Convergence of Traffic Assignments: How Much is Enough?[1]

David Boyce, M.ASCE[2]; Biljana Ralevic-Dekic[3]; and Hillel Bar-Gera[4]

Abstract: Daily traffic assignments to a large-scale road network are described for build and no-build scenarios to evaluate the addition of two proposed ramps between I-295 and SR-42 in the New Jersey part of the Delaware Valley region. The road network consists of 39,800 links connecting 1,510 zones. The user-equilibrium traffic-assignment problem was solved with a new algorithm called origin-based assignment (OBA), which can achieve highly converged solutions with reasonable computing effort. Following a description of the user-equilibrium traffic-assignment problem and the OBA algorithm, the stability of link-flow differences between the two scenarios in the vicinity of the proposed ramps are examined over a broad range of assignment convergence levels. Then, link-flow differences over this range of convergence levels are compared to link-flow differences between two very highly converged solutions. Examination of the findings reveals, in the writers' view, that a relative gap of 0.01% (0.0001) is required to ensure that the traffic assignments are sufficiently converged to achieve link-flow stability. These convergence levels are then interpreted in terms of the number of Frank-Wolfe iterations needed to achieve comparable relative gaps as well as the computational effort required.

CE Database subject headings: Traffic assignment; Convergence; Algorithms; Trip forecasting.

Problem

In performing traffic assignments on road networks for large urban areas, practitioners making travel forecasts are faced with the practical question of how many iterations to perform, or equivalently, how long to wait for the result. This question is rendered even more perplexing by the very slow rate of convergence of the Frank-Wolfe method (e.g., EMME/2 1998) now widely used in travel-forecasting software systems. Practitioners notice when applying this method that link flows fluctuate substantially from iteration to iteration, or drift gradually up or down, which gives little assurance that an adequate approximation of the equilibrium solution is achieved. Curiously, the same practitioners often use the phrase ''reaches closure'' to describe the termination of the solution procedure as if a satisfactory solution has actually been found.

In fact, the Frank-Wolfe method does not achieve a stable or highly converged solution with any reasonable amount of computational effort for large, congested urban networks. Whether the solutions obtained by the Frank-Wolfe method after 10, 50, or even 500 iterations are a useful basis for transportation planning

decisions depends upon their discrepancies from a highly converged solution that approximates the true user equilibrium. The objective of this paper is to perform a case study that examines this convergence issue.

Origin-based assignment (OBA) algorithm, devised by Bar-Gera (1999, 2002), allows the question of the desired level of convergence of solutions to the traffic-assignment problem (TAP) to be addressed meaningfully for the first time. This algorithm solves TAP up to any desired level of convergence within a reasonable computational effort up to the accuracy of the computer.

One criterion for determining the desired level of convergence is that the solution should be stable in terms of link flows. That is, link flows should not fluctuate as the solution converges further. We interpret this stability as a requirement in the comparison of two scenarios, such as in a build versus no-build analysis. A second criterion is that the computation time required to achieve this solution should be reasonable. Presently, for large-scale networks, a reasonable solution time is often considered to be an overnight computer job or up to 12–14 h.

Following the definition of TAP and related terms, and a brief description of OBA, a case study is presented based on a proposal to construct a new pair of ramps between I-295 and SR-42 in Camden County, N.J., in the Delaware Valley region. Link-flow differences are examined for the freeway links and for arterials with interchanges with I-295 near the proposed ramps. Similar results are considered for the relative errors in link-flow differences. Finally, the computational effort to achieve stable solutions is considered for OBA and for the Frank-Wolfe algorithm widely used in professional practice. A discussion of results concludes the paper.

Definitions

In order to compare build and no-build solutions, a common measure of convergence is required. To compare solutions for differ-

---

ent networks, the relative gap (RG) is proposed. This measure assures the two solutions are equally converged in terms of their respective proximity to the optimal value of the user-equilibrium objective function. TAP, for user-equilibrium route choice with fixed daily origin-destination demand, is defined as follows:

$$\min_{h} z(\mathbf{h}) = \sum_{a} \int_{0}^{f_a} t_a(x)dx \qquad (1)$$

subject to

$$\sum_{r \in R_{pq}} h_{pqr} = d_{pq} \quad \text{for all} \quad pq \qquad (2)$$

$$h_{pqr} \geqslant 0 \quad r \in R_{pq} \quad \text{for all} \quad pq \qquad (3)$$

where

$$f_a \equiv \sum_{pq} \sum_{r \in R_{pq}} h_{pqr} \delta_{pqr}^{a} \quad \text{for all} \quad a \qquad (4)$$

$t_a(f_a)$=travel time on link a at flow f_a; d_{pq}=flow of vehicles per day from origin zone p to destination zone q (given); h_{pqr}=flow of vehicles per day on route r, $r \in R_{pq}$, the set of routes from p to q; and $\delta_{pqr}^{a} = 1$, if link a belongs to route r, $r \in R_{pq}$; otherwise it is 0.

