Maaza C. Mekuria (corresponding author)
Axum Design and Engineering Corp
255 North market Street, Suite 170
San Jose, CA 95110
E-mail: mmekuria@axumcorp.com
Telephone: 617.697-1563, fax : 360-363-7034

Peter G. Furth
Department of Civil & Environmental Engineering
Northeastern University
Boston, MA 02115
E-mail: pfurth@coe.neu.edu
Telephone: 617.373.2447; fax: 617.373.4419

Melachrinoudis, Emanuel
Department of Mechanical & Industrial Engineering
Northeastern University
Boston, MA 02115
E-mail: emelas@coe.neu.edu
Telephone: 617.373.2447; fax: 617.373.4419

A discrete model of bus stop location, in which candidate stops are either selected or not, has several advantages over classical continuum models related to both practicality and realism. Using parcels as units of demand and using the street network to model walking paths between transit stops and parcels has proven both effective and realistic for evaluating different stop sets. In this framework, demand is allocated from on-off counts made at existing stops to the parcels in each stop's service area in proportion to their trip generating ability, resulting in a demand distribution that both matches existing counts and reflects variations in land use. However, with demand modeled on the street network, the simple construct of placing a service boundaries midway between neighboring stops becomes invalid because of irregularities in the access street network and curves in the transit route. The dependence of a stop on more than its immediate neighbors for determining its service area complicates the process of optimizing stop locations using dynamic programming. The proposed solution is to expand the state space, allowing a stop's service area to be dependent on two prior and two succeeding stops. The resulting dynamic programming model is tested on two Boston area bus routes, where it finds solutions that are better than both the existing stop set and better than stop sets proposed by consultants using the simplistic yet state of the art models available.

A common complaint about bus service is that it’s too slow because buses make too many stops. An industry rule of thumb is that about 30% of running time on a typical bus route is lost at stops. While time used by passengers to board and alight is inevitable, other losses – those due to deceleration, acceleration, and opening and closing doors – roughly 8 to 15 s per stop – could be reduced by consolidating stops. Reilly (1) noted the greater stop spacing found on European transit routes, and the correspondingly speedier service. In the last few years, several U.S. transit agencies have applied or considered stop consolidation as a means of making their services more efficient and competitive.

Until now, methods for making and analyzing stop spacing decisions have been extremely simplistic. This paper describes a method for optimizing stop location on an existing route that includes realistic and localized estimates of its walking, riding, and operating cost impacts. Application to two Boston area routes indicate that it finds a better solution than both the existing stop set and that derived by experts using simplistic yet state of the art methods.

Stop location decisions on an existing route (which stops to keep, drop, or insert) have three main impacts: walking time, riding time, and operating cost. As stops become further apart, walking time increases, while riding time and operating cost decrease. These countervailing impacts create a tradeoff whose optimum depends on the relative weights given to each impact.

A fourth possible impact that our analysis ignores is that demand may change. Demand response can be divided into two kinds. One is an elastic response to service becoming slightly slower or faster, and to walking distances slightly longer or shorter. While accounting for this effect will amplify impacts, it is unlikely to change the optimal decision because, roughly speaking, the best way to attract new passengers is to offer good service to existing passengers. The second type of demand response is losing riders if stops become spaced so far apart that some riders consider the stop no longer accessible. This kind of demand change can be very important if long distances between stops are considered. For this analysis, we assume that transit agencies are under a mandate not to space stops so far apart that they become inaccessible; with this constraint accounted for by means of a maximum stop spacing, this second type of demand response can also be ignored.

A general formulation for finding the optimal spacing between stops can be derived using a continuum model; good treatments are given by Wirasinghe and Ghoneim (2) and by Van Nes and Bovy (3). The street network is assumed to be rectilinear and infinitely dense, so that walking paths to a stop can be divided into a component transverse to the transit line (which can be ignored) and a component parallel to the transit line. This assumption in effect places all demand along a single line. In addition, demand density is assumed to vary slowly along a line, so that one can meaningfully speak of the demand density in the neighborhood of a point. In their model as in ours, the service area of a transit line is assumed to be an area within a fixed distance (e.g., 0.25 miles) of the transit line.

With demand collapsed to a single dimension, the service area boundary for a stop becomes simply the midpoint to neighboring stops, adjusted to account for the fact that the cost-minimizing passenger will walk a bit further downstream than upstream in order to take advantage of riding time savings (4, 5). As a result, service area boundaries shift upstream for boarding passengers and downstream for alighting passengers. Therefore, a stop's service area boundaries for boarding passengers traveling in one direction will differ from the boundaries used by alighting passengers in the same direction, and from boarding passengers traveling in the opposite direction.

With continuum approach, optimal stop spacing can be derived using calculus. While useful for establishing a general guideline for stop spacing, this approach has three main weaknesses. First, demand is often not slowly varying. There can be significant punctuations in demand due to differing land uses (compare, for example, a hospital versus a cemetery), and stop locations should be sensitive to where demand originates. Second, the street network is often not ideal; turns on the transit route and irregularities in the street network influence people's walking paths and can make certain stop locations more favorable. Third, the result of such an optimization is hard to apply, because stops should normally be located at intersections (both for shorter walk access, and for proximity to crosswalks), and the "optimal" stop spacing may not be a convenient multiple of block length. For example, suppose intersections are 200 m apart and the optimal stop spacing is 300 m; should stops be located at every stop, every other stop, or something in between?

