Convergence Properties of Some Iterative Traffic Assignment Algorithms

ALAN J. HOROWITZ

This paper examines the convergence properties of four popular traffic assignment algorithms: Frank-Wolfe decomposition for fixed-demand equilibrium assignment, an ad hoc variation of the Evans algorithm for elastic-demand equilibrium assignment, fixed-demand incremental assignment, and elastic-demand incremental assignment. The algorithms were evaluated according to errors associated with insufficient iterations, arbitrary selection of starting point, inexact theory, and small variations in data. Each of the four algorithms reached its intended solution, but did so very slowly. Elastic-demand incremental assignment emerged as the preferred technique, principally because of its more accurate response to small variations in data and its adaptability to various models of travel demand.

The most popular traffic assignment algorithms may be thought of as logical extensions to traditional iterative capacity restraint. That is, the algorithms consist of a series of all-or-nothing assignments interspersed with computations to improve estimates of link impedances and, perhaps, link volume. Some of these algorithms, such as the Frank-Wolfe decomposition method for fixed-demand assignment (*1*) or Evans's method for elastic-demand assignment (*2*), have a strong theoretical basis. Other algorithms are ad hoc. In spite of the large body of theoretical work on traffic assignment, transportation planners have had little guidance about the algorithm that yields the best performance within the usual limits on resources. In addition, there is little accurate information on how to employ an algorithm most effectively once a choice has been made. Many common rules-of-thumb are seriously misleading.

REVIEW OF THE ALGORITHMS

The purpose of this paper is to reevaluate a few existing algorithms rather than to break new theoretical ground. The following are brief descriptions of the algorithms considered:

- Iterative capacity restraint: Iterative capacity restraint is still popular, despite its terrible convergence characteristics. This algorithm is included in this comparison because it has aptly served as a "straw man" in studies by other researchers.
- Equilibrium: This fixed-demand, equilibrium assignment technique, available in most major planning packages, is an implementation of Frank-Wolfe decomposition.
- Modified Evans: Modified Evans is an ad hoc variation of the Frank-Wolfe decomposition algorithm that recalculates demand at each iteration. It resembles the Evans algorithm in both purpose and performance.
- Fixed-demand incremental: Incremental traffic assignment loads a fraction of the trip table at each iteration using all-or-nothing assignment. This technique can be implemented as a slight variation of equilibrium assignment.
- Elastic-demand incremental: At each iteration, the trip table is recalculated and a portion of it is loaded to the network. This algorithm can be implemented as a slight variation of the modified Evans algorithm.

Occasional reference will be made to Evans's precise algorithm for elastic-demand equilibrium assignment. Although the Evans algorithm is not explicitly evaluated, it is possible to determine the extent to which the other algorithms differ from the results of a true elastic-demand equilibrium assignment—the intended product of the Evans algorithm. The Evans algorithm was dropped from consideration because of its comparatively large computational requirements on multipurpose networks.

The algorithms were tested on two networks. The first was the five-zone UTOWN network, developed for testing the equilibrium assignment in the Urban Transportation Planning System (UTPS). The second was the a.m. peak-hour network for East Brunswick, N.J. The East Brunswick network contained 129 zones, and it gives a good indication of how algorithms would perform in actual practice. Both networks had five trip purposes. The tests were performed with an experimental version of QRS II running on a Zenith Z-248 (IBM PC-AT compatible).

Convergence Error

Generally, assignment error is the difference between assigned and actual volumes. Unfortunately, we can never measure actual volume with sufficient accuracy to use it as a criterion in evaluating the differences between assignment algorithms because the results are far too similar.

Total error is quite large. Various studies have shown (*3, 4*) that root mean square (RMS) errors can regularly exceed 50 percent. Established guidelines for error (*5*) take into consideration the better performance on high-volume links, but a 20 percent error is still considered acceptable.

Convergence error, a component of total error, can be measured by comparing the results of two assignments, assuming that one of the assignments is essentially perfect. For

Center for Urban Transportation Studies, University of Wisconsin—Milwaukee, P.O. Box 784, Milwaukee, Wis. 53201.

example, a network can be run through a huge number of iterations of equilibrium assignment to obtain a nearly perfect solution to the fixed-demand problem. This solution becomes a standard for comparison. Because the primary purpose of an assignment algorithm is to forecast volumes on links, it makes sense to measure convergence error as the RMS difference in volume between the test algorithm and the standard algorithm. The RMS difference is analogous to standard error and is in units of vehicles, so it is easily interpreted.