The gap at iteration k may be defined as

$$\text{gap}(k) = \sum_{a} t_a[f_a(k)] \cdot [y_a(k) - f_a(k)] \leqslant 0 \qquad (5)$$

where $y_a(k)$=vehicle flow on link a at iteration k given by an all-or-nothing assignment based on link travel times, $\{t_a[f_a(k)]\}$. Using the gap, a lower bound (LB) on the value of the objective function $z(\mathbf{h})$ is defined as

$$\text{LB}(k) = z[\mathbf{h}(k)] + \text{gap}(k)$$

$$= \sum_{a} \int_{0}^{f_a(k)} t_a(x)dx + \sum_{a} t_a[f_a(k)] \cdot [y_a(k) - f_a(k)] \qquad (6)$$

The best lower bound (BLB) and RG are then defined as

$$\text{BLB} = \max_{k}[\text{LB}(k)] \qquad (7)$$

$$\text{RG}(k) = \frac{-\text{gap}(k)}{|\text{BLB}|} \geqslant 0 \qquad (8)$$

where $|\text{BLB}|$=absolute value of the BLB.

Since the emphasis in this analysis is on comparing scenarios, after some experimentation, the differences in link flows between the two scenarios were chosen as the primary basis for assessing the stability of the solutions. Accordingly, the first analysis examines the link-flow differences between scenarios for a range of RGs. Once the link flows have been determined for a given RG $= 10^{-n}$, one may wish to compare this solution with a highly converged or true solution, say an RG of 10^{-11}. The procedure for computing this relative error of link-flow differences for link a is

$$\text{Relative error for link } a(10^{-n}|10^{-11}) = \left| \frac{\Delta f_a^{10^{-n}} - \Delta f_a^{10^{-11}}}{\Delta f_a^{10^{-11}}} \right| \qquad (9)$$

where $\Delta f_a^{10^{-n}}$=difference in vehicle flows for link a between the

build and no-build solutions at an RG of 10^{-n}; and $\Delta f_a^{10^{-11}}$=difference in vehicle flows for link a between the build and no-build solutions at an RG of 10^{-11}.

The convergence level at an RG of 10^{-11} may be regarded as a very close approximation of the true equilibrium solution to TAP at the accuracy of a Sun Microstation Ultra 10 computer.

Origin-Based Assignment

The OBA algorithm was formulated, implemented, and extensively tested on large networks by Bar-Gera (1999, 2002) starting with his PhD thesis and continuing since that time. An executable code for the algorithm has recently been made available for research and evaluation (Bar-Gera 2003). This section seeks to provide a short, relatively nonmathematical overview of the algorithm.

The Frank-Wolfe algorithm may be characterized as being link-based since the solution variables are link flows. Other algorithms are route-based in that each route flow is a solution variable. The concept of the origin-based assignment algorithm is to define the solution variables in an intermediate way between links and routes. In this way, the algorithm seeks to solve larger networks more efficiently than can be solved by either of the two approaches.

The main solution variables in this algorithm are origin-based approach proportions α_{pa} for every origin p and every link a, such that for every origin p and node i the sum of origin-based approach proportions over all links ending at node i is equal to one. Using origin-based approach proportions, route proportions are determined as the product of approach proportions of all the links along the route—that is $\gamma_{pqr} = \Pi_{a \subseteq r} \alpha_{pa}$. Route flows are determined as the product of origin-destination flow and route proportion, that is $h_{pqr} = d_{pq} \cdot \gamma_{pqr}$. Bar-Gera [2002 Eq. (14)] has shown for link a from node i to node j, if the total flow from origin p to node j is g_{pj}, then the total flow from origin p that arrives at node j through link a is $f_{pa} = \alpha_{pa} \cdot g_{pj}$. In that respect α_{pa} is indeed the proportion of flow on approach a to node j for origin p as implied from the name of these variables.

The representation of the solution by origin-based approach proportions allows storing a complete description of the route flows very efficiently. The efficiency of the representation is further enhanced using the fact that at most nodes, one link receives an approach proportion value of one, while the value of all other links ending at the same node is zero. The availability of route flows can be useful for solution analysis. It is also useful in searching for the equilibrium solution, which is a major difference from many alternative solution procedures, including the Frank-Wolfe algorithm, which stores only total link flows during the iterative process.

