White Paper

# Controlling Order Circuity in Pickup and Delivery Problems

**Douglas A. Popken***

**Systems View,** 9139 Roadrunner St., Highlands Ranch, CO 80129

## 5/18/2005

***Email:** dpopken@systemsview.com, **Phone:** 303-683-0552, **Fax:** 775-256-5547

**Abstract –** This paper examines the effects of constraining order circuity during the course of a dynamic pickup and delivery routing algorithm. Order circuity, the ratio of actual to direct shipping distance, affects shipping performance metrics such as transit time and vehicle utilization. The paper begins by exploring the tradeoff between utilization and order circuity, and then describes the application of circuity controls within an algorithm implemented for automobile shipping. We found that circuity controls are more effective than standard least-cost insertion heuristics alone in producing efficient route solutions with lower frequency of anomalous order routings.

## 1. Introduction

In pickup and delivery routing problems, a set of feasible routes must be developed to satisfy transportation requests between origin-destination pairs. The routes are covered by a fleet of vehicles with attributes such as capacity (perhaps multiple types), maximum route length, and maximum route duration. The transportation requests (hereafter referred to as *orders*) may also have a number of relevant attributes, including pickup and/or delivery time windows, size (units, weight, volume, etc), and priority.

Pickup and delivery problems find application in point-to-point transportation systems such as household goods moving, personal automobile moving, full-truck load shipping, and one-time commercial provisioning projects. They can also be applied to cargo shipping, airline scheduling, and on-demand transit services, i.e. the “Dial-A-Ride” problem. Many of these applications require the vehicles to be responsive to new requests even as they are servicing existing requests. Thus, these routes must be *dynamic*, whereby existing routes become part of the starting solution for a new problem that accounts for new orders and updated vehicle status. The vehicle status can include location, remaining capacities, and distance traveled.

The general pickup and delivery problem is difficult to solve for practically sized problems. A variety of solution methods have been proposed, most involving a combination of traditional optimization, insertion based constructive heuristics, and meta-heuristics such as tabu search (see, for example, Savelsbergh and Sol, 1995; Nanry and Barnes, 2000; Yang et al., 2004). However, one of the difficulties in formulating effective heuristic search strategies has been the lack of a reliable measure of closeness (Savelsbergh and Sol, 1995). Geographical closeness of stops is rarely a good indicator since each demand involves two stops, only one of which may be “close to” a stop for another demand. This makes it more important to develop simple measures that can quickly evaluate the quality of potential route assignments.

This paper discusses how rules based on easily computed order circuity measures improve an insertion-based solution heuristic used to solve general pickup and delivery problems. The Methodology section first defines the basic circuity measure and contrasts it to a related utilization measure. An example helps explain how placing too much emphasis on short-term utilization can reduce route efficiency to the detriment of long-term customer service. The Application section next provides a description of the auto transport application that motivated this research, and the solution algorithm in which the circuity measure operates. The Test Setup section compares the performance of three variations of order circuity routing rules incorporated within the solution algorithm. We conclude with recommendations and outline the contributions made by this research.

## 2. Methodology

Order circuity is defined as the ratio of the actual distance an order travels on a vehicle to the direct travel distance that would occur without making any intermediate stops. The correlation between circuity and transit time has been recognized in many areas of transportation network analysis and design, including LTL networks (Taylor, et al., 1995) and airline scheduling (Mathaisel, 1997). In Potvin and Rosseau (1992) circuity measures are used in the solution algorithm itself to cluster customers in an initial assignment phase of a dial-a-ride problem. Controlling order circuity is largely seen as a way to reduce customer transit times.

The quality of a single solution instance to a *dynamic* pickup and delivery problem can be readily computed with respect to operations based measures such as order circuity, and vehicle utilization. Similarly, we can easily compute some customer service based metrics such as lateness and transit time for those orders routed to their destination (in dynamic problems some orders may be deferred to later solutions). Other types of customer service based metrics, such as vehicle availability and order fulfillment lead time may not be relevant or computable from a single problem solution instance. However, as we will show, these metrics are directly affected by the easily-computed order circuity. These effects are best understood by first examining the relationship between order circuity and vehicle utilization.