Other researchers have attempted to measure convergence error by monitoring the objective function of the equilibrium assignment algorithm:

$$U = \sum_{\substack{\text{all} \\ \text{links} \\ i}} \int_0^{V_i} t_i(v)\, dv \tag{1}$$

where $t_i(v)$ is the functional relationship between travel time and volume on link i, and V_i is the assigned volume. Since the equilibrium solution is achieved when U is minimized, an experienced individual can roughly judge the progress of an algorithm by comparing U at successive iterations. However, this objective function is deceptive. Surprisingly large changes in volume can be associated with very small changes in U. It is known (at least for the fixed-demand problem) that smaller values of U are better, but it is difficult to determine how much better or how fast the solution is improving.

A related criticism applies to monitoring the RMS change in volume between successive iterations [see paper by Sheffi and Powell (*6*) for an example]. The algorithms, as a group, converge slowly. It is not possible to determine the ultimate amount of change in volume by the change from a single iteration.

A given level of convergence error can be either important or unimportant, depending on the purpose of the forecast. To understand the role of convergence error in forecast validity, it is first necessary to list various forms it can take.

1. Insufficient iterations: Solutions generally improve at each new iteration. There can be significant convergence error associated with terminating an algorithm prematurely.
2. Resolution: An algorithm should be able to reach the same solution to a given problem each time that it is run. Since an algorithm is trying to replicate real-world processes, we would also expect it to produce similar solutions to similar problems. If it cannot do this, the algorithm is flawed.
3. Starting point: An algorithm should arrive at the same solution regardless of how it is started. Practically speaking, the solutions produced by all algorithms are affected by the choice of starting point. Insensitivity of an algorithm to its starting point is an important characteristic.
4. Ad hoc algorithm: An ad hoc algorithm could fail to converge or it could converge to a solution that is inconsistent with assignment theory. The justifications for choosing an ad hoc algorithm are potentially less error due to insufficient iterations and potentially better resolution.

It is important to keep these errors in perspective. Assignment algorithms are highly imperfect models of travel behavior. Much more significant errors stem from our poor understanding of route choice behavior, limited knowledge of impedance functions, problems in collecting demographic and network data, and our inability to show the network as it actually exists. Imperfections in theory and data are much more serious than imperfections in algorithms to implement the theory.

Test Conditions

To the best of their ability, tests were representative of planning practice. Neither the UTOWN nor East Brunswick network was modified in any way. With the exceptions of the assignment algorithm and the number of iterations of the trip distribution model, all parameters were set to the defaults for QRS II.

A doubly constrained entropy-maximizing model was used for trip distribution. For the UTOWN network, the attraction-end constraints were satisfied by 10 iterations of the trip distribution model. Trip distribution on the East Brunswick network was iterated only three times.

A Fibonacci search was used to find the averaging weights for the equilibrium and modified Evans algorithms, which minimize U. The Fibonacci search was permitted to run for 21 iterations, assuring four significant digits in the weights.

Only links that would normally carry traffic were compared for error. Centroid connectors and other artificial network elements were ignored. Also ignored were links that received no volume in any of the assignments.

Relationship Between Equilibrium and Incremental Assignments

Each iteration of equilibrium assignment consists of (a) an all-or-nothing assignment, (b) an averaging of volumes, and (c) a recalculation of link travel times given the averaged volumes. The averaging step consists of finding a weighted average between the all-or-nothing assignment and the results of the previous iteration such that U is minimized. Each iteration has a different weight, and it is impossible to know ahead of time what those weights will be.