A key point in the algorithm is the following: for every origin p, an a-cyclic restricting subnetwork A_p is chosen such that approach proportions of links that are not included in A_p are restricted to 0. Using the equation for route proportions for every origin, it can be seen that under these restrictions, only routes that are limited to the links in its restricted subnetwork can be used. In particular, since A_p is a-cyclic, meaning that it does not contain a directed cycle of links, any cyclic route must contain at least one link that does not belong to A_p, and hence the flow along any cyclic route must be zero. Note that the restriction to a-cyclic subnetworks does exclude many solutions that do not use cyclic routes, which are usually considered legitimate. Bar-Gera [2002 (Lemma 3)], among others, has shown that there is always a

user-equilibrium solution that is a-cyclic by origin. Therefore, this restriction does not prevent the algorithm from converging to the true equilibrium solution.

The restriction to solutions that are a-cyclic by origin has several important advantages. First, the simple route-flow interpretation presented previously is, in fact, only valid for solutions that are a-cyclic by origin. Second, a-cyclic subnetworks allow one to definite an origin-specific topological ordering of the nodes, so that every link in the restricting subnetwork goes from a node of lower topological order to a node of higher topological order. Most computations in the proposed algorithm are done in a single pass over the nodes either in ascending or descending topological order. The time required by such computations is a linear function of the number of links in the network. This computational efficiency is the main reason for restricting to a-cyclic solutions.

In solving TAP, the algorithm starts with trees of minimum cost routes as restricting subnetworks leading to an all-or-nothing assignment. Then, the algorithm considers all origins in sequential order. For each origin, the restricting subnetwork is updated, and the origin-based approach proportions are adjusted within the given restricting subnetwork. To update a restricting subnetwork, unused links are removed, v_i the maximum cost from the origin to node i within the restricting subnetwork is computed for all nodes, and all links $[i,j]$ such that $v_i < v_j$ are added to the restricting subnetwork. Once a new restricting subnetwork is found, several computationally intensive steps are needed including reorganization of the data structure. However, restricting subnetworks tend to stabilize fairly quickly. Therefore, it is useful to update origin-based approach proportions while keeping the restricting subnetworks fixed. This is done by introducing inner iterations as described below in the flow chart.

To update origin-based approach proportions within a given restricting subnetwork, a search direction based on shifting flow from high-cost alternatives to low-cost alternatives is used. In addition to current costs, estimates of cost derivatives are used to improve the search direction in a quasi Newton fashion. When two independent routes are considered, the amount of flow shifted by the search direction equals the difference between route costs divided by the sum of route cost derivatives, which is exactly what would be obtained by considering the second-order approximation of the objective function.

This second-order search direction is viewed as desirable flow shifts, which are scaled by a step size between zero and one and then truncated to guarantee feasibility. This technique is referred to as the boundary-search procedure, since it tends to choose solutions along the boundary, although it does consider interior points as well. This technique is somewhat different from conventional line-search techniques, where shifts are first truncated to guarantee feasibility and only then scaled by a step size. The importance of the boundary search for origin-based assignment is that it is effective in eliminating residual flows, i.e., small flows on suboptimal routes. The elimination of residual flows is critical for algorithm convergence. See Bar-Gera (2002) for details.

In order to guarantee descent of the objective function, and convergence of the algorithm, the search considers step size values of 1, 1/2, 1/4, 1/8, etc. The stopping condition is based on the concept of social pressure introduced by Kupsizewska and Van Vliet (1999). The basic idea is that travelers shifted from route r_1 to r_2 apply pressure (positive or negative), which is equal to their gain (or loss) according to the difference in route costs resulting from the shift.

The total social pressure is the sum of the pressure from all the travelers. Our search direction is good in the sense that it always

enjoys positive social pressure for small step sizes. As the step size increases, the social pressure decreases, and eventually it may become negative. Our goal is to find the largest step size, i.e., the first in the sequence 1, 1/2, 1/4, 1/8... with positive social pressure. This social-pressure principle is in fact equivalent to the stopping condition of the line search in the Frank-Wolfe algorithm, only this principle is applicable in certain cases where the line-search optimization rule is not. The algorithm is described by the following flow chart:

Initialization:
    for every origin p
        let A_p be a tree of minimum cost routes under free-flow conditions from p
        let $\alpha_{pa} = 1$ for all links in A_p and 0 otherwise (all-or-nothing assignment)
    end for
Main loop:
    for $n=1$ to number of main iterations
        for every origin p
            update restricting subnetwork A_p
            update origin-based approach proportions α_{pa}
        end for
        for $m=1$ to number of inner iterations
            for every origin p
                update origin-based approach proportions α_{pa}
            end for
        end for
    end for
Update restricting subnetwork for origin p:
    remove unused links from A_p
    for every node i compute the maximum cost v_i from p to i
    for every link $a = [i,j]$
        if $v_i < v_j$ add link a to A_p
    find new topological order for new A_p
    update data structures
Update origin-based approach proportions for origin p:
    compute average costs and Hessian approximations
        for step size 1, 1/2, 1/4, 1/8...
            compute flow shifts and scale by step size
            project and aggregate flow shifts
            if social pressure is positive then stop
        end for
    apply flow shifts
    update total link flows and link costs

Case Study

Background

A large-scale hypothetical case study using origin-based assignment was presented to the Transportation Network Modeling Committee of the Transportation Research Board in 2001 (Boyce et al. 2001). As these computational studies for the Chicago Region road network were quite promising, additional networks with actual improvement proposals were sought to test the algorithm further.

