Modeling Techniques for Periodic Vehicle Routing Problems

Peter Francis and Karen Smilowitz*

Department of Industrial Engineering and Management Sciences

Northwestern University, Evanston, IL, USA 60208-3119

January 16, 2006

Abstract

This paper presents a continuous approximation model for the Period Vehicle Routing Problem with Service Choice (PVRP-SC). The PVRP-SC is a variant of the Period Vehicle Routing Problem in which the visit frequency to nodes is a decision of the model. This variation can result in more efficient vehicle tours and/or greater service benefit to customers. The continuous approximation model can facilitate strategic and tactical planning of periodic distribution systems and evaluate the value of service choice. Further, results from the continuous model can provide guidelines for constructing solutions to the discrete PVRP-SC.

Keywords: Vehicle routing; continuous approximation models; vehicle routing models

1 Introduction

This paper develops continuous approximation techniques to study periodic vehicle routing problems with service choice where service is defined by the frequency of visits to nodes. Applications arise in courier services, elevator maintenance and repair (Blakely et al., 2003), the collection of waste (Russell and Igo, 1979) and the delivery of interlibrary loan material (Francis et al., 2006).

This problem is a variant of the Period Vehicle Routing Problem (PVRP) in which routes are designed for a fixed fleet of capacitated vehicles each day of a $t-$day period to visit customers exactly

*Corresponding author. Email address: ksmilowitz@northwestern.edu; Tel.: (847) 491 - 4693; Fax.: (847) 491-8005

a pre-set number of times; see (Beltrami and Bodin, 1974; Christofides and Beasley, 1984; Russell and Igo, 1979). However, Francis et al. (2006) note that operational efficiencies and increased service benefit may be gained by allowing customers to be visited more often than their minimum required frequencies as illustrated in the example in Figure 1. In the example, nodes are assigned a preset visit frequency from the following options: daily, Mon-Wed-Fri or Tue-Thr. If Node 2 (visited twice a week) is close to Node 1 (visited three times a week), then serving Node 2 on the same schedule as Node 1 may result in considerable routing cost savings. This gain in operational efficiency is possible only when nodes may be served with greater frequency than required.

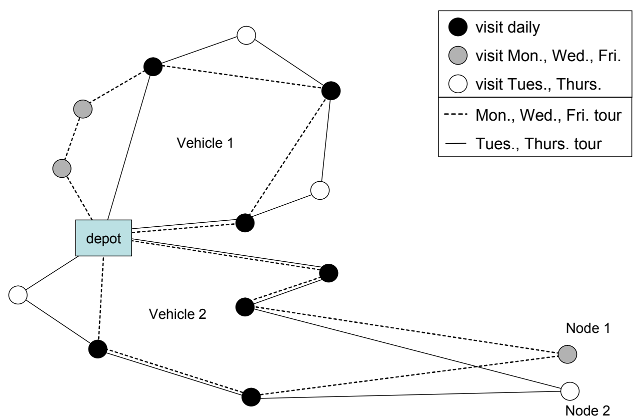

Figure 1: Efficiency improvements possible with service choice (Francis et al., 2006)

Francis et al. (2006) introduce the Period Vehicle Routing Problem with Service Choice (PVRP-SC) which allows service levels to be determined endogenously. The PVRP-SC is defined as follows:

Given: A set of nodes with known demand and minimum visit frequency requiring service over the planning period; a fleet of capacitated vehicles; a set of service schedules with headways and service benefits; and a network with travel times.

Find: An assignment of nodes to service schedules and a set of vehicle routes for each day of the planning period that minimizes the total routing cost incurred net of the service benefit accrued.

Francis et al. (2006) develop an integer programming formulation of the PVRP-SC with exact and heuristic solution methods. Due to the computational complexity of the problem, solutions to the discrete PVRP-SC are limited by instance size. Alternatively, approximate solutions for such instances may be obtained with continuous approximation models, yet the use of continuous approximation models for periodic routing problems has been limited. Daganzo (1987) presents continuous approximation modeling techniques for single period distribution problems with varying fixed service requirements. Smilowitz and Daganzo (2004) develop continuous approximation models for integrated distribution network design with two service levels. These references suggest that continuous approximations can be powerful tools for strategic and tactical decisions when service choice exists. In continuous approximation models, aggregated data are used in place of more detailed inputs. Aggregating data in this manner smooths minor dynamic and stochastic variations in input parameters which are less critical in strategic planning. Continuous approximation models can produce results for large problem instances in less time compared with discrete models. This is particularly useful if the system designer would like to experiment with multiple settings of the input parameters such as demand node distribution and service characteristics. Further, the simplicity of continuous approximation models can facilitate the development of managerial insights for system planning.

This paper presents the first continuous approximation model for the PVRP-SC, which can also be used for the PVRP as a special case. The model is used in the strategic analysis of the benefits of service choice and the sensitivity of these benefits to various parameters. Results obtained with the model also answer tactical questions relating to the service mix of customers and vehicle fleet planning. This research provides practitioners with a tool to analyze efficiencies in distribution operations arising from service choice, without requiring extensive computations and detailed data collection typical of discrete models for periodic vehicle routing problems.

Section 2 reviews the discrete formulation of the PVRP-SC from Francis et al. (2006) and introduces the continuous approximation model. Section 3 describes the solution method for the continuous approximation model. Section 4 discusses numerical studies, and Section 5 concludes with key insights and areas of future research.

2 Models of the PVRP-SC

We present the discrete formulations of the PVRP-SC from Francis et al. (2006) in Section 2.1 and the continuous formulation in Section 2.2.

In the PVRP-SC, customers are visited a preset number of times over the period with a schedule that is chosen from a menu of schedule options. Let S denote this menu of schedules, and D denote the set of days in the period. The parameter a_{sd} links schedules to days, where $a_{sd} = 1$ if day $d \in D$ is in schedule $s \in S$ and $a_{sd} = 0$ otherwise. Each schedule $s \in S$ has an associated visit frequency γ^s measured by the number of days in the schedule: $\gamma^s = \sum_{d \in D} a_{sd}$. For a given schedule option s, the headway between visits is defined in terms of the visit frequency as $H^s = 1/\gamma^s$. Each schedule has an associated benefit α^s related to the monetary benefit of more frequent service which is assumed to be stationary over the time period.