Furth and Rahbee (5) were the first to propose a discrete model for optimizing stop spacing for an existing route. It treats each intersection along the routes as a candidate stop, and then selects the optimal set of stops from this list. To model punctuations in demand density, they allowed an analyst to assign different levels of density (1, 2, 3, etc.) to each intersecting street, and to the longitudinal streets between each intersection. For the existing set of stops, demand at a stop (number of ons and offs per hour) is reflected back to the block faces in the stop's service area by allocating it in proportion to a block's "strength" – its length multiplied by its demand density. When considering new combinations of stop locations, the number of ons and offs at a stop can simply be determined by aggregating over the blocks in its service area. Walk distance can likewise be modeled as a sum of transverse and longitudinal distance from the center of each blockface.

In the discrete model, the riding time impact of a candidate stop is a fixed loss accounting for deceleration, acceleration, and opening a closing doors, multiplied by the probability of stopping, which is based on a Poisson model of passenger arrivals. Both arriving and departing passengers contribute to the probability of a stop being requested; passenger demand in turn is based on the demand in the stop's service area multiplied by one headway. The discrete modeling framework is not designed to make fine-level decisions about where at an intersection to place a stop (e.g., near-side or far-side); analysts are expected to make those choice based on local factors. One paper has analyzed effects of stop placement to recommend varying values of time lost due to stopping (6).

Furth and Rahbee's model still makes significant approximations by assuming a regular, rectilinear grid. This conveniently reduces the demand profile to a single dimension, equivalent to a combination of point demands and uniformly distributed demands along a line, like loads on a structural beam.

With demand thus reduced to a single dimension, the optimal stop location problem can be solved easily using dynamic programming (DP). Impacts are additive, and a stop's service area depends only on which candidate stops are selected to be its *j* as a candidate stop to be included, the impact of the decision of which should be the next stop after *j* is independent of what stops were chosen in the upstream of the starting point stop, except for the previous stop, since that decision affects probability that a bus would have to stop at the starting point stop.

Let stops be numbered consecutively from the start to the end of the line, and consider three successive stops *i, j,* and *, k*. With demand along a line, this (*i,j,k*) triplet defines the service area of stop *j*, and allows one to define stop-specific demand

ons (j; i, k) and *offs(j; i, k)* = demands (ons and offs per hour) at stop *j* when its predecessor is *i* and its successor is *j*

All of the impacts associated with stop *j* can then be determined using the defined model of walking routes, probability of stopping, and delay due to stopping. These impacts are:

walk (j; i, k) = walking cost for all passengers walking to or from stop *j* when its predecessor is *i* and its successor is *j*

ride (j; i, k) = riding cost for all passengers between *j* and *k* when stop *j*'s predecessor is *i* and its successor is *j*

oper (j; i, k) = operating cost between *j* and *k* when stop *j*'s predecessor is *i* and its successor is *j*

These cost functions include unit costs applied to walk time, riding time, and running time. The riding and operating cost functions include dwell time and acceleration delay at *j* and deceleration delay at *k*.

The backward recursive dynamic programming formulation is

$$f(j;i) = \min_{\min S(j) \le k \le \max S(j)} \{walk\ (j:i,k) + ride\ (j;i,k) + oper\ (j;i,k) + f(j;k)\}, \quad for \quad \min P(j) \le i \le \max P(j) \qquad (1)$$

where *f*(*j;i*), the optimal return function, is the cost of serving the route from stop *j* to the end, given that the stop before *j* is *i*. A maximum and minimum permitted stop spacing is accounted for by defining *minP(j)* and *maxP(j)* as the further and closest predecessor of each stop *j*, and *minS(j)* and *maxS(j)* as the closest and further successor of each stop *j*. In the dynamic programming algorithm, the stage variable is *j*, and there is only one state variable, *i*. The algorithm begins with *j* as the final stop, initializing *f*(*j;i*) at 0 for all legal values of *i*; then it proceeds backwards, reducing *j* by 1 and using equation 1 to solve for *f*(*j;i*) at 0 for all legal values of *i*. The algorithm thus proceeds until *j* = 1; letting stop 0 be a dummy "start" stop, the value of *f*(*1;0*), again found using equation 1, is the optimal solution for the route.

Unlike continuum modeling, discrete modeling doesn't give a formula for an optimal stop spacing; it gives a list of optimal stop locations, selected from a set of candidate stop locations. That makes the results directly applicable. However, while Furth & Rahbee proved the value of discrete modeling, the subjective assignment of "demand intensity level" required made it impractical, and its assumption of a regular, rectilinear network made it too idealistic for many situations.

In two previously published papers (7, 8), we developed the first model of demand for the stop location problem that went beyond a single dimension. The unit of demand is land parcels, which are located on the real street network. Passengers use shortest paths along the street network between their parcel and their stop. Existing geographic information system (GIS) databases of parcels, streets, and transit stops make such an analysis feasible. For consistency, passengers are assumed to minimize their combined travel time (walking time plus riding time from the stop they walk to, with the same weights used in the global optimization), creating the same offset in service area boundaries mentioned earlier.

