Properties of Vehicle Routes with Variable Shipment Sizes in Euclidean Plane

RANDOLPH W. HALL

A fundamental limitation of the literature on continuous-space routing models is that vehicles are assumed to have a capacity for a fixed number of stops. In reality, the maximum number of stops on a route may be defined by how well shipments pack into available capacities. An exploratory study of the geometric properties of vehicle routes that carry shipments of variable size (hence, capacity utilization and number of stops vary from route to route) is presented. Instead of taking a purely theoretical approach, the study relies on an empirical analysis of routing solutions to reveal the properties of near-optimal routes. Route geometry is then explained with a "theory of spokes," where a spoke is a line segment connecting the terminal to the most distant stop on a route. Because the number of spokes per unit circumference increases in the vicinity of the terminal, the incremental distance of serving a stop is a nonlinear function of distance to the terminal. Because of packing considerations, the incremental distance is also a nonlinear function of shipment size.

The classic vehicle routing problem (VRP) entails creating a set of routes of minimum total length so that each available stop is visited and vehicle loads do not violate constraints on capacity or time. The VRP has a long research history, beginning with the work of Dantzig and Ramser in 1959 (*1*). Research on the VRP is summarized by Bodin et al. (*2*), LaPorte and Nobert (*3*), and Magnanti (*4*), along with the book edited by Golden and Assad (*5*).

A number of researchers have used continuous-space models to study the theoretical behavior and geometric properties of VRP solutions. Most notable are the works by Daganzo (*6,7*) and Newell and Daganzo (*8*). These build from earlier models for the traveling salesman problem by Beardwood et al. (*9*), Few (*10*), Stein (*11*), and Verblunsky (*12*), and empirical work on the VRP by Christofides and Eilon (*13*). The accomplishment of Daganzo and Newell was to develop a theory of near-optimal route geometry and expected route length in Euclidean space. Haimovich and Rinnooy Kan (*14*); Haimovich, Rinnooy Kan, and Stougie (*15*); and Spaccamela, Rinnooy Kan, and Stougie (*16*) have also produced bounds on the length of VRP routes in Euclidean space as part of their effort to evaluate the asymptotic performance of routing heuristics.

More recently, routes with mixed pickups and deliveries have been studied by Daganzo and Hall (*17*) and Hall (*18*). The latter identifies the optimal shape and orientation of routes that begin and end at different terminals. The former studies mixed pickup and delivery routes out of a single terminal. The work by Daganzo and Hall (*17*) is significant in that it introduces the concept of line-haul spokes, which are used in this paper to explain the effects of capacity utilization on route length. Also, research has recently been completed on dynamic routing, where the presence of stops or characteristics of shipments, or both, are not revealed until the vehicle is in motion (*19–21*).

A fundamental limitation of the continuous-space literature is that vehicles are assumed to have a capacity for a fixed number of stops. In reality, the maximum number of stops may be defined by how well shipments pack into available capacities (e.g., weight of time). And because shipments do not have to be identical, the number of stops and the capacity utilization can vary from route to route. For instance, if splitting shipments among routes is not allowed, then solving the VRP depends in part on the solution to a bin-packing problem. However, unlike the classic bin-packing problem, the objective is not simply to minimize the number of bins, but is instead to minimize the product of the number of bins (i.e., routes) and the average length per route. The latter depends on the dispersion of stops on routes.

As discussed by Hall et al. (*22*), the two goals of minimizing the number of routes and the average length per route conflict with each other because efficient packings may demand noncompact routes. This trade-off is most prominent when shipments tend to be large relative to vehicle capacity.

Though the theoretical properties of bin-packing algorithms are well known (*23*) and the geometric properties of vehicle routes with a fixed number of stops are well known, little research had been completed on the geometric properties of vehicle routes with variable shipment sizes. (There is, however, an extensive body of literature on algorithms for solving vehicle routing problems with variable shipment sizes.) Hall and Daganzo (*24*) examined route characteristics when vehicles are limited by weight and volume constraints with infinitely divisible commodities. Hall (*25*) studied the trade-off between packing efficiency and average route length under a scheme whereby sets of vehicles are allowed to cove identical territories. However, this scheme is surely inferior to partially overlapped territories.

The introduction of variable shipment sizes motivates changes in route structure as well as changes in the relationship between route length and stop density. The objective of this paper is to develop an understanding of how variations in shipment size affect optimal route length and optimal route geometry.

Although this research entails testing algorithms, algorithm development is not a primary goal. The intention is to begin the development of an empirically based theory of routing

Department of Industrial Engineering and Operations Research, Institute of Transportation Studies, University of California, Berkeley, Calif. 94720; current affiliation: Department of Information and Operations Management, School of Business Administration, University of Southern California, Los Angeles, Calif. 90089–1421.

with variable shipment sizes. The hope is that by understanding the geometric behavior of routes created by good heuristics, it will be possible in the future to develop approximation algorithms that quickly produce near-optimal solutions. These algorithms may identify heuristic solutions by optimizing an approximation to the true objective function. The benefits are twofold: (a) a good initial solution, produced by an approximate algorithm, may preclude the need to apply fine-tuning algorithms or (b) if a fine-tuning algorithm is used, it may be possible to obtain a slightly better final solution or to find a good final solution in fewer iterations.