In response, Dr. W. Thomas Walker, of the Delaware Valley Regional Planning Commission (DVRPC) in Philadelphia, related difficulties he had encountered in performing 24 h traffic assignments to test a pair of proposed ramps connecting I-295 and SR-42 in Camden County. Applying widely used–travel forecast-
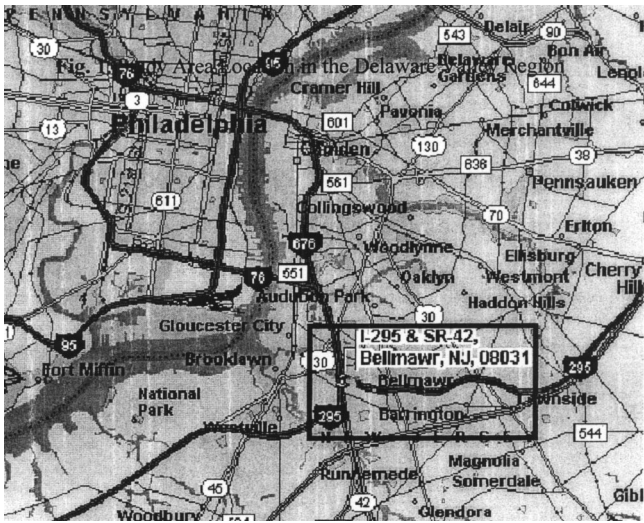
Fig. 1. Study area location in Delaware Valley region

ing software, he had been unable to obtain assignments that did not require substantial manual adjustment. In particular, in comparing build and no-build scenarios, he noticed substantial changes in link flows throughout the Delaware Valley region, even though this freeway interchange is located in one corner of the region relatively far from central Philadelphia. Refer to Fig. 1 for the freeway interchange location in the Delaware Valley region.

Through the cooperation of Walker, the road network and 24 h origin-destination table used in his analyses up until early 2000 were furnished to the writers. (Subsequent studies for the New Jersey DOT have used somewhat refined networks and a different trip table for the build alternative.) The case study presented below is based on the network and trip table provided in early 2000.

The freeway interchange examined in the case study is shown in Fig. 2. The interchange is unusual for its complexity. The north–south freeway, SR-42, approaches the Walt Whitman Bridge spanning the Delaware River just north of the area shown in Fig. 2. To the south, SR-42 provides direct access to the Atlan-
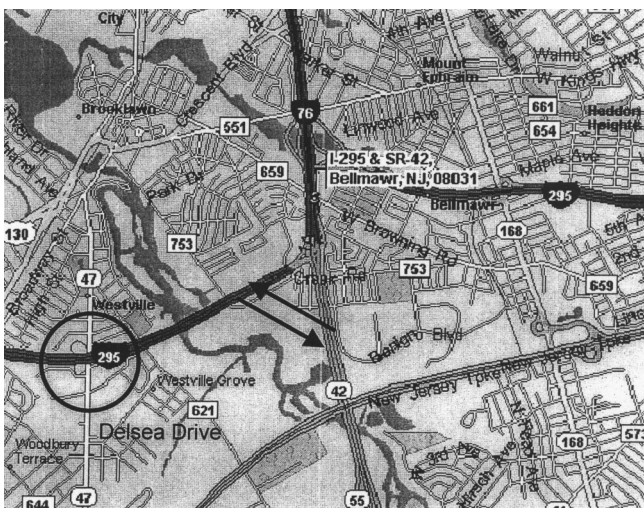

Fig. 2. Interchange of I-295 with SR-42 showing proposed ramps and interchanges with arterial streets

tic City Expressway, a toll road serving Atlantic City, N.J. I-295 is a circumferential route that parallels I-95, which is located on the opposite side of the Delaware River in Pennsylvania. As can be seen in Fig. 2, the design of I-295 in the vicinity of SR-42 facilitates traffic flows to and from the Walt Whitman Bridge and Philadelphia. No provision was made in the original design for eastbound traffic flows on I-295 to turn south onto SR-42. Therefore, vehicles desiring to make this right-hand turning movement must leave I-295 west of SR-42 and make their way over the arterial street system to SR-42. A pair of ramps, shown diagrammatically also on Fig. 2, was proposed to provide for this turning movement from eastbound I-295 to southbound SR-42 and from northbound SR-42 to westbound I-295. The objective of Walker's study was to determine the flows on these proposed ramps and their effect on the surrounding freeway and arterial roadways.

