Quantifying the impact of probe vehicle localisation data errors on signalised junction control

B. Waterson   S. Box

Transportation Research Group, University of Southampton, Highfield, Southampton SO17 1BJ, UK
E-mail: B.J.Waterson@soton.ac.uk

Abstract: This study investigates theoretical signal control algorithms based solely on probe vehicle data. Through development of a simulation system that can model urban signalised junction control using localisation probe data from all vehicles in the local area, improvements in junction operational efficiency that result from the improved input data are demonstrated for both isolated and coordinated junctions. Results from the isolated junction scenario show that the richness of the information contained within probe vehicle data means that control algorithms based just on positions and velocities of vehicles can produce 25% reductions in average delay compared to the current standard control algorithm MOVA. Results from the twin junction scenario confirm the importance of using high-level synchronisation to coordinate closely connected junctions, achieving reductions in average delays (compared to independent control approaches) of up to 40% through a process of weighting the probe vehicle data to reflect prior stage decisions of other parts of the junction. Critical to achieving these benefits, however, is the availability of high localisation accuracy probe data, with results indicating that the levels of accuracy necessary are representative of the typical performance of current in-vehicle global positioning system units, except when those vehicles are operating in urban canyon environments.

1 Introduction

Recently, a number of large European Commission-funded projects (CVIS [1], SafeSpot [2] and Coopers [3]) have focused on the development of technologies and standards for vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication systems. This has led to common European protocols being set for this type of communication [IEEE 802.11 (WAVE) and IEEE 802.11p]. Some of the most important data that may be communicated between vehicles and infrastructure are localisation data, that is, dynamic estimates of the vehicle's position. Localisation technologies that can provide these data such as global positioning system (GPS) receivers are already commonplace in many vehicles, in use for navigation.

Urban signalised junction control is a task that requires sensors to monitor the state of the network, a processing system to analyse sensor data and make control decisions and traffic lights to implement the control. Currently, sensors that are commonly used in signalised junction control are inductive loops [4], microwave emitter/detectors [5] and traffic monitoring cameras. Examples of automated control algorithms that are currently in use to process data from these sensors and set signal timings in a demand-responsive manner are MOVA [6] for isolated junctions and SCOOT [7], which can coordinate multiple connected junctions. All the sensor types currently used in urban signal control, however, collect census data, that is, counts of vehicles passing a specific point in space. The type of data that can be collected using on-board vehicle localisation sensors is

probe data and this different type of data can present a fundamentally different view of the state of the network [8]. Probe data allow an analysis of the system that tracks each vehicle individually and can provide a higher resolution of position data. Work that examines the use of V2I communications and limited localisation systems in signalised junction control is already under way, for example, the iBus project [9] which uses localisation systems on London buses to give them priority at signalised junctions.

This paper investigates theoretical signal control algorithms based on vehicle localisation data alone. Through the development of a computational simulation system that can model the (currently hypothetical) scenario of urban signalised junction control using localisation probe data from all vehicles in the local area, the improvements in junction operational efficiency that result from the improved input data are demonstrated in both single (isolated) and multiple (coordinated) junction situations. This allows the impact on algorithm performance of errors in localisation accuracy and reductions in the proportion of vehicles for which probe data are available to be investigated and an assessment made of the required levels of data availability necessary to implement such a signal control system.

2 Simulating probe data

2.1 Simulator architecture

Fig. 1 shows the architecture of the simulation test bed developed in this research. At the centre of this is a module

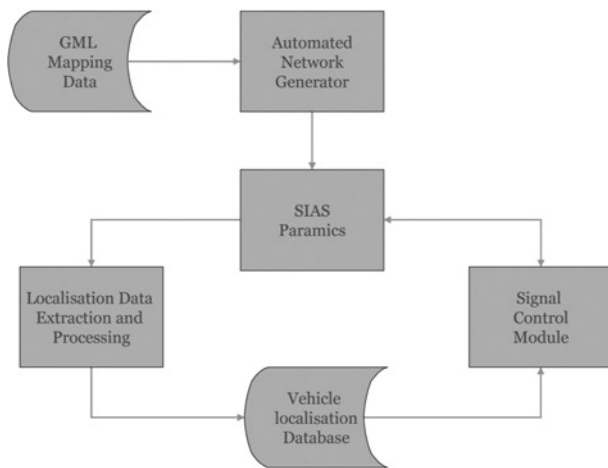**Fig. 1** *Block diagram showing the simulation software architecture*

for simulating vehicle movements and interactions through signalised junctions at the individual vehicle level (microscopic). The approach used in this research was to use an existing commercially available microscopic traffic simulator (SIAS Paramics [10]) to fill the roll of this module. Although this allows both the rapid development of test scenarios and rich graphical visualisations of the algorithm in operation, care must be taken with this approach that control strategies developed in the test bed are not too highly tuned simply to the behaviour of vehicles in Paramics that may not be completely representative of the real world. Although this paper demonstrates the potential of the new algorithms therefore on scenarios based on real-world situations, any strategies developed for specific junctions will require real-world validation to confirm their efficacy.

