

Flexibility and complexity in periodic distribution problems

Peter Francis*, Karen Smilowitz*, and Michal Tzur†

July 18, 2006

Abstract

In this paper, we explore trade-offs between operational flexibility and operational complexity in periodic distribution problems. We consider the gains from operational flexibility in terms of vehicle routing costs and customer service benefits, and the costs of operational complexity in terms of modeling, solution methods and implementation challenges for drivers and customers. The period vehicle routing problem (PVRP) is a variation of the classic vehicle routing problem in which delivery routes are constructed for a period of time; the PVRP with service choice (PVRP-SC) extends the PVRP to allow service (visit) frequency to become a decision of the model. For the periodic distribution problems represented by PVRP and PVRP-SC, we introduce operational flexibility levers and a set of quantitative measures to evaluate the trade-offs between flexibility and complexity. We develop a Tabu Search heuristic to incorporate a range of operational flexibility options. We analyze the potential value and the increased operational complexity of the flexibility levers.

*Department of Industrial Engineering and Management Sciences, Northwestern University

†Department of Industrial Engineering, Tel Aviv University

Introduction

In periodic distribution problems, vehicles visit customers over a given period of time, resulting in a schedule which may be repeated. One such problem is the period vehicle routing problem (PVRP), an extension of the vehicle routing problem (VRP) in which delivery routes are constructed over a period of time (for example, multiple days) to visit customers according to preset visit frequencies. The PVRP with service choice (PVRP-SC) extends the PVRP by making service (visit) frequency to customers a decision of the model. In these systems, customers may be heterogeneous in their demand levels, visit requirements, and willingness-to-pay for more frequent service. Operational flexibility can help to avoid under-serving customers with high service requirements and over-serving customers with low requirements.

While introducing operational flexibility in periodic distribution problems can increase efficiency in terms of vehicle routing costs and customer service benefits, it poses challenges in (i) modeling flexibility accurately, (ii) addressing the computational effort needed to solve problems with such flexibility, and (iii) implementing resulting solutions. Francis et al. [11] introduce a formulation and an exact solution method for the PVRP and PVRP-SC; however, the formulation and solution method are limited in the range of operational flexibilities that can be incorporated.

In this paper we develop a Tabu Search method that can incorporate a range of operational flexibility options, including the ability to increase the set of visit schedules, decide visit frequency, vary the drivers who visit a customer, and decide delivery amounts per visit. Further, we develop a set of quantitative measures to evaluate the trade-offs between flexibility and complexity in distribution problems. These are new and novel measures that may be used in various distribution problems. We analyze the trade-offs between the system performance improvements due to operational flexibility and the resulting increases in implementation, computational and modeling complexity as they relate to the PVRP. Our results provide insights into the value of flexibility options and the associated increase in complexity. Using the Tabu Search method, problems with other characteristics may be examined with respect to the trade-offs between operational flexibility and solution complexity.

Section 1 reviews the literature related to the periodic distribution problems discussed in this paper, including the exact method of Francis et al. [11]. Section 2 presents the flexibility options and complexity measures, and introduces a Tabu Search method to study general PVRPs and a set of performance metrics to evaluate resulting solutions. Section 3 presents a computational study of

the trade-offs between flexibility and complexity and provides insights into these results. Finally, Section 4 summarizes the paper and discusses future work.

1 Periodic distribution problems

Section 1.1 discusses related literature and Section 1.2 reviews the exact method of Francis et al. [11].

1.1 Related literature

Periodic distribution problems occur in many industries, including courier services, elevator maintenance and repair (Blakely et al. [5]), the collection of waste (Russell and Igo [17]) and the delivery of interlibrary loan material (Francis et al. [11]). The Period Vehicle Routing Problem (PVRP), introduced in Beltrami and Bodin [4] and Russell and Igo [17], finds a set of vehicle tours over a period of days that minimizes total travel time while satisfying operational constraints (vehicle capacity and pre-determined visit requirements for each customer). A set of visit schedules is available for each customer (node), and one schedule from this set must be chosen. A schedule represents the days on which a node is visited. All feasible schedule options for a node must provide the pre-determined number of visits for that node. For example, if over the period of one week, a node is to be visited three times, the feasible schedule options may include: Mon-Tue-Thu, Mon-Wed-Fri, or Tue-Wed-Fri.

Heuristic solution methods for the PVRP are presented in Tan and Beasley [18], Russell and Gribbin [16], Chao et al. [7], Cordeau et al. [8], and Angelelli and Speranza [1]. Francis et al. [11] introduce the Period Vehicle Routing Problem with Service Choice (PVRP-SC), which allows customers to be visited more often than their pre-determined frequencies. Service choice may be advantageous if, for example, two customers with different minimum requirements are located in isolation of all other customers and the depot. If the schedule options for these customers do not contain overlapping days, it may be beneficial to raise the visit frequency of one customer such that both customers are visited together. Francis et al. [11] show that this is also true in less extreme cases in which arriving at a certain region makes it beneficial to visit neighboring customers, hence increasing the frequency with which some nodes are visited. Francis and Smilowitz [10] present a continuous approximation model of the PVRP-SC and show that the value of service choice depends on the relative density of customers of different visit requirements.

When flexibility in service choice is introduced, the problem begins to resemble the Inventory Routing Problem (IRP). The IRP, like the PVRP-SC, determines visit frequency and route configuration simultaneously, but with an additional decision of how much to deliver to the customers; see Anily and Federgruen [3] and Chan et al. [6], and the surveys in Federgruen and Simchi-Levi [9], Anily and Bramel [2], and Kleywegt et al. [14]. Rusdiansyah and Tsao [15] model the IRP as an integrated IRP/PVRP with Time Windows. In the IRP, service-related costs are modeled as holding costs associated with each item unit. In the PVRP-SC, the amount delivered to a customer is determined by the schedule chosen for the customer and the adopted delivery strategy, as defined in Section 2.1. Service is modeled as a benefit term related to each customer.

1.2 Exact method from Francis et al. [11]

We review the formulation and the exact method for the PVRP-SC from Francis et al. [11], which can also be used for the PVRP. Let D denote the set of days in the period and S denote the set of visit schedules. The parameter a_{sd} links schedules to days: $a_{sd} = 1$ 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: $\gamma^s = \sum_{d \in D} a_{sd}$, and an associated benefit α^s related to a monetary benefit of the corresponding frequency.

The PVRP-SC is defined for a set of nodes, N_0 , which consists of customers nodes, N , and a depot, $i = 0$, and a set of arcs connecting the nodes, $A = \{(i, j) : i, j \in N_0\}$. Each customer node $i \in N$ has a known daily demand, W_i , and a visit requirement, f_i , measured in days per period. The demand accumulated between visits, w_i^s , is a function of the schedule $s \in S$ and the daily demand of the node, which is set at the maximum accumulation between successive visits. The stopping time at a node, τ_i^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 $(i, j) \in A$ is a known travel time, c_{ij} . There is a set, K , of vehicles, each with capacity C .

The following allocation and routing variables are used.

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

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

We introduce a parameter, $\beta \geq 0$, which weighs the service benefit relative to vehicle travel and

stopping times. The formulation of the PVRP-SC from Francis et al. [11] is:

$$Z^* = \min \sum_{k \in K} \left[\sum_{d \in D} \sum_{(i,j) \in A} c_{ij} x_{ijk}^d + \sum_{s \in S} \sum_{i \in N} \gamma^s \tau_i^s y_{ik}^s - \beta \sum_{s \in S} \sum_{i \in N} W_i \alpha^s y_{ik}^s \right] \quad (1a)$$

subject to

$$\sum_{s \in S} \sum_{k \in K} \gamma^s y_{ik}^s \geq f_i \quad \forall i \in N \quad (1b)$$

$$\sum_{s \in S} \sum_{k \in K} y_{ik}^s \leq 1 \quad \forall i \in N \quad (1c)$$

$$\sum_{s \in S} \sum_{i \in N} w_i^s a_{sd} y_{ik}^s \leq C \quad \forall k \in K; d \in D \quad (1d)$$

$$\sum_{j \in N_0} x_{ijk}^d = \sum_{s \in S} a_{sd} y_{ik}^s \quad \forall i \in N; k \in K; d \in D \quad (1e)$$

$$\sum_{j \in N_0} x_{ijk}^d = \sum_{j \in N_0} x_{jik}^d \quad \forall i \in N_0; k \in K; d \in D \quad (1f)$$

$$\sum_{i,j \in Q} x_{ijk}^d \leq |Q| - 1 \quad \forall Q \subseteq N; k \in K; d \in D \quad (1g)$$

$$y_{ik}^s \in \{0, 1\} \quad \forall i \in N; k \in K; s \in S \quad (1h)$$

$$x_{ijk}^d \in \{0, 1\} \quad \forall (i, j) \in A; k \in K; d \in D \quad (1i)$$

The objective function (1a) balances travel time, stopping time and demand-weighted service benefit. Francis and Smilowitz [10] analyze the impact of the value of α^s on the resulting solution. Constraints (1b) enforce the visit requirement for each node. Constraints (1c) ensure that one schedule and one vehicle are chosen for each node. Constraints (1d) represent vehicle capacity constraints. Constraints (1e) link the x and y variables. Constraints (1f) ensure flow conservation at each node. Constraints (1g) are the subtour elimination constraints and ensure that all routes contain a visit to the depot. Constraints (1h) and (1i) define the binary variables for allocation and routing, respectively.

