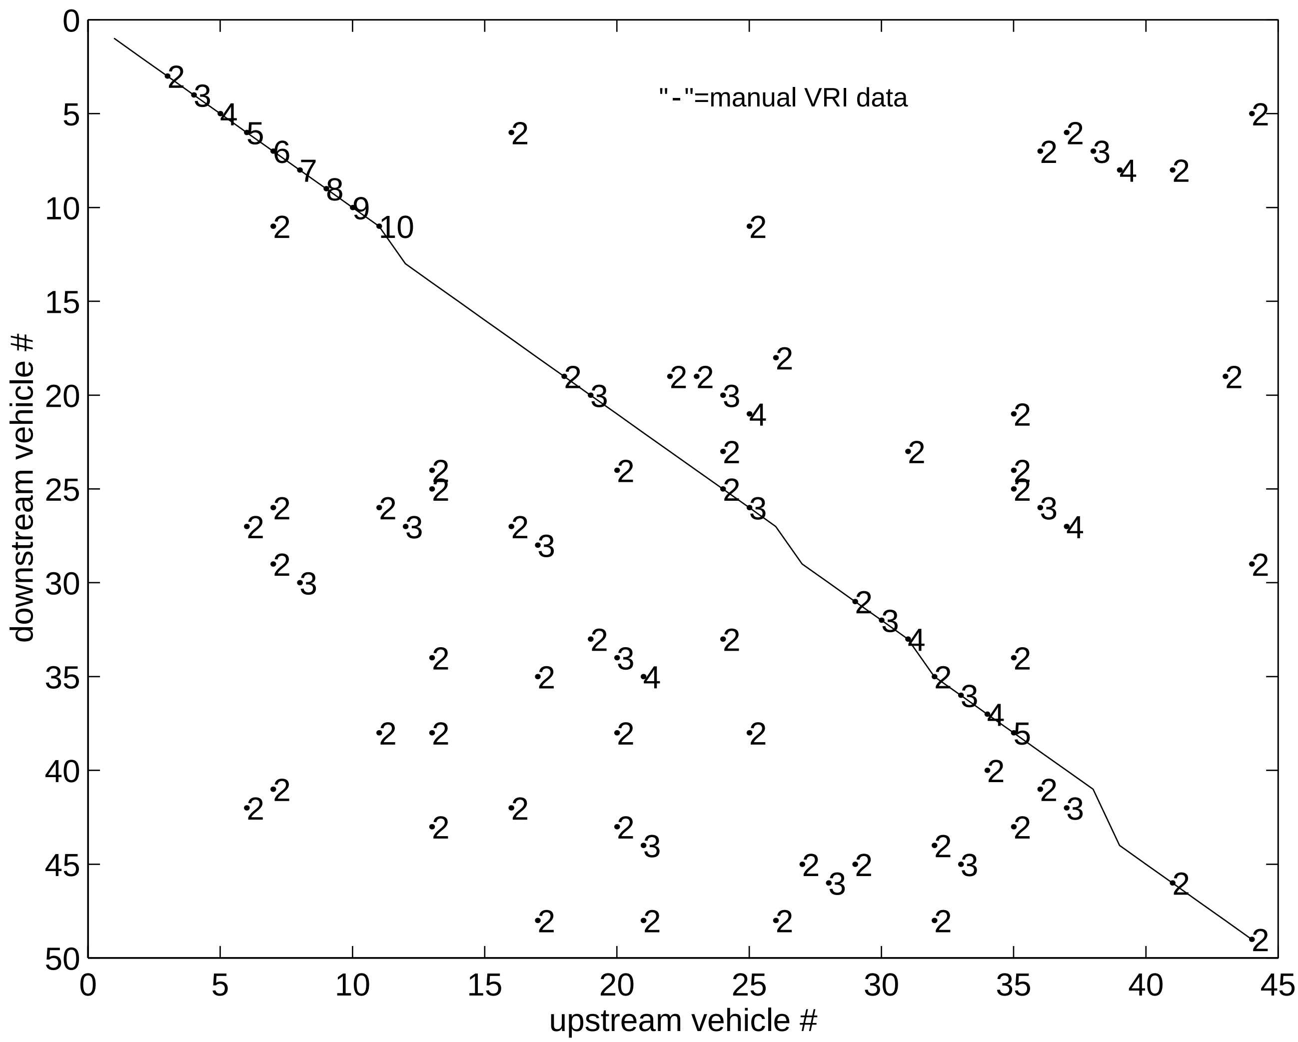

Next the algorithm allows for lane changes and/or misdetections in the sequences. Figure 7A-C shows the three lane change maneuvers searched for by the algorithm:

one vehicle exits the lane between stations or a vehicle is not detected at the downstream station, upstream vehicle n-1 in the example,

one vehicle enters the lane between stations or a vehicle is not detected at the upstream station, downstream vehicle m-1 in the example,

one vehicle enters and one vehicle leaves the lane between stations or there is a false negative in the data, vehicles m-1, n-1 in the example.

For each sequence of vehicles in the *sequence matrix*, the algorithm checks the first element to see if it could be linked to an earlier sequence (i.e., a sequence starting with a lower vehicle number) via a lane

FIGURE 7: A simple example illustrating the possible lane change maneuvers recognized by the Basic Algorithm: (A) One vehicle exits the lane between stations, (B) One vehicle enters the lane between stations, (C) One vehicle enters and one vehicle exits the lane between stations, (D) The search region for the sequence starting at element (m,n), (E) a hypothetical sequence matrix with (F) the resulting lane change matrix with a modified-sequence starting at element (m,n) shown in black.

