Summer 2014

Strategic placement of telemetry units and locomotive fuel planning

Amit Kumar Verma
University of Iowa

STRATEGIC PLACEMENT OF TELEMETRY UNITS AND LOCOMOTIVE FUEL PLANNING

by

Amit Kumar Verma

A thesis submitted in partial fulfillment of the
requirements for the Doctor of Philosophy
degree in Business Administration (Management Sciences)
in the Graduate College of
The University of Iowa

August 2014

Thesis Supervisor: Professor Ann Campbell

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Amit Kumar Verma

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Business Administration (Management Sciences) at the August 2014 graduation.

Thesis Committee: ______________________________
Ann Campbell, Thesis Supervisor

____
Barrett Thomas

____
Philip Jones

____
Pavlo Krokhmal

____
Samuel Burer

ACKNOWLEDGEMENTS

I want to thank my advisor, friends and family for being part of this journey and always keeping faith in me. My advisor deserves a special thanks for guiding me through this entire process. Also I would like to pay gratitude to my committee members for their valuable feedback.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

LIST OF ALGORITHMS

Algorithm

CHAPTER 1
STRATEGIC PLACEMENT OF TELEMETRY TO REDUCE ROUTING COSTS

1.1 Introduction

Our motivation for the study of the strategic placement of telemetry units comes from a project with NuCO_2. NuCO_2 provides carbon dioxide and nitrogen gas to more than 130,000 customers across the United States, primarily for use in beverage carbonation [1]. These customers include national chain and local restaurants, as well as convenience stores, sports venues, and theme parks. The scheduling of deliveries to these customers is not based on customer orders, but is based currently on usage forecasts made by NuCO_2. Product usage rates are estimated based on several deliveries to a customer, and these rates are used to compute a delivery frequency for each customer. NuCO_2 schedules deliveries according to these frequencies and fills up the customer's tank when the delivery occurs. NuCO_2 has observed a few problems with this practice. If customers are using product at an unusually low rate, NuCO_2 does not know this and will deliver to a customer more often than necessary. This creates expensive routing costs that could potentially be avoided. Similarly, if customers are using product at an unusually high rate, NuCO_2 runs the risk of allowing these customers to stock out of product if they receive a delivery according to their defined frequency.

Due to the location of customer tanks, the types of gauges involved, and customer preferences, it is hard to get accurate tank readings between deliveries.

Thus, $NuCO_2$ wanted to explore the use of telemetry to remotely gather accurate tank levels. $NuCO_2$, though, only wants to purchase a very limited number of telemetry units and approached us with the question of where to best place these units. To the best of our knowledge, the strategic placement of telemetry units has not been previously considered in the literature. Our goal is to identify where to place the telemetry units to maximize the savings in routing costs while preserving a given customer service level.

The use of telemetry can be beneficial to customers, distributors and suppliers of many types of products besides carbon dioxide and nitrogen. For example, suppliers of liquified petroleum gas also use telemetry to monitor levels at customer tanks [6]. Customer tanks are often located in inaccessible areas, and telemetry provides a means of monitoring the inventory levels from a control facility [3]. Vending machines also transmit inventory information to supply centers via telemetry units [8]. This helps delivery people know what products are needed and in what volume.

In all of these applications, the use of telemetry represents a significant investment. Unit costs for the telemetry units can be high [5] especially when companies consider installing them at all of their customers. Besides the cost of the units, there are costs for installation and setup, including new IT infrastructure [32], and costs to maintain and operate the units. Hence for many small and medium businesses installing telemetry at all customer locations may not be viable and thus is not pursued. If a business has the budget to install a limited number of telemetry units, though, they may still be able to observe a significant amount of gain from telemetry, but it requires careful consideration of where to place these units. This is the issue addressed in this paper.

In this paper, we present the related literature in Section 1.2. In Section 1.3, we formally define the problem we solve, discuss how the given service level can be maintained for both non-telemetry and telemetry customers, describe how we model the stochastic nature of our problem via realizations, and present a full mathematical model of the problem. The details of our solution methodology are presented in Section 1.4. This involves a local search among different assignments of customers to telemetry as well as heuristics for solving the resulting routing problems across each of the realizations. We examine the performance of our approach as well as the impact of different problem characteristics in Section 1.5. We provide our final conclusions in Section 1.6.

1.2 Literature Review

Telemetry is a technology that allows remote measurement and reporting of information by transferring data over wired or wireless networks [9]. We will first discuss routing literature that specifically addresses the adoption of technology. We will then describe related research problems and research that is related to our solution approach.

Most of the literature that involves technology adoption in conjunction with vehicle routing deals more with investments in Geographical Information Systems (GIS). For example, Mbiydzenyuy [33] evaluates different information services by estimating the reduction in fuel costs, infrastructure costs, administrative costs, and costs of missing and delayed goods. The combination of GIS with existing dynamic routing algorithms creates the capability of making dynamic interventions such as when customers experience shortages or customers require immediate service. Thus, the use of real time location data is often combined with dynamic vehicle routing (see Psaraftis [41] for a survey).

Other applications of related technology in the vehicle routing domain include the use of a mobile reader. Shuttleworth et al. [43] design optimal routes for utility companies who use mobile readers to read utility meters from a distance. Pacciarelli et al. [40] assume that the item reading capability is not error free and use estimates for the tag scanning error and the associated manpower cost for the manual correction of reading error.

Our problem has connections with the location routing problem where the goal is to choose one or many depots from a set of candidate sites and to construct delivery routes from the selected depots to the customer sites such that the total cost is minimized. Reviews of work in this area include Min et al. [34] and Nagy and Salhi [35]. A variant of the location routing problem that includes periodic deliveries to customers has also been examined in Prodhon [11].

Another closely related problem to ours is the Inventory Routing Problem (IRP) (for a review, see Campbell and Savelsbergh [15] and Coelho et al. [20]). The IRP is concerned with the distribution of a product from a facility to a set of customers over a given planning horizon [15]. The objective is to minimize the routing costs during the planning horizon without causing stockouts at any of the customers. The IRP often occurs when a vendor managed inventory replenishment policy is adopted. In such a policy, a supplier manages the inventory of his customers and decides how and when inventory should be replenished at each customer [16]. Such vendors need access to accurate and timely information about the inventory status of customers and often do so by adopting telemetry units. Much of the existing work on the advantages of telemetry adoption in an IRP setting assumes that telemetry is installed at all customer locations [15]. Our problem is more strategic in that we are trying to establish where the limited number of telemetry units should be placed.

The stochastic nature of customer usage rates has also been captured in the IRP literature. For example, Hemmelmayr et al. [27] plan routes for the supply of blood products to hospitals incorporating stochasticity associated with blood product usage at hospitals. The authors use an integer programming model and suggest several recourse models involving the use of realizations. We modeled the stochastic usage using realizations in a similar way.

Our solution technique involves transforming a subproblem into an instance of the Periodic Vehicle Routing Problem (PVRP). Christofides and Beasley [19] present several of the first heuristics for PVRP. The objective is to minimize distribution costs while satisfying periodic delivery constraints. For a review of how this problem has been modeled and solved, see Campbell and Hardin [17].

Our solution approach resembles the multiple scenario approach (MSA) of Bent and Van Hentenryck [12]. We generate routes for non-telemetry customers using their computed delivery frequencies and telemetry customers using their expected frequencies, but then remove the telemetry customers from these routes. This leaves room to accommodate telemetry customers when deliveries are actually required, which is very similar to the ideas employed in MSA. In MSA, routes are created that accommodate potential future requests, but then the future requests are removed from the plan. Bent and Van Hentenryck [12] show that MSA is a preferred method to incorporate stochastic knowledge into route planning.

1.3 Model

1.3.1 Problem Definition

The problem we want to solve is where to install M telemetry units to minimize the average routing costs across a set of R customer usage realizations over a given time horizon T while maintaining a given customer service level for all n customers. We assume there is a fleet of V capacitated vehicles available each day, and each customer can be visited by at most one vehicle on a day (no split deliveries). The set of arcs A includes directed edges between all customers, as well as edges between each customer and the depot (indexed by 0 and $n+1$). For each edge $(i, j) \in A$, we assign a distance t_{ij} and use a speed ρ to convert distance into travel time. The duration of a vehicle's tour has a given time limit D, and stops at customers require both a fixed stop time and variable stop time depending on the quantity delivered. We also assume that the daily usage rates of product for each customer follow a normal distribution, and the rates are independent of each other. We use μ_i to represent the mean daily usage rate for customer i ($i = 1$ to n) and σ_i for the standard deviation in daily usage rate for customer i ($i = 1$ to n). We also assume that M represents a fairly small proportion of the customer set. We refer to this problem hereafter as the Telemetry Location Problem (TLP).

1.3.2 Service Level

Next, we look at how to maintain a given service level at all customers and understand how this differs for customers with and without telemetry. We show that the savings from using telemetry is tied to the computation of safety stock.