In operations management, *utilization* generally refers to the proportion of system capacity being used, ranging from 0 to 1.0. In shipping applications, we can measure the system capacity in terms of order-miles. If vehicles are on average ½ full of orders as they travel their routes, then the average system utilization would be approximately 0.50. Intuitively, we like to have high values for utilization, as this appears to indicate that we are making good use of available transportation resources. However, pickup and delivery shipping is no different than many other industries, where a single-minded focus on high utilization can be detrimental; we need to make sure that capacity is also being used effectively. To define average vehicle utilization more formally, let:

$d^a_i$ = the actual travel distance (miles) of order i
$D_k$ = the total travel distance (miles) of vehicle/route k
$P_k$ = the static capacity of vehicle/route k in orders
$O_k$ = the set of orders serviced by vehicle/route k

The utilization, $U_k$, of vehicle/route k is then:

$$U_k = \frac{\sum_{i \in O_k} d_i^a}{P_k D_k} \qquad (1)$$

An order uses the minimum amount of system capacity if the carrying vehicle proceeds directly from the order's origin to its destination. However, since a vehicle may make a number of intermediate stops to service other requests, a given order may take a relatively circuitous route as it travels to its final destination. We can measure the *circuity* of the travel for that order by taking the ratio: (actual travel distance / direct travel distance). This measure takes on the value of 1.0 for a perfectly efficient (direct) route and larger values for more circuitous routes. Since a single vehicle route will likely be relatively efficient for some requests and less so for others, the circuity measure is most useful when averaged over a large number of orders. However, this implies a need to weight the order circuity values by order transit distance since the effects of high circuity are greater as the order transit distance increases in proportion to the total route distance. On a route of 12,000 miles we are likely to be more concerned with an order causing 1000 miles of extra travel than one causing 10 miles of extra travel even if the latter has higher circuity. (Assuming all delivery constraints can be satisfied in either case. Shorter transit distances are often associated with tighter pickup/delivery time windows.) We define the average distance-weighted circuity of a vehicle/route, $C_k$, as:

$$C_k = \frac{\sum_{i \in O_k} d_i^a c_i}{\sum_{i \in O_k} d_i^a} \quad \text{where:} \qquad (2)$$

$c_i = d^a_i / d^0_i$ = the *circuity* of order i, and
$d^0_i$ = the direct travel distance from the origin to the destination of order i.

Another way of looking at circuity is as a form of “excess utilization”. Suppose that an order contributes .18 to the average utilization of a vehicle on a given route, while the circuity of that order is 1.50. If the order circuity was instead 1.0, (i.e., $d^a_i$ was 2/3 of its current value) then, assuming the same total travel distance, the order would have contributed only (2/3 x .18) = .12 to the utilization. The difference, .18 - .12 = .06, is the *excess utilization.* The excess utilization is a form of opportunity cost measurable in order-miles. Potentially, the nonproductive order-miles could have been employed to move additional orders on the same vehicle had the route been more efficient. Orders that might otherwise have been serviced by a given vehicle may now have to wait for a new vehicle to become available. Accordingly, this has a direct impact on the metrics of vehicle availability and order fulfillment lead-time mentioned above.

The best situation is routes that have an average circuity and utilization both near 1.0. This can sometimes be achieved on routes between terminals with high inter-terminal order flows. But it is not always the case that high vehicle utilization implies a well-run system. While in general, vehicles *should* be well utilized; high utilization should not be pursued at the expense of higher circuity, as illustrated in the following example.