It is easy to give the algorithm a predetermined series of weights. Although it will not necessarily converge to the equilibrium solution, the algorithm runs faster, behaves more predictably, and is easier to explain to those outside the field. One particular sequence of weights yields an incremental assignment:

$$W = 1/(i + 1) \tag{2}$$

where W is the weight given to the all-or-nothing assignment that is calculated at iteration i. Regardless of the number of iterations, each all-or-nothing assignment (including the one from the 0th iteration that starts the algorithm) is weighted equally in the final average. Running the equilibrium algorithm with this particular fixed series of weights is a form of incremental assignment, in which the link travel times for the next increment are calculated from extrapolations of the partial volumes that have already been assigned (*7*).

Incremental assignment, as described in this paper, is a case of the method of successive averages (MSA) (*6, 8*). As a group, algorithms based on MSA are not as precise as purer optimization methods but have a greater range of applicability.

The close relationship between equilibrium and incremental assignment suggests that their solutions would be similar. It

TABLE 1 COMPARISON OF ITERATIVE CAPACITY RESTRAINT WITH MODIFIED EVANS ALGORITHM (UTOWN NETWORK)

Iteration	% RMS Difference in Link Volumes	% of Optimal Objective Function
0	84	413
1	156	407
2	134	559
3	104	1110
4	133	540
5	119	631
10	135	561
20	144	777

Modified Evans' algorithm was run for 200 iterations.

is expected that equilibrium assignment would converge faster when measured by iterations, but equilibrium assignment might well be slower when measured by total computer time.

Straw Man: Iterative Capacity Restraint

It is important to understand that an ad hoc algorithm can be either good or bad, depending on its design. A popular ad hoc algorithm is iterative capacity restraint. As implemented in QRS II, each iteration consists of (a) calculation of a trip table with travel times from the previous iteration, (b) an all-or-nothing assignment, and (c) a recalculation of link travel times. Thus, the algorithm can be considered "elastic demand"; it attempts to find a trip table that is consistent with link loads. Travel times are recalculated with the Bureau of Public Roads (BPR) speed-volume function. To provide some stability to the algorithm, link travel times were damped. That is, link travel times were taken as a weighted average of the results of the BPR function (25 percent) and the link travel times from the previous iteration (75 percent).

The UTOWN network was run for the 7:00 to 8:00 a.m. peak hour through 20 iterations of iterative capacity restraint. These results were compared with those of the modified Evans algorithm. The modified Evans algorithm is also ad hoc, but (as will be seen later) converges nicely. The comparison volumes were taken from the 200th iteration.

As expected, Table 1 shows that iterative capacity restraint performs poorly. Link volumes oscillate wildly. RMS error never becomes better than 84 percent; the value of the Frank-Wolfe objective function, U, never falls below its starting value.

The weaknesses of iterative capacity restraint are well documented, so these results are not totally unexpected. The especially poor performance seen in Table 1 illustrates that the UTOWN network can be hostile to ad hoc algorithms.

Ad Hoc Error of the Modified Evans Algorithm

Evans's algorithm correctly solves an elastic-demand assignment problem. It produces a solution consisting of (a) link volumes that are consistent with both link travel times and the trip table and (b) a trip table that is consistent with path travel times. In practice, the Evans algorithm looks like a variation of equilibrium assignment. Each iteration consists of computation of a trip table, an all-or-nothing assignment, an averaging step, and a recalculation of link travel time from the averaged volumes. The major obstacle to implementation of Evans's algorithm is the objective function of its averaging step. It requires far more computation and memory than regular equilibrium assignment, especially on large, multipurpose networks.

The elastic-demand equilibrium algorithm in QRS II replaces Evans's objective function with Equation 1. Consequently, QRS II is ensured of converging to a slightly wrong solution. It is possible to estimate the size of the error by the following procedure.

1. Run the modified Evans algorithm through enough iterations that link volumes are no longer changing. The selected number of iterations for the UTOWN network was 1,000. The assignment for the East Brunswick network was terminated at 100 iterations.
2. Save the trip table at the final iteration.
3. Run a fixed-demand equilibrium assignment for the same large number of iterations on this same network using the saved trip table.
4. Compare the volumes from the two assignments.

To control computation errors in the trip table, the trip distribution model was iterated 20 times (for each assignment iteration) on the UTOWN network and 10 times (for each assignment iteration) on the East Brunswick network.