In our equations, we use z to represent the z-statistic associated with the given service level. For a customer i without telemetry, the delivery frequency f_i (in work days) should be such that average usage rate over f_i days plus the safety stock will be equal to the tank capacity I_i^{max}:

$$\mu_i f_i + z\sqrt{f_i}\sigma_i = I_i^{max}. \tag{1.1}$$

The delivery frequency f_i derived from Equation 1.1 would be such that, on average, a truck would arrive at a customer just when the customer reaches its safety stock. The safety stock involves $\sqrt{f_i}$ since the customer is vulnerable for runout for the f_i days between deliveries, and the standard deviation of demand over f_i days is $\sqrt{f_i}\sigma_i$. The safety stock also involves z, so that we can be confident at a service level corresponding to z that the customer will not experience an outage during the f_i days between delivery.

The challenge with using a delivery frequency of f_i in practice is that deliveries may not always occur at the same time. For example, consider a customer with $f_i = 5$. If delivery on day 1 is scheduled at 9 a.m, then the next delivery on day 6 must be scheduled by 9 a.m. to guarantee the service level is maintained. To allow for flexibility in delivery time to promote good routing options, we suggest a delivery frequency $f_i^{'} = f_i - 1$ for customers without telemetry. With this change, equation (1.1) becomes:

$$\mu_i(f_i^{'} + 1) + z\sqrt{f_i^{'} + 1}\sigma = I_i^{max}. \tag{1.2}$$

We can solve for $f_i^{'}$:

$$f_i^{'} = \frac{z^2\sigma_i^2 + 2\mu_i(I_i^{max} - \mu_i) - z\sigma_i\sqrt{(z^2\sigma_i^2 + 4\mu_i I_i^{max})}}{2\mu_i^2}. \tag{1.3}$$

The value for $f_i^{'}$ would be rounded down to the nearest day or the nearest week depending on the level of customer usage.

With telemetry, the above changes. A customer is not vulnerable in the same way since the vendor would now know the status of each customer's inventory level every night via a telemetry reading. Every night, the vendor can see if a customer's inventory is low and can decide to go there the next day (assuming capacity is available). Thus, a customer is vulnerable for a much shorter amount of time. This allows safety stock to be greatly reduced without increasing the likelihood of stockout. A delivery must occur the next day if:

$$\mu_i + z\sigma_i \geq \text{most recent tank reading}. \tag{1.4}$$

Equation 1.4 says that if a day's expected usage plus a safety stock buffer is more than what is remaining in the tank, a delivery must occur. If the tank level is higher than this threshold, this formula implies that the delivery can be postponed while preserving the service level. Like with non-telemetry customers, this may create challenges in delivery timing. For example, for a particular customer, this may indicate a delivery is not needed on day 3, but will require a delivery early on day 4 to ensure the service level is maintained. To allow delivery flexibility to promote good routing, we also slightly modify Equation 1.4 to require a delivery the next day when:

$$2\mu_i + z\sqrt{2}\sigma_i \geq \text{most recent tank reading.} \tag{1.5}$$

Thus, we are essentially forcing the delivery to occur one day earlier than necessary, so that we have the whole day to make the delivery. We want to note that this is not the only time a delivery can be made to a telemetry customer. We can make a delivery even earlier if it leads to better routes, and we discuss this issue in detail in Section 1.4.2. Here, we are just trying to understand exactly how telemetry can lead to potential savings.

A customer with telemetry will not receive a delivery with a strict frequency, but the use of telemetry and the reduction in safety stock translates to an increase in the average time between deliveries. Based on Equation 1.5, the new average frequency $\hat{f}_i$ is computed as follows:

$$\frac{I_i^{max} - 2\mu_i - z\sqrt{2}\sigma_i}{\mu_i}. \tag{1.6}$$

The value $\hat{f}_i$ value helps us understand what average cost savings may result over a time horizon as a result of using telemetry. For most customers, $\hat{f}_i$ is much larger than f_i' due to the reduction in safety stock. Over a time horizon T, a larger frequency translates to fewer visits, and thus we get savings in the routing costs. An expected change in the number of visits over a horizon T as a result of telemetry would be computed as follows:

$$(\frac{T}{f_i'} - \frac{T}{\hat{f}_i}). \tag{1.7}$$

By manipulating the equations for f_i' and $\hat{f}_i$, we can derive that, all other parameters equal, the customers with the larger μ_i, larger σ_i, or smaller I_i^{max} values would yield the larger expected reduction in visits over the time horizon due to the adoption of telemetry. These observations guide the experiments in Section 1.5.3.

Our formulas for f_i' and $\hat{f}_i$ can help us precisely compute the M customers where the visit reduction over the time horizon should be the largest, but this does not necessarily determine which set of customers creates the biggest savings in routing costs. The other big factor is where a customer is located. For example, a customer that is expensive to serve can create a large cost reduction with only a small reduction in the number of visits. Looking just at the distance from the depot is not sufficient to derive the cost reduction associated with assigning a customer telemetry because most routes involve multiple customers. A customer may be routed with different customers on different days, making their cost contribution both dependent on the location of other customers but also dependent on the frequency of visits to the other customers. Thus, we must solve the problem defined by the mathematical model in Section 1.3.4 which considers both the number of visits as well as how the customers are routed on those visits.

1.3.3 Realizations

In practice, for a customer with a highly variable usage rate who has telemetry, the average number of days between deliveries may be $\hat{f}_i$, but it is very unlikely that deliveries will occur precisely every $\hat{f}_i$ days. Recall that for customers with telemetry, deliveries now will occur based on tank readings. To model the stochasticity of customer usage, we create a set of R realizations of daily usage at all customers over T. The realizations reflect the type of information we would get from telemetry units on a daily basis. These simulated usage values help us create realistic delivery dates for telemetry customers for each realization which helps us create a better assessment of the routing costs.

For each realization, for each customer, we simulate a daily usage value for each day of the planning horizon based on the customer's μ_i and σ_i values. Starting from an initial inventory level, the inventory level is reduced each "night" by the generated usage rate for that day for that realization, and it is set to tank capacity when a delivery occurs. The delivery dates for telemetry customers are generated in two different ways. In the first method, the vendor decides to deliver the next day if the current inventory level is too low according to Equation 1.5. This corresponds to delaying a delivery as far as possible in the planning horizon while preserving the service level and a full day of delivery flexibility. We refer to this as the 'last day' rule, because the delivery is made on the last day possible. We also consider a 'best day' rule, described in detail in Section 1.4.2, that allows telemetry customers to be visited earlier if the routing costs are lower. These rules reflect the type of choices that companies, such as $NuCO_2$, would use once telemetry is installed at the restricted set of customers to decide when deliveries are required.

1.3.4 Mathematical Model

In the following mathematical model, telemetry customers are routed based on the last day rule, while non-telemetry customers are routed based on their f_i' values. The model identifies routes for each vehicle for each day of the planning horizon for each realization. For non-telemetry customers, their delivery days will be the same across all realizations. The integer program essentially decides the start date of delivery for each non-telemetry customer.

The model helps us make the strategic decision of where to place the M telemetry units. Because it is a strategic model rather than a model tied to a particular point in time, we do not want to restrict the solutions based on a set of initial inventory levels that would yield bad routes. Thus, customers without telemetry are allowed to start their first delivery on any of the first f_i' days because initial inventories could easily be manipulated in practice if certain levels would enable better routes. Similarly, we strategically set the initial inventory levels for telemetry customers to promote the creation of good routes. Details on the initial inventory levels are provided in Section 1.4.3.

The following are the parameters and variables needed for our model:

Parameters

M	number of telemetry units
R	number of realizations
T	number of days in planning horizon
n	number of customers; nodes 0 and $n+1$ represent the depot
N	set of all customers
V	number of vehicles
P	capacity of vehicle
t_{ij}	distance between customers i and j
ρ	average speed of delivery vehicles
D	maximum route duration
e	fixed stop time at a customer
c	service time per unit delivery quantity
μ_i	average daily usage rate at customer i
σ_i	standard deviation of usage rate at customer i
$f_i^{'}$	delivery frequency of customer i without telemetry installation
q_i	estimated delivery quantity for non-telemetry customer i, computed as $f_i'\mu_i$
u_i^{tr}	the product usage at customer i on day t in realization r with telemetry
I_i^{max}	tank capacity at customer i
I_i^0	initial inventory at customer i

Variables

z_i 1 if telemetry is installed at customer i, 0 otherwise

y_i^{tv} 1 if customer i receives a delivery on day t by vehicle v, 0 otherwise

x_{ijr}^{tv} 1 if vehicle v travels from customer i to customer j on day t in realization r, 0 otherwise

I_{ir}^{t} inventory at customer i in realization r at the end of the day t with telemetry;

d_{ir}^{tv} a continuous variable indicating the quantity delivered to customer i on day t in realization r by vehicle v

Next, we define the objective, constraints, and variable restrictions for the problem we solve.

