#4

Paper No. 00-0870

Optimal Bus Stop Spacing Through Dynamic Programming and Geographic Modeling

Peter G. Furth and Adam B. Rahbee

A discrete approach was used to model the impacts of changing bus-stop spacing on a bus route. Among the impacts were delays to through riders, increased operating cost because of stopping delays, and shorter walking times perpendicular to the route. Every intersection along the route was treated as a candidate stop location. A simple geographic model was used to distribute the demand observed at existing stops to cross-streets and parallel streets in the route service area, resulting in a demand distribution that included concentrated and distributed demands. An efficient, dynamic programming algorithm was used to determine the optimal bus-stop locations. The model was compared with the continuum approach used in previous studies. A bus route in Boston was modeled, in which the optimal solution was an average stop spacing of 400 m (4 stops/mi), in sharp contrast to the existing average spacing of 200 m (8 stops/mi). The model may also be used to evaluate the impacts of adding, removing, or relocating selected stops.

A design decision in establishing or reviewing a bus line is the spacing and location of stops. The main societal costs and benefits related to stop location can be quantified, and therefore using a quantitative model to select optimal stop locations is a reasonable expectation. Even if practical and political considerations play a role in such design decisions, the political process must include results of the best available scientific analysis.

Bus-stop location decisions have three main societal impacts that involve a trade-off between the costs and benefits of more frequent stops:

1. Riding time—the more frequent the stops, the more time through passengers spend in the vehicle;
2. Operating cost—the more frequent the stops, the greater the cycle time and therefore the operating cost of the route; and
3. Walking time—the more frequent the stops, the shorter the walking time to access the route.

A fourth potential impact is that greater spacing between stops may result in longer walking distances and a loss of passengers. A primary objective of most transit systems, however, is to maximize ridership; therefore, stop location decisions are usually constrained so that they do not have a significant impact on ridership. This analysis imposed such a constraint and held ridership fixed. Also, factors were excluded that tend to dominate rail station location decisions, such as capital cost, neighborhood and development impact, and a prevalence of nonwalking (automobile and bus) access modes.

Most transit agencies recognize the trade-offs in stop spacing decisions by establishing stop spacing guidelines. A recent study found that 95 of 111 responding U.S. agencies have stop spacing guidelines, with about one-half recommending spacing of 200 to 270 m (six to eight stops/mi), and closer spacing in business districts (*1*). However, policies are not uniform, ranging from more than 400 m (fewer than 4 stops/mi) to less than 130 m (more than 12 stops/mi) in areas outside the central business district (CBD).

A motivation for this research was the Massachusetts Bay Transportation Authority's (MBTA) review of its bus lines in an effort to increase operating speed by eliminating or relocating stops. Bus stops in northern Europe, where transit has a much greater market share despite comparable levels of affluence, are generally located considerably farther apart than in the United States. One observer noted an average bus-stop spacing between 400 and 530 m, and cited a recommendation of 300 m spacing from one German city official, versus the standard U.S. practice of 160 to 230 m (7 to 10 stops/mi) (*2*). The MBTA's guideline, albeit an informal one, is that bus stops should be spaced about 200 m apart (8 stops/mi) outside commercial areas, which is the spacing used by many other large transit systems (*3*). In contrast, according to Benn, the Chicago Transit Authority recommends 400 m spacing (4 stops/mi), unless there is a major intersection or generator, and Philadelphia's transit authority specifies a minimum spacing of 320 m (5 stops/mi) on suburban and new urban routes (*1*). Benn cites a 1992 study by the New York City Transit Authority that favorably evaluated bus-stop relocations that increased average stop spacing from about 160 to 230 m (from 10 to 7 stops/mi).

The political nature of stop location decisions may account for the close spacing in the United States. The benefits to nearby residents of adding a stop are clear and concentrated, while the impacts on other riders and to the operating agency are diffuse. One purpose of stop-spacing guidelines is to give transit agencies an objective way to resist the pressure to add unnecessary stops or eliminate stops. However, the generality of the guidelines, and the large leeway planners necessarily exercise in their application, hampers the guidelines' effectiveness. A scientific framework that accounts for site-specific impacts, evaluating the social costs and benefits of stop location choices and determining optimal locations, should therefore be valuable for public relations and operations analysis.

Until now, scientific studies of stop spacing have used a continuum modeling approach that yields optimal stop spacing. Converting the recommended spacing to actual stop locations has been left for a later, subjective stage of design. By contrast, the approach

P. G. Furth, Department of Civil and Environmental Engineering, Northeastern University, 420 SN, 360 Huntington Avenue, Boston, MA 02115. A. B. Rahbee, Center for Transportation Studies, Massachusetts Institute of Technology, 1-235, 77 Massachusetts Avenue, Cambridge, MA 02139.

described herein finds optimal locations, with optimal spacing as a byproduct.

First, the continuum approach is reviewed, pointing out its shortcomings and advantages compared with the discrete approach taken in this study. Next, the discrete formulation and its dynamic programming solution are presented, as well as the geographical model used to support the discrete formulation. The method is then applied to a bus line in Boston. The application demonstrates the model's practicality and its sensitivity to various parameters, and offers a comparison with the continuum approach. Finally, some model extensions are described and conclusions are presented.

