Alternative Methods To Iterate a Regional Travel Simulation Model: Computational Practicality and Accuracy

W. THOMAS WALKER AND HAIOU PENG

The results of full-scale testing of the available methods for implementing an iterative travel simulation process are presented. These methods include simple iteration of the simulation model chain, weighting iterative model outputs by the method of successive averages, and the Evans equilibrium algorithm. Simulated travel demands for each version of the model are compared with regional highway performance monitoring system data, 1990 highway traffic counts summarized by screenline, and public transit ridership data. These accuracy checks, in concert with estimates of the computational effort needed to execute each model variation, provide a useful insight into the costs and benefits associated with implementing an iterative travel simulation model. These comparisons also give guidance with respect to the relative efficacy of each iterative approach.

The Clean Air Act Amendments of 1990 and Intermodal Surface Transportation Efficiency Act of 1991 have significantly expanded the role of travel simulation models in evaluating the efficacy of proposed transportation improvements and in projecting the impact of these improvements on progress toward achieving mandated air quality standards. To adequately fulfill this expanded role in determining conformity, the federal legislation also requires that existing travel simulation models be validated with ground counts and upgraded to reflect an acceptable level of modeling practice. Perhaps the most significant of these requirements involves starting the simulation process with observed free-flow speeds and then iterating the entire simulation until a "reasonable agreement" is achieved between the travel speeds assumed for trip distribution and modal split and the resulting restrained speeds output by the highway assignment model.

The purpose of this paper is to conduct a full-scale exploration of the available methods for implementing this required iterative simulation process using the existing regional travel simulation model for the Delaware Valley Region as a test system. Socioeconomic data based on the 1990 Census, together with highway and public transit networks that reflect the facilities open to traffic in 1990, are used as inputs to Delaware Valley Regional Planning Commission's (DVRPC) existing simulation and selected iterative reconfigurations of this model. Simulated travel demands are then compared with regional highway performance monitoring system data, 1990 highway traffic counts summarized by screenline, and public transit ridership data. These accuracy checks, in concert with the relative computational effort needed to execute the model, provide a useful insight into the costs and benefits associated with adopting an iterative travel simulation model. These accuracy comparisons also provide guidance with respect to the efficacy of alterative iterative approaches. This analysis is limited to accuracy comparisons that are based on the existing noniterative calibrations of the travel demand models—the starting point for most state and regional attempts to implement an iterative simulation procedure. The results for all iterative approaches analyzed might be improved by recalibrating the models to better replicate actual travel data within an iterative formulation.

Delaware Valley Regional Planning Commission, Bourse Building, 111 South Independence Mall East, Philadelphia, Pa. 19106-2515.

METHODS FOR ITERATING TRAVEL SIMULATION MODEL

The methods that have been proposed to iterate the travel simulation model to some degree parallel the evolution of the highway assignment from simple iteration to weighted average to equilibrium, although discussions of iterative methods to date have focused on model convergence properties rather than model accuracy in a calibration sense. Levinson and Kumar (*1*) opened the current round of discussion by proposing that the modeling chain be simply iterated, starting with free-flow speeds, until a reasonable degree of convergence is obtained between the times used as input to the gravity and modal split models, and the congested times resulting from the subsequent highway assignment. The assignment results from the last iteration of this process form the basis for plan evaluation, conformity determination, and so on. Failure to iterate was found to overestimate congestion levels resulting from long-range socioeconomic and land use forecasts.

Although the convergence of travel times is monitored in the simple iteration method, it is not clear whether the simulated link volumes converge to a stable solution. For this reason Putman (*2, 3*), in work done in association with the Southern California Council of Governments, weighted together the highway link volumes from each simple iteration of the model chain using the method of successive averages (MSA). This successive weighting technique, proposed by Sheffi (*4*), uses a fixed weighting sequence, where the weight given to the link volume difference between the current iteration (n) and the weighted average resulting from the previous iterations is $1/(n + 1)$. The link volumes resulting from this method are easily shown to converge for any pattern of highway assignments. As the overwhelming proportion of the overall weight is given to the first few iterations, this method is often started with congested rather than free-flow speeds. This algorithm is usually run for a fixed number of iterations because the degree of convergence is directly determined by the progression of the weighting sequence.

The so-called Evans algorithm also uses a successive averaging technique to weight together the results of subsequent iterations of the modeling chain. However, instead of using $1/(n + 1)$ as the weighting sequence, weights are calculated by a gradient search within a Frank-Wolf decomposition. This weighting takes the following form:

$$P'_{ijm} = (1-\lambda)P_{ijm} + \lambda Q_{ijm}$$

and

$$v'_a = (1-\lambda)v_a + \lambda w_a$$

where P_{ijm} and v_a represent the results of successive weighted averages over the previous iteration estimates of trips between zones i and j by mode m and travel over link a, respectively. Q_{ijm} and w_a represent corresponding results for the current iteration run of the simulation models.