$$\min \frac{1}{R}\sum_{r=1}^{R}\sum_{v=1}^{V}\sum_{t=1}^{T}\sum_{i=0}^{n}\sum_{j=1,j\neq i}^{n+1} t_{ij}x_{ijr}^{tv} \tag{1.8}$$

subject to

$$\sum_{i=1}^{n} z_i = M \tag{1.9}$$

$$\sum_{v=1}^{V}\sum_{t=1}^{f_i^{'}} y_i^{tv} \geq 1 - z_i \ , \ \forall i \in [1,n] \tag{1.10}$$

$$\sum_{v=1}^{V} y_i^{tv} \leq \sum_{v=1}^{V} y_i^{t+f_i^{'}\ v} + z_i \ , \ \forall i \in [1,n], \ \forall t \in [1, T - f_i^{'}] \tag{1.11}$$

$$d_{ir}^{tv} \geq q_i(\sum_{j=1,j\neq i}^{n+1} x_{ijr}^{tv} - z_i)\ ,\ \forall i \in [1,n],\ \forall t \in [1,T],\ \forall v \in [1,V],\ \forall r \in [1,R] \tag{1.12}$$

$$I_i^{max}(\sum_{v=1}^{V}\sum_{j=1,j\neq i}^{n+1} x_{ijr}^{1v} + 1 - z_i) \geq z_i(2\mu_i + \sqrt{2}z\sigma_i) - I_i^0\ ,\forall i \in [1,n], \forall r \in [1,R] \tag{1.13}$$

$$I_i^{max}(\sum_{v=1}^{V}\sum_{j=1,j\neq i}^{n+1} x_{ijr}^{tv} + 1 - z_i) \geq z_i(2\mu_i + \sqrt{2}z\sigma_i) - I_i^{t-1}\ ,\forall i \in [1,n], \forall t \in [2,T], \forall r \in [1,R] \tag{1.14}$$

$$d_{ir}^{1v} \geq (y_i^{1v} + z_i - 1)I_i^{max} - I_i^0\ ,\ \forall i \in [1,n],\ \forall v \in [1,V],\ \forall r \in [1,R] \tag{1.15}$$

$$d_{ir}^{tv} \geq (y_i^{tv} + z_i - 1)I_i^{max} - I_{ir}^t\ ,\ \forall i \in [1,n],\ \forall t \in [2,T],\ \forall v \in [1,V],\ \forall r \in [1,R] \tag{1.16}$$

$$I_{ir}^1 = I_i^0 - u_i^{1\ r} z_i + \sum_{v=1}^{V} d_{ir}^{1\ v},\ \forall i \in [1,n]\ , \forall r \in [1,R] \tag{1.17}$$

$$I_{ir}^{t+1} = I_{ir}^t - u_i^{t+1\ r} z_i + \sum_{v=1}^{V} d_{ir}^{t+1\ v},\ \forall i \in [1,n]\ , \forall t \in [2,T-1]\ , \forall r \in [1,R] \tag{1.18}$$

$$\sum_{i=0}^{n}\sum_{j=1,j\neq i}^{n+1} \frac{t_{ij}x_{ijr}^{tv}}{\rho} + \sum_{i=1}^{n} c d_{ir}^{tv} + \sum_{i=1}^{n}\sum_{j=1,j\neq i}^{n+1} e\ x_{ijr}^{tv} \leq D\ , \forall r \in [1,R], \forall t \in [1,T], \forall v \in [1,V] \tag{1.19}$$

$$\sum_{i=1}^{n} d_{ir}^{tv} \leq P\ ,\ \forall t \in [1,T]\ , \forall r \in [1,R],\ \forall v \in [1,V] \tag{1.20}$$

$$y_i^{tv} = \sum_{j=1,j\neq i}^{n+1} x_{ijr}^{tv} \ , \ \forall i \in [1,n], \ \forall t \in [1,T], \ \forall v \in [1,V] \tag{1.21}$$

$$\sum_{i=1}^{n+1} x_{0\ ir}^{tv} = 1 \ , \ \forall t \in [1,T], \ \forall v \in [1,V], \ \forall r \in [1,R] \tag{1.22}$$

$$\sum_{i=0}^{n} x_{i\ n+1\ r}^{tv} = 1 \ , \ \forall t \in [1,T], \ \forall v \in [1,V], \ \forall r \in [1,R] \tag{1.23}$$

$$\sum_{j=1,j\neq i}^{n+1} x_{ijr}^{tv} = \sum_{k=0,k\neq i}^{n} x_{jir}^{tv} \ , \ \forall i \in [1,n] \ , \ \forall t \in [1,T], \ \forall v \in [1,V], \ \forall r \in [1,R] \tag{1.24}$$

$$\sum_{i,j\in Q} x_{ijr}^{tv} \leq |Q| - 1 \ , \ \forall t \in [1,T], \ \forall v \in [1,V], \ \forall r \in [1,R], \ \forall Q \subseteq N \tag{1.25}$$

$$d_{ir}^{tv}, I_{ir}^{t} \geq 0 \ , \ \forall i \in [1,n], \ \forall t \in [1,T], \ \forall v \in [1,V], \ \forall r \in [1,R] \tag{1.26}$$

$$x_{ijr}^{tv} \in \{0,1\} \ , \ \forall i \in [0,n], \ \forall j \in [1,n+1], \ \forall t \in [1,T], \ \forall v \in [1,V], \ \forall r \in [1,R] \tag{1.27}$$

$$z_i \in \{0,1\} \ , \ \forall i \in [1,n] \tag{1.28}$$

The objective function in Equation (1.8) minimizes the average travel cost over all realizations. Constraint (1.9) captures the restriction on the total number of telemetry units. Constraints (1.10)-(1.12) focus on customers without telemetry. Constraints (1.10) and (1.11) ensure the deliveries to non-telemetry customers occur according to their computed frequencies. Non-telemetry customers are allowed to start delivery on any of the first $f_i^{'}$ days. Constraints (1.12) enforce the appropriate delivery quantity for non-telemetry customers. Constraints (1.13)-(1.18) are focused on telemetry customers. Constraints (1.13) and (1.14) require that a delivery is made to a telemetry customer the day after the tank level falls below the level specified in Equation 1.5. Constraints (1.15) and (1.16) set the delivery quantities for telemetry customers. We set the delivery volume equal to the amount possible before daily usage occurs as that should always be a feasible quantity for customers to receive. Constraints (1.17) and (1.18) ensure inventory balance from period to period for customers with telemetry. The inventory level at the end of one period equals the inventory level at the previous time period adjusted by usage and delivery. Constraints (1.13), (1.15), and (1.17) involve initial inventory levels for telemetry customers. In Section 1.4.3, we discuss how to set initial inventories for telemetry customers in a way that promotes good routing. Constraints (1.19)-(1.25) reflect traditional routing constraints involving all customers. Constraints (1.19) and (1.20) impose restrictions on the time duration of a route and vehicle tank capacity. Constraints (1.21) connect the x and y decision variables. Constraints (1.22) and (1.23) enforce the tours to start and end at the depot. Constraints (1.24) and (1.25) are the standard flow conservation and subtour elimination constraints for vehicle routing problems. Finally, Constraints (1.26) ensure that tank levels and delivery quantities are non-negative, and constraints (1.27) and (1.28) define the other decision variables as binary.

1.4 Solution Methodology

The size of the proposed integer programming model for TLP is huge even for moderate instances. The vehicle routing problem belongs to the NP hard domain [29], and the vehicle routing problem can be viewed as a subproblem here. Hence finding an exact solution to the proposed model is computationally intractable. Thus, we propose a heuristic to solve this new problem. The heuristic starts with an initial assignment of customers to telemetry, estimates the routing cost of this assignment, and iterates to a new assignment. We propose different ways to create the initial assignment of customers to telemetry, different ways to assign the delivery dates for telemetry customers, and different options for iterating the new assignments of customers to telemetry. Since this is a strategic problem, the solution methodology does not necessarily have to be fast, but finding options that yield high quality results and reduce run time is definitely desirable. Thus, we consider the impact of these choices not only on solution quality but also on run time in Section 1.5.

1.4.1 Routing Cost Algorithm

As mentioned above, our algorithm for solving the TLP involves evaluating the routing costs for different assignments of customers to telemetry. The methods we developed to initialize and create new assignments will be discussed in detail in Section 1.4.4. First, we describe how we evaluate the routing costs associated with a particular assignment. A summary of our approach is given in Algorithm 1.1.

As indicated in the discussion of the mathematical model, for non-telemetry customers, their delivery dates will be the same across all realizations. For these non-telemetry customers, their computed frequency $f_i^{'}$ can be viewed as a strict periodic frequency, as considered in periodic vehicle routing problems (PVRP). Thus, one option is to solve a PVRP involving only non-telemetry customers and try to insert the telemetry customers into this schedule in each realization. The reason why this does not work well is that there are often very little available capacity to insert the telemetry customers into such a schedule. Thus, we need a way for the PVRP to effectively "save some room" for the telemetry customers on the routes. We can do this by solving a PVRP involving all customers, using the f_i' values for the non-telemetry customers and the $\hat{f}_i$ values for the telemetry customers (the 'Initialization' phase of Algorithm 1.1). The PVRP is known to be solved well using a tabu-search algorithm, so we employ the successful tabu-search algorithm for the PVRP presented by Cordeau et al. [21]. The parameter values used here follow from the paper.

