George Lunt
TRL Ltd
glunt@trl.co.uk

Vehicle Re-Identification Using Induction Loop Data

Introduction

In this paper we report progress on the design of a vehicle tracking algorithm developed by TRL in conjunction with Bristol University. Once implemented, the algorithm will provide a comprehensive data set describing individual vehicle trajectories through a motorway. These trajectories can primarily be used to help calibrate and improve microscopic traffic simulation models.

We start by giving a brief overview of current microscopic car-following and lane changing models found in the literature, and describe the various sources of data currently available to calibrate them. We then provide a detailed description of the data source used by the algorithm described in this paper. Finally, we describe how the data can be used to re-identify individual vehicles as they travel through a motorway, and provide some results of early algorithm implementation.

Overview of car-following and lane changing models

Microscopic traffic simulation packages commonly contain two core behavioural models:

(1) a model that describes the acceleration of a vehicle at a given instant,
(2) a model that describes whether the vehicle will change lane at a given instant.

Vehicles are treated as discrete entities moving in time and space.

The first model is often of the form of a *car-following model* where the acceleration of the vehicle under consideration (subject) is dependent on the position and velocity of both the subject and the vehicle immediately ahead (target). A common relationship is that proposed by Gipps [1], which is based on the following behavioural model:

What speed should I travel at now, given the behaviour of the vehicle in front one reaction time ago? If the vehicle in front comes to a stop at what I think is its hardest rate, and one reaction time later I commence braking at my hardest rate, I must come to a stop safely.

This behavioural model can be converted into a mathematical formula which describes the new target vehicle's speed based on the speed and positions of the target and subject vehicle one reaction time ago. On implementing this car-following model on a single lane, circular road, Figure 1 shows what vehicle trajectories could be observed.

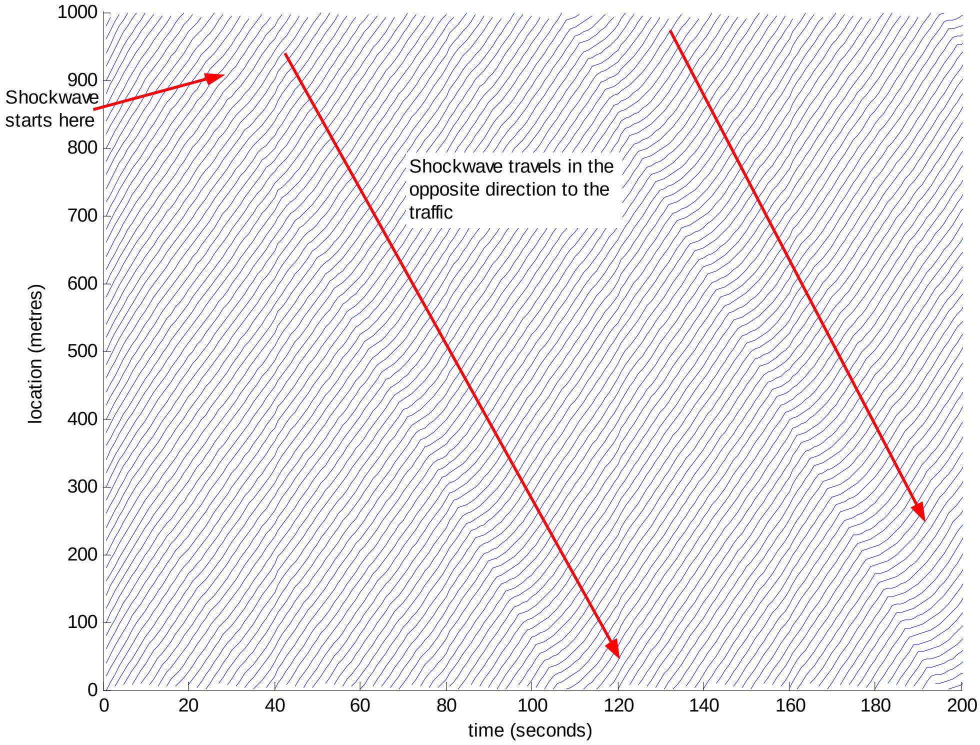

Figure 1. Vehicle trajectories on implementing the Gipps car following model.

Each vehicle at the start of the simulation has been given an initial velocity slightly perturbed around 20 metres/second. A travelling wave (or *shockwave*) is started (after about 30 seconds at 900 metres) and propagates continuously in the opposite direction to the travelling vehicles. It is widely believed that these travelling waves are also generated on busy motorways in the UK. There is currently very little data that helps describe how individual vehicle(s) initially generate these shockwaves, and also how vehicles behave on approaching and travelling through them. Application of such knowledge would help in motorway shockwave containment and prevention. The successful implementation of the algorithm described later in this paper will supplement this sparse data source.