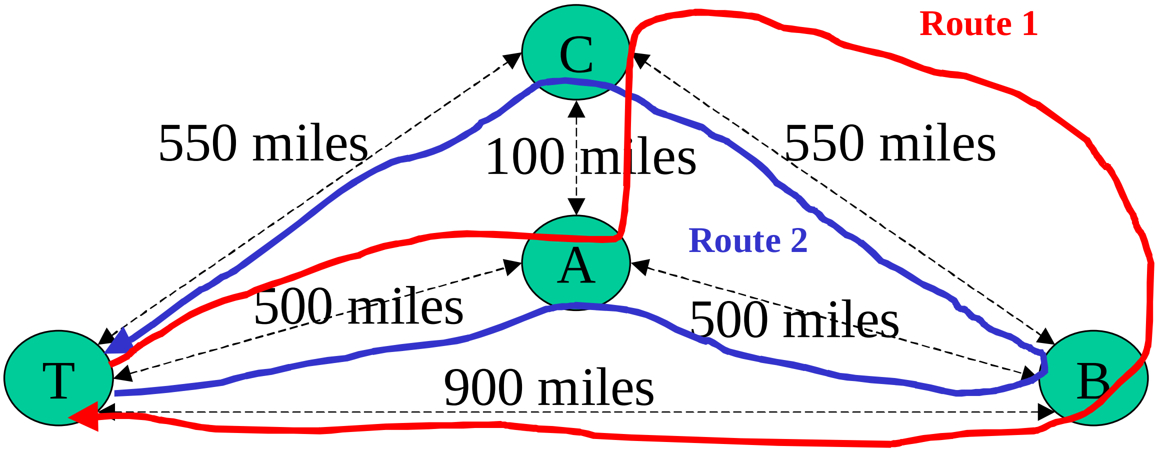

**Figure 1 – Two Possible Routes with Different Utilization and Circuity**

Figure 1 illustrates a situation where two alternative vehicle routes are being considered. Assume that the vehicle has a capacity of 10 orders and a maximum travel distance of 2100 miles. Suppose that we must satisfy the following transportation requests:

A→B: 5 orders
C→T: 5 orders

and that we must start and end at Terminal T (the home terminal for the vehicle). Route 1 (Red) takes the path T→A→C→B→T with a length of 2050 miles while Route 2 (Blue) takes the path T→A→B→C→T with a length of 2100 miles. The utilization and circuity for the routes are computed below:

<u>Route 1</u>

$$Avg.Utilization = \frac{(5\times 650)+(5\times 1450)}{10\times 2100} = \frac{10500}{21000} = .50$$

(Note that we have used 2100 in the denominator rather than 2050. The reason is that the 50 additional miles that were not traveled in Route 1 are a form of unused capacity.)

$$Circuity_{A\to B} = \frac{650}{500} = 1.30 \qquad Circuity_{C\to T} = \frac{1450}{550} = 2.64$$

$$Avg.Circuity = \frac{(5\times 650\times \frac{650}{500}) + (5\times 1450\times \frac{1450}{550})}{(5\times 650 + 5\times 1450)} = 2.22$$

Route 2

$$Avg.Utilization = \frac{(5\times 500)+(5\times 550)}{10\times 2100} = \frac{5250}{21000} = .25$$

$$Circuity_{A\rightarrow B} = \frac{500}{500} = 1.00 \qquad Circuity_{C\rightarrow T} = \frac{550}{550} = 1.00$$

$$Avg.Circuity = \frac{(5\times 500\times \frac{500}{500})+(5\times 550\times \frac{550}{550})}{(5\times 500+5\times 550)} = \frac{5250}{5250} = 1.0$$

Route 1 is shorter and exhibits a high average utilization while Route 2 has the lowest (best) weighted average circuity. A naïve observer would prefer the Route 1 outcome since the vehicle appears to be better utilized and traverses a slightly shorter route as well. However, the average circuity of 2.22 for Route 1 indicates that we have significant excess utilization:

$$ExcessUtilization = \frac{(5\times 150)+(5\times 900)}{(10\times 2100)} = .25$$

We see that the excess utilization is also equal to the difference in the average utilizations. Equivalently, Route 2 has .25 x 10 x 2100 = 5250 more order-miles of capacity available than does Route 1. If the routing problem were static, this wouldn't make any difference since both routes cover all orders. However, in dynamic situations, new demands present themselves over time. Route 2, with the lower circuity, would have a higher likelihood of accommodating additional orders within the constraints of the system.

We have seen that circuity affects both transit time and the ability to service potential *future* orders. By reducing circuity, we increase our ability to pack orders into routes, freeing up order-miles for future orders. This fact led us to believe that controlling circuity during the course of a sequential assignment algorithm could also provide superior solutions. In particular, more orders could be serviced by the same vehicle fleet. Another practical benefit we sought was the elimination of occasional routing "anomalies" in a solution, where vehicles appeared to go well out of their way to service orders that were neither high priority nor had imminent due dates. Although the anomalous looking route may provide the best (lowest penalty-weighted fleet travel distance) way to service an order, we apparently needed additional control over which order assignments *should* be made versus those that *could* be made. If the additional circuity is too great, perhaps the order should be deferred until a later time when the system configuration is more favorable.

## 3. Application

We developed a pickup and delivery solution algorithm for use by Dependable Auto Shippers, a provider of personal automobile shipping services, based in Mesquite, TX.. Their transportation network consists of approximately 85 terminals located across the United States. Customers may also request a “direct” pickup and/or drop-off from a non-terminal location (e.g. home or office). Most terminals operate 24 hours per day, implying no time window on pickups or drop-offs other than those imposed by the orders themselves. However, *direct* pickups or drop-offs may have time windows. The problem is *dynamic*, since routes are modified over time as new orders arrive. For each problem instance, all orders that have been delivered are removed from the problem, the truck is updated to its current state (position, miles traveled, loads, etc) and new orders are considered. This implies that there is not a fixed set of orders that must be fulfilled in a given solution.