Once we have the set of routes created by tabu search, we can remove the telemetry customers from these routes and preserve the routing decisions for the non-telemetry customers. These routes become the basis of the routes for each realization. For each realization, the vehicle assignments for the non-telemetry customers are allowed to be changed, but the scheduled dates for the deliveries to the non-telemetry customers cannot be changed. For each of the telemetry customers, we obtain the precise days of deliveries across the planning horizon for each realization according to the last day or the best day delivery rule (discussed in Section 1.4.2). For each of the telemetry customers, we compute the insertion cost for the delivery sequence and sum this across all realizations. We insert the telemetry customer with the highest total cost of insertion across all realizations. This customer is the most important to serve well, as reflected by the high insertion cost. Once a telemetry customer is inserted in the schedule, we recompute the insertion costs across the planning horizon to determine the telemetry customer with the next highest insertion cost. The process is repeated until all telemetry customers are inserted in the schedule. This procedure of transforming a PVRP solution into a solution for the TLP is summarized by the 'Construction' phase of Algorithm 1.1.

Once the routes are constructed, three improvement routines are run. This is indicated by the 'Improvement' section of Algorithm 1.1. All of these routines are focused on improvements within a single day or a single vehicle's tour. We do not allow the transfer of any customers across different days of the planning horizon. These limitations are made to ensure that all non-telemetry customers are visited on the same days across all realizations and to reduce run time. These improvement ideas were carefully chosen, after a series of computational experiments, to offer maximum improvement in a limited amount of time. First, we run a one-shift improvement routine that improves the routes on each day of the planning horizon (Lin and Kernighan [30]). Customers are potentially swapped to different locations on the same tour or to other tours on the same day. All possible moves are considered for each customer being served on a given day, and the move that creates the best improvement is selected. This process is repeated until no more improving moves are found on a given day, and then the improvement process moves to the next day. This process is repeated for each realization. Second, we run the standard 2-OPT improvement heuristic [30] to improve individual routes. The best improvement move is chosen among all the candidates at each step and the process is repeated until no improvement is found. This process is repeated for each route on each day of each realization. Third, we run VRP crossover improvement heuristics, allowing exchanges of portions of two different same day routes with each other. For example, the last three positions on the first vehicle for day one of the planning horizon are swapped with the last three positions on the second vehicle. The first move which achieves savings in routing cost is performed. This heuristic is motivated by two point crossover ideas in genetic algorithms (Gwiazda [26]).

Once the improvement process is completed, we can now compute the average routing cost associated with a potential assignment of customers to telemetry. We next consider a different assignment of customers to telemetry, as discussed in Section 1.4.4, and repeat the routing cost algorithm.

1.4.2 Last Day vs. Best Day

As previously mentioned, for telemetry customers, we use a set of usage realizations to reflect the type of information we would get from a telemetry unit. An initial inventory level, these usage values, and a decision rule determine the dates when telemetry customers are visited. It is important that we approximate how telemetry would be used in practice so that we get a good estimate of the routing costs of a particular assignment. The decision rule should be both cost-effective in terms of routing costs, but also quick to evaluate since we want to evaluate many potential assignments.

For the last day delivery rule, we use Equation 1.5 as the governing equation to determine the exact delivery dates for telemetry customers for each realization. On each delivery date selected by the last day rule, we evaluate the insertion cost of the telemetry customer at all points in the current routes for that day. For example, if customer i is served between $i-1$ and $i+1$ on a particular day, the insertion cost for this delivery is

$$(t_{i-1,i} + t_{i,i+1} - t_{i-1,i+1}). \tag{1.29}$$

We choose the location that creates the lowest insertion cost. Such an approach should yield a savings in routing costs due to the reduction in the number of visits across the planning horizon. However, fixing the day of delivery for a customer according to the last day possible has a shortcoming. The total routing cost may be somewhat high because the customer may not combine well with the other customers on the selected days.

Another option is to allow more flexibility in the choice of delivery day for the telemetry customer. We refer to this as the best day rule, where the delivery date for a telemetry customer can occur earlier if it yields better routes. If the customer is within one week of the projected last day, the best day rule evaluates the cost of making the next delivery on each of the possible days within that week. In terms of equations, a delivery is considered when

$$I_{ir}^{t} - (w-1)\mu_i \leq 2\mu_i + z\sqrt{2}\sigma_i \tag{1.30}$$

where w represents the number of days per week where usage and deliveries occur. When the day is identified that creates the lowest insertion costs, the delivery is tentatively scheduled on that date. In doing so, the method yields savings in routing costs in the short term. For example, if customer x must receive a delivery by day 5 of the planning horizon but combines well with non-telemetry customers y and z on day 4, it is a good choice to deliver to customer x on day 4 instead of day 5. For each realization, our decision to schedule a delivery for a telemetry customer will adjust for the observed usage rates during the days leading to the tentative delivery date. If we observe a significantly low usage rate for the next few days in one realization, the projected last day will move further into the planning horizon. This would allow us to postpone the delivery day for the customer with telemetry in that realization. Thus, the realizations of daily amounts of usage for customers with telemetry allow us to reevaluate our delivery decisions and potentially improve them. The downside of the best day rule is the long term impact, in that it may create a few more visits over the time horizon. We compare the value of a best day and last day approach in our computational experiments in Section 1.5.

1.4.3 Initial Inventory

To be able to decide the precise delivery dates under any delivery rule, an initial inventory level is required. As indicated in the model in Section 1.3.4, we want to use the same initial inventory level for a customer across each realization. This is so that the associated delivery dates will reflect the variety in dates that a particular customer and their usage variation could create. At the same time, these initial values should be chosen to provide the best chance at finding good routes, since this is a strategic problem and initial levels could easily be modified in practice. It is also important that the telemetry customers have different initial inventory levels (i.e., do not set all customer's inventory level at a full tank), so the telemetry customers will not all need a delivery at the same time. After careful experimentation, we found that the lowest cost solutions are found when the initial inventory levels are updated based on the set of customers currently assigned to telemetry. For each potential assignment that is considered, we update the initial inventory values based on a PVRP solution, again using the tabu-search algorithm presented by Cordeau et al. [21]. As in the first part of the algorithm, the current set of telemetry customers are routed in a PVRP with a frequency of $\hat{f}_i$ and non-telemetry customers are routed with their f_i' values. For each customer in the telemetry set, we look at the first delivery date for that customer in the PVRP solution and calculate the time between this delivery date and start of planning horizon e.g., g_i days. We set the initial inventory level at each telemetry customer equal to the mean usage plus safety stock, at a service level corresponding to z, for this time duration or

$$I_i^0 = g_i\mu_i + z\sqrt{g_i}\sigma_i. \tag{1.31}$$

The choice of this initial inventory level makes the solution to the PVRP feasible for customers with low variation, but creates big changes for customers with high variation in usage.

1.4.4 Initialization and Telemetry Set Improvement

Next, we consider ways to both initialize the set of customers that have telemetry and strategic ways to iteratively improve this set. These choices have a relatively small impact on the final solution cost, but have a big impact on the run time. Most improvements we consider replace one telemetry customer with one non-telemetry customer at a time, evaluating the cost of this new assignment using the procedure described in Section 1.4.1 and keeping the new assignment if an improvement in cost

Algorithm 1.1 RoutingCost

Input: S the set of customers assigned to telemetry

Initialization:

x= 0

Solve PVRP using delivery frequency $\hat{f}_i$ for $i \in S$ and $f_j^{'}$ for $j \in N \setminus S$ to create set of routes $ROUTES$

Construction:

Extract telemetry customers from $ROUTES$ and replicate for each realization

while $S \neq \emptyset$ **do**

  for $i \in S$ **do**

    for $r = 1$ to R **do**

      $IC_i^r =$ lowest insertion cost for i into $ROUTES_r$ based on best day or last day rule

    end for

    $IC_i = \sum_r IC_i^r$

  end for

  $k =$ customer $i \in S$ with largest IC_i value

  for $r = 1$ to R **do**

    Insert customer k into $ROUTES_r$

  end for

  $S = S \setminus k$

end while

for $r = 1$ to R **do**

  $x = x+$ cost of $ROUTES_r$

end for

$x = \frac{x}{R}$

Improvement:

Update x and $ROUTES_r$ ($r = 1$ to R) using one-shift improvement heuristics

Update x and $ROUTES_r$ ($r = 1$ to R) using 2-OPT improvement heuristics

Update x and $ROUTES_r$ ($r = 1$ to R) using VRP crossover improvement heuristics

Return x

is found. Preliminary experiments show that a first-improving search is much faster than a best-improving search while achieving similar quality solutions. This process is repeated until no improvement can be found, as summarized in Algorithm 1.2. The one-exchange methods we consider vary in how customers are ordered for the one-exchange. We describe the different ordering methods for the initialization and one-exchange here and evaluate them in Section 1.5.

1.4.4.1 Method 1

The first method, which we refer to as Method 1, uses f_i' and $\hat{f}_i$ values to estimate the impact of telemetry on each customer and is based on the same ideas as found at the end of Section 1.3.2. The initialization values are created using the tabu search heuristic from Cordeau et al. [21] to solve the PVRP with all customers assigned to their original $f_i^{'}$ values. With this solution, we compute C_i values for each customer, where C_i reflects the total cost to serve customer i without telemetry in the PVRP solution. The total cost for a customer is computed as follows. Each time that customer i is served, we compute the insertion cost based on its location relative to preceding and succeeding customers. The insertion cost represents the extra cost required to add i to an existing route. These insertion costs are summed over the planning horizon to create the total insertion cost C_i which is a reflection of the cost to serve customer i without telemetry. The following expression captures the estimated reduction in C_i associated with assigning a customer to telemetry and

Algorithm 1.2 Pseudocode for Initial Allocation and Improvement

Initialization:

$e(i) = COST(i, m)$, $i = 1, \ldots, n$ and m=1, 2, or 3.

