INDUSTRY STUDIES ASSOCIATION

INDUSTRY STUDIES ASSOCIATION
WORKING PAPER SERIES

Flexibility and Complexity in Periodic Distribution Problems

By

Peter Francis
Department of Industrial Engineering and Management Sciences
Northwestern University
Evanston, IL 60208

Karen Smilowitz
Department of Industrial Engineering and Management Sciences
Northwestern University
Evanston, IL 60208

Michal Tzur
Department of Industrial Engineering
Tel Aviv University
Israel

2005
Industry Studies Association
Working Papers

WP-2005-09
http://isapapers.pitt.edu/

Flexibility and complexity in periodic distribution problems

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

March 3, 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 implementation difficulty. Periodic distribution problems arise in a number of industries, including food distribution, waste management and mail services. The period vehicle routing problem (PVRP) is a variation of the classic vehicle routing problem in which driver 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. While introducing operational flexibility in periodic distribution systems can increase efficiency, it poses three challenges: the difficulty of modeling this flexibility accurately; the computational effort required to solve the problem as modeled with such flexibility; and the complexity of operationally implementing the resulting solution. This paper considers these trade-offs between the system performance improvements due to operational flexibility and the resulting increases in operational and computational complexity as they relate to periodic vehicle routing problems. In particular, increasing the operational complexity of driver routes can be problematic in industries where some level of system regularity is required. As discussed in the paper, recent work in the literature suggests that dispatching drivers consistently to the same geographic areas results in driver familiarity and improved driver performance. Additionally, having the same driver visit a customer on a continual basis can foster critical relationships. According to UPS, such driver-customer relationships are a key competitive advantage in its package delivery operations, attributing 60 million packages a year to sales leads generated by drivers. In this paper, we develop a set of quantitative measures to evaluate the trade-offs between flexibility and complexity.

*Department of Industrial Engineering and Management Sciences, Northwestern University

†Department of Industrial Engineering, Tel Aviv University

Introduction

In periodic distribution problems, vehicles visit customers along routes over a given period of time. Customers may be heterogeneous in their demand levels, visit requirements, and willingness-to-pay for more frequent service (visits). In this paper, we consider methods to develop and evaluate strategies to serve heterogeneous customers in periodic distribution systems. In such systems, operational flexibility can avoid under-serving customers with high service requirements and over-serving customers with low requirements. Operational flexibility may be the ability to increase schedule options, decide visit frequency, have different drivers visit customers, and/or decide delivery amounts per visit.

While introducing operational flexibility in periodic distribution systems can increase efficiency, it poses three challenges: the difficulty of modeling this flexibility accurately; the computational effort required to solve the problem as modeled with such flexibility; and the complexity of operationally implementing the resulting solution. This paper considers these trade-offs between the system performance improvements due to operational flexibility and the resulting increases in operational and computational complexity as they relate to periodic vehicle routing problems.

The period vehicle routing problem (PVRP) is an extension of the classic vehicle routing problem in which delivery routes must be constructed over a period of time (for example, multiple days). The PVRP with service choice (PVRP-SC) extends the PVRP to allow service frequency to become a decision of the model. We show that as one introduces flexibility in service choice and delivery strategies into the PVRP, measured by visit frequency, the problem begins to resemble the inventory routing problem (IRP). We investigate the trade-offs mentioned above with respect to operational flexibility and the resulting increases in operational and computational complexity.

We develop a Tabu Search method that can incorporate a range of operational flexibility options and a set of quantitative measures to evaluate the trade-offs between flexibility and complexity in distribution problems.

Section 1 introduces the periodic distribution problems considered in this paper. Section 2 reviews the formulation and exact method for the PVRP-SC from Francis *et al.* (2005) and introduces a new Tabu Search solution method to study more general problems. Section 2 also introduces performance metrics that capture the implementation complexity. Section 3 presents a computational study of the trade-offs between flexibility and complexity. Finally, Section 4 summarizes the paper and discusses future work.

1 Periodic distribution problems

Periodic distribution problems occur 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.* (2005)). The Period Vehicle Routing Problem (PVRP), introduced in Beltrami and Bodin (1974) and Russell and Igo (1979), finds a set of vehicle tours over a period of t 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 (1984), Russell and Gribbin (1991), Chao *et al.* (1995), and Cordeau *et al.* (1997). Francis *et al.* (2005) 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 nodes with different minimum requirements are located in isolation to all other nodes and the depot. If the schedule options for these requirements do not contain overlapping days, it may be beneficial to raise the visit frequency of one node such that both nodes are visited together. Francis *et al.* (2005) show that this is also true in less extreme cases in which arriving at a certain region makes it beneficial to visit neighboring nodes, hence increasing the frequency with which the neighboring nodes are visited.

When flexibility in service choice is introduced, the problem begins to resemble the Inventory Routing Problem (IRP). The IRP also determines visit frequency and route configuration simultaneously, with an additional decision of how much to deliver to nodes; see Anily and Federgruen (1990) and Chan *et al.* (1998), and the surveys in Federgruen and Simchi-Levi (1995), Anily and Bramel (1998), and Kleywegt *et al.* (2002). 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 node is simply the demand accumulated since the last visit (consistent with the PVRP literature), and service is modeled as a benefit term related to each node.

Francis *et al.* (2005) highlight the difficulties in formulating and solving the PVRP-SC that result from the introduction of service choice. Several assumptions are made regarding schedule options and visit conditions to accommodate service choice in their formulation. Exact solution methods can yield optimal solutions to the PVRP-SC for moderate-sized problems with these assumptions.

In this paper, we develop a Tabu Search heuristic to solve more general cases of 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 a changes to some or all operating conditions. The following levers of operational flexibility are examined:

   (a) *Service choice.* The ability to determine customer visit frequency subject to stated visit minima. 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 stipulated by the customer. If there is no service choice flexibility, customers are served at their visit requirements. This is mode of flexibility is modeled by Francis *et al.* (2005).

   (b) *Schedule options.* The set of different schedules available that can be chosen by the service provider to serve a customer.

   (c) *Visit condition.* The ability to have any driver visit a customer during the period. If the visit condition is enforced, each customer is visited by one driver throughout the period.

   (d) *Delivery strategy.* The ability to choose the amount delivered during each customer visit, rather than being restricted to deliver the accumulated demand at each visit.

2. *Operational complexity:* The difficulty of solution implementation, from the perspective of the distribution 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 complexity:

   (a) *Driver coverage:* The portion of the total service region visited by a driver over the period. Zhong *et al.* (2004) 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.

(b) *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. (2006) 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.

(c) *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.