This method is related to the well-known equilibrium assignment method; however, the Evans algorithm incorporates the results of the entire simulation process from trip distribution to transit and highway assignment within the gradient search. The method is known to converge rapidly. It requires that only one iteration of highway assignment be conducted between successive runs of the simulation model chain. These two factors result in large potential savings in computation vis-à-vis simple iteration of the traditional modeling chain. This algorithm is based on work done by Evans as part of her Ph.D. dissertation in the early 1970s (*5*). Recently, Boyce et al. suggested using this method to satisfy the federal iterative modeling requirements (*6,7*).

Convergence criteria can be rigorously defined for this algorithm using the difference at a given iteration between the numerical value associated with the primal and dual of the underlying nonlinear impedance minimization problem (primal and dual are equal at convergence). Because these criteria are difficult to calculate, a convergence criterion similar to the one applied in most implementations of the equilibrium assignment is used. After weighting the current iteration trip interchanges and link volumes together with the composite results from previous iterations, the new capacity restrained link times together with transit travel times, fares, and parking charges are used to project system total impedance (S1). This value is then compared with the total impedance resulting from the next iteration of the simulation models (S2). The difference (error) between these two estimates expressed as a fraction of current impedance (S2) is taken as a measure of convergence. This assumes that the impact of reiterating the travel simulation becomes progressively smaller as convergence nears. This definition of convergence has proved to be adequate in practice.

ITERATIVE FORMULATION

The existing DVRPC travel simulation model is a classic implementation of the four-step process. All aspects of the model produce estimates of daily travel. Trip generation is based on constant trip rates imbedded in a cross-classification structure. The trip distribution, modal split, and highway assignment models are based on average daily highway travel speeds. Bus speeds are taken from the existing a.m. peak transit operating schedule and held constant throughout the simulation. Trip distribution uses a doubly constrained gravity model, stratified into three-person (home-based work, home-based nonwork, and non-home–based) and four-vehicle trip purposes. The person-trip gravity models utilize a combined highway/transit network interzonal impedance measure based on a relative highway/transit service level bias adjustment (*8*). Modal split utilizes a binary probit-like formulation stratified by trip purpose, transit submode, and automobile ownership. The highway assignment is based on the equilibrium method using minimum travel time paths. Initial highway speeds are input through a table lookup stratified by functional class and density of development (area type). The transit assignment is unrestrained. It uses minimum paths that are based on the modal split model definition of impedance.

The DVRPC model is among the largest in existence, covering a densely developed area of about 10 400 km^2 (4,000 mi^2) that is subdivided into 1,449 traffic zones. The highway network contains about 34,000 one-way links, and the transit network contains about 360 routes, including commuter rail, rapid transit, light rail, and bus facilities. Overall, the model has been stable over time, achieving validation with counts for 1960, 1970, 1980, and 1990 with minimal structural or parameter changes. A more detailed description of this model is given in Walker (*9*). The DVRPC model was originally developed on a mainframe using the PLANPAC/UTPS packages but recently has been converted to a microcomputer environment using TRANPLAN. All sensitivity tests reported in this paper were done using the TRANPLAN microcomputer system.

Incorporating Actual Highway Speeds into Simulation Model

The DVRPC model has a fundamental problem that prevents it from being used directly in an iterative framework. Input highway speeds are unrealistically low, particularly on freeways. Furthermore, the output speeds from the assignment (via the BPR restraining curve) are even more unrealistic, perhaps half the true average daily highway operating speeds. This is common in simulation models developed during the 1970s. Although these speeds cannot be used for emissions calculations, they generally improve the accuracy of the highway assignment, which responds favorably to a bias against freeways and severe capacity restraint. A postprocessor methodology is used to reestimate highway operating speeds on the basis of assigned volumes before it estimates emissions.

The most straightforward way to correct this problem is to insert "actual" congested speeds into the highway network through a revised speed lookup table. However, this substitution increased the simulation error to an unacceptable level. Clearly, a more sophisticated method is needed to incorporate actual operating speeds into the travel simulation model. It was always obvious that some of the values in the original highway speed lookup table were not real speeds but rather a crude form of impedance. The phenomenon being addressed was that drivers consider distance (or operating cost) as well as travel time when choosing routes. Freeways move faster than arterials, but there is a limit to the route circuity that drivers will accept to achieve a savings in travel time.

