27th Australasian Transport Research Forum, Adelaide, 29 September – 1 October 2004

Paper title: **Evaluation of lane changing and merging in microsimulation models**

Author(s) name(s): Peter Hidas

Organisation(s): University of New South Wales

Contact details:
Postal address: School of Civil and Environmental Engineering
University of New South Wales, Sydney, 2052, Australia

Telephone: +61 (0)2 9385 5988
Facsimile: +61 (02) 9385 6139
email: P.Hidas@unsw.edu.au

Abstract (200 words):

This study investigated and evaluated the lane changing and merging abilities of three traffic simulators, AIMSUN, VISSIM and ARTEMiS, under congested flow conditions. Several hypothetical traffic scenarios were constructed that require a large proportion of vehicles to merge or change lanes. These scenarios were implemented in each model and simulated under increasing traffic flow rates with various model parameter combinations. All three simulators have produced a reasonably similar range of output results from the same input values. All three models seem to have a number of parameters that can be used to calibrate the models to observed data. While the aggregate, macroscopic performance indicators seem to be reasonably close to each other and realistic, for each simulator a number of weaknesses and model limitations were identified that require further investigation and development. One common weakness in all three models is the occurrence of 'lost vehicles' indicating failures in the model procedures. Overall, this study has shown that there are inconsistencies between the simulation models and the results need to be treated with caution, especially when modelling highly congested traffic scenarios.

Introduction

Microscopic traffic simulation models are becoming increasingly important tools in modelling transport systems. They can be used to model complex transport networks and evaluate various traffic management alternatives in order to determine the optimum solution for traffic problems that cannot be studied by other analytical methods. Several traffic simulation tools are widely used in many countries, including AIMSUN, PARAMICS and VISSIM. While these packages are regarded in practice as the state-of-the-art, several problems were also identified including computational performance, the accuracy of models in representing the traffic flow, and the difficulty of integration with advanced traffic management and traffic information systems. A topic of increasing concern is the validity of the microscopic sub-models, such as the car following and lane changing models. It is often claimed that the visualisation and animation capabilities of microsimulation models are more advanced than their ability to reproduce measured/observed traffic flow characteristics. In a recent literature review and user opinion survey (Hidas, 2004a) several reports were found on lane changing and merging problems in congested flow conditions.

The aim of this study is to investigate, compare and evaluate the lane changing and merging abilities of some traffic simulators, especially under congested flow conditions. Three simulators are included in the study: AIMSUN (versions 4.1 and 4.2), VISSIM (version 3.70) and ARTEMiS (version 1.50). While PARAMICS is the most widely used model in Australia, unfortunately it was not available for this study. Several hypothetical traffic scenarios are constructed that require a large proportion of vehicles to merge or change lanes. These scenarios are then implemented in each model and simulated under increasing traffic flow rates with various model parameter combinations. The model outputs are compared and evaluated and conclusions are drawn on the quality of the simulators in terms of their ability to model lane changing and merging in congested conditions. This study is not aimed at calibrating the models to any particular real traffic scenario, but the aim is rather to discover the realistic range of output performance indicators, such as flow-speed relationships, that the models are able to produce in the tested scenarios, and to identify any limitations or irregularities in the models that require further research and development work.

The paper is structured as follows. After this introduction, the next section briefly introduces the lane changing models of each simulator and the parameters allowing calibration of the models. The following section describes the hypothetical traffic scenarios used to test the capabilities of the models. The next section presents the results of the simulation runs, and the final section summarises the findings and formulates recommendations for further research and model development.

Description of the lane changing models

This section presents the lane changing models of each simulator based on the information available from the user manuals, especially focusing on the model parameters that allow the user to calibrate the models.

AIMSUN

AIMSUN (TSS 2004) uses a development of the Gipps (1986) lane changing model. Lane change is modelled as a decision process, analysing the necessity of the lane change (eg. for required turning manoeuvres), the desirability of the lane change (to reach the desired speed), and the feasibility of the lane change (the availability of gaps in the target lane). Once a decision was made to change lane, the manouevre is executed in AIMSUN as an instantaneous switch from one lane to the other.

The decision process depends on how far the vehicle is from the next required turning point. AIMSUN divides this distance into 3 zones with distinct considerations in each zone:

- *Zone 1*: this is the zone farthest from the next turning point. While the vehicle is in this zone, it only considers lane changing desirability to reach its desired speed.
- *Zone 2*: the intermediate zone. In this zone, lane changing decisions are made to reach the required turning lane, but the behaviour of vehicles in the turning lane is not affected.
- *Zone 3*: the final section before the turning point. In this zone the vehicles are forced to reach their required turning lane, by reducing their speed or even stopping to wait for a gap, and vehicles in the target lane may modify their behaviour in order to allow the other vehicles to move into the target lane.

These zones are defined for each link in the network by two distance parameters: Distance Zone 2 and Distance Zone 1, measured in seconds of travel time, and they can extend beyond the length of the link due to the 'look ahead' model, so vehicles may start moving into their required turning lane before reaching the last link if necessary. The values must be calibrated for each location as they have a significant influence on the output results.

AIMSUN also has an optional 'Two-lanes Car Following Model', that can be used to force vehicles in the faster lanes of multilane roads to slow down in order to create better opportunities to change lane for vehicles in the slower lane(s) or in an on-ramp. The model is controlled by four user-defined parameters: number of vehicles considered, maximum distance ahead, maximum speed difference, and maximum speed difference on ramp.

If a vehicle is unable to reach its required turning lane, it will become stuck in the wrong lane, and after a given waiting time it will become a *lost vehicle* by continuing on the wrong route. The waiting time is determined by the user in the *maximum give-way time* vehicle parameter.

VISSIM

In VISSIM (PTV 2003)**,** basic traffic flow model parameters are grouped into a 'driver behaviour' setting. Several driver behaviours are defined, e.g. for urban and interuruban (motorway) conditions, and each link is associated with a driver behaviour set. This allows different parameters to be used in the same network model. The car following and lane changing models are based on the continuing work of Wiedemann (1991). Lane changing is controlled by the following parameters:

Lane Change Behaviour: can be *Free Lane Selection* (vehicles allowed to overtake in any lane) or *Right/Left Side Rule* (in the fast lane only).

Min. Headway (front/rear): minimum distance that must be available for a lane change in standstill condition.

Two distance parameters associated with each connector (i.e. turning point) determine when and how drivers make lane changing decisions in VISSIM:

- *Lane change* defines the distance at which vehicles will begin to attempt to change into lanes that allow moving into the connector. This distance can extend beyond the start of the link.
- *Emergency Stop* defines the last possible position for a vehicle to change lanes. If a vehicle could not change lanes due to high traffic flows but needs to change in order to stay on its route, it will stop at this position to wait for an opportunity to change lanes.

