ITS Journal, 7:79–102, 2002

1024-8072/02 $12.00 + .00
DOI: 10.1080/10248070190048691

DEVELOPMENT AND PERFORMANCE EVALUATION OF AN ITS-READY MICROSCOPIC TRAFFIC MODEL FOR IRVINE, CALIFORNIA

B. Abdulhai
Department of Civil Engineering, University of Toronto, Toronto, Ontario, Canada

J.-B. Sheu
Department of Transportation, Warehousing and Logistics, National Kaohsiung First University of Science and Technology, Kaohsiung, Taiwan, ROC

W. W. Recker
Department of Civil and Environmental Engineering, Institute of Transportation Studies, University of California, Irvine, California, USA

The research in this paper presents the detailed development and on-line evaluation of a microscopic traffic flow model for the city of Irvine in Southern California. This effort is the first stage of evaluating micro-simulators in terms of their ability to model and analyze Intelligent Transportation Systems (ITS) under faster-than-real-time conditions. We utilize "Paramics," a particularly promising ITS-capable advanced traffic flow simulator and visualization tool, as one of an array of newly emerging ITS-capable simulation tools and we apply it to the Irvine network as part of a staged effort to model the much larger Southern California network. The driver and vehicle models

This research was funded by the California Partners for Advanced Transit and Highways (PATH) headquartered at the University of California, Berkeley, and the California Department of Transportation (Caltrans). The work was facilitated by the California ATMIS Testbed, headquartered at the University of California Irvine.

Address correspondence to B. Abdulhai, Department of Civil Engineering, University of Toronto, 35 St. George Street, Toronto, Ontario, Canada M5S 1A4. E-mail: baher@ecf.utoronto.ca

and parameters were developed to reflect U.K. driver and vehicle characteristics. In this effort we explain our procedure used the calibrate these parameters to reproduce local U.S. traffic behavior. We built a model of a conventional U.S. freeway/arterial network in Southern California and calibrated its parameters using on-line field data. The calibrated models are validated, both at the section and network levels, and evaluated relative to their potential application in Advanced Traffic Management and Information Systems (ATMIS). Based on obtained results, the calibrated model performed well during validation on a freeway link. On the full network, the vehicle release mechanism showed some time-lag in releasing demand onto the network. This is potentially due to stacking of vehicles in memory before adequate headways are found on the road to release the vehicles. Although the problem itself is simple, its effects on the results were notable.

Keywords: microscopic traffic simulation; intelligent transportation system; traffic control and management

INTRODUCTION

Intelligent Transportation Systems (ITS) appear to provide a set of intuitively promising tools to improve the increasingly complex and congested transportation systems of today. Potential success of such ITS is implicit and derived to a large extent from the success of the underlying technology in other fields such as the defense and aerospace industries. Comprehensive research tools for quantifying the expected benefits from ITS are, however, either still absent or in their infancy. To quantify potential benefits prior to any major investment in development and deployment, the use of traffic simulation is regarded inevitable, at present at least. Simulators are needed, not only to assess the benefits of ITS in a planning mode, but also to generate scenarios, optimize control, and predict network behavior at the operational level.

Although a number of traffic simulators are currently under development, covering wide ranges of complexity, comprehensiveness, and potential usefulness, these simulators have to be closely calibrated and evaluated before their results can be locally applied. In this regard, researchers involved in simulation models calibration usually encounter three major obstacles: (1) the unavailability of a network/corridor of a manageable size and that is adequately instrumented with *functioning* vehicle detectors to supply the necessary data, (2) the absence of a comprehensive test bed that is capable of modeling and testing a variety of ITS components and of acquiring real-time traffic data, and (3) the unavailability of a *non-limiting* traffic simulator that is easy to use and yet capable of ITS modeling in a corridor/network environment of any realistic size. Non-limiting refers to being capable of handling large networks.

The prime objective of using traffic simulators *within an ITS context* is to serve as tools for dynamic transport management. More specifically, simulators can play two distinct roles: (1) as an off-line evaluation/design tool and (2) as an on-line control/guidance tool. Both roles cover numerous ATMIS applications, including: provision of traveler information and route guidance, a wide variety of surface street and freeway adaptive control (adaptive signal control, adaptive ramp metering, lane use control, etc.), incident detection and management, automated toll collection, and assessment of environmental impact of transport design and management, to name but a few. The off-line role is the easier of the two roles, as real-time operation is not as pressing a need as it is the case for the on-line role. If a simulator is fast enough, however, it could be used for both functions—off-line as well as on-line. Since both off-line and on-line functionality are inherent objectives of ATMIS, the efficiency and scalability of the simulator are regarded as key factors.

Transportation networks are, by default, dynamic relative both to supply/performance and demand. Such unexpected events as incidents, for example, inevitably occur and therefore change the supply side of the network. Any intervention by the responsible Transportation Management Center (TMC), possibly in the form of updating control measures, further change the supply side. Such changes can motivate drivers to change their behavior in several ways, including enroute and/or pre-trip route adjustments, that may occur within the day, or from day to day. Similar dynamics take place regarding drivers' choice of departure times in response to the dynamically changing supply conditions. The collective user behavior and response in this fashion give rise to dynamic demand profiles. Therefore, for any simulator to prove useful for dynamic transport management, it should be capable of:

1. Capturing the dynamics of supply, in terms of the detailed configuration of the transportation network and its performance in response to demands and TMC control functions implementation.

2. Capturing the dynamics of demand, in terms of dynamic user behavior in response to observed supply, either directly or via traveler information systems.
3. Capturing the complex dynamic interaction between supply and demand.
4. Performing faster than real-time to allow for proactive (based on predicted conditions) rather than reactive (based on observed conditions) dynamic transport management.

OBJECTIVE AND SCOPE

The research in this paper presents the detailed development, calibration, and on-line evaluation of a microscopic traffic flow model for the city of Irvine in Southern California. Calibration results reported herein concentrate on driver behavior parameters (i.e., those controlling car following, lane changing, gap acceptance, awareness, and aggressiveness), the heart of any microscopic simulator. This effort is the first stage of evaluating micro-simulators in terms of their ability to model and analyze Intelligent Transportation Systems (ITS) under faster-than-real-time conditions. Towards that goal, we utilize "Paramics," a particularly promising ITS-capable advanced traffic flow simulator and visualization tool, as one of an array of newly emerging ITS-capable simulation tools and we apply it to the Irvine network as part of a staged effort to model the much larger Southern California network. Paramics was chosen mainly because of its scalability and high performance in terms of modeling large networks at high processing speed. The driver and vehicle models and parameters in the simulator, which are developed to reflect U.K. driver and vehicle characteristics, are calibrated to reproduce local behavior on a conventional U.S. freeway/arterial network in Southern California using on-line field data. The calibrated models are validated, both at the section and network levels, and evaluated relative to their potential application in Advanced Traffic Management and Information Systems (ATMIS).