With demand originating at parcels on the street network, service area boundaries for each stop are no longer defined by 1-dimensional constructs such as the perpendicular bisector of a line segment. Service areas are not modeled explicitly; rather, they simply fall out as an outcome of each passenger selecting the stop that minimizes their weighted walk plus ride time from the stop they walk to, in a process that can be considered a Voronoi diagram on a network. Walking distance is directly measured along the street network from each parcel to each stop.

To assign demand to a parcel, service areas for the historic stop set are first determined using the shortest path / Voronoi process. Then, for each historic stop, historic boarding demand (an input) is allocated over all parcels in the service area in proportion to a parcel's production strength. A parcel's production strength is the product of its size variable (e.g., gross floor area) and a production coefficient reflecting the parcel's land use (e.g., single family residential, multifamily residential, commercial). Production coefficients were based on production factors found in the ITE handbook *Trip Generation*. Likewise, historic alighting demand is allocated over parcels in proportion to each parcel's attraction strength. Additional factors account for diminishing demand as one gets further from a stop, and for competition with other transit routes.

By using allocation logic rather than directly estimating demand for each parcel, demand along the route is completely consistent with existing counts. However, by using parcel attributes (land use type and size) to allocate demand, its spatial distribution within each stop's service area reflects the variations in land use within the service areas of historic stops. By relating demand to parcel attributes, candidate stop locations closer to parcels with higher propensity for transit use are "rewarded" relative to locations further from likely demand concentrations.

As we pointed out in (8), when demand points are distributed spatially and passengers use the real street network to access transit stops, a stop's service area can be affected by more than the location of its neighboring stops. Real street networks are often have diagonal or curving streets and discontinuities (streets that do not continue uninterrupted within the route's influence area) and transit lines often include curves and turns, which together make it that a stop's service area can have boundaries with more than one upstream or downstream stop. Figure 1 shows a service area analysis for a transit stop in which one stop, stop 11, has inbound service boundaries with not only stops 10 and 12, but also with stops 13 and 14, and even a small boundary with stop 7.

This finding, which we call the "curve effect," makes a model in which demand is distributed on the real street network violate a key assumption of prior optimization models, including Furth & Rahbee as well as the continuum models. That is the assumption that demand at a stop – and therefore the impacts associated with a stop – depend only on the location of its neighboring stops. Due to the curve effect, it is impossible to define impacts based on (i,j,k) triplets. The demand at stop j – and therefore the walking cost and other impacts associated with j – also depend on which stops precede i and which follow k.

