STOCHASTIC VARIABILITY IN MICROSIMULATION MODELING RESULTS AND CONVERGENCE OF CORRIDOR-LEVEL CHARACTERISTICS

Sarah Ellie Ziems
School of Sustainable Engineering and the Built Environment, Arizona State University
Room ECG252, Tempe, AZ 85287-5306. Tel: (480) 965-3589; Fax: (480) 965-0557
Email: sarah.ziems@asu.edu

Bhargava Sana
School of Sustainable Engineering and the Built Environment, Arizona State University
Room ECG252, Tempe, AZ 85287-5306. Tel: (480) 965-3589; Fax: (480) 965-0557
Email: bsana@asu.edu

Joseph Plotz
School of Sustainable Engineering and the Built Environment, Arizona State University
Room ECG252, Tempe, AZ 85287-5306. Tel: (480) 965-3589; Fax: (480) 965-0557
Email: jplotz@asu.edu

Ram M. Pendyala
School of Sustainable Engineering and the Built Environment, Arizona State University
Room ECG252, Tempe, AZ 85287-5306. Tel: (480) 727-9164; Fax: (480) 965-0557
Email: ram.pendyala@asu.edu

Submitted for Presentation and Publication to
Committee ADB50: Transportation Planning Applications
90th Annual Meeting of the Transportation Research Board

Word count: 7066 (text) + 6 (tables/figures) x 250 = 8566 equivalent words

August 2010

ABSTRACT

Planning agencies around the world are increasingly turning to disaggregate and often more flexible microsimulation models to simulate behaviors of individual agents in urban contexts. The practical application of microsimulation models is often affected by the stochastic variability in results obtained across multiple runs of the model system, as the random number seed sequence that drives the simulation is varied from one run to the next. The analysis presented in this paper aims to quantify the stochastic variation arising from repeated runs of a microsimulation model system. This is accomplished by repeatedly running the router and microsimulator modules of TRANSIMS, a microsimulation model system capable of simulating activity-travel and traffic patterns for large regions. The modules are applied to the Greater Phoenix metropolitan area in Arizona, and traffic characteristics on two short heavily traveled stretches of highway corridors are examined with respect to variability over 20 simulation runs. The convergence of the successive average for various traffic characteristics is examined over the 20 simulation runs, and conclusions are drawn with respect to the extent of variability present, the factors that might influence the variability, and the number of simulation runs needed to achieve stable and robust results. In general, it is found that variability is small (particularly in uncongested conditions), the number of simulation runs needed depends on the phenomenon being modeled and is context-specific, and 20 simulation runs may be sufficient for some, but not all, microsimulation model applications.

Keywords: planning applications, microsimulation models, stochastic variability, random processes, TRANSIMS application, model convergence and stability

INTRODUCTION

Modeling of urban systems has steadily progressed into the era of microsimulation-based approaches. There is a vast body of literature on the development of microsimulation models of activity-travel demand, land use dynamics, and time-dependent traffic patterns on networks (Axhausen and Garling, 1992; Vovsha et al, 2002; Waddell et al, 2003; Mahmassani, 2001). The primary motivation behind the adoption of microsimulation approaches is that such model systems are capable of simulating decision-processes and choice behaviors at the level of the individual decision-maker or behavioral unit while accounting for the interactions, constraints, opportunities, linkages, and dependency structures that inevitably shape individual behavior (Goulias and Kitamura, 1992; Bhat et al, 2004; Salvini and Miller, 2005; Balmer et al, 2006). By adopting such approaches, one is able to avoid aggregation biases associated with more traditional models of urban land use and transportation that do not simulate behaviors of individual agents. As one is able to simulate behaviors at the level of the individual decision-maker, microsimulation approaches offer a robust framework for analyzing distributional impacts of policy interventions (Castiglione et al, 2006).

Microsimulation approaches are increasingly being implemented in practice due to advances from a methodological, computational, and data availability standpoint. However, the issue arising with microsimulation model systems is that they are invariably comprised of numerous probabilistic choice and decision models. Analysts accustomed to obtaining the same outcome every time they run traditional aggregate model systems are often uncertain as to how to interpret the stochastic variation associated with multiple runs of a microsimulation model system employing different random number seeds. This is one issue associated with the application of microsimulation model systems in practice.

Running a microsimulation model system multiple times offers the ability to obtain stable and robust model outputs that facilitate comparisons across scenarios. The output of a single simulation run represents a realization of the system conditions on a day. Consistent with the fact that there is considerable day-to-day variability in system conditions, it is possible to interpret the stochastic variability in microsimulation outputs as representative of the day-to-day variability observed in the real world. However, stochastic variability in simulation model systems may mask the true impacts of policy and modal investment scenarios. When simulation models are run only once for each scenario, it is not possible to isolate differences due to changes in scenario inputs from differences attributable to stochastic random variability in the microsimulation process itself. In running a model just once for multiple scenarios, the critical question that arises is to what extent each of these sources of variability contributes to differences in outputs. Unless one can be sure that sufficient runs of the model have been performed, such that a comparison of stable output values can be made across scenarios, it is not possible to conduct accurate scenario analysis. Although hardware and software advances have facilitated the implementation of microsimulation models, the model run times can be quite long. Planning agencies that are under tight schedules to produce long range transportation plans and policy studies in response to local and federal regulations, and numerous requests from policy makers, cannot afford the luxury of making multiple runs. In this context, the potentially long run times, and the desire to know how many times a stochastic microsimulation model system needs to be run to obtain stable averages, constitutes the second issue motivating this paper.

There are other issues that certainly merit attention in the application of microsimulation models, but issues related to stochastic variability of microsimulation model outputs constitute the focus of this paper.

This paper aims to contribute to a greater understanding of the stochastic variability in microsimulation modeling processes by quantifying the variability arising from repeated runs of a traffic microsimulation model system that is capable of developing route plans for trips in an urban area, and then microsimulates vehicle movements on the network as the route plans are executed. The model components used in this study are those specifically embedded in the TRANSIMS (Transportation Analysis and Simulation) model system, although the experiment reported in this paper may be undertaken using any microsimulation model. Repeated runs of the Router and Microsimulator modules of TRANSIMS are performed using different random number seeds driving the simulation of choice behaviors and vehicular movements in these model components. The resulting stochastic variability, and the extent to which stability in results is achieved at the end of a certain number of runs, are evaluated and reported.