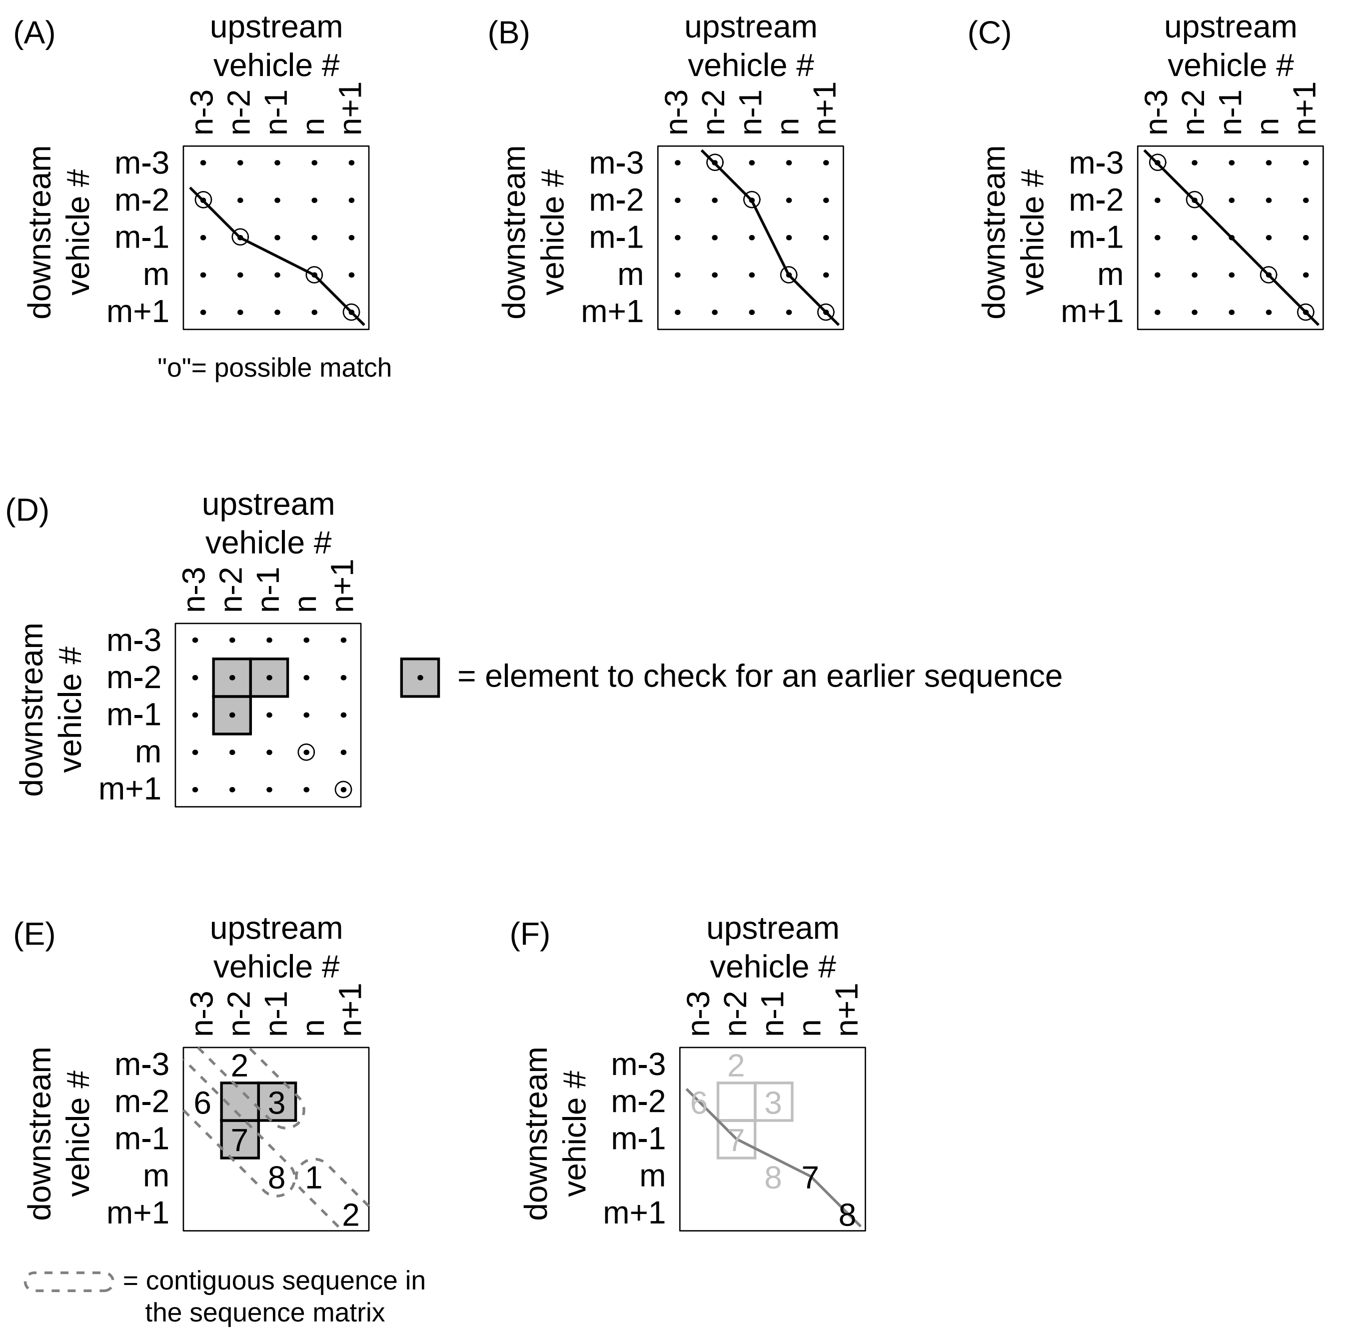

FIGURE 8: Lane change matrix, allowing for modified-sequences containing a single lane change maneuver