Procedure

In this case study, attention is focused on the freeway links east and west of the junction of the ramps with I-295 and the freeway links north and south of the junction of the ramps with SR-42. These links are labeled in the following figures using this terminology. The effect of the proposed ramps on the traffic flows on the north–south arterials with interchanges with I-295 lying west and east of its interchange with SR-42 was also examined. These arterials are Delsea Drive on the west and Black Horse Pike on the east. Representative results are given for the Delsea Drive interchange shown by a circle on Fig. 2.

Using OBA, the 24 h link flows were found for the build and no-build scenarios for the above links for RGs ranging from 10 to $10^{-11}\%$. As described in the Definitions section, the RG may be defined as the ratio of the gap to the BLB. It may also be defined as the ratio of the gap to the objective function. The former definition is more conservative—especially for less-converged solutions. In the following, the RG is expressed as a percentage rather than as a ratio since this convention is used in some travel-forecasting software and may be more intuitive. Also, the number of Frank-Wolfe iterations needed to achieve the indicated RG is shown. The origin-based and Frank-Wolfe solutions are not equivalent; rather, they only have the same RG.
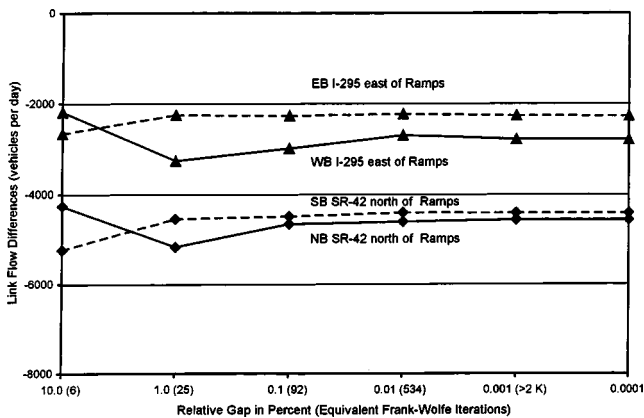
Freeway Links

First, the results for freeway links and ramps are presented. Because of the wide range of link-flow differences, both positive and negative, these results are divided into two figures. Fig. 3 shows the freeway links with positive flow differences (build less no-build), including the proposed ramps, and Fig. 4 shows the links with negative flow differences. Note the scale of the link-flow differences is the same on these two figures. On the horizontal axis, the RG is shown from 10 to 0.0001%. This range was selected to illustrate how the link-flow differences stabilize as the solutions converge. The individual plots are labeled according to the direction of flow and the location relative to the proposed ramps.

In Fig. 3, one may see that the link-flow differences for the two ramps decrease rather substantially from an RG of 10% to stable values in the range of 0.01 to 0.001%. Since changes of less than one vehicle per day occur beyond an RG of 0.0001%, the figure is truncated at this level. For the I-295 links west of the pair of ramps, some fluctuations also occur between RGs of 0.01 and 0.001%. In contrast, the flow differences for SR-42 south of the ramps are essentially stable once an RG of 0.01% is reached.
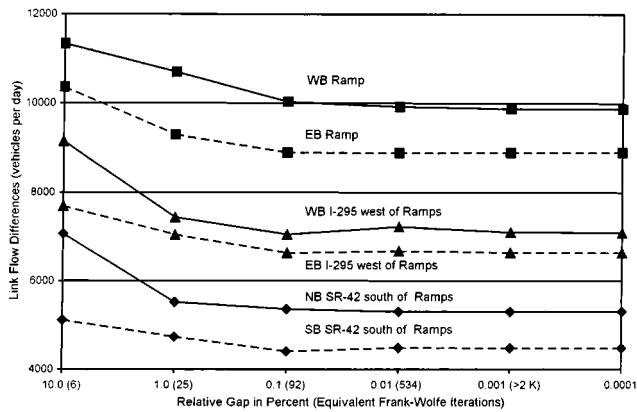
Fig. 3. Link flow differences (build less no-build) versus relative gap New Jersey freeway links with positive flow differences
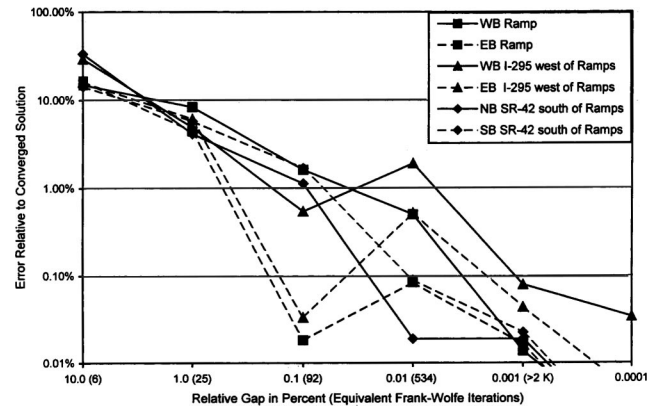