The issues addressed in this paper are not new; there is a body of literature that has examined stochastic variability in microsimulation modeling contexts and proposed approaches to address it. The intent here is to further add to the body of knowledge on this topic so that practitioners can be more informed on ways to handle stochastic variability in outcomes of microsimulation models. The next section offers a brief review of the literature to highlight the progress that has already been made. The third section describes the TRANSIMS modeling process adopted in this study, while the fourth section describes the geographic area used for this study and the details of the experiment conducted. The fifth section offers detailed results while the sixth section offers concluding thoughts.

STOCHASTIC VARIATION IN MICROSIMULATION MODELING

The literature that speaks to the issues of stochastic variability in microsimulation modeling is constantly growing. Gibb and Bowman (2007) consider the notion of simulation error, i.e., the Monte Carlo simulation error arising from the discretization of choice behaviors in microsimulation model systems of activity-travel demand. They term simulation error as the random noise problem, and describe techniques to establish convergence in these models. They report that minor adjustments in the method of successive averaging, where outcomes from multiple runs are successively averaged to convergence, can substantially improve computational efficiency. Walker (2005) proposed a microsimulation modeling environment to make such approaches more accessible to planners. She notes that the simulation error is actually an appealing feature of microsimulation models in that it allows one to estimate the size of the error associated with point forecasts and generate confidence intervals based on the distribution of outcomes from multiple runs. Vovsha et al (2008) also describe some practical approaches to achieve equilibrium in activity-based microsimulation models. In addition to successive averaging of outcomes over multiple runs, they suggest enforcement techniques such as reusing random number seeds and gradual freezing of travel choice dimensions once they exhibit stability in outcomes.

A more detailed study was undertaken by Castiglione et al (2003) who investigated the amount of stochastic variation arising from random simulation error in the San Francisco County Transportation Authority (SFCTA) activity-based travel demand model. Their experiments were also aimed at determining the number of runs required to obtain stable and reliable results. They ran the SFCTA model 100 times; each time the model was run, only the sequence of random numbers used to simulate individual choices in the model system was changed. The variability in the output was quantified based on two factors – the type of sub model (i.e., tour generation, destination choice, and so on); and the geographic resolution (such as zone or county level) at which the variability in outcomes is measured. For each combination of the two factors, the percent difference between the successive average of a particular output and the final mean (after 100 simulation runs) was computed and plotted after each simulation run. They found that all model components demonstrated a high level of stability even at the highest geographic resolution (zone level). The variability at lower geographic resolutions (county and neighborhood levels) was relatively lower, suggesting that aggregation over space reduces (masks) variability in outcomes. They also found that a relatively small number of runs was sufficient for the outputs to converge to a stable value. However, they do note the potential pitfall associated with running a microsimulation model only once. The outputs from individual model runs could vary as much as 10 to 25 percent from the successive average computed after 100 simulation runs. The authors also indicate that the number of simulation runs required to achieve stability in model outcomes is dependent on the model system and the particular planning application. Finally, they note that their findings apply in the context of the SFCTA activity-based model and it would be useful to conduct similar analysis (such as the one in this paper) for other model systems.

Lawe et al (2009) implemented TRANSIMS for Chittenden County in Vermont. They conducted various sensitivity tests, including tests to assess the sensitivity of model results to changes in the random number seed. Five model runs, each with a different random number seed, were performed and the variation in results (traffic volumes and average speeds) for 10 links in the network was examined carefully. The coefficient of variation (CV) was computed for every hour of the day for each of the 10 links, and the average CV value was computed for a full 24 hour period. It was found that there was very little variation in the average daily CV among the five different runs for both traffic volumes and average speeds. Overall, it was concluded that, for medium-sized areas with little or no congestion, microsimulation models may not be that sensitive to variations in the random number seed. Similar results were also reported by Veldhuisen et al (2000), where the effects of simulation error on travel demand estimates were found to be negligibly small. They found a very high correlation across outcomes from successive runs of the model system and also note, similar to Castiglione et al (2003), that the Monte Carlo error is higher at higher geographic resolutions.

Overall, while there has been some evidence on variability due to repeated runs of a microsimulation model system with different random number seeds, additional tests are needed to accumulate a larger body of knowledge on this issue. Practitioners need such information to help them apply microsimulation models in planning studies appropriately – too few runs could result in reporting results that have not yet reached a stable and reliable set of values, while too many model runs could result in computational inefficiency and unnecessary expense. Additional evidence on this issue can help practitioners better explain microsimulation outputs to policy makers. This paper also offers insights into stochastic variability that occurs in the context of trip routing and traffic simulation, as opposed to some of the previous efforts that have focused on the activity-based travel demand model components.

THE TRANSIMS ROUTER AND MICROSIMULATOR PROCESSES

In response to legislative action in the United States in the 1990's that called for a more disaggregate approach to transportation planning models, the Los Alamos National Laboratory developed TRANSIMS, the TRansportation ANalysis and SIMulation System, under contract to the US Department of Transportation. Since the initial development effort, TRANSIMS has evolved into an open-source modular software package made available freely to the research and practitioner community for transportation systems analysis (http://www.transims-opensource.net/).

TRANSIMS incorporates a series of modules that allow one to simulate activity schedules, travel demand, and network dynamics in a unifying framework. TRANSIMS includes a population synthesizer for generating a regional synthetic population, an activity scheduler that synthesizes activity schedules for each individual in the synthetic population using a classification tree approach, a router that creates trip route plans for the activity schedules, and a microsimulator that simulates vehicular movements on the links of the network as travelers execute their plans. TRANSIMS is also capable of accepting regional origin-destination trip tables as input; when such trip tables from a traditional four-step travel demand model are provided to TRANSIMS, the first two steps (population synthesis and activity schedule generation) can be bypassed. The router and microsimulator modules take the origin-destination trip tables and parse them, both in space and time, to create synthetic trips that must be routed and simulated on the network. In this particular study, this approach is used, i.e., regional origin-destination trip tables are directly input for routing and microsimulation. Thus, the assessment of stochastic variability in this paper focuses exclusively on the network supply modeling side of the enterprise, keeping demand fixed.