Approximately 50-90 trucks are active at any one time. Trucks typically travel an average of 12,000 miles over two to three weeks on routes that begin and end at their home terminal. These high travel rates are made possible by the use of driver teams and by the fact that terminals are open 24 hours per day. Trucks are characterized by:

a) Current Geo-coordinates
b) Miles Traveled
c) Current Orders Assigned or Loaded
d) Route Start/End Time
e) Home Terminal
f) Maximum Route Length
g) Maximum Miles/Day
h) Maximum Load Weight
i) Maximum Vehicles
j) Maximum Oversized Vehicles
k) Maximum Top Loaded Vehicles
l) Maximum Height of Top Loaded Vehicles
m) Maximum Bottom Loaded Vehicles
n) Maximum Height of Bottom Loaded Vehicles
o) Enclosed: Yes/No

Hundreds of new orders are received every day, and are characterized by:

a) Pickup location
b) Min Pickup Time
c) Max Pickup Time (for direct pickups)
d) Drop-off location
e) Min Drop-off Time (for direct drop-offs)
f) Max Drop-off Time
g) Weight (lbs.)
h) Truck Type: Enclosed/NonEnclosed, Either

i) Position on Truck: Top Rack, Bottom Rack, Either
j) Height: may determine Position on Truck and "Oversize" Status
k) Priority
l) Assignment Status: Unassigned, Assigned (Not Loaded), Assigned (Loaded)
m) Assigned Truck (if Assigned)
n) Transship Eligibility: Yes/No

We use a "rolling horizon" approach (Psaraftis, 1988) to determine which orders should even be considered for assignment in the current solution. The dynamics of the problem are such that many or all of the trucks could be loaded to capacity (or near) if the entire set of outstanding orders were allowed to enter the problem. However, since orders are distinguished by priority and due date, allowing *all* orders into the solution would likely require foregoing new orders in subsequent runs with higher priority and more imminent due date. We use a simple preprocessing algorithm to implement the rolling horizon approach to order filtering. Let $t^0$ be the current clock time, and $t_i^+$ be the latest clock time that order ***i*** can be picked up so that it will not be late to its destination, assuming that the order travels directly from origin to destination. Another definition for $t_i^+$ is a "ship no later than" date for order ***i***. We define solution parameters $\Delta t_{min}$ and $\Delta t_{max}$ with:

"Must Ship" Cutoff Time = $\Delta t_{min} + t^0$
"May Ship" Cutoff Time = $\Delta t_{max} + t^0$

If $t_i^+$ is less than or equal to the "May Ship Cutoff Time", order ***i*** may be considered for assignment in the current run. If $t_i^+$ is also less than or equal to the "Must Ship Cutoff Time", then order ***i*** is given the highest assignment priority. The latter ensures that low priority orders are always eventually assigned and are not continually "bumped" by the high priority orders.

To control the problem size and provide solutions in a reasonable amount of time, we allow a maximum of 10 stops per route in a given solution. The solution provides a current routing from the current truck position to its home terminal. Although the algorithm creates a full sequence of stops to the final truck destination, these stops should only be a considered a "skeleton" for a future solution. Each problem instance removes stops already completed from previous solutions so that room is made for new stops to be added. Additional orders can also be added in the future to the current stop sequence.

Our general objective is to satisfy as many assignable orders as possible, in as efficient a manner as possible, while respecting hard constraints and penalizing deviation from soft constraints. We use penalty weighted truck mileage as our objective. Penalties accrue for the miles each truck must travel beyond its maximum, and for projected lateness of orders.

The assignment algorithm is made up of a controlled series of order insertions and swaps in a parallel search of trucks. In a parallel search, each order is operated upon separately, and assigned to its best feasible truck and best feasible pickup/drop-off sequence within that truck's route.