DESCRIPTION OF THE TRAFFIC SIMULATION ENVIRONMENT

Because of its high processing performance and ITS modeling capabilities, Paramics was acquired and adopted for ATMIS modeling at the University of California Irvine. Individual vehicles are modeled in fine detail for the duration of their entire trip, providing accurate traffic flow, transit time and congestion information, as well as enabling the modeling of the interface between drivers and ITS. Parameters affecting vehicle and driver performance are accessible to modelers. The software is portable across computing platforms and scalable, allowing a unified approach to traffic modeling across the whole spectrum of network sizes, from single intersections up to national networks. Performance is a trade-off between number of vehicles simulated and the processing power available. The model interfaces to standard macroscopic data formats, as well as to individual vehicle counts from induction loops and optical sensors data. The model is also capable of interfacing with emission models. Finally, it is equipped with advanced visualization tools that adds a complementary subjective traffic behavior evaluation capability under different ITS scenarios.

Computing speed/performance affects the maximum size of network that can be processed. Measured in terms of number of vehicles that can be simulated in real-time (vehicle in a snap shot of the network) using a 0.5 second time-step, a single processor version of the simulator has been found capable of simulating 20,000 vehicles on a 500 Mhz machine. It is noteworthy that in microscopic simulation the number of vehicles in a network at any instant is more important that the physical network size in terms of numbers of links and nodes.

Of particular importance to the development and use of the simulation set of tools in an ATMIS environment is a "framework" version of the software in which the simulator could be used as a shell, providing researchers with the ability to test their own models through a series of application programming interfaces (APIs). Such APIs have a dual role; first to allow override of the simulator's default models (car following, lane changing, and route choice models, for instance), and second, to interface complementary modules to the simulator. Complementary modules include such ITS applications as signal optimization, adaptive ramp metering, and incident detection and management. Current APIs provide a two-way functional interface with (1) control functions, and (2) callback functions. The control functions allow the user to pass external information to the simulator, controlling such items as speed of each vehicle, likelihood of changing lanes, and so on. Callback functions achieve the opposite, allowing the user to extract information from the simulator, for example, the attributes of a vehicle and its environment. Current APIs[1] allow the user to override the default car-following and lane-changing logic by specifying alternative dynamics to control both longitudinal and lateral vehicle movements. Broader functionality of the APIs, including dynamic traffic assignment capabilities, are nonexistent in the tested 1.5 version.

[1] Version 1.5

Calibration of Key Model Parameters

At the fundamental level of the simulator, a number of functions are used to control vehicle-following and lane-changing behaviors. These functions operate on data structures that describe individual vehicles, and take, as input, an array of vehicles that describe the environment around a particular vehicle. A vehicle's data structure contains information that describes the physical characteristics of the vehicle, the behavior parameters of the driver, the vehicle's current position and dynamics, its destination, and transit time.

Calibration results reported herein concentrate on driver behavior parameters (i.e., those controlling car following, lane changing, gap acceptance, awareness, and aggressiveness), the heart of any microscopic simulator. Only subjective description of the driver behavior models incorporated in the simulator is presented here; actual model details can be found in Paramics documentation.

The Car-Following Model

In terms of car following, each Driver-Vehicle Unit (DVU) has a target headway that varies around a mean value depending on such factors as weather, highway type, vehicle type, driver's aggressiveness, and awareness. High aggressiveness, for instance, would cause drivers to adopt smaller headways. Similarly, high awareness value would cause drivers to use longer headways near lane drops as a sort of yielding to merging traffic. Alternatively, if a merging DVU is aggressive, it accepts a smaller gap. DVUs either accelerate, cruise, or brake to maintain the target headway stimulated by the perceived relative speed, acceleration, and brake lights of the vehicle ahead. Perception-reaction lag is taken into account. Under light traffic conditions, DVU flow unconstrained by other vehicles, limited by lane-specific mean target speed.

The Lane-Changing Model

Lane use is affected by the vehicle's position relative to its target range of lanes, the latter being consistent with upcoming routing decisions, and the overtaking interactions between nearby vehicles. The target range of lanes is also affected by the DVU aggressiveness and awareness. A higher level of aggression causes a DVU to use the outer high-speed lanes, and a higher level of awareness causes a DVU to adopt the target lane for an impending turn sooner. Overtaking is controlled by varying two tables of thresholds. The lane-awareness threshold table specifies when a vehicle should move out to let a vehicle behind pass. If a DVU's awareness is greater than or equal to the threshold applicable to its current lane, it will move out. The lane-aggressiveness threshold is similar; if a DVU's aggressiveness is greater than the threshold applicable to its lane, it will attempt to overtake. DVU's propensity to change lanes can be controlled by varying the length of time that a vehicle must receive a positive stimulus, that is, a suitable gap must exist in the target lane for *n* seconds in order for the vehicle to change lanes. To dampen lane-changing oscillations, a waiting period is specified. A vehicle receiving a negative stimulus must wait for the "wait on failure" period before attempting to change lanes again. A vehicle successfully changing lanes must wait for the "wait on success" period before attempting another lane-changing maneuver. Lane changing is therefore affected by both the availability and stability of acceptable gaps in the target lane as well as the history of lane changes of the DVU itself. Gap acceptance itself is a function of DVU aggressiveness and awareness. The user defines driver's aggressiveness and awareness by selecting a distribution for each, across the driver population (normal distribution for instance).

Baseline Model Results

As a first step, a phase of preliminary testing of the simulator without calibration was conducted. The objectives of this phase included: 1) evaluating the uncalibrated model's performance on a small freeway section prior to network-wide modeling, to better isolate errors from demand matrices associated with large networks, 2) assessing the transferability of the model's U.K.-calibrated parameters for use in the U.S., and 3) identifying critical parameters for calibration.

Field data employed in this series of tests were collected via an on-line intertie to inductance loop detectors in the field. Thirty-second traffic data (volume and occupancy) were collected over two hours in a morning-peak period on the southbound I-405 freeway, California, between the Culver and Jeffrey off-ramps. This site is 1,495 feet long, and instrumented with loop detectors as shown in Figure 1.

Two types of demand-related data files are necessary: 1) *demand*, which records the Origin-Destination (O-D) traffic flows in the network, and 2) *profile*, which determines the percentage of total traffic flow released onto the network in each time interval (e.g., during consecutive five-minute intervals). Field data were processed to generate the input data files, *demands* and *profile*, as well as output data figures, which were used for comparisons to simulated outputs.