The modal split model already had a highway impedance measure that considered both highway time and operating cost. A theoretically appealing way to incorporate actual congested speeds is to extend this impedance measure to the gravity model and highway assignment as well. The entire simulation model would then be based on a uniform definition of impedance. This impedance definition is similar to the one found in most disutility-based modal split models:

$$Z = k_1\,ET + k_2\,RT + k_3\,C$$

where

Z = impedance for given travel mode;
ET = excess or out-of-vehicle time (i.e., terminal time for highway, sum of walk and wait time for transit; transit impedance also includes a supplemental transfer penalty);
RT = running or in-vehicle time;
C = monetary cost (i.e., fare for transit; out-of-pocket operating cost plus tolls and parking for highway); and
k_1, k_2, k_3 = calibration constants.

To test this approach, highway trees were built using the modal split impedance definition with actual congested times in the lookup table. The resulting impedance skims were found to be perfectly collinear with the minimum time skims from the original speed lookup table. Only a simple-scale factor was required to make these impedance skims usable with the original gravity model friction factor curves and terminal and intrazonal times, and so on. Highway assignment path building also was based on this impedance definition. However, the capacity restraint calculation was limited to the travel time portion of the impedance. To improve the highway travel speeds produced by the model, the exponent of the Bureau of Public Roads (BPR) restraint curve was reduced from 4.0 to 3.0.

The accuracy of the resulting assignment was checked on the basis of 1990 traffic counts summarized through a series of 14 screenlines. These screenlines form the basis for FHWA model validation within the DVRPC region. Included are circumferential central business district and intermediate suburban cordon lines, all crossings of the Schuylkill and Delaware rivers, and a series of radial cutlines.

The use of a highway impedance model increased the total volume error for all screenline counts from 2.2 percent for the original model to 3.8 percent. The number of screenlines with volume errors greater than 10 percent increased from two, with the worst screenline having 12 percent error, to four, with the worst being 13 percent. The R^2 between predicted and actual traffic volumes for all screenline crossings was reduced from 0.89 to 0.85 by the highway impedance model. The simulated highway speeds produced by the impedance-based model, although almost 10 percent low on average, were judged to be sufficiently accurate to test iterative simulation methods.

Implementing Simple Iteration and MSA Approaches

The simple and MSA methods of iterating the travel simulation model are straightforward to incorporate into an existing travel simulation model. The simple iterative method requires only that a feedback loop be inserted into the model that inputs the highway link speeds output from highway assignment of the current iteration into the network before rebuilding and reskimming the minimum impedance paths, so that trip distribution and modal split model runs of the next iteration step will be based on the current iteration's congested link travel times. Most, if not all, travel simulation model software packages incorporate link travel times. Most, if not all, travel simulation model software packages incorporate provisions for this feedback loop. The estimates of link volumes produced by the simple iterative approach are taken directly from the highway and transit assignments of the last iteration that is executed.

The MSA approach builds on these simple simulation model iterations by combining the link volumes of each simple iteration into a composite volume, using the weighting scheme outlined earlier. This composite link volume is calculated by a postprocessor computer program that processes the output of each successive model iteration. The MSA software used in this analysis was prepared by the Urban Analysis Group as part of the TRANPLAN package.

Implementing Evans Algorithm

The Evans algorithm is not difficult to implement in a four-step travel simulation model that includes a highway assignment model based on the equilibrium method, although some extension of the modal split and highway assignment models, as well as the associated computer code is required (Figure 1). Evans reexecutes the gravity and modal split models after each iteration of highway assignment. Therefore, a restart procedure must be available in the highway assignment program to access the weighted average highway link volumes from the previous iteration, load the network for the current iteration, calculate the weight for the current iteration (λ), and prepare a convex combination of the link volumes for the current iteration and previous weighted average. This is not a fundamental departure from the way things are normally done in the equilibrium assignment, and the restart option already exists in TRANPLAN and perhaps other packages.

The second required extension is to include the impedance implications of the highway and transit trip tables into the gradient calculation that is used to determine λ. This requires an estimate of transit impedance and off-network highway impedance (terminal