The dispatching staff at DAS would sometimes observe that the routing application would produce a few routes that seemed overly-circuitous, or for some reason did not “make sense” to them. There would usually be some obvious reason for this, for example, particularly late orders that had to be serviced immediately at the expense of significant additional truck travel miles. In other cases, there seemed to be no apparent rationale. We describe a recent example that arose during testing the algorithm on actual live data. At the time of the run, the truck is currently located in Kansas, and loaded with 8 orders heading west (destinations in parenthesis): 763680 (UT), 765060 (UT), 764600 (UT), 763956 (ID), 765442 (ID), 765882 (WA), 758222 (WA), 764423 (WA). The algorithm produced the following route for this particular truck:

| **Load** | **Unload** |
|---|---|
| **Stop 1: MO1 (ETA: 3/26 11 pm)** | |
| 766800 | |
| **Stop 2: IL (ETA: 3/27 6 am)** | |
| | 766800 (Due 3/31) |
| **Stop 3: MI (ETA: 3/27 7pm)** | |
| 766089 | |
| **Stop 4: UT (ETA: 3/29 6am)** | |
| 766337 | 763680 (L[a]) (Due 3/30) |
| | 765060 (L) (Due 3/23 - **Late**) |
| | 764600 (L) (Due 3/23 – **Late**) |
| **Stop 5: MO10 (ETA: 3/30 6pm)** | |
| 765059 | |
| **Stop 6: ID (ETA: 4/1 12 pm)** | |
| | 763956 (L) (Due 3/24) |
| | 765442 (L) (Due 3/30) |
| **Stop 7: WA (ETA: 4/1 11pm)** | |
| | 765882 (L) (Due 3/29 – **Late**) |
| | 766089 (Due 4/2) |
| | 765059 (Due 3/30 – **Late**) |
| | 758222 (L) (Due 3/30 – **Late**) |
| | 764423 (L) (Due 3/23 – **Late**) |
| **Stop 8: CO (ETA: 4/3 4am)** | |
| 765716 | |
| **Stop 9: FL (ETA: 4/5 4am)** | |
| | 765716 (Due 4/2 – **Late**) |
| | 766337 (Due 4/1 – **Late**) |
| **Stop 10: TX (ETA: 4/6 7am)** | |

a. “ (L)” indicates an order that was already loaded at the time of the run.

**Table 1. An Anomalous Route**

What seemed puzzling was how the truck seemed to reverse course and head back east to make stops in Missouri, Illinois, and Michigan. It then headed west to Utah, but then returned again to Missouri (at a different terminal) before heading back out to deliver orders in Idaho and Washington. Based on the sequential logic of the algorithm, the most likely explanation is that the algorithm first assigned order 766089, originating in

Michigan (Stop 3), because it had a destination in common (Washington- Stop 7) with several already loaded orders. Since Missouri and Illinois were on the way, it could also easily service order 766800 (Stops 1 and 2). At this point, the orders going to Utah (Stop 4) were running late, so they had to be taken care of next. Next, although the truck was eventually heading to Idaho and Washington, there was still time to go back to Missouri (Stop 5 - a different Missouri terminal) to pick up another order going to Washington. The final delivery in Florida (Stop 9) results in a circuitous route for order 766337 (picked up at Stop 4, UT), but not for order 765716 (picked up at Stop 8, CO).

While the majority of routes produced by the algorithm were well behaved, there were perhaps two or three trucks per solution that had somewhat “anomalous” behavior similar to that described in the extreme example above. It was hoped that adding some form of circuity control to the algorithm would prevent this. Because the assignment algorithm is a parallel search (rather than one truck at a time), it is not possible to make “before” and “after” comparisons of a single route to determine change effects. One small change in the algorithm affects the future chain of events so that a completely different assignment is produced. Instead we have to look at the performance measures for the entire solution. One of these measures, the number of orders with anomalous (high circuity) routings, is examined in our tests below.

## 4. Test Setup

We developed an “Instance Generator” to generate randomized input data sets to the algorithm. System parameters were set to the specific size and physical characteristics of the truck fleet and terminal network of DAS. Orders were generated so that the distribution of origin-destination pairs is in accordance with the empirical distribution over a year of activity. Similarly, order classes and priorities are randomly generated according to historical distributions. Trucks are initialized to a position at the geo-coordinates of a randomly selected terminal. Travel distances are computed using a great-circle distance between geo-coordinates multiplied by a road distance factor of 1.17 (Ballou, 1992). (The full software implementation uses a distance table generated by PCMiler®, a commercial travel management software that computes actual road mileages).