The comparison is not a tautology. The modified Evans algorithm does not converge to the exact solution because the averaging weights disregard information about trip distribution. As the algorithm progresses, an inconsistency develops between the averaged volumes and the trip table, which is recomputed at each iteration. If this inconsistency is small, then final path travel times and, thus, the final iteration trip table are at the equilibrium solution. However, the final assigned volumes partially come from trip tables that were not at the equilibrium solution. The inconsistency can be measured by locking the trip table at its known equilibrium solution and running an exact, fixed-demand equilibrium assignment.

With UTOWN the link volumes differed (RMS) by 1.1 percent. With East Brunswick, the link volumes differed by

TABLE 2 PERCENT RMS ERROR FROM INSUFFICIENT ITERATIONS (UTOWN NETWORK)

Iteration	Modified Evans'	Equilibrium	Fixed-Demand Incremental	Elastic-Demand Incremental
1	64.4	46.8	70.4	84.9
2	55.3	35.6	48.0	53.6
3	34.2	32.5	39.7	41.3
4	28.1	24.9	29.5	34.3
5	25.1	21.8	25.2	27.6
10	15.9	14.8	15.6	18.7
20	9.1	9.3	10.7	11.3
50	3.7	4.0	4.5	4.3
100	1.4	1.3	2.0	1.5

TABLE 3 EQUILIBRIUM OBJECTIVE FUNCTION BY ITERATION (UTOWN NETWORK)

Iteration	Modified Evans'	Equilibrium	Fixed-Demand Incremental	Elastic-Demand Incremental
1	21.533	14.466	178.018	208.939
2	13.398	11.291	32.121	36.710
3	11.188	10.553	15.356	16.801
4	10.679	10.147	11.732	12.947
5	10.484	9.985	10.652	11.226
10	9.940	9.745	9.817	9.965
20	9.708	9.616	9.632	9.710
50	9.578	9.525	9.505	9.563
100	9.546	9.485	9.466	9.535
200	9.532	9.464	9.447	9.525

Units are 100,000 vehicle-minutes. Fixed-demand trip tables were taken from the 20th iteration of modified Evans'.

0.4 percent. Some of this error may be due to rounding. The small differences in assigned volumes indicate that the ad hoc error of the modified Evans algorithm, when used with a doubly constrained trip distribution model, is unimportant.

These comparisons were repeated using elastic-demand incremental assignment. For the UTOWN network the RMS difference in link volumes was 0.8 percent. When the East Brunswick network was subjected to the same comparison, the RMS difference in link volumes was 0.7 percent. The ad hoc error of elastic-demand incremental assignment is similar to that of the modified Evans algorithm.

Convergence Rates of Iterative Algorithms

An important attribute of an algorithm is its speed of convergence—often measured as the number of iterations necessary to reach a convergence criterion. Convergence speed was tested on four algorithms: equilibrium, modified Evans, fixed-demand incremental, and elastic-demand incremental. The first tests concerned performance on the UTOWN network. The volumes from various iterations of each algorithm were compared with volumes from 200 iterations of the same algorithm. The RMS differences are summarized in Table 2.

The convergence rates of all the algorithms were remarkably slow. Regardless of the algorithm, it took approximately 20 iterations before the convergence error fell below 10 percent. A convergence error of less than 5 percent required nearly 50 iterations. Interestingly, the variable-weight algorithms (equilibrium or modified Evans) did not significantly outperform either incremental assignment algorithm. Some of the slow convergence can be attributed to the hostility of the UTOWN network.

Table 3 gives the values of the equilibrium objective function, U, at each iteration. Note that fixed-demand and elastic-demand assignments approach slightly different values of the objective function, as expected. Table 3 illustrates the deceptive nature of the objective function. By the fifth iteration, U is changing only by about 1 percent per iteration, but the link volumes are nowhere near their equilibrium values.

The incremental algorithms did surprisingly well; after 20 iterations their objective functions were lower than their variable-weight counterparts (equilibrium and modified Evans). Furthermore, the incremental assignments required considerably less time to reach the same number of iterations. For example, 10 iterations of the modified Evans algorithm took 406 sec of elapsed time; 10 iterations of elastic-demand incremental assignment took just 225 sec.