Choose M largest $e(i)$ values as S

x= RoutingCost(S)

Improvement:

$e(i) = COST(i, m)$, $i = 1, \ldots, n$ and m=1, 2, or 3.

Sort $i \in S$ from lowest to highest $e(i)$ value

Sort $j \in N \setminus S$ from highest to lowest $e(i)$ value

```
for i = from lowest e(i) to highest e(i) value in S do
    S' = T \ i
    for j= from highest e(j) to lowest e(j) value in N \ S do
        S' = S' ∪ j
        x' = RoutingCost(S')
        if x' < x then
            S = S'
            x = x'
            Goto start of Improvement
        end if
        S' = S' \ j
    end for
    S' = S' ∪ i
end for
Return x
```

is fairly simple to implement:

$$\frac{C_i}{\frac{T}{f_i'}}\left(\frac{T}{f_i'} - \frac{T}{\hat{f}_i}\right). \tag{1.32}$$

The average cost per visit is captured by the first term in Equation 1.32, and the expected reduction in number of visits is given by the second term. Thus, Equation 1.32 captures the estimated reduction in cost associated with adopting telemetry for customer i.

We can evaluate Equation 1.32 for each customer and choose the M customers that have the maximum value as our initial set of telemetry customers. This approach has been shown to greatly reduce run times as compared with starting with a random initial allocation.

Once we have an initial set of customers with telemetry, we need to identify an ordering of how we consider alternative solutions. We want to order the non-telemetry customers so that we consider adding customers who are likely to create cost improvements with the addition of telemetry. For the telemetry customers, we want to remove customers that are not creating much actual savings as a result of telemetry. For both sets of customers, we simply adapt the above calculations based on the current solution. For each customer, we compute its insertion cost C_i^r for each realization and average these to create a new value C_i'. This value represents the current cost to serve customer i on average over the R realizations. For each customer, we also compute the actual number of visits a_i^r for each realization r and average them over all realizations to create a value a_i. For non-telemetry customers, the projected savings from telemetry can be computed by adapting Equation 1.32 as:

$$\frac{C_i'}{a_i}(a_i - \frac{T}{\hat{f}_i}). \tag{1.33}$$

The first part represents the updated average cost per visit, and the second part represents the updated projection of change in number of visits. The higher this value is, the higher the projected savings from using telemetry. Thus, we consider non-telemetry customers from highest to lowest value in our one-exchange. Similarly, for telemetry customers, the projected cost increase from removing telemetry is:

$$\frac{C_i'}{a_i}(\frac{T}{f_i'} - a_i). \tag{1.34}$$

The higher this value is, the higher cost increases are expected from removing telemetry. Thus, we sort this list from low to high, so the customers with the lowest value are considered first for removal of telemetry.

These cost estimates are used in Algorithm 1.2 as indicated with $m = 1$. Each time an improvement is found, the telemetry set is changed, and the new solution is used to recompute the above values.

1.4.4.2 Method 2

Method 2 is based on ideas from Ohlmann et al. [39] which determines delivery frequencies for customers in a lean production system. Though our problem is much different in terms of the more limited delivery options for customers, we can still use some of their ideas on how to identify which customers are not being served efficiently. Specifically, we look at which customers are being served on routes that do not use a full truck's capacity, as these do not represent efficient use of vehicles. In terms of initialization, we again look at the PVRP solution where all customers are routed at their $f_i^{'}$ values. For each customer, we look at each route where they are served, compute the volume being delivered, and subtract it from the truck's capacity to compute a remainder value. We can average these over all routes where a customer is being served to create a value rem_i. The higher the value of rem_i, the poorer quality the routes used to serve customer i. We can choose the customers with the M highest values of rem_i as those most needing a change, which here is the placement of telemetry.

Once we have an initial set of customers with telemetry, we need to identify an ordering of how we consider alternative solutions. As with Method 1, we want to order the non-telemetry customers so that we consider adding customers who are likely to create cost improvements with the addition of telemetry. For the telemetry customers, we want to remove customers that are not creating much actual savings as a result of telemetry. For both sets of customers, we simply update the above calculations based on the current solution. For each customer, we compute the average remaining volume on its deliveries rem_i^r for each realization and average these to create a new value $rem_i^{'}$. This value represents the average volume remaining on routes visiting customer i across all realizations. Again, if it is high, it reflects the customer needs a change. For non-telemetry customers, we consider these values from highest to lowest, and for telemetry customers, we consider this value from lowest to highest.

As with Method 1, these cost estimates are used in Algorithm 1.2 as indicated for $m = 2$. Each time an improvement is found, the telemetry set is changed, and the new solution is used to recompute the above values.

1.4.4.3 Method 3

The obvious drawback of Method 2 is that it does not consider the cost of the routes, but only the remaining volume. Our last method is based on a combination of Methods 1 and 2. We want to consider customers for telemetry that are not currently being served well in terms of truck volume, but we want to prioritize those where this creates the highest costs. Our hybrid initialization combines the two methods above as follows:

$$C_i(1 - \frac{f_i^{'}}{\hat{f}_i})\frac{rem_i}{P}. \tag{1.35}$$

Our initialization method for Method 3 chooses the M customers with the largest value defined by Equation 1.35.

Next, we again need a method to evaluate the candidates for the one-exchange. For customers without telemetry, we evaluate them using Equation 1.36:

$$\frac{C_i^{'}}{a_i}(a_i - \frac{T}{\hat{f}_i})\frac{rem_i^{'}}{P}. \tag{1.36}$$

This yields a large value for customers who are not currently being served well and have a high projection of cost savings from adding telemetry. We consider non-telemetry customers for the one-exchange based on the value defined by Equation 1.36 from highest to lowest. For telemetry customers, we use Equation 1.37:

$$\frac{C_i^{'}}{a_i}(\frac{T}{f_i^{'}} - a_i)\frac{rem_i^{'}}{P}. \tag{1.37}$$

We consider telemetry customers for the one-exchange from lowest to highest values based on Equation 1.37. These cost estimates are used in Algorithm 1.2 as indicated for $m = 3$.

The downside of Method 3 is that the estimate we compute above does not provide a particular translation in terms of units, where Method 1, for example, gives an expected cost reduction. The advantage of Method 3 is how it performs computationally, as will be demonstrated in Section 1.5. Like the previous two methods, the values for Method 3 are recomputed for each customer each time an improving solution is found.

1.4.4.4 Pair-Exchange

Once the one-exchange is complete, a two-exchange can be performed to create additional improvements in the telemetry set. For this improvement method, we can order the customers in the telemetry set and the customers without telemetry using the methods described above. To increase the speed of the two-exchange, we simplify the process by grouping each customer set into pairs and only considering the exchange of these pairs in a process similar to that described in Algorithm 1.2. For example, for the customers in the telemetry set, we first consider removing the pair with the two lowest $e(i)$ values, then the next two lowest $e(i)$ values, etc. For the customers without telemetry, we first consider the pair with the two highest $e(i)$ values, then the two with the next two highest $e(i)$ values, etc. The exchanges are performed until no improvement can be found. Because this is not a full two-exchange procedure, we refer to it as a pair-exchange. We also implement this improvement method in a first-improving manner for speed purposes.

1.5 Computational Study

A variety of computational experiments were performed to assess the impact of the proposed solution strategies. We also wanted to assess the impact of different problem characteristics on the solutions, including the number of telemetry units, the customer tank capacities, average daily usage rates, and amount of variance in the usage rates.

The algorithms were implemented and executed in MATLAB Version 7.8.0.347. The experiments were performed on a 2.40 GHz Intel Core 2 Quad processor running the LINUX operating system [4].

Our computational analysis uses eight datasets for the symmetric traveling salesman problem from TSPLIB [7]. These include dantzig42, st70, eil101, tsp225, gil262, rat783, dsj1000 and pr1002. The number associated with the name represents the number of customers in the dataset. For example, dantzig42 has 42 customers. The distances between customers for four datasets (dantzig42, st70, and gil262) were computed directly from the location data in miles with no adjustments, while for four others (tsp225, rat783, dsj1000 and pr1002), the distances were scaled so that no two customers were more than eight hours travel time apart with the use of a travel speed of 45 miles per hour. The depot was fixed at the mean latitude and mean longitude position for each dataset. We chose these datasets because they represent a nice variety in terms of both the numbers of customers and how the customers are distributed across the geographical area.

In our experiments, we examine the average costs across 100 usage realizations over a period of 12 weeks. This time period is chosen so that every customer will have at least one delivery given our choice of usage rates. For each dataset, we begin with routing customers with no telemetry units to get a baseline cost. Because each potential local search move involving a new assignment of customers to telemetry involves a run of the routing algorithm described in Section 1.4.1, it can add up to a significant computational time for a large set of customers. To reduce the run time for larger sets of customers, we reduce the candidate set of customers that are considered for possible telemetry placement to those with the top 50% (for 200-500 customers) or 25% (for more than 500 customers) values for the selected initialization method (described in Section 1.4.4).