As can be seen in Fig. 1 the simulation test bed was completed through the development of three additional bespoke modules that are built around Paramics. The 'network generator' module [11] is used to automatically encode the structure of road networks in Paramics using a database containing mapping data. The 'localisation data extraction and processing' module continually interrogates the Paramics simulation to obtain localisation data for all the vehicles in the simulation and stores the processed data in a database. The 'signal control' module extracts the relevant localisation data from the database and uses this to make decisions about signal control, with the decisions then implemented directly in the Paramics simulation. The Paramics module, the localisation module and the signal control module are all synchronised to allow real-time simulation of signal control using localisation probe data.

2.2 Sources of localisation data

There are a number of existing on-board vehicle technologies that can provide dynamic data on vehicle position. These include mobile telephone, or cellular network localisation [12], GPS [13], inertial measurement systems (IMU) [14], laser range-finding systems (LIDAR) [15] and computer vision systems [5]. In addition to these hardware technologies other software technologies can be used to improve localisation estimates. These include map matching software, which constrains the vehicle's position to the road network [16] and Bayesian recursive filtering techniques,

such as the Kalman filter [17]. The latter allow data from more than one sensor and data from other sources such as dynamic data and vehicle control data to be fused to provide a probabilistic estimate of position. The performance of localisation systems is a function of positioning accuracy, frequency of position measurements and reliability (e.g. latency), with high levels of accuracy only usually achievable through (high cost) configurations which utilise more than one localisation sensor in tandem.

2.3 Localisation data accuracy

As the localisation module records vehicle position and velocity data directly from Paramics, these data are perfectly accurate at the time of sampling. This is unrepresentative of the data that would be obtained from a real localisation system. Therefore the localisation module must process these data to make them more realistic. To this end the localisation vector describing the vehicle's longitudinal position, lateral position (lane) and velocity obtained from Paramics is made stochastic by the addition of Gaussian noise to each element. The variance of the noise is chosen to be representative of the performance of a given localisation system, and thus the performance of signal control systems can be tested for different localisation systems and different levels of localisation performance. Within this paper a range of localisation accuracies are tested, ranging from standard deviations in the range $1-2$ m [representative of very good differential GPS (DGPS) in open areas], through standard deviations of $4-8$ m (representative of accuracies from current standard GPS units), to standard deviations of $16-32$ m (representative of GPS systems operating in urban canyon environments).

For the longitudinal position accuracies a random number drawn (for each vehicle at each time point) from a $N(0, \sigma^2)$ distribution is added to the exact distance between the vehicle and the junction stop line. For lateral positioning the algorithm simply needs to know which lane the vehicle is in. To reflect lateral positioning error in the analysis a random number to represent the lateral error is drawn from a $N(0, \sigma^2)$ distribution for each positioning estimate. This value is compared with the lane width and, if the lateral error is greater than half the lane width, the lane that the vehicle is adjudged to be in is changed accordingly. The critical issue, in terms of algorithm performance, is whether the position of the vehicle can be resolved accurately enough to place it in the correct lane with a high probability. For low standard deviations ($\sigma \le 2$ m) the probability of the vehicle being placed in the correct lane is high (hence this level of accuracy can be referred to as 'lane resolution'), but for higher standard deviations the probability of incorrect lane placement is high with consequential impacts on the performance of the signal control algorithms.

In reality the velocity measurement for each vehicle would be calculated on-board from a series of recent position estimates, each with their own error distribution. To replicate this situation within the simulation model, however, where exact measurements of velocity are available, it is necessary to reduce the accuracy of velocity measurements in a way similar to the addition of errors for positional data. For each observation therefore a random number drawn from a $N(0, \sigma_v^2)$ distribution is added to the actual velocity, where $\sigma_v^2 = 2\sigma^2 t^{-2}$ relates the error variance for velocities to the variance used for positions
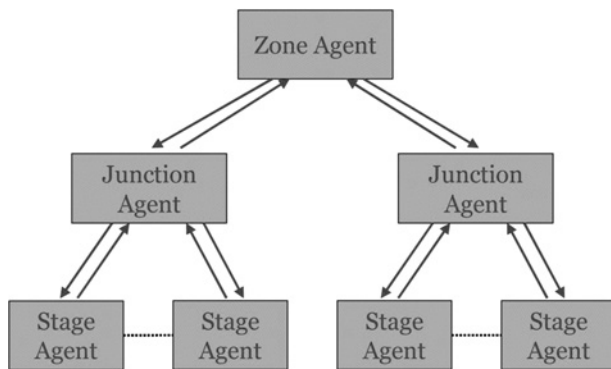
Fig. 2 *Structure of the agent hierarchy tree*

(σ^2) and the nominal sampling rate (t) of the simulated measuring system.