Similar tests were performed on the East Brunswick network. The comparison assignments were obtained from the 50th iteration of each algorithm. These results are shown in Table 4. Convergence rates, as measured by percent RMS difference in link volumes, were twice as fast as with the UTOWN network. Nonetheless, it took approximately 10 iterations to achieve a 10 percent error. Usually a 10 percent computational error is considered unacceptable.

TABLE 4 PERCENT RMS ERROR FROM INSUFFICIENT ITERATIONS (EAST BRUNSWICK NETWORK)

Iteration	Modified Evans'	Equilibrium	Fixed-Demand Incremental	Elastic-Demand Incremental
1	34.8	36.3	34.8	33.6
2	22.9	23.2	29.3	32.4
3	16.6	17.4	25.4	22.4
4	15.7	15.4	17.5	16.2
5	13.1	11.9	13.7	13.1
10	8.7	7.6	7.1	6.4

TABLE 5 EQUILIBRIUM OBJECTIVE FUNCTION BY ITERATION FOR FIXED-DEMAND ASSIGNMENTS (EAST BRUNSWICK NETWORK)

Iteration	Equilibrium	Fixed-Demand Incremental
1	2.373	3.013
2	2.448	2.496
3	2.359	2.465
4	2.286	2.325
5	2.247	2.282
10	2.161	2.178
50	2.057	2.078

Units are 100,000 vehicle-minutes.

This research did not evaluate methods of accelerating equilibrium assignment (*9*), so these tests may somewhat understate its potential. Similar acceleration techniques would also apply to the original Evans algorithm; however, one would guess that the amount of acceleration is insufficient to overcome the algorithm's large computational requirements on meaningfully complex networks.

Ad Hoc Error of Incremental Assignment

The previous results show that the two incremental algorithms run at about the same rate (measured by iterations) as equilibrium assignment. As a further comparison, the East Brunswick network was run with fixed-demand incremental assignment for a total of 50 iterations. These results were compared with 50 iterations of equilibrium. The RMS difference in link volumes was 1.7 percent. Table 5 shows that the values of *U* for the two assignments were also close after the third iteration.

A similar comparison has already been seen in Table 3. The last line shows that at 200 iterations on the UTOWN network, incremental assignment actually outperformed equilibrium assignment. Incremental assignment was slightly closer to the equilibrium solution. The RMS difference in link volumes was 1.0 percent. The superior performance of incremental assignment on this network should be considered unusual.

Resolution Error

In many planning situations, a serious concern is the ability of an algorithm to produce similar results from similar networks. For example, a small change in a single zone's trip production should have just a small effect on volume. Table 6 shows the behavior of the several assignment algorithms when 1,000 dwelling units are added to a single zone of the UTOWN network. Each line in the table compares the volumes obtained from the base network with the volumes from the modified network when run through the same number of iterations of the same algorithm.

The first line in Table 6 should be considered the correct answer. It compares the two networks after 200 iterations of the modified Evans algorithm. It is seen that the addition of 1,000 dwelling units causes a 4.2 percent RMS change in assigned volumes.

The other algorithms, if they are working properly, should always show a smaller RMS change than all-or-nothing assignment. The other algorithms are inherently multipath, so the additional trips are split among a greater number of links. As expected, the comparison using all-or-nothing assignment (line 4) is larger than that obtained with 200 iterations of the modified Evans algorithm.

The remaining lines in Table 6 show that the other algorithms are not working properly. They all overestimate the amount of change. The most accurate was the modified Evans algorithm at 20 iterations (overestimating the change by 1.3 percent of average link volume); the least accurate was elastic-demand incremental at 10 iterations (overestimating the change by 6.9 percent of average link volume).

Given these disturbing results, a more elaborate series of tests was run on the East Brunswick network; the results are summarized in Table 7. As with the tests of the UTOWN network, each cell in the table represents a comparison of two slightly different networks, which were run on exactly the same algorithm. Each pair of networks differed by the addition of 84 dwelling units to a single zone of one network. Five separate zones were arbitrarily chosen for investigation. The iterative assignment algorithms were run for just 10 iterations.

The RMS difference using all-or-nothing assignment gives a slight overestimate of the expected change. At most, the addition of 84 dwelling units to Zone C resulted in an (RMS) impact of 2.6 percent. Three of the five zones had impacts of less than 1 percent.