Several simulation runs were conducted to identify critical variables that may significantly influence the performance of the model. Any

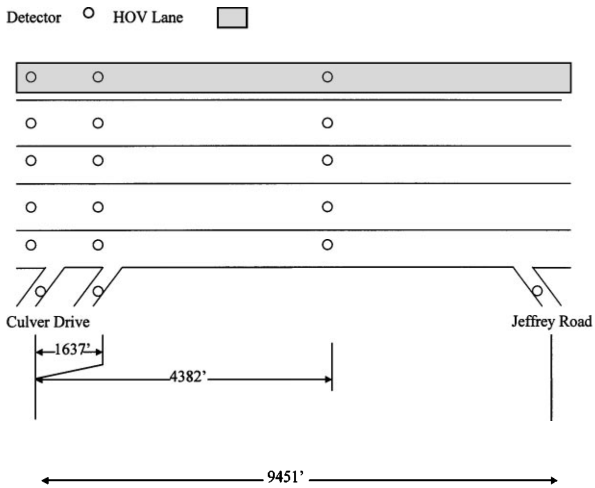

FIGURE 1. Detector layouts at the study site.

critical variables would be regarded as target variables in the calibration phase. In this procedure, the variables related to driver behavior (e.g., the distributions of aggressiveness and awareness), compositions of vehicles, driver's familiarity, and so on, were examined one-by-one. The numerical and graphical outputs from the simulation were then compared to corresponding real data.

Outputs used for comparison were classified into two groups: Origin/Destination related (total demand generated at each zone and total flow reaching the destination) and lane-specific measures (lane outflow, lane use, and lane-specific headway distribution). Default values for all of the variables in files *Vehicles* and *Driver* were kept unchanged (including vehicle population composition, carpool-eligible vehicles, and driver aggressiveness and awareness distributions). Other demand data, including *demands* and *profile*, were formulated on the basis of field data. Thirty simulation runs were executed and outputs were aggregated and analyzed. Outputs aggregated in this step for further comparison include: 1) total generation (total vehicles released in a two-hour period), 2) total attractions (total number of vehicles completing the short journey in a two-hour period), 3) lane-by-lane outflow (the number of vehicles collected from each lane in a two-hour period), and 4) percent lane-by-lane usage.

Results in Table 1 indicate substantial discrepancies in the number of vehicles being released to the section, the number of vehicles traversing the section, and lane use distribution. To bring the simulated values more in line with field observation on this simple section, a two-stage

TABLE 1. Output Data Comparison for Base and Stage 1 Calibration of the Simulator

Output parameter	Observed field data	Simulator prior to stage 1 calibration		Simulator after stage 1 calibration	
		Value	Error (%)	Value	Error (%)
Total demand	12869	11906	−7	12623	−1.9
Lane outflow (proportion)					
Lane1	2170 (0.18)	1052 (0.12)	−52	2092 (0.18)	−3.6
Lane 2	2900 (0.24)	1960 (0.22)	−32	2517 (0.22)	−13.2
Lane 3	3330 (0.28)	2695 (0.30)	−19	3399 (0.29)	−2.1
Lane 4	3625 (0.30)	3293 (0.36)	−1	3785 (0.31)	4.4
Total outflow	12025	9000	−25	11793	−1.9

procedure was employed. First, parameters within the model that represent aggregate characteristics of the driver population (e.g., the types of vehicles using the network and their percentages, the types and percentages of vehicles eligible for using the HOV lane, the distributions of aggressiveness and awareness levels across the population of drivers, and the proportions of vehicles leaving the freeway) were adjusted by trial-and-error to best match the observed data. Then, a systematic two-dimensional search process was conducted to determine values for the mean headway and reaction time that minimizes the discrepancy between the simulation output and field observations. This latter step specifically addressed the car-following and lane-changing models that are the heart of the microsimulator.

Stage 1: Results Based on Calibration of Aggregate Model Parameters

The primary purpose of this stage of the calibration was to try to both isolate and control factors that were primarily an artifact of the test section and its population of drivers and not explicitly related to the microscopic car-following and lane-changing models employed in the simulation; although they can have some influence on actual headway distributions, these factors primarily are associated with lane usage. In this way, any contamination on the part of such global factors to the systematic optimization of the parameters of the microscopic driver behavior models could be minimized.

In the *vehicles* file, the user can specify the types of vehicles using the network and their percentages, including the types and percentages of vehicles eligible for using the HOV lane; this latter value does not equal the observed percentage from field data, since vehicles "eligible" for using the HOV lane do not necessarily use it. Consequently, the observed field data are not sufficient to determine the appropriate parameter value. Rather, the percentage of eligible vehicles (i.e., multiple-occupancy vehicles in the traffic stream) had to be inferred from empirical experimentation based on the target number of vehicles actually observed in the HOV lane. In addition, the default values for parameters controlling for the types of distributions of aggressiveness and awareness levels across the population of drivers, based on research at the Transportation Research Laboratory (Jeffreys, 1994), were varied. Both the number and configuration of destination zones beyond the immediate vicinity of the short freeway section were varied to better capture the boundary effects of downstream off-ramps (not included in the test section) on driver lane choice within the section. Numerous runs with different combinations of distributions of the appropriate parameters were conducted until the lane use discrepancies were improved to an acceptable level. Although the actual changes to the parameter values are unimportant to this stage of the investigation, the effect of this stage of calibration was to bring the simulated total demand, outflow, and lane use distributions into general agreement with field observation (Table 1).

Graphically, Figures 2 and 3 present the distributions of lane-by-use generated using observed (from field Vehicle Detection Stations or

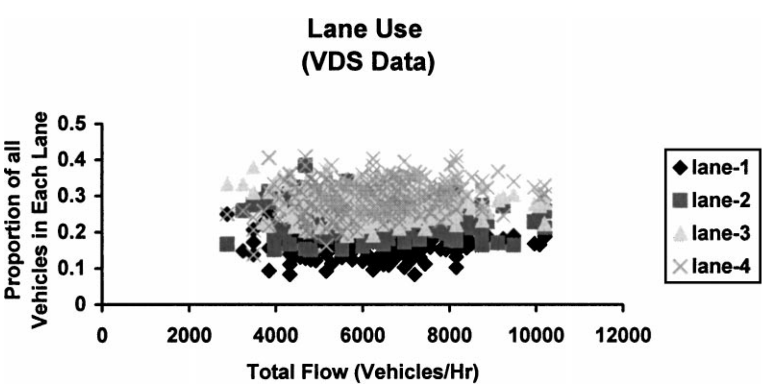

FIGURE 2. Lane use on 4-lane freeway segment (loop data).

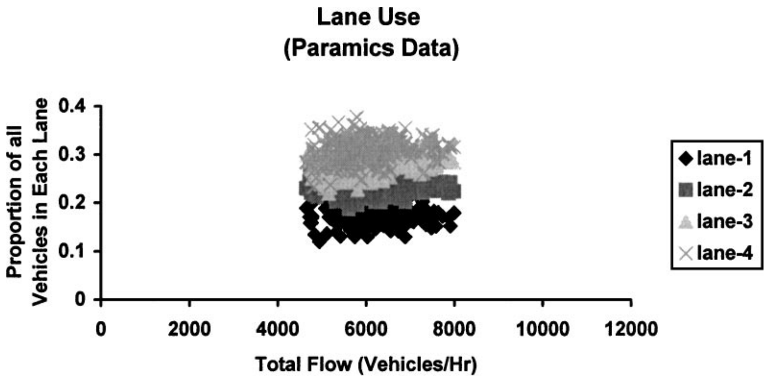

FIGURE 3. Lane use on 4-lane freeway segment (simulator).