In conjunction with the car-following model, microscopic simulation packages also contain a lane changing model which governs drivers' desire, availability, and mechanics of changing lane. Some reasons a driver may desire to change lane are to travel faster (the speed in a lane may be faster than the current lane), to reach a dedicated lane required for diverging, or in order to move into a 'preferred' lane. Once a driver has a desire to change lane then the lane changing model will assess the gaps available for the driver to move in to. Drivers with a high desire to change lane will accept smaller gaps than drivers with a low desire. Once a gap is deemed acceptable, some models also attempt to replicate the physical process of changing lane, which is thought to last – on average – about 5 seconds. There is very little traffic data available that could be used to research the various lane changing dynamics described above. Behavioural assumptions in most models are often implemented using arbitrary quantities and thresholds. Once again, the successful implementation of the algorithm described later in this paper will help fill this data-void.

Some very basic microscopic modelling ideas have been presented here, and in practice simulation packages include many more considerations in both the car following and the lane changing model. However, this overview has illustrated the very complicated nature of traffic simulation packages, and the importance of being able to test and calibrate the various behavioural assumptions outlined above with 'real' data. There is very little data available to perform such calibration and testing at the required level of detail, and most comparisons are performed using averaged, macroscopic traffic data.

Motorway Traffic Data

Traffic data can usually be partitioned into two distinct categories – macroscopic and microscopic traffic data. Macroscopic data describes averaged traffic statistics, or traffic measurements made across a large distance and/or time. For example, traffic measurements made at a specific point on the motorway averaged into minute-by-minute statistics would be classed as macroscopic data, as would matched vehicle data (using automatic video analysis of the number plate) obtained from two consecutive sites five miles apart. In the UK, these are the two most commonly used forms of data for performing quantitative analysis of motorway performance and operation. They are perfectly sufficient to carry out such tasks.

However, this data-class is often used to calibrate and test microscopic traffic models, once the model's output has been sufficiently 'averaged'. This means that very few microscopic models are strictly tested using microscopic traffic data, and so the fundamental properties of the model do not get calibrated. Such examples of microscopic traffic data are individual vehicles' positions every second created using in-car equipment (also called probe vehicle data) or vehicle trajectories calculated using automatic image processing of video footage. One of the problems in using probe vehicles is that the data set (typically of the order of 10 vehicle trajectories recorded over an entire day) does not capture the complete spectrum of driver behaviour. Furthermore, drivers of the probe vehicles know that they are being monitored, and so may drive differently to how they would under 'normal' conditions. It is possible to track vehicles using intelligent image processing techniques of video footage taken from a strategic location (like an overbridge). However, the technology is not quite adequate to provide a suitable and reliable data source. Furthermore, the operational cost of installing the equipment is quite high.

This paper is concerned with an alternative source of traffic data, that can be processed in such a way to provide tracked vehicle statistics, and is also available at a relatively low cost as it makes use of the existing infrastructure installed on the major UK motorways. The level of detail is much finer than the macroscopic data explained above, and does not carry the problems associated with the microscopic data.

The UK Highways Agency (HA) has developed an automatic speed-control environment. At its core is the Motorway Incident Detection and Automatic Signalling (MIDAS) system which monitors current traffic conditions. Alerts are communicated to a central control computer, which in turn relays messages back to motorists via message signs. In particular, the system at present uses an algorithm known as HIOCC to set temporary (and sometimes mandatory) speed limits upstream of queues in an attempt to improve safety. This programme is being rolled out across the UK.

The MIDAS detection system consists of outstations every 500m down the length of the motorway. Each outstation consists of a pair of inductance loops buried in the surface of the road and a smart signal processing and communications box. The magnetic field of passing vehicles induces a current in each loop and the system is wired so that the carriageway and lane number of the vehicle can be identified.

Furthermore, for each passing vehicle, the times at which each half of the inductance loop pair is activated and deactivated are captured. The microcontroller then uses this information to give each vehicle a timestamp (from which, by comparison with the timestamp of the preceding vehicle, the front-to-front time headway may be calculated). Further, the difference in activation times and longitudinal separation of each half of the loop pair may be used to estimate the speed of each vehicle.

The activation and deactivation times may also be used to calculate the *time over loop*, which combined with the speed estimate may be used to estimate the vehicle's length. All measurements are rounded at various stages in the microcontroller's algorithm, and the nominal accuracy of recorded measurements are 1s, 0.1s, 1km/h and 1cm for the timestamp, time headway, speed and length respectively.

In the past, outstations have not stored the full details of individual vehicles, but have bundled the data into one minute averages which is then sent to the central control computer. However, TRL has now intercepted the individual vehicle data (IVD) before it is lost, and consequently has access to comprehensive IVD sets from across the UK. This allows us to perform highway traffic analysis at a level of detail which has not previously been possible [2]. For example, individual speed-headway statistics can be obtained, and probabilistic speed-headway distributions derived.

Vehicle Re-Identification - Motivation

The IVD data described above can be used to gain a better understanding of traffic behaviour at the microscopic level. We now introduce the concept of re-identifying vehicles recorded at consecutive IVD detector sites.

The ability to track these vehicles would provide an additional dimension to the microscopic IVD data. The tracked vehicle statistics could be used to research how individual vehicles' speed, headway, and lane changed through the course of a journey. Furthermore, if vehicles did change lane, then the data would provide information on the types of gaps the vehicles accepted, and how this acceptance changed at various locations. All these factors would be incredibly useful in microscopic model calibration and validation; providing a level of detail and representation not previously available.