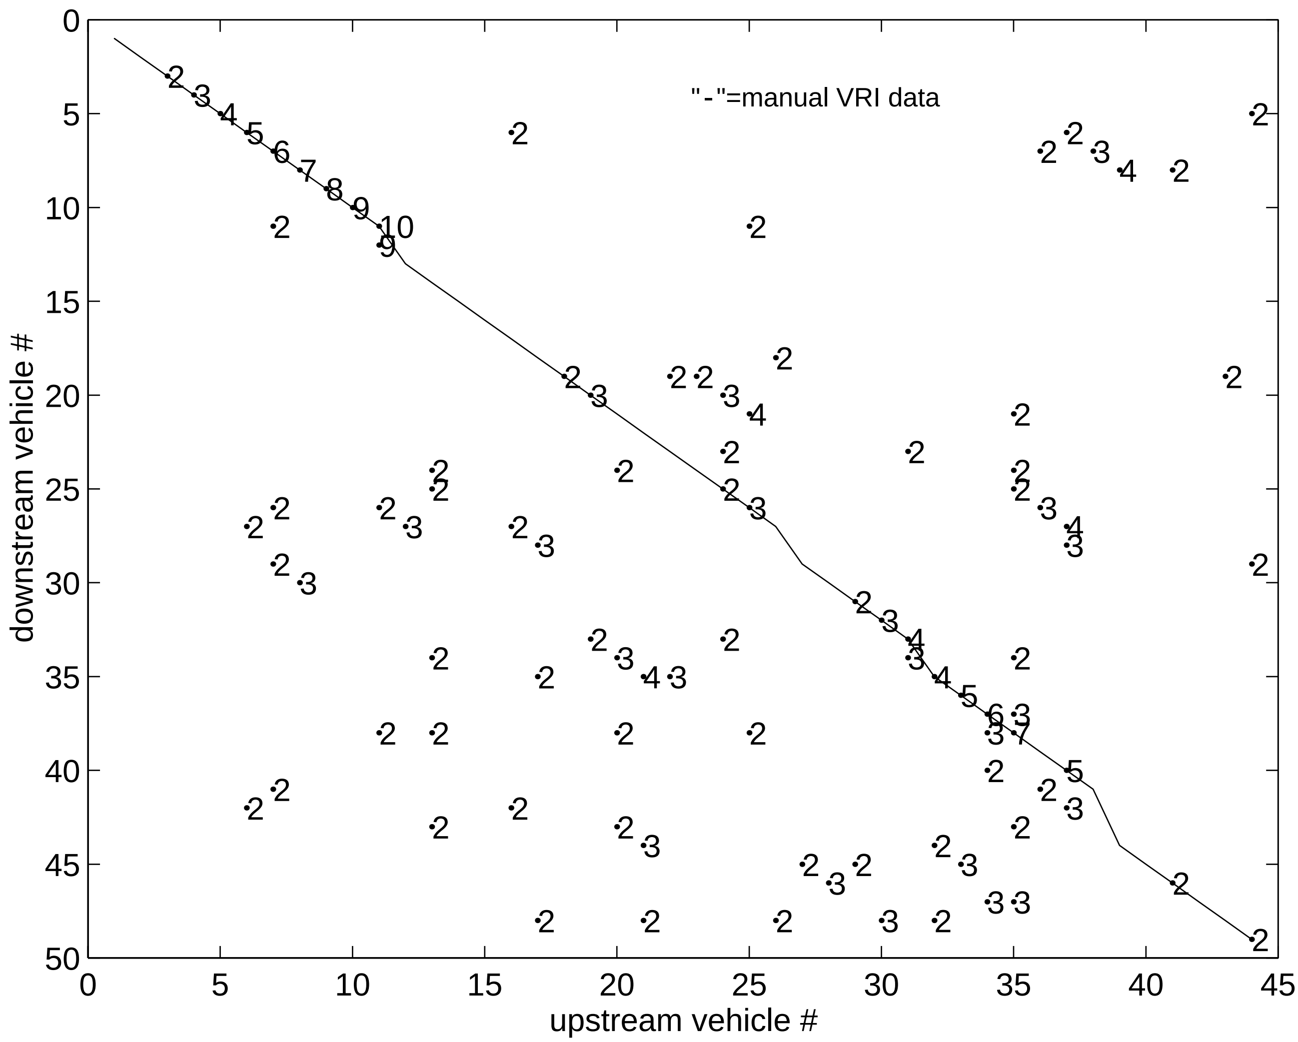

change maneuver. The procedure is demonstrated using the sequence starting with element (m,n) in Figure 7D, the algorithm checks the *sequence matrix* to see if there are any earlier sequences passing through one of the three shaded elements, where each element corresponds to one of the lane change maneuvers. If so, the algorithm increments all elements in the sequence starting at (m,n) by the highest value from the shaded elements in the *sequence matrix*, less a penalty of one vehicle for the lane change, and places the modified-sequence in the *lane change matrix*. The penalty gives contiguous sequences a slight advantage in the final step of the algorithm. Otherwise, if there are no preceding sequences in the shaded elements, then the algorithm simply copies the entire sequence unchanged from the *sequence matrix* to the *lane change matrix*.

For example, Figure 7E shows a hypothetical *sequence matrix* with three sequences, two of which start before downstream vehicle m-3 and are not shown in their entirety. When the algorithm reaches the sequence starting at (m,n), it finds that there are two earlier sequences that pass through the search area (shown in gray). It takes the highest value in the search area, 7, subtracts 1, adds the result

FIGURE 9: Threshold matrix, retaining only those sequences longer than a threshold length