The exact solution method in Francis et al. [11] consists of a Lagrangian relaxation phase which relaxes constraints (1e) to decompose the problem into a capacitated assignment subproblem in the y variables and a prize-collecting traveling salesman subproblem in the x variables. If a gap remains after the Lagrangian relaxation phase, it is closed using a branch-and-bound phase that incorporates information from the earlier phase. A heuristic variation of this approach truncates nodes of the branch-and-bound tree that are within $\delta\%$ of the lower bound, obtaining solutions

within $\delta\%$ of the optimal. Using this variation, PVRP-SC instances with up to 50 nodes are solved to within $\delta = 2\%$ of optimality.

2 Operational flexibility and complexity in the PVRP and the PVRP-SC

Section 2.1 presents the flexibility options and complexity measures for the PVRP. Section 2.2 introduces a Tabu Search heuristic capable of incorporating all levers of operational flexibility discussed in Section 2.1. Section 2.3 defines performance metrics to quantify the operational complexity of resulting solutions.

2.1 Flexibility options and complexity measures

Francis et al. [11] highlight the difficulties in formulating and solving the PVRP-SC that result from introducing service choice. Several assumptions are made regarding schedule options and visit conditions to accommodate service choice in their formulation. Their exact solution method yields optimal solutions to the PVRP-SC for moderate-sized instances with these assumptions. In this paper, we develop a Tabu Search heuristic to solve more general cases of the PVRP and the PVRP-SC. As a result, we can relax some modeling assumptions of the exact method and evaluate the value and increased complexity of additional levers of operational flexibility. Throughout this paper, we use the following terminology to discuss flexibility and complexity:

1. *Operational flexibility*: The ability to make changes to operating conditions. We focus on the following levers of operational flexibility that are commonly found in distribution systems:
 - (a) *Service choice*. The ability to determine customer visit frequency subject to the stated visit requirement. A customer's *visit frequency* is the number of times the customer is visited in the period. A customer's *visit requirement* is the minimum number of visits allowed. Without service choice flexibility, customers are served at their visit requirements.
 - (b) *Crew flexibility*. The ability to have multiple drivers visit a customer during the period. Without crew flexibility, each customer is visited by one driver throughout the period.
 - (c) *Schedule options*. The ability to offer a greater number of schedule options of different visit patterns that can be chosen by the service provider to serve a customer.

(d) *Delivery strategy*. The ability to choose the amount delivered during each customer visit by allowing inventory or shortage, assuming all demand is eventually served.

2. *Operational complexity*: The difficulty of solution implementation, from the perspective of the service provider and its customers. Solutions with high operational complexity may be difficult to convey (e.g. no simple rules characterize the service selection decision), may involve a high learning cost for drivers, and/or may cause dissatisfaction to customers. We consider three measures of operational complexity:

(a) *Arrival span*: The variability in the time of day when customers are visited over the period. In applications where staffing at customer locations is tied to vehicle visits, high variability in visit time can increase customer staffing complexity.

(b) *Driver coverage*: The portion of the total service region visited by a driver over the period. Zhong et al. [20] model a learning/forgetting behavior for drivers and show that dispatching drivers consistently to the same geographic areas results in driver familiarity and improved driver performance.

(c) *Crewsize*: The number of different drivers visiting a customer over the period. Smaller crewsize indicates consistent dispatching of drivers to customer locations, building relationships between drivers and customers. UPS Corp. [19] cites driver-customer relationships as a competitive advantage in its package delivery operations, attributing 60 million packages a year to sales leads generated by drivers.

Using the Tabu Search heuristic, we consider the four dimensions of operational flexibility in periodic routing problems and explore their impact on the three measures of operational complexity. Next, we show how these measures of flexibility and complexity are modeled in the periodic routing problems.

Results from Francis et al. [11] indicate that the magnitude of the savings obtained by introducing service choice in the PVRP for a given instance depends on geographic distribution of nodes (in particular, nodes of highest visit requirements). In this paper, we explore how additional levers of operational flexibility impact the magnitude of savings and the complexity of the resulting solutions, and how the impact of these levers depends on problem characteristics such as node distribution. We make the following observations regarding these levers in the context of the model and solution method of Francis et al. [11]:

1. *Service choice:* Service choice flexibility can be restricted by modeling the problem with constraint (1b) fixed at equality.
2. *Crew flexibility:* The allocation variables, y_{ik}^s , are defined such that nodes are always visited by the same driver. In the motivating example in Francis et al. [11], this is required due to access restrictions. Assigning nodes only to schedules, y_i^s , relaxes this assumption and may reduce routing costs. However, since the capacity constraints of the PVRP-SC depend on the vehicle index and service level at each node, crew flexibility requires either a non-linear capacity constraint or a fifth index on the routing variables for schedule choice.
3. *Schedule options:* Computational limitations may restrict the number of schedule options considered by the exact method. Choosing schedule options carefully can offer more discriminating choices with limited variation in driver routes. Francis et al. [11] observe that for any set of schedules S consisting of $|S|-1$ disjoint schedules (schedules that do not share any common days) and a schedule that is the union of all disjoint schedules, there are at most $|S|-1$ different routes for each vehicle. Thus, the number of routing variables, x_{ijk}^d , is reduced significantly since it is not necessary to model each day $d \in D$, but rather only one unique delivery day for each disjoint schedule which is repeated each day of that schedule.
4. *Delivery strategies:* It is assumed in most periodic distribution problems that the amount delivered to a node is fixed a priori. Relaxing this assumption may improve routing costs and service benefits; however, adding a set of decision variables for the delivery amount increases the difficulty of the problem significantly.

As discussed above, modeling the levers of flexibility (apart from service choice) using the exact solution method is difficult and the resulting computational effort is significant. Therefore, we develop a Tabu Search heuristic that can incorporate all levers of flexibility.

2.2 Tabu Search

Cordeau et al. [8] implement a Tabu Search heuristic for the PVRP and obtain solutions equal to or better than the best solutions for PVRP test cases in the literature. Angelelli and Speranza [1] also successfully use a Tabu Search method to solve an extension of the PVRP. We develop a Tabu Search method based on that of Cordeau et al. [8], with suitable extensions to model the PVRP-SC and incorporate operational flexibility. The principal change to the Tabu Search algorithm is

the definition of a move (in Step 2(b) below), considering moves from one schedule to another, rather than from one route to another. This change also allows for the exploration of the different flexibility levers.

Tabu Search is a local search improvement method in which neighbors of the current solution are explored at each iteration; see Glover and Laguna [13]. For the PVRP-SC, a solution is a complete specification of the allocation variables (either y_{ik}^s or y_i^s) and a set of routes for each vehicle on each day (the x_{ijk}^d variables), such that each node $i \in N$ is assigned a schedule that satisfies or exceeds its visit requirement, f_i . An attempt is made to improve the solution by changing the schedule allocation of a given node at each iteration. Routes are constructed based on these schedule allocations, using the GENI heuristic of Gendreau et al. [12] which evaluates various tour configurations through a limited number of insertions and reinsertions. The solutions are allowed to be infeasible with respect to capacity but not with respect to visit requirements. Capacity infeasibilities are penalized in the objective function using a penalty term as in the TABUROUTE procedure of Gendreau et al. [12]. Briefly, the algorithm is as follows:

1. Construct an initial solution:
 - (a) Allocate each node $i \in N$ to the lowest-frequency schedule that satisfies the visit requirement f_i (choosing randomly if more than one schedule is a candidate).
 - (b) Construct routes to visit nodes for each day with the GENI heuristic. Without crew flexibility (i.e. y_{ik}^s), each node is always allocated to the vehicle chosen for the first day of the schedule.
 - (c) Create a tabu list (initially empty) to store moves that are temporarily prohibited.
2. Construct a set of possible moves:
 - (a) Randomly select a set of nodes as possible candidates for movement.
 - (b) For each node, consider all moves from its present schedule allocation to another (frequency-feasible) allocation which contains at least one of its p -closest geographic neighbors.
 - (c) Calculate the change in the objective function for each candidate move using the GENI heuristic to evaluate changes in routing costs with penalties for capacity infeasibility.
3. Identify the best move and check its tabu status from the tabu list. A tabu move may be accepted only if its solution is feasible and better than the best feasible solution; otherwise,

the best non-tabu move is accepted (according to standard Tabu Search acceptance criteria for feasible and infeasible solutions) and the solution is updated accordingly.