Waiting time before diffusion (default 60 sec) defines the maximum amount of time a vehicle can wait at the emergency stop position waiting for a gap to change lanes in order to stay on its route. When this time is reached the vehicle is taken out of the network (diffusion) and a warning message will be written to the error file denoting the time and location of the removal.

For necessary lane changes, the user can define the decelerations that the lane changing driver accepts for himself (*Own*) as well as for the vehicle he is moving ahead of (*Trailing*). The range of these decelerations is defined by the *Maximum and minimum (Accepted) Deceleration*, and a *reduction rate* (as metrer per 1 m/s²) is used to increase the Maximum Deceleration with decreasing distance to the emergency stop position. The concept is illustrated in Figure 1.

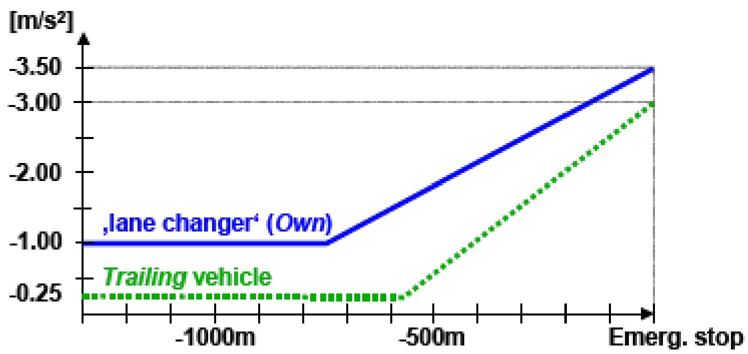

Figure 1 Accepted deceleration parameters in VISSIM (from PTV, 2003)

ARTEMiS

The lane changing model of ARTEMiS is developed by the author and described in detail in previous publications (Hidas, 2002, Hidas, 2004b). The model uses intelligent agent techniques to simulate the complex behavioural decision-making processes in vehicle interactions during lane changing and merging under congested flow conditions. The feasibility conditions are influenced by the following model parameters:

Urgency Start and *Max Urgency* distance: these distances, measured backwards from the turning point, determine the rate of the accepted speed and deceleration parameters defined below. When a vehicle is further from the turning point than the Urgency Start distance, the minimum parameters are used. When the vehicle reaches the Max Urgency distance, the maximum parameters are used, and in between the two, the rate increases linearly from minimum to maximum. The concept is similar to the one used in VISSIM (see Figure 1).

Min/Max Speed Reduction (FSR): these are the minimum/maximum speed reduction values that a follower vehicle in the target lane is willing to accept in a forced or cooperative lane changing situation.

Min/Max Foll Deceleration Rate Factor (FDR): the minimum/maximum deceleration rates, as a factor of the vehicle's normal deceleration rate, that a follower vehicle in the target lane is willing to accept in a forced or cooperative lane changing situation.

Min/Max Own Deceleration Rate Factor: the minimum/maximum deceleration rates, as a factor of the vehicle's normal deceleration rate, that the lane changing vehicle is willing to use in a lane changing situation.

Max Wait Time(MWT): the maximum time that a vehicle spends waiting for an opportunity to change lane. When this time elapsed and the vehicle was unable to change lane, it is removed from the simulation and becomes a *lost vehicle*.

The case study scenarios

In order to test the capabilities of the selected simulators, three hypothetical traffic scenarios were constructed: a motorway weaving section, a motorway on-ramp terminal, and an urban signalised intersection. The physical parameters and the demand flow levels for the three scenarios are as follows.

Weaving section

The weaving section network model consists of two input sections and two exit sections joining in a V-shape, and a common weaving section in the middle. Each road section is 150 m long, and the speed limit is 110 km/h. The input and exit sections each have 2 lanes, while the weaving section link has 4 lanes.

A period of 90 minutes was simulated, using increasing 15-minute demand flow rates. The demand flow rate at the start was set to approximately 50 % saturation level and gradually increased to full saturation (see Table 1). Both input sections had the same input flow rate, and the same proportion of weaving vehicles set to 40 percent. The demand flow composition is 95 % cars and 5 % trucks.

On-ramp terminal

The on-ramp terminal network consists of a 300 m long motorway section and a 150 m long on-ramp section, ending in a 100 m long ramp terminal (i.e. one acceleration lane along the motorway lanes). The motorway has 2 lanes and the speed limit is 110 km/h, while the ramp has one lane and 80 km/h speed limit.

A period of 120 minutes was simulated, using increasing 15-minute demand flow rates, again starting at a low saturation level and going up to full saturation (see Table 1). The ramp flow rate is always 30 percent of the total flow rate. The demand flow composition is 95 % cars and 5 % trucks.

Urban road

The urban road network consists of two input sections joining in a V-shape, followed by a common section which ends in a signalised T-junction. The input sections have one lane, the common section has 2 lanes, and all roads have a 50 km/h speed limit. The signalised intersection is operated by fixed-time control, with a 60 sec cycle time, and two phases of 30 sec each with 3 sec amber time.

A period of 90 minutes was simulated, using increasing 15-minute demand flow rates. Both input sections had the same input flow rate (see Table 1), and the same proportion of weaving vehicles set to 50 percent. The demand flow composition is 95 % cars and 5 % trucks.

Table 1 Demand flow rates used in the experiments

Time Interval	Flow rate (veh/h)			
	Weaving section	On-ramp terminal		Urban road
		Main flow	Ramp flow	
1	2500	1050	450	700
2	3000	1400	600	850
3	3250	1750	750	1000
4	3500	2100	900	1100
5	3750	2380	1020	1200
6	4000	2520	1080	1200
7		2940	1260	
8		3220	1380	

Results of the case studies

This section presents the results obtained from the simulation case studies. The simulators can produce a wide range of performance indicators in various time and space aggregation levels, but due to space constraints, only the speed-flow relationships are presented here (based on 5-minute averages) with a few other basic output values, such as the number of lost vehicles, as indicators of the success of the simulation model.

In each case, first results from each model with different parameter values are shown and compared, then the ‘best’ results of the 3 models are compared.

Weaving section case

AIMSUN: In AIMSUN, the effects of the following parameters were tested: the optional ‘Two-lane car following’ model (2LCF), the Distance Zone 2 (Z2, measured in seconds of travel time), and the reaction time (RT). In the simulation runs all parameters were used with their default values, and only one parameter was modified in each run to observe its effects on the output results. The average results (from the weaving section link only) for the whole simulation period are given in

The results show that when the Two-lane car following model is not used, the mean flow and speed are only slightly lower. The Distance Zone 2 has a more significant impact on the results. When Zone 2 = 2 seconds (by default), which is equivalent to about 61 metres, only about half of the 150 m long weaving section is available for the vehicles to find an opportunity to change lane; when they reach that distance from the end of the link, the model will stop them and force the vehicles in the target lane to slow down (or stop) and let them in. This causes the whole system to slow down and reduces the throughput. When Z2 = 3 seconds, the effect is even stronger, as expected. As the Z2 distance is reduced, the lane change opportunities are better distributed along the weaving section and consequently the mean flow and speed are higher. However, at the same time, the number of lost vehicles increases significantly. At Z2 = 1 sec, the 17 lost vehicles can still probably be neglected, but when Z2 becomes 0, the number of lost vehicles becomes exceedingly high, therefore this situation is unrealistic.