Throughout this paper, the routing objective will be to minimize total length of all routes, which will be defined by the Euclidean metric. Stop locations will be independently distributed over a circle of radius R according to a uniform distribution. Routes will consist of either deliveries only or pickups only from or to a single terminal, or both, and all routes will be constructed simultaneously. Capacity and shipment sizes will be deterministic and defined by a single attribute, such as weight or volume.

PREDICTED ROUTE LENGTH

This section considers two formulations of the VRP with variable shipment sizes. The first formulation allows shipments to be split among vehicles. The second formulation assumes that each shipment is assigned to a single vehicle, as is customarily done in vehicle routing algorithms. To achieve an efficient use of vehicle capacity, some route districts must, consequently, overlap. (Conceptually, a routing district is the convex hull of the collection of stops on a vehicle tour, absent the terminal.) In practice, split assignments can enable efficient capacity utilization without much overlap (*26*).

In contrast, the literature on continuous-space models considers neither the issue of overlapping districts nor split shipments. Because shipments are assumed to be identical in size, overlap and splitting is not needed to fill vehicles to capacity. Therefore, continuous-space models need to be adjusted to accurately account for the extra travel distance from overlap and splitting that are needed in practice.

Split Routing

To illustrate the fundamental difference between split and nonsplit routing, this section first casts the split routing problem in the context of Fisher and Jaikumar's (*27*) generalized assignment methodology. As with Fisher and Jaikumar's methodology, routing is viewed as a two-stage process: (a) an assignment of stops to vehicles and (b) the routing of individual vehicles. Let c_{ij} be an approximation for the incremental cost of serving stop i on route j (*27*). Let I be the total number of stops, and let J be the total number of routes. Then the assignment of stops to vehicles with split routing amounts to

$$\min_{X_{ij}, Y_{ij}} \Sigma c_{ij} Y_{ij} \tag{1}$$

such that

$$\Sigma_j X_{ij} = q_i \qquad i = 1, \ldots, I \tag{1a}$$

$$\Sigma_i X_{ij} \le s_j \qquad j = 1, \ldots, J \tag{1b}$$

$$(1 - Y_{ij}) X_{ij} = 0 \qquad \forall\, i,j \tag{1c}$$

$$X_{ij}, Y_{ij} \ge 0 \qquad \forall\, i,j \tag{1d}$$

$$Y_{ij} = 0 \text{ or } 1 \qquad \forall\, i,j \tag{1e}$$

where

q_i = shipment quantity for stop i,
s_j = size of vehicle j,
c_{ij} = incremental cost of assigning stop i to vehicle j,
X_{ij} = quantity assigned from stop i to route j, and

$$y_{ij} = \begin{cases} 1, & X_{ij} > 0 \\ 0, & X_{ij} = 0 \end{cases}$$

The primary difference between this formulation and the customary generalized assignment problem (GAP) formulation is that two decision variables are needed for each stop/vehicle pair: (a) the shipment quantity assigned and (b) Y_{ij}, which indicates whether route j is used for stop i. The formulation also resembles the classic transportation problem. The primary difference is that Equation 1 measures the circuitry cost of diverting a vehicle to a stop, which is independent of the quantity assigned to the vehcile, and therefore requires integer variables.

Now, consider how the average route length can be estimated from a continuous space model. If vehicles are filled to 100 percent capacity (i.e., $\Sigma_i q_i = \Sigma_j s_j$), there exists an optimal solution to Equation 1 such that the number of nonzero values of X_{ij} does not exceed $I + J$, the sum of the number Type a and b constraints. (If vehicles are not filled to 100 percent of capacity, there will be fewer tight constraints and fewer nonzero values). Suppose that stops are partitioned into nonoverlapping districts, such that if a customer is split among routes, then it must fall on the boundary(ies) dividing the districts.

Let

$\bar{d}$ = average distance from terminal to stop,
ρ = stop density,
s = vehicle capacity (assumed to be identical), and
N = number of stops per route (based on fixed shipment size).

Daganzo's (*6*) approximation for route length (for uniformly and independently distributed stops over a large region) is composed of a line-haul component, dependent on $\bar{d}$ and N, and a local component dependent on ρ. A simple adjustment to Daganzo's approximation would be to estimate the line-haul cost from the number of routes (J) and the local cost from the maximum number of times stops are visited (I + J):

$$\text{Total length} \approx 2\bar{d}J + .57(I + J)/\sqrt{\rho} \tag{2}$$

On a per stop basis, Equation 2 becomes

$$\overline{D} = \text{mean route length per stop}$$

$$= 2\overline{d}(J/I) + .57[(I + J)/I]/\sqrt{\rho} \qquad (3)$$

The substantive difference from Daganzo's result is the inclusion of the multiplier $(I + J)/I$ in the second term to account for stops that are visited more than once. This adjustment is negligible if N is large (i.e., $I \gg J$). J is interpreted as the minimum number of vehicles needed to accommodate the freight because of split shipments.