CONTINUUM APPROACH

The best-developed exposition of the continuum approach is found in Wirasinghe and Ghoneim (*4*). The bus route is modeled as a continuum in one dimension, with x describing the distance from the start of the route. Demand for boarding and alighting is modeled as a continuous function of x, which is assumed to be smooth (slowly changing) except for concentrations at transfer points, assumed to be locations of bus stops. Every point on the route is a candidate stop location. The three main impacts—passenger riding time, operating cost, and passenger walking time—are formulated as functions of parameters describing the demand density and cruising speed in the neighborhood of x, unit costs associated with the three impacts, constants such as the bus acceleration rate, and the decision variable $s(x)$ equals the stop spacing in the neighborhood of x.

Using calculus, optimal conditions were derived for the value of $s(x)$ that minimized the sum of the societal costs in the neighborhood of x. If buses were assumed to stop at every bus stop whether passengers were waiting or not, there was a closed form solution; otherwise, the solution may be obtained by a simple one-dimensional search.

The continuum approach to stop spacing was pioneered by Vuchic and Newell, whose focus was on rail station spacing, assuming feeder bus as the access mode and end-of-the-line demand (*5*). Vaughn and Cousins extended the continuum approach to bus lines, with many-to-many demand and walking access (*6*). The focus was the effect of stop spacing on the competition between walking for the entire trip and taking the bus, using an origin-destination (O-D) matrix as an input. Vaughn and Cousins found little competition between the travel modes, without which the O-D specification of demand is unnecessary. They also recognized that practical constraints do not allow stops to be spaced exactly at the regular intervals, and thus the average walking distance will be one-quarter longer than the stop spacing. However, their assumption of Poisson-distributed stops, which effectively double the average walking distance, seems hardly plausible. Lesley also applied the continuum approach to bus lines, but an assumed radial service area around each stop, with the radius equal to one-half of the stop spacing, led to the peculiar result that demand grows as the stops move farther apart, distorting the walking distance comparisons (*7*).

The main advantage of the continuum approach is that it readily shows the sensitivity of optimal stop spacing to various parameters. For example, this approach demonstrates that, in contrast to most planning guidelines, optimal stop spacing should vary with local conditions. For example, spacing should be greater on sections of a route with high through volume and little local on-and-off traffic, as often occurs in the middle of a route, and smaller if through volume is small and on-and-off traffic is relatively heavy, which is typical at the end of a route. If buses only stop for waiting passengers, the continuum model gives the interesting, and correct, result that stop spacing should be very small—even to the point of stopping on demand (i.e., every point is a bus stop)—on sections with sufficiently low demand. These insights can be used qualitatively, and perhaps quantitatively, by planners to adapt stop spacing to local conditions. Another possible use of the continuum model is to derive an optimal spacing guideline based on "typical" input.

The continuum approach, however, does have shortcomings. After optimal spacing is determined, the process of applying it to a realistic street network, in which stops usually are at intersections, is hardly trivial. For example, if intersections are spaced every 200 m, and the optimal spacing is 300 m, should stops be placed at every intersection or every other intersection? Or should they be located at two out of three intersections, consistent with the optimal stop density?

The other major shortcoming of the continuum approach is that demand is not a smooth, continuous function of location along the route, but is sharply punctuated as each cross-street brings its demand to a specific point on the line—the intersection of the cross-street with the line. Optimal stop locations that recognize those demand concentrations will tend to align themselves with the demand concentrations and may vary considerably from optimal locations, assuming continuously distributed demand. In fact, it should not be surprising if the optimal stop density, based on an approach that recognizes demand concentrations, varies considerably from the optimal density based on a continuum approach.

In conclusion, compared with the continuum approach, the discrete approach is more realistic and provides a more readily applicable solution. Another advantage is its use in evaluating specific sets of stop locations that may be suggested by planners, elected officials, or citizens.

DISCRETE FORMULATION

The discrete approach differs from the continuum approach in two respects. First, a discrete set of candidate stop locations along the route is used (i.e., normally all of the intersections along the route). This approach is consistent with industry practice and with the well-known result of network theory that locating facilities (in this case, bus stops) at intersections minimizes average access distance (*8*).

Second, a geographic model is used to distribute demand to the blocks in the route's service area, without any restriction that the demand should vary smoothly. Because some of the demand is distributed to cross-streets, the punctuated nature of demand along the route is accurately modeled.

Candidate Stop Locations

In the simplest formulation, a bus route is analyzed in a single direction. The street along which the bus operates is called the main street. Initial and terminal stops, numbered 1 and N, respectively, are given. A given set of candidate stop locations is used, normally including all the intersections along the route, numbered $2, \ldots, N-1$ in the direction toward the terminal stop. Candidate stop locations are assumed to be at intersections; midblock stops can still be accommodated in the general framework of the discrete model but require more complex formulas for market splitting and walking distance. For some of the following formulas, it is also helpful to define stops 0 and $N+1$ as artificial stops just beyond the end stops.

As mentioned earlier, if certain intermediate stop locations are fixed (e.g., transfer points), the problem decomposes into finding optimal stop locations between the fixed stop locations, yielding several smaller problems structurally identical to the original (*4*).