The TRANSIMS router is a software module that provides a path to each vehicle traveling from an origin (activity location) to a destination (another activity location). The path is chosen based on the impedance assigned to every link in the system, with vehicles choosing the path of least impedance. As travelers do not have perfect information about travel impedances on links/paths, each traveler in the real-world context is assumed to have a perceived impedance. The router module is equipped with the ability to assign random impedance values that vary across travelers, reflecting the real-world situation. The analyst must provide an input value to TRANSIMS specifying how much perceived impedance may vary randomly around the true impedance for each traveler. A random number seed, provided by the user, dictates the assignment of perceived impedance in each router module run, thereby causing travelers to have different perceived impedance values in each router run. In turn, this may cause travelers to choose different paths between the same origin-destination pairs across runs.

The TRANSIMS microsimulator module is an agent-based cellular automata model that simulates the movements of individual vehicles on the network in one-second time steps. The algorithm begins by segmenting each link into a series of cells, where each cell is approximately equal to the length of a single vehicle. A vehicle's forward motion occurs when the cell directly in front of it, with respect to its intended path, becomes available. As traffic is routed through the network, the microsimulator determines vehicle position and speed characteristics in each time step while accommodating for the random nature of driver behavior. For example, the microsimulator uses a random rounding procedure to calculate reaction time for each driver and uses a probabilistic procedure to determine whether a vehicle will allow another vehicle to change lanes into its path. Essentially, in this module, the random number seed plays a role in determining the individual vehicle trajectories as traffic propagates through the network.

As one would expect, simulating traffic flow through a network using the router and microsimulator modules of TRANSIMS is an iterative process. The router and microsimulator have to be run iteratively until convergence is achieved in the outputs (from a single run). In the router, the process begins by assigning a path to each vehicle (based on an initial set of input speeds) and then executing trips on the assigned paths to find a base level of impedance. Vehicles that are greatly delayed due to the presence of a congested link can be chosen to be rerouted, or assigned a new, less congested path in the subsequent iteration. The selection of vehicles for rerouting is based on a user-specified minimum percent travel time decrease that must be met. Among those for which the minimum threshold is met, vehicles can be chosen randomly. This selection process depends on the random number seed used for the simulation run. In spirit, this process may be considered as representative of the learning process that drivers go through as they experience congestion along different routes; some drivers will choose to try alternate paths, while others may be less willing to do so and remain on their less optimal travel path.

A common practice in TRANSIMS implementations is to first stabilize the router using a simple shortest path algorithm (based on length and free flow speeds), prior to initiating and stabilizing the microsimulator. Figure 1 illustrates the process in which the router and microsimulator modules are run sequentially while achieving stabilization in each component. A simulation begins with an initial run of the router. Once this is complete, the router stabilization begins by calculating link delays (due to congestion). Based on potential improvements in travel time that can be achieved, trips are randomly selected for rerouting in a subsequent iteration. Trips chosen for rerouting are rerouted and combined with all other trips (that were not chosen for rerouting) to complete an iteration of the router stabilization process. The iterative process continues until the number of trips chosen for rerouting falls below a certain user-specified threshold value. For this study, it was found that 12 iterations of the router stabilization process and 8 iterations of the microsimulator stabilization process were needed to achieve convergence, amounting to a total of 20 iterations for each trial run (using a different random number seed).

DESCRIPTION OF EXPERIMENT

The experimental investigation was conducted in this study for the Maricopa County region of Arizona, which corresponds to the Greater Phoenix metropolitan area. This metro region has a population of more than four million people and is the twelfth most populated metropolitan area in the country. The area, similar to many other large urban regions in the country, experiences heavy traffic volumes and congestion during peak travel hours, particularly on major freeway commute corridors. As such, the application of microsimulation approaches for transportation demand forecasting and network modeling is very relevant in the region.

The Maricopa Association of Governments (MAG), the planning agency in the region, provided much of the data and networks necessary to conduct a TRANSIMS application. The model network consists of approximately 13,000 links and 11,000 nodes, connecting 1,995 internal zones and 11 external traffic analysis zones (TAZs). In the context of this study, the router module is applied to the entire regional network, but the microsimulator module is applied only to a subarea constituting a 5-mile buffer region around the 20-mile light rail line that commenced operations in December 2008. As the focus of an ongoing TRANSIMS research project is to apply the model system to the analysis of light rail corridors, this subarea was chosen for microsimulation. However, the results reported here are based on outputs of running simulations for the automobile mode (private vehicle) trips only. In other words, the router is run for the entire region. Once the router is stabilized, the trips that flow through or have an origin and/or destination within the light rail corridor buffer region are microsimulated within the subarea.

Traffic characteristics on two very short and very specific stretches of corridors are analyzed in the context of this study. The first is a short stretch of US Highway 60 in Mesa, Arizona, which is a heavily traveled east-west freeway corridor in the county. The second is a similar short stretch of State Route 51, which is a heavily traveled north-south corridor. Both stretches of highway are less than two miles in length. Corridors of such short length were considered deliberately to have a homogeneous stretch of highway in both cases with a limited number of links over which to compute traffic characteristics of interest.

MAG furnished all trip tables by mode and purpose for the most recent validated four-step travel demand model of which only the auto trip tables were used. Time of day distributions of travel derived from the National Household Travel Survey (NHTS) data set were used to temporally distribute the 24-hour travel demand represented in the origin-destination tables. MAG also provided detailed network files, zonal data files, and traffic count files (that could be used for model validation).

The goal of the current investigation is to determine the quantity of stochastic variation in various corridor traffic characteristics resulting from varying the input random number seed only in the router, only in the microsimulator, or in both modules of TRANSIMS. This is accomplished by first producing a base scenario of the simulation and repeating the simulation using exactly the same random number seed to ensure that nothing but the random number seed contributes to variation in outputs across model runs. Once this was established, 60 additional model runs, each with a different random number seed, were undertaken. It should be noted that each model run involves 12 iterations of the router stabilization process and 8 iterations of the microsimulator stabilization process. As such, even running 60 trials of the microsimulation is quite computationally cumbersome. In order to make the experiment more computationally manageable, only those trips planned to start between 6 AM and 7 AM on an average weekday were included. The investigation identifies the volumes, travel times, average travel speeds, and other traffic characteristics (e.g., density) resulting from the execution of travel plans for this subset of trips. As many of the trips planned to begin in this hour are likely to remain on the network during the following hour of 7 AM to 8 AM, traffic characteristics are compiled for specified corridors for this time period. It should be noted that this feature of the experiment underestimates the potential traffic on the corridors, as there may be other trips begun earlier than 6 AM or after 7 AM that should be part of the traffic during 7 AM to 8 AM. However, for the sake of computational efficiency, this simplification was adopted.