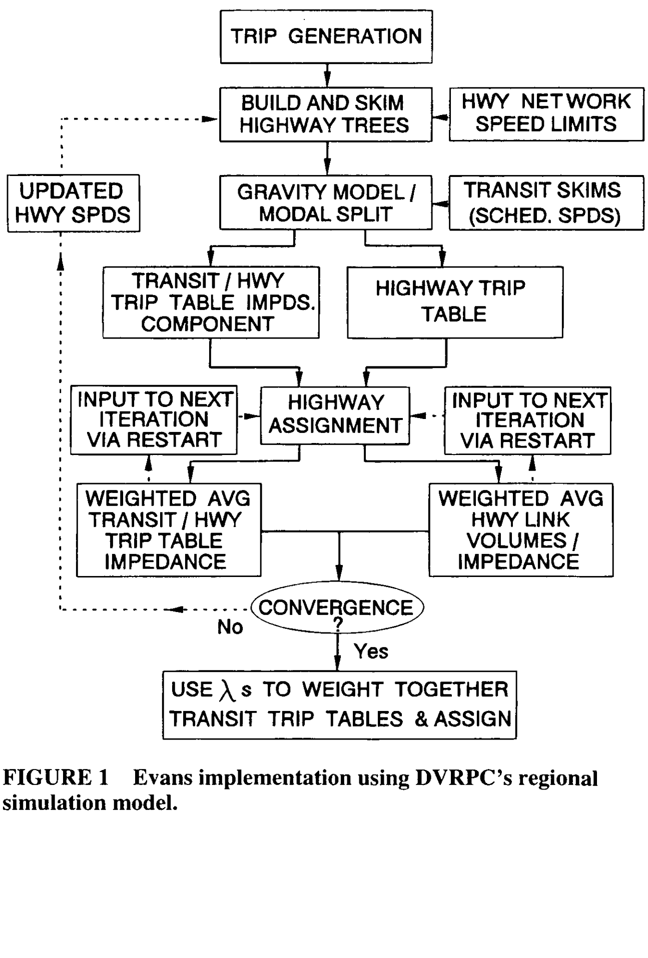

FIGURE 1 Evans implementation using DVRPC's regional simulation model.

times and parking charges) for the trip tables of the current iteration and the weighted average of the previous iterations. Transit impedance is assumed to be independent of the highway link restraining process and is calculated as the sum of the products of interzonal transit impedances and transit volumes. It may be theoretically desirable to also include the effect of highway congestion on bus and trolley travel times. However, this enhancement requires massive changes to the highway assignment computer program and is beyond the scope of this study. In any case, only about 4 percent of the region's total travel is made by transit.

In this implementation, it is assumed that weighted average totals of transit and off-network highway impedance are linear in λ and can be calculated directly from the system totals for the current and weighted average of the previous iterations. The alternative would be to calculate a new λ-weighted trip table and multiply this new table by the interzonal impedance matrix. This simplification has little effect on the accuracy of the calculation. It greatly reduces the computational effort in the search routine that is used to determine λ and the complexity of the required program code changes. For the current iteration, the system total for both the off-network highway and transit impedance are calculated in the modal split model and passed in a scratch file to the highway assignment for inclusion in the gradient calculation. Similarly, the weighted transit and off-network highway impedance calculated in the highway assignment is passed from iteration to iteration in a scratch file.

In the Evans algorithm, trip tables are weighted together from iteration to iteration using λ-based successive averages in exactly the same way as highway link volumes. Thus, the transit trip table must be calculated with this method before assignment to the transit network.

COMPARISON OF RESULTS

This section compares the results of the impedance version of the DVRPC simulation model under three alternative methods for iterating the model: simple iteration, MSA, and the Evans algorithm. All iterative simulation model runs were started with highway speed limits, which are assumed to represent the "free-flow speeds" recommended in the federal guidance. Congested speeds are unacceptable as a starting point in iterative processes using the BPR restraining curve shown later because the restrained link times can never be lower than the input times T_0.

$$T = T_0\left[1.0 + 0.15\left(\frac{V}{C}\right)^f\right]$$

where

T = adjusted link time,
T_0 = initial input link time,
V/C = ratio of volume to capacity in current assignment, and
f = exponent on V/C; 3.0 in these runs, 4.0 default.

The fact that T_0 is not increased leads to errors in mobile source emissions because speed increases above current congested speeds, perhaps resulting from highway capacity improvements or land use changes, cannot be modeled. To standardize the comparisons, all three methods were iterated five times after an initial iteration (0) execution of the travel simulation models.

Convergence and Computational Efficiency

Table 1 compares the systemwide convergence criteria for the simple and Evans method. Both the simple and Evans algorithms converged to the neighborhood of 0.01 error after five iterations. The error statistic in this table refers to highway link impedance only in the simple model but also reflects the trip table impedance components in the Evans results. For this reason the Evans model estimates of S1 and S2 are somewhat larger. The difference between S2 in the simple and Evans cases gives an indication of the relative impact of the trip table impedance component within the Evans gradient calculation. Overall, the highway links provide about 90 percent of the influence in the determination of λ. The effect of the trip table gradient component is usually to reduce the weight given to the first two or three iterations.