Using the Tabu Search heuristic, we consider the four dimensions of operational flexibility in periodic routing problems and explore the impact on the three measures of operational complexity. In the next section, we show how these measures of flexibility and complexity are modeled in the context of the PVRP-SC.

2 Approaches to the PVRP-SC

Section 2.1 presents the exact solution method for the PVRP-SC from Francis *et al.* (2005) and Section 2.2 introduces a Tabu Search solution method capable of incorporating all levers of operational flexibility discussed in Section 1. Section 2.3 defines performance metrics to quantify the operational complexity of PVRP-SC solutions.

2.1 Exact method from Francis *et al.* (2005)

We review the formulation of the PVRP-SC from Francis *et al.* (2005). Let D denote the set of days in the period and S denote the menu of service 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: $\gamma^s = \sum_{d \in D} a_{sd}$. Each schedule has an associated benefit α^s related to a monetary benefit of the corresponding frequency. We introduce a parameter, $\beta \geq 0$, which weighs the monetary benefit relative to vehicle travel and stopping times.

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 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. This is approximated by the maximum possible 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.

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

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

The formulation for PVRP-SC by Francis *et al.* (2005) is:

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

subject to

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

$$\sum_{s \in S} \sum_{k \in K} y_{ik}^s \leq 1 \qquad \forall i \in N \tag{1c}$$

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

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

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

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

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

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

The objective function (1a) balances arc travel times, stopping times and demand-weighted service benefit. Constraints (1b) enforce the visit requirements 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 for the 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.

The exact solution method in Francis *et al.* (2005) 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 of the exact method, PVRP-SC instances with up to 50 nodes are solved to within $\delta = 2\%$ of optimality.

Results from Francis *et al.* (2005) 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 various levers of operational flexibility further impact the magnitude of savings and the complexity of the resulting solution, 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 modeling and solution methods:

1. *Service choice:* Service choice flexibility can be restricted by modeling the problem with constraint (1b) fixed at equality.

2. *Visit condition:* The allocation variables, y^s_{ik}, are defined such that nodes are always visited by the same driver. In the motivating example in Francis *et al.* (2005), nodes must be visited by the same driver due to access restrictions. Allocating nodes only to schedules, y^s_i, may reduce routing costs. However, since the capacity constraints of the PVRP-SC depend on the specific vehicle and service level at each node, allowing nodes to be visited by multiple vehicles 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.* (2005) 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 $|S|$ that is the union of all disjoint schedules, there are at most $|S|-1$ different routes for each vehicle. The number of routing variables, x^d_{ijk}, is reduced significantly with this set of schedules since it is not necessary to model each day $d \in D$, only the unique delivery days which are repeated each day of the schedule.

4. *Delivery strategies:* It is assumed in most periodic distribution problems that the amount, w^s_i, delivered to a node is the amount accumulated since the previous visit. Relaxing this assumption may lead to improvements in the routing and inventory costs; however, adding an additional 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 to solve such models is significant. In the next section, we describe a Tabu Search algorithm for the PVRP-SC that can incorporate all levers of flexibility. The Tabu Search algorithm is applied to a series of test cases in Section 3 to explore the tradeoffs between operational flexibility and operational complexity.

2.2 Tabu Search

Cordeau *et al.* (1997) implement a Tabu Search algorithm for the PVRP and obtain solutions equal to or better than the best solutions for PVRP test cases in the literature. We develop a Tabu Search method based on that of Cordeau *et al.* (1997), with suitable extensions to model the PVRP-SC and incorporate operational flexibility. In what follows, we describe these changes to the Tabu Search method.

Tabu Search is a local search improvement method in which neighbors of the current solution are explored at each iteration (see Glover and Laguna (1997)). For the PVRP-SC, the solution obtained at each iteration is a complete specification of the allocation variables (either y^s_{ik} or y^s_i) and a set of routes for each vehicle on each day (the x^d_{ijk} 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.* (1994) which evaluates various tour configurations by attempting a limited number of insertions and reinsertions. The solutions are allowed to be infeasible with respect to capacity but not with respect to the visit requirement. Capacity infeasibilities are penalized in the objective function using a penalty term as in the TABUROUTE procedure of Gendreau *et al.* (1994). Briefly, the algorithm proceeds 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. If visit conditions are enforced (i.e. y_{ik}^s), then each node is always allocated to the vehicle chosen for the first day of the schedule.
   (c) Create an empty tabu list 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 selected, consider all possible 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 some number of iterations.

Suitable values for the number of candidate nodes to be chosen in Step 2 and the value of p are discussed in Gendreau *et al.* (1994). 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 to examine, the algorithm randomly chooses from all frequency-feasible choices to ensure diversity of moves.

When the visit condition in Step 2 is relaxed (i.e. y_i^s), we pick the least-cost vehicle assignments for individual days of each candidate schedule. To enforce the visit condition, we explore all possible vehicle-schedule combinations for nodes in the set, constructing candidate solutions such that the same vehicle visits the node on each day.

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. The Tabu Search method is used to solve the PVRP by not allowing service choice in Step 2. We implement the Tabu Search heuristic in C++ and execute on a Sun Fire 150 workstation with two UltraSPARC IIi processors.

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[1]. The first set of metrics, related to routing efficiency and service benefit, is explicitly considered in the objective function of both the exact solution method and Tabu Search method and the solution methods attempt to optimize solutions for these metrics.

The second set of metrics, related to operational complexity, may be implicitly considered in both solution methods. Operational complexity may be constrained through the visit condition and through limited schedule choice (only disjoint schedules and their union), which limits the driver coverage and arrival span. We relax these constraints with the Tabu Search in an effort to increase solution efficiency. We quantify the increased complexity with metrics for driver coverage, crewsize and arrival span.

Measuring the variability of 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 a node $i \in N$ is allocated to a schedule $s \in S$ and visited by vehicle $k \in K$ on day $d \in D_s$, the time at which it is visited is:

$$T_i^d = \sum_{(m,j)\in A(i)} (c_{mj} + \tau_m^s)\hat{x}_{mjk}^d \tag{2a}$$

[1] All metrics described in this section apply to both PVRP and PVRP-SC solutions.

where $A(i)$ is the set of arcs traversed before node i on the route. 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 each node can be calculated as:

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

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

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

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

Recall that solutions with higher arrival span are considered less desirable from the perspective of the customer.