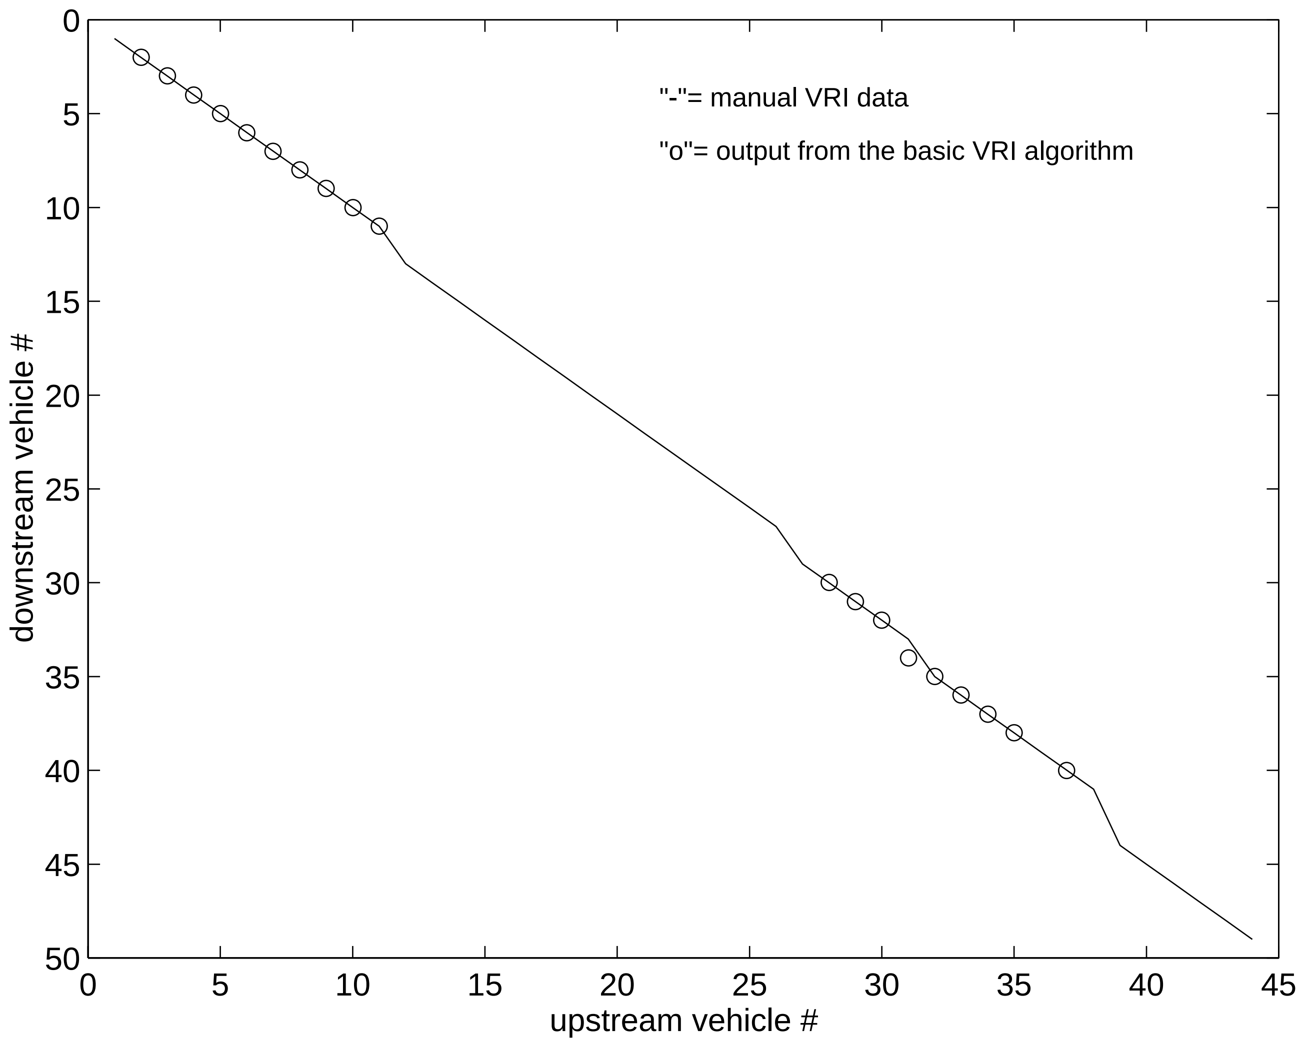

to all of the elements in the current sequence and then places the modified-sequence in the *lane change matrix*, Figure 7F. Figure 8 shows the *lane change matrix* for the on-going example, again, all elements of length one are omitted for clarity.

Finally, the algorithm extracts all sequences from the *lane change matrix* longer than a pre-specified threshold, yielding the *threshold matrix*. Entire sequences (and modified-sequences) are selected from the *lane change matrix*, successively from longest to shortest. Once a given match has been identified, the corresponding row and column of the *lane change matrix* are removed from further considerations. In the on-going example, a threshold level of five matches for a sequence yields the two platoons shown in Figure 9. Note that both platoons fall on the manually calibrated data and almost half of the vehicles that passed the detector stations were reidentified (i.e., matched).

Travel time for a reidentified vehicle can then be measured by taking the difference in known arrival times at the two stations. To estimate travel time during the short periods with no reidentified

FIGURE 10: The Subsampling VRI Algorithm, steps 1-4 apply the Basic algorithm to all vehicles over a pre-specified length