In order to isolate the impact of the random number seed being varied between the router process and the microsimulator process, three sets of 20 runs each were performed. One set of 20 runs involved varying the random number seed of the router module, but keeping the random number seed of the microsimulator module constant. A set of 20 runs involved varying the random number seed of the microsimulator module, but keeping that in the router module constant. A final set of 20 runs involved varying the random number seeds in both the router and microsimulator modules. By comparing outcomes across these three sets, it is possible to isolate the stochastic variation arising from the router module from that of the microsimulator module, as well as determining the cumulative impact of varying the random number seed in both modules. There is no specific reason for the choice of 20 simulation runs. It was felt that this is a reasonable number that practitioners can potentially adopt in balancing considerations of computational burden versus obtaining an adequate measure of stochastic variation.

The random impedance perception threshold is set to 20 percent in all 60 model runs. In other words, a traveler may perceive the travel impedance to be up to 10 percent higher or 10 percent lower than the true objective impedance value. Trips selected for rerouting must be able to reduce travel time by at least five percent and no more than 80 percent of the total trips on the network may be selected for rerouting in any single iteration. All of these user-specified parameters are kept constant to ensure that variation across model runs is purely due to random number seed changes. The random number seed associated with the traveler rerouting selection process is set to a constant value as well.

The stochastic variation in outcomes over the 20 simulation runs is measured using traditional statistical indicators such as range, standard deviation, and coefficient of variation. In order to determine whether the successive average of the outcomes reached stability after 20 simulation runs, the successive average was computed after each run and plotted to see if the degree of oscillation virtually vanishes by the time all runs are completed. The cumulative average at the end of the n^{th} trial is a simple arithmetic mean of all outcomes obtained up to and including the n^{th} trial run:

$$\bar{X}_n = \frac{(X_0 + X_1 + \cdots + X_n)}{n}$$

where $\bar{X}_n$ is the cumulative average after the n^{th} trial, and $X_0, X_1, \ldots, X_n$ are the corridor characteristics of interest at the end of each trial run. In addition, consistent with the computations in Castiglione et al (2003), a second measure of stability was calculated as the percent difference between the cumulative average up to a certain trial run and the final cumulative average obtained after all 20 trial runs were completed. Essentially, the percent difference is measured as:

$$\%Diff = \frac{\bar{X}_n - \bar{X}_{20}}{\bar{X}_{20}} \times 100$$

where $\bar{X}_{20}$ is the cumulative average at the end of the 20th trial run (i.e., the last one).

EXPERIMENTAL RESULTS

The results of the simulation experiment are summarized in Table 1 for each group of trials conducted. For a variety of traffic characteristics, the table summarizes statistical indicators so that the stochastic variability due to random number seed variations across multiple simulation runs can be assessed. The first part provides variability measures as a result of changing random number seed in the router, the second provides variability due to changing random number seed in the microsimulator, and the third part provides variability due to changing random number seed in both modules.

In general, it is seen that the stochastic variation due to changing random number seeds over multiple runs is quite small for all sets of runs. The coefficient of variation, computed as the ratio between the standard deviation and the mean of the traffic characteristic of interest, is uniformly very small across the entire table. The measures presented in the table include the number of vehicles that began their journey between 6 and 7 AM that are on the specific sections of highway during the 7 to 8 AM hour, the average travel time per vehicle to traverse the sections of highway considered in the paper, and the corresponding vehicle miles and hours of travel. The traffic volume represents the average number of vehicles per link (in the specific stretches of corridor considered) corresponding to the definition noted in this paragraph. The travel times and density values are also averages. Thus the means reported in the table for these three measures are averages (computed over 20 runs) of 10-link averages calculated within each run. The vehicle miles and hours of travel measures, on the other hand, are 20-run averages of cumulative total miles and hours of travel calculated by summing up these measures over the 10 links in each stretch of corridor considered.

An examination of the table reveals that, not only is the coefficient of variation small, but the coefficient variation is larger for the router than for the microsimulator module. It is not readily clear whether this is due to the nature of the experiment in which the router is applied region-wide while the microsimulator is applied for a small subarea of the region, thus producing potential variability in the microsimulator application portion of this study. One would naturally expect a greater level of variability across a larger network/system than when focusing on a smaller subarea. In any event, in this context, the random variability due to random seed variation is higher in the router module, and the cumulative effect of changing random number seeds in both modules, is very similar to the sole effect of the router. Thus, while it appears that stochastic variability from multiple modules does accumulate (as opposed to canceling out), the results suggest that the variability does not accumulative in a purely additive way. It is likely that the simulation variation arising from the accumulation of variability across multiple modules of a microsimulation model system is close to that of the single module that exhibits the greatest amount of variability. While it is true that the coefficient of variation is small, one should also note that the minimum and maximum values realized from the 20 simulation runs do deviate substantially (in a qualitative assessment) from the overall mean. In other words, it is not sufficient to run a simulation model just once as the stochastic realization from a single run could deviate considerably from the value that would be obtained from multiple runs. Thus, the study confirms previous literature that suggests multiple runs of microsimulation models should be undertaken in any planning context.

Four plots are shown to further illustrate the nature of the convergence process in this study for the router and microsimulator modules. Two plots depict the nature of convergence using the percent difference approach adopted by Castiglione et al (2003). Figure 2 shows the convergence pattern of volume on SR51 to the 20-run average, while Figure 3 shows the corresponding pattern for US60. In general, both graphs show rather similar patterns, but an important difference is that the pattern appears to converge to the 20-run average by about the 14^{th} or 15^{th} simulation run for SR51. In the case of US60, the pattern appears to converge to the 20-run average only at about the 19^{th} iteration, suggesting that the US60 simulation needs more simulations to reach a stable result. This may be a manifestation of the fact that US60 is a more heavily traveled freeway corridor (note the higher density on this stretch of corridor reported in Table 1), and therefore, greater stochastic variability may be present when simulating corridors of higher traffic volume. When similar plots were generated for travel time, and vehicle miles and hours of travel, similar results were found. The convergence appears to occur faster for SR51 than for US60. Another noteworthy finding corroborating what was mentioned earlier is that the graph tracking the convergence of simulation run averages from varying both module random number seeds tracks closely with that obtained from varying the random number seed of the router module only.