2.1 Discrete formulation of the operational/tactical PVRP-SC

The discrete formulation of the PVRP-SC is defined for a set of nodes, N_0, which consists of customers nodes, N, and a depot, $l = 0$, and a set of arcs connecting nodes, $A = \{(l, m) : l, m \in N_0\}$. Each customer node $l \in N$ has a known daily demand, W_l, and a minimum service frequency, F_l, measured in days per period. The demand accumulated between visits, w_l^s, is a function of the schedule $s \in S$ and the daily demand of the node. The stopping time at a node, τ_l^s, is a function of the frequency of the schedule since more items accumulate with less frequent service and, therefore, require more time to load/unload. Associated with each arc $(l, m) \in A$ is a known travel cost, c_{lm}. There is a set K of vehicles, each with capacity C.

The following allocation and routing variables define the solution to the discrete formulation.

$$y_{lk}^s = \begin{cases} 1 & \text{if node } l \in N \text{ is visited by vehicle } k \in K \text{ on schedule } s \in S \\ 0 & \text{otherwise} \end{cases}$$

$$x_{lmk}^d = \begin{cases} 1 & \text{if vehicle } k \in K \text{ traverses arc } (l, m) \in A \text{ on day } d \in D \\ 0 & \text{otherwise} \end{cases}$$

The discrete formulation for PVRP-SC developed in Francis et al. (2006) is:

$$\min Z = \sum_{k\in K}\left[\sum_{d\in D}\sum_{(l,m)\in A} c_{lm}x^d_{lmk} + \sum_{s\in S}\sum_{l\in N}\left(\gamma^s\tau^s_l - W_l\alpha^s\right)y^s_{lk}\right] \tag{1a}$$

subject to

$$\sum_{s\in S}\sum_{k\in K}\gamma^s y^s_{lk} \ge F_l \qquad \forall l\in N, \tag{1b}$$

$$\sum_{s\in S}\sum_{k\in K} y^s_{lk} \le 1 \qquad \forall l\in N \tag{1c}$$

$$\sum_{s\in S}\sum_{l\in N} w^s_l a_{sd} y^s_{lk} \le C \qquad \forall k\in K; d\in D \tag{1d}$$

$$\sum_{m\in N_0} x^d_{lmk} = \sum_{s\in S} a_{sd}y^s_{lk} \qquad \forall l\in N; k\in K; d\in D \tag{1e}$$

$$\sum_{m\in N_0} x^d_{lmk} = \sum_{m\in N_0} x^d_{mlk} \qquad \forall l\in N_0; k\in K; d\in D \tag{1f}$$

$$\sum_{l,m\in Q} x^d_{lmk} \le |Q|-1 \qquad \forall Q\subseteq N; k\in K; d\in D \tag{1g}$$

$$y^s_{lk}\in\{0,1\} \qquad \forall l\in N; k\in K; s\in S \tag{1h}$$

$$x^d_{lmk}\in\{0,1\} \qquad \forall (l,m)\in A; k\in K; d\in D \tag{1i}$$

The objective function (1a) balances arc travel times, stopping times and demand-weighted service benefit. Constraints (1b) enforce the minimum frequency of visits for each node. Constraints (1c) ensure that one schedule and one vehicle are chosen for each demand node. Constraints (1d) represent vehicle capacity constraints. Constraints (1e) link the x and y variables for the demand nodes. Constraints (1f) ensure flow conservation at each node. Constraints (1g) are the subtour elimination constraints and ensure that all tours contain a visit to the depot. Constraints (1h) and (1i) define the binary variables for allocation and routing, respectively.

2.2 Continuous approximation of the strategic PVRP-SC

In this section, we develop a continuous approximation model for the strategic PVRP-SC.

2.2.1 Continuous decision and data functions

In the continuous model, we approximate discrete variables and parameters with continuous functions. It is assumed that these approximating functions are smooth, continuous and vary slowly over the service region $\mathcal{R}$.

First, approximations replace exact data for node locations and demand volumes. Let $\delta^i(\mathbf{x})$ denote the spatial density of nodes with minimum service schedule $i \in S$ about a point $\mathbf{x}$, measured in nodes per unit area. For a subregion $\mathcal{A}$ of $\mathcal{R}$, the number of nodes in the region, $\mathcal{N}(\mathcal{A})$, is: $\mathcal{N}(\mathcal{A}) = \sum_{i \in S} \left(\int_{\mathbf{x} \in \mathcal{A}} \delta^i(\mathbf{x}) d\mathbf{x} \right)$. Demand density rates replace the exact demand volumes associated with customer nodes. Let $\lambda^i(\mathbf{x})$ denote the demand density rate about a point $\mathbf{x}$ of nodes with minimum service schedule $i \in S$, measured in items per unit time-area.

Continuous functions are introduced to describe service allocations and vehicle routes. Let $f^{is}(\mathbf{x})$ denote the fraction of nodes about a point $\mathbf{x}$ with minimum schedule i being served by service schedule s. Nodes cannot be served with a frequency lower than the minimum specified. The vehicle routes to serve the selected schedules are determined with two auxiliary decision functions. Let $\Delta^d(\mathbf{x})$ denote the spatial density of nodes about a point $\mathbf{x}$ to be visited on day d, measured in nodes per unit area, and $\Lambda^d(\mathbf{x})$ denote the demand density on day d, measured in demand per unit area. The effective density of nodes visited per unit area on day d is:

$$\Delta^d(\mathbf{x}) = \sum_{s \in S} a_{sd} \sum_{i \in S} \delta^i(\mathbf{x}) f^{is}(\mathbf{x}) \tag{2}$$

The effective demand density collected per unit area on day d is:

$$\Lambda^d(\mathbf{x}) = \sum_{s \in S} a_{sd} H^s \sum_{i \in S} \lambda^i(\mathbf{x}) f^{is}(\mathbf{x}) \tag{3}$$

The routing solutions are described further by two additional decision functions. Let $n^d(\mathbf{x})$ denote the number of stops on a route on day d and $v^d(\mathbf{x})$ denote the shipment size collected at a node.