VDS) and simulation data, respectively. Each data point in either figure represents the ratio of the lane traffic count in a given lane to the sum of lane traffic counts collected from the midstream detectors during a given 30-second time step. Compared to the observed data, the simulated data points shown in Figure 3 distribute more densely, mostly in the range between 5000 veh/hr and 7000 veh/hr. However, cluster centers match in both figures. The flow-density relationship determined from the actual loop data, shown in Figure 4, closely approximates the simulation, shown in Figure 5, indicating that the simulator properly captures the fundamental flow-density relationship.

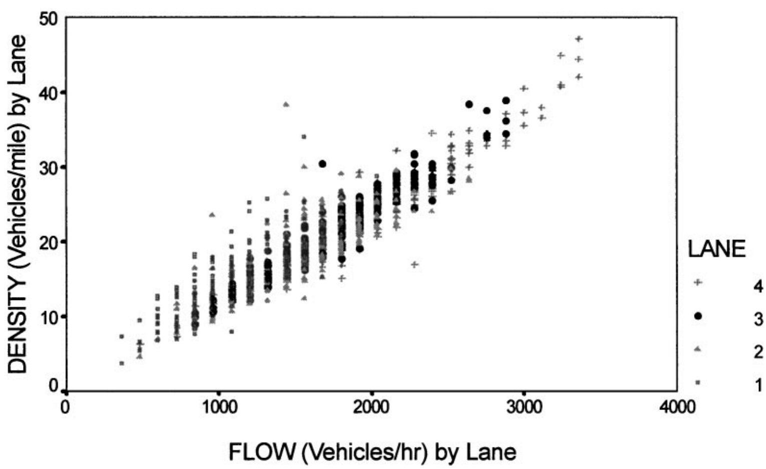

FIGURE 4. Flow-density characteristics determined from VDS data.

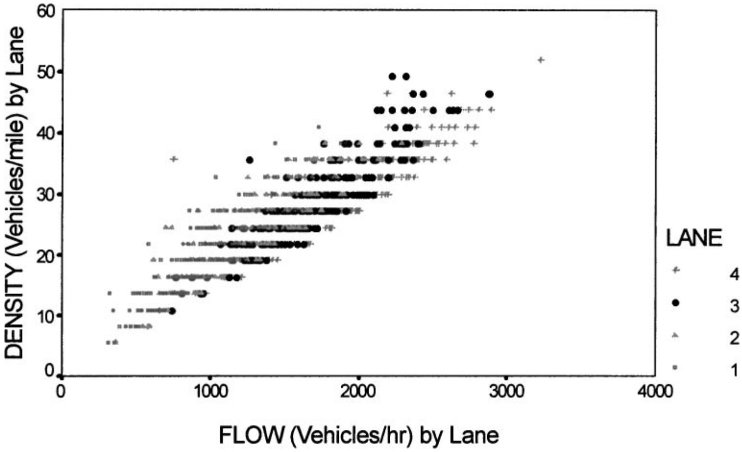

FIGURE 5. Flow-density characteristics determined from the simulator.

Stage 2: Validation Results Based on Calibration of Microscopic Driver Behavior Parameters

In this section, we examine the heart of the microsimulator, more specifically, the car-following and lane-changing models. The objective is to identify key parameters for calibration and validation using California data. Other parameters already adjusted up until this stage are kept the same.

The default values for the mean headway and the mean reaction time are both set to one second. Early related studies (Greenshields, 1947; King & Wilkinson, 1973; Kunzman, 1978), recommend headway values of approximately two seconds in the fields of urban traffic control and related applications. The value suggested by the Highway Capacity Manual, 1985 (HCM) is also two seconds. For reaction time, numerous early experiments had been conducted to measure the range of reaction time. The range of 0.3–2.0 seconds (0.66 seconds is the median) was suggested in a study by Johansson and Rumar (1971).

Determining appropriate values for the mean headway and reaction time is a two-dimensional search process that minimizes the discrepancy between the simulation output and field observations, and that requires numerous simulation runs for every point in the search space. Because the shape of the search space is unknown, a simple empirical iterative search process was employed that alternately searches in one direction, holding the other value fixed. One parameter, say β, is fixed at its best-known value, possibly from the previous search iterations, while one-directional search is performed on the other parameter, say α, within reasonable upper and lower bounds. Once the optimal α value is found, it gets fixed and its range empirically reduced around this value, and the search for β starts. The process continues until improvements diminish.

In this particular application, the appropriate parameters are the mean headway (H) and the mean reaction time (R), which were calibrated in six iterations of the procedure described above. The average absolute error of simulated traffic volumes was employed as an index to evaluate the system's performance. Based on this procedure, the calibrated values for mean headway (H) and the mean reaction time (R) were found to be 1.65 seconds and 0.42 seconds, respectively. Table 2 presents aggregated comparison results from ten simulations using these calibrated parameter values; the results from the Stage 1 calibration procedure are included for comparison purposes.

Table 2 shows a slight improvement in the total demand generated by the simulator, and the corresponding outflow. The probable mechanism for this improvement is the reduction in the virtual queue through adjustments to the mean headway and reaction time, allowing more of the actual demand to be generated and passed through the system.

Table 3 presents a summary of the comparison between field observation and model results on an independent validation data set. Results indicate that the average error in the validation phase is greater than the case of testing after calibration, which is expected. Nevertheless,

TABLE 2. Output Data Comparison for Stage 1 and Stage 2 Calibration of the Simulator

Output parameter	Observed field data	Simulator after stage 1 calibration		Simulator after stage 2 calibration	
		Value	Error (%)	Value	Error (%)
Total demand	12869	12623	−1.9	12949	0.6
Lane outflow (proportion)					
Lane 1	2170 (0.18)	2092 (0.18)	−3.6	2160 (0.17)	−0.5
Lane 2	2900 (0.24)	2517 (0.22)	−13.2	2849 (0.24)	−1.7
Lane 3	3330 (0.28)	3399 (0.29)	−2.1	3351 (0.28)	0.6
Lane 4	3625 (0.30)	3785 (0.31)	4.4	3772 (0.31)	4.0
Total outflow	12025	11793	−1.9	12132	0.9