Although not quite reaching the 0.01 criteria, the Evans convergence rate was particularly impressive because it is based on only six executions of the highway assignment. The DVRPC network is slow to converge in equilibrium assignment, requiring 12 to 15 iterations to reach this level of error. For this reason, the simple iteration results required a total of 90 executions of the highway assign-

TABLE 1 Convergence Statistics After Five Iterations from Speed Limits

ITERATIVE METHOD	PROJECTED TOTAL IMPEDANCE $\times 10^4$ (S1)	ACTUAL TOTAL IMPEDANCE $\times 10^4$ (S2)	ERROR (S1-S2)/S1	HIGHWAY ASSIGNMENT ITERATIONS	APPROXIMATE COMPUTATION TIME [a]
SIMPLE ITERATION	99,508	98,498	0.010	90	78 HRS.
MSA	NA	NA	NA	90	79 HRS.
EVANS ALGORITHM	111,120	109,636	0.014	6	15 HRS.
EVANS ALGORITHM FULL RESTRAINT ITERATION 0	109,486	109,484	0.000	20	26 HRS.

[a] 66 MHZ 486 UNDER OS/2.

ment. The total computation time for the simple method is about 78 hr on a 66-MHZ 486 microcomputer under OS/2. This is an impractical running time for most planning applications. For example, current federal guidance for the DVRPC region requires a total of seven simulations (1990 base year plus build and no-build alternatives for 1996, 2005, and 2015) to demonstrate conformity. Including the time needed for program setups and output checking, the simple method would require somewhere between 25 and 35 days to complete the computation. Five iterations of the Evans approach will run overnight (15 hr per alternative), or about 7 to 10 days to complete the required conformity simulations. Despite the ongoing advances in microprocessing speed, this is an overwhelming computational advantage.

The MSA method weights together the results of the five simple method iterations with a special postprocessor program. This MSA weighting operation requires something less than 1 additional hr to complete (79 hr total for five iterations). The MSA approach does not lend itself to the calculation of S1 and S2 parameters. Furthermore, these parameters reflect only the system total of impedance and do not directly measure the variation in link volumes from iteration to iteration. To directly measure link level convergence, the percent root mean square (RMS) difference in link volumes, from iteration to iteration, was also calculated for each of the three iterative methods. The results of the calculation are shown graphically in Figure 2. As one might expect, the MSA approach had the fastest rate of link volume convergence. It significantly improved the convergence rate of the simple method, which tended to level out at about 5 percent RMS difference per iteration. With the MSA method it would seem that it is possible to terminate computation after Iteration 2, a savings of 50 percent (40 hr of computation per simulation). The Evans approach demonstrated a high rate of convergence, closing all the way from 78 percent RMS difference between Iterations 0 and 1 to 18 percent between Iterations 4 and 5. However, it is clear from Table 1 and Figure 2 that additional iterations of the simulation model are required for the Evans algorithm to reach the level of convergence of the simple approach. This lack of link-level convergence is also reflected in the error statistics.

To achieve complete convergence, the Evans algorithm was restarted and run for five additional iterations. Convergence to the 0.01 criteria was achieved on Iteration 7 (20 hr of computation time), although link-level convergence to 5 percent RMS difference was not achieved. On Iteration 7, this measure dropped to just below 10 percent, and stayed around 10 percent throughout Iterations 8 to 10. It seems that the 5 percent RMS difference level of link convergence requires multiple iterations of the highway assignment within each Evans iteration to smooth out the highway assignment though traditional capacity restraint. DVRPC's highway network is dense in terms of link topology, and the Evans results might also be improved by the creation of additional traffic zones.

When starting from speed limits, a particularly critical point in terms of the smoothness of the traffic assignment was Iteration 0. For this reason, it seemed probable that the convergence properties of the Evans algorithm could be improved by executing a full traditional capacity restraint (15 iterations) in Evans Iteration 0, thence continuing with the standard single iteration of restraint within each Evans iteration. The results of the test are also reported in Table 1. In terms of total impedance, this variation on the Evans model significantly improved the rate of convergence. After five iterations of Evans, the error term was reduced to less than 0.001, although the link-level convergence did not go below 10 percent RMS difference. As the 0.01 level of convergence was achieved in Iteration 3 the last two iterations could be eliminated saving about 5 hr of computation time over the 26-hr required. An alternative to a full restraint may be to use congested speeds as the BPR curve T_0 value in Iteration 0 and then switch to speed limits in subsequent Evans iterations.