Fig. 5. Error relative to converged solution versus relative gap—New Jersey freeway links with positive flow differences

Fig. 4. Link flow differences (build less no-build) versus relative gap—New Jersey freeway links with negative flow differences

Fig. 6. Error relative to converged solution versus relative gap—New Jersey freeway links with negative flow differences
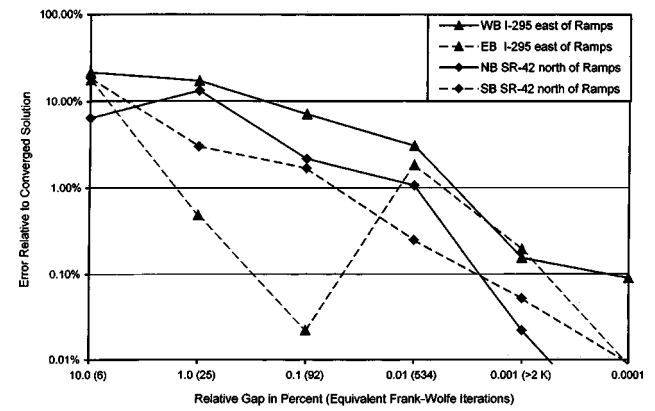
Links with negative link-flow differences are shown in Fig. 4. These are the links downstream from the proposed ramps. As with Fig. 3, some fluctuations in flow differences occur up to an RG of 0.001%. Note that the less converged solutions are rather different from the converged solutions.

Next, the flow differences for the same links are considered on a relative basis. As explained in the Definitions section, the link-flow differences for a given convergence level are compared with the most highly converged solution. Actually, there is little change in the solutions for these links beyond an RG of 0.0001%. As one may see in Fig. 5, all links but one achieve an error relative to the converged solution of less than 1% at an RG of 0.01%. However, for higher RGs, the relative errors (REs) can be quite large. For example, the links with positive link-flow differences have REs of 4 to 8% at an RG of 1%, which corresponds to 25 Frank-Wolfe iterations, and as high as 33% for one link at an RG of 10% corresponding to 6 Frank-Wolfe iterations. For the links with negative link-flow differences as shown in Fig. 6, the results are somewhat more problematic. Three of the four links have REs greater than 1% at an RG of 0.01%, and the REs are even higher for less-converged solutions. An exception is eastbound I-295 east of the proposed ramp.

The conclusion for freeway links and ramps based on Figs. 3–6, as well as details not presented here, is that an RG of 0.01% is necessary to ensure that the changes in link-flow differences are less than 200 vehicles per day from the solutions at 0.1%, and the

REs are less than 3% of the highly converged solutions. Similar findings were reported for a hypothetical study of adding lanes to the I-290 Expressway in western Cook County by Boyce et al. (2001) and Ralevic-Dekic (2000).

Arterial Links

In the case of arterial links with interchanges near the proposed ramps, results for four links are presented. Fig. 7 shows link-flow differences for Delsea Drive west of the proposed ramps. Note the scale of the vertical axis is only 1/8 of the scale of Figs. 3 and 4.

As can be readily observed in Fig. 7, an RG of 0.01% is necessary to achieve stability of link-flow differences. Although the flow differences are substantially smaller than for freeway links, the capacities of these links are also much smaller. In the writers' view, such stability of flow differences is important and needs to be achieved for planning of arterial links as well as freeways. Note that the change in link-flow differences from an RG of 1 to 0.01% can be as large as several hundred vehicles per day.