We considered the effects of adding circuity considerations to the order assignment process. We used one of three logical tests when considering an order, $\boldsymbol{i}$, for assignment to truck route, $\boldsymbol{k}$. Assume that $c_i^*$ and $C_k^*$ are the individual order and average route circuities, respectively, resulting from the “best” insertion of order $\boldsymbol{i}$ into truck route $\boldsymbol{k}$. The three rules are:

Rule 1 (individual): Accept assignment of $\boldsymbol{i}$ to $\boldsymbol{k}$ if $c_i^* < C_T$
Rule 2 (route average): Accept assignment of $\boldsymbol{i}$ to $\boldsymbol{k}$ if $C_k^* < C_T$ OR $C_k^* < C_k'$
Rule 3 (combined): Accept assignment of $\boldsymbol{i}$ to $\boldsymbol{k}$ if $c_i^* < C_T$ AND ($C_k^* < C_T$ OR $C_k^* < C_k'$)

where $C_T$ is a threshold circuity level that was varied during testing, and $C_k'$ is the average route circuity *prior* to insertion. The purpose of the latter is to handle situations where routes are initialized with a high circuity due to orders assigned prior to the algorithm. Rule (3) is the strictest, requiring the order to pass rules (1) AND (2) to be accepted for assignment. During testing we varied $C_T$ over the set {1.5, 2.0, ∞} to get an idea of the magnitude of the circuity effect. The rationale for using the *average* route circuity, $C_k$, when evaluating a new order is that adding a new order to a route can change the circuity for existing orders. As an example, observe order 765059 in the "sample anomaly" above. The resulting circuity for that order is quite low; it travels from Missouri to Washington with one intermediate stop in Idaho. The assignment would likely be accepted by any Rule 1 type test. However, adding that order to the route causes a significant increase in the circuity of the other currently loaded orders. Rule 2 is meant to reduce the frequency and severity of those situations. On the other hand, use of average circuity could allow new orders into the route that have relatively high circuity provided the average remains below a threshold.

To determine the influence of order intensity, we varied the number of initial orders over the set {1000, 1500, 2000, 2500}. Note that the number of *initial* orders is not the same as *assignable* orders. Any order with the "must ship date", $t_i^+$, that is greater than the user defined limit, L, will be excluded. But the number of assignable orders will increase by approximately the same proportion as the initial orders.

## 5. Results

| Order Intensity | Circuity Measure | Circuity Threshold | Avg. Circuity[a] | Avg. Utilization[a] | Avg. Orders Carried[a] | Cost / Order Carried[a] |
|---|---|---|---|---|---|---|
| 1000 | $c_i$ | 1.5 | [1.11; .015] | [.20; .01] | [401.84; 25.87] | [853.60; 98.27] |
| 1000 | $c_i$ | 2.0 | [1.19; .018] | [.24; .02] | [427.48; 27.12] | [741.66; 75.69] |
| 1000 | $C_k$ | 1.5 | [1.16; .011] | [.25; .01] | [429.40; 4.09] | [731.44;121.49] |
| 1000 | $C_k$ | 2.0 | [1.24; .022] | [.26; .02] | [428.24; 25.70] | [706.44; 83.38] |
| 1000 | $c_i+ C_k$ | 1.5 | [1.11; .008] | [.21; .01] | [412.40; 28.15] | [808.06; 71.85] |
| 1000 | $c_i+ C_k$ | 2.0 | [1.19; .021] | [.26; .02] | [424.88; 30.23] | [752.33; 83.25] |
| 1000 | NA | None | [1.35; .022] | [.28; .01] | [431.48; 33.34] | [733.30; 81.77] |
| 1500 | $c_i$ | 1.5 | [1.11; .026] | [.24; .01] | [506.16; 33.58] | [796.73; 61.22] |
| 1500 | $c_i$ | 2.0 | [1.18; .020] | [.28; .02] | [528.92; 36.26] | [728.21; 64.51] |
| 1500 | $C_k$ | 1.5 | [1.17; .013] | [.28; .02] | [532.08; 35.37] | [712.00; 80.72] |
| 1500 | $C_k$ | 2.0 | [1.24; .019] | [.30; .01] | [541.08; 31.08] | [732.84; 82.14] |
| 1500 | $c_i+ C_k$ | 1.5 | [1.10; .011] | [.25; .02] | [512.28; 31.94] | [799.00; 84.11] |
| 1500 | $c_i+ C_k$ | 2.0 | [1.18; .018] | [.28; .02] | [523.56; 22.14] | [719.26; 84.78] |
| 1500 | NA | None | [1.34; .041] | [.32; .02] | [527.52; 27.15] | [693.62; 78.86] |
| 2000 | $c_i$ | 1.5 | [1.10; .015] | [.27; .01] | [600.68; 31.31] | [744.55; 76.59] |
| 2000 | $c_i$ | 2.0 | [1.18; .021] | [.31; .02] | [609.36; 33.12] | [707.89; 66.21] |
| 2000 | $C_k$ | 1.5 | [1.17; .013] | [.31; .01] | [619.00; 41.09] | [708.62; 74.35] |
| 2000 | $C_k$ | 2.0 | [1.25; .023] | [.33; .02] | [614.40; 42.63] | [710.40; 74.35] |
| 2000 | $c_i+ C_k$ | 1.5 | [1.10; .011] | [.28; .01] | [599.68; 33.33] | [707.05; 70.97] |
| 2000 | $c_i+ C_k$ | 2.0 | [1.17; .016] | [.31; .02] | [604.16; 27.01] | [696.40; 50.68] |
| 2000 | NA | None | [1.36; .040] | [.36; .01] | [610.68; 46.64] | [721.77; 91.09] |
| 2500 | $c_i$ | 1.5 | [1.09; .013] | [.29; .02] | [651.16; 43.02] | [723.87; 63.51] |
| 2500 | $c_i$ | 2.0 | [1.16; .015] | [.33; .01] | [674.40; 40.36] | [686.17; 63.89] |
| 2500 | $C_k$ | 1.5 | [1.16; .013] | [.33; .01] | [670.08; 40.89] | [684.85; 62.85] |
| 2500 | $C_k$ | 2.0 | [1.25; .019] | [.36; .02] | [677.00; 31.84] | [666.06; 76.35] |
| 2500 | $c_i + C_k$ | 1.5 | [1.10; .011] | [.30; .02] | [665.08; 41.49] | [668.66; 71.32] |
| 2500 | $c_i + C_k$ | 2.0 | [1.16; .016] | [.33; .02] | [669.84; 45.75] | [683.84; 92.95] |
| 2500 | NA | None | [1.36; .033] | [.37; .02] | [672.32; 40.69] | [689.91; 71.76] |