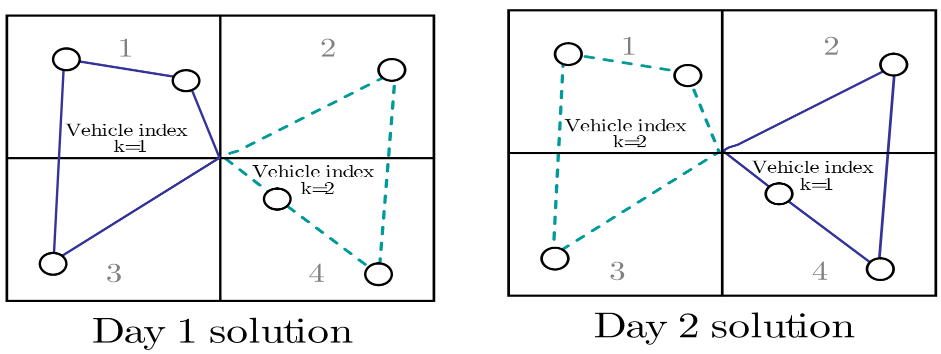

Figure 1: Same solutions may be assigned different vehicle indices

On the other hand, measuring driver coverage and crewsize requires additional processing of PVRP-SC solutions when the allocation variables are y_i^s rather than y_{ik}^s. The vehicle index $k \in K$ assigned to a route is arbitrary for any PVRP-SC solution. Figure 1 shows an example of 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. When the visit condition is enforced, the same vehicle index is given to the left region (and to the right region) on both days; however, when the visit condition is relaxed, 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 solution. In this case, 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 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 vehicle dispatcher may allocate service areas to drivers familiar with certain neighborhoods and/or customers.

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 $l = 1, 2, 3, 4$. Let N_l denote the set of nodes contained in cell $l \in L$. A driver visits a cell if he visits at least one node within that cell. The assignment problem minimizes the number of cells that each driver covers. Let V be the set of drivers to be assigned to vehicles. Given a PVRP-SC solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$, we define a parameter b_{kld} as:

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

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:

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

subject to

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

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

$$\sum_{v \in V} W_{vkd} \leq 1 \qquad \forall k \in K, d \in D \tag{3d}$$

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

$$U_{vl} \geq 0 \qquad \forall v \in V, l \in L \tag{3f}$$

The objective (3a) minimizes the total 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 $d \in D$. Constraints (3d) ensure that only one driver is assigned to each vehicle index on each day. Constraints (3e) and (3f) define the decision variables (note that U_{vl} is binary, given binary values for W_{vkd}). Once we obtain a solution $(\hat{\mathbf{U}}, \hat{\mathbf{W}})$ to the assignment problem, we can calculate driver-dependent metrics.

Driver coverage measures 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, calculated as:

$$\theta_v = \frac{\sum_{l \in L} \hat{U}_{vl}}{|L|} \tag{4a}$$

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

$$\theta = \frac{\sum_{v \in V} [\sum_{l \in L} \hat{U}_{vl}/|L|]}{|V|} = \frac{Z_a}{|L||V|} \tag{4b}$$

where Z_a is the solution to Formulation (3). Clearly, the number of cells and vehicles affects the possible values of θ. We would expect $\theta \approx \frac{1}{|V|}$ in solutions that equally partition neigborhoods 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.

The second metric measures average crewsize. Using the solution $\hat{\mathbf{W}}$ from Formulation (3), and the PVRP-SC solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$, we determine which driver visits a node on any given day. Let indicator $e_{iv} = 1$ if node $i \in N$ is visited by driver $v \in V$ during the period and 0 otherwise. Then, for each node $i \in N$ and driver $v \in V$, we have:

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

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

$$\phi_i = \sum_{v\in V} e_{iv} \tag{5b}$$

The average crewsize in the PVRP-SC solution is:

$$\phi = \sum_{i\in N} \phi_i/|N| \tag{5c}$$

Accordingly, ϕ ranges from 1 to $|V|$. A high value of ϕ indicates that many different drivers visit nodes, which may be undesirable in applications which require drivers to have knowledge/training specific to customer locations (configuration of the customer's 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 (equation (1a)); average arrival span σ (equation (2d)); average driver coverage θ (equation (4b)); and average crewsize ϕ (equation (5c)). Additionally, we consider the computational complexity of the solution methods by comparing solution times.

3 Numerical analysis

We evaluate tradeoffs between operational flexibility and complexity in this section. Section 3.1 examines the quality of the Tabu Search method relative to optimal solutions from Francis *et al.* (2005). Section 3.2 introduces the test cases used for the numerical studies. Section 3.3 introduces the measures used in the numerical analysis. Section 3.4 analyzes operational flexibility from service choice, visit condition, and schedule options. Section 3.4.4 analyzes operational flexibility from delivery strategies.

3.1 Evaluation of Tabu Search

In order to measure the quality of the Tabu Search method, we compare the exact solutions to problem instances from Francis *et al.* (2005) to solutions obtained using the Tabu Search method, in terms of the objective and operational complexity. Further, we compare the solution times for the two methods. The test cases range in size from 12 node to 40 node, with 3 or 4 vehicles, and various capacity levels. The schedule set is restricted to three options: {Mon-Wed-Fri, Tue-Thr, daily}. We solve the test cases used in Francis *et al.* (2005) with our Tabu Search method with the same assumptions as the exact method used in their paper, by imposing the visit condition and the above schedule set.

Test case	Change in Z	Change in σ	Change in θ	Change in solution time
12-node	0.0%	0.3%	0.1%	81.2%
16-node	0.1%	0.2%	0.4%	-5.2%
20-node	0.0%	0.3%	0.4%	-39.0%
28-node	0.1%	0.4%	0.3%	-59.3%
36-node	0.3%	0.7%	0.5%	-71.9%
40-node	0.4%	0.9%	0.7%	-84.6%

Table 1: Comparison of Tabu Search solutions with exact solutions (all figures are % difference)

Table 1 compares the average difference (aggregated over vehicles and capacity levels) between the two solution methods, examining the percentage difference between the objective, the performance measures σ and θ, and the solution time (Note that the crewsize measure $\phi = 1$ as the visit condition is imposed). The avreage objective values of solutions obtained with the Tabu Search are within 0.4% of the optimal solution. We conclude that the Tabu method produces solutions that are quite close to the optimal solutions in their objective values. The operational complexity of solutions obtained by the two methods also differ by less than 1% across all above measures. Thus, the solutions of the Tabu Search method provide a good representation of the operational complexity of the optional solutions. Finally, 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 problems in roughly 40 CPU minutes. Additionally, the Tabu Search obtains solutions for far larger problems, that cannot be solved with the exact method, within a reasonable amount of time.

While the heuristic version of the exact method could be used with the precision (δ) set to large values, we expect that the Tabu Search method would outperform such solutions, both in terms of solution quality as well as solution times. Given the ability of the Tabu Search method to incorporate flexibility levers, as well as its speed and quality, we use it to produce solutions for our analysis.