Accuracy and Usability for Emissions Calculations

The effect of iterating the travel simulation models on assignment accuracy is indicated in Table 2. Although the total of predicted and counted volumes for all screenline links remains well below a 5 percent difference, individual screenline accuracy is degraded versus the noniterated travel simulation under all three iterative approaches. The R^2 between predicted and counted volumes for all screenline links is also significantly reduced by the iterative simulations. The biggest factor in this error increase is the use of speed limits rather than congested speeds as the starting point for the assignment. The

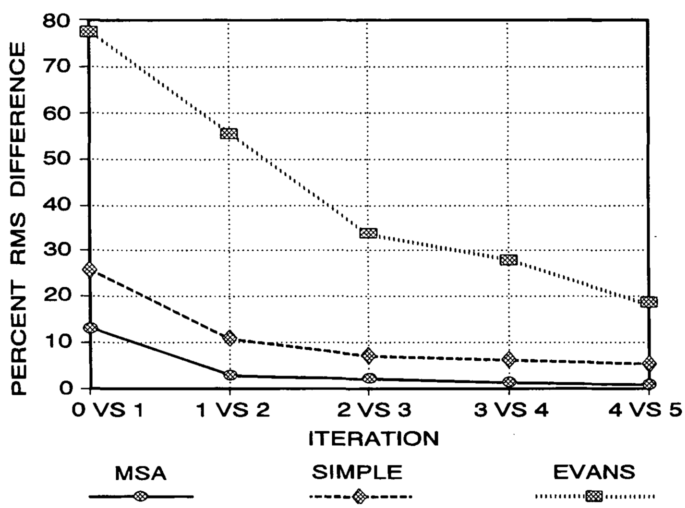

FIGURE 2 Comparative rate of convergence of highway link volumes.

TABLE 2 Highway Screenline Error Statistics for Simple Iteration, MSA, and Evans Algorithm from Speed Limits

ITERATIVE METHOD	OVERALL ERROR	# OF SCREEN LINES > 10% ERROR (WORST)	AVG. ABS. SCREEN LINE ERROR	R^2 ALL LINKS
SIMPLE ITERATION	2.2%	4 (19%)	7.17%	0.75
MSA	3.6%	5 (18%)	7.54%	0.75
EVANS ALGORITHM FIVE ITERATIONS	-3.5%	6 (23%)	8.57%	0.66
EVANS ALGORITHM COMPLETE CONVERGENCE	-3.6%	4 (24%)	7.61%	0.67
EVANS ALGORITHM FULL RESTRAINT ITERATION 0	-2.9%	3 (22%)	7.38%	0.67

results for the MSA approach are slightly worse than for the simple iteration method but comparable overall.

The Evans method showed a somewhat larger reduction in accuracy. In part, this resulted from executing the highway assignment only six times. Restarting Evans for two additional iterations resulted in some improvement in accuracy of the screenline volumes, but the full-capacity restraint in Evans Iteration 0 almost achieved screenline validation in terms of volume totals. Only 1 of the 14 screenlines and cutlines checked had a total traffic volume error greater than 11 percent, with the worst (22 percent) being a small suburban circumferential cutline. However, this variation of the Evans algorithm continued to have a significantly smaller link-level R^2 than either simple iteration or MSA. The trip table and restrained link volumes rapidly converge to a hand-in-glove fit in the Evans approach. This tends to magnify the effect of network topological and model calibration/specification deficiencies. All three modeling approaches will require some degree of simulation model enhancement and recalibration to achieve screenline validation. This recalibration is beyond the scope of this investigation.

As indicated in Table 3, all iterative approaches produced acceptable estimates of regional highway vehicle kilometers of travel (VKMT) and transit ridership; however, all significantly overestimated highway operating speed (by 12.4 to 17.6 percent). None of these methods can be used to estimate mobile source emissions without first reestimating congested speeds with a postprocessor.

Alternative Capacity Restraining Functions

In an attempt to improve the accuracy of the speed estimates produced by the iterative simulations, four variations of the capacity restraining function were tried. Because the computation associated

TABLE 3 Selected Regional Travel Statistics for Simple Iteration, MSA, and Evans Algorithm from Speed Limits

ITERATIVE METHOD	HWY. VKMT [a] $\times 10^6$ (% DIFF. FROM HPMS)	HIGHWAY AVG. SPEED KM/H (% ERROR)	TRANSIT BOARDING $\times 10^6$ (% ERROR)
SIMPLE ITERATION	143.2 (-2.5%)	50.5 (17.6%)	1.26 (7.7%)
MSA	145.1 (-1.2%)	48.3 (12.4%)	1.26 (7.7%)
EVANS ALGORITHM FIVE ITERATIONS	142.2 (-3.2%)	48.8 (13.5%)	1.24 (5.9%)
EVANS ALGORITHM COMPLETE CONVERGENCE	141.4 (-3.7%)	48.9(13.9%)	1.24 (5.9%)
EVANS ALGORITHM FULL RESTRAINT ITERATION 0	140.8 (-4.2%)	49.1 (14.2%)	1.26 (7.7%)

[a] VKMT = VEHICLE KILOMETERS OF TRAVEL;
VEHICLE MILES OF TRAVEL = VKMT ÷ 1.6093

with the simple iteration method is excessive, these tests were limited to the Evans algorithm. All three iterative methods produce similar estimates of regional VKMT and operating speed in the earlier comparisons.