a. [Mean; Std. Dev]

**Table 2. Average Route Characteristics by Circuity Measure and Order Intensity**

The results are shown in Table 2. We see that as the circuity threshold is decreased: (1) average circuity decreases significantly, (2) average orders carried decreases slightly in general, but not significantly, and (3) cost /order carried increases slightly in general, but not significantly. However, the percentage decrease in utilization is greater than any decreases in orders carried – thus the routes use capacity more efficiently. These effects occur at each level of order intensity. Any of the three circuity controls will significantly reduce circuity. However, $c_i$ and $c_i + C_k$ appear to provide slightly better performance than $C_k$ alone. To better distinguish their relative performance, we examine the distribution of individual order circuitries occurring over 25 runs of the simulation. We set the order intensity at 2500 and the circuity threshold to 1.5 to produce the distribution in Table 3 below.

| Measure | Circuitry Range | | | | | | | | | | Total |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | [1,2) | [2,3) | [3,4) | [4,5) | [5,6) | [6,7) | [7,8) | [8,9) | [9,10) | [10,inf) | |
| **C** | 16255 | 431 | 58 | 7 | 0 | 0 | 0 | 0 | 0 | 1 | 16752 |
| 1-F | 1 | 0.029668 | 0.00394 | 0.000478 | 5.97E-05 | 5.97E-05 | 5.97E-05 | 5.97E-05 | 5.97E-05 | 5.96944E-05 | |
| **c** | 16228 | 34 | 7 | 6 | 1 | 1 | 0 | 0 | 0 | 2 | 16279 |
| 1-F | 1 | 0.003133 | 0.001044 | 0.000614 | 0.000246 | 0.000184 | 0.000123 | 0.000123 | 0.000123 | 0.000122858 | |
| **C+c** | 16593 | 26 | 6 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 16627 |
| 1-F | 1 | 0.002045 | 0.000481 | 0.00012 | 6.01E-05 | 0 | 0 | 0 | 0 | 0 | |
| **None** | 14833 | 1367 | 376 | 132 | 53 | 27 | 10 | 7 | 2 | 1 | 16808 |
| 1-F | 1 | 0.117504 | 0.036173 | 0.013803 | 0.00595 | 0.002796 | 0.00119 | 0.000595 | 0.000178 | 5.94955E-05 | |

**Table 3 – Circuity Distribution (count and proportion ≥ range)**

We see that overall, the results are very similar. However, the combined control, $c_i$+ $C_k$ was better at reducing anomalous values in the right hand tail of the distribution. For example, the proportions of the values ≥ 4 are: .0005 for $C_k$, .0006 for $c_i$, and .0001 for $c_i$+ $C_{k.}$ The latter control appears best at scrubbing out the high circuity order routings.