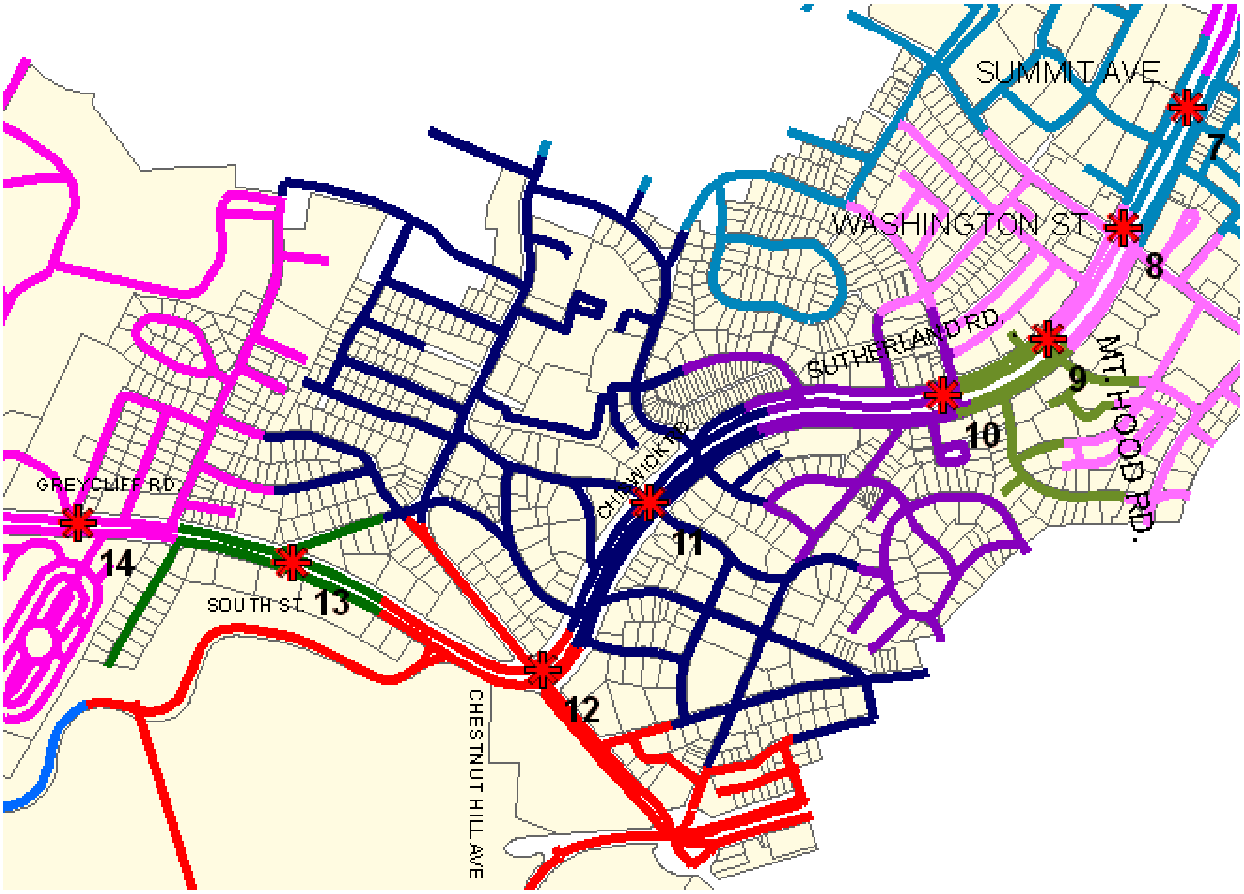

(source: (7))

The curve effect is accounted for when using the discrete model to do *evaluation*, that is, when evaluating the costs associated with a particular stop set, as shown by examples in (5) and (6) involving a streetcar line in Boston and a bus route in the Albany area. However, it was not accounted for by dynamic programming using equation 1, which will give spurious results depending on which stops before *i* and after *k* were assumed when evaluating the impact for a triplet (*i,j,k*).

It was a goal of this research to develop an optimization method using the very realistic model of parcel-level demand located on the street network. Ultimately, the chosen approach to overcome the curve effect was to expand the dynamic programming state space to three variables by assuming that a stop's service area will be defined by its previous two stops and its following two stops, creating the paradigm of a quintuplet (*i,j,k,l,m*) for which impacts associated with stop *k* can be determined. This is an assumption that can be violated in theory, but is unlikely to be violated in most practical problems. Its violation would mean, for example, that stop 10 has a service area boundary not only with stop 9 and stop 8, but also with stop 7 – that people along that last boundary choose either stop 7 or stop 10, skipping stops 8 and 9.

The formulation based on quintuplets has the following recursion formula:

$$f(k;i,\ j,l) = \min_{\min S(k)\le m\le \max S(k)} \{walk(k;i,\ j,l,m) + ride(k;i,\ j,l,m) + oper(k;i,\ j,l,m) + f(l;\ j,k,m)\} \quad (2)$$

where all of the variables are as defined previously, with *k* as the stop whose impacts are being evaluated when the previous two stops are *i* and *j* and the succeeding two stops are *l* and *m*.

To check whether the assumption of dependence on stops outside the quintuplet is violated, one can track the demand (ons and offs) associated with each quintuplet in the optimal solution, and check it against the given total demand. A violation of the assumption would mean that some parcels have been double counted, or not counted at all in the solution; and given that the method seeks to minimize cost and that walking cost will be underestimated if some passengers are not accounted for, it is likely to lean toward solutions in which some slivers of demand are unaccounted for, if indeed such solutions exist. Boston, unlike most American cities, has a street network that is not at all rectilinear or regular; yet in the applications we have done on Boston bus and streetcar routes, there have either been no violations, or – as in the case of Figure 1 – the amount of demand unaccounted for was small enough (no more than 1%) that it would not affect the optimal solution appreciably, and would be readily detectable.

The state space expansion greatly increased the algorithm's computational burden. Due to space limitations, those computational issues and their solutions will not be discussed. In the end, the algorithms that were coded are very efficient, solving the problems reported in the examples section in less than 10 seconds on a standard desktop PC.

As a practical matter, transit demand changes across the day, as do running times; however, it is impractical for stops to vary across the day. Multiple periods can easily be accounted for by summing impacts over all periods in equation 2. That is, for example, a variable such as *walk(k; i,j,l,m)* should be a sum of walking cost impacts associated with stop *k* over all the periods of the day for a stop set including the quintuplet *(i,j,k,l,m)*.

If it is permissible to have different stop sets in the two directions of a route, each direction can be optimized separately. However, if policy is such that a route should have the same stops in each direction, then each of the terms within equation 2 should include a sum over both directions.

The Massachusetts Bay Transportation Authority (MBTA) has shown an interest in stop consolidation on its bus and streetcar routes for several years. Recently, it has embarked on a "Key Routes" program to improve the service quality and image of its 15 most heavily used routes. Part of the Key Routes program is considering stop consolidation as one of several ways to increase service speed and reliability. Routes 1 and 57 were recently studied as part of this program.

Route 1 is a crosstown route in Boston and Cambridge that runs between two major transfer points (Harvard Square and Dudley Square), crosses three rapid transit lines, and serves such destinations as the Prudential Center, Boston Medical Center, and numerous colleges and universities. Figure 2 depicts the three essential geographic datasets (street network, parcels represented by dots, and bus stops). The service area was assumed to extend 0.25 miles to either side of the route.

Available demand data included on and off volumes by direction in the a.m. and p.m. peak hour. For the purpose of this example, a representative "day" is treated as 5 hours of a.m. peak service and five hours of p.m. peak service. The two directions were optimized separately, consistent with MBTA policy that allows differing stops by direction.