Splitting Disallowed

When split shipments are disallowed, each stop is visited exactly once. However, the optimal number of routes must certainly exceed $\Sigma q_i/s$ because vehicles cannot economically or feasibly be filled to 100 percent of capacity. If J is the actual number of routes employed, Daganzo's model could be interpreted as

$$\overline{D} = 2\overline{d}(J/I) + .57/\sqrt{\rho} \qquad (4)$$

Unfortunately, for nonsplit shipments, J is itself a product of the optimization process, for it depends on the optimal capacity utilization. Therefore, the model is useful only if capacity utilization can be predicted in advance. Equation 4 may also underestimate local distance because it does not account for the fact that, for a given number of stops, overlapping districts will be less compact than nonoverlapping districts.

Validation

The models were validated by comparing route length predictions to actual routes constructed for a series of test problems. Although the routing methods used are, by necessity, heuristic, they replicate the logic underlying the analytical model. Although this allows the route-length approximation to be validated, it does not allow validation of optimality. The basic structure of the heuristics is as follows:

- Form an initial feasible solution with a heuristic based on continuous-space approximations and
- Fine tune the initial solution with a heuristic that accounts for discrete stop locations.

Routes were created for a series of 160 test problems found in work by Hall et al. (*22*), with 20 to 170 total stops. In each case, stops were randomly and independently located according to a uniform distribution over a circle of radius R with the terminal in the center. This radius increased as the number of stops increased to maintain a uniform density of approximately 1. Twenty problems were solved within each category, which was defined by the number of stops and the coefficient of variation of the shipment sizes (always a uniform random variable). In each case, the expected shipment size was one-third of the vehicle capacity (a constant value s). Large shipment sizes are used for two reasons: (a) to ensure that vehicle packing is an important factor and (b) to ensure that larger problems generate multiple rings of routing districts.

In the case of split-shipment routing, initial assignments were found with the continuous-space initialization heuristic of Hall et al. (*22*), which partitions the service region into districts with a combination of dynamic programming and a sweep algorithm (the dynamic program creates ring boundaries by optimizing a continuous-space approximation). The sweep algorithm terminated a district as soon as vehicle capacity was reached, which forced some stops to be split among adjacent districts. The initial assignment was then adjusted by applying the heuristic of Dror and Trudeau (*26*). This adjustment stage produced reductions over the initial solution on the order of 6 to 8 percent for 20 stop problems and 1 to 2 percent for larger problems.

In the case of nonsplit routing, an initial partition was found in the same manner as the split case, except that the sweep algorithm allowed partial overlap among districts (*22*). Initial assignments were updated by applying a generalized assignment algorithm (*22*). Once final assignments were made, individual vehicles were routed with Little et al.'s (*28*) traveling salesman optimization algorithm.

Table 1 presents average results along with predictions for the split case. The predicted length/stop assumes that all routes are filled to 100 percent of capacity (hence $J = \Sigma d_i/s$). These predictions tend to be slightly less than observed (up to 3 percent). This discrepancy may be due to the heuristic nature of the solution. Perhaps more important, it may be due to the fact that actual aggregate capacity utilization was only 95 to 98 percent, slightly less than the assumed 100 percent. With this in mind, the adjusted prediction factors actual capacity utilization into the line-haul cost (i.e., the line-haul distance was multiplied by the factor $100/P$, where P is the percentage capacity utilization). The latter accurately predicted route length for the large problems ($I = 115$ and $I = 170$). Predictions are not as accurate for smaller problems, possibly because fewer than $I + J$ total stops are made and possibly because Daganzo's model is an asymptotic result. In any case, there is no reason to doubt that simple adjustments to Daganzo's model produce reasonable predictions for route length when split shipments are allowed (Ideally, P would be determined endogenously. Later in the paper, insights will be provided into how this might eventually be accomplished.)

Results for the nonsplit case are provided in Table 2. Predictions are based on observed capacity utilizations. With the exception of the 20-stop case, predictions are also reasonably close to observed values. Furthermore, test results (*22*) show

TABLE 1 Route Lengths with Split Shipments

	I=20/R=2.45		I=75/R=4.90		I=115/R=6.12		I=170/R=7.35	
	CV=.5	CV=1.0	CV=.5	CV=1.0	CV=.5	CV=1.0	CV=.5	CV=1.0
Average	36.8	36.6	227.0	228.0	414.8	414.2	708.2	707.6
Average/ Stop	1.84	1.83	3.03	3.04	3.61	3.60	4.16	4.16
Predicted/ Stop	1.83		2.94		3.49		4.03	
Adjusted Predicted/ Stop	2.07		3.08		3.61		4.17	

Each problem class contains 20 randomly generated problems. J stops are uniformly distributed over a circle of radius R. Shipment size has uniform distribution with mean of 1/3 vehicle capacity and coefficient of variation (CV) indicated.