3.2 Test cases

Francis *et al.* (2005) find that the routing efficiency gains from service choice in the PVRP-SC are impacted by the geographic distribution of nodes in the test cases. We show that, in addition to geographic distribution, the levers of operational flexibility impact the magnitude of savings from service choice.

We create a set of randomized instances for our computational study in abstracted patterns of various city configurations. Figure 2 displays the four configurations of the service region for test instances with characteristics of demand patterns and node distribution 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. The service region is divided by customer visit requirements (either daily, 3-day, or 2-day) which are determined by demand levels. In all configurations, the depot is located in the center of the region.

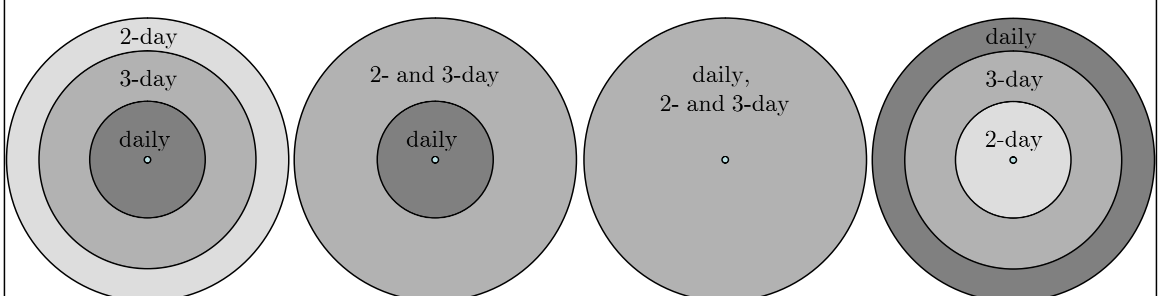

a. Traditional city model, TC

b. Traditional city with sprawl model, TCSP

c. Sprawl model, SP

d. Vanishing city model, VC

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. Problem instances for each configuration are randomly generated with 200 nodes each. Demands are drawn from a normal distribution, and nodes are uniformly scattered within each subregion.

For each test case, we consider the schedule sets 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 possible visit requirement values (f_i) for the schedule set. The schedule set is shown in the final column. Note that set A includes only disjoint schedules and their union, which is the set of schedules used with the exact solution method in Francis *et al.* (2005). 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 respectively[2].

Schedule Set	$\lvert S\rvert$	Min frequencies	Schedule set
A	3	2,3,5 days	$(MWF), (TR), (MTWRF)$
B	5	2,3,5 days	$(MWF), (MWR), (TR), (TF), (MTWRF)$
C	5	1-5 days	$(W), (TR), (MWF), (MTRF), (MTWRF)$
D	7	1-5 days	$(W), (TR), (TF), (MWF), (MWR),$ $(MTRF), (MTWRF)$
E	10	1-5 days	$(W), (R), (F), (MR), (TR), (TF),$ $(MWF), (MWR), (MTRF), (MTWRF)$

Table 2: Schedule sets for five-day test cases

We can now study the effect of operational flexibility levers on these test cases. We study the node visitation flexibility (service choice, visit conditions, and service options) and delivery strategy flexibility in Section 3.4.

3.3 Measuring efficiency and complexity

Numerical studies from Francis *et al.* (2005) suggest that test cases resembling the Traditional City configuration observe significant improvements in routing efficiency and customer service from

[2]Benefits specific to each schedule option can be incoporated in the Tabu Search procedure easily.

service choice flexibility, while cases resembling the Sprawl configuration experience lower routing efficiency (while still increasing customer service). The city configurations facilitate the study of the interplay between geographic node distribution and operational flexibility in periodic vehicle routing problems. This section examines this interplay through computational tests on the 200-node city configurations with a broader representation of operational flexibility, testing 10 randomly generated instances for each city configuration.

The absolute values of the metrics cannot be aggregated across instances as each test instance considers a different set of nodes. Instead, we consider the change in the metrics induced by introducing flexibility. We examine the impact of the flexibility levers as follows: The subscript *cons* denotes the value corresponding to the constrained solution without the flexibility lever, and the subscript *flex* denotes the solution value when the flexiblity lever is used.

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

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

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

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

We separate the objective improvements by routing cost and the service benefit: the contribution due to the routing cost component, $\Delta_{Z(c,\tau)}$, and the contribution due to the service benefit component $\Delta_{Z(\alpha)}$. Note that the driver coverage metric θ is bounded by $\frac{1}{|K|} \leq \theta \leq 1$. Hence, the corresponding rise in driver coverage is bounded by $0 \leq \Delta_\theta \leq (|K|-1)$. Also, the value of crewsize ϕ is bounded by $1 \leq \phi \leq |K|$, and the corresponding rise by $0 \leq \Delta_\phi \leq (|K|-1)$. Note that the denominators of complexity measures are small compared to those of the objective measure(Z).

Note that although the above measures are being expressed as percentages, they are of very different orders of magnitude. Therefore, we use the following example to illustrate how the magnitude of the measures may be interpreted. Figure 3 shows one instance of a Sprawl configuration. For this instance, the PVRP-SC solution with schedule set A and the visit condition enforced yields an objective value of $Z =$\$4,254 (routing cost: \$5,938; service benefit: \$1,684), $\sigma = 2.01$ hours, $\theta = 16.8$ cells, and $\phi = 1$ driver.

When we introduce visit condition flexibility, the objective value decreases by \$28, or $\Delta_Z = -0.6\%$. Comparing the solutions, we see no change in the schedule allocations, and hence the