Fig. 8 presents the error relative to converged solution for Delsea Drive. For this arterial, errors well in excess of 100% are observed for an RG of 10%. To achieve an RE of less than 10%, an RG of 0.01% is required. In the writers' view, an RE of 10% is acceptable for arterials, which implies convergence to an RG of 0.01%.
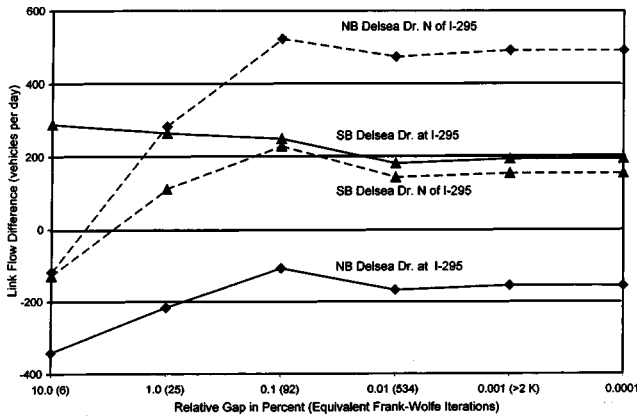
Fig. 7. Link flow differences (build less no-build) versus relative gap—Delsea Drive crossing I-295 west of SR-42
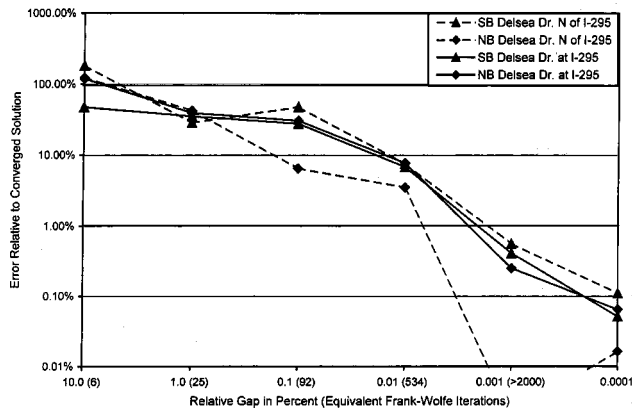

Fig. 8. Error relative to converged solution versus relative gap—Delsea Drive crossing I-295 west of SR-42

Comparison of Computational Effort

Figs. 3–8 show on the horizontal axis the equivalent number of iterations of the Frank-Wolfe algorithm, which is widely used in professional practice to achieve the indicated RG. This section explains and elaborates on this information.

Several algorithms are available for solving TAP to a specified RG. Given this information, these solutions can be compared. However, the solutions are not equal or necessarily of the same quality. Using *EMME/2*'s capability (EMME/2 1998) to solve TAP with the Frank-Wolfe algorithm, 2,000 iterations were performed and the RGs of the solutions were computed. The number of iterations needed to reach the indicated RG is shown in Table 1 for a Sun Microsystems Ultra 10 with 1 gb of memory using the Solaris operating system.

As noted, *EMME/2* reached the recommended RG of 0.01% in 534 iterations. To check that this level of convergence is stable, OBA was used to compute solutions for smaller RGs, also as shown in Table 1. The number of Frank-Wolfe iterations necessary to reach these levels of convergence is unknown but very large. The Frank-Wolfe algorithm proceeds by very gradually shifting flows from more costly routes to the current least costly route. However, the algorithm never succeeds in removing all of the flow from the more costly routes (in a finite number of iterations). In contrast, OBA removes such residual flows, which is one of the reasons that such fine convergence is possible.

Table 1. Computational Effort for OBA and Frank-Wolfe Algorithm

Relative gap (%)	Equivalent number of Frank-Wolfe iterations	Computation effort by algorithm (h)		Excess costs of origin-based solutions	
		Frank-Wolfe (*EMME/2*)	OBA	Total (h)	Average (min)
1.0E+01	6	0.18	0.96	384,067.88	1.61E+00
1.0E+00	25	0.64	1.53	28,254.71	1.18E−01
1.0E−01	92	2.31	2.35	3,129.54	1.31E−02
1.0E−02	534	13.28	3.41	338.12	1.42E−03
1.0E−03	>2,000[a]	>50.00	5.80	30.07	1.26E−04
1.0E−04	—	—	8.04	3.22	1.35E−05
1.0E−05	—	—	10.82	0.29	1.20E−06

[a]Relative gap at 2,000 iterations=0.0024% (2.4E−03%).

A characteristic of TAP respected by OBA is that the equilibrium solution does not include any cycles. A simple example of a cycle is a route that circles a city block, which is obviously not a least-cost route. More subtle examples of cycles are possible, however, such as two routes from a given origin to a given destination that traverse the same street segment in opposite directions. It can be shown that one of these routes is not a least-cost route and therefore should be removed. The Frank-Wolfe algorithm can easily generate such a cycle, which is never completely eliminated. Janson and Zozaya-Gorostiza (1987) suggest this may be one reason for its very slow convergence. As noted above, OBA adds routes in a manner that a-cyclic solutions are always maintained.

The computational effort required by *EMME/2* and OBA to reach the indicated RGs is also shown in Table 1 for the Sun Ultra 10. Even with this older 333 MHz computer, the recommended RG of 0.01% can be reached by *EMME/2* in an overnight job. However, professionals do not often solve TAP with as many as 500 iterations. OBA achieves this recommended RG in 3.4 h. However, it requires somewhat more effort than *EMME/2* for less converged solutions.