TABLE 2 Route Lengths Without Split Shipments

	I=20/R=2.45 CV=.5	I=20/R=2.45 CV=1.0	I=75/R=4.90 CV=.5	I=75/R=4.90 CV=1.0	I=170/R=7.35 CV=.5	I=170/R=7.35 CV=1.0
Average	36.8	36.8	228.8	230.2	728.2	726.0
Average/ Stop	1.84	1.84	3.05	3.07	4.28	4.27
Predicted/ Stop	1.94		3.08		4.15	

Each problem class contains 20 randomly generated problems. J stops are uniformly distributed over a circle of radius R. Shipment size has uniform distribution with mean of 1/3 vehicle capacity and coefficient of variation (CV) indicated.

close agreement when stop density is a slowly varying function of the distance from the terminal. So, again, there seems to be no reason to doubt that simple adjustments to Daganzo's model produce reasonable predictions for route length.

There remains the possibility that the heuristics produce nonoptimal solutions, in which case the approximation would overestimate the true optimal route length.

PREDICTED ROUTE GEOMETRY FOR NONSPLIT ROUTING

Just because route length conforms to model predictions, it does not follow that route geometry conforms to model predictions. This is especially true for the nonsplit shipment case, which must contain some overlap among districts. With this in mind, this section examines the observed geometry of routes that prohibit split shipments. The hope is that a better understanding of optimal route geometry will enable the development of better approximation based heuristics, which may reduce the need for fine-tuning algorithms.

Number of Rings

Daganzo (*6*) and Newell and Daganzo (*8*) represent optimal route geometry with a series of circular rings centered at the terminal and split into routing districts by line segments radiating from the terminal.

Ring depth is defined as the depth of a ring that partitions the service region into districts. (The depth is the radial separation between the two concentric circles that bound the ring.) District length is defined as the radial distance between the closest stop to the terminal within a district and the furthest stop to the terminal within the district.

Asymptotically, as the number of stops per district becomes large, the optimal ring depth and the optimal district length are both approximated by $\ell^* = N/\sqrt{6.7\rho}$, where N is the number of stops per route (*6*).

Within the initialization algorithm of Hall et al. (*22*), the optimal ring depth is approximated by solving a dynamic program that incorporates a ring-radial continuous-space approximation. The author's concern is whether the solutions produced from this approximation are similar to the near-optimal solutions found at termination after the algorithm has been applied.

To address this issue, 20 problems were solved within each of the three classes of test problems:

- Class 1: 170 stops, mean shipment size = ⅓ capacity, shipment size coefficient of variation (CV) = 0;
- Class 2: 170 stops, mean shipment size = ⅓ capacity, shipment size CV = .5; and
- Class 3: 170 stops, mean shipment size = ⅓ capacity, shipment size CV = 1.0.

In all cases, stops were uniformly and independently distributed over a circle of radius 7.35. For each problem, the VRP was solved approximately using the following algorithm:

1. Initialization: Set n = desired number of rings (n = 2, 3, 4, 5, 6, or 7).
2. Divide service region into exactly n rings. Determine boundaries between annuli with dynamic program in Hall (*18*) 1991, which optimizes continuous-space approximation.
3. Partition each ring into routing districts. Perform partition with modified sweep algorithm in Hall (*18*) 1991.
4. For each district, find the optimal traveling salesman route with the branch-and-bound algorithm of Little et al. (*28*).

This algorithm is analogous to the initialization steps of (*22*) (i.e., the solution is not adjusted with a GAP algorithm), except that n is constrained rather than optimized.

Table 3 provides the average route length as determined from the actual routes. Table 3 also provides the estimated route length as derived from the ring-radial continuous-space approximation imbedded in the dynamic program. There are two surprising results:

- The number of rings that minimizes actual route length is consistently less than the optimum determined by the approximation.
- The actual route length is insensitive to the number of rings.

Both results raise doubt as to the validity of the continuous-space theory for predicting optimal route geometry with small N. But there is an important caveat: ring depth (R/n) and district length are identical only when the number of stops/route is large. When N is small, the average length of a district should be less than the depth of a ring. Therefore, it may be that the continuous-space theory accurately predicts optimal district length but not optimal ring depth.

TABLE 3 Route Lengths with Variable Rings, Without Split Shipments (Total Length Among All Routes)

	RINGS 2	3	4	5	6	7
CV=0						
Estimated*	738	705	693	690	691	693
Actual	665	654	655	659	663	665
CV=.5						
Estimated*	740	706	694	690	690	692
Actual	750	744	750	757	761	767
CV=1.0						
Estimated*	741	706	694	690	690	692
Actual	753	760	771	778	785	790

* Estimated based on ring-radial metric with 100% capacity utilization. Actual is based on initial partition of stops into routes, with application of Little et al's (1983) algorithm.

† Each problem class contains 20 randomly generated problems. 170 stops are uniformly distributed over a circle of radius 7.35. Shipment size has uniform distribution with mean of 1/3 vehicle capacity and coefficient of variation (CV) indicated.

District Length

To investigate district length, the complete algorithm of Hall et al. (*22*) was applied to the problems in Classes 1 and 2. Based on the final routes obtained, the following functions were derived:

- $F_o(r)$ = proportion of routes whose furthest stop is within a distance r of terminal,
- $F_i(r)$ = proportion of routes whose closest stop is within a distance r of terminal, and
- $F(r)$ = proportion of stops that are located within a distance r of terminal.