Figures 4 and 5 track the actual successive average over the 20 simulation runs for different experiments. Once again, results are shown for the average traffic volume, but the findings are consistent for the other traffic measures as well. The average traffic volume itself is found to converge rather well for the SR51 stretch of highway corridor by the 14^{th} or 15^{th} iteration. In the case of this corridor, it appears that a modest number of runs is sufficient to provide a stable successive average value for various traffic characteristics. On the other hand, for the US60 highway corridor, the graph depicts greater levels of variability with the successive average not appearing to converge even at the 20^{th} simulation run. In other words, it does not appear that 20 runs of the router and microsimulator modules are sufficient to provide stable, converged, and robust results for this highway corridor. These graphs also show that the variability emanating from the router module is greater than that emanating from the microsimulator module. It should be noted that a finding of greater stochastic variability (simulation error) does not necessarily mean that a certain module is inferior in any way to another module (that offers less variability), in terms of ability to replicate stochastic processes at play. Such differences may truly reflect the fact that some choice processes of behavior depict greater variability than others, and the extent of such variability can vary across individuals in a population, and across links in a network. The findings should be interpreted in the context of the number of simulation runs needed to achieve stable and robust results; naturally, modules representing behavioral phenomena of greater variability would need to be run a greater number of times to achieve such stability.

SUMMARY AND CONCLUSIONS

The ushering in of the microsimulation era in urban systems modeling has made it possible for planning agencies to model and forecast location choice decisions, activity-travel demand patterns, and network dynamics under a wide range of policy and system scenarios at the level of the individual behavioral unit. Microsimulation approaches involve running a series of probabilistic models that characterize the behavior of individual agents in the system. As each run of a microsimulation model system results in outputs that represent a possible realization of the underlying stochastic choice process, one invariably obtains slightly different results every time the model system is run with a different sequence of random number seeds. This stochastic simulation error or variability in microsimulation model outputs raises a couple of important questions that merit investigation from a planning practice standpoint. First, there is a need to understand the number of model system runs that need to be executed in order to achieve stability in model outputs for a range of output measures of interest at various levels of spatial and temporal resolution. Second, there is a need to understand and quantify the extent of stochastic variability that arises in the application of microsimulation model systems. Fulfilling these two research needs will considerably aid practitioners in their quest to move towards the deployment and application of microsimulation model systems for a range of planning studies in a variety of contexts.

In this paper, an attempt is made to add to the body of knowledge regarding the stochastic variation arising from the application of microsimulation model systems by conducting an experiment that involves repeatedly running a traffic microsimulation model using different random number seeds. The router and microsimulator modules of TRANSIMS (Transportation Analysis and Simulation System) were used to conduct the experiment. TRANSIMS is a comprehensive microsimulation model system capable of simulating activity-travel patterns of individuals in time and space over multimodal networks. In this particular study, fixed demand in the form of origin-destination trip tables from a traditional four-step travel demand model is input to the router and microsimulator modules to examine the stochastic variability associated with the application of these two modules. Three sets of 20 simulation runs each (for a total of 60 simulation runs) are conducted in the experiment; one set in which the random number seed of the router is varied, another set in which the random number seed of the microsimulator is varied, and a third set in which the random number seed of both modules is varied. The experiment is conducted for the Maricopa County (Greater Phoenix) region of Arizona in the United States. The router is applied at the regional scale, routing all trips in a morning one hour interval, while the microsimulator is applied to a subarea encompassing a five-mile buffer region around a 20-mile light rail line.

An analysis of the results shows that stochastic variability is quite small, at least in the context of this experiment in which the router and microsimulator results were examined for two short stretches of heavily traveled highway corridors during a one hour morning period. As only a limited set of trips were simulated, congested conditions did not exist on the corridors considered and it is possible that the results are sensitive to the level of congestion experienced on the roadways (one would expect greater stochastic variability in highly congested conditions than in uncongested conditions). Successive averages for different traffic characteristics (such as volume, travel time, and vehicle miles and hours of travel) were plotted for the 20 simulation runs to see if the average output reached a stable value within this number of simulations. In addition, the percent difference between the successive average at the end of each iteration was compared against the overall successive average obtained at the end of all 20 iterations. Thus, this experiment showed the extent to which 20 simulation runs are sufficient to achieve stable results, at least under the highway system conditions simulated in this study.

The stochastic variability associated with the router module was consistently higher than for the microsimulator module, but this may be because the router module is applied to the entire regional network, while the microsimulator module is applied only to the limited subarea around the light rail line. As the scope of geographical coverage decreases, one would expect the extent of stochastic variability to decrease as well. This was indeed found to be the case in this study. The coefficient of variation values for the various traffic characteristics of interest, were found to be very small suggesting that variability is quite small. Plots of successive averages, with the number of trials represented along the x-axis, showed that the successive averages of various traffic characteristics reached fairly stable and robust values within the 20 simulation runs for one corridor, but more than 20 runs were warranted in the more heavily traveled roadway corridor. This finding suggested that the number of simulation runs required to achieve stable results is not only dependent on the nature of the phenomenon being modeled, but is also location- and context-specific. There is clearly an additive impact of varying random number seeds across multiple probabilistic model components in a microsimulation model system, with the cumulative effect generally tracking closely with the model component that produces the greatest amount of variability. In the case of the experiment in this paper, the additive effect was found to be similar to the stochastic variability exhibited by the router module, consistent with the finding that the stochastic variability arising from repeated applications of the microsimulator module was considerably smaller than that arising from repeated applications of the router module.