The number of vehicles for each experiment is the minimum number required to create a feasible solution to the PVRP with no telemetry units. This minimum is largely driven by the restriction that routes are no more than eight hours long. We set the fixed stop time to 2.67 minutes and a variable delivery time to .04 minutes per unit. The time limit on routes is set to eight hours. We assume usage and deliveries occur five days per week. These values are based on real values at $NuCO_2$. For most experiments, we solve for the best solution with 5, 10, and 20 telemetry units for each dataset.

For the base case, we assumed a tank capacity of 350 units at every customer.

The mean usage rate μ_i for each customer is randomly assigned a value between 10 units per day and 30 units per day. We fix the variance in usage rate as a fixed proportion of the mean usage rate such that $\sigma = 0.25\mu$. We assume all trucks have a capacity of 6000 units and use a service level of 95%. Again, these values are based on real data at $NuCO_2$.

1.5.1 Rule Comparison

We begin with a comparison between the proposed last day and best day delivery rules for telemetry customers. Both rules were used with Method 1 for initialization and improvement of the telemetry set. We report the average cost across all the realizations.

The results are found in Table 1.1. The columns "Dataset", "Tele", "Veh", "Last Day", "% Imp 1" and "Time 1" represent the name of the dataset, number of telemetry units used, number of vehicles used, best solution cost for the last day rule, percentage change in cost with respect to the case without any telemetry units, and CPU time taken for last day rule in seconds. "Best Day", "% Imp 2" and "Time 2" are the corresponding parameters for the best day rule. The "% Diff " field represents the difference in the objective value of the Best Day with respect to Last Day in terms of average routing costs with the same number of telemetry units.

Looking at Table 1.1, we see that the best day rule leads to better solution quality in all of our tests. The percentage improvement ranges from 0.12% to 2.15% vs. the last day rule.

Table 1.1: Base Case: $\mu_i \in [10, 30], \sigma_i = 0.25\mu, I_i^{max} = 350$

Dataset	Tele	Veh	Last day	% Imp 1	Time 1	Best Day	% Imp 2	Time 2	% Diff
dantzig42	0	1	9698.10			9698.10			
	5		9585.24	1.16	3008.53	9539.03	1.64	3622.54	0.48
	10		9536.82	1.66	5285.74	9486.92	2.18	6005.28	0.51
	20		9480.92	2.24	13564.93	9382.46	3.25	14735.94	1.02
st70	0	1	8930.43			8930.43			
	5		8795.02	1.52	4436.94	8718.34	2.37	4930.62	0.86
	10		8716.14	2.40	10175.25	8637.90	3.28	11006.28	0.88
	20		8554.03	4.21	20541.96	8452.62	5.35	21752.70	1.14
eil101	0	1	7184.78			7184.78			
	5		7009.25	2.44	5973.43	7000.52	2.56	6712.84	0.12
	10		6687.96	6.91	13385.41	6658.93	7.32	14263.27	0.40
	20		6427.73	10.54	20292.67	6324.72	11.97	21545.82	1.43
tsp225	0	6	70150.21			70150.21			
	5		67952.35	3.13	10002.53	67817.49	3.33	11104.62	0.19
	10		67158.41	4.26	18562.92	67004.96	4.48	20034.84	0.22
	20		65993.36	5.93	31428.58	65791.44	6.21	33752.65	0.29
gil262	0	4	43854.59			43854.59			
	5		43172.32	1.56	9277.42	42299.52	3.55	10472.02	1.99
	10		42915.21	2.14	24621.15	41973.73	4.29	26911.74	2.15
	20		41262.75	5.91	48027.17	40342.84	8.01	50107.21	2.10
rat783	0	6	57433.95			57433.95			
	5		56962.93	0.82	19755.93	56285.63	2.00	21662.84	1.18
	10		55975.25	2.54	35182.02	55212.57	3.87	37006.29	1.33
	20		54971.46	4.29	60154.55	54182.11	5.66	63418.28	1.37
dsj1000	0	8	109259.28			109259.28			
	5		106288.35	2.72	24093.42	105297.65	3.63	25935.74	0.91
	10		103872.67	4.93	37548.92	102673.52	6.03	41053.63	1.10
	20		103105.70	5.63	75133.05	101387.42	7.20	78209.85	1.57
pr1002	0	8	88575.32			88575.32			
	5		85913.62	3.01	21038.04	85595.03	3.36	23694.55	0.36
	10		84778.25	4.29	36628.72	84454.94	4.65	39005.02	0.37
	20		84002.94	5.16	68425.56	83247.55	6.01	72221.52	0.85

This was predictable as the best day rule encompasses the last day rule. We notice that there is also an increase in solution time associated with using the best day rule. This comes from the fact that the last day rule is less computationally intensive than the best day rule. Since the best day rule provides better solution quality with a slight increase in computational time, we use it in the rest of the tests in this section.

In Table 1.1, it is very interesting to see the impact of different numbers of telemetry units on the different datasets. One might expect that for a given number of telemetry units we would see a larger percentage change on the smaller datasets since the telemetry units would make up a larger proportion of the total customer set. This definitely is not true. With 5 telemetry units, we see the smallest change with the best day rule, 1.64% is actually with the smallest dataset we tested (dantzig42), and the largest change, 3.63% is with one of the largest datasets (dsj1000). This shows that it is hard to predict the impact of telemetry without a tool such as the one presented here. For all datasets, we see a steady improvement in the percentage savings with increasing number of telemetry units. For most datasets, the savings with 20 telemetry units is roughly double with 5 telemetry units. The notable exception, though, is eil101 with a savings of 2.56% with 5 units and a remarkable 11.97% with 20 telemetry units. Many of the customers in this dataset that require frequent deliveries are located far from the depot. This translates to higher cost savings from the adoption of telemetry units.

In terms of other observable trends, we notice that the run time increases as the size of the dataset increases. This happens because the complexity of the VRP subproblem increases. This is depicted by solution times for dantzig42 compared to dsj1000. The solution times can also depend on the distribution of customers.

1.5.2 Telemetry Set Initialization and Improvement Method Comparison

Next, we compare the proposed telemetry set initialization and improvement methods. The results are summarized in Table 1.2. The columns “Dataset”, “Tele”, “Method 1”, “Method 2”, “% Diff 2, “Method 3” , “% Diff3”, “Pair-Ex”, and “% Diff” represent the name of the dataset, number of telemetry units used, best solutions found using Method 1, best solutions found with Method 2, percent improvement of Method 2 relative to Method 1, best solution found with Method 3, percent improvement of Method 3 relative to Method 1, best solution with Method 3 and pair-exchange, and percent improvement relative to Method 3.

Since all three methods are based on the same one-exchange idea, it is not surprising that all three methods yield similar solution quality. We do observe that the Method 3 rule is slightly better than the other methods in terms of solution quality. It yields the same solutions as Method 1 for the smaller datasets, but it yields slightly better results for bigger datasets like rat783,dsj1000 and pr1002. Method 2 also does well for dsj1000 and pr1002 but is worse for st70 and eil101.

One key reason for developing different methods for the one-exchange process was to see if different orderings yield a faster convergence. Thus, the run times for these same tests are found in Table 1.3. The columns “Dataset”, “Tele”, “Method 1”, “Method 2”, “% Diff 2, “Method 3” , “% Diff3”, “Pair-Ex”, and “% Diff” represent the name of the dataset, number of telemetry units used, run time using Method

1, run time with Method 2, percent improvement of Method 2 relative to Method 1, run time with Method 3, percent improvement of Method 3 relative to Method 1, run time with Method 3 and pair-exchange, and percent improvement relative to Method 3. We see that Method 2 and Method 3 both offer an improvement in run time relative to Method 1 for all tests, but Method 3 offers a larger improvement across all tests. The run time improvements for Method 3 range from 4.32% for gil262 with 10 telemetry units to 26.30% for dantzig42 with 20 telemetry units. The large improvements found in run time with Method 3 do not seem to follow a pattern in terms of dataset size or number of telemetry units.

Since Method 3 is the best in terms of solution quality and run time, we ran the pair-exchange on the Method 3 solutions to see if additional improvements are possible. We see in Table 1.2 that the pair-exchange does improve solution quality slightly for bigger datasets ($\geq$ 200 customers) but the maximum is only an additional 1.18%. As Table 1.3 reflects, these additional improvements come at a significant increase in computational time. Based on our experiments, we recommend Method 3 for one-exchange and the addition of a pair-exchange if time permits. All of the remaining tests in this section will use the combination of Method 3 with pair-exchange.

We observed with Table 1.1 that all datasets show notable cost reductions from the adoption of telemetry. With the choice of Method 3 for telemetry set initialization and Method 3 plus pair-exchange for telemetry set improvement, we examine what happens as the number of telemetry sets increases further with the larger datasets.