2.4 Signal control module

As the signalised junction controller uses localisation probe data from all vehicles in the local area it may have to process significant amounts of data to set signal timings. Previous research on signal control strategies, where a large amount of loop data need to be processed, has demonstrated the advantage of a hierarchical 'agent' structure [18]. Here individual software agent's process small amounts of the raw data, which they then pass on in a significantly refined form to another agent above them in the hierarchy. This research has adopted an agent hierarchy very similar to the one presented in [18], the structure is shown in Fig. 2. The lowest-level 'stage agents' receive the vehicle localisation data relating to vehicles whose approach relates to a single signal stage only. These data are refined by the stage agents into a simplified form which constitutes a bid for priority. These bids are received by the junction agent, which will then assign priority to the stage with the winning bid. In a situation where a number of closely connected signalised junctions need to coordinate signal timings the junction agent will communicate with a zone agent above them in the hierarchy before assigning priority.

3 Isolated junctions

3.1 Specification of algorithms

This paper compares two prototype stage agent algorithms, which differ in the level of localisation detail about each vehicle approaching the junction that is included in the bid for priority.

Bidding algorithm 1: Each stage agent (s) has a set (N_s) of vehicles to consider, consisting of all vehicles that would receive a green light if stage s achieves the winning bid (determined on the basis of their localisation). This set of vehicles will therefore likely be different each time a stage agent makes a bid (especially if it previously had the winning bid whereupon some of the vehicles forming the previous set would now have passed through the junction and no longer be considered) and it should also be noted that the sets are not exclusive, with potential for a vehicle to be contained in (e.g.) both the set for a 'straight ahead' stage and the set for a 'straight ahead and right turn' stage. To calculate its bid the stage agent simply counts the number of vehicles in the set N_s that are stationary (or more technically those that have a velocity below a threshold

value of V_t) and this count is used as the bid

$$\text{Bid}_s = \sum_{i \in N_s} \begin{cases} 1 & \text{if } V_i < V_t \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Bidding algorithm 2: In this case the bid is calculated as a linear function of the number of vehicles in the set N_s, the speed of each vehicle (V_i) and the distance of each vehicle from the junction X_i as in (2), where α and β ($\alpha = 0.01$ and $\beta = 0.001$ in the results presented in this paper) are coefficients that can be tuned for individual junctions [19]

$$\text{Bid}_s = \sum_{i \in N_s} (1 - \alpha V_i - \beta X_i) \quad (2)$$

Having received bids from all stage agents the junction agent simply needs to select the stage with the highest bid and assign priority, as this is an isolated junction there is no need for the junction agent to communicate with a zone agent higher in the hierarchy. To avoid changing the stages too rapidly the junction agent performs these auctions only at a fixed time interval ($\delta t = 10$ s in the results presented in this paper) called the auctioning rate [19]. Both of the stage agent algorithms tested use vehicle speed in the calculation. It should be noted here that there are many practical considerations in signal control which are not reflected in this theoretical design, but which would need to be satisfied in an equivalent 'production' system. For example, the metrics used do not account for the scenario of a single vehicle being held at a red light indefinitely (something i.e. undesirable in the real world but possible under the bidding system used here) and there are other interesting metrics (such as CO_2 emissions or differential priority by vehicle type) which it is desirable to optimise but which are not considered in this paper.

3.2 Test scenario

Simulation tests were carried out on the simple isolated T-junction example illustrated in Fig. 3, where the east–west road is two lanes in each direction and the joining north–south arm is a single lane in each direction. The T-junction has three signal stages: stage 1 gives priority to vehicles approaching from the west and east arms of the junction, stage 2 is a right turn priority stage coming from the west and stage 3 gives priority to vehicles from the south. It should be noted that because of the opposed turning movement, vehicles approaching the junction from the west in the right-hand lane only contribute to the bid for stage 2 (where their turning movement is not dependent on a gap in the flow of vehicles from the east), whereas those in the left-hand lane contribute to the bids for both stages 1 and 2.