Stop Shed Lines

Consistent with the objective of minimizing social costs, including walking and riding times, passengers are assumed to use the stop that minimizes a weighted sum of their walking and riding times, with weights c_w, c_r equal to the values of the walking time and riding times, respectively; and v_w, v_r equal to the average walking and riding times, respectively.

This assumption provides a basis for determining the shed lines between each stop's market. As shown in Figure 1*a*, if the distance between adjacent stops is L, the shed line for boarding passengers is located a distance $(1 - r)L/2$ from the upstream stop, with r chosen so that the weighted travel time from the shed line to the downstream stop is the same whether a passenger walks to the upstream or to the downstream stop. The shed line for alighting passengers is shifted identically in the opposite direction, that is, toward the downstream stop. For boarding passengers, this is expressed as

$$(1 + r)\frac{L}{2}\left(\frac{c_w}{v_w}\right) = (1 - r)\frac{L}{2}\left(\frac{c_w}{v_w}\right) + L\left(\frac{c_r}{v_r}\right) \qquad (1)$$

leading to the shed line formula

$$r = \frac{c_r/c_w}{v_r/v_w} \qquad (2)$$

Similar results have been derived by others (*5, 9*). A typical value of r for an urban application is 0.1, which follows from a typical disutility ratio $c_r/c_w = 0.4$ and typical riding and walking speeds of 20 km/h (12 mph) and 5 km/h (3 mph), respectively. This value of r puts the stop shed line for boarding passengers at 45 percent of the distance from the upstream stop, a rather small departure from the simpler assumption that passengers walk to the nearest stop. With a smaller disutility ratio and slower riding speed, the departure from midstop shed lines is greater; for example, with a disutility ratio of 1 and an average operating speed of 15 km/h (9 mph), $r = 0.3$ and the boarding shed line is about one-third of the distance from the upstream stop.

Geographic Model for Distributing Demand

Demand is assumed to be fixed and along streets within a fixed distance w from the main street. To the extent that stop spacing decisions affect passenger travel time, they can be expected to contribute to marginal changes in demand in keeping with a general level-of-demand elasticity. However, it is reasonable to ignore this demand effect—even when ridership is an important institutional priority—on the grounds that with small service changes such as changing stop locations, the best way to attract new passengers and retain existing passengers is usually to offer the best possible service for the existing passengers. A simple way to explicitly account for demand changes is to apply a multiplier to the passenger-related aspects of the objective function. The small effect of competition with walking is also neglected.

Existing demand data for a bus route are assumed to come from available on-and-off and transfer counts. The purpose of the geographic model is to distribute the demand observed at the existing bus stops to the blocks of the main street, parallel streets, and cross-streets in each stop's service area. As alternative stop locations are examined, the distributed demand is allocated to those stops based on natural shed lines, enabling the walking distance to each stop to be determined. In effect, the geographic model is a rational means for redistributing demand from existing stops to alternative stop locations.

Ideally, the demand observed at a stop should be distributed to the blocks in its service area in proportion to "opportunities" along each block, as determined from either a detailed passenger survey or a geographic information system with data on the block, including data on the population, jobs, and retail space. Because these detailed data sources were available for the example application in Boston, a method was developed for allocating demand based on simple map information and subjective assessments of development intensity along the different streets in the service area.

The service area was assumed to be a rectilinear network with streets lying parallel and perpendicular to the main street, as illustrated in Figure 1*a*. This assumption, which may require some abstraction of the street network, captured the essential features that affect walking distance without requiring the coding of a network of access streets. For each block along the main street, parameters were assigned to indicate the relative density of opportunities for generating and attracting trips along the main street, a typical parallel street, and the cross-street.

The on-and-off demand observed at an existing stop, excluding transfer demand, was allocated to all the block faces in the stop's service area in proportion to the product of a block face's density parameter and its length in the stop's service area. Density parameters for trip generation were used to distribute boardings, and attraction density parameters were used to distribute alightings. Even if the density parameters were determined subjectively (e.g., for trip generation, 1 for detached homes, 3 for low-rise multifamily, 5 for mid-rise development), this model enabled the demand to be redistributed from existing stops to alternative stops in a consistent manner that recognized the influence of the underlying street network and development patterns.

Walking distance perpendicular to the bus line is independent of stop-spacing decisions and therefore was omitted from the formulation. All demand effectively was projected onto the main street. Demands from the main street and parallel streets became distributed demands, and demands from the cross-street, along with transfer demands, became concentrated or point demands, as illustrated in Figure 1*b*. Demand originating before Stop 1 and ending after Stop N was similarly projected to Stops 1 and N, respectively. For computer implementation, the modeled boarding and alighting distributions are most easily represented as cumulative arrays.

A special case of geographical allocation occurs when the network of cross-streets and parallel streets is regular and uninterrupted, with an average block length L_x for the cross-street and L_m for the main street and parallel streets. If the opportunity density parameters are the same for every block type, the fraction of demand allocated to the cross-streets will be $L_x/(L_x + L_m)$. For example, if cross-streets blocks are twice as long as main-street blocks, two-thirds of the demand will be allocated to the cross-streets and represented by point, rather than distributed, demands.