4. Update the tabu list to include the implemented move; the move is declared tabu for a random number of iterations.
5. Return to Step 2 and repeat until no improvements in the best feasible or infeasible solutions are found for 60 iterations.

Suitable values for the number of candidate nodes chosen in Step 2 and the value of p are discussed in Gendreau et al. [12]. In cases with many schedule options, requiring the presence of a geographic neighbor in any candidate schedule limits the complexity of the evaluation phase; in cases where this requirement results in very few schedule choices, the algorithm randomly chooses from all frequency-feasible choices to ensure diversity of moves.

The Tabu Search method is used to solve the PVRP by not allowing service choice in Step 2. With crew flexibility in Step 2 (i.e. y_i^s), we pick the least-cost vehicle assignments for each individual day, given the chosen schedule. Without crew flexibility, we explore all possible vehicle-schedule combinations. Schedule options are controlled by the set S , which can also model delivery strategy, as described in Section 3.4.4.

Note that unlike some Tabu Search implementations, no post-optimization is attempted on the routes after each movement as numerical tests show that resulting improvements are minimal and the post-optimization improvement routines are computationally expensive.

2.3 Performance metrics

We use two sets of performance metrics to quantify the trade-offs between operational flexibility and operational complexity in periodic distribution problems, all of which apply to PVRP and PVRP-SC solutions. Metrics in the first set, related to routing cost and service benefit, are explicitly considered in the objective function of both the exact solution method and Tabu Search method; therefore, the solution methods attempt to optimize these metrics. Metrics in the second set are related to operational complexity (driver coverage, crewsize and arrival span). These metrics are calculated after solutions are obtained and, therefore, are not optimized by the solution method. It is possible, however, to consider these metrics implicitly during the solution phase through restrictions on variable and parameter definitions. For example, without crew flexibility, we define allocation variables by y_{ik}^s and the crewsize is always one driver y_{ik}^s at each node. Further, restricting schedule

choices in the set S (for example, only disjoint schedules and their union) reduces complexity in driver coverage and arrival span by limiting the number of different routes each driver performs. In the following, we describe how the metrics for operational complexity are calculated.

Obtaining the arrival span is straightforward given a solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ to the PVRP-SC. All routes are assumed to be performed in a counter-clockwise direction so that visit times are not affected by the choice of route direction. If node i is allocated to schedule s and visited by vehicle k on day d , the time at which it is visited is:

$$T_i^d = \sum_{(m,j) \in A_{kd}(i)} \left[c_{mj} \hat{x}_{mjk}^d + \sum_{r \in S} \tau_m^r \hat{y}_{mk}^r \right] \quad (2a)$$

where $A_{kd}(i)$ is the set of arcs traversed before node i by vehicle k on day d . For each $s \in S$, we define D_s as the set of days $d \in D$ where $a_{sd} = 1$. Note that $|D_s| = \gamma^s$. The mean and standard deviation of the visit times for a node, given a chosen schedule s , are:

$$T_i = \frac{1}{\gamma^s} \sum_{d \in D_s} T_i^d \quad (2b)$$

$$\sigma_i = \begin{cases} 0 & \text{if } \gamma^s = 1 \\ \sqrt{\frac{\sum_{d \in D_s} [T_i^d - T_i]^2}{\gamma^s - 1}} & \text{if } \gamma^s > 1 \end{cases} \quad (2c)$$

We define the average arrival span of a solution over all nodes $i \in N$ as:

$$\sigma = \frac{1}{|N|} \sum_{i \in N} \sigma_i \quad (2d)$$

Calculating driver coverage and crewsize is straightforward without crew flexibility (i.e., with y_{ik}^s); however, calculating these metrics requires additional processing of solutions with crew flexibility (i.e., with y_i^s). With crew flexibility, the vehicle index $k \in K$ assigned to a route is arbitrary for any PVRP-SC solution, as shown with the example in Figure 1. The figure depicts a PVRP-SC solution for an instance with six nodes, two vehicles, and a period of two days. On day 1, the vehicle indexed by $k = 1$ visits the nodes on the left side of the service region and the vehicle indexed by $k = 2$ visits the nodes on the right side. On day 2, the indices are reversed. Without crew flexibility, the solution must assign the same vehicle index to the left region (and to the right region) on both days; however, with crew flexibility, there is no incentive to assign the same vehicle index to the left region on both days. It would be an overestimation of the operational complexity to say that the drivers serve different regions on the two days, when the indices may be switched without affecting the objective function. In this example, the complexity-minimizing assignment of indices

is obvious. However, one can envision many instances in which the assignments of indices are not straightforward, particularly with multiple vehicles and multiple days. Therefore, we introduce a mathematical programming approach to assign the driver indices to the arbitrary vehicle indices of the PVRP-SC solution. The goal of this approach is to minimize total driver coverage. We focus on areas rather than nodes since the set of nodes visited changes by day. Such a policy corresponds to an industry practice in which a dispatcher may allocate service areas to drivers familiar with certain neighborhoods and/or customers.

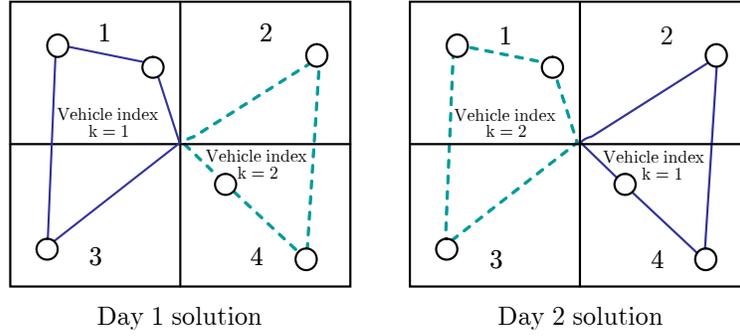


Figure 1: Same solutions may be assigned different vehicle indices

We partition the service region into a set L of cells (cells may represent city blocks), indexed by l , such that each cell contains at least one node. In Figure 1, the service region is divided into four cells: $l = 1, \dots, 4$. Let N_l denote the set of nodes contained in cell $l \in L$ and let V be the set of drivers. We assume that $|V| = |K|$. A driver covers a cell if he visits at least one node in that cell. The assignment problem minimizes the number of cells that each driver covers. Given a PVRP-SC solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$, we define a parameter b_{kld} as:

$$b_{kld} = \begin{cases} 1 & \text{if vehicle index } k \in K \text{ visits cell } l \in L \text{ on day } d \in D; \text{ i.e. if } \sum_{i \in N_l} \sum_{j \in N} \hat{x}_{ijk}^d \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

We define two decision variables:

$$U_{vl} = \begin{cases} 1 & \text{if driver } v \in V \text{ visits cell } l \in L \text{ at least once during the period} \\ 0 & \text{otherwise} \end{cases}$$

$$W_{vkd} = \begin{cases} 1 & \text{if driver } v \in V \text{ is assigned to vehicle index } k \in K \text{ on day } d \in D \\ 0 & \text{otherwise} \end{cases}$$

The assignment problem is formulated as:

$$Z_a = \min \sum_{v \in V} \sum_{l \in L} U_{vl} \quad (3a)$$

subject to

$$U_{vl} \geq \sum_{k \in K} b_{kld} W_{vkd} \quad \forall v \in V, l \in L, d \in D \quad (3b)$$

$$\sum_{k \in K} W_{vkd} \geq 1 \quad \forall v \in V, d \in D \quad (3c)$$

$$\sum_{v \in V} W_{vkd} \leq 1 \quad \forall k \in K, d \in D \quad (3d)$$

$$W_{vkd} \in \{0, 1\} \quad \forall v \in V; k \in K; d \in D \quad (3e)$$

$$U_{vl} \geq 0 \quad \forall v \in V, l \in L \quad (3f)$$

The objective (3a) minimizes the number of cells covered by the drivers. Constraints (3b) set U_{vl} to 1 if driver $v \in V$ is assigned to a vehicle index $k \in K$ that visits cell $l \in L$ on at least one day. Constraints (3c) ensure that each driver is assigned to a vehicle index on each day. Constraints (3d) ensure that only one driver is assigned to a vehicle index on a given day. Constraints (3e) and (3f) define the decision variables (note that U_{vl} is binary, given binary values for W_{vkd}).