To provide a baseline for the tests of the new algorithms, loop detectors have been included in the simulated T-junction shown in Fig. 3 using the facilities included in Paramics. This enables a test to be carried out where the junction was controlled by the MOVA algorithm [6]. The loop detectors used have separate sensing loops for each lane, as with the more accurate probe localisation systems, MOVA has lane resolution within the input data.

Each test covered a simulated 1 h period, during which the level of demand was constant with the demand matrix (vehicles per minute) as shown in Table 1. The overall levels of demand were set so that the junction was performing slightly below its theoretical saturation point,
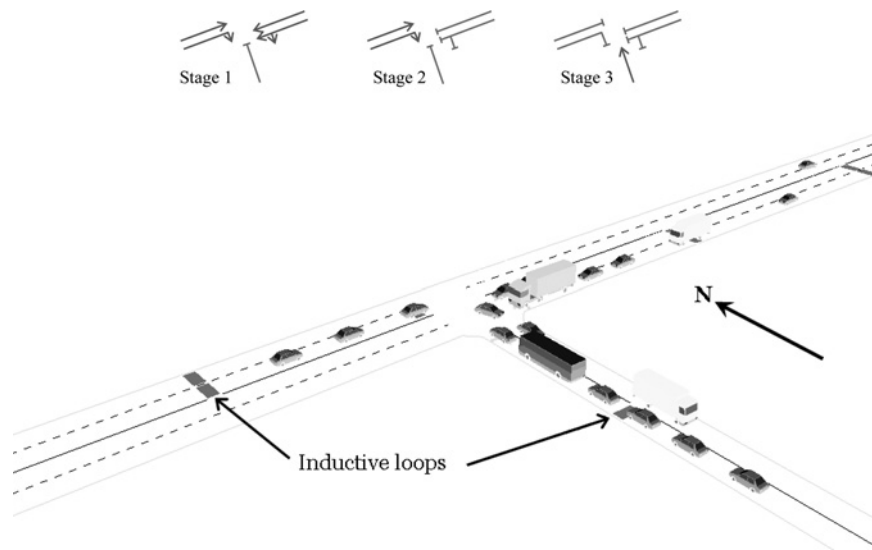
Fig. 3 *Isolated Paramics T-junction scenario*

Table 1 Demand matrix for isolated junction tests

	West	East	South
West	–	12.5	3.3
East	15.8	–	0.8
South	2.7	2.7	–

with the precise arrival distribution of the vehicles being generated by Paramics.

3.3 Comparative algorithm performance

Fig. 4 compares the performance of bidding algorithm 1 (BA1) and bidding algorithm 2 (BA2) across a range of positioning accuracies. The delay value calculated is the travel time (in seconds) through the junction for each vehicle minus the corresponding free-flow travel time for that turning movement, averaged across all vehicles in ten independent modelled periods.

These results clearly show the benefits of using probe vehicle data over census (inductive loop) data with the equivalent mean delay value for MOVA being 20.3 s (although it should be noted that this is based on the assumption that MOVA receives perfect data from the inductive loops and is therefore actually a lower bound on the true performance). At low levels of positional error ($\sigma \leq 4$ m) the BA1 and BA2 algorithms perform very similarly and consistently better than
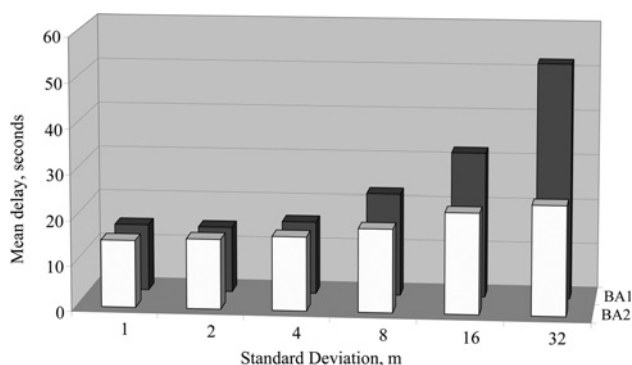

Fig. 4 *Mean delay values for isolated junction algorithms*

MOVA, achieving over 25% reduction in delay per vehicle comparing MOVA and the BA2 algorithm with a standard deviation of positioning accuracy of 1 m). At higher levels of positional error Fig. 4 shows that the performance of the BA1 and BA2 algorithms diverge, with BA2 being substantially less affected by the errors as it is based on the richer dataset of both positions and velocities, and by $\sigma = 16$ m for the BA2 algorithm ($\sigma = 8$ m for BA1) the overall performance is similar to MOVA.