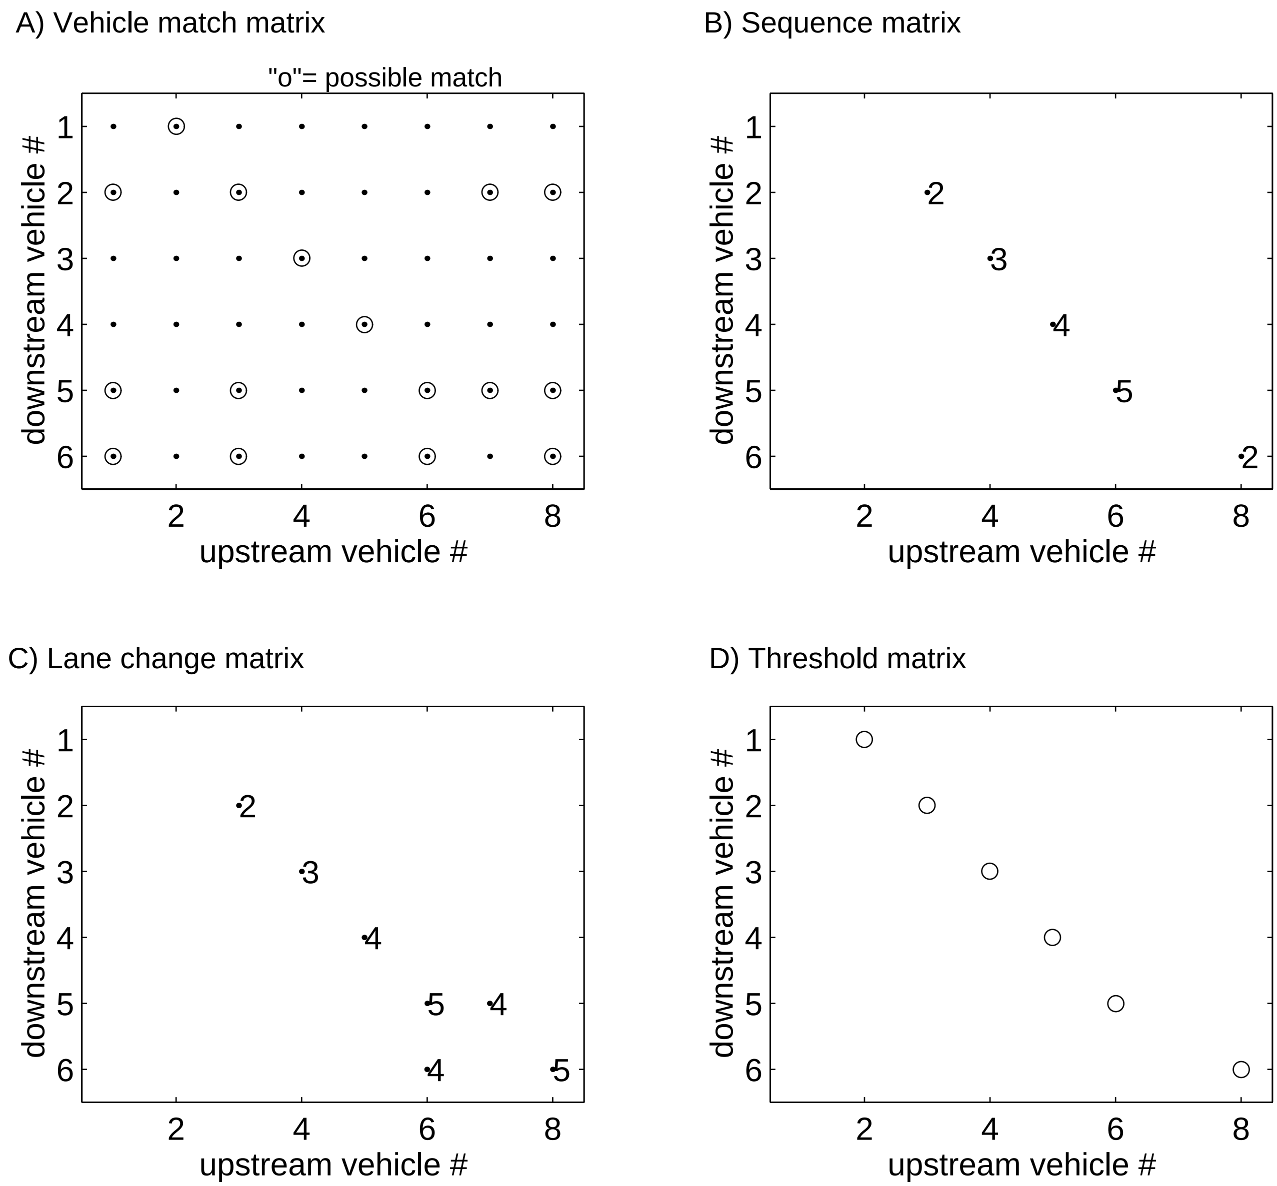

vehicles, the reidentification process can be approximated by matching vehicles based on the cumulative number to pass each station after the last matched sequence.

Extending the Method to Low Measurement Resolution: 'The Subsampling VRI Algorithm'

Under free flow conditions, the length measurement resolution degrades causing the number of *possible matches* to increase in the *Basic Algorithm*. Furthermore, vehicles may be less likely to maintain their order between detector stations in free flow conditions due to frequent opportunities to overtake one another. Subsampling a distinct segment of the total sample can overcome these problems.