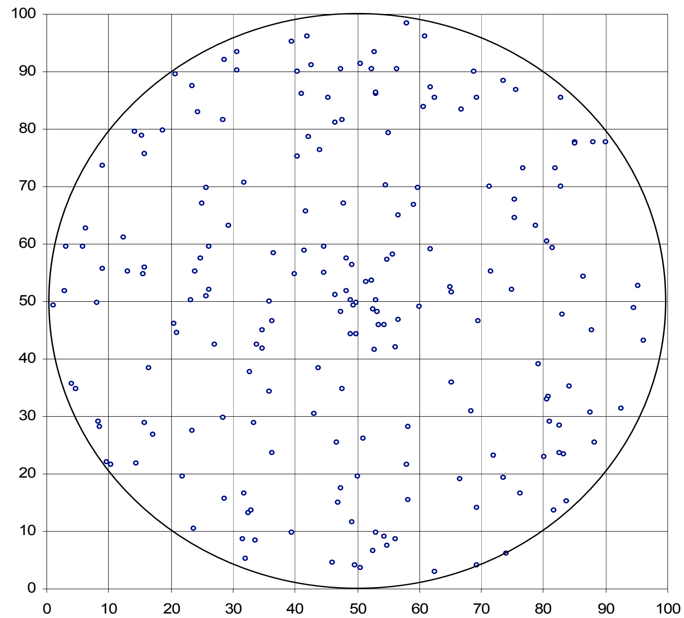

Figure 3: An instance of the Sprawl (SP) City Configuration set of test cases

improvement in the objective value is entirely due to change in the routing. Note that differences in the generated routes can be partly attributed to the use of a heuristic in route construction. Visit time σ is almost the same at 2.05 hours, or $\Delta_\sigma = 2\%$. Similarly, there is a slight rise in θ to 17.6 cells, or $\Delta_\theta = 4.7\%$. In the flexible solution, 2 customers are visited by 3 drivers, 12 customers are visited by 2 drivers and the remaining 186 customers by one driver. Thus, $\phi = 1.08$ drivers per customer, or $\Delta_\phi = 8\%$. A 1% change in Δ_Z is change of about \$42 in the objective, while a 1% change in Δ_σ is an average variation of about 1 minute in visit time. (*Delete the last sentence?*)

Now introduce another lever of flexibility, expanding the set of schedule options from set A to set E. Comparing this solution with the first, we obtain $\Delta_Z = 5.87\%$, with 4.17% savings due to routing, and 1.70% savings due to service benefit increases (19 nodes more are served at different frequencies). Arrival span σ rises to 2.24 hours, yielding $\Delta_\sigma = 11.4\%$. Similarly, driver coverage θ rises another 0.6 cells to 18.2, yielding $\Delta_\theta = 8.3\%$. In the flexible solution, 4 nodes are visited by 3 drivers, 20 nodes by 2 drivers, and the remaining by 1 driver, yielding $\phi = 1.14$ drivers per customer and $\Delta_\phi = 14\%$. In this case, we can see when additional schedule options are introduced, it becomes desirable to serve more nodes at higher frequencies resulting in more savings from increased benefit and improved routing efficiency. As more nodes are served at higher levels, the resulting routes become complex (higher σ and ϕ) and drivers cover larger areas (higher

θ).

3.4 Flexibility impacts on performance metrics

In what follows, we examine the effect of the flexibility levers across the city configuration test cases.

3.4.1 Visit condition

Table 3 shows the average percentage change in the metrics, for both PVRP and PVRP-SC solutions. In the constrained solution ($cons$), there are no flexible visits (visit condition enforced). In the flexible solution ($flex$), we allow flexible visits (visit condition relaxed). In both cases, we use schedule set E which provides the most schedule options.

	PVRP					PVRP-SC, routing only					PVRP-SC, routing & service benefit				
	$\Delta_{Z(c,\tau)}$	$\Delta_{Z(\alpha)}$	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	-0.4	0.0	0.4	2.8	2.2	-1.8	0.0	1.5	7.8	6.1	-1.9	0.0	1.6	7.8	6.1
TCSP	-0.4	0.0	0.4	2.6	2.3	-1.3	-0.1	0.8	6.6	5.4	-1.3	-0.1	0.9	6.6	5.4
SP	-0.3	0.0	0.3	2.2	2.1	-1.2	-0.1	1.0	6.7	5.4	-1.1	-0.2	1.0	6.7	5.4
VC	-0.1	0.0	0.4	2.4	1.8	-1.7	0.0	1.4	7.0	5.9	-1.6	-0.1	1.4	7.0	5.9

Table 3: Effects of relaxing the visit condition (all values are in %)

The average objective improvement for each of the configurations in Table 3 is consistently less than 2%. Further, the relaxation of the visit condition appears to improve the PVRP-SC objectives more significantly than PVRP objectives. PVRP-SC solutions have a larger feasible region (more schedule choices are possible); therefore, this additional flexibility expands the feasible region more when used in conjuction with service choice flexibility. The relative magnitude of 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. Service allocations appear to change significantly in the TCSP and SP configurations when service benefit is included in the objective ($\beta \geq 0$). This effect is to be expected because these two configurations have nodes of differing demand densities in close proximity with each other, allowing nodes to be served at higher service allocations when the visit condition is relaxed.

Adding flexibility by relaxing the visit condition has a noticable impact on the complexity measures θ and ϕ. This result is to be expected as these metrics are directly affected by the relaxation. Relaxing the visit condition expands the set of solutions by those solutions that are characterized by higher values of individual θ (visiting more nodes increases U_{vl} for drivers) and ϕ. 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 not significantly affected by changing vehicle routes (particularly in dense delivery areas).

System regularity can be enforced by using the visit condition, without significantly affecting the objective function. These results suggest that imposing the visit condition guides solutions to create driver delivery districts that have fewer overlapping areas, which reduces complexity for customers and drivers. This restriction does not have a large impact on the objective function.

3.4.2 Schedule options

Table 4 shows the average percentage change in the metrics, comparing solutions with schedule set A (*cons*) with solutions with schedule set E ($flex$) for both PVRP and PVRP-SC. We relax the visit condition in both cases to allow the system to choose the best vehicle assignments for all days. This allows us to examine the unrestricted change in crewsize under schedule option flexibility.

	PVRP					PVRP-SC, $\beta = 0$					PVRP-SC, $\beta \geq 0$				
	$\Delta_{Z(c,\tau)}$	$\Delta_{Z(\alpha)}$	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	0.0	3.3	1.7	2.1	-5.4	-1.6	8.3	3.1	4.9	-5.3	-2.9	9.8	4.9	6.2
TCSP	-2.5	0.0	2.9	0.7	1.6	-4.8	-1.5	7.1	1.9	3.6	-4.7	-2.7	8.5	2.1	4.8
SP	-2.1	0.0	2.3	1.0	1.8	-3.9	-0.8	6.8	2.8	4.0	-3.8	-2.4	7.2	3.0	7.9
VC	-1.5	0.0	1.0	1.1	1.9	-3.2	-0.6	3.1	2.1	4.2	-2.9	-1.4	3.3	3.3	5.6

Table 4: Effects of increasing number of schedules (all values in %)

Table 4 suggests that the geographic distribution of nodes has an effect on the change in efficiency and complexity when increasing the number and type of schedule options. 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 the TC and TCSP cases (centralized demand configurations) and the VC and SP cases (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 the VC and SP configurations. 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 high level of service since the marginal cost of including them on the route is relatively low (Francis *et al.* (2005) illustrate this principle). In fact, when the fixed portion of the stopping cost τ is negligible, then nodes receive higher service as long as the benefit increase offsets this low marginal cost, subject to capacity constraints on the routes. Hence, adding additional flexibility by extending the set of schedule options has lesser value for the dispersed configurations as opposed to the centralized configurations.