The data could also – in theory – be used to improve OD (origin-destination) information on a given motorway as vehicles would be tracked from their entry point to their exit point. This information is currently obtained using survey data (expensive and potentially un-representative) and OD estimation techniques (not always reliable and dependent on prior knowledge).

There is also scope to use the algorithm in real-time to detect incidents. Currently, incidents are raised by detecting the queue that is formed from the location of the incident. The traffic sensor immediately upstream from the incident is the first to detect the queue, which typically travels at a speed of 30 kph during periods of high demand. However, if the re-identification algorithm were running in real time, then it would instead detect the forward travelling wave that propagates from an incident, and travels at a much faster speed (about 100 kph). This wave is essentially a wave of sparse/no traffic, and represents the few/no vehicles able to drive through the incident. The algorithm would raise an incident alert if it were unable to re-identify a percentage of the vehicles it was expecting to arrive at the detector site immediately downstream from the loop. Therefore, if successfully implemented this algorithm would help improve the response time to incidents, and hence improve safety.

One of the great benefits of the algorithm is that the data it uses (IVD data) is relatively inexpensive to obtain, and makes use of the existing infrastructure (MIDAS). Other techniques used to obtain microscopic vehicle data are – in general – expensive to obtain, and require additional infrastructure to implement.

Vehicle Re-Identification – Initial Development

The inductance loop sites operated under MIDAS (which provide the mechanism to collect IVD data) are normally installed at regular intervals along the motorway. At consecutive MIDAS sites where there is no on or off ramp, all vehicles that pass the initial site will also pass the second site further downstream. If IVD data is being recorded at both consecutive sites, it is theoretically possible to re-identify the same vehicle from one location to the next.

Initially, two consecutive IVD sites were available for analysis at an interval of 1km. Initial algorithm development was slow as this is a relatively long distance over which to re-identify vehicles, and some of the key algorithm ideas had not been fully developed. However, these initial ideas are still discussed in this paper.

The initial algorithm generates a forecast ‘window’ for all vehicles detected at the upstream site. The window represents a time interval and length tolerance in which it would expect to detect the same vehicle at the downstream site. The time the vehicle is expected to arrive at the downstream site is dependent on the vehicle’s speed. Vehicles measured at the downstream site whose timestamp and length lies within the forecast window of the vehicle measured at the upstream site are considered as a ‘potential match’. Similarly, each potential match at the downstream site throws a back-forecast to records held at the upstream site. Figure 2 illustrates this idea further.

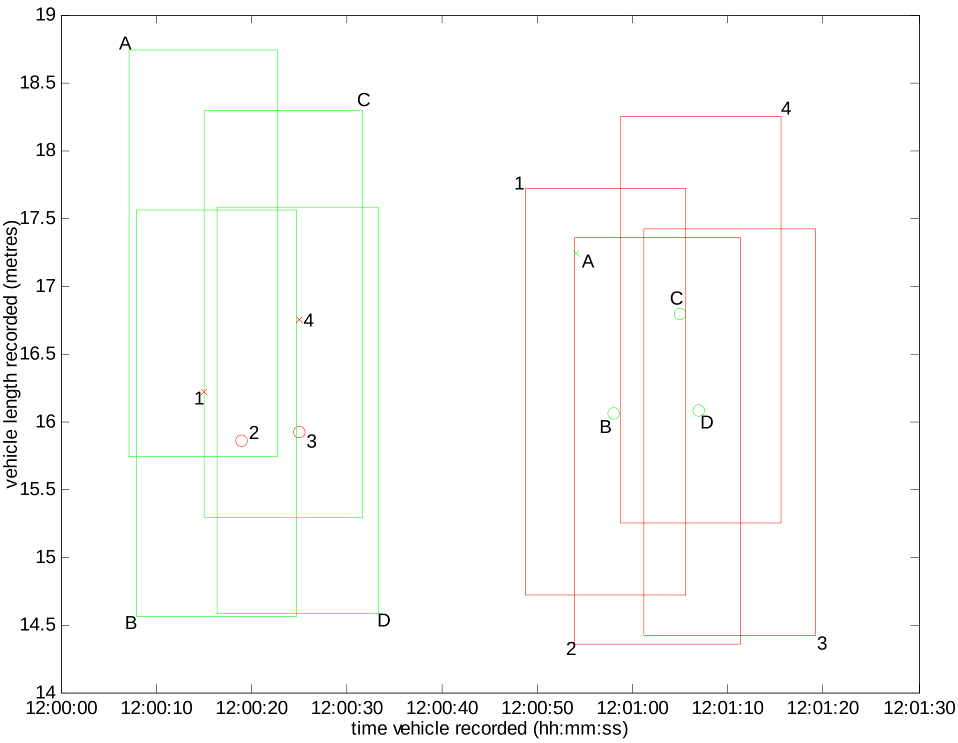

Figure 2. Vehicle records and forecast windows at two consecutive sites 1 km apart