Parameters used in the cost function were as follow:

walk speed = 1.2 m/s (4 ft/s)

unit walk cost = $12/pax-h

unit ride cost = $6/pax-h

unit operating cost = $143/veh-h

unit boarding time = 2 s/pax

unit alighting time = 2 s/pax

lost time per stop = 8.5 s

headway = 8.75 min

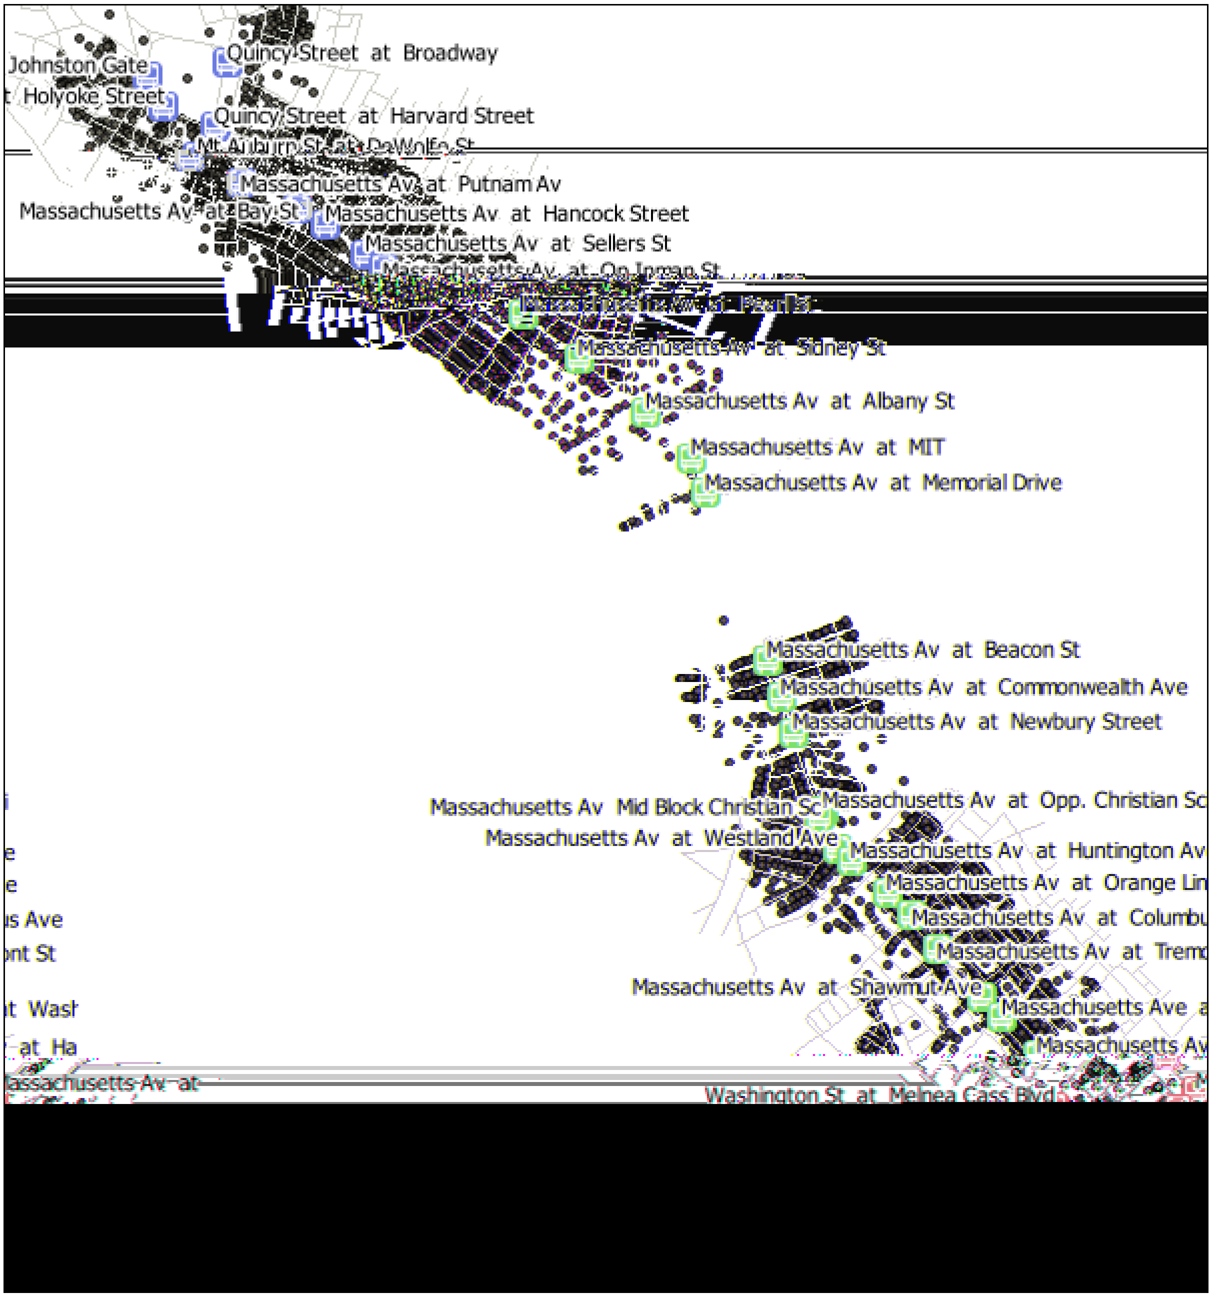
Johnston Gate
Holyoke Street
Quincy Street at Broadway
Quincy Street at Harvard Street
Massachusetts Av at Putnam Av
Massachusetts Av at Bay St
Massachusetts Av at Hancock Street
Massachusetts Av at Sellers St
Massachusetts Av at Sidney St
Massachusetts Av at Albany St
Massachusetts Av at MIT
Massachusetts Av at Memorial Drive
Massachusetts Av at Beacon St
Massachusetts Av at Commonwealth Ave
Massachusetts Av at Newbury Street
Massachusetts Av Mid Block Christian Sc
Massachusetts Av at Westland Ave
Massachusetts Av at Huntington Av
Massachusetts Av at Shawmut Ave
Washington St at Melnea Cass Blvd