Given a solution $(\hat{\mathbf{U}}, \hat{\mathbf{W}})$ to the assignment problem, we can calculate driver-dependent metrics. Recall that driver coverage measures the portion of the geographic area covered by drivers. For each driver $v \in V$, the number of cells visited is $\sum_{l \in L} \hat{U}_{vl}$. Driver coverage is defined as the ratio of the number of cells visited to the total number of cells:

$$\theta_v = \frac{1}{|L|} \sum_{l \in L} \hat{U}_{vl} \quad (4a)$$

The average driver coverage for a given PVRP-SC solution is:

$$\theta = \frac{1}{|V||L|} \sum_{v \in V} \sum_{l \in L} \hat{U}_{vl} = \frac{Z_a}{|L||V|} \quad (4b)$$

where Z_a is the objective value for Formulation (3). Clearly, the number of cells and vehicles affects the possible values of θ . We would expect $\theta \approx \frac{1}{|V|}$ for solutions that equally partition neighborhoods between drivers (with no overlap). High values of the average driver coverage, $\theta \gg \frac{1}{|V|}$, indicate a complex solution in which drivers may visit many neighborhoods.

Using the solution $\hat{\mathbf{W}}$ from Formulation (3) and the PVRP-SC solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$, we can determine which drivers visit a node during the period. Let indicator $e_{iv} = 1$ if node $i \in N$ is visited by driver

$v \in V$ during the period and 0 otherwise. For each node $i \in N$ and driver $v \in V$, we have:

$$e_{iv} = \begin{cases} 1 & \text{if } \sum_{d \in D} \sum_{k \in K} \sum_{j \in N} \hat{x}_{ijk}^d \hat{W}_{vkd} \geq 1 \\ 0 & \text{otherwise} \end{cases} \quad (5a)$$

We calculate the crewsize for a node i over the period as:

$$\phi_i = \sum_{v \in V} e_{iv} \quad (5b)$$

The average crewsize in the PVRP-SC solution is:

$$\phi = \frac{1}{|N|} \sum_{i \in N} \phi_i \quad (5c)$$

Accordingly, ϕ ranges from 1 to $|V|$. A high value of ϕ indicates that, on average, many drivers visit a node, which may be undesirable in applications which require drivers to have knowledge/training specific to customer locations (configuration of the facility layout, security clearance, etc.).

In the numerical analysis in Section 3, we evaluate PVRP and PVRP-SC solutions relative to the following metrics: objective function Z ; average arrival span σ ; average driver coverage θ ; and average crewsize ϕ . Additionally, we consider the computational complexity of the solution methods with solution times.

3 Numerical analysis

Section 3.1 evaluates the Tabu Search heuristic relative to the exact method from Francis et al. [11]. Section 3.2 introduces the test cases in the numerical studies. Section 3.3 introduces the measures used in the numerical analysis and Section 3.4 presents and analyzes the results.

3.1 Evaluation of Tabu Search

We implement the Tabu Search heuristic in C++ and execute on a Sun Fire 150 workstation with two UltraSPARC III processors. To evaluate the Tabu Search method, we use problem instances from Francis et al. [11] solved with the exact method. The test cases range in size from 12 to 40 nodes, with 3 and 4 vehicles, and various capacity levels. We impose the same assumptions as Francis et al. [11] (i.e., allowing service choice, no crew flexibility, a schedule set of {Mon-Wed-Fri, Tue-Thr, daily}, and delivery of accumulated demand).

Table 1 compares the average difference (aggregated over vehicle and capacity values) in the two solution methods, examining the percentage differences between the objective, the performance

Nodes	Optimality Gap	σ Gap	θ Gap	Change in sol. time
12	0.0%	0.3%	0.1%	81.2%
16	0.1%	0.2%	0.4%	-5.2%
20	0.0%	0.3%	0.4%	-39.0%
28	0.1%	0.4%	0.3%	-59.3%
36	0.3%	0.7%	0.5%	-71.9%
40	0.4%	0.9%	0.7%	-84.6%

Table 1: Comparison of Tabu Search solutions with exact solutions

measures σ and θ , and the solution time. Since there is no crew flexibility, the crewsize measure $\phi = 1$ in all cases. The average objective values of Tabu Search solutions are within 0.4% of optimality. The operational complexity of solutions obtained by the two methods differ by less than 1% across all above measures. Thus, the Tabu Search finds solutions that are close to optimal and provide a good representation of the operational complexity of the optimal solutions given the assumptions of the exact method.

An important advantage of the Tabu Search method is its significantly lower computation times compared to the exact method. The exact solution method takes more than 8 hours for instances of the PVRP-SC with 40 nodes, yet the Tabu Search can solve these instances in roughly 40 CPU minutes. Additionally, the Tabu Search obtains solutions within a reasonable amount of time for larger instances that cannot be solved with the exact method. While the heuristic version of the exact method could be used with the precision (δ) set to high values for larger instances, we expect that the Tabu Search method would outperform such solutions, both in terms of solution quality and solution times. Further, even with large δ values, the heuristic version cannot be used to test all flexibility levers due to modeling and computational limits. Given the ability of the Tabu Search to incorporate flexibility levers, as well as its speed and quality, we use the Tabu Search to produce solutions for our analysis.

3.2 Test cases

Francis et al. [11] find that the routing efficiency gains from service choice in the PVRP-SC are impacted by the geographic distribution of customer nodes. We examine how geographic distribution impacts the gains from other flexibility levers. Further, we examine how these flexibility levers in turn impact the magnitude of savings from service choice.

We create a set of randomized instances in idealized patterns of various city configurations. Figure 2 displays the four configurations of the service region for instances with characteristics of demand patterns and node distributions that mimic trends in cities today. Each city covers a circular area of 100 miles diameter, and the service region is divided into square cells ($10mi \times 10mi$) for calculation of the driver-dependent metrics. Cells that do not contain nodes are eliminated from the set L . In all configurations, the depot is located in the center of the region. The service region is divided by customer visit requirements over a five-day period (either 2 days, 3 days, or 5 days).

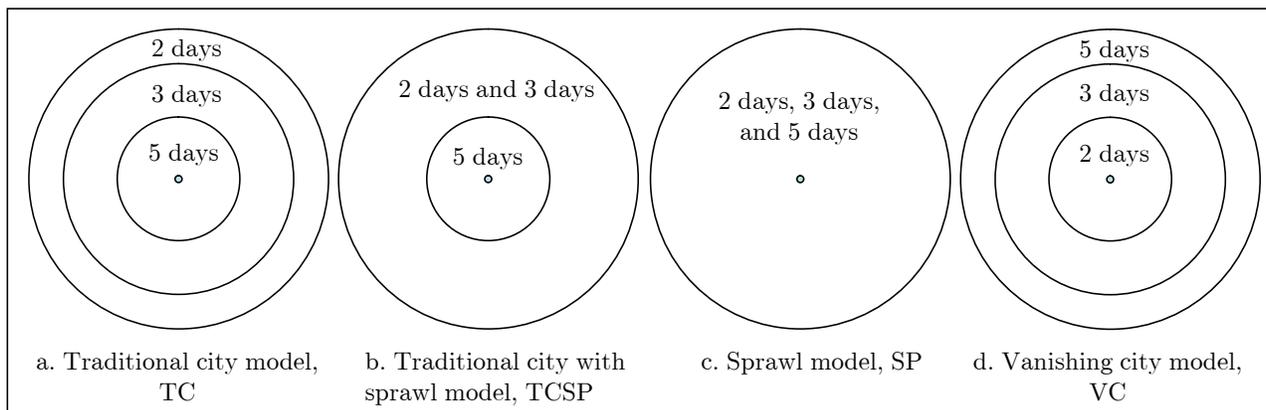


Figure 2: Test case configurations

Configuration TC in Figure 2(a) represents a traditional city in which high demand customers are located near the center of the region and demand density decreases with distance from the center. Configuration TCSP in Figure 2(b) is a variation of the traditional city in which there is a mix of low and moderate demand levels beyond the region of high demand in the center. Configuration SP in Figure 2(c) represents modern sprawl in which demand levels are scattered throughout the region with no central business district. Configuration VC in Figure 2(d) represents a city in which high demand has left the central business district and moved to the outlying areas.