## 6. Conclusions

We examined three variations of circuity control – an individual order circuity control, an average route circuity control, and a combined control that uses both. In all cases, using a circuity control provides a more efficient utilization of vehicle capacity. Average circuity and utilization are both reduced significantly at the expense of only slightly fewer orders carried and slightly higher costs per order carried. We did not observe increases in the number of orders carried within a single run of the algorithm as anticipated. We believe that this is because our simulation generated uncorrelated orders, but in practice there tends to be small order clusters with identical origin-destination pairs. We might also have simply increased the order intensity above typical levels. However, since the algorithm is implemented in a dynamic setting, the lower utilization implies more assignments can be achieved in subsequent runs. Of the three circuity controls examined, the average route circuity control was overall slightly less effective than the other two.

The circuity controls also significantly reduced the number of orders with unusually high circuitries, which we referred to as "anomalies". Of the three circuity controls examined, the combined circuity control provided the best results. Since anomaly reduction is very important to dispatching staff, this effect combined with the overall results above, make the combined control superior.

An alternative approach to our use of circuity constraints would be an assignment algorithm in which insertions are evaluated on the basis of marginal additional *order-miles* versus the current *vehicle miles*. This would eliminate the need for an explicit constraint. We did not try this approach since in our situation it was important to be able to enforce the soft constraints with penalties whose units are vehicles miles. It is also not clear how effective such an algorithm would be at minimizing anomalies in the right hand side of the order circuity distributions. Such an evaluation would be an obvious extension of this research.

This paper's contribution to the literature is to provide a simple but effective control mechanism for improving routes in general pickup and delivery problems. The results can be applied to a variety of pickup and delivery situations, including household goods moving and "Dial-a-Ride" type people moving systems. Using route rejection rules based on the circuity measure provides a way to quickly eliminate nonsensical routes, and improve routes overall, within a large solution space. We have focused on measurable benefits such as higher long-term truck space availability. Another less measurable benefit is greater customer satisfaction. Routes with high circuity would be intolerable for passenger systems. However, this is also becoming more the case for freight systems, where customers increasingly have visibility into the progress of their orders as they move from origin to destination (this is true for DAS). If customers see their order taking a very indirect route, perhaps even backtracking at times, they would tend to question the competence of the carrier. The same is true for the vehicle drivers, who would feel they are not being utilized effectively.

## Acknowledgements

The ideas in this article benefited greatly from the input and experience of DAS operations and IT personnel including Eric Weinkam, Andy Duong, and especially, Roque Moreira.

## References

Ballou R.H. 1992. *Business Logistics Management*. Prentice-Hall, Englewood Cliffs, NJ.

Mathaisel, D. 1997. Decision support for airline schedule planning. *Journal of Combinatorial Optimization*, **1**(3), pp. 251-75.

Nanry, W.P., Barnes, J.W. 2000. Solving the pickup and delivery problem with time windows using reactive tabu search. *Transportation Research*, **34B**, pp. 107-121.

Potvin, J.Y., Rousseau, J.M. 1992. Constraint-directed search for the advance request dial-a-ride problem with service quality constraints. In: *Computer Science and Operations Research: New Developments in their Interfaces*, Balci, O., Shrada, R., Zenios, S.A. (Eds.), Pergamon Press, Oxford, England, pp. 457-474.

Psaraftis, H.N. 1988. Dynamic vehicle routing problems. In: *Vehicle Routing: Methods and Studies*, Golden, B.L., and Assad, A.A. (Eds.), North-Holland, Amsterdam.

Savelsbergh, M.W.P., Sol, M. 1995. The general pickup and delivery problem. *Transportation Science*, **29**(1), pp. 17-29.

Taylor G.D., Harit, S., English, J.R., Whicker, G.L. 1995. Hub and spoke networks in truckload trucking: configuration, testing, and operational concerns. *The Logistics and Transportation Review*, **31**(3), pp. 209-237.

Yang, J., Jaillet, P., Mahmassani, H. 2004. Real-Time multivehicle truckload pickup and delivery problems. *Transportation Science*, **38**(2), pp. 135-148.