Parcel data was obtained from the city assessor's offices. Relevant land uses and their production and attraction coefficients for the a.m. period are shown in Table 1. For the p.m. period, production and attraction coefficients are reversed.

Description	Code	measure of size	attraction coefficient	production coefficient
Residential / commercial	RC	Living_area	0.05226	0.08774
Single family				

	Distance from start	Departing volume	Ons + Offs (pax/h)			Avg walk time (min)			Total Cost ($/Pd)		
Stop	(mi)	(pax/h)	Historic (H)	Recom-mended (R)	Optimal (O)	H	R	O	H	R	O
1	0	41	41	41.0	58.6	3.3	3.3	2.4	$97	$97	$190
2	0.08	56	15	15.0		4.9	4.9		$147	$147	
3	0.4	63	9	9.0	3.2	2.0	2.0	3.0	$149	$149	$234
4	0.58	68	5	5.0		2.8	2.8		$95	$95	
5	0.7	73	5	5.0		3.0	3.0		$90	$90	
6	0.86	84	13	13.0	30.4	4.9	4.9	5.6	$134	$133	$293
7	1.03	94	14	14.0		3.9	4.0		$114	$115	
8	1.12	100	8	8.4	18.3	2.1	2.2	3.4	$87	$107	$182
9	1.26	101	5			2.3			$75		
10	1.32	105	12	16.6	16.6	3.5	3.7	3.7	$103	$158	$158
11	1.52	148	75	75.0	75.0	2.4	2.4	2.4	$238	$239	$237
12	1.72	161	23	23.0	25.5	3.3	3.4	3.5	$212	$214	$280
13	1.96	164	13	13.0		5.7	5.7		$204	$204	
14	2.14	173	39	41.0	51.5	2.8	2.7	3.5	$184	$384	$507
15	2.24	175	2			0.8			$249		
16	2.73	162	17	22.0	22.0	2.5	2.9	2.9	$272	$355	$353
17	2.84	154	12			2.5			$92		
18	2.94	123	81	89.1	88.0	3.6	3.5	3.4	$286	$343	$334
19	3.18	120	15		15.0	2.2		2.2	$110		$140
20	3.2	120		32.7			4.1			$208	
21	3.29	113	29		29.0	2.7		2.7	$97		$104
22	3.34	116	21	31.1	21.0	1.6	2.4	1.6	$81	$133	$81
23	3.47	133	47	47.0	47.0	2.5	2.5	2.5	$149	$149	$149
24	3.57	132	7	7.0	7.0	3.2	3.2	3.2	$94	$94	$93
25	3.68	125	11	14.5	21.7	3.6	4.0	4.3	$132	$167	$217
26	3.86	118	9			3.5			$113		
27	3.94	110	16	21.1		3.7	3.7		$101	$174	
28	4.09	86	78	78.1	112.3	4.8	4.8	4.7	$288	$289	$525
29	4.22	81	21	13.0		4.7	5.7		$120	$204	
30	4.28	80	3			4.1			$83		
31	4.45	74	6	8.0	9.7	5.5	5.2	5.7	$113	$142	$238
32	4.59	72	2	2.0		2.8	2.8		$81	$82	
33	4.7	70	2	2.4		4.7	5.1		$65	$67	
34	4.79	65	5	5.0	9.4	1.1	1.1	2.1	$53	$54	$173
35	4.86	61	4	4.0		2.0	2.0		$67	$67	
36	5.03	0	61	61.0	60.0	0.0	1.8	1.9	$116	$116	$141
TOTAL			**726**	**717.1**	**721.1**	**3.0**	**3.3**	**3.3**	**$4,693**	**$4,775**	**$4,631**

Southbound, there are 35 stops in the historic stop set. Stop 20 was added as a suggestion by the consultant. The consultant recommended reducing the number of stops to 29; the optimization model recommends reducing that to 20. By more judiciously choosing stops to eliminate, the optimal solution results in slightly less average walk time than the consultant recommendation, in spite of eliminating more stops, and therefore has a better total cost.