Although it is desirable to accumulate additional evidence on stochastic variability arising from the repeated application of microsimulation model systems before definitive conclusions are drawn, the findings in this study, coupled with those reported previously in the literature, lend credence to the notion that stochastic variability in microsimulation models is rather small and should not be a deterrent to the use of such disaggregate modeling tools. There is no question, however, that multiple runs should be undertaken, as any one stochastic realization may considerably deviate from the stable successive average value that one would obtain from running the simulation model system multiple times. The range of values obtained in the experiment conducted as part of this study is evidence of this. Nevertheless, even with a reasonably modest number of runs (such as 20 as used in this study), it appears that analysts can be confident that they have (at least nearly) converged on a stable and robust set of values for traffic characteristics of most interest. Future research efforts should further attempt to analyze the extent of stochastic variability associated with simulating congested (unstable) network conditions, and how such variability differs across alternative microsimulation model systems that are composed of different probabilistic choice components.

REFERENCES

Axhausen, K. and T. Garling. Activity-based Approaches to Travel Analysis: Conceptual Frameworks, Models, and Research Problems. *Transport Reviews*, Vol. 12, No. 4, 1992, pp. 323-341.

Bhat, C.R., J. Guo, S. Srinivasan, and A. Sivakumar. Comprehensive Econometric Microsimulator for Daily Activity-Travel Patterns. *Transportation Research Record, Journal of the Transportation Research Board*, No. 1894, Transportation Research Board of the National Academies, Washington, D.C., 2004, pp. 57-66.

Balmer, M., K. Nagel, and B. Raney. Agent-based Demand Modeling Framework for Large Scale Micro-simulations. *Transportation Research Record, Journal of the Transportation Research Board*, No. 1985, Transportation Research Board of the National Academies, Washington, D.C., 2006, pp. 125–134.

Castiglione, J., J. Freedman, and M. Bradley. Systematic Investigation of Variability Due to Random Simulation Error in an Activity-Based Microsimulation Forecasting Model. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 1831, Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 76-88.

Castiglione, J., R. Hiatt, T. Chang, and B. Charlton. Application of Travel Demand Microsimulation Model for Equity Analysis. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 1977, Transportation Research Board of the National Academies, Washington, D.C., 2006, pp. 35-42.

Gibb, J. and J. L. Bowman. Convergence of an Activity-based Travel Model System to Equilibrium: Experimental Designs and Findings. Proceedings of the 11th National Transportation Planning and Applications Conference of the Transportation Research Board, Daytona Beach, Florida, 2007.

Goulias, K.G. and R. Kitamura. Travel Demand Forecasting with Dynamic Microsimulation. In *Transportation Research Record* No. 1357, Transportation Research Board of the National Academies, Washington, D.C., 1992, pp. 8-17.

Lawe, S., J. Lobb, A.W. Sadek, S. Huang, and C. Xie. TRANSIMS Implementation in Chittenden County, Vermont: Development, Calibration, and Preliminary Sensitivity Analysis. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 2132, Transportation Research Board of the National Academies, Washington, D.C., 2009, pp. 113-121.

Mahmassani, H.S. (2001) Dynamic Network Traffic Assignment and Simulation Methodology for Advanced System Management Applications. *Networks and Spatial Economics*, Vol. 1. No. 3-4, pp. 267-292.

Salvini, P.A. and E.J. Miller. ILUTE: An Operational Prototype of a Comprehensive Microsimulation Model of Urban Systems. *Networks and Spatial Economics*, Vol. 5, No. 2, 2005, pp. 217-234.

Veldhuisen, J., H. Timmermans, and L. Kapoen. Microsimulation Model of Activity-Travel Patterns and Traffic Flows: Specification, Validation Tests, and Monte Carlo Error. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 1706, Transportation Research Board of the National Academies, Washington, D.C., 2000, pp. 126-135.

Vovsha, P., E. Petersen, and R. Donnelly. Microsimulation in Travel Demand Modeling: Lessons Learned from the New York Best Practice Model. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 1805, Transportation Research Board of the National Academies, Washington, D.C., 2002, pp. 68-77.

Vovsha, P., R. Donnelly, and S. Gupta. Network Equilibrium with Activity-Based Microsimulation Models: The New York Experience. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 2054, Transportation Research Board of the National Academies, Washington, D.C., 2008, pp. 102-109.

Waddell, P.A., A. Borning, M. Noth, N. Freier, M. Becke, and G. Ulfarsson. Microsimulation of Urban Development and Location Choices: Design and Implementation of UrbanSim. *Networks and Spatial Economics*, Vol. 3, No. 1, pp. 43-67.

Walker, J. L. Making Households Microsimulation of Travel and Activities Accessible to Planners. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 1931, Transportation Research Board of the National Academies, Washington, D.C., 2005, pp. 38-48.

TABLE 1 Summaries of Traffic Characteristics on Simulated Corridors

Measure	Mean		Minimum		Maximum		Median		Std. Deviation		Range		CV	
	US60	SR51	US60	SR51	US60	SR51	US60	SR51	US60	SR51	US60	SR51	US60	SR51
Results of Changing RNS in Router Module Only														
Volume (veh)	175.88	82.80	170.20	80.20	181.70	85.60	176.20	82.80	2.55	1.29	11.50	5.40	0.014	0.016
Travel Time (min)	2.06	1.69	2.05	1.55	2.06	1.72	2.06	1.72	0.00	0.05	0.02	0.17	0.00	0.03
Density (veh/lane-m)	0.51	0.32	0.46	0.30	0.65	0.36	0.50	0.31	0.04	0.02	0.19	0.06	0.077	0.050
VMT (mi)	538.65	213.94	520.81	207.58	557.60	221.11	539.36	213.79	8.03	3.42	36.79	13.53	0.015	0.016
VHT (hr)	5.78	2.29	5.59	2.22	5.98	2.37	5.79	2.29	0.09	0.04	0.39	0.14	0.015	0.016
Results of Changing RNS in Microsimulator Module Only														
Volume (veh)	176.12	82.67	170.20	81.60	176.90	83.00	176.70	82.70	1.47	0.25	6.70	1.40	0.008	0.003
Travel Time (min)	2.06	1.71	2.05	1.66	2.06	1.72	2.06	1.72	0.00	0.01	0.01	0.06	0.00	0.01
Density (veh/lane-m)	0.51	0.31	0.46	0.31	0.53	0.33	0.51	0.31	0.02	0.00	0.07	0.02	0.031	0.013
VMT (mi)	539.39	213.66	520.81	210.23	541.51	214.60	541.39	213.79	4.64	0.81	20.70	4.37	0.009	0.004
VHT (hr)	5.78	2.29	5.59	2.25	5.81	2.30	5.80	2.29	0.05	0.01	0.22	0.05	0.008	0.004
Results of Changing RNS in Both Router and Microsimulator Modules														
Volume (veh)	175.29	82.91	168.20	80.20	181.70	85.60	175.60	82.90	3.01	1.38	13.50	5.40	0.017	0.017
Travel Time (min)	2.06	1.69	2.05	1.55	2.06	1.72	2.06	1.72	0.00	0.05	0.01	0.17	0.00	0.03
Density (veh/lane-m)	0.51	0.33	0.45	0.30	0.65	0.36	0.50	0.33	0.04	0.02	0.20	0.06	0.086	0.049
VMT (mi)	537.17	214.22	516.03	207.58	557.60	221.11	537.39	213.81	9.25	3.65	41.57	13.53	0.017	0.017
VHT (hr)	5.76	2.30	5.53	2.22	5.98	2.37	5.76	2.29	0.10	0.04	0.45	0.14	0.017	0.017