Table 1.2: Solution Costs

Dataset	Tele	Method 1	Method 2	% Diff 2	Method 3	% Diff 3	Pair-Ex	%Diff
dantzig42	0	9698.10						
	5	9539.03	9539.03	0	9539.03	0	9539.03	0
	10	9486.92	9486.92	0	9486.92	0	9486.92	0
	20	9382.46	9382.46	0	9382.46	0	9382.46	0
st70	0	8930.43						
	5	8718.34	8754.92	-0.42	8718.34	0	8718.34	0
	10	8637.90	8685.55	-0.55	8637.90	0	8637.90	0
	20	8452.62	8501.83	-0.58	8452.62	0	8452.62	0
eil101	0	7184.78						
	5	7000.52	7062.47	-0.88	7000.52	0	7000.52	0
	10	6658.93	6684.09	-0.38	6658.93	0	6658.93	0
	20	6324.72	6376.65	-0.82	6324.72	0	6324.72	0
tsp225	0	70150.21						
	5	67817.49	67817.49	0	67817.49	0	67601.28	0.32
	10	67004.96	67004.96	0	67004.96	0	66582.58	0.63
	20	65791.44	65791.44	0	65791.44	0	65297.01	0.75
gil262	0	43854.59						
	5	42299.52	42299.52	0	42299.52	0	42116.25	0.43
	10	41973.73	41973.73	0	41973.73	0	41583.92	0.93
	20	40342.84	40342.84	0	40342.84	0	40002.75	0.84
rat783	0	57433.95						
	5	56285.63	56285.63	0	56285.63	0	56062.52	0.40
	10	55212.57	55212.57	0	55183.96	0.05	54728.95	0.82
	20	54182.11	54182.11	0	54135.06	0.09	53845.02	0.54
dsj1000	0	109259.28						
	5	105297.65	105185.75	0.11	105185.75	0.11	105118.95	0.06
	10	102673.52	102602.98	0.07	102557.01	0.11	102255.69	0.29
	20	101387.42	101182.55	0.20	101095.65	0.29	99903.75	1.18
pr1002	0	88575.32						
	5	85595.03	85362.52	0.27	85337.63	0.30	85304.89	0.04
	10	84454.94	84227.36	0.27	84201.96	0.30	84103.56	0.12
	20	83247.55	83005.77	0.29	82985.52	0.31	82683.44	0.36

Table 1.3: Solution Time

Dataset	Tele	Method 1	Method 2	% Diff 2	Method 3	% Diff 3	Pair-Ex	%Diff
dantzig42	0							
	5	3622.54	3383.46	6.60	3123.59	13.77	4003.65	-28.17
	10	6005.28	5418.71	9.77	5095.63	15.15	6245.32	-22.56
	20	14735.94	12528.63	14.98	10860.35	26.30	17535.35	-61.46
st70	0							
	5	4930.62	4582.33	7.06	4178.65	15.25	5379.25	-28.73
	10	11006.28	10683.07	2.94	10127.83	7.98	15572.36	-53.76
	20	21752.70	19045.64	12.44	16584.69	23.76	28095.73	-69.41
eil101	0							
	5	6712.84	6492.63	3.28	6196.93	7.69	9003.36	-45.29
	10	14263.27	13825.76	3.07	11475.03	19.55	17905.72	-56.04
	20	21545.82	20053.46	6.93	17565.09	18.48	24827.25	-41.34
tsp225	0							
	5	11104.62	9907.53	10.78	9853.44	11.27	13085.36	-32.80
	10	20034.84	19635.93	1.99	15795.25	21.16	22756.68	-44.07
	20	33752.65	31759.2	5.91	30052.00	10.96	35535.72	-18.25
gil262	0							
	5	10472.02	10175.88	2.83	9935.02	5.13	15835.70	-59.39
	10	26911.74	25905.64	3.74	25750.45	4.32	35324.08	-37.18
	20	50107.21	45767.13	8.66	42362.86	15.46	56629.92	-33.68
rat783	0							
	5	21662.84	18655.72	13.88	17759.06	18.02	28936.36	-62.94
	10	37006.29	34925.56	5.62	31083.42	16.01	44863.93	-44.33
	20	63418.28	60528.93	4.56	53866.93	15.06	67835.52	-25.93
dsj1000	0							
	5	25935.74	23062.24	11.08	21002.52	19.02	28952.92	-37.85
	10	41053.63	40075.65	2.38	35132.63	14.42	48036.03	-36.73
	20	78209.85	76936.38	1.63	61859.94	20.91	80268.42	-29.76
pr1002	0							
	5	23694.55	21969.52	7.28	20003.34	15.58	25762.52	-28.79
	10	39005.02	37045.63	5.02	31926.45	18.15	42965.16	-34.58
	20	72221.52	68016.12	5.82	60075.57	16.82	76836.42	-27.90

Figure 1.1 plots the percentage savings from the use of telemetry vs. the number of telemetry units for dsj1000 and pr1002. We ran additional experiments using 30 and 50 telemetry units for these two datasets. We can see from both datasets that most of the benefit is from first 20 telemetry units with little improvement after that. This shows the value of installing a very limited number of telemetry units. Small and medium businesses can benefit from this insight.

1.5.3 Sensitivity Analysis

Based on Section 1.3.2, we were interested in varying the parameters μ_i, σ_i and I_i^{max} to see how they alter the results. These tests are performed by changing only one parameter value at a time. There is not much to be learned from comparing the costs of the final solutions, but it is interesting to analyze the changes in the set of customers chosen for telemetry placement. Based on Section 1.3.2, customers with high σ_i or μ_i should have a higher chance of being picked as telemetry candidates. We changed these parameters for a random 25% of the customers. For the first test, μ_i is selected between 30 units per day and 50 units per day for the selected customers. In the second test, σ_i is increased at 25% customer sites such that $\sigma_i = 0.75\mu_i$. We also expect customers with a smaller tank size to be preferred for telemetry. Thus, in our third test, we wanted to reduce the tank capacity for selected customers. Reducing tank capacity below the 350 value in our base case creates feasibility issues, so the tank capacity at 25% of the customer sites will be 350 units, while the rest of the customers will have a larger tank capacity of 500 units.

In our experiments, we examine how many of the customers with a different

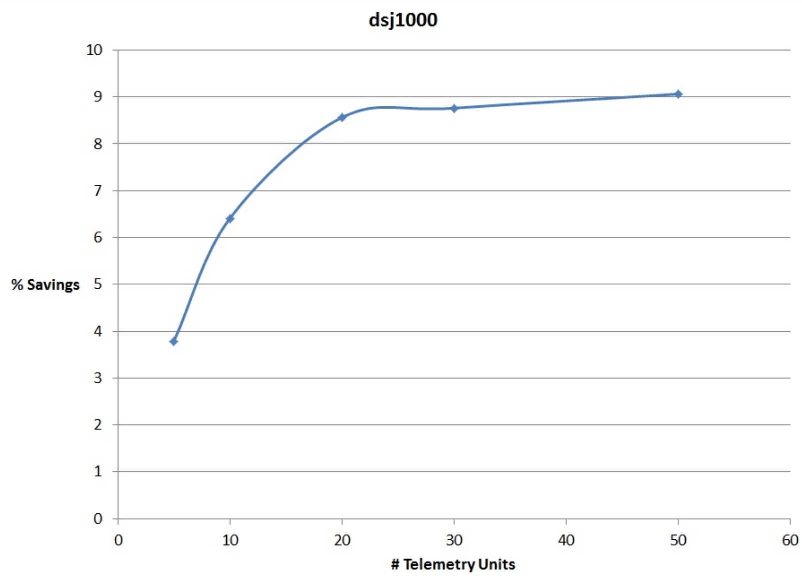

(a) Savings for dsj1000

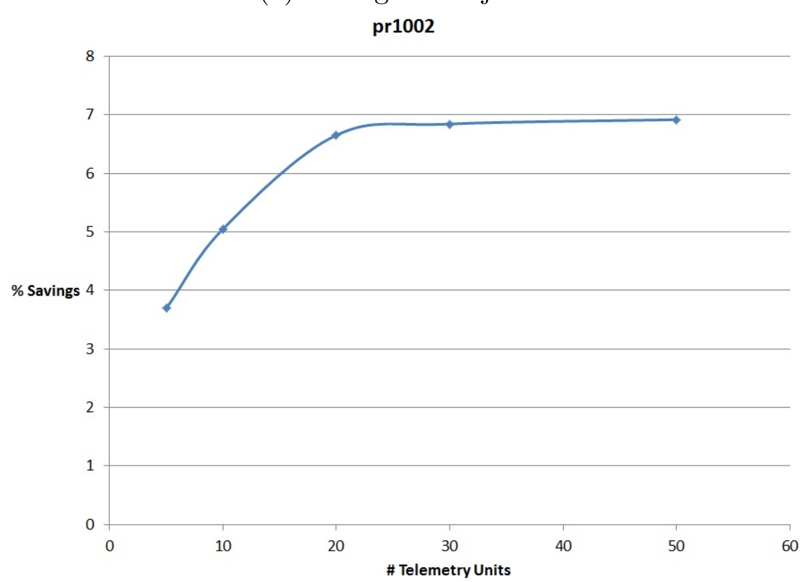

(b) Savings for pr1002

Figure 1.1: Savings vs. Number of Telemetry Units

parameter value are selected for telemetry. We look at all datasets but restrict our test to involve 20 telemetry units. The results are shown in Table 1.4. The columns "Dataset", "Impacted Customers" and "Selected" represent the name of the dataset, the number of customers whose control parameters are changed, and the total number of customers among the control set that are picked in the final 20 customer telemetry set.