TABLE 3. Output Data Comparison for Validation Runs of Calibrated Model

	Validation runs														
	Case 1: 07:00–09:00 (08/04/97)			Case 2: 16:00–18:00 (09/09/97)			Case 3: 11:00–13:00 (10/22/97)			Case 4: 07:00–09:00 (07/17/97)			Case 5: 16:00–18:00 (09/26/97)		
Output parameter	Field	Model	Error (%)	Field	Model	Error (%)	Field	Model	Error (%)	Field	Model	Error (%)	Field	Model	Error (%)
Total demand	13613	14150	3.9	19511	17818	−8.7	11279	11561	2.5	14440	14936	3.4	20426	18038	−11.7
Lane outflow (proportion)															
Lane 1	2105 (0.17)	2031 (0.16)	−3.5	3745 (0.22)	2599 (0.17)	−30.6	2358 (0.17)	1766 (0.17)	−25.1	2211 (0.17)	2248 (0.17)	1.6	3781 (0.22)	2619 (0.17)	−30.7
Lane 2	3095 (0.25)	2920 (0.23)	−5.6	4085 (0.24)	3670 (0.24)	−10.2	2573 (0.24)	1974 (0.19)	−23.3	3121 (0.24)	2909 (0.22)	−6.8	4125 (0.24)	3697 (0.24)	−10.4
Lane 3	3467 (0.28)	3555 (0.28)	2.5	4426 (0.26)	4434 (0.29)	0.0	2894 (0.27)	2702 (0.26)	−6.6	3641 (0.28)	3701 (0.28)	1.7	4469 (0.26)	4467 (0.29)	0.0
Lane 4	3714 (0.30)	4189 (0.33)	12.8	4767 (0.28)	4587 (0.30)	3.8	2894 (0.27)	3948 (0.38)	−36.4	4031 (0.31)	4363 (0.33)	8.2	4812 (0.28)	4621 (0.30)	−4.0
Total outflow	12381	12695	2.5	17023	15290	−10.2	10719	10390	−3.1	13004	13221	1.7	17187	15404	−10.4

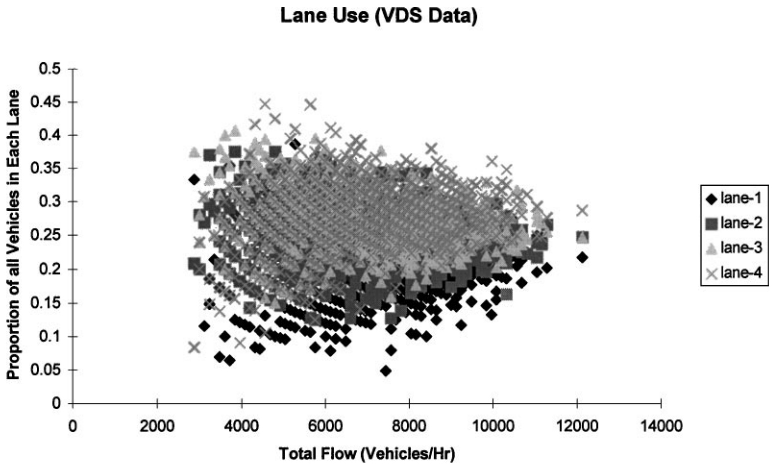

FIGURE 6. Validation: lane use (loop data).

error magnitudes compare favorably with models of this type. There is some evidence that error magnitudes increase as demand increases, pointing to the issue of vehicle stacking in memory awaiting a suitable gap to be released onto the network. This occurs when the network is overly congested as the vehicle release mechanism, which is based on gap acceptance, fails to release the vehicles onto the network, keeping them in memory until a suitable gap appears. In real networks, this phenomenon does not happen because of the presence of a significant number of local and collector streets for vehicles to be released on. Streets smaller than arterials are usually ignored in simulation analysis, which is the case in this investigation. Results of total outflow also support this observation. The lane use percentages compare favorably to field observations; errors are slightly higher than in the case of testing on calibration data. Figures 6 to 13 present graphical comparisons between the microscopic output of the simulator and field (VDS) data during the validation phase.

Validation on a Freeway/Arterial Network

Validation on a single segment, as above, has the advantage that it allows for the evaluation of the underlying car-following/lane-changing models while minimizing exogenous errors from Origin-Destination (O-D) estimation, signal timing coding, and the like. In this phase of validation, the performance of the simulator on a full network, including freeways as well as surface streets, is evaluated. Using the parameters as calibrated previously, the model is applied to a portion of the California ATMS Testbed network, shown in Figure 14. The network is located in

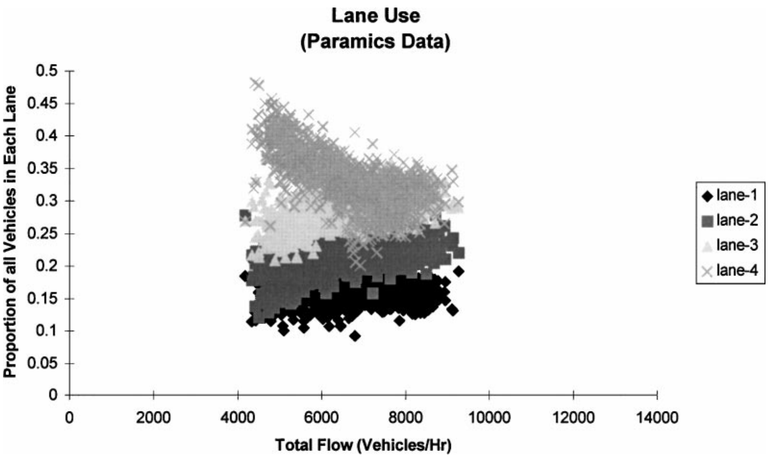

FIGURE 7. Validation: lane use (simulator).

the city of Irvine, bounded by a triangle of two freeways (I-405 and I-5) and one major arterial (Jeffrey Road).

The network was coded with very specific detailing of the geometry of freeways (including radii of curvature, the lengths of on-ramps and off-ramps, lane width, lane restrictions for the HOV lanes, and lane configurations) and surface streets (including the numbers of lanes for

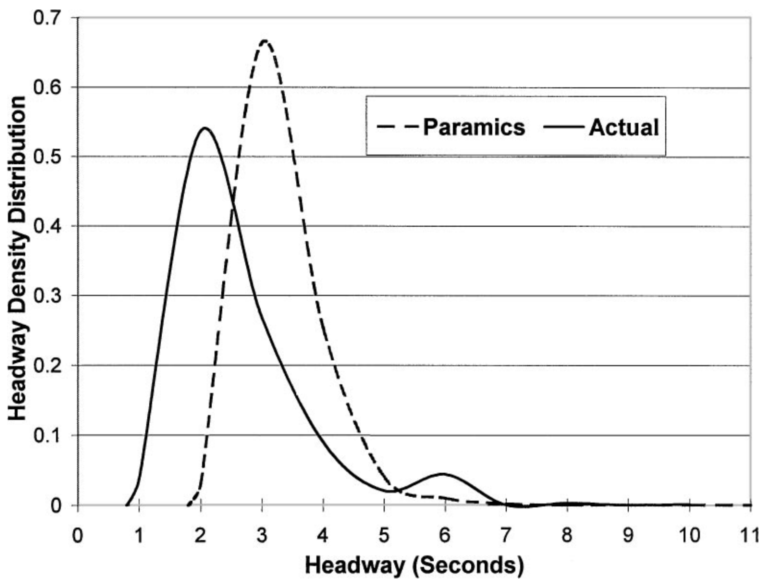

FIGURE 8. Validation: headway distributions for lane 1.