2.2.2 Continuous cost model

Routing costs are based on VRP approximations in Daganzo (1999). Routing costs are divided into a linehaul cost from the depot to the vicinity of the nodes on the route and a detour cost to visit individual nodes.[1] Let the parameter $r(\mathbf{x})$ denote the distance from the depot to a point $\mathbf{x} \in \mathcal{R}$. The average distance between nodes is approximated by the inverse of the square root of node density, $(\Delta^d(\mathbf{x}))^{-1/2}$, and a metric-dependent constant, $\hat{k}$. The average cost per distance, $\bar{c}$, is the average of c_{lm} divided by the distance between nodes l and m over all $(l, m) \in A$. The stopping cost τ_l^s is decomposed into a fixed stopping cost, $\dot{\tau}$, and a variable cost, $\ddot{\tau}$, that increases with the demand accumulated at node l. The expression for routing cost per item on day $d \in D$ is as follows:

$$z^d(\mathbf{x}) = \frac{2r(\mathbf{x})\bar{c}}{n^d(\mathbf{x})v^d(\mathbf{x})} + \frac{\bar{c}\hat{k}(\Delta^d(\mathbf{x}))^{-1/2} + \dot{\tau} + \ddot{\tau}v^d(\mathbf{x})}{v^d(\mathbf{x})} \tag{4}$$

The first term represents the linehaul travel between the depot and nodes, $2r(\mathbf{x})\bar{c}$. This cost is prorated to all items in the vehicle, $n^d(\mathbf{x})v^d(\mathbf{x})$. The second term represents the cost of visiting demand nodes: the detour distance cost to reach the nodes, $\bar{c}\hat{k}(\Delta^d(\mathbf{x}))^{-1/2}$, and the stopping cost at the nodes, $\dot{\tau} + \ddot{\tau}v^d(\mathbf{x})$. This cost is prorated by items per stop, $v^d(\mathbf{x})$.

The routing costs per item are summed over all days and integrated over all demand to obtain the total cost over the planning period.

[1] In cases where the routes are adjacent to the depot, the linehaul component may be ignored.

$$\min Z_{\mathcal{R}} = \int_{\mathbf{x}\in\mathcal{R}} \left(\sum_{d\in D} \Lambda^d(\mathbf{x}) z^d(\mathbf{x}) - \sum_{s\in S} \alpha^s \sum_{i\in S} \lambda^i(\mathbf{x}) f^{is}(\mathbf{x}) \right) d\mathbf{x} \tag{5a}$$

subject to

$$n^d(\mathbf{x}) v^d(\mathbf{x}) \leq C \qquad \forall d \in D, \mathbf{x} \in \mathcal{R} \tag{5b}$$

$$v^d(\mathbf{x}) = \frac{\Lambda^d(\mathbf{x})}{\Delta^d(\mathbf{x})} \qquad \forall d \in D, \mathbf{x} \in \mathcal{R} \tag{5c}$$

$$\Delta^d(\mathbf{x}) = \sum_{s\in S} a_{sd} \sum_{i\in S} \delta^i(\mathbf{x}) f^{is}(\mathbf{x}) \qquad \forall d \in D, \mathbf{x} \in \mathcal{R} \tag{5d}$$

$$\Lambda^d(\mathbf{x}) = \sum_{s\in S} a_{sd} H^s \sum_{i\in S} \lambda^i(\mathbf{x}) f^{is}(\mathbf{x}) \qquad \forall d \in D, \mathbf{x} \in \mathcal{R} \tag{5e}$$

$$\sum_{s\in S} f^{is}(\mathbf{x}) = 1 \qquad \forall i \in S, \mathbf{x} \in \mathcal{R} \tag{5f}$$

$$0 \leq f^{is}(\mathbf{x}) \leq 1 \qquad \forall i, s \in S : \gamma^i \leq \gamma^s, \mathbf{x} \in \mathcal{R} \tag{5g}$$

$$f^{is}(\mathbf{x}) = 0 \qquad \forall i, s \in S : \gamma^i > \gamma^s, \mathbf{x} \in \mathcal{R} \tag{5h}$$

$$n^d(\mathbf{x}) \geq 0 \qquad \forall d \in D, \mathbf{x} \in \mathcal{R} \tag{5i}$$

The objective function (5a) sums the routing cost over all days and the service benefit over all schedules. Constraints (5b) ensure that vehicle routes do not exceed capacity. Constraints (5c) define the items collected per stop as the demand density about a point divided by node density about a point. Since all accumulated demand between visits is collected, we have $v^d(\mathbf{x}) = \frac{\Lambda^d(\mathbf{x})}{\Delta^d(\mathbf{x})}$. Constraints (5d)-(5e) define the auxiliary decision functions according to equations (2) and (3). Constraints (5f) ensure that all nodes are assigned to a schedule. Constraints (5g) and (5h) ensure that no node is served with a lower frequency than the minimum specified. Constraints (5i) are non-negativity constraints on the decision function $n^d(\mathbf{x})$.

2.3 Discussion of models

Francis et al. (2006) develop exact and heuristic solution methods for the discrete formulation for the tactical/operational PVRP-SC. These methods are limited to problem instances of moderate size (50 nodes, 3 service choices, and 5 days). The number of decision variables and constraints expand exponentially as the number of nodes, service choices and days increase. As a result, it is difficult to conduct comprehensive strategic analysis on large instances. Hence, the method is well-suited for operational and tactical problems when detailed solutions are necessary. Alternatively, the continuous model shown in Section 2.2 is an approximation of the discrete formulation. In Section 4, we show that the continuous approximation is quite accurate in estimating the objective value for a test case in the literature. A comparison of the two formulations is shown in Appendix A for the PVRP without service choice. Ouyang and Daganzo (2006) present an analytical discussion of the accuracy of continuous approximation models.

In the next section we show that the continuous approximation model can be solved easily with a few modifications. The modified continuous approximation model can yield solutions for large instances which may arise in the strategic planning phases of periodic distribution systems. Thus, the continuous model is not suggested as a replacement for the discrete modeling approach; rather, it is as a complimentary model that can be used to estimate costs and develop design guidelines. Hall (1986) presents examples illustrating the complementary nature of discrete and continuous models.

This paper focuses on the use of the continuous approximation method for strategic decision making, estimation, and the selection of parameters. After parameters have been chosen, the operational/tactical decisions can be made using the discrete method. Ouyang and Daganzo (2006) and Ouyang (2006) propose methods to discretize solutions from continuous models.

3 Solution method for the continuous approximation models

This section describes the solution method for the continuous approximation of the PVRP-SC. We use geographic decomposition and variable substitution to reduce formulation (5) to a simple problem that can be solved easily.