All of the iterative assignment techniques estimated the impact badly. For example, we know from the all-or-nothing assignments that the correct impact for Zone A is less than 0.5 percent. However, the iterative assignment algorithms yielded impacts between 1.6 and 7.2 percent. The elastic-demand incremental algorithm behaved best for every zone.

It appears that resolution error is largely a consequence of error due to insufficient iterations. This convergence error has both random and systematic components. The systematic

TABLE 6 PERCENT RMS DIFFERENCE IN VOLUME AFTER 1,000-DWELLING UNIT INCREASE IN ONE ZONE (UTOWN NETWORK)

Algorithm	Iterations	Percent RMS Difference
A. Modified Evans'	200	4.2
B. Modified Evans'	10	7.9
C. Modified Evans'	20	5.5
D. All-or-Nothing	0	6.1
E. Equilibrium	10	6.5
F. Equilibrium	20	7.5
G. Elastic-Demand Incremental	10	11.1
H. Elastic-Demand Incremental	20	7.6

TABLE 7 PERCENT RMS DIFFERENCE IN VOLUME AFTER 84-DWELLING UNIT INCREASE IN SINGLE ZONE (EAST BRUNSWICK NETWORK)

Zone	All-or-Nothing	Modified Evans'	Equilibrium	Elastic-Demand Incremental
A	0.5	7.2	2.0	1.6
B	0.6	7.0	2.6	2.4
C	2.6	7.1	3.3	2.1
D	2.5	7.9	4.3	3.0
E	0.7	9.4	2.2	1.6

TABLE 8 PERCENT RMS DIFFERENCE IN VOLUME FROM VARIOUS STARTING POINTS (UTOWN NETWORK)

Total Iterations	Modified Evans'	Equilibrium	Elastic-Demand Incremental
10	10.6	15.7	10.8
20	6.3	12.1	6.2

component vanishes in the comparison; the random component does not. As seen here, large amounts of random error can mask the actual impact. Comparing the errors in Table 4 with those in Table 7 shows that the convergence error in the modified Evans algorithm is almost entirely random, whereas the convergence error in elastic-demand incremental assignment has a large systematic component.

The distinction between random convergence error and systematic convergence error is critical to the selection of an assignment algorithm. The nature of transportation planning is to compare alternatives. During such comparisons the only important errors are random. Random convergence error can be attenuated only by running additional iterations.

Starting Point Error

All iterative assignment algorithms require an initial estimate of link travel times. In practice, the results of assignment algorithms depend on this estimate.

Table 8 shows the effect of the starting point on the UTOWN network. Each cell in Table 8 compares two assignments for the identical network on an identical algorithm. The two assignments differ only by the method of estimating the initial link travel times. One assignment uses free travel time; the other assignment uses travel times estimated from volumes resulting from an all-or-nothing assignment.

Starting point errors are almost as large as errors due to insufficient iterations. Interestingly, the two ad hoc algorithms (modified Evans and elastic-demand incremental) were shown to be far less sensitive to the starting point than equilibrium assignment.

There exists a rule of thumb that a good initial estimate of link travel times will produce a better assignment than an inaccurate initial guess. Although partially correct, this rule of thumb is not very helpful. Table 9 shows the effect of an optimal set of initial travel times on the objective function (U) of the modified Evans algorithm. The optimal link travel times were taken from the 200th iteration of the same algorithm. A comparison of Table 9 with the first column of Table 3 shows that optimal link travel times were essentially useless. Any early advantage was erased by the 20th iteration. Similar results were obtained with the other algorithms.

CONCLUSIONS

All algorithms tested, with the exception of iterative capacity restraint, are derived from Frank-Wolfe decomposition. For practical purposes, they all converge to their intended solu-

TABLE 9 EQUILIBRIUM OBJECTIVE FUNCTION FOR OPTIMAL STARTING POINT OF MODIFIED EVANS ALGORITHM (UTOWN NETWORK)

Iteration	Optimal Start
1	11.582
2	10.610
3	10.314
4	10.188
5	10.082
10	9.911
20	9.736

Units are 100,000 vehicle-minutes.

tions at about the same rate, as measured by iterations. However, this convergence rate is unexpectedly slow. An unacceptable 10 percent convergence error remains after 20 iterations on the UTOWN network and after 10 iterations on the East Brunswick network. A more reasonable error of 5 percent is reached after about 50 iterations on the UTOWN network. Given these slow convergence rates, it is more appropriate to refer to "near-equilibrium" solutions, that is, solutions within some acceptable error limit.