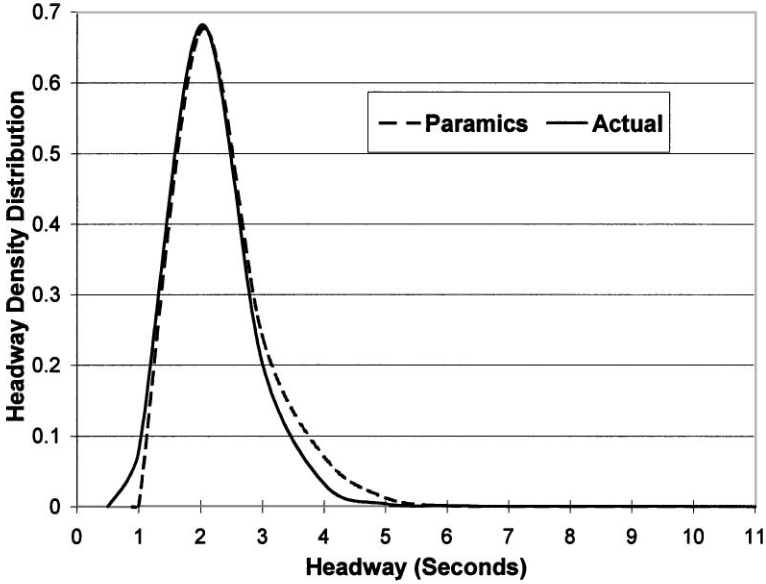

FIGURE 9. Validation: headway distributions for lane 2.

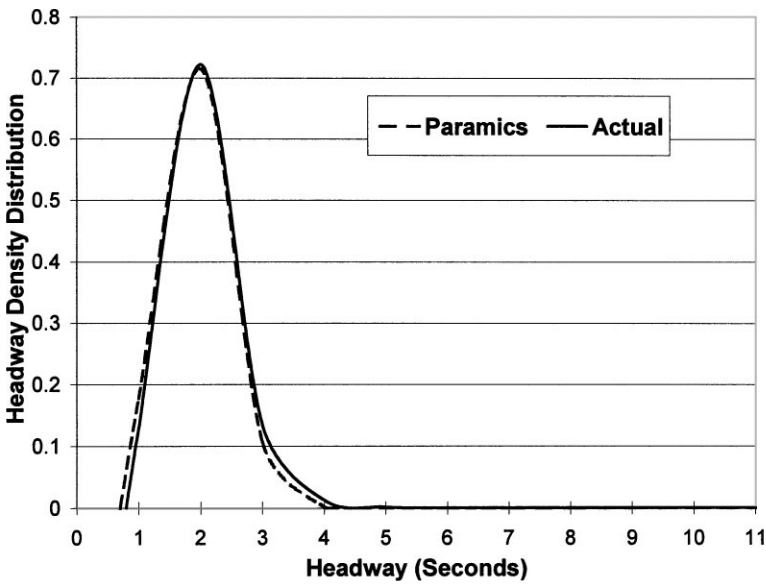

FIGURE 10. Validation: headway distributions for lane 3.

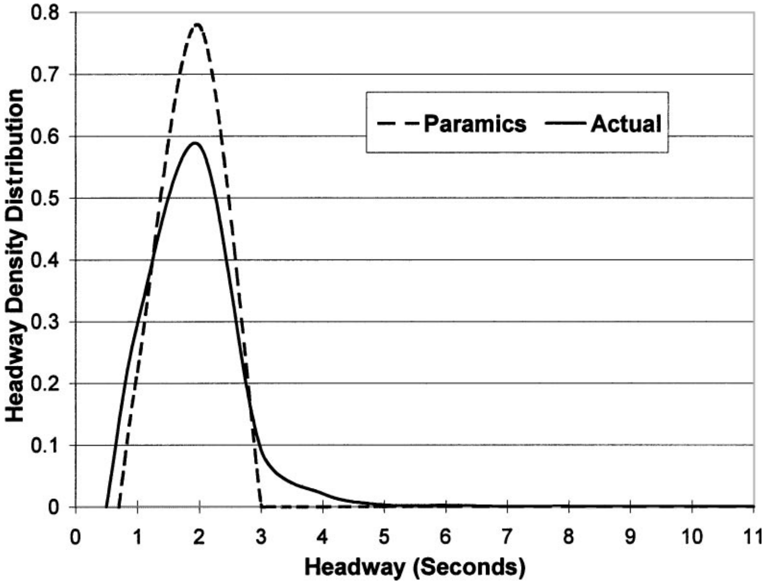

FIGURE 11. Validation: headway distributions for lane 4.

all the approaches to intersections, curb location within each intersection, lane composition for each approach, stop-line location, and signal timing plans).

The O/D data were based on PM peak O/D data obtained from the Irvine Transportation Analysis Model (ITAM); 15-minute dynamic O/D

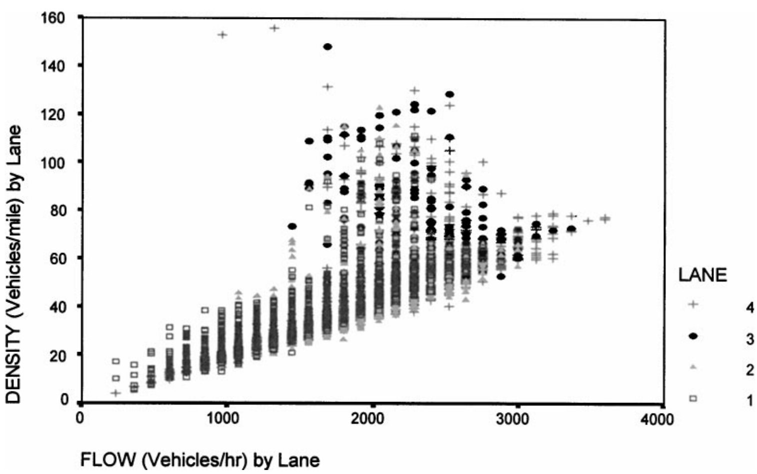

FIGURE 12. Validation: flow-density characteristics (loop data).