Walking Impacts

Passenger impacts were measured relative to a base in which passengers ride the bus without any unnecessary stopping delays from

(a)

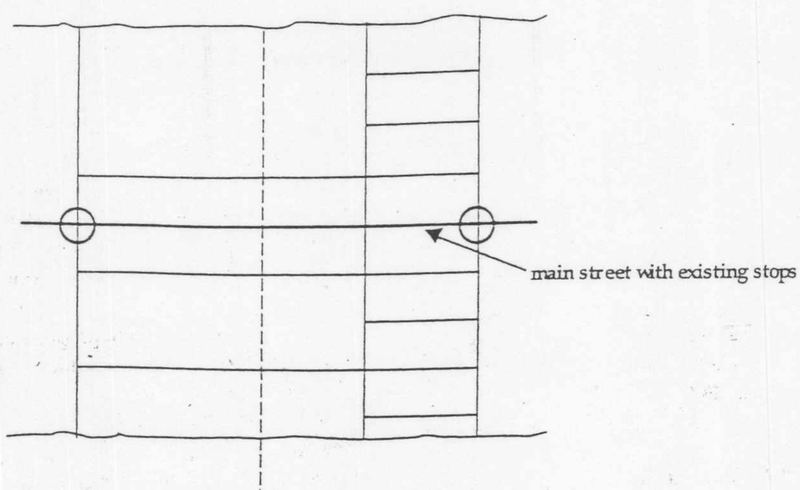

(b)

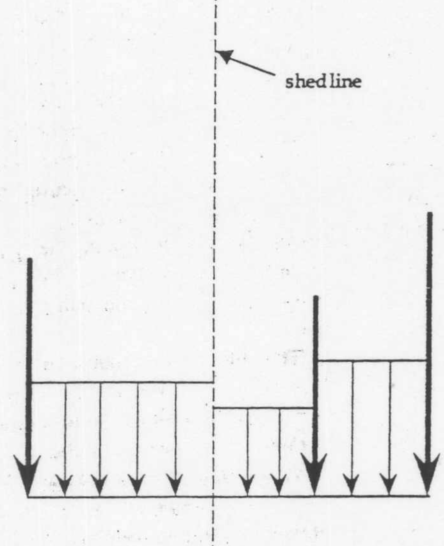

FIGURE 1 Service area assumptions: (a) idealized street network, (b) geographic demand distribution.

the point on the main street closest to their ultimate origin to the point on the main street nearest their ultimate destination. To make the optimization process efficient, all impacts were separated into contributions attributed to the individual stops.

Given three consecutive stops i, j, and k, the shed lines for stop j's service area were determined. Total passenger demand at stop j could be determined straightforwardly from the cumulative boarding (B) and alighting (A) profiles generated from the geographic modeling of the existing demand, expressed as $B(j; i, k)$, $A(j; i, k)$ equals boardings and alightings, respectively, per unit time at stop j, given that the preceding stop is stop i and the succeeding stop is stop k.

To determine the time spent walking to stop j, the portions of its service area lying upstream and downstream of the stop must be distinguished and the demand centroid of each portion must be determined for both boarding and alighting passengers. To be consistent with a passenger strategy of minimizing the weighted sum of walking and riding times, the net walking time, rather than walking time, should be tracked. Net walking time is defined as the walking time parallel to the bus route minus the change in riding time resulting from that walk, with the riding time weighted to reflect the disutility ratio between walking and riding times. If the demand centroids for the upstream and downstream portions of a stop's market are distances x_u and x_d, respectively, from the stop, it can easily be shown that, for boarding passengers, the corresponding net walking times per person are as follows:

$$\frac{x_u}{v_w}(1 - r) \text{ and } \frac{x_d}{v_w}(1 + r) \tag{3}$$

For alighting passengers, the formulas are the same, except that $(1 - r)$ and $(1 + r)$ are reversed. These net walking times, expanded by the number of boardings or alightings in the corresponding portion of the market and then summed, yield total net passenger walking time for a stop. To avoid the complexity of aggregation and centroid formulas, a simple notation is used: $W(j; i, k)$ equals the total net passenger walking time per unit time for passengers boarding and alighting at stop j, given that its neighboring stops are i and k.

Riding Time and Operating Cost Impacts

If a bus stops at location j, it will be delayed, affecting through riders and operating cost. The delay can be divided into independent and variable components, with the variable component depending on the number of passengers boarding and alighting. Let $D(j)$ equal the independent component of delay at stop j, given that a bus stops there. The independent component of delay can be further subdivided into two parts. The first part is a fixed amount $\alpha(j)$ for the opening and closing of doors and returning to traffic, which can vary between stops depending on street geometry and traffic levels. The second part is the delay incurred while decelerating and accelerating, which can also vary between stops depending on the cruise speed, defined as the average speed of a bus not stopping as it passes a specific stop, and the grade, which influences the acceleration rate. The cruise speed at a stop depends not only on the local speed limit, but also on traffic control and possible intersection queuing. Low cruise speed at signalized intersections is a factor that favors locating stops at signalized rather than unsignalized intersections. If the cruise speed at stop j is $v(j)$ and the local deceleration and acceleration rates are $a_1(j)$ and $a_2(j)$, respectively, fixed delay is given as follows (*4*):

$$D(j) = \alpha(j) + 0.5\, v(j)\left[\frac{1}{a_1(j)} + \frac{1}{a_2(j)}\right] \tag{4}$$

If the variable component of stop delay is assumed to be a linear function of boarding and alighting movements (e.g., requiring n_1 seconds per alighting passenger plus n_2 seconds per boarding passenger), its impacts are independent of stop location and therefore may be omitted from the formulation. The time needed for passengers to board and alight may change from one stop location to another, but the net impact on riding and cycle times will not change.