Figure 2 illustrates how the measured length and timestamps are used to create possible matches at the associated upstream/downstream loop. The leftmost group of circles (vehicles recorded in lane 1) and crosses (vehicles recorded in lane 2) labelled 1,2,3,4 represent vehicles recorded at the upstream site and the group of rectangles labelled 1,2,3,4 towards the right of the figure represent the forecast window inside which possible matches at the downstream site must lie. Likewise, the leftmost group of rectangles A,B,C,D represent the back-forecast window inside which possible matches must lie for the respective downstream records. The size of the forecast window is determined by the length error tolerance (±1.5 metres), and the timestamp tolerance (±20% deviation in average speed between the loops, compared to what was measured). By determining whether each vehicle lies within a given forecast window, possible downstream matches of 1 are A,B,C; of 2 are A,B,C,D; of 3 are C,D; of 4 are C,D. Likewise, the possible upstream matches of A are 1,2; of B are 1,2; of C are 2,3,4; of D are 2,3,4.

Potential matches are only considered further if both the upstream and downstream matches agree. Taking this into account, this then reduces the problem to matching vehicle 1 with A,B; 2 with A,B; 3 with C,D; and 4 with C,D.

Given the forecast windows presented in Figure 2, it is possible to compose a *scoring matrix* which corresponds to the likelihood of a match between the upstream and downstream records. Suppose we take the score for a match to be dependent on the proximity of each record to the centre of the corresponding forecast window. Table 1 shows a possible scoring matrix for the situation presented in Figure 2. A higher number represents a more acceptable match. For the purpose of this example the scores have been approximately calculated and do not correspond to any rigorous geometrical analysis. In practice the scores would be issued in a much more precise manner.

		Downstream records			
		A	B	C	D
Upstream records	1	5	6	0	0
	2	2	5	0	0
	3	0	0	4	6
	4	0	0	6	4

Table 1. Possible Scoring Matrix of upstream and downstream matches

From this matrix, it can be seen that the best solution is found by matching 1 with A, 2 with B, 3 with D, and 4 with C. Under this matching scenario, vehicle 4 has been undertaken by vehicle 3, and vehicle 4 has subsequently pulled in to lane 1.

For much larger matrices, the optimal pairing strategy may not be obvious, particularly when all vehicle types are considered (the example shown above only considers matching vehicles over 14 metres in length). An example of a much larger scoring matrix is given in Figure 3, where the colours represent the likelihood of a match.

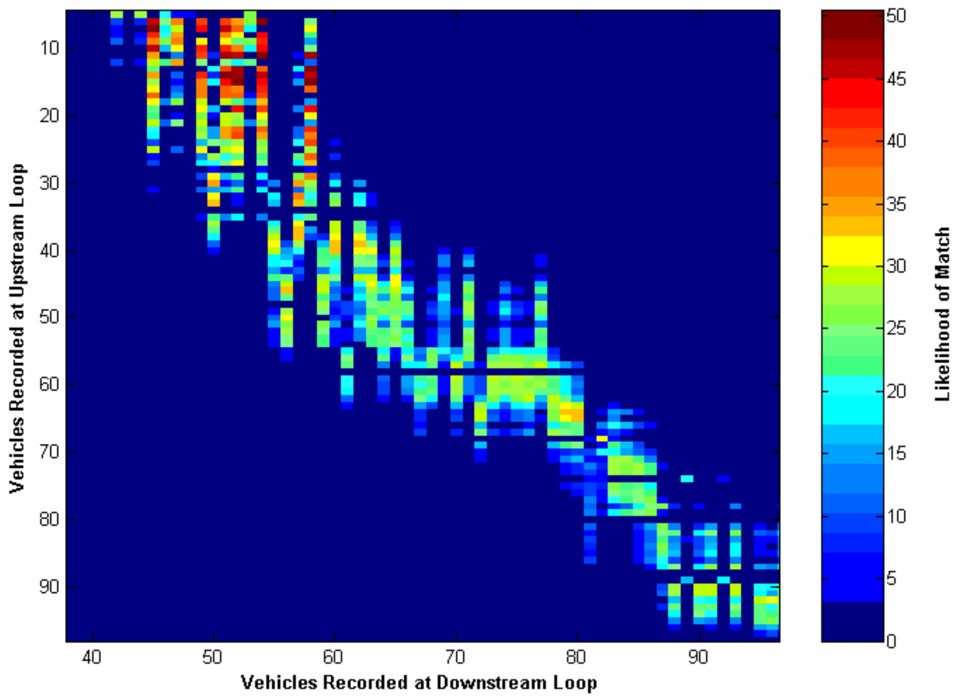

Figure 3. Scoring matrix for 100 vehicles recorded at consecutive sites 1 km apart

The formal optimisation problem for this particular situation is known as *weighted bipartite matching* (as in mathematical terms the graph we seek is a bipartite graph with weighted vertices). The process can be likened to finding the best one-to-one match of employees to jobs within a working environment, or best playing positions for footballers on a team. The common algorithm used to solve this problem is known as the *Hungarian Algorithm* [3].

The method described above is very effective when matching long vehicles. However, the primary application of this algorithm is to test and calibrate microscopic traffic models for all vehicle types. To do this, there should be no bias associated with the types of vehicles re-identified. The next section describes how this is possible.