The first and second alternative restraining function involved resetting the exponent of the BPR curve *V/C* ratio to 4.0 and 7.0, respectively, and then running the iterative simulation from speed limits. The standard value of the V/C exponent is 4.0, but recent research has suggested that larger values, perhaps 7.0, may produce better results. The third and fourth variations involved direct use of the speed curves from DVRPC's emissions postprocessor methodology as the restraining function. These speed curves are much more complex than the BPR function, being related to the methods contained in the Highway Capacity Manual. The exact formulation of these curves may be found in Walker (*9*). Because the times output by these curves are not limited by the input T_0 values, the postprocessor speed curves were used in two ways: one using speed limits as the starting point of the simulation process and the other using congested speeds.

The results produced by these tests are presented in Tables 4 and 5. Resetting the BPR exponent to 4.0 significantly improved the screenline accuracy of the Evans algorithm, although the results were still not as good as the simple or MSA results shown earlier. The exponent value of 7.0 improved the screenline results even further, being comparable with those of simple iteration and MSA shown earlier. The regional VKMT and transit ridership estimates for both exponent values were comparable with those produced by the 3.0 case, but average speed estimates produced by the 7.0 exponent value had less than 1 percent error, raising the possibility of eliminating the speed estimation postprocessor. This version of the Evans model seems to be able to produce reasonably accurate estimates of both VKMT and speed. However, the 7.0 BPR exponent slowed down the rate of algorithm convergence. Ten Evans iterations were required to achieve 0.01 convergence.

Use of the postprocessor speed curves generally degraded the accuracy of travel volumes produced by the Evans algorithm. This occurred in part because the modal split model went out of calibration, leading to severe overestimation of center-city transit ridership and corresponding underestimation of some highway screenline totals and of regional VKMT. This restraining function did produce significantly more accurate estimates of simulated highway speeds, however (about 3.3 percent overestimated).

The postprocessor curves produced about the same error statistics, whether the simulation was iterated from speed limits or congested speeds. The highway link volumes produced by these alternative starting points had about a 13.5 percent RMS difference after five iterations. This version of the Evans algorithm seemed to produce relatively unique results at both the regional and link level, regardless of the initial speeds, although as one might expect, convergence was significantly faster when the algorithm was started from congested speeds.

CONCLUSIONS

It is clear from the results presented in this paper that converting the DVRPC travel simulation model to an iterative formulation on the basis of initial free-flow speeds is not a trivial undertaking. Simple iteration of the modeling chain requires days of computation to complete the simulation for a single alternative. The draft federal guidance also requires disaggregating the simulation process into separate peak and off-peak models. Implementing this requirement would effectively double all computing times reported in this paper. Furthermore, the off-peak time period is far from homogeneous in terms of congestion. Midday congestion resembles the peak period in many suburban areas, whereas evening travel in these areas in virtually free flow. Three or four time periods may be required. For this reason the computational efficiencies resulting from the MSA and Evans algorithms are essential to the continued computational practicality of the travel simulation process.

The Evans algorithm required the least amount of computer time to achieve convergence in terms of systemwide total impedance, reducing the time required by 80 percent versus simple iteration. This time savings is dependent on the number of iterations of restraint that are required for the highway assignment in the simple method. DVRPC's network requires 15 iterations in a normal assignment. Other regions whose network converges faster may receive a smaller time savings from the Evans algorithm.

The MSA procedure allows the number of iterations (and associated computation) required to achieve link-level convergence to be

TABLE 4 Highway Screenline Error Statistics for Alternative Restraining Functions

ITERATIVE METHOD	OVERALL ERROR	# OF SCREEN LINES > 10% ERROR (WORST)	AVG. ABS. SCREEN LINE ERROR	R^2 ALL LINKS
BPR EXP. 4.0 FROM SPEED LIMITS	-3.4%	3 (21%)	8.16%	0.70
BPR EXP. 7.0 FROM SPEED LIMITS	-6.6%	4 (19%)	7.20%	0.74
POST-PROCESSOR CURVES FROM SPEED LIMITS	-7.1%	4 (26%)	10.53%	0.74
POST-PROCESSOR CURVES FROM CONGESTED SPEEDS	-7.3%	4 (21%)	10.89%	0.73

TABLE 5 Selected Regional Travel Statistics for Alternative Restraining Functions

ITERATIVE METHOD	HWY. VKMT [a] × 10^6 (% DIFF. FROM HPMS)	HIGHWAY AVG. SPEED KM/H (% ERROR)	TRANSIT BOARDING × 10^6 (% ERROR)
BPR EXP. 4.0 FROM SPEED LIMITS	141.4 (-3.7%)	47.1 (9.7%)	1.25 (6.8%)
BPR EXP. 7.0 FROM SPEED LIMITS	140.8 (-4.2%)	42.6 (-0.7%)	1.24 (6.0%)
POST-PROCESSOR CURVES FROM SPEED LIMITS	134.4 (-8.5%)	44.4 (3.3%)	1.34 (14.5%)
POST-PROCESSOR CURVES FROM CONGESTED SPEEDS	134.2 (-8.7%)	44.4 (3.3%)	1.34 (14.5%)