Numerical studies from Francis et al. [11] suggest that test cases resembling Configuration TC observe significant improvements in routing efficiency and customer service with service choice flexibility, while cases resembling Configuration SP experience lower routing efficiency gains (while still increasing customer service). We use the city configurations here to study the interplay between geographic node distribution and a broader range of operational flexibility.

Ten problem instances with 200 nodes are randomly generated for each configuration. Nodes are uniformly scattered within each subregion. Demands are drawn from a truncated normal

distribution (mean 125, standard deviation 100), and the resulting demand value is accepted if it fits the subregion characteristics, given as follows: 0 to 75 items/day for $f_i= 2$ days, 76 to 150 for $f_i= 3$ days, and 151 to 250 for $f_i= 5$ days. Otherwise, the demand value is discarded and another realization is created.

Label	$ S $	Visit requirements	Schedule set
<i>A</i>	3	2,3,5 days	$(TR), (MWF), (MTWRF)$
<i>B</i>	5	2,3,5 days	$(TR), (TF), (MWF), (MWR), (MTWRF)$
<i>C</i>	5	1-5 days	$(W), (TR), (MWF), (MTRF), (MTWRF)$
<i>D</i>	7	1-5 days	$(W), (TR), (TF), (MWF), (MWR),$ $(MTRF), (MTWRF)$
<i>E</i>	10	1-5 days	$(W), (R), (F), (MR), (TR), (TF),$ $(MWF), (MWR), (MTRF), (MTWRF)$

Table 2: Schedule options for five-day test cases

For each test case, we consider the sets of schedule options listed in Table 2. The first column lists the label for the schedule set. The second column lists the number of schedules in the set S and the third column lists the visit requirement values (f_i) for the nodes. The schedule set is shown in the fourth column. The set *A* includes only disjoint schedules and their union, which is used with the exact solution method in Francis et al. [11]. All other sets include non-disjoint schedules which cannot be easily incorporated into the exact method. The service benefit (α^s) values are 0.05, 0.1, 0.15, 0.175, and 0.2 for schedules with γ^s values of 1, 2, 3, 4 and 5 respectively¹.

3.3 Measuring efficiency and complexity

We examine the impact of introducing flexibility levers in pairwise comparisons with and without flexibility. The superscript *cons* denotes the value of the constrained solution and the superscript *flex* denotes the value of the flexible solution, both with respect to a certain flexibility lever.

¹Benefits specific to each schedule option can be incorporated in the Tabu Search procedure easily.

$$\text{Objective improvement } \Delta_Z = \frac{Z^{flex} - Z^{cons}}{Z^{cons}} \quad (6)$$

$$\text{Arrival span complexity rise } \Delta_\sigma = \frac{\sigma^{flex} - \sigma^{cons}}{\sigma^{cons}} \quad (7)$$

$$\text{Driver coverage complexity rise } \Delta_\theta = \frac{\theta^{flex} - \theta^{cons}}{\theta^{cons}} \quad (8)$$

$$\text{Crewsize complexity rise } \Delta_\phi = \frac{\phi^{flex} - \phi^{cons}}{\phi^{cons}} \quad (9)$$

The objective improvement is decomposed into two parts: the contribution due to the routing cost component, $\Delta_{Z_{c,\tau}}$, and the contribution due to the service benefit component Δ_{Z_α} . The driver coverage metric θ is bounded by $\frac{1}{|V|} \leq \theta \leq 1$. Hence, the corresponding change in driver coverage is bounded by $-(|V| - 1) \leq \Delta_\theta \leq (|V| - 1)$. The value of crewsize ϕ is bounded by $1 \leq \phi \leq |V|$, and the corresponding change by $-(|V| - 1) \leq \Delta_\phi \leq (|V| - 1)$.

While we cannot derive such bounds for Δ_Z and Δ_σ , we use the example in Table 3 to illustrate how the magnitude of the measures may be interpreted. For a Traditional City instance with 200 nodes and 5 vehicles, the PVRP-SC solution with schedule set A and without crew flexibility yields an objective value of $Z = \$3,961$ (routing cost: \$5,082; service benefit: \$1,121), $\sigma = 1.86$ hours, $\theta = 0.22$ (15.4/70 cells), and $\phi = 1$ driver, shown in the base case column of the table.

Metric	Base case	Variation 1		Variation 2	
	Value	Value	Δ	Value	Δ
Z	\$3,961	\$3,945	-0.4%	\$3,640	-8.1%
σ	1.86 hours	1.9 hours	2.2%	2.2 hours	18.3%
θ	0.22 (15.5/70)	0.23 (16.1/70)	4.5%	0.25 (17.5/70)	13.6%
ϕ	1 driver	1.04 drivers	4.0%	1.18 drivers	18.0%

Table 3: Example of metric changes

Adding crew flexibility (Variation 1 in the table), the objective function decreases by \$16 and $\Delta_Z = -0.4\%$. Further, we observe $\Delta_\sigma = 2.2\%$, corresponding to a rise in arrival span σ to 1.9 hours. Similarly, $\Delta_\theta = 4.5\%$, corresponding to a rise in θ to 0.23 (16.1/70 cells). Finally, $\Delta_\phi = 4.0\%$ and $\phi = 1.04$ drivers per customer (2 customers are visited by 3 drivers, 5 customers are visited by 2 drivers and the remaining 193 customers by a single driver). This is representative of a small change in the solution.

Increasing the set of schedules from A to E (Variation 2 in the table), we obtain $\Delta_Z = -8.1\%$,

with 6.1% savings due to routing, and 2.0% savings due to service benefit increases (12 more nodes are served at higher frequencies). We observe $\Delta_\sigma = 18.3\%$, with arrival span σ rising to 2.2 hours. We observe $\Delta_\theta = 13.6\%$ as θ rises to 0.25 (17.5/70 cells). In the flexible solution, $\Delta_\phi = 18\%$. Here, 2 nodes are visited by 3 drivers, 32 nodes by 2 drivers, and the remaining 166 nodes by 1 driver, yielding $\phi = 1.18$ drivers per customer. This is representative of a large change in the solution.

3.4 Effect of flexibility

In what follows, we examine the effect of the flexibility levers on the performance metrics. Section 3.4.1 shows the effect of crew flexibility and Section 3.4.2 shows the effect of flexibility in schedule options. Section 3.4.3 considers the tradeoffs between flexibility and complexity when schedule options and service choice are considered together. Section 3.4.4 considers the effect of flexibility in delivery strategies. Finally, Section 3.4.5 presents managerial insights from our analysis.

3.4.1 Crew flexibility

Table 4 shows the average percentage change in the metrics between solutions without crew flexibility (*cons*) and with crew flexibility (*flex*). In both cases, we use schedule set E which provides the most schedule options. We show the metrics under three conditions: without service choice (PVRP); with service choice and $\beta = 0$, (PVRP-SC_{routing}); and with service choice and $\beta = 1$ (PVRP-SC). The PVRP-SC_{routing} option allows for routing improvements due to frequency changes, but these changes do not affect the objective function through the service benefit term.

	PVRP					PVRP-SC _{routing}					PVRP-SC				
	$\Delta_{Z_{c,\tau}}$	Δ_{Z_α}	Δ_σ	Δ_θ	Δ_ϕ	$\Delta_{Z_{c,\tau}}$	Δ_{Z_α}	Δ_σ	Δ_θ	Δ_ϕ	$\Delta_{Z_{c,\tau}}$	Δ_{Z_α}	Δ_σ	Δ_θ	Δ_ϕ
TC	-0.4	N/A	0.4	2.8	2.2	-1.8	N/A	1.5	7.8	6.1	-1.9	0.0	1.6	7.8	6.1
TCSP	-0.4	N/A	0.4	2.6	2.3	-1.3	N/A	0.8	6.6	5.4	-1.3	-0.1	0.9	6.6	5.4
SP	-0.3	N/A	0.3	2.2	2.1	-1.2	N/A	1.0	6.7	5.4	-1.1	-0.2	1.0	6.7	5.4
VC	-0.1	N/A	0.4	2.4	1.8	-1.7	N/A	1.4	7.0	5.9	-1.6	-0.1	1.4	7.0	5.9

Table 4: Effects of introducing crew flexibility (all values are in %)

The average objective improvement in all cases in Table 4 is consistently less than 2%, with uniform variance of less than 1% in individual observations. While crew flexibility appears to improve the PVRP-SC objectives more than PVRP objectives, the improvement is still small. PVRP-SC solutions have a larger feasible region (more schedule choices are possible); therefore, crew flexibility

expands the feasible region more when used in conjunction with service choice flexibility. For the PVRP-SC, the relative magnitude of the change in the two components of the objective indicate that most of the savings are due to improved routing efficiency without significant change in the service allocation of the nodes. Overall the geographic configuration does not have a large effect on the objective improvement.