FIGURE 11: Transpose the subsample matches back to the original vehicle indices

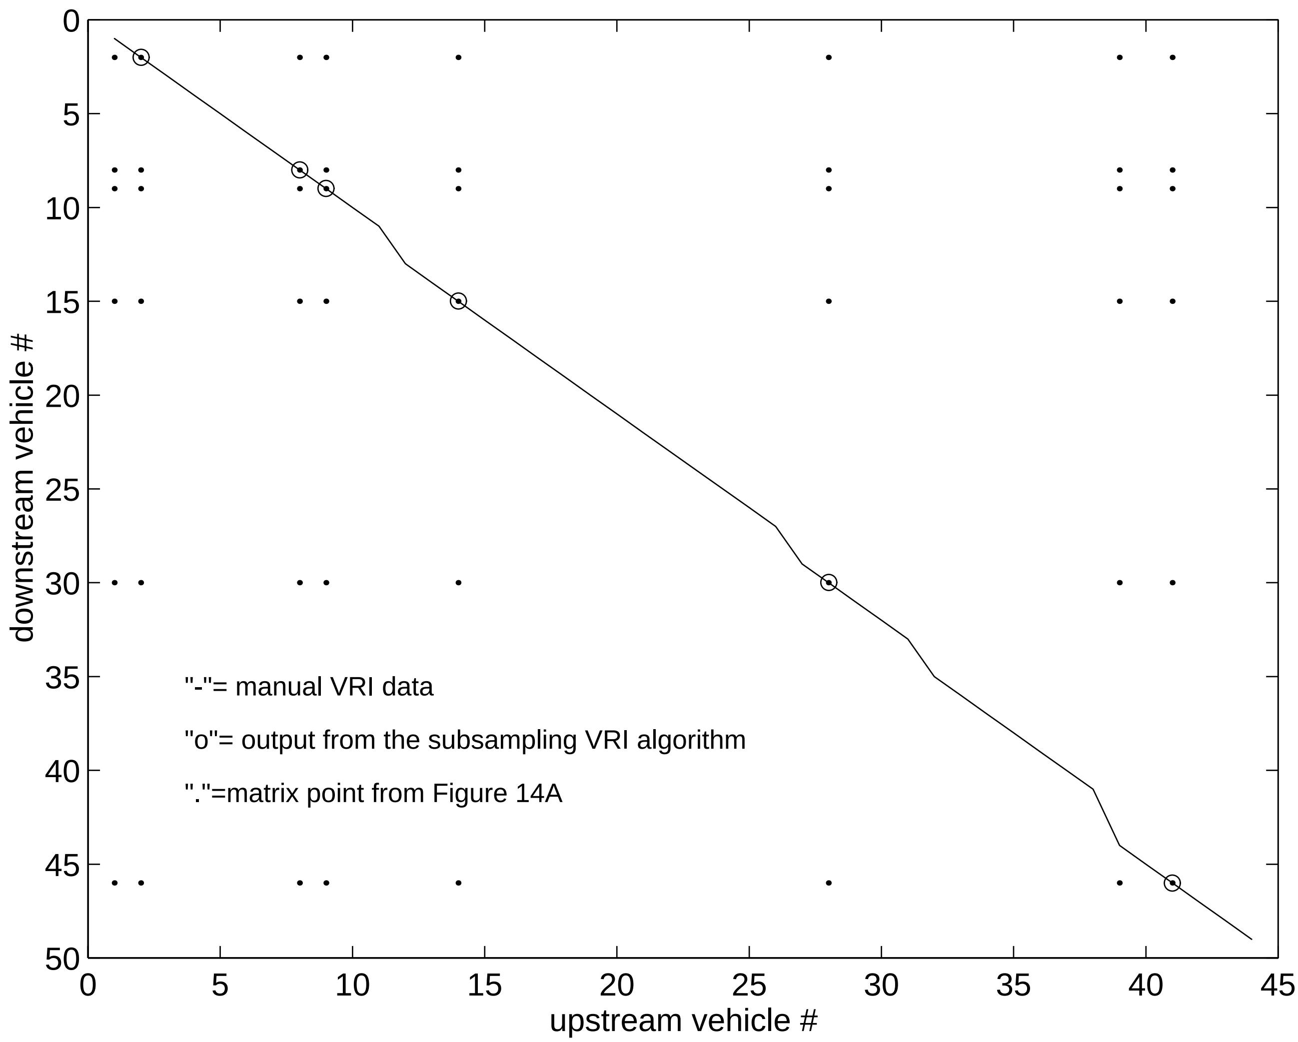

Most vehicles on the highway (e.g., sedans, pickup trucks, etc.) are small and have effective lengths of about 5-7 m (16-22 ft), i.e., the range of these effective lengths is on the order of 2 m (6 ft). The length measurement resolution may be as poor as 60 cm (2 ft) under free flow conditions, making it difficult to differentiate one small vehicle from another using the resolution test. The effective length for long vehicles, on the other hand, can range from 7 m to over 24 m (22 ft to over 80 ft) (upper limit is a semi truck with two trailers). By restricting the *Basic Algorithm* exclusively to long vehicles, the large range of lengths can offset the degraded measurement resolution. Because the long vehicles make up a small portion of the population, there will frequently be large headways between two successive observations. The large headways reduce the opportunity for overtaking and increase the probability of maintaining the vehicle sequence between detector stations.

Before comparing measurements from two stations, the algorithm “subsamples” all vehicles longer than some pre-specified minimum length at each station and assigns sequential integers according to their arrival. Using the data in Figure 3 and a minimum length of 6.4 m (21 ft), the algorithm subsamples about 20 percent of the vehicles at each station. The *Subsampling Algorithm* applies the *Basic Algorithm* only to the subsamples, and follows the same steps to reidentify vehicles. First, the algorithm generates a *vehicle match matrix* (Figure 10A), second, it identifies sequences of potential matches (Figure 10B), third, it allows for lane change maneuvers (Figure 10C), fourth, it keeps only those sequences over a given threshold (Figure 10D). Finally, the matches from Figure 10D are transposed back to the original sample as shown in Figure 11. Note that the *Subsampling Algorithm* has correctly reidentified two vehicles, downstream #15 and #46, that were not matched using the *Basic Algorithm* in Figure 9.

A second example of the subsampling VRI algorithm

This section illustrates the *Subsampling Algorithm* under the adverse conditions of Figure 12A-B; traffic at the upstream station is free flowing until the very end of the sample while traffic at the downstream station starts out with slower velocities then collapses into stop and go traffic after a shock wave passes through the segment. The two time series are plotted with different time scales. Over 250 vehicles pass the two stations in these figures. All vehicles over 6.7 m (22 ft) are subsampled, yielding 21 upstream vehicles and 13 downstream vehicles. The total run time, from cleaning up the raw detector data to generating the matches was less than 6 seconds for this sequence.

Figure 12C shows the measured travel times versus downstream arrival time using the same time scale as Figure 12B. The manual VRI for the same subsample of long vehicles are indicated by "X"'s. The travel time is slowly increasing before the downstream station breaks down, the last measured travel time before the breakdown is about 1.5 minutes, indicating the average velocity between stations is around 60 km/h (40 mph). The increasing travel times for the last two vehicles reflect the fact that the shock wave continues to propagate upstream (Although it is not shown in this figure, the shock wave eventually reaches the upstream detector station at t = 500 seconds).

FIELD IMPLEMENTATION

The VRI/TTM system presented in this paper is still in the development stages. All of the analysis was conducted on personal computers in the laboratory. A field deployable system would require the existing loop detector controller, an inexpensive PC (486 generation) and a low bandwidth communication link between detector stations. The loop detectors would require better tuning than current practice dictates, but the necessary tools are already included in the system. We currently use a speed trap diagnostic program on the PC to examine loop detector performance in the field. It simply requires 15 minutes of free flow traffic data to evaluate the loop calibration. The analysis tools have already identified a critical flaw in a large batch of loop amplifier cards in use by a state DOT. The flaw had not been detected by the manufacturer or the DOT.