$F_o(r)$ and $F_i(r)$ are plotted for shipment size CVs of 0 and 0.5 in Figure 1. From $F_o(r)$ and $F_i(r)$, the following statistics were derived:

$$\bar{o} = \text{mean distance to furthest stop} = \int_0^\infty [1 - F_o(r)]dr$$

$$\bar{i} = \text{mean distance to closest stop} = \int_0^\infty [1 - F_i(r)]dr$$

$$\bar{d} = \text{mean radial distance to stop} = \int_0^\infty [1 - F(r)]dr$$

$$\bar{\ell} = \text{mean district length} = \bar{o} - \bar{i}$$

Daganzo (*6*) predicts that $(\bar{o} + \bar{i})/2$ equals the mean radial distance to a stop, or $(2/3)R$ (R = radius of region). The model also predicts that $\bar{\ell} = N/\sqrt{6.7\rho}$. The following table compares the predictions to the observations.

	Identical	*CV = 0.5*	*Theory*
$\bar{o}$	5.35	5.36	5.48
$\bar{i}$	4.41	4.10	4.22
$\frac{\bar{o} + \bar{i}}{2}$	4.88	4.73	4.9
$\bar{\ell}$	0.93	1.25	1.16 (N = 3)

Unlike ring depth, the continuous-space model appears to overestimate, not underestimate, district length when shipment sizes are identical. As illustrated by Figure 1, $F_o(r)$ and $F_i(r)$ do not exhibit a staircase pattern, which would be expected if districts neatly fit within rings. Instead, $F_o(r)$ and $F_i(r)$ are smooth functions, indicating that district boundaries are randomly (but not uniformly) scattered over the service region. Nevertheless, route characteristics are similar to the theory in two important respects: (a) when the CV = 0, $(\bar{o} + \bar{i})/2$ (an approximation for the mean location of district centroids) is almost identical to $(2/3)R$, and (b) mean district length is comparable to (though less than) $\bar{\ell}$ for small N.

Results are different for nonidentical shipment sizes. Although $\bar{o}$ is nearly the same for the CV = 0 and the CV = 0.5 cases, $\bar{i}$ is not. The inner edges of districts are drawn closer to the terminal when CV = 0.5.

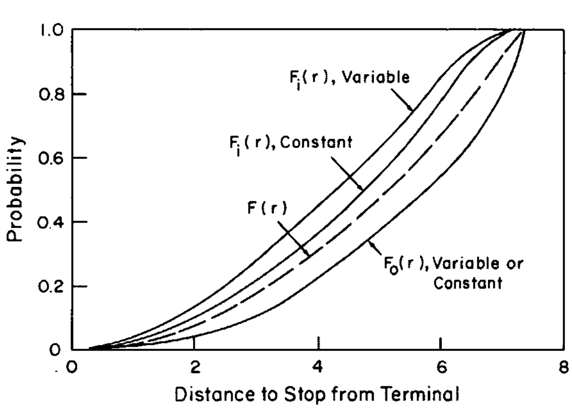

FIGURE 1 Probability distributions for location of inside stop and outside stop (variable and identical shipment sizes).

Theory of Spokes

This section provides a preliminary way to measure the effects of excess vehicle capacity due to imperfect packings and explains the previous finding that the inner edges of districts are drawn closer to the terminal. It borrows from a theory of spokes, introduced by Daganzo and Hall (*17*) in a paper on routing with pickups and deliveries. For each route, a spoke can be envisioned as the line segment connecting the terminal to the most distant stop on the route. The angular position of spokes will be ignored, but radial length and the assigned load size will be incorporated.

Let

o_j = radial distance to most distant stop from terminal on route j (denoted outside stop),

i_j = radial distance to closest stop to terminal on route j (denoted inside stop),

$O(r)$ = number of spokes that end outside circle of radius r centered on depot,
= number of routes for which $o_j \geq r$, and

$I(r)$ = number of spokes whose inside stop is outside r,
= number of routes for which $i_j \geq r$.

According to Daganzo (*6*), districts are rectangular (with dimensions that are independent of location relative to the terminal), and districts do not overlap. Consequently, district width, denoted w, is invariant to r:

$$W \approx \frac{2\pi r}{O(r) - I(r)} \tag{5}$$

where $O(r) - I(r)$ is a constant multiple of r. When shipment sizes are not identical, districts must overlap to attain an efficient packing, and a 100-percent capacity utilization is neither optimal nor (usually) feasible. Hence, the question is: What is the optimal pattern for overlapping routes?

Overlap Within Rings

In the work by Hall et al. (*22*), the service region was partitioned into rings exactly as though stops were identical in size, but districts were allowed to overlap within rings. Though reasonable as a first-cut analysis, observed values of $\bar{i}$ indicate that the inside edges of routes are pulled toward the terminal, which suggests that districts should overlap in the radial direction.

Overlap Between Rings

To understand the process by which routes overlap, the incremental distance for serving a stop with a shipment of size v located at a distance r from the terminal is discussed.

Let

$S(v,r)$ = number of spokes that cross circle with radius r and carry a load size less than or equal to $s - v$.