3.1 Geographic decomposition

We propose a solution approach for formulation (5) based on geographic decomposition. Since parameters are assumed to vary slowly, we can decompose the problem geographically into a set of subregions. In designing the decomposition, the following trade-off must be considered. A subregion $\mathcal{A}$ of $\mathcal{R}$ should be small enough such that $\delta^i(\mathbf{x})$ and $\lambda^i(\mathbf{x})$ are nearly constant over the subregion, as well as any other parameters that vary by location. It can then be assumed that the values for all data functions for points $\mathbf{x} \in \mathcal{A}$ are equal to the average over the region. On the other hand, the subregion should be large enough to contain at least one route; therefore, the objective function and constraints for each subregion can be treated as independent of the other subregions. As a result, the problem decomposes by subregion. Let $\zeta_{\mathcal{A}}$ denote the cost in subregion $\mathcal{A}$ of $\mathcal{R}$:

$$\zeta_{\mathcal{A}} = \sum_{d \in D} \Lambda^d z^d - \sum_{s \in S} \alpha^s \sum_{i \in S} \lambda^i f^{is} \tag{6}$$

Note that we drop the point coordinates $\mathbf{x}$ for parameters and decision functions since averages across all points in the subregion are used. For simplicity of notation, we include the index $\mathcal{A}$ only for the cost $\zeta_{\mathcal{A}}$, and not on subregion-specific parameters and decision functions. We minimize $\zeta_{\mathcal{A}}$ for each subregion, multiply $\zeta_{\mathcal{A}}$ by the area of the subregion, and sum over all subregions to obtain the total cost.

3.2 Modified cost model

We simplify formulation (5) such that the entire model can be expressed as the function of a single decision function. Consider the routing component of $\zeta_{\mathcal{A}}$:

$$\sum_{d \in D} \Lambda^d z^d = \sum_{d \in D} \Lambda^d \left[\frac{2r\bar{c}}{n^d v^d} + \frac{\bar{c}\hat{k}(\Delta^d)^{-1/2} + \dot{\tau} + \ddot{\tau} v^d}{v^d} \right]$$

Since there are no limits on the number of vehicles or the distance traveled by a vehicle, an optimal solution will choose to transport full vehicle loads: $n^d v^d = C$. As a result, within each subregion the capacity constraints (5b) are always binding. Additionally, we can replace v^d with auxiliary variables Δ^d and Λ^d using definitional constraints (5c). Further, since the sum over all days of Λ^d is equal to the total demand over the planning period and all demand must be served, we obtain the following constant: $\beta_0 = \left(\frac{2r\bar{c}}{C} + \ddot{\tau}\right) \sum_{i \in S} \lambda^i$. The modified routing component of $\zeta_{\mathcal{A}}$ becomes:

$$\sum_{d \in D} \Lambda^d z^d = \beta_0 + \dot{\tau} \sum_{d \in D} \Delta^d + c\hat{k} \sum_{d \in D} \sqrt{\Delta^d}$$

The node densities Δ^d can be expressed in terms of f^{is} using definitional constraints (5d). Recall that $\sum_{d \in D} a_{sd} = \gamma^s$ by definition. Incorporating the service benefit term, which is already a function of only f^{is}, we can write $\zeta_{\mathcal{A}}$ as:

$$\zeta_{\mathcal{A}} \quad = \quad \beta_0 + \sum_{i \in S} \sum_{s \in S} \beta_1^{is} f^{is} + \sum_{d \in D} \sqrt{\sum_{i \in S} \sum_{s \in S} \beta_2^{isd} f^{is}}$$

where:

$$\begin{aligned} \beta_1^{is} &= \dot{\tau} \delta^i \gamma^s - \alpha^s \lambda^i \\ \beta_2^{isd} &= (\bar{c}\hat{k})^2 a_{sd} \delta^i \end{aligned}$$

In this form, the optimization problem to be solved for each subregion is:

$$\min \zeta_{\mathcal{A}} = \beta_0 + \sum_{i \in S} \sum_{s \in S} \beta_1^{is} f^{is} + \sum_{d \in D} \sqrt{\sum_{i \in S} \sum_{s \in S} \beta_2^{isd} f^{is}} \tag{7a}$$

subject to

$$\sum_{s \in S} f^{is} = 1 \quad \forall i \in S \tag{7b}$$

$$f^{is} = 0 \quad \forall i, s \in S : \gamma^i > \gamma^s \tag{7c}$$

$$0 \leq f^{is} \leq 1 \quad \forall i, s \in S : \gamma^i \leq \gamma^s \tag{7d}$$

The problem is a generalized assignment problem with a non-linear objective function, and can be solved using commercially available non-linear solvers. The continuous approximation solution method is implemented using AMPL (Fourer et al., 2003) with the KNITRO solver (Waltz, 2004) on a Sun workstation with two UltraSparc III processors and 2 GB of memory. As the KNITRO solver only guarantees a locally-optimal solution, we iteratively call AMPL to solve the problem within a multi-start algorithm. The starting points correspond to a sample of the extreme points of the solution hull. Computational studies on a series of test cases for which the global optimal solution can be found through enumeration suggest that the multistart algorithm obtains the optimal solution. As a result, no special numerical methods are required for the test cases in this paper.