Table 2, and the 5-minute flow-speed relationships are shown Figure 2.

The results show that when the Two-lane car following model is not used, the mean flow and speed are only slightly lower. The Distance Zone 2 has a more significant impact on the results. When Zone 2 = 2 seconds (by default), which is equivalent to about 61 metres, only about half of the 150 m long weaving section is available for the vehicles to find an opportunity to change lane; when they reach that distance from the end of the link, the model will stop them and force the vehicles in the target lane to slow down (or stop) and let them in. This causes the whole system to slow down and reduces the throughput. When Z2 = 3 seconds, the effect is even stronger, as expected. As the Z2 distance is reduced, the lane change opportunities are better distributed along the weaving section and consequently the mean flow and speed are higher. However, at the same time, the number of lost vehicles increases significantly. At Z2 = 1 sec, the 17 lost vehicles can still probably be neglected, but when Z2 becomes 0, the number of lost vehicles becomes exceedingly high, therefore this situation is unrealistic.

Table 2 AIMSUN results for the weaving section case

Experiment	Model Parameters			AIMSUN Version 4.2			AIMSUN Version 4.1		
	2LCF*	Z2**	RT***	Mean Flow	Mean Speed	Lost	Mean Flow	Mean Speed	Lost
Run 1	OFF	2	0.75	1324	44.7	0	1623	71.9	3
Run 2	ON	2	0.75	1358	46.6	0	1589	65.0	3
Run 3	ON	3	0.75	1289	43.0	0	1518	61.3	2
Run 4	ON	1	0.75	1467	53.9	17	1631	68.4	8
Run 5	ON	0	0.75	1674	59.2	329	1618	69.9	30
Run 6	ON	2	1.00	1319	43.9	1	1417	67.7	0
Run 7	ON	2	0.50	1642	76.2	2	1672	84.9	9
Run 8	ON	2	0.25	1680	90.9	4	N/A		
Run 9	ON	2	0.10	1699	92.6	29	N/A		
*2LCF: Two-lane car following model **Z2: Distance Zone 2 (sec) ***RT: reaction time									

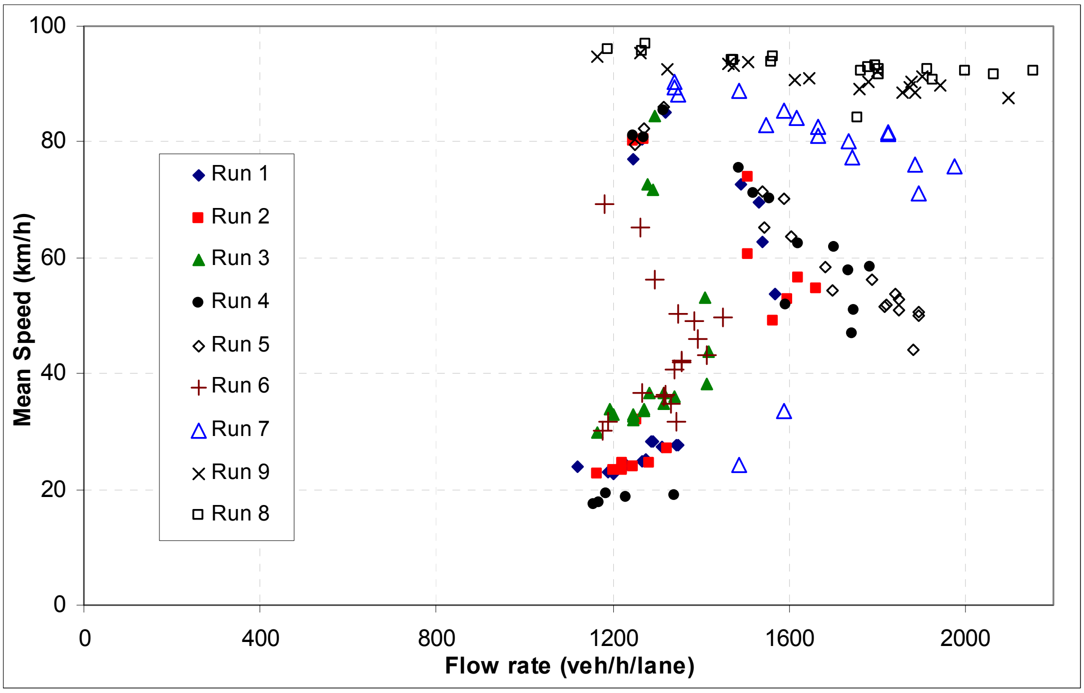

Figure 2 AIMSUN v4.2 results for the weaving section case

The global parameter ‘Reaction Time’ has the most significant impact on the results. According to the user manual ‘for implementation reasons the reaction time is also taken as the Simulation Time Step’, and in the latest version 4.2 its value can vary from 0.1 to 1.25 seconds. Using the default Distance Zone 2 = 2 seconds, four simulation runs were made with different values of reaction time. The results show a significant increase of the mean flow and speed as the ‘reaction time’ decreases. The speed-flow curves in Figure 2 show that the relationships become unrealistic below 0.5 seconds, as the model does not replicate the expected congestion effects. These results lead to the question of what might be the purpose of this range of ‘reaction time’ when it produces totally unrealistic flow results. If the simulation time step is tied to the reaction time then its range should not go below what is humanly possible.

The first experiments of this study were run in AIMSUN version 4.1. As version 4.2 became available, the experiments were repeated in the new version and this gave an opportunity to compare the outputs.

The results show that when the Two-lane car following model is not used, the mean flow and speed are only slightly lower. The Distance Zone 2 has a more significant impact on the results. When Zone 2 = 2 seconds (by default), which is equivalent to about 61 metres, only about half of the 150 m long weaving section is available for the vehicles to find an opportunity to change lane; when they reach that distance from the end of the link, the model will stop them and force the vehicles in the target lane to slow down (or stop) and let them in. This causes the whole system to slow down and reduces the throughput. When Z2 = 3 seconds, the effect is even stronger, as expected. As the Z2 distance is reduced, the lane change opportunities are better distributed along the weaving section and consequently the mean flow and speed are higher. However, at the same time, the number of lost vehicles increases significantly. At Z2 = 1 sec, the 17 lost vehicles can still probably be neglected, but when Z2 becomes 0, the number of lost vehicles becomes exceedingly high, therefore this situation is unrealistic.