One additional advantage of OBA is that route-based measures of convergence can be readily computed. Such measures are useful in assessing the quality of the solution. Two basic measures found to be especially helpful are the total and the average excess cost, which are equivalent to the gap. These measures are defined below using the notation in the Definitions section. For a given solution of TAP, which is not necessarily fully converged, define the following terms: u_{pq}=cost of the shortest route from zone p to zone q (minutes); $C_{pqr}=\Sigma_{a \in A} t_a(f_a)\delta_{pqr}^a$=cost of route r from zone p to zone q (minutes); and $E_{pqr}=C_{pqr}-u_{pq}$=excess cost of route r from zone p to zone q (minutes).

The measures may be defined as follows:

$$\text{Total excess cost} = \sum_{pq} \sum_{r \in R_{pq}} h_{pqr} E_{pqr} \qquad (10)$$

$$\text{Average excess cost} = \text{Total excess cost} / \sum_{pq} d_{pq} \qquad (11)$$

The total and average excess costs for a given solution are intuitive measures of the additional travel time of nonoptimal routes, weighted by the user route flows of the current solution. The total excess cost is simply the total additional route time over the minimum OD time of all route flows in the network. All costs are measured in vehicle minutes. The values of total and average excess costs for the solutions corresponding to each RG are shown in Table 1. In Table 2, the average congested and free-flow

Table 2. Delaware Valley Regional Planning Commission Road Network Indicators

Indicator	Congested	Free-flow[a]
Average travel time (min)	23.42	17.84
Average travel distance (km)	12.64	12.64
Space-mean-speed (kph)	32.41	42.54

Note: Total flow=14,336,062 vehicles per day. Total vehicle hours of travel=4,334,532 h per day.

[a]At user-equilibrium flows.

travel times and distances are shown together with the total interzonal flow per day and total vehicle hours per day on the network.

Conclusions

For the DVRPC case study, the following conclusions are offered:

1. For the build and no-build networks examined, an RG of 0.01% (or a ratio of 10^{-4}) is required to reach stability of link-flow differences for both freeway and arterial links. This level of convergence is reached by OBA in about 3.4 h on a Sun Ultra 10. Achieving the same RG requires over 500 iterations of the Frank-Wolfe algorithm widely used in professional practice, which takes 13.3 h on the same computer using *EMME/2*.

2. These results pertain to the assignment of the daily (24 h) trip table provided. Although the writers believe that peak-period assignments provide more meaningful results for both evaluation and design, the findings of this case study should be applicable to both hourly and daily assignments. Generally, the assignment of a congested, peak-period trip table requires somewhat more computational effort.

3. Utilizing OBA to solve TAP to this or higher levels of convergence will be even more worthwhile as improvements are made in network representation, such as improved flow-delay functions at intersections, as well as other inputs to the travel-forecasting procedure.

4. The results presented in this paper are for one network and a single trip table representing one level of daily congestion. Additional studies are needed to validate these results for other networks and various congestion levels.

Several software developers are currently adding OBA to the assignment options available in their systems. Practitioners who appreciate that assignment results of this quality are necessary to distinguish among competing scenarios should require improved options from their software vendor. In view of the findings of this research, practitioners urgently need traffic assignment accuracy standards and application guidelines in order to produce meaningful results, particularly for scenario analyses.

Acknowledgments

References

Bar-Gera, H. (1999). "Origin-based algorithms for transportation network modeling." PhD thesis, Univ. of Illinois at Chicago; *Technical Rep. No. 103*, ⟨http://www.niss.org/⟩.

Bar-Gera, H. (2002). "Origin-based algorithm for the traffic assignment problem." *Transp. Sci.,* 36(4), 398–417.

Bar-Gera, H. (2003). "Executable code for origin-based assignment." ⟨http://www.openchannelsoftware.org/projects/origin-based-assignment⟩.

Boyce, D., Ralevic-Dekic, B., and Bar-Gera, H. (2001). "Convergence of traffic assignments: How much is enough?" *Annual Meeting, Transportation Research Board*, Washington, D.C., presentation.

EMME/2 release 9. (1998). INRO Consultants, Inc., Montréal.

Janson, B., and Zozaya-Gorostiza, C. (1987). "The problem of cyclic flows in traffic assignment." *Transp. Res.,* 21B(4), 299–310.

Kupsizewska, D., and Van Vliet, D. (1999). "101 uses for path-based assignment." *Transport Planning Methods: Proc., Seminar C, PTRC Planning and Transport Summer Ann. Meeting*, Univ. of Sussex, U.K.

Ralevic-Dekic, B. (2000). "Convergence of link flow solutions by origin-based and link-based traffic assignment algorithms." MS thesis, Univ. of Illinois at Chicago.