$S(v,r)$, which will be called the number of surplus spokes, is a nonincreasing function with respect to v and r.

The author's hypothesis is that the incremental distance of inserting a stop in an existing route is approximately proportional to the inverse of the ratio (number of surplus spokes per unit circumference):

$d'(v,r)$ = incremental distance for stop of size v located at r

$$\approx k \cdot \frac{2\pi r}{2S(v,r)} \tag{6}$$

The coefficient 2 in the denominator accounts for both the vehicle's forward and reverse trips (in essence, each route creates two spokes). The coefficient k reflects the spatial distribution of spokes. If spokes serve equal sized and nonoverlapping arcs and stops are served by ring-radial paths, then k would equal ⅓. However, because new spokes are continuously introduced, it is unrealistic to maintain nonoverlapping arcs. Alternatively, the polar positions of spokes might be independent uniform $[0,2\pi]$ random variables.

Then

$F(x)$ = probability nearest spoke is ring distance of x or greater

$$= \left[\frac{2\pi r - 2x}{2\pi r}\right]^{2 \cdot S(v,r)}$$

$$d'(v,r) = \int_0^{\pi r}\left[\frac{2\pi r - 2x}{2\pi r}\right]^{2 \cdot S(v,r)} dx = {}^1\!/_2\left(\frac{2\pi r}{2S(v,r) + 1}\right) \tag{7}$$

An important feature of Equations 6 and 7 is that $d'(v,r)$ is nonlinear. The incremental distance to retrieve a shipment located close to the terminal is quite small, both because r is small and because the number of surplus spokes is large. And because $S(v,r)$ is nonincreasing, $d'(v,r)$ increases at an increasing rate as r increases.

The relationship between $d'(v,r)$ and weight is also nonlinear. Because optimally routed vehicles tend to be filled close, but not completely, to capacity, $S(v,r)$ will be large for small values of v. Hence, the incremental distance of serving a small shipment can be negligible. On the other hand, for large values of v, $S(v,r)$ may be as small as 0, in which case it may be impossible to serve the shipment without reassigning stops or introducing a new route. In either case, the incremental distance is large. Overall, the relationship between incremental distance and shipment size is unlikely to be a smooth linear function, but instead something more akin to a threshold function with a low cost below the threshold and a high cost above (*25*).

As illustrations, Figures 2 and 3 show examples of $S(v,r)$ and $d'(v,r)$ (Equation 7) as averaged from the twenty 170-stop problems with a shipment size CV of 0.5 and a vehicle capacity of s = 1,000. For example, Figure 2 shows that approximately 50 of the total 63 spokes terminate outside the circle of radius 4. Of these 50 spokes, 30 are filled to no more than 90 percent of capacity [the remaining space equals or exceeds shipment size (v) of 100], 11 are filled to no more than 60 percent of capacity [the remaining space equals or exceeds shipment size (v) of 400], and so on. Figure 3 uses Equation 6 (with k = 0.5) and the data in Figure 2 to estimate incremental distance. The figure demonstrates the nonlinear relationship between incremental distance and shipment size and distance, as discussed earlier.

The incremental distances predicted by Equation 7 have not been verified, an effort that would entail a massive computational effort and repeated solution of VRPs with and without stops inserted into routes. Nevertheless, the implication that incremental cost is a nonlinear function of shipment size and shipment distance, with increasing marginal cost, appears highly plausible.

The theory of spokes may also explain why route-length predictions are accurate even when districts are known to overlap each other. The existence of surplus capacity effectively reduces the local distance serving a stop. It may be that this reduction is adequate to compensate for the increased separation between stops when districts overlap (and by necessity cover larger areas per stop).

CONCLUSIONS

This paper presented an exploratory study of vehicle routing with shipments of variable size, where shipment size is large

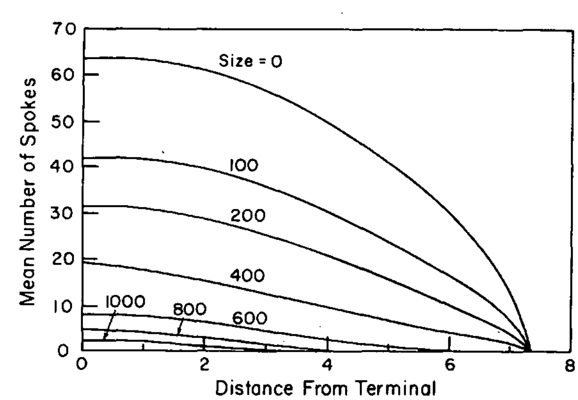

FIGURE 2 Mean number of spokes as function of distance from terminal and shipment size. (Spoke is counted if remaining space exceeds shipment size and spoke extends beyond indicated distance.)