Table 2 shows that version 4.1 produced consistently higher flows and speeds with the same model parameters, and less lost vehicles in the extreme cases. These results are difficult to explain, as details of the model component changes are not described in the user manual, but the magnitude of the changes from one version to the next is certainly surprising. A discussion with the developers (Barcelo 2004) revealed that the difference is due to the changes in the revised car following model implemented in version 4.2. Only version 4.2 was used for the other case studies.

VISSIM: In VISSIM, the effects of the simulation time step and the accepted deceleration parameters of the trailing vehicle were tested. Table 3 shows very consistent mean flow and speed results on the weaving section for the 0.10-0.50 second time step range, but at 1 second time step there is a significant decrease in throughput and mean speed. When the deceleration accepted by the trailing vehicle is reduced (using the default 0.10 sec time step), a gradual decrease can be observed in the mean flow and speed.

Table 3 VISSIM results for the weaving section case

Experiment	Model parameters				Results		
	Time step (s)	Trailing deceleration (m/s^2)			Mean Flow	Mean Speed	Lost
		Max	rate (m)	Min			
Run 1	0.10	-3.0	200	-0.5	1442	30.3	1
Run 2	0.25	-3.0	200	-0.5	1434	31.7	0
Run 3	0.50	-3.0	200	-0.5	1427	33.3	0
Run 4	1.00	-3.0	200	-0.5	1211	29.3	1
Run 5	0.10	-2.0	200	-0.5	1446	31.8	2
Run 6	0.10	-1.0	200	-0.5	1423	26.7	0
Run 7	0.10	-1.0	200	0.0	1395	25.5	0
Run 8	0.10	0.0	200	0.0	1218	18.7	0

Most experiments had no lost vehicles at all, only three of the 8 produced one or two lost vehicles. However, when checking the details of the simulation display, it was found that these vehicles were removed immediately when a collision occurred between two vehicles involved in a lane changing manoeuvre, and not after the set period waiting for an opportunity to change lane.

Figure 3 shows the speed-flow curves for the eight experiments. The relationships are very similar for all but two cases: the 1.0 second time step (Run 4) and the minimum accepted trailing deceleration (Run 8) show significantly lower flow rates and speeds than the others. In all cases, the curves indicate that the weaving section operates in fully saturated (forced flow) conditions almost from the beginning of the simulated time period.

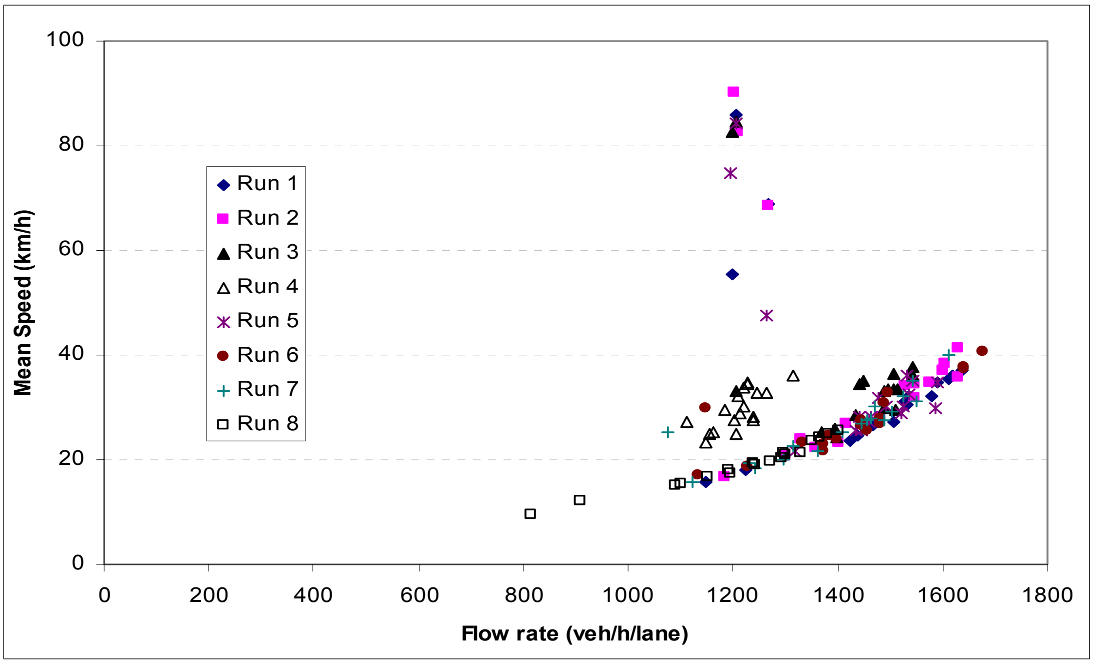

Figure 3 VISSIM results for the weaving section case

ARTEMiS: ARTEMiS runs at a fixed 1.0 second simulation time step, so the experiments were constructed with a range of model parameters representing different levels of willingness to cooperate by the follower vehicle. These include the speed reduction and deceleration rate accepted by the follower vehicle and the maximum waiting time. Higher values of speed reduction (FSR) and deceleration rate (FDR) accepted by the follower represent higher level of cooperation. As expected, Table 4 shows decreasing mean speed and flow results as the willingness to cooperate is reduced. The number of lost vehicles also shows an increasing tendency, with an unexpected higher value in the second experiment. Although the number of lost vehicles remains below the level that would have a perceivable impact on the overall results, it is definitely more than negligible, requiring further work.

Table 4 ARTEMiS results for the weaving section case

Experiment	Model parameters					Results		
	Min FSR[1]	Max FSR[2]	Min FDR[3]	Max FDR[4]	MWT[5]	Mean Flow	Mean Speed	Lost
MaxCoop	10	30	0.8	1.5	60	1475	52.8	2
Mid1	10	20	0.5	1.2	60	1393	51.9	16
Mid2	0	10	0.2	1.0	60	1420	46.9	7
Mid3	0	5	0.0	0.2	60	1380	48.7	12
MinCoop	0	0	0.1	0.2	60	1262	40.8	18

1: Min FSR: Min Speed Reduction accepted by the Follower vehicle (km/h)
2: Max FSR: Max Speed Reduction accepted by the Follower vehicle (km/h)
3: Min FDR: Min Deceleration Rate Factor accepted by the Follower vehicle
4: Max FDR: Max Deceleration Rate Factor accepted by the Follower vehicle
5: MWT: Max Wait Time (sec)

The speed-flow curves in Figure 4 show the same properties mentioned above: the higher the level of cooperation, the higher the throughput and speed per time interval. The ARTEMiS experiments show the first third of the time period as under-saturated and only the rest of the time as fully saturated.