FIGURE 12: Time series velocity and vehicle length data during challenging conditions at the onset of congestion

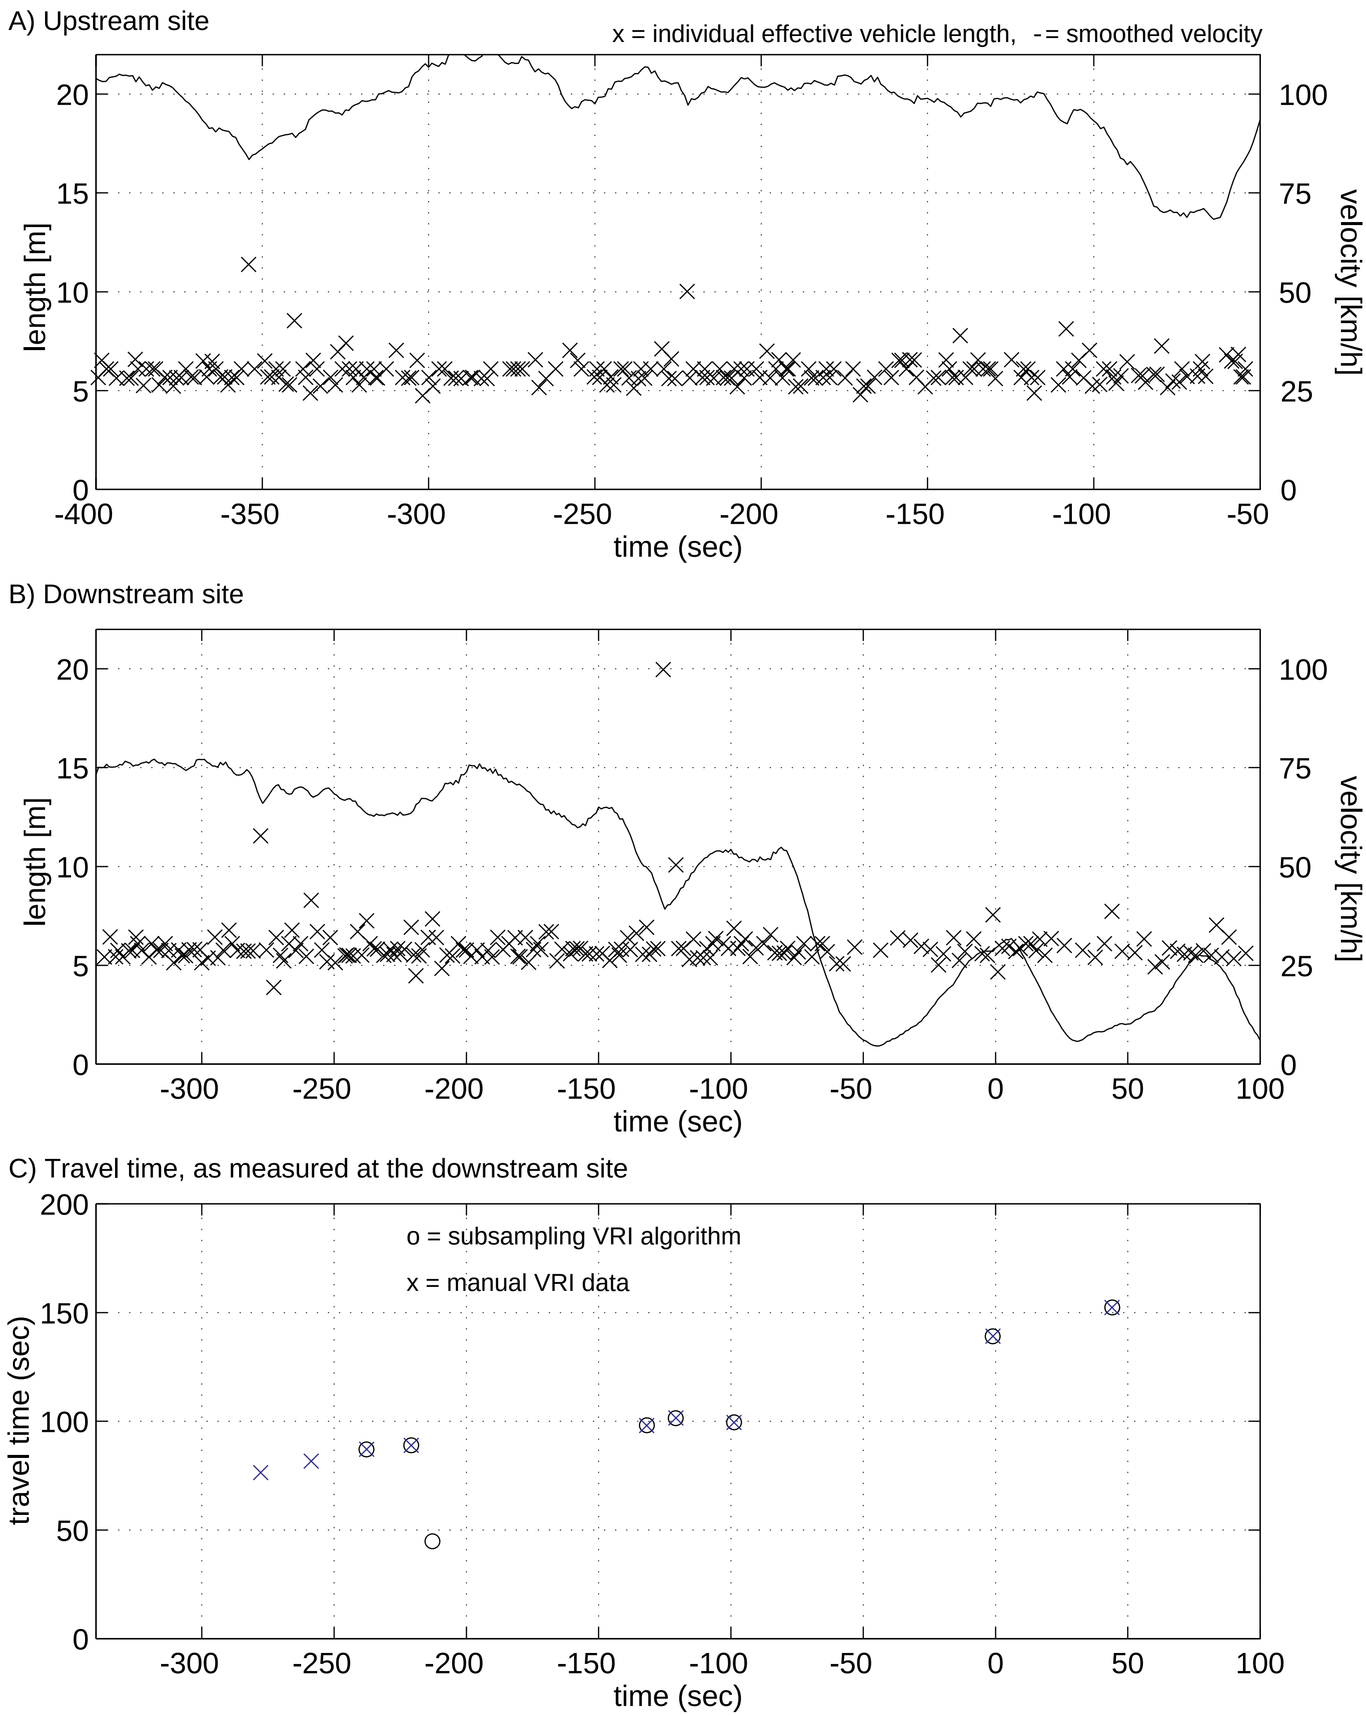

It is important to note, however, that the detectors must be distributed wisely. There will be short blind spots where the VRI system will not work (e.g., over a major diverge or a heavy weaving section). Intelligent detector placement will avoid these blind spots and could even allow for delay detection inside the blind spots [23].

CLOSING

This paper has presented a new methodology for VRI and TTM using a very simple signature: vehicle length. The illustration has shown that it is possible to reidentify vehicles between two speed traps, 1.6 km (1 mi) apart, using nothing more than measured vehicle lengths and temporal order. Unlike existing TTM systems that require new hardware to reidentify vehicles or simply estimate travel time from aggregate parameters, the new VRI/TTM system uses off-the-shelf 170-controllers and yields a vehicle level reidentification.

Although the accuracy of VRI based on speed trap length measurements can already be surpassed by emerging signature extraction technologies, the lower percentage of reidentified vehicles may be sufficient for many applications. Furthermore, the new system has one distinct advantage: it can be implemented using the existing detector infrastructure and controller hardware. Thus, it is possible to investigate applications of travel time data and quantify the benefits, off-line, without field deployment or making a major financial/institutional commitment to a particular technology.

If an application proves cost effective, the speed trap based system could be implemented in the field or the reidentification algorithms could be generalized to another signature based detector system.

REFERENCES

[1] Lindley, J. *Quantification of Urban Freeway Congestion and Analysis of Remedial Measures*, Federal Highway Administration, 1986.

[2] Schrank, D., Turner, S., Lomax, T. *Urban Roadway Congestion- 1982 to 1992*, Vol 1 & 2, Texas Transportation Institute, 1995.

[3] Goolsby, M. E. “Influence of Incidents on Freeway Quality of Service”, *Highway Research Record Number 349*, Transportation Research Board, 1971

[4] Messer, C., Dudek C., "Method for Predicting Travel Time and Other Operational Measures in Real-Time During Freeway Incident Conditions", *Highway Research Record 461*, Highway Research Board, 1973, pp 1-16

[5] Urbanek, G., Rogers, R. Alternative Surveillance Concepts and Methods for Freeway Incident Management, Volume 1: Executive Summary, Federal Highway Administration, 1978.

[6] Palen, J. "The Need for Surveillance in Intelligent Transportation Systems", *Intellimotion*, Vol 6, No 1, University of California PATH, Berkeley, CA, 1997, pp 1-3, 10.