4 Coordinated junctions

4.1 Test scenario

The twin junction scenario in Fig. 5 has two signalised intersections, each of which has the same staging as the isolated T-junction scenario in Fig. 3 above. The simulated distance between the two junctions is 55 m, giving a maximum internal storage capacity (for the two lanes in total) of approximately 30 passenger car units. Similar to the demand levels used for the isolated junction, demand levels for this scenario were set so that the junction as a whole was performing only slightly below its theoretical saturation point. It should be noted, however, that with the precise arrival distribution of the vehicles being generated by Paramics this leads to the junction demand actually varying between periods of slightly over and under capacity.

4.2 Specification of algorithms

Two approaches to operating the two junctions are possible. The first approach ('independent control') is where the two junctions are operated separately using the BA2 algorithm described above. As the bidding algorithm is based on the positions and velocities of vehicles approaching the junction on each arm this causes inefficiencies as the algorithm for each junction can only consider vehicles as far upstream as the other junction.

Alternatively, 'synchronised control' is achieved through invoking the zone agent in Fig. 2. The junction agents pass their independent preferred stage choices to the zone agent along with the bid from the corresponding stage agent. The zone agent then compares these bids and assigns the junction with the highest overall stage bid to be the 'lead
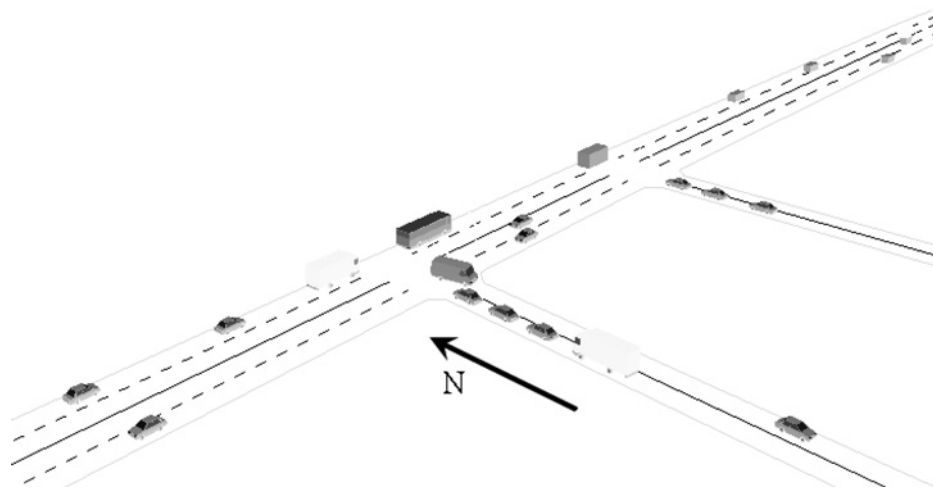
Fig. 5 *Twin Paramics T-junction scenario*

junction' and allows it to set its choice of stage, essentially assigning priority to the junction with the highest need in the same way that the junction agents assign priority to the stage with the greatest need.

The losing junction then has to revise its individual stage bids to reflect the additional knowledge of the stage selected by the lead junction. This process is done through weighting the current bids to approximate a scenario where the losing junction is treated as an isolated junction (running the BA2 algorithm) and the bids represent the changed road structure created by the lead junction stage being chosen. For example, suppose that the eastern junction is chosen by the zone agent to be the lead junction; three possible scenarios exist dependent on the precise stage chosen by the eastern junction as in Fig. 6, where B_{W1} and RB_{W1} represent the original and revised bids, respectively, by the stage agent for stage W1 etc.

The weighing coefficients (C_1 and C_2 in Fig. 6) must be tuned to reflect the expected turning proportions from upstream of the lead junction, as not all traffic would usually be expected to reach the losing junction. It should be noted, however, that if the localisation accuracy of the probe data is sufficient to allow lane resolution then vehicle turning movements can be assessed more accurately for many combined junction layouts and hence the bids adjusted more accurately. Rules similar to those in Fig. 6 are applied if the zone agent selects west as the lead junction, with RB_{E2} needing to be adjusted through appropriate weights.

4.3 Comparative algorithm performance

Fig. 7a shows the benefit of the synchronisation process being carried out by the zone agent, with the performance of the synchronised control algorithm outperforming the independent control algorithm at all levels of positional error. What is also apparent, however, is that vehicles using the junction are experiencing delays on average in excess of 1 min and that, even averaged over ten simulation runs, there remains substantial variability in the results when compared across different levels of positional error. Both these issues are related to the total demand for this scenario being only slightly below the saturation point of the junction.