The relative composition of the objective function shows that including service benefit in the objective has the expected effect of inducing nodes to be visited more often than they would if routing efficiency was the only concern. The objective function improves from 3.8-7.0% when service benefit is not considered, to 4.3-8.2% when it is considered. The routing costs rise as there are more visits to nodes when service benefit is considered. The relatively small difference in the routing cost when using different objectives can be attributed to the relatively high density of nodes in these test cases. In sparser instances, the marginal cost of inserting nodes into routes has a greater effect. Further, some noise may be expected due to the use of a heuristic to solve the problem.

With respect to the operational complexity, recall that drivers on schedule set A perform at most two different routes which significantly reduces complexity. The arrival span σ changes significantly, which suggests that adding more flexibility in determining the number of schedules is at the expense of consistency in customer visit times. This is to be expected as the day-to-day composition of routes has more variance when more schedule choices are available. While the other measures θ and ϕ are also affected, the effect is not consistent across instances of each type. As before, there is a greater rise in complexity measures for the PVRP-SC solutions than in those for PVRP solutions as introducing service choice increases the feasible region and more visits are made. The complexity metrics that relate to number of visits σ and ϕ show corresponding increase as service benefit is included in the objective function and more nodes are being visited. The driver coverage θ also rises for two reasons. First, the number of stops in the period increases, forcing each vehicle to cover a wider area. Second, the routes generated are not as efficient as those created when routing cost is the only objective and tend to have more overlapping regions.

We also analyze the effect of marginal changes in the set of schedule options. Comparisons of solutions with schedule set D with those obtained with schedule set E yield a 1.4% change in Δ_Z on the average, with significant variance in individual observations ranging from 0% to 4%. This suggests that adding a small set of additional schedules can provide some incremental benefit on the average; however, the high variance in the savings obtained indicates that the benefit of adding schedules varies widely from instance to instance, even for instances of the same configuration type.

Finally, solution times for the Tabu Search are sensitive to the size of the schedule set $|S|$; however, the number of nodes affects its solution time quite moderately. 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 $|S| = 10$.

3.4.3 Interaction of service choice and schedule options

The preceding analysis suggests that introducing flexibility in schedule options appears to have a significant impact on the objective function, with or without service benefit in the objective. Therefore, we examine the trade-offs between flexibility and complexity when increased schedule options and service choice are considered together. As before, the visit condition is relaxed to allow the system the flexibility to pick the best vehicle assignment for each node.

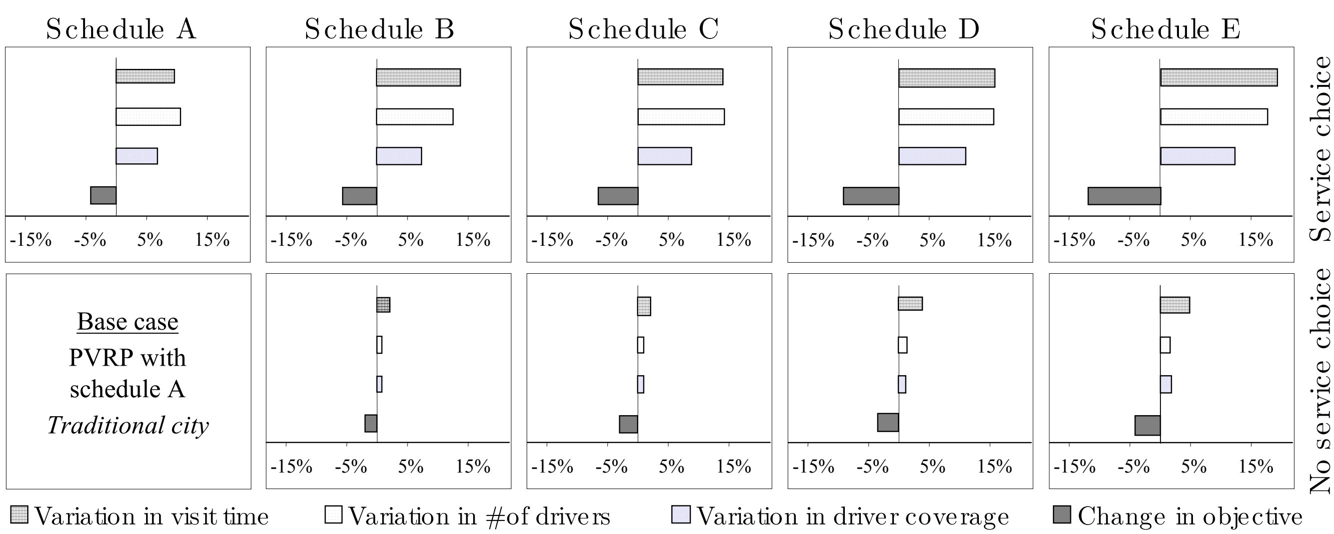

Figure 4: Increasing flexibility in node visitation: Traditional city configuration *(update figure)*

Figures 4 and 5 illustrate how the performance metrics change as both service choice and schedule option flexibility are introduced for the TC and VC city configurations, respectively. Along the horizontal direction, from left to right, we change the schedule options from A to E;

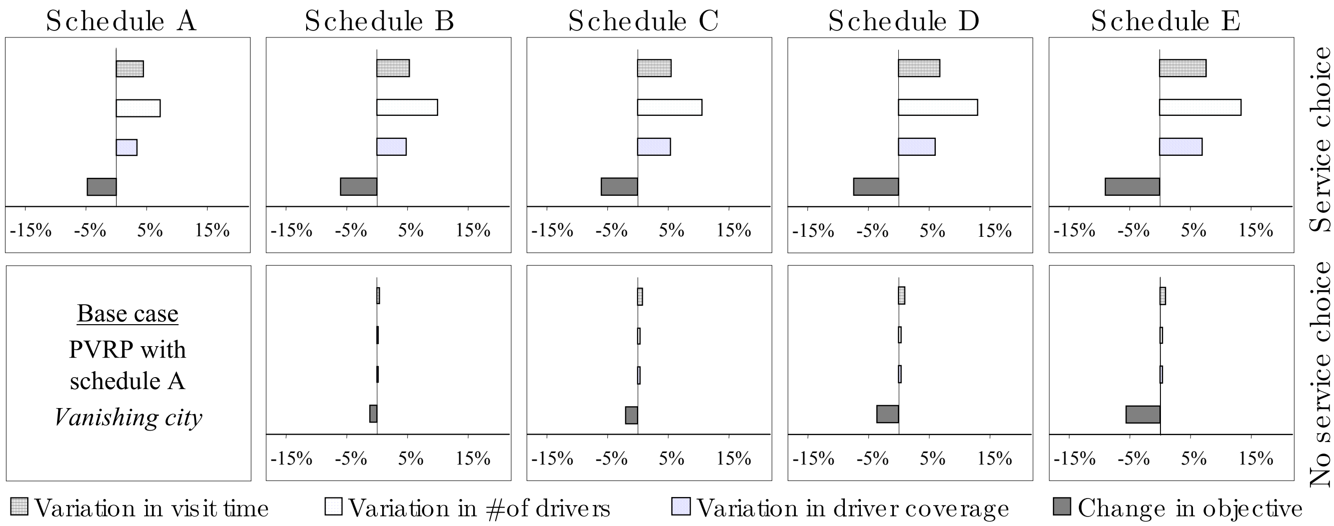

Figure 5: Increasing flexibility in node visitation: Vanishing city configuration *(update figure)*