As expected, adding crew flexibility has a noticeable impact on the complexity measures θ and ϕ since crew flexibility expands the set of solutions by specifically those solutions that are characterized by higher values of individual θ (visiting more nodes increases U_{vl} for drivers) and ϕ (by definition). Again, this impact is larger when service choice is allowed since more solutions are feasible. The change in σ is not significant since nodes are visited in an order which is affected mostly by their position relative to the depot, and less by changing the vehicle routes (particularly in dense delivery areas).

From a practical standpoint, these results suggest that system regularity can be enforced by removing crew flexibility, without significantly affecting the objective function. The benefit from removing crew flexibility is in creating solutions with driver delivery districts that have fewer overlapping areas, which reduces complexity for customers and drivers. Solution times for the Tabu Search are not significantly affected by allowing or removing crew flexibility, with the average solution time of 190 minutes with a standard deviation of less than 8 minutes across all instances.

3.4.2 Schedule options

Table 5 shows the average percentage change in the metrics, comparing solutions with schedule set A (*cons*) with solutions with schedule set E (*flex*). Crew flexibility is used in both cases to allow the system to choose the best vehicle assignments for all days which allows us to examine the unrestricted change in crewsize under schedule option flexibility.

	PVRP					PVRP-SC _{routing}					PVRP-SC				
	$\Delta_{Z_{c,\tau}}$	$\Delta_{Z_{\alpha}}$	Δ_{σ}	Δ_{θ}	Δ_{ϕ}	$\Delta_{Z_{c,\tau}}$	$\Delta_{Z_{\alpha}}$	Δ_{σ}	Δ_{θ}	Δ_{ϕ}	$\Delta_{Z_{c,\tau}}$	$\Delta_{Z_{\alpha}}$	Δ_{σ}	Δ_{θ}	Δ_{ϕ}
TC	-2.8	N/A	3.3	1.7	2.1	-5.4	N/A	8.3	3.1	4.9	-5.3	-2.9	9.8	4.9	6.2
TCSP	-2.5	N/A	2.9	0.7	1.6	-4.8	N/A	7.1	1.9	3.6	-4.7	-2.7	8.5	2.1	4.8
SP	-2.1	N/A	2.3	1.0	1.8	-3.9	N/A	6.8	2.8	4.0	-3.8	-2.4	7.2	3.0	7.9
VC	-1.5	N/A	1.0	1.1	1.9	-3.2	N/A	3.1	2.1	4.2	-2.9	-1.4	3.3	3.3	5.6

Table 5: Effects of increasing number of schedules (all values in %)

Table 5 suggests that the geographic distribution of nodes impacts efficiency and complexity when increasing the number and type of schedule options. While the individual difference between instances results in variances ranging from 0 to 4% for $\Delta_{Z_{c,\tau}}$, increasing schedule options appears to be more beneficial for configurations with high frequency nodes distributed close to the depot. In other cases, such as Configurations SP and VC, limiting the set of schedules is less costly. There are two reasons for this variation between centralized demand configurations and dispersed demand configurations. First, the magnitude of routing costs differs between the configurations. Routing costs tend to be higher when the most frequently visited nodes are far away from the depot, as is the case in Configurations VC and SP. Hence, the same absolute improvement in objective appears smaller for dispersed configurations as opposed to the centralized configurations. Second, the presence of nodes with high visit requirements at the outer periphery requires vehicles to serve neighborhoods near such nodes on each day of the week. Nodes lying on the path of such routes, or near the outer periphery, can receive higher levels of service since the marginal cost of including these nodes on the routes is relatively low (Francis et al. [11] illustrate this principle). In fact, when the fixed portion of the stopping cost τ is negligible, nodes receive higher service as long as the increased benefit offsets the marginal cost, subject to vehicle capacity constraints. Hence, adding additional flexibility by extending the set of schedule options has less value for the dispersed configurations as opposed to the centralized configurations.

Including service benefit in the objective of the PVRP-SC has the expected effect of raising visit frequency as more schedule options are allowed. The improvement from raising frequencies ranges from 1.4 - 2.9%. In PVRP-SC_{routing} solutions, visit frequencies are raised as well in some instances to increase routing efficiency. If these changes were rewarded in the objective function, they would contribute an improvement of 0.6 - 1.6% in Z_α (this is not shown in the table since service benefits in PVRP-SC_{routing} are not rewarded).

Schedule set A is designed specifically to limit the complexity of driver routes; i.e., drivers perform at most two different routes. As expected, complexity measures rise when schedule set E is used. The increase in arrival span σ suggests that adding more flexibility in schedule options is at the expense of consistency in customer visit times. The day-to-day composition of routes has more variance when more diverse schedule choices are available and hence the number of nodes to be visited changes from day to day. A similar effect corresponds to ϕ and θ . Further, these metrics rise with service choice since the number of customers visited each day increases, which may force vehicles to cover wider areas. When service benefit is considered in the objective, the

metrics increase more since the routes generated are not as efficient as those created when routing cost is the only objective.

Finally, solution times for the Tabu Search are sensitive to the size of the schedule set $|S|$. For our 200-node data sets, the solution times are found to increase from 168 CPU minutes for $|S| = 3$ to about 192 CPU minutes for $|S| = 10$ with a standard deviation of less than 12 minutes across all instances.

3.4.3 Interaction of service choice and schedule options

The preceding analysis suggests that introducing flexibility in schedule options has an impact on the objective function, with or without service benefit in the objective. Next, we examine the trade-offs between flexibility and complexity when increased schedule options and service choice are considered together. We include crew flexibility to allow the system the flexibility to pick the best vehicle assignment for each visit to a node.

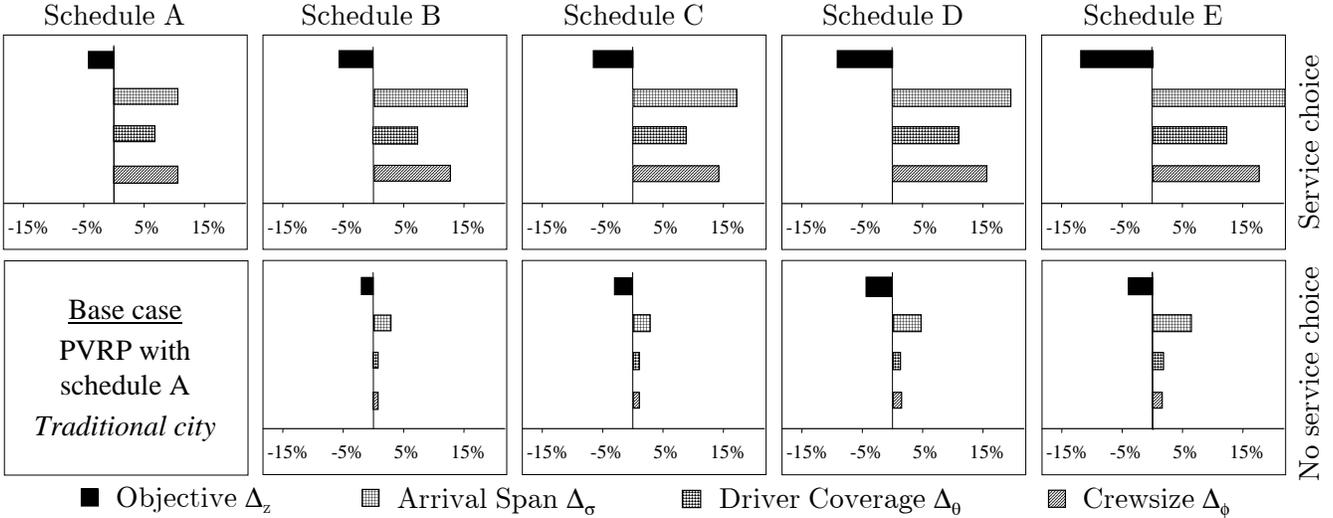


Figure 3: Increasing flexibility in node visitation: Traditional city configuration