The most disturbing aspect of convergence error is its random component. Even a small amount of random error can completely invalidate comparisons of close alternatives; the only proven method of reducing random error is to run more iterations. Incremental assignment algorithms appear to have much smaller random components in their convergence errors, suggesting that fewer iterations are required.

The existence of convergence error should force planners to adopt innovative methods of assignment. For example, it is sometimes possible to forecast only the increment of traffic due to site development. Such a forecast will have more validity if the random error can be confined to the increment, while treating any errors in background volumes as entirely systematic.

Ad hoc algorithms are not necessarily bad. It is possible for an ad hoc algorithm to greatly outperform a rigorously derived algorithm, given the same computer budget. Because ad hoc algorithms do not come with a pedigree, confidence in an ad hoc algorithm must be established through extensive testing.

If the results of a simulation are to be readily accepted, its algorithms must be lucid. Given the choice, planners should pick an assignment algorithm that can be easily explained to decision makers. The elastic-demand incremental algorithm is conceptually simple; Evans's algorithm is conceptually complex. Both algorithms produce essentially the same answer.

The existence of several algorithms that can consistently produce near-equilibrium solutions to a given traffic model should enhance prospects of improving the model. Model developers should concentrate on incorporating better traffic theory and not be overly concerned with finding an algorithm that delivers the intended solution. The algorithm appearing to adapt most easily to different traffic models is elastic-demand incremental assignment.

Overall, the tests indicate that elastic-demand incremental assignment produces the best solutions. The method is easy to implement, it can be quickly modified to handle a variety of demand models, and it converges reasonably well. Its speed of convergence is no worse than that of more precise algorithms; its ad hoc error is insignificant; it is relatively insensitive to the starting point; it has the best resolution among the tested algorithms; and it is easy to understand. The relative success of elastic-demand incremental assignment contributes evidence of the resiliency of incremental (or successive average) methods.

REFERENCES

1. L. Leblanc, E. Morlok, and W. Pierskella. An Efficient Approach to Solving the Road Network Equilibrium Traffic Assignment Problem. *Transportation Research*, Vol. 9, 1975, pp. 309–318.
2. S. P. Evans. Derivation and Analysis of Some Models for Combining Trip Distribution and Assignment. *Transportation Research*, Vol. 10, 1976, pp. 37–57.
3. G. R. M. Jansen and P. H. L. Bovy. The Effect of Zone Size and Network Detail on All-or-Nothing and Equilibrium Assignment Outcomes. *Traffic Engineering and Control*, Vol. 23, 1982, pp. 311–317.
4. B. R. Wildermuth, D. J. Delaney, and K. E. Thompson. Effect of Zone Size on Traffic Assignment and Trip Distribution. In *Highway Research Record 392*, HRB, National Research Council, Washington, D.C., 1972, pp. 58–75.
5. N. J. Pedersen and D. R. Samdahl. *NCHRP Report 255: Highway Traffic Data for Urbanized Area Project Planning and Design.* HRB, National Research Council, Washington, D.C., 1982.
6. Y. Sheffi and W. Powell. A Comparison of Stochastic and Deterministic Traffic Assignment over Congested Networks. *Transportation Research*, Vol. 15B, 1981, pp. 53–64.
7. C. Fisk. Some Developments in Equilibrium Traffic Assignment. *Transportation Research*, Vol. 14B, 1980, pp. 243–255.
8. W. B. Powell and Y. Sheffi. The Convergence of Equilibrium Algorithms and Predetermined Step Sizes. *Transportation Science*, Vol. 16, 1982, pp. 45–55.
9. L. J. LeBlanc, R. V. Helgason, and D. E. Boyce. Improved Efficiency of the Frank-Wolfe Algorithm for Convex Network Programs. *Transportation Science*, Vol. 19, 1985, pp. 445–462.