In standard operations, buses will stop at a bus stop only if passengers want to board or alight there. Assuming that passengers arrive in a Poisson process and that the service operates with a regular headway, the probability that a bus will stop at stop j is a function of the expected number of passenger movements at the stop (*4*):

$$P(j; i, k) = 1 - \exp\{h[B(j; i, k) + A(j; i, k)]\} \tag{5}$$

where $P(j; i, k)$ equals the probability of stopping at j, given that the straddling stops are i and k, and h equals the service headway.

Of course, if buses stop, according to transit policy, at stop j, $P(j; i, k)$ equals 1. This will typically be the case at the initial and terminal stops, and perhaps at other stops en route.

Delays due to stopping are responsible for two social costs. One cost is the time lost for through passengers. Let $Q(j; i, k)$ equal the through volume at stop j, given that the neighboring stops are i and k. $Q(j; i, k)$ equals the cumulative boardings from the start of the route to the (i, j) boarding shed line, minus cumulative alightings through the (j, k) alighting shed line.

The second social cost is a higher operating cost from increasing the cycle time and thus the amount of vehicle-hours required for the route. The increase in vehicle-hours per unit time because of a stop is the expected delay at the stop multiplied by the service frequency. The unit cost is given as c_{op}, which equals the marginal operating cost per unit time delay.

If a short-term view is taken, the argument could be made that changes in cycle time because of stop location changes are likely to be too small to affecting the operating cost, thus c_{op} equals 0. However, in the long run, changes in cycle time from changing stop locations will be combined with changes from other factors, such as traffic and demand, and thus can be expected to have a proportional influence on the necessary fleet size and operating cost. For the long run, then, it is appropriate for c_{op} to equal the marginal cost of a vehicle-hour, which may vary with the time of day.

Objective Function and Constraints

A decision must be made about which subset S is to be used from a set of candidate stops. It is helpful to define predecessor and successor functions for each stop $j \in S$ as $p(j)$ for the stop preceding j and $s(j)$ for the stop succeeding j.

The objective of the discrete model is to select the members of S to minimize

$$\begin{aligned} Z_{\text{discrete}} &= \sum_{j\in S}\Big\{c_w W[j; p(j), s(j)] + c_r Q[j; p(j), s(j)] \\ &\quad \times P[j; p(j), s(j)]\, D(j) + \frac{c_{op}}{h} P[j; p(j), s(j)]\, D(j)\Big\} \\ &= \sum_{j\in S}\{Z_w[j; p(j), s(j)] + Z_r[j; p(j), s(j)] \\ &\quad + Z_{op}[j; p(j), s(j)]\} \end{aligned} \tag{6}$$

where

Z_w = net walking time cost per unit time,
Z_r = riding delay cost per unit time, and
Z_{op} = operating cost per unit time (all functions attributed to stop j).

The chief constraint on S membership is a maximum-allowed stop spacing. A general way of expressing this constraint, which allows for different stop-spacing policies on different portions of the route and for exceptions because of topographical features such as rivers and parks, is to specify the most distant neighboring stop locations allowed for each potential stop location: $s_{\text{allow}}(j)$ equals the most distant stop that may succeed stop j, and $p_{\text{allow}}(j)$ equals the most distant stop that may precede stop j.

DYNAMIC PROGRAMMING ALGORITHM

Because the discrete optimization problem can be separated by stop, it lends itself well to a solution using dynamic programming (*10*). The stage variable is j, a candidate stop location. At stage j, the objective is to find the least-cost solution for serving the route from j to its end (i.e., to stop N). The stage variable is the index of the preceding stop, the only information needed about decisions made upstream of stop j to optimize the remainder of the route. The optimal return function is $f(j; i)$, which equals the minimum cost of serving stop j to stop N, given that the stop preceding j is i.

The optimization progresses backward, beginning at stage N with

$$\begin{aligned} f(N; i) = \{&Z_w(N; i, N+1) + Z_r(N; i, N+1) \\ &+ Z_{op}(N; i, N+1)\} \quad \text{for } i = p_{\text{allow}}(N), \ldots, N-1 \end{aligned} \tag{7}$$

Stages $j = N - 1$ to $j = 1$ are processed in order, with the recursion at each stage being

$$\begin{aligned} f(j; i) = \min_{k=j+1,\ldots,s_{\text{allow}}(j)} \{&Z_w(j; i, k) + Z_r(j; i, k) + Z_{op}(j; i, k) \\ &+ f(k; j)\} \quad \text{for } i = p_{\text{allow}}(j), \ldots, j-1 \end{aligned} \tag{8}$$

The optimization process ends when $f(1; 0)$ is determined, which is the minimum cost of serving the entire route. The solution can then be derived by tracing back through the optimal decisions made at each application of Equation 8.