[a] VKMT = VEHICLE KILOMETERS OF TRAVEL; VEHICLE MILES OF TRAVEL = VKMT ÷ 1.6093

reduced by one-half. Although converging very rapidly, the Evans algorithm did not achieve the degree of link-level convergence of the simple iteration or MSA approach in the test applications. Running Evans for two additional iterations improved the link and system level convergence (and accuracy) but reduced the computer time savings versus MSA somewhat. However, the Evans algorithm has considerable theoretical appeal, in that the weights on successive simulation model iterations are based on a Frank-Wolf decomposition rather than the arbitrary sequence used by MSA.

All three iterative approaches significantly degraded the accuracy of the travel simulation model, making validation of screenline volumes and congested speed much more difficult to achieve. The use of speed limits rather than congested speeds as a starting point for the iterative process was a major factor in this accuracy loss. The Evans approach was somewhat less accurate in part because of the drastic reduction in the number of iterations of the highway assignment required for five iterations. However, the rapid convergence between trip table and congested link volumes in this approach may also magnify the effect of certain deficiencies in the travel simulation model. Simulation model enhancement or recalibration may be necessary to optimize the accuracy of the results from any of the three iterative approaches.

Almost all iterative formulations tested tended to significantly overestimate congested highway link speeds and will require post-processor-based reestimation of speeds before mobile source emissions calculation. Only the Evans algorithm with a BPR restraint curve exponent of 7.0 seems to produce estimates of both highway VKMT and congested operating speed when starting the iterative process from highway speed limits.

The motivation for implementing an iterative simulation is to be able to accurately assess the impact of future land use patterns and proposed transportation facilities. It is interesting to note that the highway travel speed lookup table and other model parameters in the existing DVRPC model have remained almost unchanged for the last 30 years, despite repeated intervening forecasts of increased highway congestion. Furthermore, budget constraints and intense citizen opposition have limited the region's ability to build new freeways and improve existing roadways. Potential excessive congestion resulting from population and employment growth and increasing dependance on automobiles has been controlled by highway peak spreading and decentralization of urban activity into suburban and rural areas of the region. From this perspective, it would seem more likely that a significant projected imbalance between input and output speeds in the simulation model would be caused by an underestimate of decentralization and peak spreading than a failure to iterate. Iterative travel simulation models should include a feedback loop that incorporates the impact of localized projected congestion levels on the underlying land use and socioeconomic forecast. This feedback could utilize formal land use models, if sufficiently sensitive to localized congestion conditions, or might be accomplished through ad hoc adjustments.

REFERENCES

1. Levinson, D., and A. Kumar. Integrating Feedback into the Transportation Planning Model: Structure and Application. In *Transportation Research Record 1413,* TRB, National Research Council, Washington, D.C., 1993, pp. 70–77.
2. Putman, S. H. *Integrated Urban Models 2: New Research and Applications of Optimization and Dynamics.* Pion Limited, London, 1991.
3. *MSACOM Computer Program Documentation.* Southern California Association of Governments, 1994.
4. Sheffi, Y. *Urban Transportation Networks.* Prentice-Hall, Englewood Cliffs, N.J., 1985.
5. Evans, S. P. Derivation and Analysis of Some Models for Combining Trip Distribution and Assignment. *Transportation Research,* Vol. 10, 1976, pp. 37–57.
6. Boyce, D. E., M. Tatineni, and Y. Zhang. *Scenario Analysis for the Chicago Region with a Sketch Planning Model of Origin-Destination Model and Route Choice.* Urban Transportation Center, University of Illinois at Chicago, 1992.
7. Boyce, D. E., M. R. Lupa, and Y. Zhang. Introducing Feedback into the Four-Step Travel Forecasting Procedure versus the Equilibrium Solution of a Combined Model. In *Transportation Research Record 1443,* TRB, National Research Council, Washington, D.C., 1994.
8. Walker, W. T. Testing and Adjusting Regional Travel Simulation Models with 1980 Census Data. *Transportation Quarterly,* Jan. 1988, pp. 63–88.
9. Walker, W. T. The Impact of Preaggregation of Highway Network Travel Data on the Accuracy of Mobile4-Based Emissions. In *Transportation Research Record 1366,* TRB, National Research Council, Washington, D.C., 1992, pp. 51–59.

Publication of this paper sponsored by Committee on Passenger Travel Demand Forecasting.