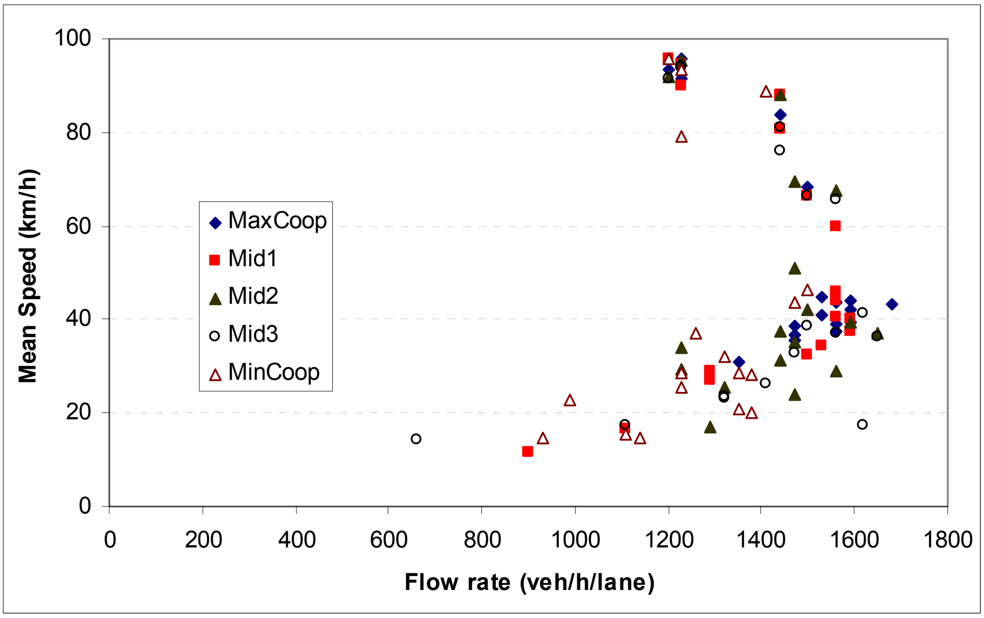

Figure 4 ARTEMiS results for the weaving section case

Comparison: Based on the output results, the experiment considered as the 'best' (i.e. most realistic with the highest throughput) was selected from each model for comparison purposes. From the AIMSUN runs this was the case with Zone 2 = 1 second, with 0.75 sec reaction time and Two-Lane Car Following model on. From VISSIM, the default parameters with maximum trailing deceleration (Run 1) was selected, and from ARTEMiS, the maximum cooperation case was chosen.

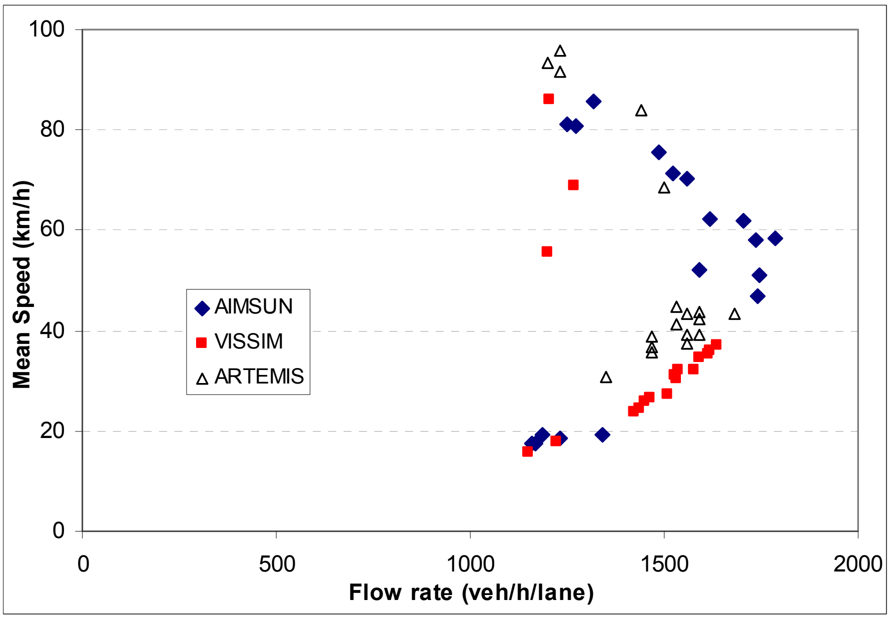

Figure 5 Comparative results for the Weaving section case

Figure 5 shows the speed-flow curves for the three simulators. There are significant differences between the models. AIMSUN shows the highest maximum throughput and the longest under-saturated period. The ARTEMiS curve has a similar shape with somewhat lower maximum flow, but higher speed range in the congested period. VISSIM shows a very different picture with the lowest maximum throughput and most of the time in the low-speed fully congested conditions.

It is important to note that these comparisons are purely relative and are not meant to indicate any qualitative evaluation of the models, because the given demand scenario does not represent a 'real' situation and therefore the 'true' flow conditions that would occur in a real case are not known. However, the characteristic differences between the models in the outputs resulting from the same inputs are worth noting.

On-ramp terminal case

AIMSUN: In the case of an on-ramp, the Zone 2 time distance parameter in AIMSUN is not relevant. There is a separate parameter, Time Distance On-Ramp, to indicate the length of the acceleration lane to be considered by vehicles for merging into the continuous lanes, but this is only relevant on longer ramp terminals. This leaves the Two-lanes Car Following model and the reaction time as parameters that can be used to calibrate the on-ramp model. In the last experiment, the effect of the Maximum Give-Way Time (MGWT) vehicle parameter was tested: the mean MGWT was reduced for cars from 15 (default) to 5 seconds, for trucks from 50 to 10 seconds. Note that as there is only one destination in the on-ramp scenario, there can be no lost vehicles in this case in AIMSUN. When a vehicle at the end of the on-ramp has been stopped more than the MGWT, the lane changing rules of Zone 3 are applied, that is vehicles in the adjacent lane are forced to slow down and allow the vehicle to merge.

Table 5 AIMSUN results for the ramp terminal case

Experiment	2LCF	RT	MGWT	Total		Lost	Main stream		Ramp stream	
				Flow	Speed		Flow	Speed	Flow	Speed
Run 1	OFF	0.75	15/50	2998	56.8	NA	1081	76.7	836	42.9
Run 2	ON	0.75	15/50	3009	53.0	NA	1095	72.2	819	40.8
Run 3	ON	0.50	15/50	3067	65.7	NA	1082	86.1	901	52.8
Run 4	ON	0.25	15/50	3166	76.8	NA	1103	92.4	958	65.7
Run 5	ON	1.00	15/50	2861	46.4	NA	1029	61.4	802	36.6
Run 6	ON	0.75	5/10	2978	52.6	NA	1054	70.8	869	42.1

Table 5 shows the results of the AIMSUN experiments for the on-ramp terminal case. The use of the Two-lanes Car Following model slightly increases the mean flow but the mean speed becomes a little slower, as expected. However, the distribution by streams is somewhat surprising. One would expect that the Two-Lanes Car Following model should create better opportunities for the vehicles from the ramp to merge, but both the mean flow and speed from the ramp are slightly lower when the model is used. Again, the reaction time (i.e. simulation time step) has the greatest influence on the results: smaller time steps lead to higher mean flow and speed, and vice versa. In Run 6 the reduction of the MGWT (with the default 0.75 time step) has very little effect on the total results compared with the default case (Run 2), but this is the result of opposite effects on the main and ramp stream values: the mean flow and speed increase on the ramp while they decrease on the main road, as expected.