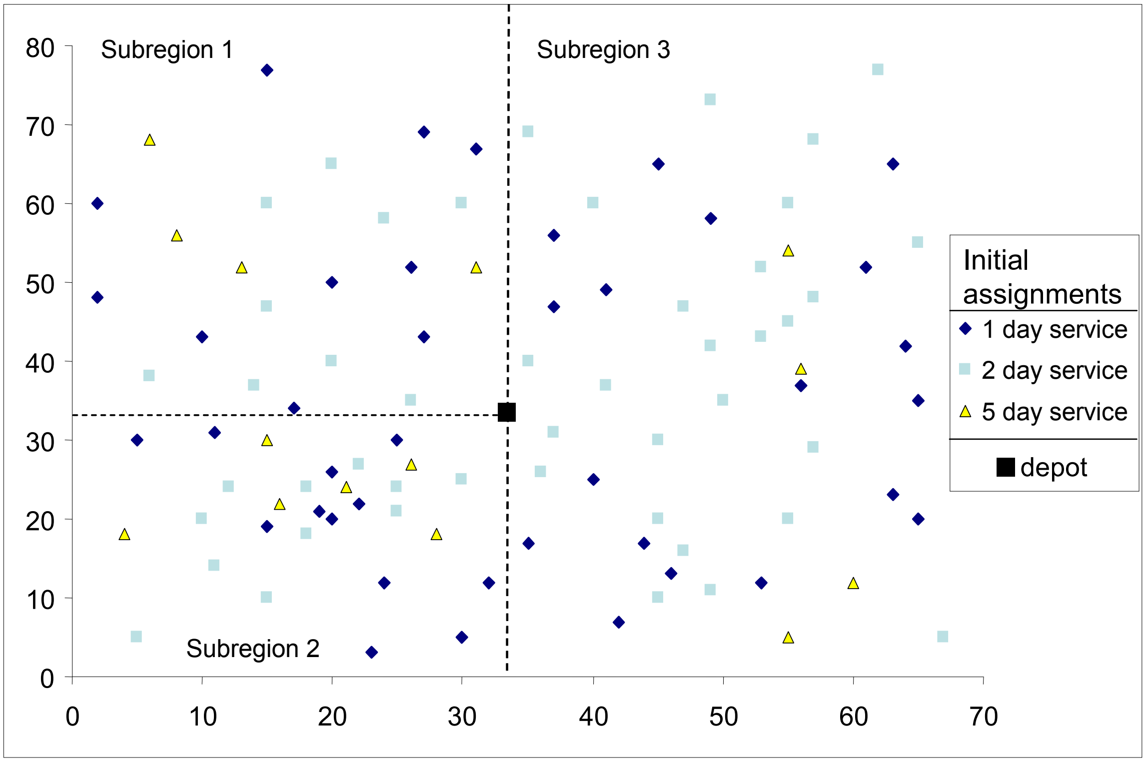

Figure 2: Distribution of nodes in the 100b dataset

4 Computational study

In this section, we present numerical studies to validate the continuous approximation model and demonstrate its use in the strategic and tactical analysis of the PVRP-SC. The experiments are conducted with the 100b PVRP dataset from Christofides and Beasley (1984). This data set is commonly used in the PVRP literature; therefore, we can compare the best known discrete solution with the solution obtained with the continuous approximation model (without service choice) and explore the benefits when service choice is introduced.

The 100b dataset contains 100 nodes randomly distributed across the region as shown in Figure 2. The nodes are visited over the five-day period in the PVRP according to demand levels. Nodes with daily demand less than 10 units are visited once; nodes with daily demand between 11 and 25 units are visited twice; and nodes with daily demand of more than 25 units are visited daily. Vehicle capacity C is 1100 units.

Subregion	$i \in S$	δ^i	λ^i
	1	0.0064	0.0406
Subregion 1	2	0.0058	0.0864
	3	0.0026	0.0774
	1	0.0110	0.0634
Subregion 2	2	0.0101	0.1754
	3	0.0055	0.1680
	1	0.0061	0.0409
Subregion 3	2	0.0088	0.1537
	3	0.0014	0.0462

Table 1: Parameters for 100b dataset

To solve this problem instance with continuous approximation, the entire region is divided into three subregions such that the node densities and demand rates are approximately uniform within each region. These parameters are listed in Table 1 by service level assignment and subregion. The subregions are large enough to contain at least one vehicle tour. The illustrations of 100b solutions in the literature suggest that all tours fan out from the depot. As a result, the linehaul component is not considered. To account for the travel distance to/from the depot, the depot is added in the calculation of node density for daily service. Although previous studies of instance 100b in the literature have considered other schedule options, we consider three service schedules: one-day per week ($s = 1, \gamma^1 = 1$); two-days per week ($s = 2, \gamma^2 = 2$); and daily service ($s = 3, \gamma^3 = 5$). Since instance 100b does not contain stopping costs, both $\dot{\tau}$ and $\ddot{\tau}$ are set to zero.

4.1 Model validation

The continuous approximation model is applied to instance 100b without service choice to evaluate the routing cost results from the model relative to discrete results for the PVRP that appear in the literature. The best known solutions for 100b are due to Chao et al. (1995) and Cordeau et al. (1997). Chao et al. (1995) obtain a routing cost of 2,075 in 13 CPU minutes (a cost of 2,042 is obtained, but no solution time is provided). Cordeau et al. (1997) obtain a routing cost of 2,055 in 10 CPU minutes (a cost of 2,042 is obtained, but no solution time is provided). The continuous approximation model yields a solution of 2,164 in under 1 CPU second which represents roughly a 5% deviation from the best known solutions. One reason for the discrepancy may be that the continuous model considers fewer schedule options. Additionally, solving the three subregions independently in the continuous model can increase the objective function.

As discussed in Francis et al. (2006), the complexity of the discrete PVRP increases significantly when service choice is introduced. As a result, instance 100b cannot be solved with existing solution methods as a discrete PVRP-SC. Although current research is underway on solution methods for larger instances of the discrete PVRP-SC, it is important to assess whether the operational improvements due to service choice warrant the increased complexity. Fortunately, the continuous approximation model allows us to answer this question for a problem instance without solving the discrete formulation. In the next section, we demonstrate how the continuous approximation model can address this and other strategic and tactical questions regarding service choice in periodic distribution operations.

4.2 Value of service choice

We examine the value of introducing service choice in the PVRP using three basic scenarios. Scenario I is the traditional PVRP in which nodes must be served at their initially assigned service level. Scenario II allows service choice, but only considers the impact of this choice on routing efficiency; the objective function does not include the service benefit term. Scenario III allows service choice and considers both routing cost and service benefit in the objective function. The following default service benefit parameters are used: $\alpha^1 = \$1$, $\alpha^2 = \$2$, and $\alpha^3 = \$5$.

Table 2 presents the results for the three subregions and the total service area. For each scenario, the routing cost is displayed for each subregion. In addition, the service benefit is presented even when this term is not explicitly considered in the model. In these cases, the optimal f^{is} values are used to calculate service benefit with the default parameter values. Finally, we present the total objective value of routing cost net service benefit.

The results in Table 2 for Scenario I represent the routing costs discussed in model validation, with the service benefit of the initial assignments calculated with the default values of α^s for all

	Subregion 1	Subregion 2	Subregion 3	Total region
I. No service choice				
Routing	559	530	1075	2164
Benefit	-931	-703	-1715	-3349
Total	-372	-173	-640	-1185
II. Service choice; benefit not considered in optimization				
Routing	559	530	1041	2130
Benefit	-931	-703	-1836	-3470
Total	-372	-173	-795	-1340
III. Service choice; benefit considered in optimization				
Routing	722	677	1430	2829
Benefit	-1585	-1552	-3564	-6701
Total	-863	-875	-2134	-3872

Table 2: Impact of service choice on objective function

$s \in S$. Introducing service choice in Scenario II does not affect subregions 1 and 2, suggesting that the initial assignments yield highly efficient routing solutions. In these subregions, all demand is served at its initial level: $f^{11} = 1$, $f^{22} = 1$, and $f^{33} = 1$. However, in subregion 3, the routing solution improves with service choice, and demand from level $i = 1$ is served at level $s = 2$: $f^{12} = 1$. Consequently, the service benefit improves as well. These effects are consistent with those observed for other data sets tested with the discrete model in Francis et al. (2006). The relative densities of the initial service levels can explain why these changes are observed in subregion 3 and not in the other subregions. This effect is studied empirically in Section 4.2.1 and analytically in Appendix B.