Results at a daily level are shown Table 3. The two directions were optimized separately, considering impacts on the two periods combined. Results are shown for the two directions summed together. The number of stops recommended for the two directions combined falls from 71 in the historic case to 55 (consultant recommendation) and 45 (DP optimum). Compared to the historic case, both other recommendations involve a little more walking on average, and a little less riding time. In terms of round trip running time, the consultant recommendation saves 1.1 minutes and the DP optimum 2.4 minutes.

	Total 2-way stops	Average walk time (min)	Average ride time (min)	Avg two-way running time (min)	Change in Walk Cost ($/day)	Change in Ride Cost ($/Day)	Change in Oper Cost ($/Day)	Change in Total Cost ($/Day)
Historic	71	3.46	36.4	88.5	0	0	0	0
Analyst recommendation	55	3.77	35.7	87.4	705	-267	-113	326
DP Optimal	45	3.57	35.5	86.1	243	-218	-156	-132

In spite of a large difference in number of stops, the difference in overall impact is small, indicating that the total cost curve is rather flat near its optimum. Compared to the historic stop set, the optimal solution has a net societal savings of $130 per day, which is 0.66% of total walking, riding, and operating cost. Still, the value of an optimizing method can be seen by comparing it against the consultant's recommendation, which actually has a negative effect overall for the set of unit costs used. Its main weakness is that its stop location changes increase average walking time quite a bit, by 0.3 minutes, without comparable savings in riding time or operating cost.

A second case study is MBTA bus Route 57, which runs from Kenmore Station to the Watertown Yard running through Boston, Newton, and Watertown, and whose service area also includes parts of Brookline and Cambridge, as shown in Figure 3. It has 45 inbound and 43 outbound stops, and runs every 8 minutes in the a.m. peak, and every 10 minutes midday.

The analysis for Route 57 was similar to that for Route 1, except that the Route 57 analysis involved five inbound periods and six outbound periods. Table 4 shows results for the inbound direction.

	number of stops	Change in Walk Time (Pax-Min/Day)	Change in Running Time (min)	Change in Walk Cost ($/Day)	Change in Ride Cost ($/Day)	Change in Operating Cost ($/Day)	Change in Total Cost ($/Day)
Existing	45						
Consultant recommendation	31	7669	-1.9	$3,672	-$2,509	-$488	$674
DP Optimum	32	3955	-1.8	$1,761	-$2,211	-$469	-$920

The overall impacts are greater on Route 57 than on Route 1, with the optimum indicating a daily savings of $920, coming from savings in running and ride time, offset by small increases in walk time. In comparison, the consultant recommendation, with nearly a identical number of stops eliminated and similar running time savings, is not as judicious in its choice of stops to eliminate and ends up increasing societal cost overall, at least for the parameter set chosen.

A check on total demand for all of the Route 57 and Route 1 results shows that no demand was unaccounted for in any of the solution, supporting the practicality of the quintuplet-based DP approach.