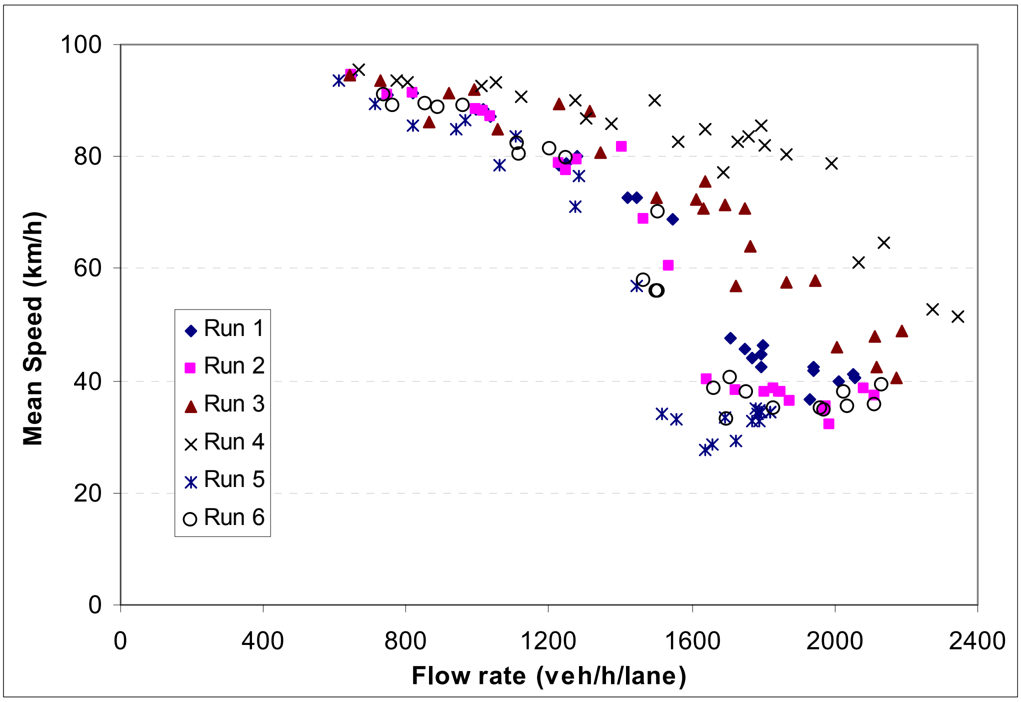

Figure 6 AIMSUN results for the ramp terminal case

From the speed-flow curves in Figure 6 the 0.25 s 'reaction time' case produces the highest speeds and throughput, and the 1.0 s curve shows the most congested conditions. The other four curves are fairly close to each other with little characteristic difference between them.

VISSIM: In VISSIM, the effects of the simulation time step and the trailing deceleration were investigated. Table 6 shows that 4 of the 6 runs produce fairly similar outputs, while two experiments had significantly lower results: the 1.0 second time step (Run 3) and the zero accepted trailing deceleration case (Run 6). It seems these parameters are beyond the satisfactory range for the model. These two runs also produced a few lost vehicles, while the other runs had no lost vehicles. Again, it was observed that some of the lost vehicles were removed immediately from the network when they reached the end of the ramp rather than waiting for the set maximum waiting time.

Table 6 VISSIM results for the Ramp terminal case

Experi-ment	Model parameters			Total Results			Main stream		Ramp stream	
	Time step	Trailing dec.		Mean Flow	Mean Speed	Lost	Flow	Speed	Flow	Speed
		Max	Min							
Run 1	0.10	-3.0	-0.5	2995	76.1	0	1075	89.4	900	76.6
Run 2	0.50	-3.0	-0.5	2923	72.1	0	1054	86.1	874	73.1
Run 3	1.00	-3.0	-0.5	2599	41.6	4	934	50.8	774	44.7
Run 4	0.25	-3.0	-0.5	2972	74.3	0	1074	88.4	892	75.2
Run 5	0.25	-2.0	-0.5	2968	68.7	0	1068	85.7	894	73.8
Run 6	0.25	0.0	0.0	2746	37.0	5	1072	89.9	712	39.1

Figure 7 also indicates that Runs 3 and 6 are less than realistic, while the other curves produce much higher speeds and maximum flows. However, all experiments in VISSIM show a fraction of the simulated period as fully congested with very low speeds.