The dynamic programming algorithm is computationally efficient. With 100 candidate stop locations, for example, the number of alternative sets of stops is too large for exhaustive enumeration. In contrast, the dynamic programming algorithm requires on the order of NR^2 calculations, where R equals the average range, in the number of candidate stop locations, between a stop and the farthest allowable adjacent stop. Because the number of calculations is linear in N, the algorithm places no practical limit on route length. Storage requirements are also very modest.

CORRESPONDING CONTINUOUS FORMULATION

To compare the continuum and discrete approach solutions, a continuum model was developed using logic corresponding to the discrete model. In most respects, the continuum model follows that of Wirasinghe and Ghoneim (*4*). Demand density at any point x was determined by using the geographic model to find the total demand rate in an 800-m (0.05-mi) interval centered on x, excluding transfer volumes, and dividing by 800 m. In determining the impacts on walking, passengers were assumed to minimize net walking time, consistent with the discrete model and with the objective. With this assumption, and with the evenly distributed demand of the continuum model, the ratio of average net walking time to the walking time between neighboring stops is $(1 - r^2)/4$, which only slightly differs from the ratio of 1:4 that exists when passengers minimize walking time only. Assuming that buses stop only if passengers are waiting, the optimal solution could be obtained by a one-dimensional search.

APPLICATION AND RESULTS

Both the discrete and continuous models were applied to MBTA Route 39, one of Boston's busiest bus routes, running from Forest Hills to Back Bay Station. Route 39 is 7.2 km (4.5 mi) long and can be divided into three sections: a 3.3-km (2-mi) section from Forest Hills to Heath Street in the dense residential neighborhood of Jamaica Plain; a 2.6-km (1.6-mi) section, ending at Northeastern University (NU), through the even denser Mission Hill and East Fenway areas, with residences and large employment centers, including several hospitals and universities; and a 1.3-km (0.8-mi) section in the Back Bay, an extension of Boston's CBD. A streetcar line parallels Route 39 along the Heath-NU section of the route. The only major transfer point is Forest Hills, although smaller transfer demands exist at several other points along the route. The inbound direction was modeled during the morning peak period, with a headway of 3 minutes.

Demand was determined from historic on-and-off counts. Because the MBTA does not count transfers routinely, transfer volumes were estimated subjectively. The service area was assumed to extend 400 m (0.25 mi) on either side of the line. Every intersection on the main street was treated as a potential stop. For each main-street block, the number of parallel streets in the service area was determined from maps. Opportunity density parameters were assigned subjectively on the basis of development density. Cruise speed was set at 48 km/h (30 mph) at unsignalized intersections and 24 km/h (15 mph) at signalized intersections.

Values of walking time, riding time, and marginal operating cost were set at \$10/h, \$4/h, and \$80/h, respectively. Bus operating speed and walk speed were 20 km/h (12 mph) and 5 km/h (3 mph), respectively. Bus deceleration and acceleration rates of 1.33 m/s^2 (3 mph/s) and a constant lost time of 9 s were assumed for every stop. The maximum allowed interstop spacing was 530 m (0.33 mi).

The dynamic programming procedure was implemented in a Microsoft Excel workbook using Visual Basic. In this approach, the user enters parameters in several worksheets and the program calculates the optimal solution or, at the user's option, evaluates a given solution. Graphical output includes the stop location–stop density diagram shown in Figure 2 and the demand distribution diagrams shown in Figure 1*b*, each drawn to scale.

Figure 2 shows the optimal stop locations versus existing stop locations. Many existing stops are eliminated, and some are shifted to locations that do not have a stop. Although the existing route has 37 stops (average spacing equals 202 m, about 8 stops/mi), the optimal solution has 19 stops (average spacing equals 404 m, about 4 stops/mi). The striking difference between the existing and optimal solutions suggests that eliminating bus stops is the appropriate direction for U.S. transit agencies.

A numerical comparison between the optimal solution, also called the base case, and the existing situation is found in Table 1. Compared