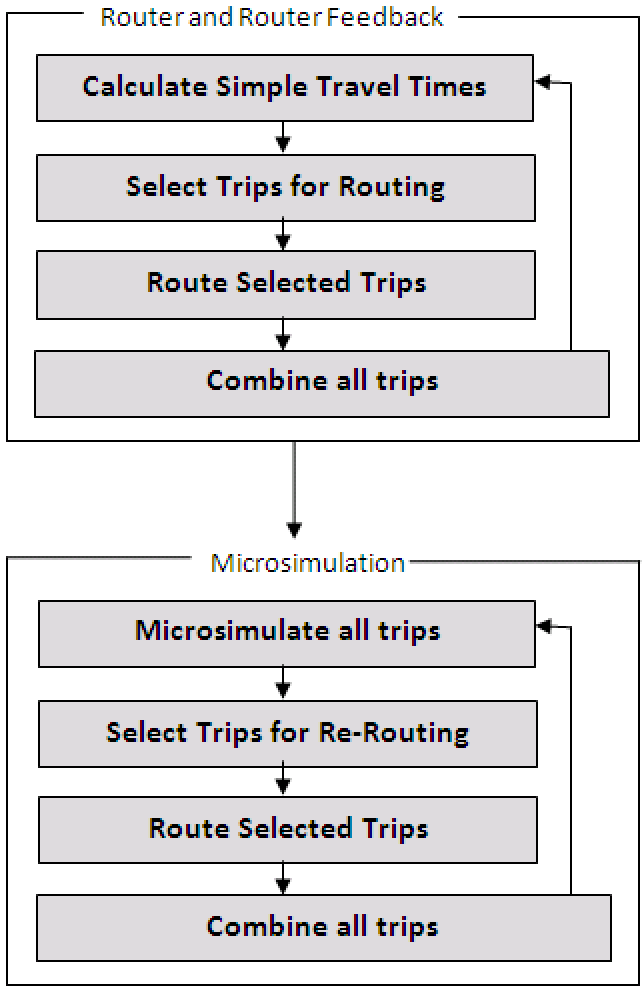

FIGURE 1 Router and microsimulator stabilization processes.

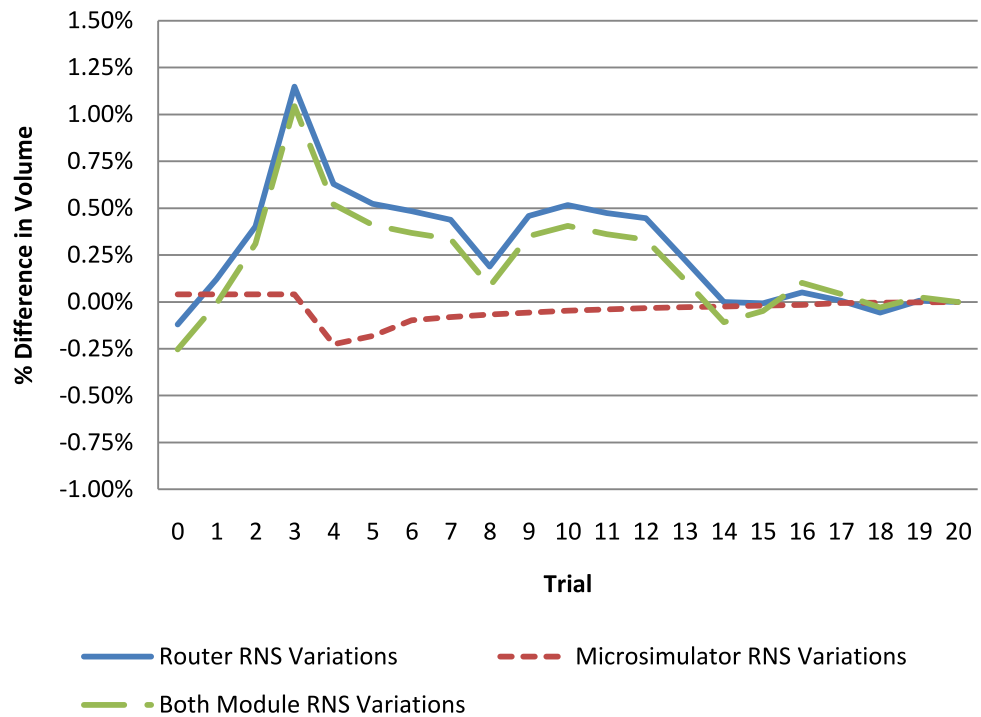

FIGURE 2 Percent difference from final average volume on SR 51.