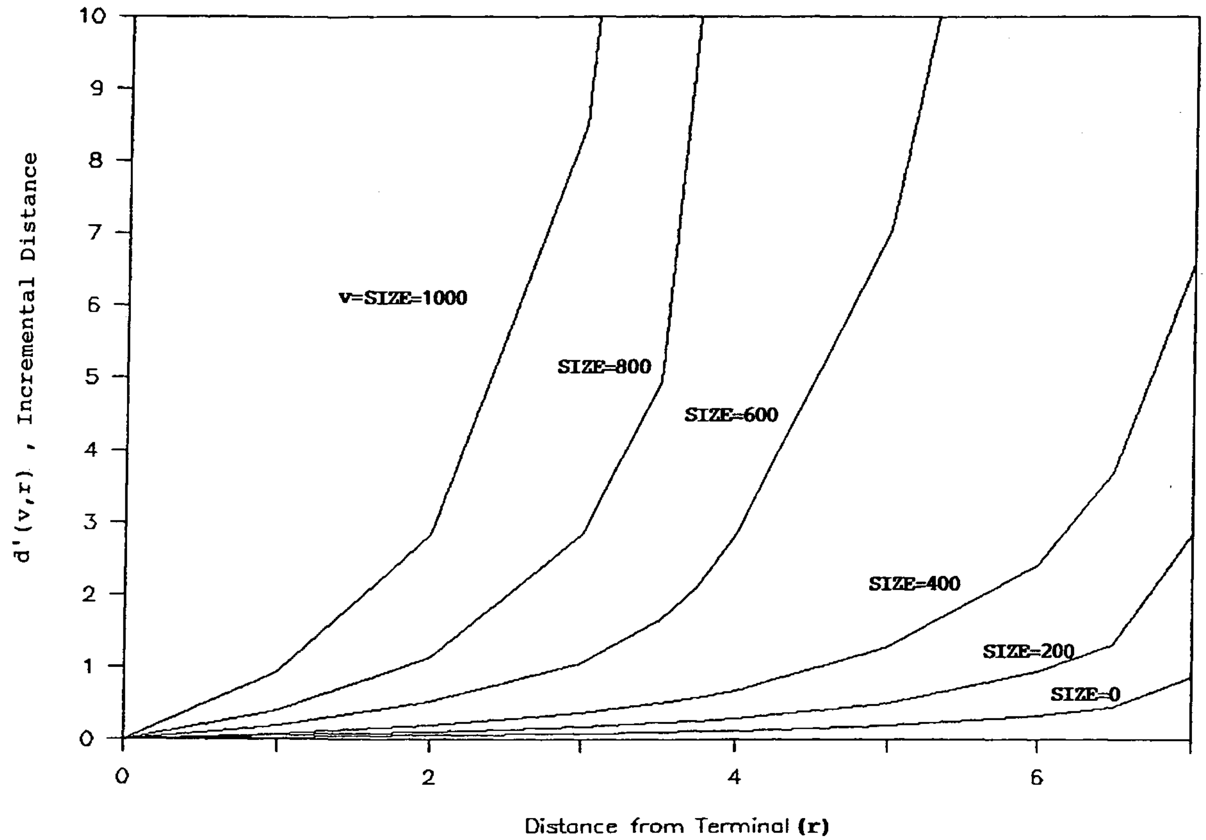

FIGURE 3 Incremental distance to serve a stop as a function of shipment size (ν) and distance from stop to terminal (r).

relative to vehicle capacity. Empirical results suggest that simple modifications to Daganzo's model lead to reasonable predictions for average route length, when split shipments are allowed and disallowed. Despite the accuracy of the route length predictions, route geometry does not match continuous-space theory. It differs in the important respect that districts do not neatly fit within rings—whether or not shipment sizes are identical. Instead, routes seem to be randomly scattered across the service region according to a continuous probability distribution.

This said, district characteristics—such as the location of centroids—are still similar to model predictions, especially when shipment sizes are identical. However, when shipment sizes are not identical, the location of the "inner stop" is pulled toward the terminal. This phenomenon is explained in terms of a theory of spokes, which also serves to explain why the existence of surplus capacity reduces the incremental cost of serving small stops and stops located close to the terminal. As of yet, the theory of spokes has not been developed to the point where it can be used to predict optimal route length or capacity utilization. This is the subject of future research.

The author of this paper hopes that an improved understanding of route geometry will lead to better approximation-based heuristics. The GAP algorithm is computationally expensive in both memory and time. It would be highly desirable to obtain good solutions without resorting to repeated application of GAP. One idea that the author has examined is to use a random sample of stops as a collection of seed points and approximate the assignment cost by the incremental distance function of Fisher and Jaikumar (*27*). Unlike Fisher and Jaikumar, the author proposes that seed points be based on the empirically derived function $F_o(r)$. Specifically, randomly select J seed points without replacement from the set of I stops, with the probability of selecting stop i given by

$$P_i = \left[\frac{dF_o(r_i)}{dr}\right] \Big/ \left[\frac{dF(r_i)}{dr}\right] \alpha \tag{8}$$

where α is a normalizing constant. The author's tests of this and other approaches have not yet produced substantial improvements over prior methods. As of yet, it remains an open research question whether empirically derived results can be the basis for effective routing heuristics.