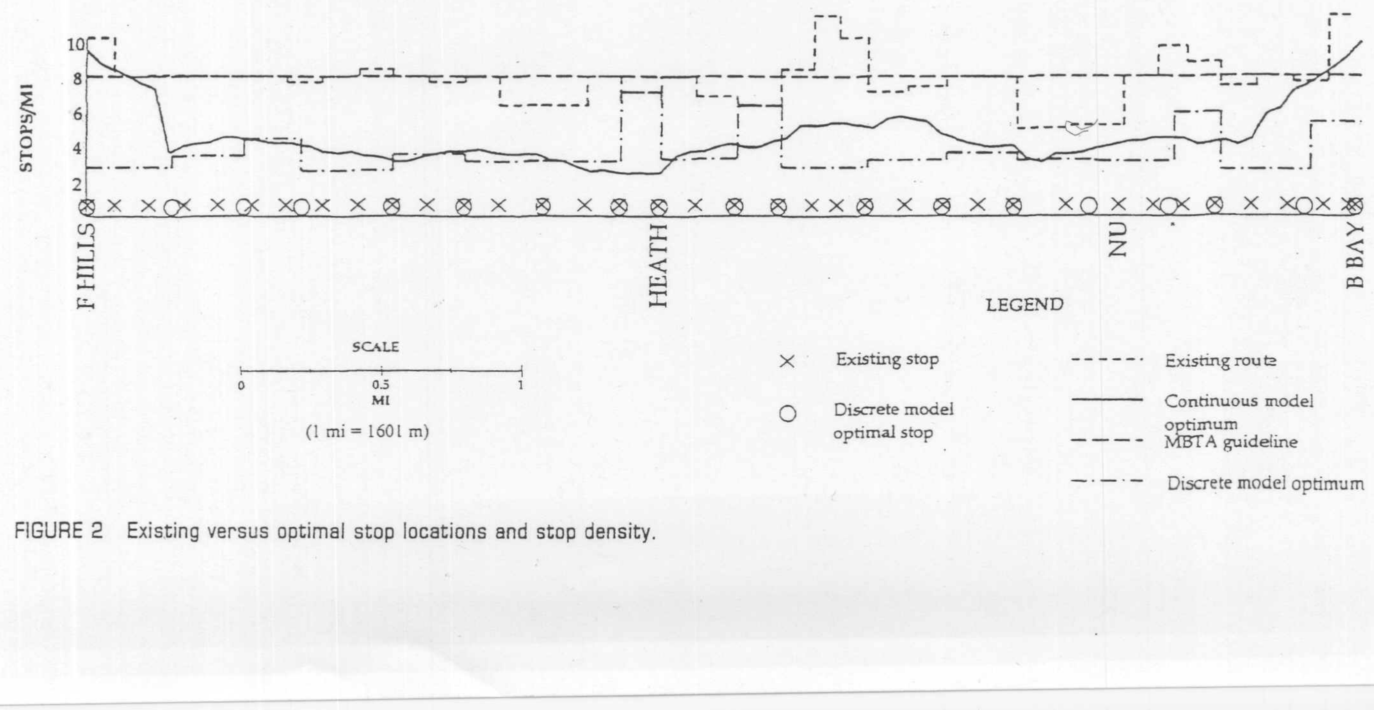

FIGURE 2 Existing versus optimal stop locations and stop density.

TABLE 1 Summary of Results

Scenario	Number of Stops	Average Spacing	Total Cost	Average Walking Time	Average Riding Delay	Incremental Running Time
		(m)	($/hr)	(min)	(min)	(min)
Existing situation	37	202	734	0.77	3.58	9.06
Base case optimum	19	404	602	1.41	1.77	4.85
No point demand	20	382	680	1.69	1.84	5.16
Zero unit operating cost	28	269	441	0.90	2.68	7.18
No walk time premium	19	404	429	1.64	1.74	4.76
Delete one stop	36	208	726	0.80	3.48	8.85

with the existing situation, the optimal solution saves $132/h in total social costs, with passengers walking 0.6 min more and riding 1.8 min less on average, and each bus saving 4.3 min of operating time per trip.

The optimal solution is strongly constrained by the maximum allowed stop spacing. Ten of the 18 segments reach the most distant intersection permissible. More sample applications are needed to determine whether or not this is the general situation.

Figure 2 also compares the optimal stop density as determined by the discrete model, continuum model, existing situation, and the MBTA's guidelines. Optimal stop density was calculated for the continuum model at sample points every 80 m (0.05 mi). The density for the existing situation and the optimal discrete solution is simply the inverse of the segment length, changing for each segment.

The discrete solution optimum has the same general magnitude as the continuum optimum, but does not track the latter very closely. The center of the Heath-NU segment highlights the difference between the discrete and continuum models. In response to high demand in this region, the continuum model calls for an increase in stop density, while the discrete model does not. The discrete model recognizes the large point demands coming from cross-streets with hospitals and other major generators, and although it locates stops at the major cross-streets, it still allows considerable distances between stops.

SENSITIVITY ANALYSIS

Several variations of the base case and existing solution were also run to indicate the sensitivity of the solution to various parameters. Table 1 summarizes these runs. The first sensitivity run explored the effect of the geographic model by setting the opportunity density on the cross-streets to 0 and making the opportunity densities on all main-street blocks equal. The result is a small change from the base case—one stop is added and two stops are relocated. The difference would probably be greater if the solution were not so tightly constrained by the maximum allowed spacing.

A substantially different solution results if the marginal operating cost is set to 0, effectively minimizing only passenger walking time (weighted) and riding time. The optimal number of stops rises to 28, with an average spacing of 269 m (six stops/mi). However, removing the walking time premium (i.e., walking time valued the same as riding time) resulted in almost no change in the optimal solution, which is to be expected, because the base case solution is already strongly constrained by the maximum allowed spacing. The only difference is a slight relocation of two stops from unsignalized intersections to signalized intersections, with the marginal cost of stopping being smaller and resulting in a slightly longer walking distance and shorter riding time.