Figures 3 and 4 illustrate how the performance metrics change as both service choice and schedule option flexibility are introduced for Configurations TC and VC, respectively. The results for Configurations TCSP and SP (not shown here) lie between the Configurations TC and VC in terms of changes in operational performance and operational complexity. Hence, Configurations TC and VC are representative of the extreme conditions of geographic configurations. Along the horizontal direction in the figures, from left to right, we change the schedule options from A to E; along the vertical direction, we introduce service choice (PVRP-SC solutions) in the upper row and

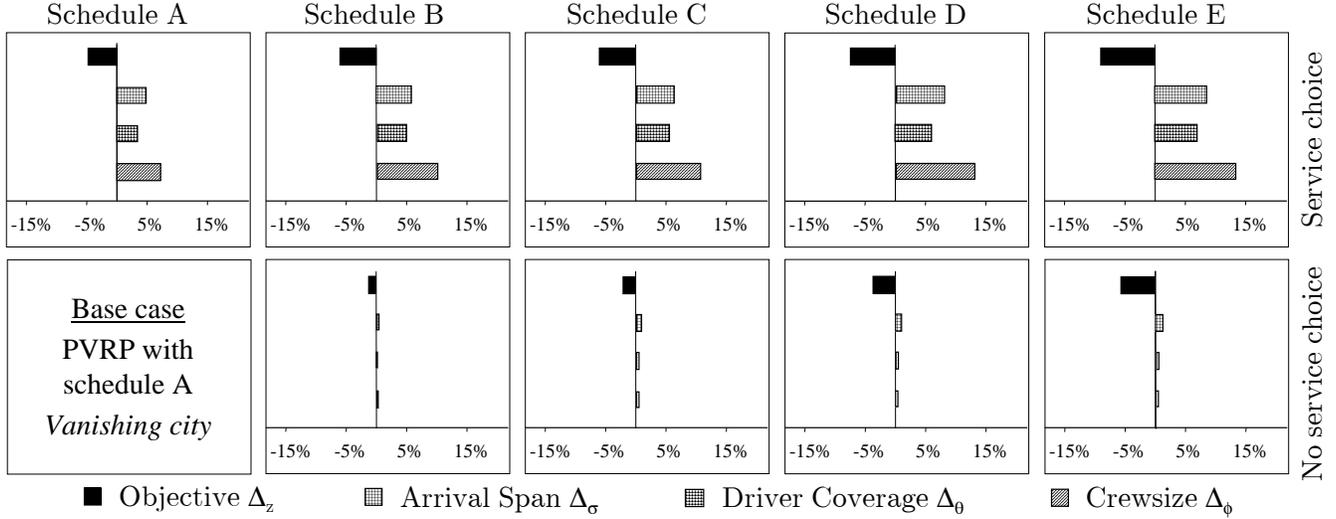


Figure 4: Increasing flexibility in node visitation: Vanishing city configuration

restrict service choice (PVRP solutions) in the lower row. Recall that schedule set C has the same number of schedule options as B , but a wider range of visit requirements. All metrics are measured with respect to a common base case, the PVRP with schedule set A (lower left corner).

The figures indicate steady improvement in the objective function as schedule option flexibility is introduced. Improvements in the objective function are accompanied by increases in operational complexity in most cases. The effects are larger for Configuration TC than any other configuration as in the previous section. Also, as before, the PVRP with VC configuration has the least increase in complexity measures. Adding three schedules from schedule set D to E results in a 1.4% change in Δ_z on the average, with variance in individual observations ranging from 0% to 4%. The additional schedule options increase the solution time of the Tabu Search, on average, by 7 minutes.

The rises in crewsize and driver coverage indicate that adding flexibility through service choice and/or new schedule options affects the number of different drivers to train for operations specific to each customer location. There are also significant differences in the arrival span. While all city configurations appear to be affected, Configuration TC, which shows the highest efficiency gains, also has the greatest change in arrival span with increased flexibility.

Displaying the metrics for varying levels of schedule options and service choice relative to a common base case facilitates comparisons of the relative costs and benefits of flexibility levers. Using Figures 3 and 4, it is possible to compare the relative complexity increases for a given improvement in the objective function for flexibility in service choice versus flexibility in schedule options. The figures indicate that flexibility in schedule options can produce efficiency gains comparable to

service choice with lower rises in complexity. For example, in Configuration TC, using service choice flexibility with schedule set A yields a 5% improvement in the objective function over the base case with complexity increases between 5% and 10%. Using set E without service choice also yields an objective function improvement of 5%, but with smaller rises in complexity.

Service providers should evaluate the relative gains from increased operational efficiency against operational costs such as driver training and possible customer dissatisfaction. These metrics provide a way of quantifying the changes. For instance, consider a distribution operation offering schedule set A to customers in a city of Configuration VC. If transportation costs are high compared to the cost of training drivers to visit different customers and regions, then a 5% increase in operational efficiency by introducing service choice may justify a 4% increase in average driver coverage and 8% increase in crewsize. A relative weighting of these metrics is likely to be application specific depending on the costs associated with increased complexity.

3.4.4 Delivery strategies

The PVRP/PVRP-SC literature assumes that the amount delivered at each customer visit is equal to the demand accumulated since the last visit. In this section, we explore how flexibility in delivery amounts can improve operational efficiency of the PVRP-SC. In order to model this flexibility, we first look at the way in which demand accumulation is modeled in periodic routing problems.

In the PVRP literature, demand accumulation is modeled as the average accumulation between visits. In the PVRP-SC model of Francis et al. [11], the accumulation is modeled as the maximum demand accumulation between visits. With both approximations, the delivery quantities, w_i^s , can be determined for node $i \in N$ and schedule $s \in S$, and the day is not needed in the capacity constraints (1d). The true accumulation includes the day, w_i^{sd} , which significantly increases the complexity of the model (constraint (1d) in particular) for the exact method of Francis et al. [11].

Using average accumulation may lead to capacity-infeasible solutions if capacity is tight and the time between visits is not uniform in certain schedules. This approximation is reasonable if at least one of the following conditions is satisfied:

1. Demand at each node does not vary significantly over the period, and the time between visits is uniform for all schedules. Further, there is sufficient slack in the vehicle capacity to accommodate the existing variability.
2. Customers are willing to accept average delivery amounts rather than the requested delivery

amount (thereby incurring shortages or carrying additional inventory).

Similarly, the maximum accumulation approximation is reasonable when Condition 1 above holds or if customers are willing to accept more deliveries in excess of the accumulated demand. Using the maximum accumulation guarantees feasibility but may lead to suboptimal solutions. In practice, when operating under the maximum accumulation modeling, only the required demand is delivered, but vehicle capacity is reserved for the maximum accumulation, which is used fully at least once during the period.

Other solution methods (and in particular our Tabu Search method) may consider the true demand accumulation between visits, which can incorporate non-uniform separation between visits, as well as non-stationary demand and service choice.

Delivery flexibility allows the delivery amount to become a decision variable so that efficiency is increased. As such, the PVRP-SC begins to resemble the Inventory Routing Problem (IRP), in which the amount delivered is a separate decision variable. We consider two ways in which the PVRP-SC can be modeled as a special case of the IRP with deterministic demand. The first is an IRP where no shortages are allowed, a zero inventory policy is followed, and there exists a limited set of visit frequencies. As a result, a vehicle always delivers an amount exactly equal to the demand accumulated between visits. However, unlike the traditional PVRP-SC, the service benefit term is modeled as the cost of holding inventory between visits rather than as the benefit of increased frequency. The service benefit term depends on the demand at the node, as well as the time that each unit of demand is held. In the second case, we relax the assumptions of a zero inventory policy and allow shortages. Allowing shortages guarantees feasibility in cases that are not feasible in the first case. Note that continuous IRP models may choose any amount to be delivered at nodes; however, we significantly limit the delivery choices to solve this variation in a reasonable amount of time using the Tabu Search method. In particular, we allow either the delivery of the average demand or the true accumulated demand. The maximum-demand strategy is excluded because it reserves more capacity than required on the vehicle, resulting in less efficient routing solutions.

Two delivery options are compared to measure the benefit of modeling delivery strategy: (1) no delivery flexibility (*cons*) – a PVRP-SC in which the service benefit is modeled as holding and shortage cost savings and exactly the true demand is delivered every time; (2) delivery flexibility (*flex*) a PVRP-SC in which service benefit is modeled as holding and shortage cost savings, and the system can choose between delivering the average-demand amount or the true-demand amount. For each node, a cost is assigned to each schedule based on the holding and shortage costs of the

amount delivered, the demand, and the visit days. We use a holding cost of \$0.05 per item per day, and a shortage cost of \$0.1 per item per day.