This demand was simulated because it is at this level of saturation that the benefits of synchronisation are most apparent, but a detailed analysis comparing between the

individual simulation runs, however, reveals that in some cases the algorithms lead the junction operation into a completely congested situation from which they cannot recover. If the averages are recalculated using only the best eight performances from each series of ten simulation runs to exclude instances when this situation occurs, a more consistent pattern of algorithm performance is evident in Fig. 7b.

It is clear from Fig. 7b that the synchronised control approach provides substantial improvements in junction performance, achieving reductions in average vehicle delay of between 25% (at $\sigma = 1$ m) and 42% (at $\sigma = 32$ m). The synchronised control algorithm performance is also more robust to higher levels of positional error with average delays increasing by 10% (when the standard deviation of positional accuracy increases from 1 to 32 m), compared to an increase of 41% for the independent control algorithm. Although this demonstrates the theoretical capabilities of the synchronised control algorithm, it must be noted that the combined effect of the positional errors and near saturation demand levels leading the algorithms to push the junction into a completely congested regime. This highlights a key issue in the deployment of such algorithms in that it suggests probe vehicle data based on current GPS localisation accuracies alone will be insufficient for optimal performance of the algorithms.

5 Probe vehicle proportions

Although the results above have shown the clear benefits of undertaking signal control based on probe data, with the proportion of vehicles in the traffic flow equipped to provide such data currently being at 0%, it is important to consider the performance of the system based on limited proportions of probe vehicles. To quantify this source of error the twin junction synchronised control scenarios were carried out keeping the standard deviation of positioning accuracy constant (at 2 m), but allowing the (randomly selected) proportion of vehicles for which probe data were available to range between 5 and 100%. The vehicles for which data were available were randomly selected for each simulation run and all other vehicles were assumed to be invisible to the algorithm.

Fig. 8 shows results averaged over ten simulation runs and gives a clear picture of the impact of reducing the proportion of vehicles for which probe data are available, with three ranges of algorithm performance being evident:

$$RB_{W1} = B_{W1} + C_1 B_{E1}$$
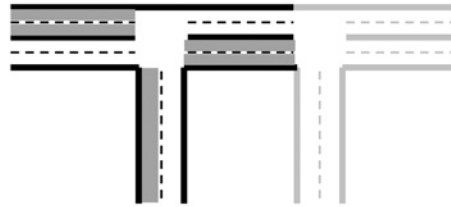$$RB_{W2} = B_{W2}$$
$$RB_{W3} = B_{W3}$$
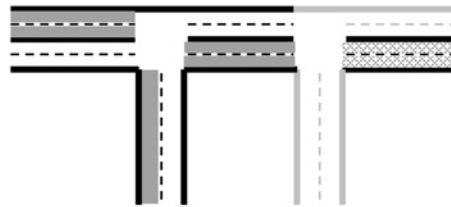
a

$$RB_{W1} = B_{W1}$$
$$RB_{W2} = B_{W2}$$
$$RB_{W3} = B_{W3}$$
(There is no need to adjust the bids in this case where the chosen east stage is leading away from the west junction)
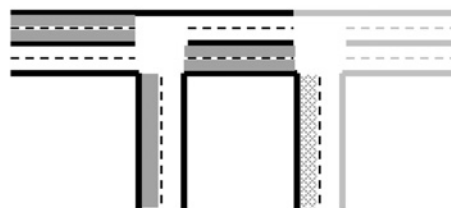
b

$$RB_{W1} = B_{W1} + C_2 B_{E3}$$
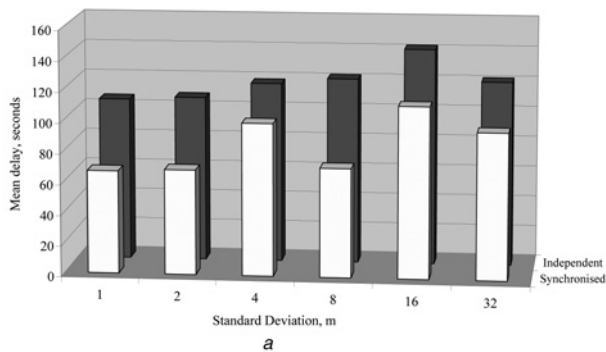$$RB_{W2} = B_{W2}$$
$$RB_{W3} = B_{W3}$$

c

Fig. 6 *Example revised bids during synchronised junction control*

a Eastern junction chooses stage E1
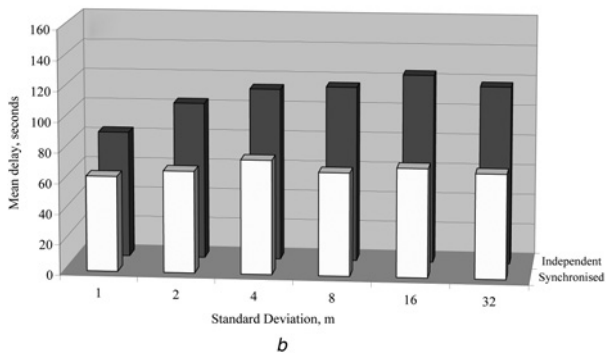*b* Eastern junction chooses stage E2
c Eastern junction chooses stage E3