Vehicle Re-Identification – Further Development

Original work described in the previous section did not provide a satisfactory correct re-identification rate for shorter vehicles. Methods described in this section are built on the work originally performed, and provide encouraging results.

An additional set of IVD data was collected that recorded individual vehicle statistics at intervals of 100 metres on the M42 (close to Birmingham). A schematic of these six sites is shown in Figure 4. This distance between sites will be the standard spacing on the M42 motorway (close to Birmingham), as it is a requirement for the operation of ATM (Active Traffic Management [4]). This provides a much finer level of detail and a much more acceptable distance over which to re-identify vehicles. Previous data was only available at an interval of 1 km.

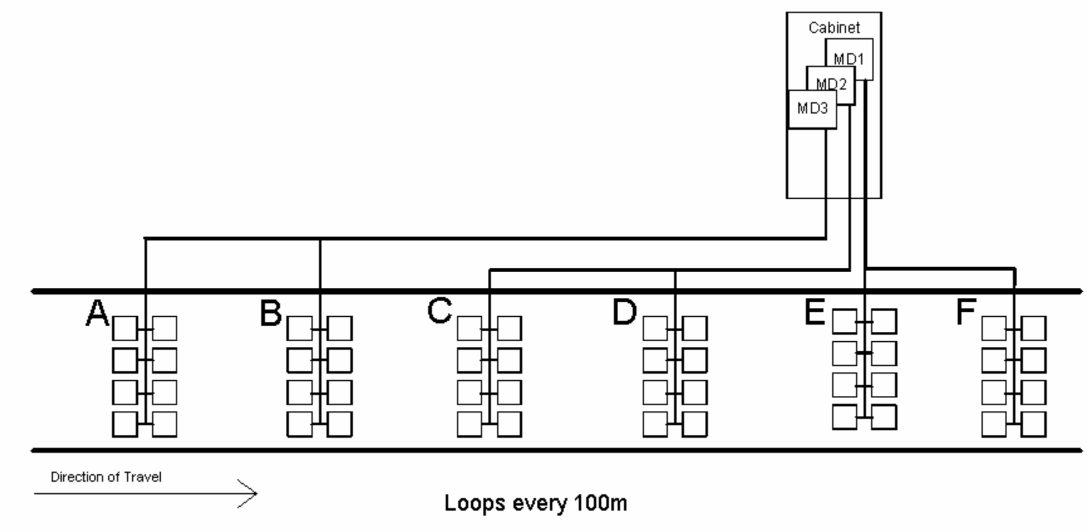

Figure 4. Schematic of IVD collection site

Figure 5 plots this IVD data recorded at all six sites in a unique way. The six sites are labelled A,B,C,D,E, F ordered from the most upstream to the most downstream (so site F is the most downstream). Each site is represented by a new axis, or *subplot*. For each site, the vehicle information recorded in any of the three lanes is displayed, as described below.

The six separate axes (or subplots) represent the IVD recorded at the six different sites, as labelled at the bottom of each horizontal axis. A vehicle record is depicted as a filled rectangle. The timestamp of a vehicle is represented by the top-most point of each filled rectangle. The height of each rectangle is a measure of the vehicle's length. The lane the vehicle was detected in is represented by the horizontal position of each filled rectangle. The number by each rectangle is the speed of the vehicle recorded in kph. The initial time for each subplot (represented by the top most point on each subplot) has been continuously offset by 3 seconds in order keep the same vehicles 'in view' for each plot. However, the only timescale given is that of the first subplot, as written on the leftmost axis.

The plot can be thought of as an overhead view of the motorway, with each subplot a snapshot at a slightly different location and time. The filled rectangles represent vehicles on the motorway, travelling in an “upwards” direction. Bearing this in mind, it is thought possible to manually track each of the vehicles by hand from one sub-plot to the next, as described below.

Figure 6 shows what is believed to be a matched set of vehicles across all six sites. Vehicles surrounded by rectangles represent vehicles that have not changed lane or have not had their ordering changed across all the sites. For this example, this is all the vehicles bar one. The one vehicle that does change lane is circled, and is recorded as being in lane 2 at the 1st site, lane 3 at the 2nd, 3rd, 4th and 5th site, and then in lane 2 at the 6th site.

On producing these plots for different times within the recording period, the vehicles can be re-identified with the same level of confidence as performed on the example shown in Figure 6. Current research is at a stage where a user can use simple point and click actions on a screen to submit matches to the computer. This hand-matched set will help provide a preliminary test for any algorithms developed in the future.

The details of such algorithms have yet to be determined. Experience of matching the vehicles manually has shown that re-identification can be performed by comparing the headways of the vehicles from one loop to the next. If an unexpected headway is found then it is likely that a lane change has occurred, and headways are compared again using this new assumption.

The matching method described in this section is different to those described in the previous section, and to any methods described in the literature. (Furthermore, no methods described in the literature [5, 6] use the Hungarian algorithm or relate the problem to a bipartite graph). The matching method described here uses as the primary vehicle signature the vehicle headway as opposed to the vehicle length.