The Tabu Search is modified to solve these special cases of the IRP as follows. Without delivery flexibility, we model $|S|$ schedule options. With delivery flexibility, we create copies of each schedule for each delivery strategy. In this case, there are two delivery strategies (average amount and required amount), yielding $2|S|$ schedule options. Thus, when considering candidate moves in each Tabu Search iteration, both the frequency and the accumulation option of a candidate schedule are simultaneously evaluated. We solve the problem for the 200-node, city configuration test cases assuming crew flexibility and schedule set E. The average solution times increases from 194 minutes for $|S| = 10$ to about 486 minutes for $|S| = 20$. While this approach can be used to consider a wider range of delivery options, the increase in solution time limits the number of options that can be practically considered.

Table 6 shows the effect of adding delivery flexibility on the objective function and the complexity measures. The objective improvements are aggregated over all instances of each configuration and vary slightly by city configuration type. However, there exists a wide variance in individual observations which suggests that the choice of delivery strategy is partially dictated by other factors such as route design and capacity utilization of vehicles. Nevertheless, the ability to control the amount of demand delivered is shown to be a useful means of flexibility.

	$\Delta_{Z_{c,\tau}}$	Δ_{Z_α}	Δ_σ	Δ_θ	Δ_ϕ
TC	-4.8%	-2.6%	6.8%	2.4%	4.2%
TCSP	-4.1%	-2.4%	6.4%	2.2%	4.0%
SP	-4.0%	-2.6%	6.4%	2.6%	4.4%
VC	-3.6%	-1.8%	6.6%	2.4%	4.0%

Table 6: Impact of introducing delivery flexibility

Overall, for the given inventory costs, adding the delivery strategy flexibility is found to be beneficial. The objective function improves by 2-13% across individual instances when the system is allowed to choose between delivering the average amount and true amount. These savings are achieved with changes in delivery quantity for only a small number of nodes. In all cases, fewer than 20 nodes are served using average demand rather than true demand, with the average number of such nodes ranging between 12 and 20. The resulting change in the objective appears to be partly due to more efficient routing and vehicle assignments made possible by the demand adjustments,

and partly due to the savings in holding/shortage costs. The exact contribution of these two components to the objective improvements varies widely among instances, even within instances of the same configuration type.

3.4.5 Managerial observations

We make the following managerial observations based on the findings in this section:

1. *Tradeoff between flexibility and complexity.* As expected, introducing operational flexibility increases the operational complexity of the solutions. In most cases, the increase in the complexity is related to the efficiency gains obtained; however, certain levers of operational flexibility (such as crew flexibility) tend to increase complexity without corresponding efficiency gains. The complexity measures facilitate the choice of flexibility levers to maximize efficiency gains with allowable complexity increases.
2. *Significance of geographic distribution.* In general, the results confirm earlier results in Francis et al. [11] on the significance of geographic distribution on value of service choice. Further, the results extend to other flexibility levers as well. The results indicate that introducing flexibility is more beneficial when high frequency nodes are located near the depot (as in Configurations TC and TCSP).
3. *Effect of crew flexibility.* Restricting crew flexibility is often required by customers (e.g., inter-library loan application of Francis et al. [11]). It is found to have a limited effect on the objective, which suggests that reducing operational complexity in this way may be desirable.

4 Conclusions and future research

We provide insights from both a managerial and a modeling perspective on the trade-offs between operational flexibility and complexity in periodic vehicle routing problems. Specifically, we quantify the gains from operational flexibility in terms of vehicle routing costs and customer service benefits, along with the costs of additional complexity in terms of modeling and implementation difficulty. We identify four levers of operational flexibility – service choice in determining customer visit frequency, crew flexibility that expands the number of drivers visiting nodes, schedule options offered by the service provider, and the delivery quantity at each visit. We show how these levers can be modeled and analyze their effect on the efficiency and complexity of resulting solutions.

We provide three metrics to quantify the operational complexity of the resulting solution – arrival span, driver coverage, and crewsize. These are the first known metrics in the literature to quantify the desirability of routing solutions in a periodic distribution context.

We introduce a Tabu Search method that can incorporate a wide range of flexibility options. The Tabu Search method obtains solutions within 3% of optimality for test cases from the literature. We quantify the operational savings from adding flexibility to periodic distribution as a function of geographic dispersion of nodes using the Tabu Search method.

The complexity measures considered in this paper are either considered endogenously through variable and parameter definitions or exogenously in post-processing. Future work could focus on adding complexity measures into the objective function of the PVRP-SC, thereby allowing the solution method to choose the appropriate balance between complexity and flexibility. In the routing literature, time windows for node visits have been incorporated with soft penalties for violations, which could form the basis for adding soft penalties for variations in visit times for nodes across days in the PVRP-SC. Further, variation in driver routes could be incorporated in the objective function. Such extensions would involve parametric analysis of the relative weighting of complexity costs to operational benefits. Further, when the complexity costs are known, we could determine a frontier of efficient solutions for different levels of complexity.

Acknowledgment

This research has been supported by grant DMI-0348622 from the National Science Foundation, the Northwestern University Transportation Center, and the Alfred P. Sloan Foundation.

References

- [1] E. Angelelli and M. G. Speranza. The periodic vehicle routing problem with intermediate facilities. *European Journal of Operational Research*, 137(2):233–247, 2002.
- [2] S. Anily and J. Bramel. Vehicle routing and the supply chain. In R. G. S. Tayur and M. Magazine, editors, *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, 1998.
- [3] S. Anily and A. Federgruen. One warehouse multiple retailer systems with vehicle routing costs. *Management Science*, 36(1):92–114, 1990.

- [4] E. Beltrami and L. Bodin. Networks and vehicle routing for municipal waste collection. *Networks*, 4:65–94, 1974.
- [5] F. Blakely, B. Bozkaya, B. Cao, W. Hall, and J. Knolmayer. Optimizing periodic maintenance operations for Schindler Elevator Corporation. *Interfaces*, 33(1):67–79, 2003.
- [6] L. Chan, A. Federgruen, and D. Simchi-Levi. Probabilistic analysis and practical algorithms for inventory-routing models. *Operations Research*, 46(1):96–106, 1998.
- [7] I.-M. Chao, B. Golden, and E. Wasil. An improved heuristic for the period vehicle routing problem. *Networks*, 26:25–44, 1995.
- [8] J.-F. Cordeau, M. Gendreau, and G. Laporte. A Tabu Search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, 30:105–119, 1997.
- [9] A. Federgruen and D. Simchi-Levi. Analysis of vehicle routing and inventory-routing problems. In *Handbooks in Operations Research and Management Science*, volume 8, pages 297–373. Elsevier Science B.V., 1995.
- [10] P. Francis and K. Smilowitz. Modeling techniques for periodic vehicle routing problems. *Transportation Research: Part B*, 40:872–884, 2006.
- [11] P. Francis, K. Smilowitz, and M. Tzur. The period vehicle routing problem with service choice. *Transportation Science*, 2006. Forthcoming.
- [12] M. Gendreau, A. Hertz, and G. Laporte. A Tabu Search heuristic for the vehicle routing problem. *Management Science*, 40(10):1276–1290, 1994.
- [13] F. Glover and M. Laguna. *Tabu Search*. Kluwer Academic Publishers, Norwell, MA, USA, 1997. ISBN 079239965X.
- [14] A. Kleywegt, V. Nori, and M. Savelsbergh. The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36(1):94–118, 2002.
- [15] A. Rusdiansyah and D. Tsao. An integrated model of the periodic delivery problems for vending-machine supply chains. *Journal of Food Engineering*, 70(3):421–434, 2005.
- [16] R. Russell and D. Gribbin. A multiphase approach to the period routing problem. *Networks*, 21:747–765, 1991.

- [17] R. Russell and W. Igo. An assignment routing problem. *Networks*, 9:1–17, 1979.
- [18] C. Tan and J. Beasley. A heuristic algorithm for the period vehicle routing problem. *Omega*, 12(5):497–504, 1984.
- [19] UPS Corp. Driving success: Why the UPS model for managing 85,000 drivers is a competitive advantage. <http://pressroom.ups.com/mediakits/popups/factsheet/0,1889,1201,00.html>, January 2006.
- [20] H. Zhong, R. Hall, and M. Dessouky. Territory planning and vehicle dispatching with driver learning. *Working paper*; <http://www-rcf.usc.edu/maged/publications/TerritoryPlanning.pdf>, 2004.