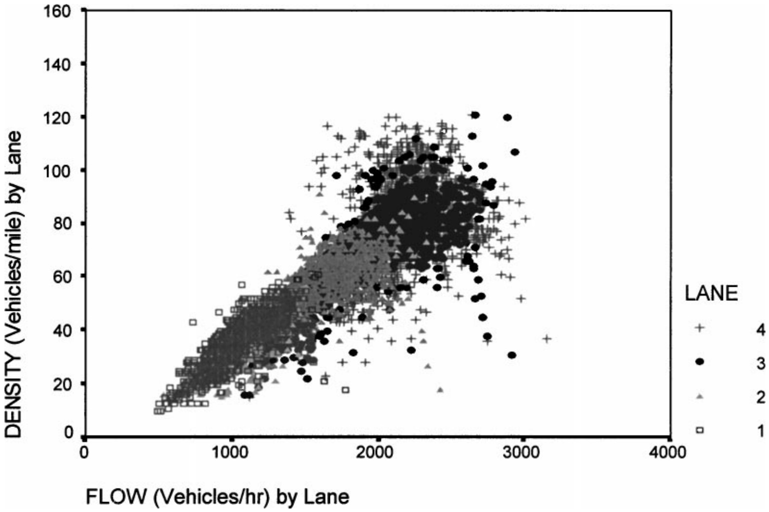

FIGURE 13. Validation: flow-density characteristics (simulator).

data were developed from the PM peak O/D data through an iterative process using the Contram and Comset traffic models. Contram assigns the O/D demand onto the network and outputs the number of vehicles (referred to as packets) using specific routes. Output from Contram serves

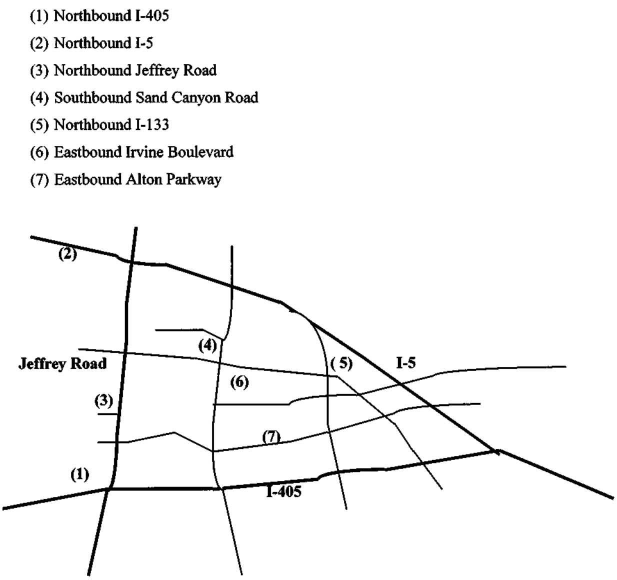

FIGURE 14. The detector layouts for collection of transition link volumes.

as input into Comset, together with observed network link count information, to provide updated 15-minute O/D data for the PM peak. These updated data are re-input to Contram, and the cycle is repeated until convergence is attained. The final output is: 15-minute O/D data, and packet information containing departure times and identified routes. A total of 106 O/D zones were coded.

Validation of the network performance of the model was based on comparisons involving: 1) total demand and supply volumes associated with each zone, and 2) volumes collected at specific links within the network. The total volumes attracted to and generated by each zone were measured by encoding two detectors for each of the 106 zones in the simulated network. Link volumes were extracted at seven key locations within the network corresponding to the location of field loop detectors. Locations of these detectors within the network are identified in Figure 14.

Using the detector layout described above, three simulation runs, each for a duration of three hours, were executed. For each three-hour simulation run, lane volumes were measured from each loop detector and updated every one hour. Data associated with each run were aggregated, and the average volumes were used for comparison to field data.

Figures 15(a)–(b) and 16(a)–(b) show the errors between the simulated and real (VDS) zonal traffic generation (departure) and attraction

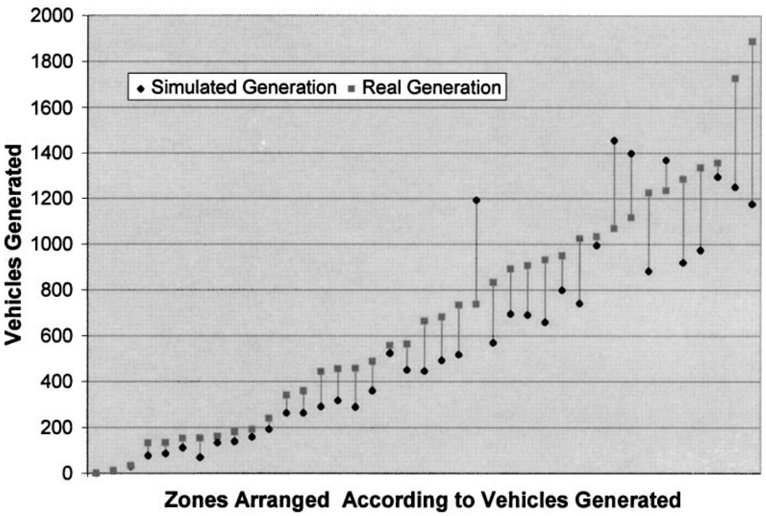

FIGURE 15(a). Comparison between real (loop) and simulated (model) zonal vehicle generation.

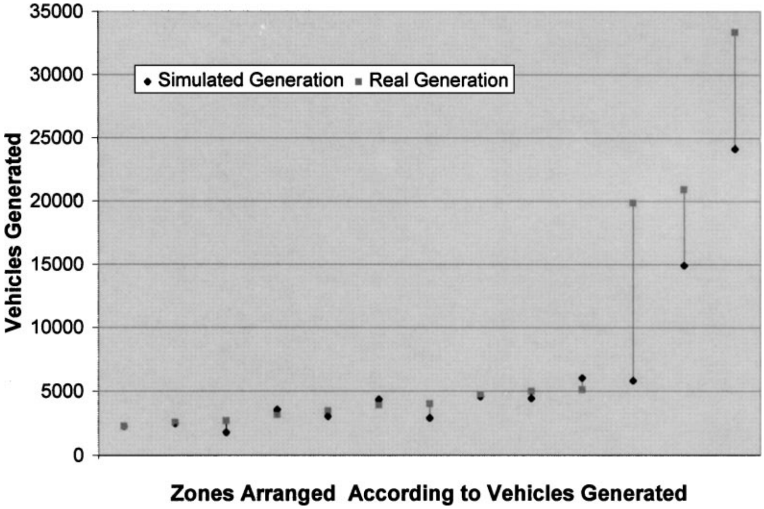

FIGURE 15(b). Comparison between real (loop) and simulated (model) zonal vehicle generation.

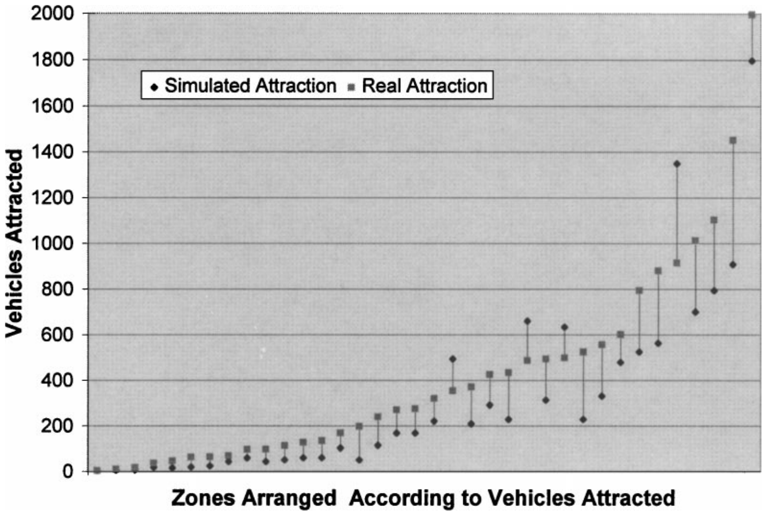

FIGURE 16(a). Comparison between real (loop) and simulated (model) zonal vehicle attraction.