In Scenario III in which service benefit is implicitly considered in the optimization, we observe significant changes in service allocation. With the default service parameter values, it is optimal to serve all the demand at the daily level. A moderate increase in routing costs is observed which is more than offset by a large improvement in service benefit. We explore the impact of service parameter values in Section 4.2.2.

4.2.1 Impact of node density

In the above analysis, we observe that the value of service choice is sensitive to the relative node densities. In Scenario II, the allocation of demand from initial level $i = 1$ to a higher level $i = 2$ in subregion 3 is due to the relative densities of all three initial service levels. We explore this further with the following experiment. We incrementally change the relative densities of δ^2 and δ^3, while keeping δ^1 constant, and monitor the changes in the allocation of nodes initially assigned to $i = 1$. Figure 3 plots the decision threshold between serving nodes at their initial level ($f^{11} = 1$) and raising the service level ($f^{12} = 1$). As δ^2 increases, it becomes optimal to set $f^{12} = 1$; the increase in the second term of the objective function of problem (7) is offset by a decrease in the third term (routing cost) due to the resulting economies of scale. As the density of level 3 nodes, δ^3, increases, the critical density δ^2 also increases (greater savings are required to balance increases in the non-separable third term).

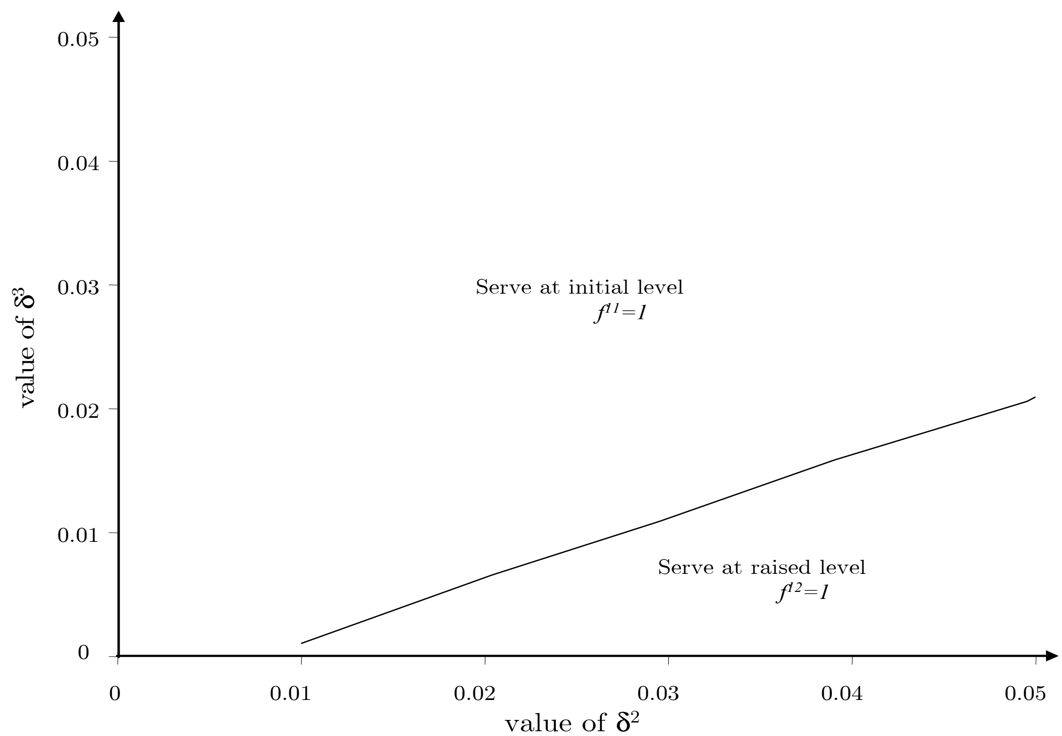

Figure 3: Sensitivity of allocation of nodes at initial level to the densities of higher service schedules

4.2.2 Impact of service benefit parameters

It is intuitive that at high values of the service benefit parameters relative to routing costs, nodes will be visited more frequently than required. To examine the sensitivity of the solution to different values of service benefit, we test instance 100b at various levels of service benefits. We maintain the ratios of the α^s values: $\alpha^1 = \frac{1}{5}\alpha^3$ and $\alpha^2 = \frac{2}{5}\alpha^3$, and gradually raise the value of α^3 from 0 to the default value of 5. Figure 4 plots the routing costs, service benefit and total objective of instance 100b for the different service benefits. From this plot, we can determine the breakpoints where changes in service benefit cause nodes to be served at higher service levels.

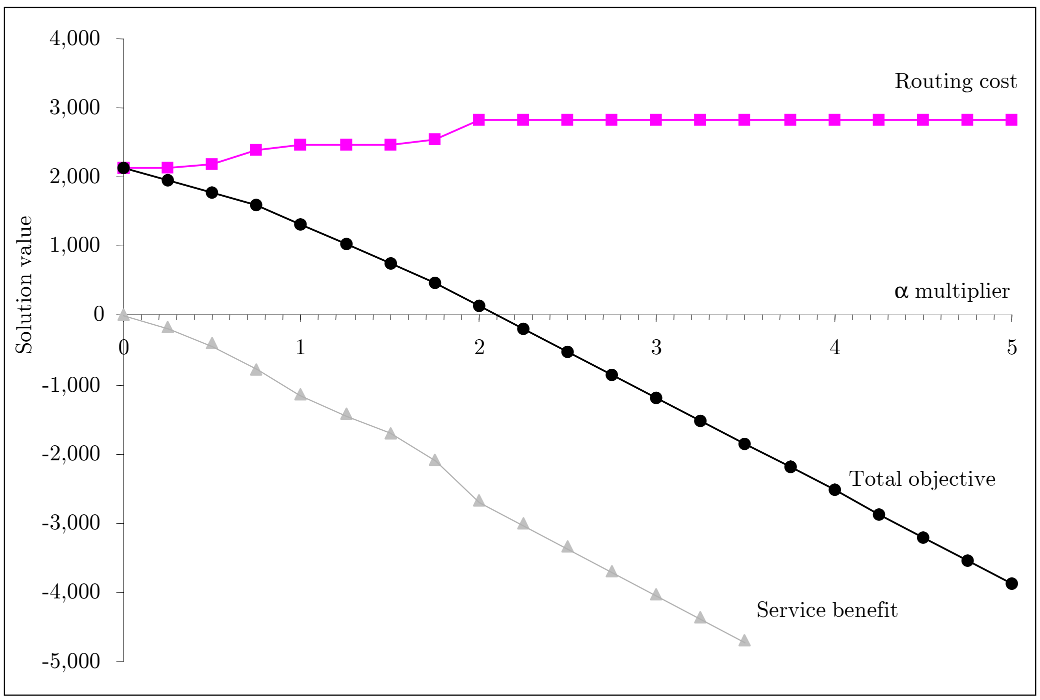

Figure 4: Sensitivity of solution to service benefit term

Note that it is possible to have $f^{12} = 1$, even if none of the demand initially allocated to level 2 is being served by level 2. This property can be generally stated as follows. For some schedules $i, s \in S$, such that $\gamma^i < \gamma^s$, it is possible to have $f^{is} = 1$ and $f^{ss} = 0$; that is, service to a region at a particular schedule s need not include any of the demand initially allocated to s.