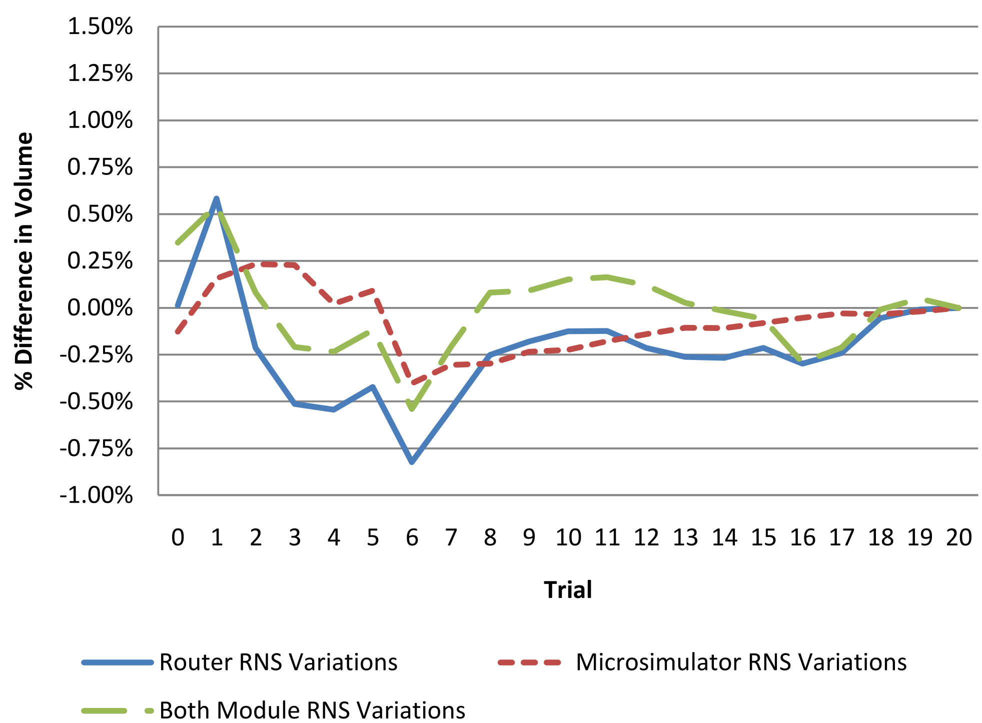

FIGURE 3 Percent difference from final average volume on US60.

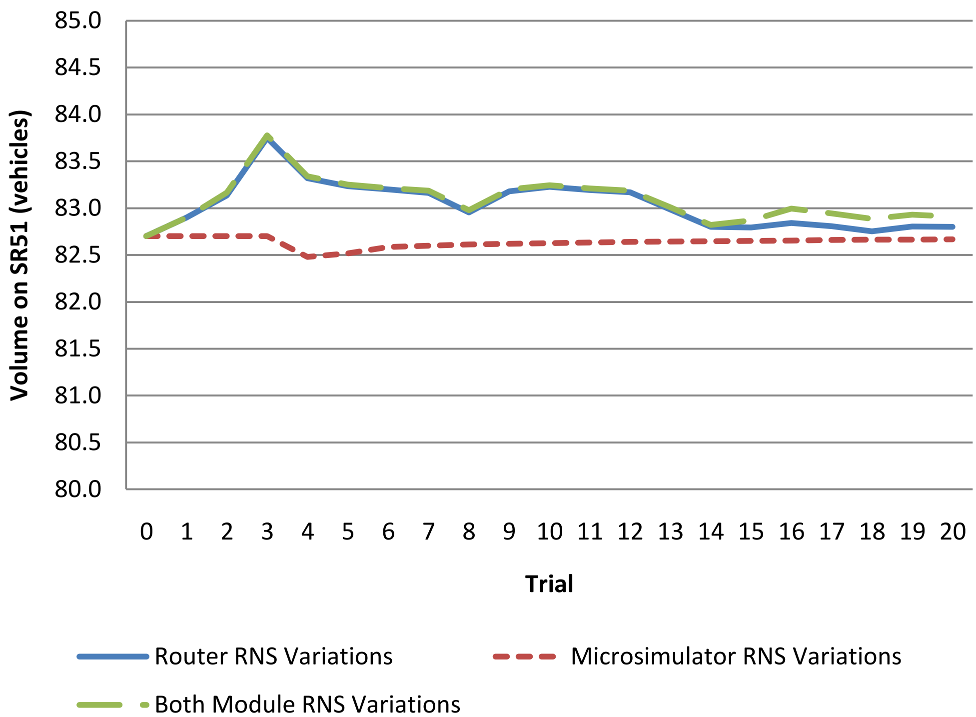

FIGURE 4 Successive Average Traffic Volume on SR51.

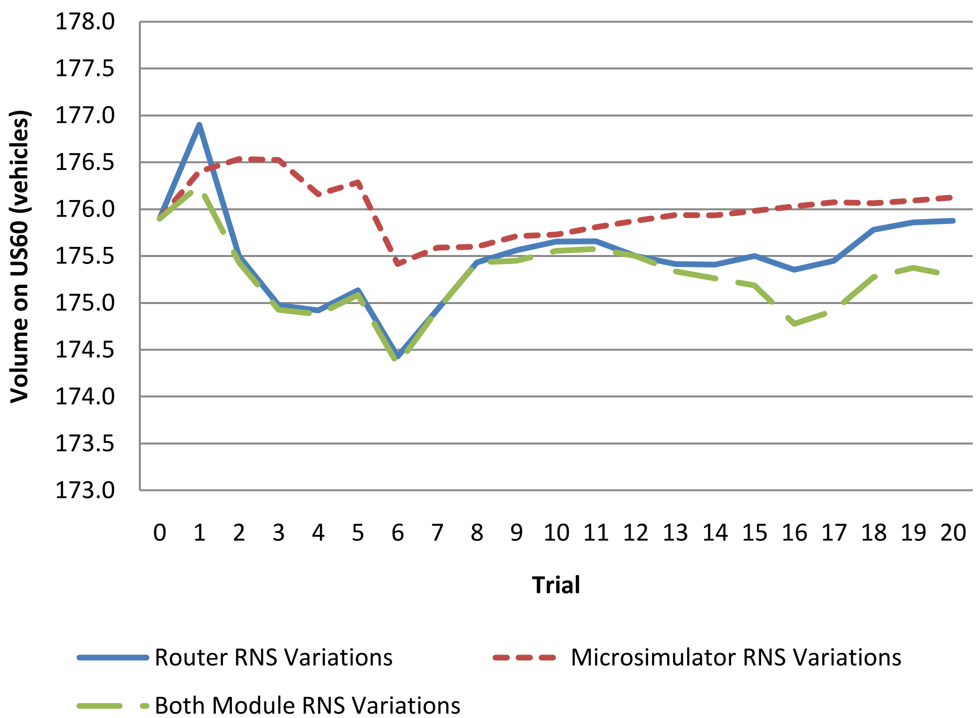

FIGURE 5 Successive Average Traffic Volume on US60.