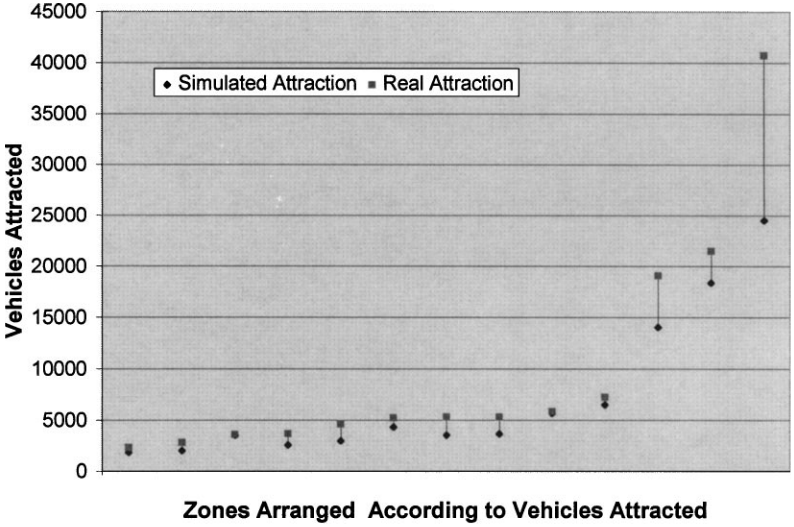

FIGURE 16(b). Comparison between real (loop) and simulated (model) zonal vehicle attraction.

(arrival) over the three-hour validation period.[2] Figure 17 compares loop detector readings from the simulator to those from the field. The errors are clearly biased toward under-estimation of the actual volumes. Most likely, these errors are caused by the vehicle release mechanism onto the network. If vehicles get stacked in memory as explained earlier, the simulation period terminates without releasing all the vehicles, which directly affects the number of vehicles propagating through the network and ultimately reaching the destination. It is noted that all vehicles eventually get generated onto and traversed through the network, but slightly after the three-hour simulated time period. The length of time required to process the additional vehicles depends on the level of "virtual congestion" (the number or vehicles in memory awaiting release) at each zone; some zones take longer than others. Visual inspection of the simulator during operation indicated that the routing of vehicles, and their behavior near and within intersections are plausible, lending additional support to the conclusion that the above error levels are due mainly to the vehicle releasing problem—a problem that should be correctable with greater attention to the defining of zones and their links to the network.[3]

[2]Only those zones for which the generation (or attraction) is non-zero are included.

[3]In this particular study, the selection of zones was restricted by the source of the O-D information.

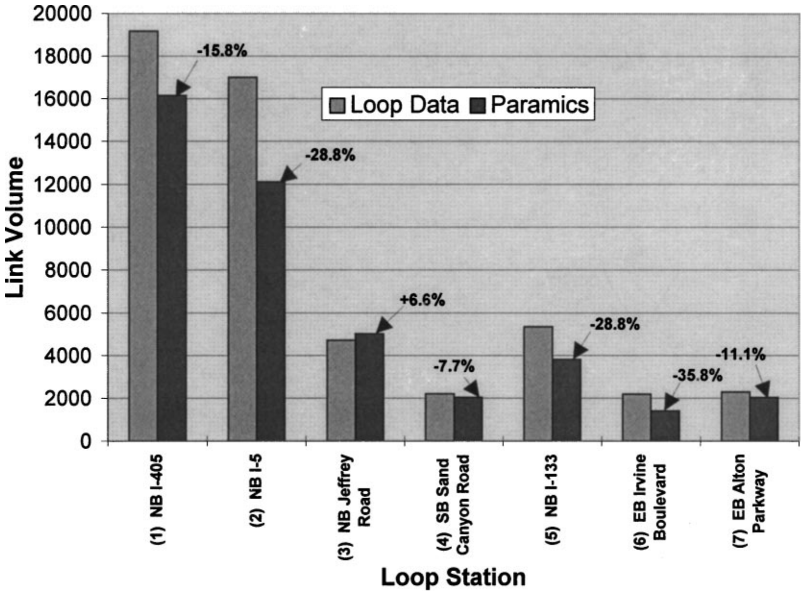

FIGURE 17. Comparison between real (loop) and simulated (model) link volumes.

SUMMARY, CONCLUSIONS AND RECOMMENDATIONS

In this study a new and potentially successful ITS-oriented micro-simulator is evaluated relative to its reproduction of field observation. The evaluation was based on calibration of selected key parameters underlying the simulator's car-following and lane-changing models, followed by its validation both on a small freeway section and on a complete network of freeways and surface streets in Southern California.

Features of the software potentially make it a plausible "shell" or "framework" for a comprehensive and extensive transportation simulation laboratory. It offers two very important features: high performance and scalability. In particular, the development of a series of the Application Programming Interfaces (APIs) is regarded to be an important capability of the software, providing the gateway for researchers to the heart of the software, without having to deal with the underlying proprietary code, both allowing researchers to override the simulator's default models, and also to interface complementary modules to the simulator. Such complementary modules could be any ITS applications as signal optimization, adaptive ramp metering, incident detection and management, or any other conceivably important ITS applications.

In terms of objective performance, the calibrated model performed well on the small freeway section, but errors were biased toward underestimation of traffic volumes when applied to an entire network. It was concluded that these errors are caused by a single coding problem—the stacking of vehicles in memory before finding a suitable gap to be released onto the network; it is conjectured that this problem could be greatly alleviated by including the minor arterials and local streets in the network coding. This "virtual congestion" causes the simulator to release fewer vehicles than in real life, and, consequently, less vehicles traverse the network and reach the destination within the simulation period. If the simulator is allowed to run for longer than the demand period, the remaining vehicles eventually make it to and through the network, taking longer than the specified demand period.

REFERENCES

Greenshields, B. D. (1947). *Traffic performance at intersections*. Yale Technical Report #1, Yale University, New Haven, CT.

Jeffreys, W. (1994). *Motorway lane discipline: Motorway traffic modeling study*. Technical report, Transportation Research Laboratory, Crowthorne, Berkshire, UK, March 1994.

Johansson, G., & Rumar, H. (1971). Drivers' brake reaction times. *Human fact.,* 13(1), 23–27.

King, G., & Wilkinson, M. (1973). *Relationships of signal design to discharge headway*. Report 615, Transportation Research Laboratory, Crowthorne, Berkshire, UK.

Kunzman, W. (1978). Another look at signalized intersection capacity. *ITE Journal*, August 1978.