along the vertical direction, we introduce service choice (PVRP-SC solutions) in the upper row and restrict service choice (PVRP solutions) in the lower row. All metrics are measured with respect to a 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 increased. Improvements in the objective function are accompanied by increases in operational complexity in most cases. The effects are larger for the TC configuration than any other configuration as in the previous section. Also, as before, the PVRP with VC configuration has the least increase in complexity measures. The results for the TCSP and the SP configurations (not shown here) lie between the TC and VC configurations in terms of change in both operational performance and operational complexity. Hence, the TC and VC configurations are representative of the extreme conditions of geographic configuration.

The complexity measures are also affected by the schedule option flexibility. Increasing flexibility by introducing service choice and/or new schedule options affects the number of different drivers that may need to be trained for operations specific to each customer location. There are also significant differences in the arrival span. Again, while all city configurations appear to be equally affected, the TC configuration which shows the highest efficiency gains, also has the greatest change in arrival span with increased flexibility.

On analyzing both figures, it becomes immediately apparent that service choice is more valuable in terms of efficiency, but it also correspondingly increases complexity to a greater extent. However, schedule option flexibility can produce comparable gains with lower complexity rise. For instance, in Traditional City configurations, a 5% efficiency gain can be observed when using service choice flexibility, offering schedule options of set A. Comparable efficiency gains can be obtained without using service choice flexibility and increasing the set of schedule options to set E, but with much smaller rise in complexity. The choice of modes of flexibility will depend on the relative importance of operational efficiency and complexity to the decision maker.

A relative weighting of these metrics is likely to be application specific depending on the costs associated with increased complexity. Service providers should evaluate the relative gains from increased operational efficiency against operational costs such as driver training and possible customer dissatisfaction arising from the increased complexity of interacting with the service provider. These metrics provide a way of quantifying the change. For instance, consider a distribution operation providing schedule set A schedules to customers in a city of the VC configuration. If transportation costs are very 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.

3.4.4 Delivery strategies

As discussed in Section 2.1, 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 vehicle routing problems.

In the PVRP literature, demand accumulation is modeled as the average accumulation between visits. In the PVRP-SC literature, the accumulation is modeled as the maximum demand accumulation between visits. With both these approximations, the delivery quantities, w_i^s, can be determined for any given nodes $i \in N$ and schedule $s \in S$, and the day is not needed in the formulation in the capacity constraint (1d). As discussed below, approximating the true accumulation by either the average or the maximum may be problematic. However, with the true accumulation, the accumulation includes the day, w_i^{sd}, significantly increasing the complexity of the problem to be solved by the exact method of Francis *et al.* (2005), since it requires the addition of another index on the routing variables or a non-linear constraint. For consistency of analysis, we use maximum accumulation for the Tabu Search method and analysis in Section 3.4.

Using average accumulation may lead to capacity-infeasible solutions if capacity is tight and the time between visits is not uniform in certain schedules. On the other hand, using the maximum accumulation guarantees feasibility but may lead to suboptimal solutions. Using average accumulation to approximate true demand accumulation in periodic delivery is reasonable if at least one of the following conditions is satisfied:

1. Demand levels at each node do 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 accomodate 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 the first condition holds or if customers are willing to accept more deliveries in excess of the accumulated demand.

However, other solution methods (and in particular our Tabu Search method) may consider the true demand accumulation between visits. True accumulation can incorporate non-uniform separation between visits, as well as non-stationary demand and service choice. In the Tabu Search, the delivery amount to node $i \in N$ on schedule $s \in S$ and day $d \in D$ is given by the parameter w_i^{sd}.

We use these demand modeling methods on our city configuration test cases to study the effect of using demand approximations instead of the true accumulation. When operating under the maximum accumulation modeling, we deliver only the required demand, but reserve a vehicle capacity for the maximum accumulation, which is completely used at least once during the period. We study 10 randomly generated instances for each city configuration type using the Tabu Search method. All the instances are run with schedule set E and the visit condition relaxed. On the average, the Tabu Search solves 200-node data sets in about 196 minutes, with no significant difference in solution time between the various accumulation options. Table 5 presents the average percentage difference in objective function values when the true accumulation is approximated either by the average accumulation or the maximum accumulation.

The computational study indicates that using the average approximation tends to produce solutions with a lower objective than if the true accumulation is used. This decrease in objective value can be attributed to the fact that the capacity usage is underestimated by the model. Many solutions obtained under the average accumulation case may be capacity-infeasible. For our set of test cases about 23% of such solutions are found to be capacity-infeasible. Converting these

	PVRP				PVRP-SC			
	TC	TCSP	SP	VC	TC	TCSP	SP	VC
Average vs True	-0.3%	-0.2 %	-0.2 %	-0.1%	-0.9%	-0.6%	-0.8%	-0.6%
Maximum vs True	1.4 %	1.3 %	1.4 %	1.3 %	3.8%	3.2%	3.6%	3.1%

Table 5: Impact of accumulation assumptions on objective function Δ_Z

infeasible solutions to feasible solutions is non-trivial.

Using maximum accumulation to approximate true accumulation tends to produce solutions with slightly higher objective values; however, the feasibility of these solutions is guaranteed as the capacity usage is overestimated by the model. Under maximum accumulation modeling, vehicle capacity may be reserved, but not completely utilized on all days. Thus, while the potential exists to improve the routing efficiency of the solution, the model cannot take advantage of unutilized capacity as the maximum approximation does not recognize distinctions between days.