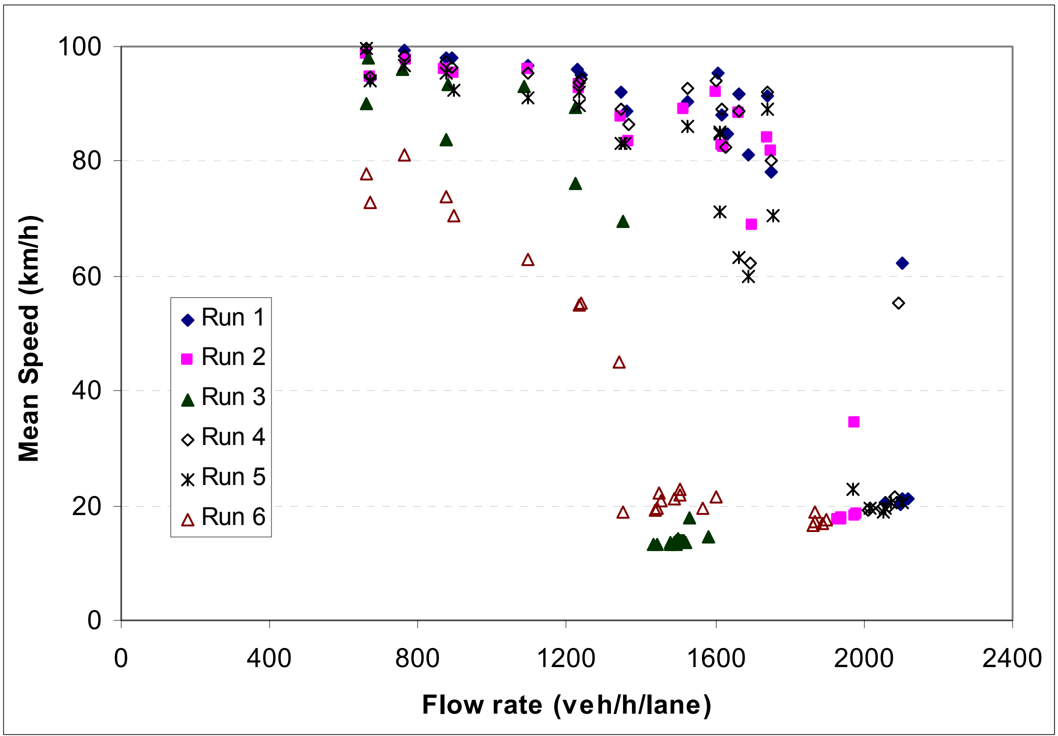

Figure 7 VISSIM results for the ramp terminal case

ARTEMiS: In ARTEMiS, the effects of the maximum speed reduction accepted by the follower (Max FSR) and the maximum waiting time (MWT) were investigated. As Table 7 indicates, it was found that the total average flow and especially speed decrease as Max FSR increases. The results separated by streams are interesting: the mean flow on the main road decreases by about 6 %, while the ramp stream flow remains constant. The mean speed on the main road also decreases more (18 %) than on the ramp (10 %). This indicates that the increased willingness of the follower vehicles on the main road to cooperate does not result in any perceivable benefit in the system. Also, the 5 second MWT with minimum cooperation produced an unacceptably high number of lost vehicles, but an increased MWT eliminated this problem. In the other experiments the number of lost vehicles was reasonably low. Figure 8 also shows that more cooperation leads to more congestion and lower maximum throughput values.

Table 7 ARTEMiS results for the ramp terminal case

Experi-ment	Parameters		Total results			Main stream		Ramp stream	
	Max FSR	MWT	Flow	Speed	Lost	Flow	Speed	Flow	Speed
Run 1	**0**	**5**	3015	78.0	37	1060	89.2	930	71.1
Run 2	**0**	**60**	3015	71.0	0	1058	86.3	905	67.7
Run 3	**5**	**5**	2880	68.6	4	993	75.9	910	62.2
Run 4	**10**	**5**	2870	62.8	3	990	71.0	905	61.1

Comparison: Figure 9 presents the selected 'best' alternatives from the 3 models: 2carFoll ON for AIMSUN, 0.10 sec time step with maximum trailing deceleration values for VISSIM, and the minimum cooperation case with 60 sec MWT for ARTEMiS. The figure shows that the 3 models produce similar maximum throughput values, but the shape of the curves are quite different. VISSIM shows the highest speed values for most of the flow range, then the curve suddenly drops to the fully congested conditions at the maximum flow with very low speed. ARTEMiS produces a fairly similar curve with slightly more continuous change between the unsaturated and saturated parts of the curve. AIMSUN predicts significantly lower speeds for the same mid-range flow rates than the other two models. Again, it is important to note that these curves do not represent any real flow conditions, and that the variance of the figures (that would result from several model runs with different random number seeds) was not investigated.