REFERENCES

1. Dantzig, G., and J. Ramser. The Truck Dispatching Problem. *Management Science*, Vol. 6, 1959, pp. 80–91.
2. Bodin, L., B. Golden, A. Assad, and M. Ball. Routing and Scheduling of Vehicles and Crews: The State of the Art. *Computers and Operations Research*, Vol. 10, 1983, pp. 105–121.
3. Laporte, G., and Y. Nobert. Exact Algorithms for the Vehicle Routing Problem. *Annals of Discrete Mathematics*, Vol. 31, 1987, pp. 147–184.
4. Magnanti, T. L. Combinatorial Optimization and Vehicle Fleet Planning, Perspectives and Prospects. *Networks*, Vol. 11, 1981, pp. 179–213.

5. *Vehicle Routing Methods and Studies* (B. J. Golden and A. A. Assad, eds.), North-Holland Elsevier, Amsterdam, The Netherlands, 1988.
6. Daganzo, C. F. The Distance Traveled to Visit N Points with a Maximum of C Stops per Vehicle: An Analytical Model and an Application. *Transportation Science*, Vol. 18, 1984, p. 331–350.
7. Daganzo, C. F. The Length of Tours in Zones of Different Shapes. *Transportation Research*, Vol. 18B, 1984, pp. 135–146.
8. Newell, G. F., and C. F. Daganzo. Design of Multiple Vehicle Delivery Tours—I, a Ring Radial Network. *Transportation Research*, Vol. 20B, 1986, p. 345–364.
9. Beardwood J., J. H. Halton, and J. M. Hammersley. The Shortest Path Through Many Points. *Proceedings of the Cambridge Philosophical Society*, Vol. 55, 1959, pp. 299–327.
10. Few, L. The Shortest Path and the Shortest Road Through n Points. *Mathematika*, Vol. 2, 1955, pp. 141–144.
11. Stein, D. M. Asymptotic, Probabilistic Analysis of a Routing Problem. *Mathematics of Operations Research*, Vol. 3, 1978, pp. 89–101.
12. Verblunksy, S. On the Shortest Path Through a Number of Points. *Proceedings of the American Mathematical Society*, Vol. 2, 1951, pp. 904–913.
13. Christofides, N., and S. Eilon. Expected Distances in Distribution Problems. *Operational Research Quarterly*, Vol. 20, 1969, pp. 437–443.
14. Haimovich, H., and A. H. G. Rinnooy Kan. Bounds and Heuristics for Capacitated Routing Problems. *Mathematics of Operations Research*, Vol. 10, 1985, pp. 527–542.
15. Haimovich, H., A. H. G. Rinnooy Kan, and L. Stougie. Analysis of Heuristics for Vehicle Routing Problems. In *Vehicle Routing: Methods and Studies* (B. L. Golden and A. A. Assad, eds.), North-Holland, Amsterdam, The Netherlands, 1988, pp. 47–62.
16. Spaccamela, M., A. H. G. Rinnooy Kan, and L. Stougie. Hierarchical Vehicle Routing Problems. *Networks*, Vol. 14, 1984, pp. 571–586.
17. Daganzo, C. F., and R. W. Hall. A Routing Model for Pickups and Deliveries: No Capacity Restrictions on the Secondary Items. *Transportation Science* (in press).
18. Hall, R. W. Characteristics of Multi-stop/Multi-terminal Delivery Routes, with Backhauls and Unique Items. *Transportation Research*, Vol. 25B, 1991, pp. 391–403.
19. Bertsimas, D. J., and G. van Ryzin. *A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane*. OR 210-90. Operations Research Center, Massachusetts Institute of Technology, Cambridge, 1990.
20. Hall, R. W. Pickup and Delivery Strategies for Overnight Carriers. *Transportation Research* (in press).
21. Simchi-Levi, D. *Hierarchical Planning for Probabilistic Distribution Systems in Euclidean Spaces*. Department of Industrial Engineering and Operations Research, Columbia University, New York, 1990.
22. Hall, R. W., Y. F. Du, and J. Lin. Integration of Discrete and Continuous Models for Routing Vehicles. *Networks* (in press).
23. Ong, H. L., M. J. Magazine, and T. S. Wee. Probabilistic Analysis of Bin Packing Heuristics. *Operations Research*, Vol. 32, 1984, pp. 983–998.
24. Hall, R. W., and C. F. Daganzo. Vehicle Miles for a Freight Carrier with Two Capacity Constraints. In *Transportation Research Record 1038*, TRB, National Research Council, Washington, D.C., 1985, pp. 34–40.
25. Hall, R. W. Vehicle Packing. *Transportation Research*, Vol. 23B, 1989, pp. 103–121.
26. Dror, M., and P. Trudeau. Split Delivery Routing. *Naval Research Logistics Quarterly*, Vol. 37, 1990, pp. 383–402.
27. Fisher, M. L., and R. Jaikumar. A Generalized Assignment Heuristic for Vehicle Routing. *Networks*, Vol. 11, 1981, pp. 109–124.
28. Little, J., K. Murty, D. Sweeney, and C. Karel. An Algorithm for the Traveling Salesman Problem. *Operations Research*, Vol. 11, 1963, pp. 972–989.

Publication of this paper sponsored by Committee on Transportation Supply Analysis.