While the traditional city configuration cities appear to be most affected when the maximum approximation is used to model demand accumulation, there is significant variance in the results obtained for individual instances. This indicates that the extent to which the maximum approximation affects the objective is tied to route design and the capacity utilization of vehicles on routes rather than to the geographic configuration of nodes. Finally, we note that operational complexity measures show no significant correlation with the different accumulation modeling methods used. (*Showing all the measures in the table is very distracting. If we want to show it, we should reserve it for an appendix.*)

The ability to make the delivery amount w_i^s a decision variable presents a new possibility. In periodic distribution problems, the delivery amount is set to the demand accumulation. However, introducing choice in the amount delivered to a node can increase efficiency. 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 exist a limited set of visit frequencies. A vehicle always delivers an amount exactly equal to the demand accumulated between visits under these conditions. However, unlike the traditional PVRP-SC, the service benefit term is modeled as the cost of holding inventory between visits rather than a benefit of increased frequency. The service benefit term depends on the demand at the node, as well as the time that each unit of the demand is being held. In the second case, we relax the assumptions of a zero inventory policy and allow shortages. Allowing shortages guarantees feasibility in cases that were not feasible previously. Note that continuous IRP model 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 can envision two strategies (delivery options): average amount and required amount. The maximum-delivery strategy can be excluded because it requires us to reserve more capacity than required on the vehicle, resulting in less efficient routing solutions. Under the average-delivery strategy, the delivery amount at each node is set to its average. Similarly, we deliver all accumulated demand under the required-delivery strategy.

Two delivery options are compared to measure the benefit of modeling delivery strategy: (1) With no delivery flexibility ($cons$) – a PVRP-SC in which the service benefit is modeled as holding and shortage cost savings and exactly the required amount is delivered every time; (2) With 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-delivery amount or the required-delivery amount. For each node, a cost is assigned to each schedule based on the holding and shortage costs depending on the amount delivered, the demand, the delivery strategy, 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 method is modified to solve these special cases of the IRP as follows: We create copies of each schedule for each delivery strategy. In this case, there are two copies for the average-delivery and required-delivery strategies, yielding a total of $2 \times |S|$ available schedule options. When modeling only the true demand accumulation, we exclude all $|S|$ schedules that are modeled with average demand accumulation. When we allow the system to choose demand accumulation, we allow all $2|S|$ schedules. 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 our 200-node, city configuration test cases. The average solution times increased from 194 minutes for $|S| = 10$ to about 486 minutes for $|S| = 20$. Note that this approach can be used to consider a wider range of delivery options; however, the increase in solution time limits the number of options that can be practically considered.

For the given inventory costs, adding the delivery strategy flexibility is found to be beneficial.

The objective function improves by 2-13% across all instances when the system is allowed to choose between delivering average demand and true demand. 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 from instance to instance, even within instances of the same configuration type.

PVRP-SC

	Δ_Z	Δ_σ	Δ_θ	Δ_ϕ
TC	7.4%	6.8%	2.4%	4.2%
TCSP	6.5%	6.4%	2.2%	4.0%
SP	6.6%	6.4%	2.6%	4.4%
VC	5.4%	6.6%	2.4%	4.0%

Table 6: Impact of introducing delivery flexibility

The overall objective improvements vary slightly by city configuration type, consistent with the results obtained in Section 3.4. However, the wide variance in the individual observations again suggests that the choice of delivery strategy is partially dictated by other factors such as route design and capacity utilization of vehicles. In this case, there may be temporary demand shortages at a node until the next visit. These results indicate that the ability to control the amount of demand delivered is a useful means of flexibility. Finally, operational complexity is found to rise as twice the number of schedules are added, but there does not appear to be a significant relationship between the increase in complexity and utilized flexibility in delivery strategy.

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 visit condition flexibility) tend to increase complexity without correspondingly significant efficiency gains. In this manner, the complexity measures can be used as a means to evaluate the different flexibility levers, maximizing efficiency gains against allowable complexity increases.

2. *Significance of geographic distribution.* In general, the results confirm earlier results findings in Francis *et al.* (2005) on the significance of geographic distribution on the savings obtained from introducing service choice flexibility and expand the results for additional flexibility levers as well. The results indicate that introducing flexibility is more beneficial when high frequency nodes are located near the depot (as in Configuration TC and TCSP).

3. *Effect of visit condition.* Enforcing the visit condition is often required by customers (e.g., inter-library loan application of Francis *et al.* (2005)). 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, visit 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 of flexibility can be modeled and analyze their effect on the efficiency and complexity of resulting solutions.

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. As mentioned in the introduction, simplifying assumptions are often required to model complex distribution systems with operational flexibility. The results presented in this paper provide insight as to when these assumptions may significantly limit the value of operational flexibility.

We develop a series of performance metrics to measure operational complexity. These are the first known metrics in the literature to quantify the desirability of routing solutions in a periodic distribution context. The complexity measures considered in this paper are either considered endogenously as constraints in problem modeling 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 as $U_{ki} \geq \sum_{j \in N} x^d_{ijk} \forall\ k \in K, i \in N, d \in D$. 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 an 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

Anily, S. and Bramel, J. (1998). Vehicle routing and the supply chain. In R. G. S. Tayur and M. Magazine, editors, *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers.

Anily, S. and Federgruen, A. (1990). One warehouse multiple retailer systems with vehicle routing costs. *Management Science*, **36**(1), 92–114.

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

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

Chan, L., Federgruen, A., and Simchi-Levi, D. (1998). Probabilistic analysis and practical algorithms for inventory-routing models. *Operations Research*, **46**(1), 96–106.

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

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

Federgruen, A. and Simchi-Levi, D. (1995). Analysis of vehicle routing and inventory-routing problems. In *Handbooks in Operations Research and Management Science*, volume 8, pages 297–373. Elseiver Science B.V.

Francis, P., Smilowitz, K., and Tzur, M. (2005). The period vehicle routing problem with service choice. *Transportation Science*. Forthcoming.

Gendreau, M., Hertz, A., and Laporte, G. (1994). A tabu search heuristic for the vehicle routing problem. *Management Science*, **40**(10), 1276–1290.

Glover, F. and Laguna, M. (1997). *Tabu Search*. Kluwer Academic Publishers, Norwell, MA, USA.

Kleywegt, A., Nori, V., and Savelsbergh, M. (2002). The stochastic inventory routing problem with direct deliveries. *Transportation Science*, **36**(1), 94–118.

Russell, R. and Gribbin, D. (1991). A multiphase approach to the period routing problem. *Networks*, **21**, 747–765.

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

Tan, C. and Beasley, J. (1984). A heuristic algorithm for the period vehicle routing problem. *Omega*, **12**(5), 497–504.

UPS Corp. (2006). 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.

Zhong, H., Hall, R., and Dessouky, M. (2004). Territory planning and vehicle dispatching with driver learning. *Working paper; http://www-rcf.usc.edu/ maged/publications/TerritoryPlanning.pdf*.