1. Proportion ≥ 60%: In this range the impact on algorithm performance is minimal with the system still able to appropriately balance the conflicting demands and to appropriately synchronise the two junctions.

Fig. 7 *Mean delay values for coordinated junction algorithms*

a All simulation runs
b With over-saturated runs removed

2. Proportion ≤ 10%: In this range the volume of probe data is insufficient for the algorithm to operate at close to optimal performance, with consequential increases in the average delay. As with the high standard deviations of positioning accuracy above, considering the individual simulation runs it is noticeable that under these low proportions of probe vehicle data the algorithms again sometimes lead the junction operation into a completely congested situation (usually beginning at times corresponding to higher demand levels and lower equipped vehicle levels in the random arrival patterns of vehicles in the simulation) from which the junction cannot recover.

3. 20% ≤ proportion ≤ 40%: In this range the difference in performance of the algorithm between simulation runs is much higher. While the overall average delay falls between that of the other two ranges, a more detailed assessment of the results shows that the system performance varies over
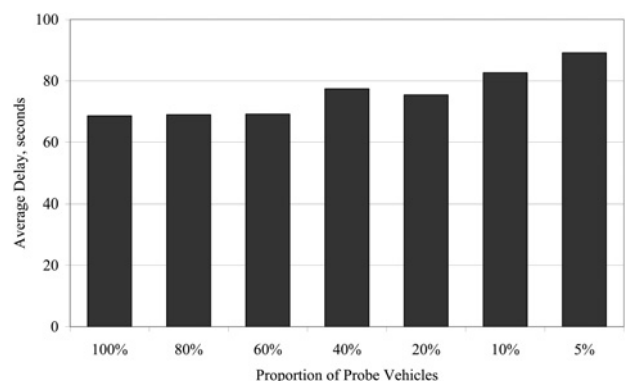

Fig. 8 *Mean delay values for differing probe vehicle proportions*

time as the number of (randomly selected) probe vehicles approaching the junction varies. In periods where the proportion of probe vehicles is low the system performs as for the low proportion range, but during periods of higher proportions the system is able to recover the situation and return the junctions to a more optimal operating state. The slightly higher average delay value for the 40% probe vehicle proportion than for the 20% proportion is a reflection of more instances of the algorithm leading the junction into an unrecoverable congested situation.

6   Conclusions

This paper has described new urban signalised junction control algorithms that make use of localisation probe data as their input source and quantified the impacts that errors in the probe data have on the performance of the algorithms.

The results from the tests on an isolated T-junction scenario have shown that the richness of the information contained within individual vehicle localisation probe data means that even a very simple control algorithm (BA1) that just accurately counts the number of stationary vehicles on each stage approach can produce better performance than the current standard control method using the MOVA algorithm. Adding vehicle velocities and distances from the junction to the bid calculation (BA2) enabled the performance to be improved further with over 25% reductions in average delay per vehicle (compared to MOVA) being suggested. Critical to achieving these benefits is the availability of low positional error probe data, with the results indicating that standard deviations of at most 8 m are necessary. This level of accuracy is representative of the current performance of standard in-vehicle GPS units, but not when they are operating in urban canyon environments suggesting that additional support to improve accuracy of vehicle positioning and velocity estimation would be necessary in these situations.

The tests on a twin T-junctions scenario confirmed the importance of using zone agents at the top of the agent hierarchy to coordinate junctions that are closely connected. This paper has shown that the simple auctioning process developed for isolated junctions can be easily adapted for zone agent control by allowing the zone agent to weight the bids for the individual junction agents to reflect known stage decisions and that this can reduce average delays (compared to the independent control approach) by up to 40%. As with the isolated junction scenarios, the level of accuracy in the input probe vehicle position and velocity data was again highlighted as an issue, but unlike for the isolated junction where increasing error standard deviations simply reduced the algorithm performance, these errors had less impact on the underlying synchronised control algorithm performance. Instead, these errors for the coordinated junctions scenario could cause the algorithms to lead the junction into a completely congested state from which it could not recover (at the constant high demand level), with more detailed analysis of this situation suggesting that the likelihood of this happening increases with increasing standard deviation in the localisation error (as it increases the probability of the algorithms making repeated poor signal stage decisions).