Concluding Remarks

This paper has illustrated the need for more microscopic traffic data, and has presented the means of developing an algorithm to help nourish this void. Although the algorithm is incomplete, preliminary research has suggested that close to 100% of vehicles can be re-identified.

References

[1]
Gipps P.G. (1981). *A behavioural car-following model for computer simulation*. Transportation Research-B 15 pp105-111

[2]
Lunt G.M., Wilson R.E. (2003). *New Data Sets and Improved Models of Highway Traffic*. UTSG 35th Annual Conference, Volume 2

[3]
Galil Z. (1986). *Efficient Algorithms for finding maximum matching in graphs*. ACM Computing Surveys (CSUR) Volume 18, Issue 1

[4]
Highways Agency. 2005. *Active Traffic Management (ATM) Project M42 Junctions 3A-7*
http://www.highways.gov.uk/knowledge/tcc/atm/index.htm

[5]
Coifman B. (2002). *Estimating travel times and vehicle trajectories on freeways using dual loop detectors*, Transportation Research-A 36 pp351-364.

[6]
Sun, C., Ritchie, S., Tsai, K., Jayakrishnan, R. (1999). *Use of Vehicle Signature Analysis and Lexicographic Optimization for Vehicle Reidentification on Freeways*. Transportation Research-C, vol 7, pp 167-185.