The final sensitivity run reduced demand on a 1.6-km (1-mi) segment after Forest Hills to 0.5 percent of its original value. The displaced boardings were replaced as transfer volume at Forest Hills, and the displaced alightings were replaced as transfer volume at the stop following the test segment. In the optimal solution, two stops were added to the middle of the segment with little demand, resulting in stops at three consecutive intersections with interstop spacings of only 64 m and 200 m. This result demonstrates the phenomenon—shared by the discrete and continuum models—that as demand becomes small, the optimal solution is to locate stops as frequently as possible. The reason for this conclusion is that by spacing stops far apart (i.e., collecting demand at discrete points), losses associated with stopping are distributed over more passengers. At low-demand levels, however, a bus will rarely serve more than one boarding or alighting passenger per stop, even if stops are widely spaced, so the bus might as well stop wherever a passenger is waiting.

EVALUATION OF GIVEN STOP LOCATIONS

To demonstrate its usefulness in evaluating stop location decisions, the model was used to evaluate the elimination of one stop (at 590 Huntington Avenue in the Heath-NU segment). As indicated in Table 1, this change resulted in a societal benefit of $8/h compared with the existing solution, $5.60 of which was in operating cost savings. Extending this change over 1 year easily might represent an annual societal benefit of $16,000 (with an operating cost savings of $11,000), strongly justifying removal of the stop.

MODEL EXTENSIONS

As indicated by Wirasinghe and Ghoneim, stop spacing should be optimized over the entire day, not one time period, because changing stop locations for specific times of the day is impractical (*4*). Also, symmetry in stop locations for both directions is probably beneficial to help orient passengers. Optimizing spacing for all periods of the day requires only a direct extension of the discrete model, in which costs in the optimal return function are simply summed over all periods. Periods may be defined as narrowly as desired, to the point of treating each trip as a period. Optimizing over both directions simply requires a choice set consisting of pairs of stops, one for each direction, and that costs in the optimal return function be summed over both directions, with one direction processed forward and the other backward. Variable definitions and formulas will require some adjustment but will not complicate the solution. The optimal stop locations for the entire day are not likely to vary much from a period-specific optimum because the optimal solution varies little with respect to a scaling of demand, if the service frequency is similarly scaled. In practice, service frequency is related to demand, although when shifting from a peak to an off-peak period, service frequency is typically scaled down less than demand. As a result, off-peak vehicle loads tend to be smaller, which makes the optimal stop spacing greater. However, the practice of scheduling for lower average loads in the off-peak period partly results from the smaller marginal cost of a bus hour during the off-peak hours, which, if properly considered, would tend to reduce optimal stop spacing. The net effect of these opposing tendencies will most likely be little or no change.

CONCLUSIONS

A discrete approach to modeling bus stops, passenger demand, and operations is a practical and highly efficient method for evaluating and making the best stop-location decisions. This approach is superior to a continuum model in most ways.

The results of the example application support stop spacing in busy urban corridors that is closer to the standard European value of 320 to 400 m (4 to 5 stops/mi) than to the standard U.S. value of 160 to 230 m (7 to 10 stops/mi). However, the model needs further testing before general conclusions can be reached. The model also needs to be refined for practical applications, including the consideration of each travel direction, multiple periods, and calibration of some model constants.

ACKNOWLEDGMENTS

The authors thank the Central Transportation Planning staff for assisting in data collection for the case study.

REFERENCES

1. Benn, H. P. *TCRP Synthesis of Transit Practice 10: Bus Route Evaluation Standards.* TRB, National Research Council, Washington, D.C., Jan. 1995.
2. Reilly, J. M. Transit Service Design and Operation Practices in Western European Countries. In *Transportation Research Record 1604*, TRB, National Research Council, Washington, D.C., 1997, pp. 3–8.
3. MacDorman & Associates, and N. H. Wilson. *Design of Service Quality Measures and Service Evaluation Standards.* Final Report to the Massachusetts Bay Transportation Authority and Conservation Law Foundation, Boston, 1995.
4. Wirasinghe, S. C., and N. S. Ghoneim. Spacing of Bus Stops for Many-to-Many Travel Demand. *Transportation Science*, No. 15, 1981, pp. 210–221.
5. Vuchic, V. R., and G. F. Newell. Rapid Transit Interstation Spacings for Minimum Travel Time. *Transportation Science,* No. 2, 1968, pp. 303–339.
6. Vaughn, R. J., and E. A. Cousins. Optimum Location of Bus Stops on a Bus Route. *Proc., 7th International Symposium on Transportation and Traffic Theory* (T. Sasaki and T. Yamaoka, eds.), 1977, pp. 698–716.
7. Lesley, L. J. S. Optimum Bus-Stop Spacing: Parts 1 and 2. *Traffic Engineering and Control,* 17, 1976, pp. 399–401, 472–475.
8. Hakimi, S. L. Optimum Locations of Switching Centers and the Absolute Centers and Medians of a Graph. *Operations Research*, 12, 1964, pp. 450–459.
9. Hurdle V. F., and S. C. Wirasinghe. Location of Rail Stations for Many-to-One Travel Demand and Several Feeder Modes. *Journal of Advanced Transportation*, 14, 1980, pp. 29–46.
10. Bellman R. E., and S. E. Dreyfus. *Applied Dynamic Programming.* Princeton University Press, Princeton, N.J., 1962.

Publication of this paper sponsored by Committee on Bus Transit Systems.