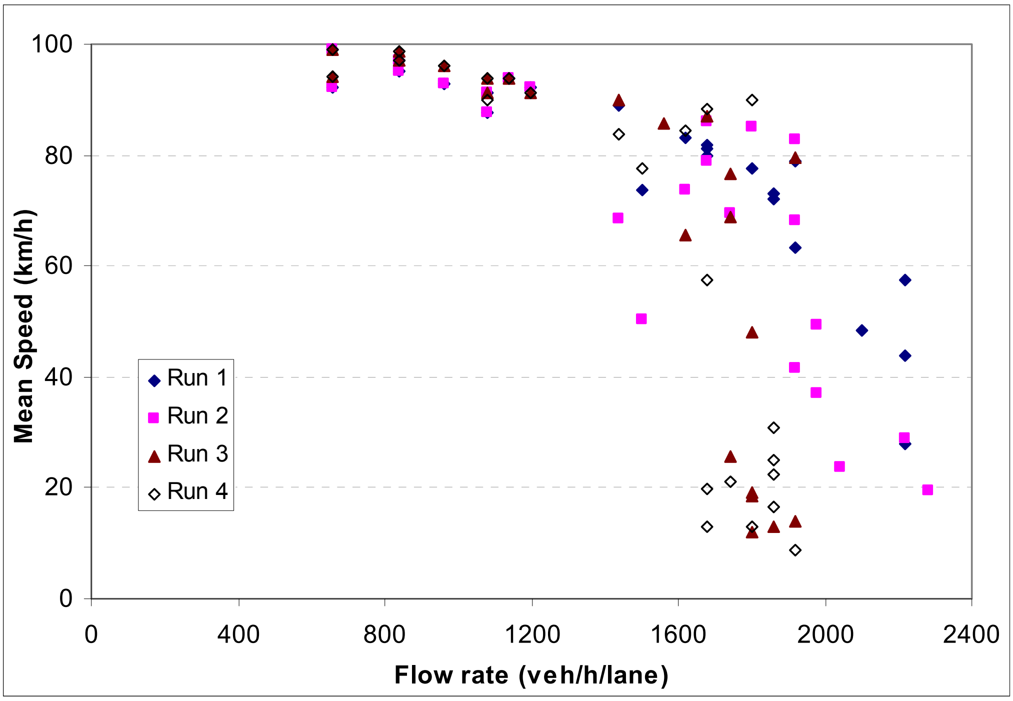

Figure 8 ARTEMiS results for the ramp terminal case

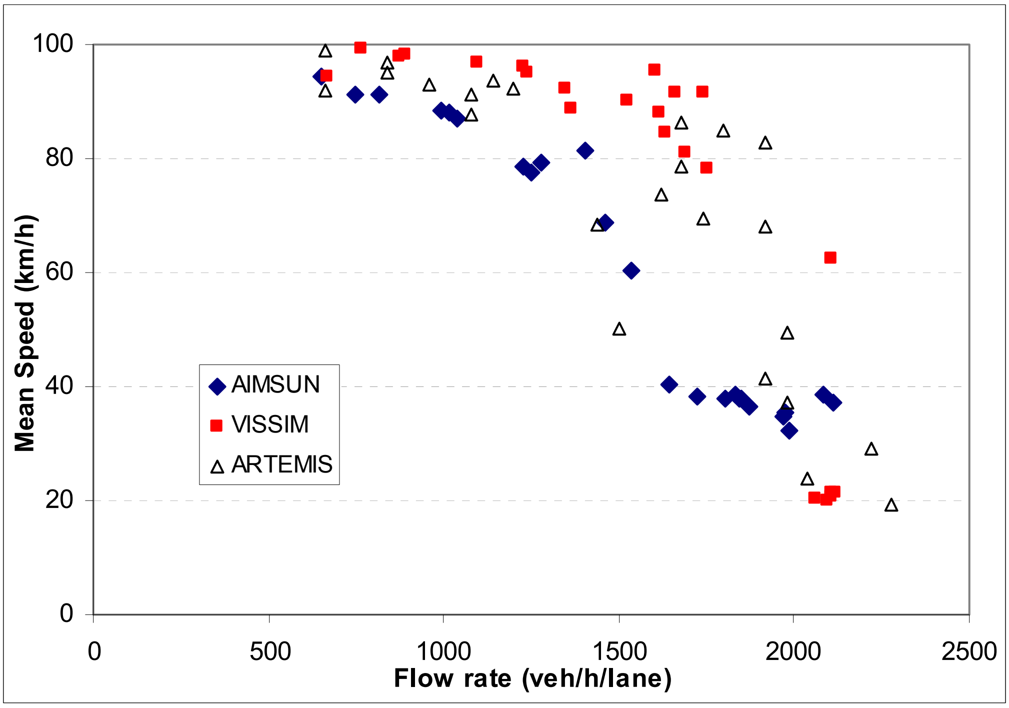

Figure 9 Comparative results for the ramp terminal case

Urban road case

In the urban network scenario, the traffic signal limits the capacity and the mean speed of the road link under investigation, therefore the speed-flow relationships are not relevant in this case. Table 8 shows the overall average results for all 3 simulators. The mean flows and speeds are very close in each case. VISSIM produces slightly higher mean flow, and AIMSUN (with Distance Zone 2 = 1 sec) has slightly higher mean speed, but the differences are not significant. The very strong effect of the 'reaction time' is again evident from the unrealistically high output results. The number of lost vehicles is also consistently high in AIMSUN. VISSIM produces very consistent results with the exception of the 1.0 sec time step, which seems too long for the model to be working satisfactorily. The results from ARTEMiS show that the model cannot produce satisfactory results without a reasonable level of cooperation.

Conclusions

The following conclusions can be drawn from the case studies. All three simulators have produced a reasonably similar range of output results from the same input values. All three models seem to have a number of parameters that can be used to fine-tune the results to a particular set of observed data. While the aggregate, macroscopic performance indicators seem to be reasonably close to each other and realistic, for each simulator a number of weaknesses and model limitations were identified that require further investigation and development.

Table 8 Results for the urban road case

Experiment	Model parameters				Results		
AIMSUN	2LCF	RT			Mean Flow	Mean Speed	Lost
Run 1	OFF	0.75			1415	11.4	9
Run 2	ON	0.75			1415	11.4	9
Run 3	ON	0.50			1729	14.7	9
Run 4	ON	0.25			1913	19.9	4
Run 5	ON	1.00			1317	8.7	12
VISSIM	Sim Time step (sec)	Trailing deceleration m/s^2					
		Max	rate (m)	Min			
Run 1	0.10	-3.0	100	-1.0	1455	8.8	0
Run 2	0.50	-3.0	100	-1.0	1454	9.3	0
Run 3	1.00	-3.0	100	-1.0	1431	9.1	4
Run 4	0.50	-3.0	100	-1.0	1445	9.1	0
Run 5	0.50	-0.5	100	0	1446	8.7	0
ARTEMiS	Max FSR	Max FDR	MWT				
Run 1	10	1.0	60		1407	8.1	0
Run 2	5	0.5	60		1407	7.8	0
Run 3	0	0	60		1093	6.0	7

AIMSUN was found to be very sensitive to the value of 'reaction time' which is taken as the simulation time step. Values of 'reaction time' at and below 0.5 second are unrealistic and indeed, produce unrealistically high vehicle throughputs without congestion effects. It seems illogical to combine the reaction time with the simulation time step if its range goes beyond what is humanly possible. The significant differences found between versions 4.1 and 4.2 must be considered by users when upgrading to the new version. The number of lost vehicles in the urban road scenario was consistently high, indicating some weaknesses in the lane changing model at low speeds.

VISSIM produced very consistent results in all three scenarios. It was found that while the model works well with simulation time steps up to 0.5 seconds, at 1.0 second time step the results become unsatisfactory. In general, the model produced very few lost vehicles, but the circumstances in which the vehicles were lost are more serious than what is described in the user manual: in fact some vehicles were removed when a collision occurred between lane changing vehicles or when a vehicle reached the end of the acceleration lane. While these occurrences are rare enough for not having a significant impact on the overall performance indicators, they highlight the need for further improvements in the lane changing and merging algorithms.

ARTEMiS was found to produce satisfactory results with its fixed 1.0 second simulation time step. The number of lost vehicles was acceptable in most cases but high enough to require further model improvements. The on-ramp scenario indicated that the increasing willingness of the follower vehicles on the main road to cooperate did not result in any perceivable benefit in the system. This shows the need for a review of the lane changing cooperation algorithm in the high speed range.

Overall, this study has shown that there are inconsistencies between simulation models and the results need to be treated with caution, especially when modelling highly congested traffic scenarios. This study was based on three very simple hypothetical traffic scenarios and a limited number of simulation experiments in order to investigate the basic lane changing and merging abilities of the three simulators. While a number of issues for further research were identified, more work is required to map out the capabilities and limitations of each model. Several simulation runs with different random seeds should be done to show the variance of the results from the same input data and model parameters. There are other model parameters and value combinations to be investigated in each model. There are other, more complex traffic scenarios to be tested, such as lane selection related to downstream short lanes and multiple turning movements from the same lane. Other widely used microsimulation models, especially PARAMICS, should also be tested as they are likely to have similar deficiencies. These experiments will be considered in the next phase of this research study.

Acknowledgements

The author is grateful to Martin Fellendorf (PTV AG) for providing a free copy of the VISSIM simulator for research purposes.

References

Barcelo, J (2004) Discussion on the differences in AIMSUN between versions 4.1 and 4.2. Personal communication.

Gipps, P.G. (1986) A model for the structure of lane-changing decisions. Transportation Research B. Vol. 20-B, No. 5, pp. 403-414.

Hidas, P (2002) Modelling Lane Changing and Merging in Microscopic Traffic Simulation, Transportation Research, Part-C: Emerging Technologies, Volume 10, Issues 5-6, October-December 2002, Pages 351-371, Elsevier Science Ltd. ISSN: 0968-090X.

Hidas, P (2004a) Comparative Evaluation of Micro-simulation Models, Final Report, prepared for the Roads And Traffic Authority of NSW, Unisearch Ltd.

Hidas, P (2004b) Modelling Vehicle Interactions in Microscopic Simulation of Merging and Weaving, paper accepted for publication in Transportation Research Part C: Emerging Technologies, Elsevier Science Ltd. ISSN: 0968-090X.

PTV (2003) VISSIM User Manual, Version 3.70. Planung Transport Verkehr AG

TSS Transport Simulation Systems (2004) AIMSUN Version 4.2 User Manual.

Wiedemann, R. (1991) Modelling of RTI-Elements on multi-lane roads. In: Advanced Telematics in Road Transport edited by the Commission of the European Community, DG XIII, Brussels.