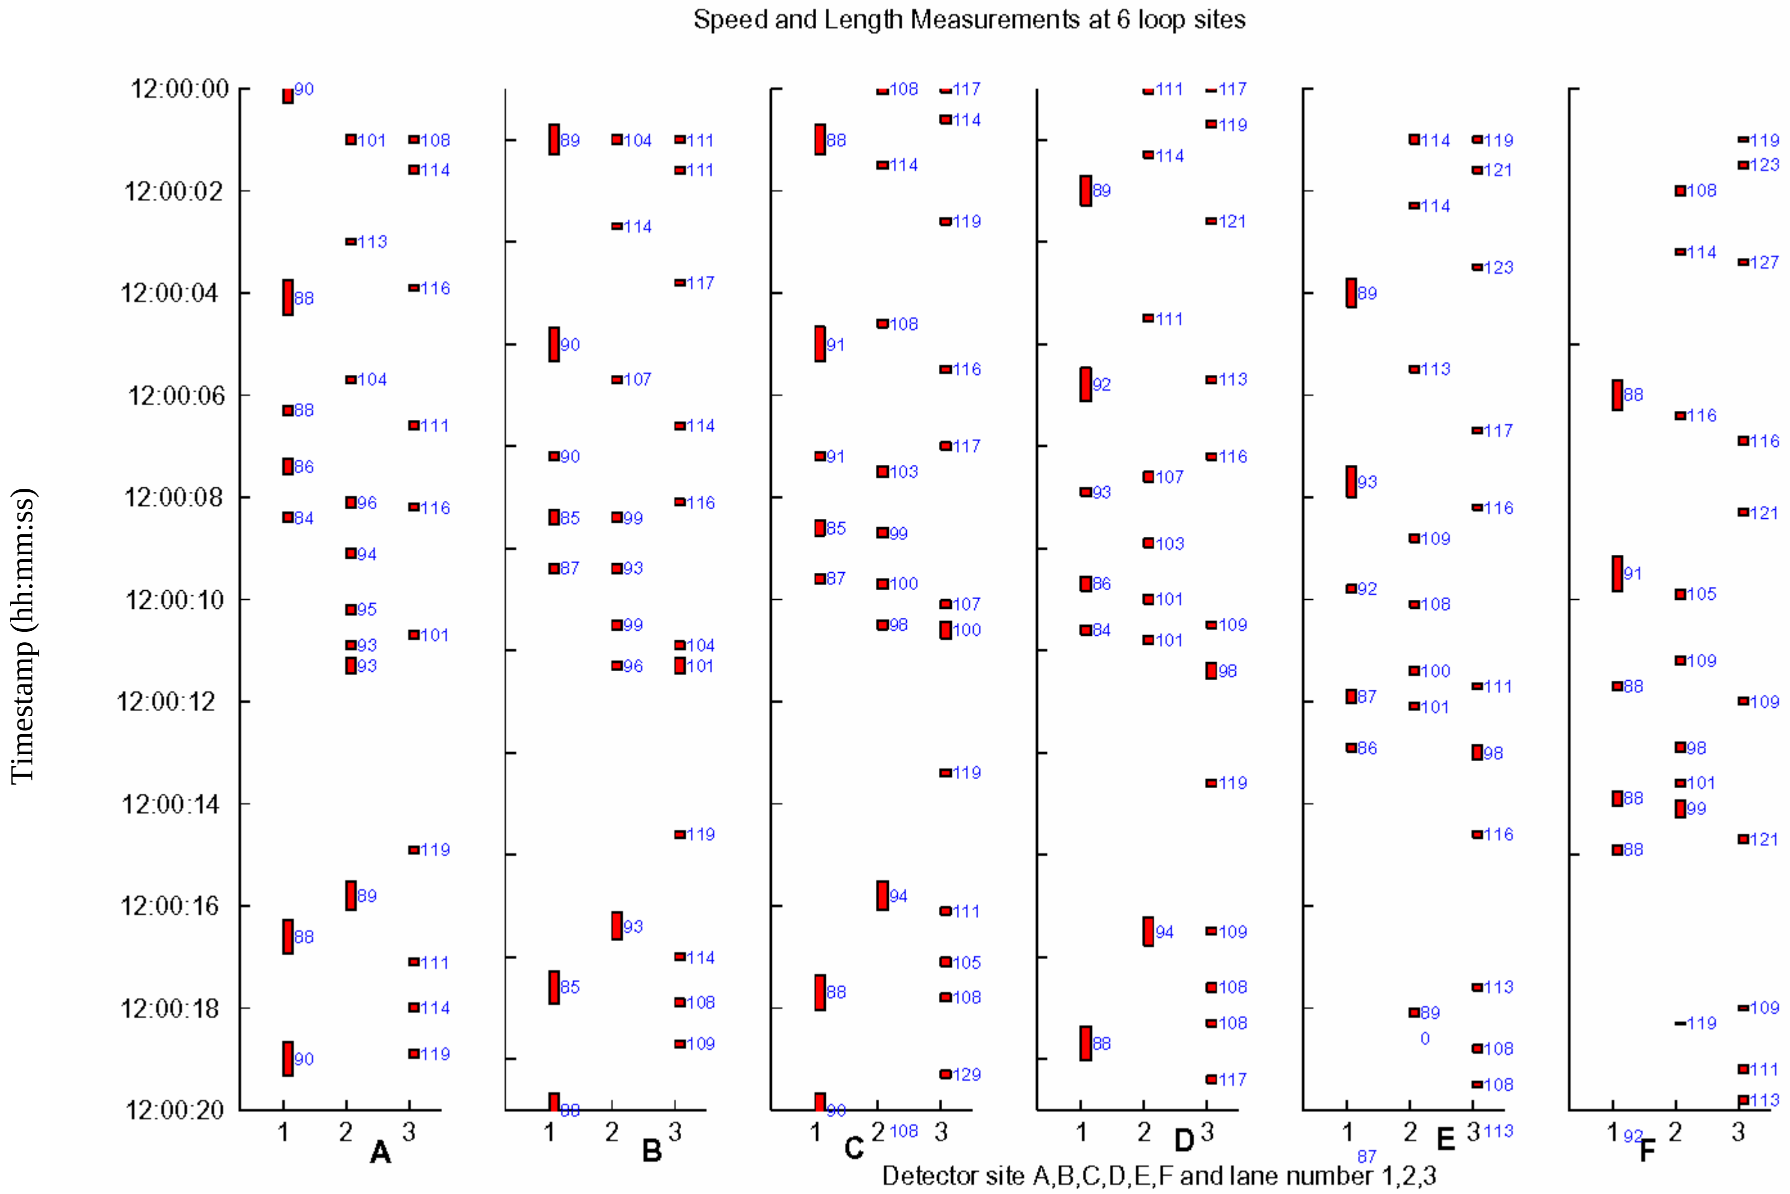

Figure 5 Vehicle Records at six consecutive loop sites, 100 metres apart

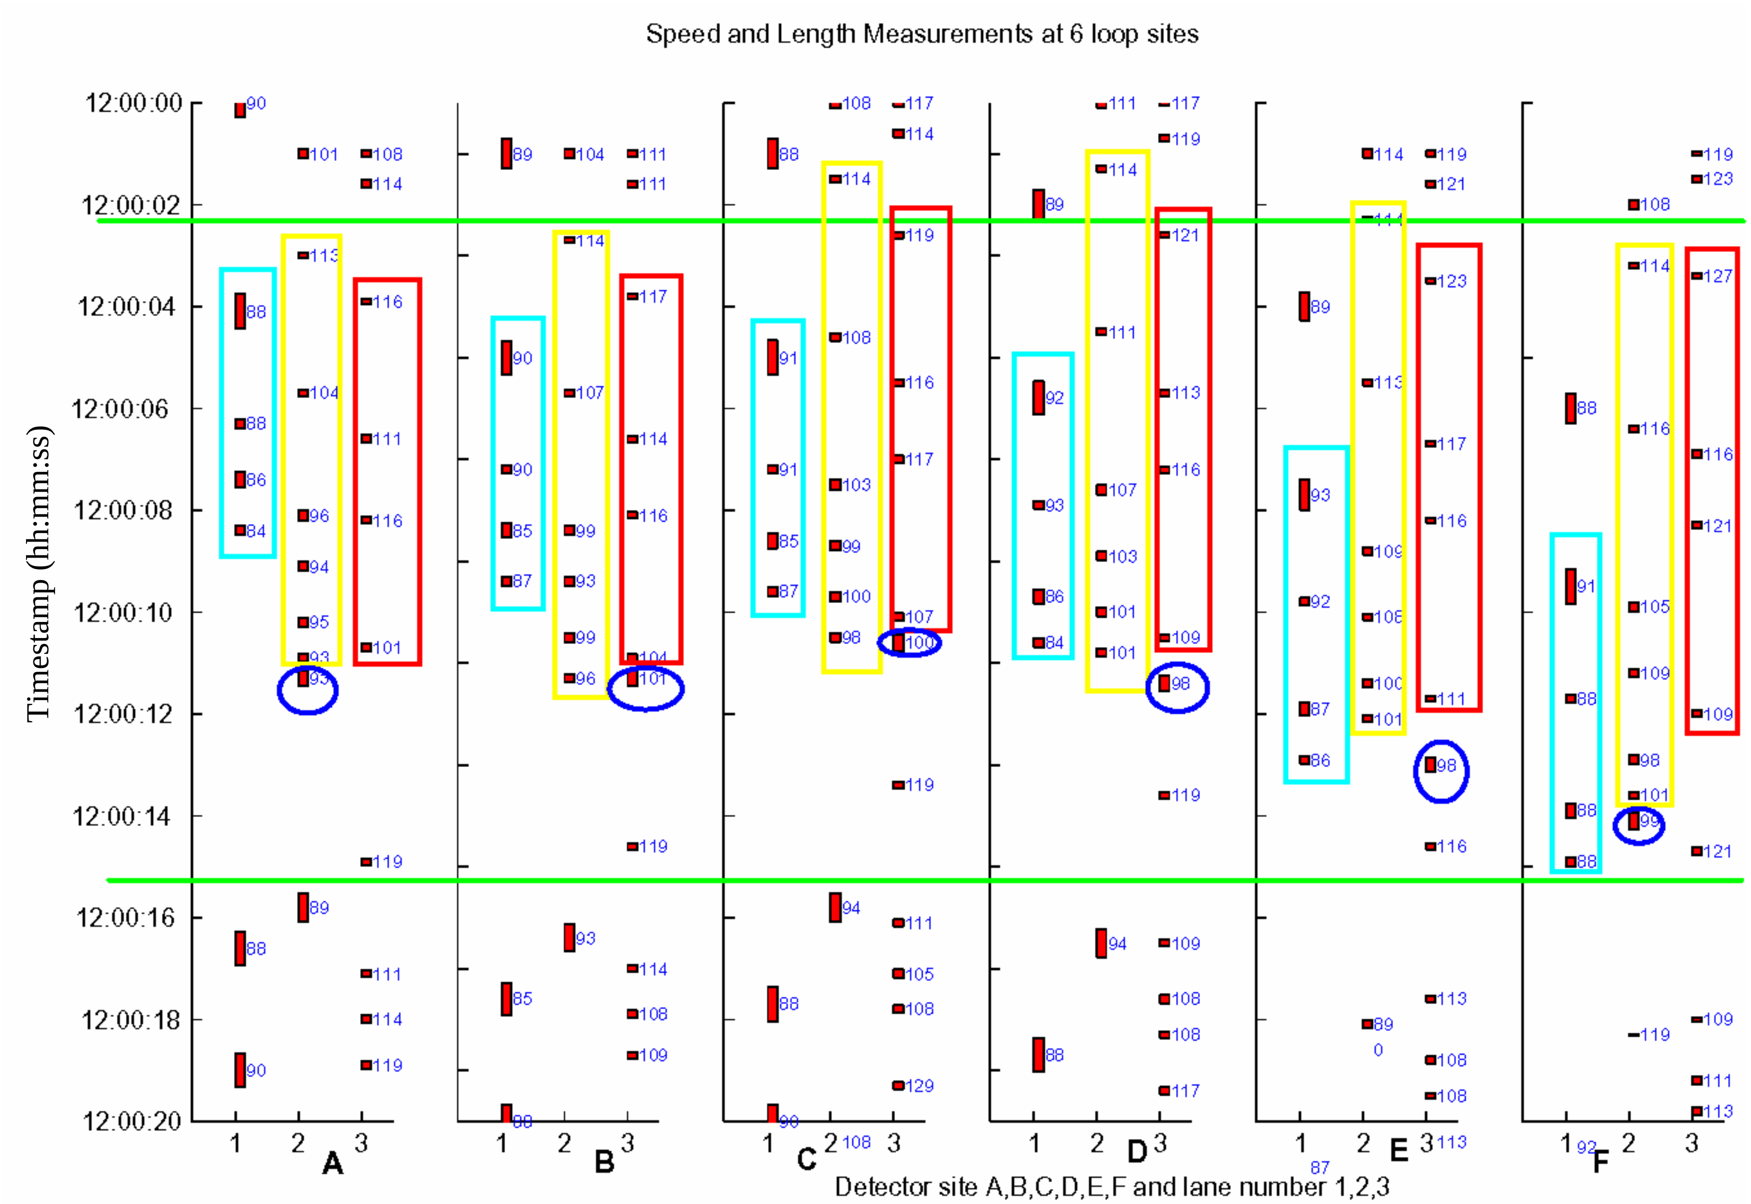

Figure 6. Vehicle records and possible matching at six consecutive loop sites 100 metres apart. The vehicle circled blue has performed two lane changes over the 500 metre stretch