5 Conclusions

This paper presents a continuous approximation model for the PVRP-SC and demonstrates the validity and usefulness of the model with numerical studies using a test instance from the literature. The results of the continuous approximation model can help distribution service providers design valuable service options. The results can also be used to guide discrete solutions to determine exact vehicle routes.

Due to the general nature of the model, it is easy to extend the model for specific conditions. For instance, if customers in schedule i are willing to be served at schedule s, but unwilling to pay for higher level of service, we merely correct the benefit term in the β_1^{is} coefficient to α^i rather than α^s. As the model can be solved quickly, one can do a parametric analysis of future conditions. It may be useful to test the sensitivity of solutions to future "what if" scenarios before committing to higher service levels for some demand nodes.

Acknowledgment

The authors would like to Richard Waltz for help with the KNITRO. This research has been supported by grant DMI-0348622 from the National Science Foundation.

References

Beltrami, E., Bodin, L., 1974. Networks and vehicle routing for municipal waste collection. Networks 4(1), 65–94.

Blakely, F., Bozkaya, B., Cao, B., Hall, W., Knolmajer, J., 2003. Optimizing periodic maintenance operations for Schindler Elevator Corporation. Interfaces 33(1), 67–79.

Chao, I.M., Golden, B., Wasil, E., 1995. An improved heuristic for the period vehicle routing problem. Networks 26(1), 25–44.

Christofides, N., Beasley, J., 1984. The period routing problem. Networks 14(2), 237–256.

Cordeau, J.F., Gendreau, M., Laporte, G., 1997. A Tabu Search heuristic for periodic and multi-depot vehicle routing problems. Networks 30(2), 105–119.

Daganzo, C. F., 1987. Modeling distribution problems with time windows: Part II. Transportation Science 21(3), 180–187.

Daganzo, C. F., 1999. Logistics Systems Analysis. Springer, New York.

Erera, A. L., 2000. Design of Logistics Systems for Uncertain Environments. PhD dissertation, University of California, Berkeley, Institute of Transportation Studies, 2000.

Fourer, R., Gay, D., Kernighan, B., 2003. AMPL: A Modeling Language for Mathematical Programming. Brooks/Cole – Thomson Learning, Pacific Grove, CA.

Francis, P., Smilowitz, K., Tzur, M., 2006. The period vehicle routing problem with service choice. Transportation Science, in press.

Hall, R.W., 1986. Discrete Models/Continuous Models. Omega Int. J. of Management Science 14(3), 213-220.

Ouyang, Y., Daganzo, C.F., 2006. Discretization and validation of the continuum approximation scheme for terminal system design. Transportation Science, in press.

Ouyang, Y., 2006. Discretizing vehicle routing zones for large-scale distribution systems with customer uncertainty. Forthcoming presentation at 85th Annual Meeting of the Transportation Research Board, Washington, DC.

Robuste, F., Daganzo, C. F., Souleyrette II, R. R., 1990. Implementing vehicle routing models. Transportation Research Part B 24(4), 263–286.

Russell, R., Igo, W., 1979. An assignment routing problem. Networks 9(1), 1–17.

Smilowitz, K., Daganzo, C. F., 2004. Cost modeling and solution techniques for complex transportation systems. IEMS Working paper 04-006, Northwestern University .

Waltz, R., October 2004. Knitro 4.0. User manual, Ziena Optimization, Inc., Evanston, IL, USA. (Available online at http://www.ziena.com/documentation.html).

Appendix A

This appendix compares the cost models for the discrete and continuous approaches for the PVRP. Recall that the objective function for the discrete model is:

$$Z = \sum_{k \in K} \left[\sum_{d \in D} \sum_{(l,m) \in A} c_{lm} x^d_{lmk} + \sum_{s \in S} \sum_{l \in N} \gamma^s \tau^s_l y^s_{lk} - \sum_{s \in S} \sum_{l \in N} W_l \alpha^s y^s_{lk} \right]$$

The objective function for the continuous model is:

$$Z_{\mathcal{R}} = \int_{\mathbf{x} \in \mathcal{R}} \left[\left(\frac{2r(\mathbf{x})\bar{c}}{n^d(\mathbf{x}) v^d(\mathbf{x})} + ck \sum_{d \in D} \sqrt{\Delta^d(x)} \right) + \left(\ddot{\tau} \sum_{i \in S} \lambda^i(\mathbf{x}) + \dot{\tau} \sum_{d \in D} \Delta^d(x) \right) - \sum_{s \in S} \alpha^s \sum_{i \in S} \lambda^i(\mathbf{x}) f^{is}(\mathbf{x}) \right] d\mathbf{x}$$

We examine the case when service choice is not allowed, i.e. $f^{ii}(x) = 1, \ \forall i \in S$. First we consider the two service benefit terms:

$$\sum_{k \in K} \sum_{s \in S} \sum_{l \in N} W_l \alpha^s y^s_{lk} \iff \int_{\mathbf{x} \in \mathcal{R}} \sum_{s \in S} \alpha^s \sum_{i \in S} \lambda^i(\mathbf{x}) f^{is}(\mathbf{x}) dx$$

We note that the service benefit terms is simply the total service benefit, and it is the same in the discrete and continuous case. In the discrete case, the service benefit is $\sum_{k \in K} \sum_{l \in N} W_l \alpha^{i_l}$, where

$i_l \in S$ is the minimum schedule to which node $l \in N$ is allocated. In the continuous case, the service benefit is $\int_{x\in\mathcal{R}} \left[\sum_{i\in S} \alpha^i \lambda^i(\mathbf{x})\right] dx$.