The final source of error considered was that related to the proportion of equipped vehicles, a key issue in the future deployment of probe vehicle-based signal control system where not all vehicles could be expected to be equipped in advance of system deployment. The algorithm performance was found to be robust to reductions in the proportion of equipped vehicles, suggesting that information on only a proportion of vehicles in the traffic stream is necessary to appropriately balance the signal stages and respond to the traffic demand. The issue of poor-quality input data causing the algorithms to lead the junction into an unrecoverable congested situation was again highlighted, with the impact of the random arrival nature of vehicles equipped or not equipped to provide probe data an additional factor, especially at low proportions of equipped vehicles.

This paper has therefore investigated the benefits of using probe vehicle data as a source for signal control algorithms and determined that while the potential benefits are substantial there remains an issue over the levels of accuracy in localisation data that can currently be achieved and the consequential impacts that this will have on the performance of the control algorithms.

7   Acknowledgments

8   References

1   CVIS – Cooperating for mobility. Available at http://www.cvisproject.org/download/cvis_brochure_May2008_Final.pdf, accessed July 2010

2   SAFESPOT Integrated Project. Available at http://www.safespot-eu.org, accessed July 2010

3   COOPERS. Available at http://www.coopers-ip.eu, accessed July 2010

4   Sreedevi, I.: 'Intelligent transport systems – traffic surveillance', available at http://www.calccit.org/itsdecision/serv_and_tech/traffic_surveillance/road-based/in-road/loop_summary.html, accessed July 2010

5   Wang, C., Hu, Z., Kusuhara, S., Uchimura, K.: 'Vehicle localisation with global probability density function for road navigation'. Proc. IEEE Intelligent Vehicles Symp., June 2007, pp. 1033–1038

6   Vincent, G., Peirce, J.: 'MOVA: traffic responsive, self-optimising signal control for isolated intersections'. TRRL Research report, RR170, 1988

7   Hunt, P., Bretherton, R., Robertson, D., Royal, M.: 'Scoot on-line traffic signal optimisation technique', *Traffic Eng. Control*, 1982, **23**, pp. 190–192

8   Rose, G.: 'Mobile phones as traffic probes: practices, prospects and issues', *Transp. Rev.*, 2006, **26**, (3), pp. 275–291

9   Hounsell, N., Shrestha, B., Head, J., Palmer, S., Bowen, T.: 'The way ahead for London's bus priority at traffic signals', *IET Intell. Transp. Syst.*, 2008, **2**, (3), pp. 193–200

10 http://www.sias.com/ng/sparamicshome/sparamicshome.htm, accessed July 2010

11 Box, S., Waterson, B.: 'Comparison of signalized junction control strategies that use localization probe data'. Proc. IET Road Transport Information and Control Conf., May 2010

12 Kos, T., Grgic, M., Sisul, G.: 'Mobile user positioning in GSM/UMTS cellular networks'. Proc. 48th Int. Symp. ELMAR-2006 Focused on Multimedia Signal Processing and Communications, 2006, pp. 185–188

13 Trimble.: 'Lassen iq gps module'. Trimble datasheet, available at http://www.trimble.com, 2007

14 OXTS.: 'Rt3000 inertial and gps navigation system'. OXTS Brochure, available at www.oxts.co.uk, 2009

15 Levinson, J., Montemerlo, M., Thrun, S.: 'Map-based precision vehicle localisation in urban environments'. Proc. Robotics: Science and Systems, Atlanta, USA, June 2007

16 Li, W., Leung, H.: 'Constrained unscented Kalman filter based fusion of GPS/INS/digital map for vehicle localization'. Proc. IEEE Intelligent Transportation Systems, 2003, vol. 2, pp. 1362–1367

17 Grewal, M., Weill, L., Andrews, A.: 'Global positioning systems, inertial navigation and integration' (John Wiley & Sons, 2001)

18 Choy, M.C., Srinivasan, D., Cheu, R.: 'Cooperative, hybrid agent architecture for real-time traffic signal control', *IEEE Trans. Syst. Man Cybern. A, Syst. Humans*, 2003, **33**, (5), pp. 597–607

19 Box, S., Waterson, B.: 'Signal control using vehicle localization probe data'. Proc. 42nd Annual Conf. Universities' Transport Study Group, January 2010