[7] Levine, S., McCasland W., "Monitoring Freeway Traffic Conditions with Automatic Vehicle Identification Systems", *ITE Journal*, Vol 64, No 3, March 1994, pp 23-28.

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

[9] Christiansen, I, Hauer, L. "Probing for Travel Time: Norway Applies AVI and WIM Technologies for Section Probe Data", *Traffic Technology International*, Aug/Sep 1996, UK & International Press, Surrey, UK, 1996, pp 41-44.

[10] Oda, T, Takeuchi, K. and Niikura, S. "Travel Time Measurement Using Infrared Vehicle Detectors," *Proc of the 8th International Conference on Road Traffic Monitoring and Control*, Report 422, 1996, pp 178-182.

[11] Seki, S. "Travel-Time Measurement and Provision System Using AVI Units", Proc. 2nd World Congress on Intelligent Transport Systems, VERTIS, 1995, pp 50-55

[12] Tanaka, Y., Nishimura, F. "Multiroute Travel-Time Data Provision Systems Operating in Osaka", *Proc. 1994 Vehicle Navigation and Information Systems*, IEEE, 1994, pp 351-356.

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

[14] Malik, J., et. al., *Traffic Surveillance and Detection Technology Development: New Traffic Sensor Technology Final Report*, University of California PATH, Berkeley, CA, 1997.

[15] Pfannerstill, E. “A Pattern Recognition System for the Re-identification of Motor Vehicles”, *Proc. 7th International Conference on Pattern Recognition*, Montreal, IEEE, New Jersey, 1984, pp 553-555

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

[17] Kühne, R., Palen, J., Gardner, C., Ritchie, S. “Section-Related Measures of Traffic System Performance”, paper presented at the 76th annual TRB meeting, Transportation Research Board, 1997.

[18] Dailey, D. “Travel Time Estimation Using Cross Correlation Techniques,” *Transportation Research Part B*, Vol 27B, No 2, 1993, pp 97-107.

[19] Petty, K., et. al. “Accurate Estimation of Travel Times From Single Loop Detectors”, paper presented at the 76th annual TRB meeting, Transportation Research Board, 1997.

[20] Westerman, M., Immers, L. “A Method for Determining Real-Time Travel Times on Motorways,” *Road Transport Informatics/Intelligent Vehicle Highways Systems*, ISATA, 1992, pp 221-228.

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

[22] Coifman, B. *Automated Vehicle Reidentification and Travel Time Measurement Using Effective Vehicle Lengths From Loop Detector Speed Traps - Part I: A Framework for Vehicle Matching and Travel Time Measurement on Congested Freeways*, working paper, University of California, Institute of Transportation Studies, Berkeley, CA, 1997.

[23] Coifman, B. “A New Bottom-up Algorithm for Vehicle Reidentification and Travel Time Measurement on Freeways”, submitted to ASCE-AATT 1998.