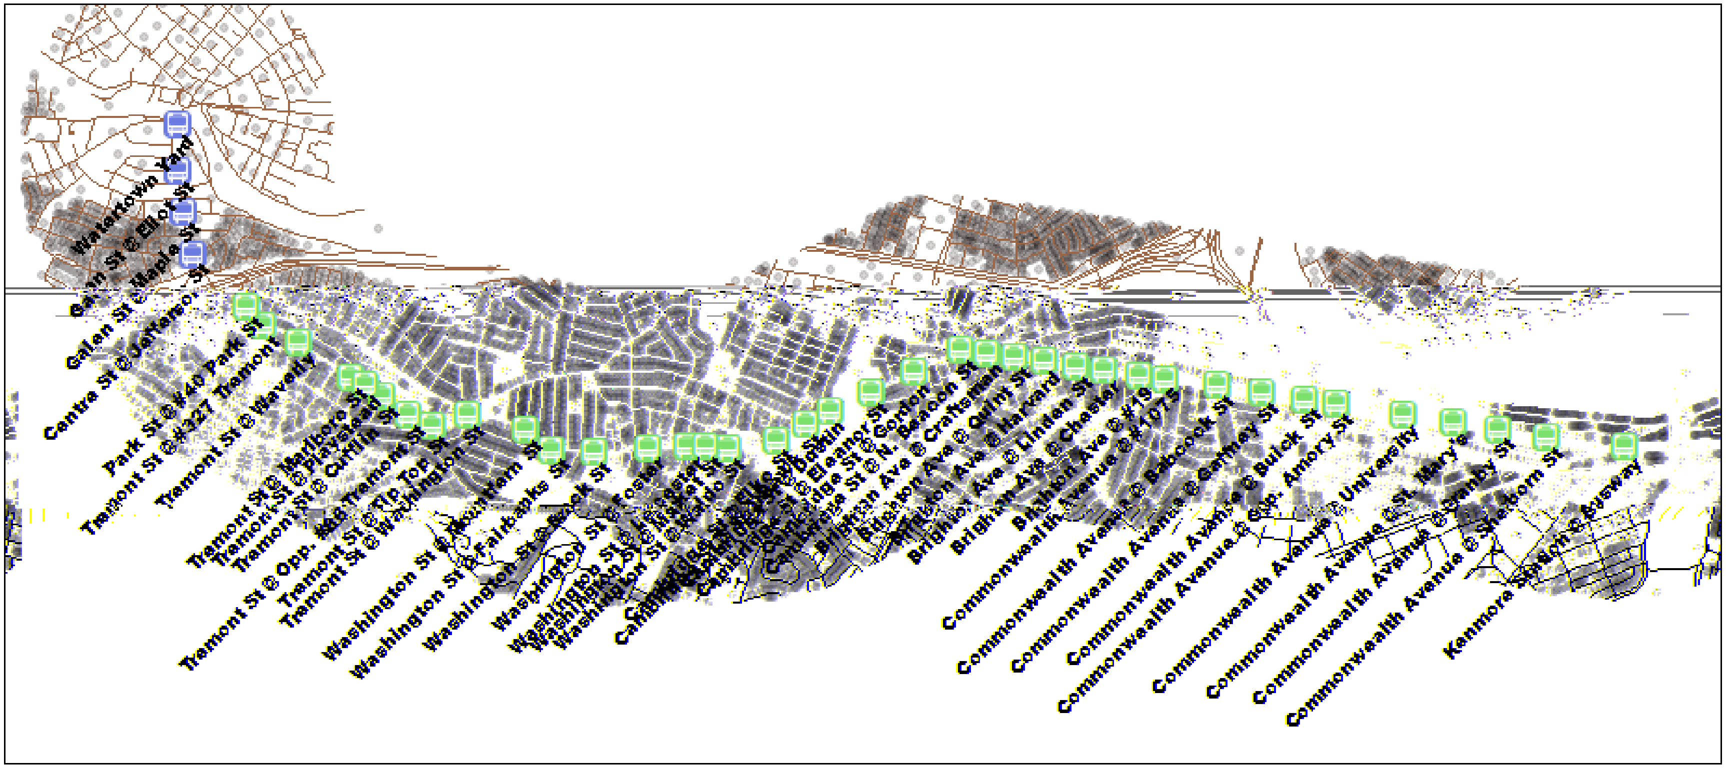
Watertown Yard
Galen St @ Eliot St
Centre St @ Jefferson
Park St @ #40 Park St
Tremont St @ Waverly
Tremont St @ Opp. Tremont St
Washington St @ Fairbanks St
Commonwealth Avenue @ Babcock St
Commonwealth Avenue @ Amory St
Commonwealth Avenue @ St. Mary's
Kenmore Station @ Busway

This paper introduces a stop-spacing optimization method that is realistic. Demand is modeled at the parcel level, sensitive to both parcel attributes and historic on-off patterns. Walking distance to stops takes place along the actual street network, with passengers choosing shortest paths. The “curve effect” that causes stops’ service areas to be sometimes bordered by more than one upstream or downstream stop’s service area and which renders single state variable dynamic programming unusable has shown to be overcome, at least for typical routes, using a 3-state formulation in which a stop’s service area and impacts depend only on the two prior and two succeeding stops. All of the necessary geographical databases were readily available.

Case studies of two Boston area bus routes revealed that on one route, the tradeoffs involved in stop spacing appeared to be so flat that gains related to speed are often countered by similar losses related to walking. On the other route, significant positive impacts were predicted from using an optimal solution with 32 instead of 45 stops. In both cases, we found that a consultant, using state of the art logic, recommended similar cuts in the optimal number of stops as the DP, but was unable to select them in a way that minimized the net impact. This result highlights the underlying complexity that makes stop-spacing tradeoffs difficult to do by hand or by rule of thumb.

1. J.M. Reilly. Transit Service Design and Operation Practices in Western European Countries. In Transportation Research Record 1604, Transportation Research Board, 1997.
2. S.C. Wirasinghe and N.S. Ghoneim. Spacing of bus-stops for many to many travel demand. Transportation Science 15, pp. 210-221, 1981.
3. Van Nes, R. and P.H.L. Bovy. The importance of objectives in urban transit network design. Transportation Research Record 1735, 2000, pp. 25-34.

4. Vuchic, V.R. and G.F. Newell. Rapid Transit Interstation Spacings for Minimum Travel Time. Transportation Science 2, pp. 303-339, 1968.
5. Furth, P. G. and A.B. Rahbee. Optimal Bus Stop Spacing Through Dynamic Programming and Geographic Modeling. Transportation Research Record 1731, 2000, p. 15-22.
6. Furth., P.G. and J.L. SanClemente, "Near-Side, Far-Side, Uphill, Downhill: Impact of Bus Stop Location on Bus Delay," Transportation Research Record 1971, 2006, pp. 66-73.
7. Furth., P.G., M.C. Mekuria, and J.L. SanClemente, "Stop Spacing Analysis Using GIS Tools with Parcel and Street Network Data," Transportation Research Record, 2034, 2007, pp. 73-81.
8. Furth., P.G., M.C. Mekuria, and J.L. SanClemente, "Parcel Level Modeling to Analyze Transit Stop Location Changes." Journal of Public Transportation 10, no. 2, 2007, pp. 73-91.