We consider the stopping time terms:

$$\sum_{k\in K}\sum_{s\in S}\sum_{l\in N} \gamma^s \tau_l^s y_{lk}^s \iff \int_{\mathbf{x}\in\mathcal{R}} \left(\ddot{\tau} \sum_{i\in S} \lambda^i(\mathbf{x}) + \dot{\tau} \sum_{d\in D} \Delta^d(x) \right) dx$$

When service choice is not allowed, we can use the definition of $\Delta^d(\mathbf{x})$ to rewrite the above as:

$$\sum_{k\in K}\sum_{s\in S}\sum_{l\in N} \gamma^s \tau_l^s y_{lk}^s \iff \int_{\mathbf{x}\in\mathcal{R}} \left(\ddot{\tau} \sum_{i\in S} \lambda^i(\mathbf{x}) + \dot{\tau} \sum_{d\in D}\sum_{i\in S} a_{id} \delta^i(\mathbf{x}) \right) dx$$

From the above comparison, we can see that the accuracy of the approximation depends largely on the accuracy of the approximation of the stopping time parameters, $\dot{\tau}$ and $\ddot{\tau}$, which can be derived from empirical analysis of discrete stopping times. Similarly, we can compare the travel times:

$$\sum_{k\in K}\sum_{d\in D}\sum_{(l,m)\in A} c_{lm} x_{lmk}^d \iff \int_{\mathbf{x}\in\mathcal{R}} \left(\frac{2r(\mathbf{x})\bar{c}}{n^d(\mathbf{x})v^d(\mathbf{x})} + ck \sum_{d\in D} \sqrt{\Delta^d(x)} \right) dx$$

Again, the accuracy of the approximation depends on the quality of the approximating functions for distance. See Hall (1986) for a discussion of continuous approximations for discrete data. The vehicle routing cost expressions have been validated by Erera (2000) and Robuste et al. (1990) who show a gap of less than 5% between continuous approximations and costs from discrete simulations for their problem sets.

Appendix B

The following appendix develops analytic expressions to determine changes in service levels based on node densities when only routing efficency is considered in the objective function of the PVRP-SC. For the 100b test case, we have no stopping costs ($\ddot{\tau} = 0,\ \dot{\tau} = 0$) and no service benefit ($\alpha^s = 0,\ \forall s \in S$). Three service levels are offered with a_{sd} vectors given by:

$$a_{1d} = [\ 10000\ ] \ \ ;\ a_{2d} = [\ 01010\ ] \ \ ;\ a_{3d} = [\ 11111\ ]$$

Recall that the objective function in the PVRP-SC is:

$$\zeta_{\mathcal{A}} = \beta_0 + \sum_{i\in S}\sum_{s\in S} \beta_1^{is} f^{is} + \sum_{d\in D} \sqrt{\sum_{i\in S}\sum_{s\in S} \beta_2^{isd} f^{is}}$$

where β_0 is a constant and can be ignored, $\beta_1^{is} = \dot{\tau}\delta^i\gamma^s - \alpha^s\lambda^i = 0$ since $\dot{\tau} = 0$ and $\alpha^s = 0$, and $\beta_2^{isd} = (\bar{c}\hat{k})^2 a_{sd}\delta^i$. Since days 2 and 4 are equal (and days 3 and 5), the terms can be combined in instance 100b. Dividing by the constants $\bar{c}$ and $\hat{k}$, the cost model becomes:

$$\frac{\ddot{\zeta}_{\mathcal{A}}}{\bar{c}\hat{k}} = \sqrt{\sum_{i\in S}\delta^i(f^{i1}+f^{i3})} + 2\sqrt{\sum_{s\in S}\delta^i(f^{i2}+f^{i3})} + 2\sqrt{\sum_{i\in S}\delta^i f^{i3}}$$

According to the a_{sd} vectors above, the first term represents costs for day 1, the second for days 2 and 4, and the third for days 3 and 5. For the 100b case, the full optimization problem is :

$$\begin{aligned}\min \sqrt{\delta^1\left(f^{11}+f^{13}\right)+\delta^2 f^{23}+\delta^3 f^{33}}& \\ &+2\sqrt{\delta^1\left(f^{12}+f^{13}\right)+\delta^2\left(f^{22}+f^{23}\right)+\delta^3 f^{33}} \\ &+2\sqrt{\delta^1 f^{13}+\delta^2 f^{23}+\delta^3 f^{33}}\end{aligned}$$

subject to:

$$\begin{aligned}\textstyle\sum_{s\in S} f^{is} &= 1 \quad \forall i\in S \\ f^{is} &= 0 \quad \forall i\in S, s\in S : \gamma^i > \gamma^s \\ f^{is} &\geq 0 \quad \forall i\in S, s\in S : \gamma^i \leq \gamma^s\end{aligned}$$

Note from the constraints that $f^{33} = 1$ and $f^{22} + f^{23} = 1$. Hence:

$$\begin{aligned}\frac{\ddot{\zeta}_{\mathcal{A}}}{\bar{c}\hat{k}} &= \sqrt{\delta^1\left(f^{11}+f^{13}\right)+\delta^2 f^{23}+\delta^3} \\ &\qquad +2\sqrt{\delta^1\left(f^{12}+f^{13}\right)+\delta^2+\delta^3} \\ &\qquad +2\sqrt{\delta^1 f^{13}+\delta^2 f^{23}+\delta^3}\end{aligned}$$

It is intuitive that setting $f^{22} = 1$ and $f^{23} = 0$ minimizes the cost by not incurring the additional costs δ^2 in the first and third terms in the objective function. For similar reasons, setting $f^{13} = 1$ is clearly suboptimal, as the resulting cost is always greater than $f^{12} = 1$. Thus, demand at $i = 1$ will be served either at $s = 1$ or $s = 2$ in the optimal solution. For the nodes initially assigned to $i = 1$ to be served at the higher level $s = 2$ ($f^{12} = 1$) the following condition must be true:

$$\sqrt{\delta^1+\delta^3} + 2\sqrt{\delta^2+\delta^3} + 2\sqrt{\delta^3} > 2\sqrt{\delta^1+\delta^2+\delta^3} + 3\sqrt{\delta^3}$$

$$\Rightarrow \sqrt{\delta^1+\delta^3} + 2\sqrt{\delta^2+\delta^3} > 2\sqrt{\delta^1+\delta^2+\delta^3} + \sqrt{\delta^3}$$

If the density of daily nodes in the region, δ^3, is very large compared to δ^1 and δ^2, then serving nodes at higher levels does not result in any significant improvement in routing costs. Alternatively, if the density of daily nodes is relatively small, and $\delta^2 >> \delta^1$, then the level 1 nodes, $i = 1$, should be merged into the level 2 routes, $s = 2$, for greater efficiency. However, if $\delta^1 >> \delta^2$, a diseconomy of scale results as many more nodes must now be served at the higher level $s = 2$, which is relatively inefficient.