Table 1.4: Sensitivity Analysis

Dataset	Impacted Customers	Selected σ	μ	I^{max}
dantzig42	11	7	6	6
st70	18	12	12	12
eil101	26	13	12	12
tsp225	57	17	16	16
gil262	66	15	16	15
rat783	196	18	16	17
dsj1000	250	18	17	17
pr1002	251	19	18	19

As we can observe from Table 1.4, the impacted customers are more likely to be selected for telemetry. At least 12 out of 20 customers picked to have telemetry are from the independent control sets for all datasets except dantzig42. These results support our claim that high μ_i or σ_i values or low I_i^{max} values tend to be good candidates for telemetry. It is also interesting that this pattern becomes more pronounced with larger customers. With smaller datasets, the location of individual customers plays a bigger part in the telemetry selection.

1.5.4 Graphical Example

Last, we examine the customers that are selected for telemetry in more detail for one dataset. We analyze the results for pr1002. Figure 1.2 depicts the location of the 5, 10 and 20 customers selected for telemetry. We observe that as the number of telemetry units increases, more parts of the graph include customers assigned to telemetry. For example with 5 telemetry units, there are no telemetry customers in the lower right of the graph, but there are many there with 20 telemetry units. There is also, though, a noticeable clustering of telemetry units as the numbers gets bigger, so that the customers with the telemetry can be served on the same routes. Thus, it appears that the interaction between telemetry customers becomes more important in the selection process as more are added. The result is that the customers selected to receive 5 telemetry units are not necessarily included within the set picked with 10 telemetry units. Here, two customers (4 and 936) picked in the set of 5 do not get picked in the set of 10. Similarly, four customers among the 10 customers selected for telemetry do not appear in the 20 customers selected for telemetry. This is important because it suggests that we cannot simply add customers to the optimal set of telemetry customers to yield a new optimal set of telemetry customers for a larger number of telemetry units.

1.6 Conclusions

In this paper we have presented the TLP where we decide the best allocation of telemetry units to minimize routing costs, while preserving a given service level for customers. We have developed the following insights:

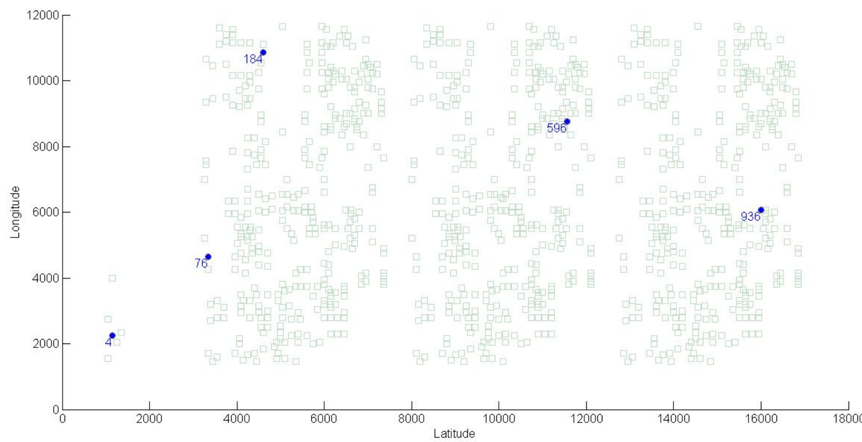

(a) 5 Telemetry Units

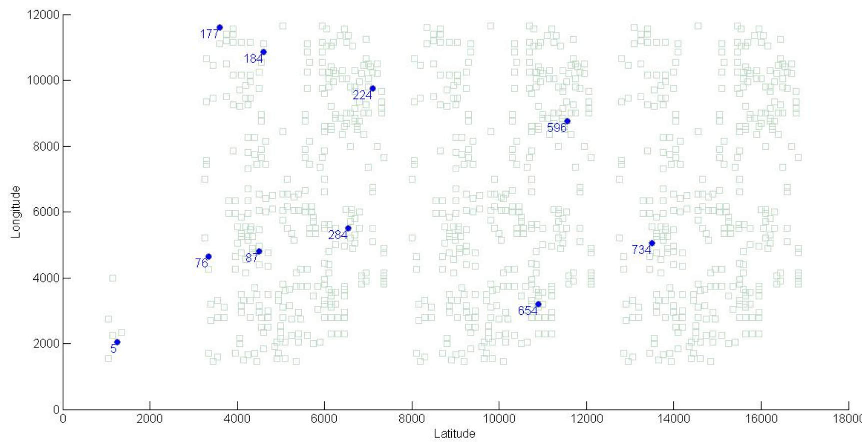

(b) 10 Telemetry Units

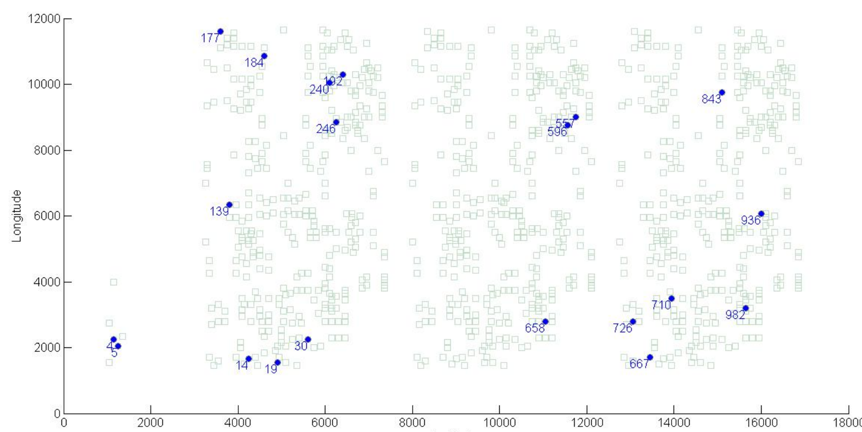

(c) 20 Telemetry Units

Figure 1.2: Distribution of Telemetry Units

- Our results demonstrate that there can be significant savings from a small number of telemetry units.
- The best day rule for visiting telemetry customers creates a lower average routing cost while meeting or exceeding customer service levels.
- The choice of initialization and ordering of customers in the one-exchange improvement can impact both solution quality and run time. The improvement in solution quality is rather small and limited to larger datasets, but the savings in run time is significant for all dataset sizes.
- The savings from telemetry tends to level off at a certain number of datasets. Running a decision support tool such as ours can be helpful in deciding the number of telemetry units to install as well as the locations.
- Customers with a higher than average daily usage rate, a higher than average variance in usage, or a smaller tank capacity are more likely to be chosen for telemetry adoption.
- The number of telemetry units has an impact on the customers selected due to potential interactions on routes.

The TLP offers many potential areas for future research. Beside alternative solution methods, it may also be useful to consider different objective functions. We are concerned here with routing costs only. Even if we observe that the telemetry units quit creating big savings in routing costs at some point, they can still reduce the number of stops due to reduced number of visits, which may also translate to big savings for the company. In the future, we also plan to examine the impact that customer usage correlation can have on the choice of where to place telemetry units. For example, if two customers in our set both represent fast food restaurants of a given brand, a spike in demand at one location may mean that there is also a spike in demand at the other, such as a result of a popular new product. If we place a telemetry unit at one of these two customers, it will give us information about the other if the demand is highly correlated.

CHAPTER 2
STRATEGIC PLACEMENT OF TELEMETRY UNITS CONSIDERING CUSTOMER USAGE CORRELATION

2.1 Introduction

Telemetry units relay daily customer demand information to the supplier. This information can be used to create good estimates of future delivery dates, which is useful for route planning. Our interest in telemetry stems from a project with NuCO2. NuCO2 provides carbon dioxide and nitrogen gas to more than 130,000 customers across the United States, mostly for use in beverage carbonation [1]. These customers include national chains and local restaurants, as well as convenience stores, sports venues, and theme parks. The company was considering the use of telemetry units to track inventory levels to get a better idea of when customers needed to be visited instead of using precomputed delivery frequencies. However, telemetry units are expensive. The company approached us with the problem of where to place a limited number of these units. We found that significant savings in routing costs can be achieved with a careful choice of the location of telemetry units [45].

Our algorithm in [45] did not explicitly capitalize on the demand correlation that may be present among customer locations. In our analysis of company data, we observed that more than half of the customers in the cities we studied were part of a chain of stores or restaurants. Some of these chains included 20 or more locations. The usage rates among the stores belonging to the same chain were found to be highly correlated as defined by ([23]). This means that if, for example, the sales were high one day at a restaurant because of a popular new product, the sales were likely to be high at the other restaurants belonging to the same chain on the same day.

This practical observation suggested that we should carefully explore the effect of correlation on the strategic choice of where to place telemetry units. For example, if the algorithm in [45] suggested we place telemetry units on customers 1 and 2, the extra knowledge of perfect positive correlation between the two customers would render one of these telemetry units redundant. To the best of our knowledge, the impact of correlation on the strategic placement of telemetry units has not been previously studied in the literature. Our goal is to identify where to place a limited number of telemetry units incorporating usage correlation that exists between some customers to minimize expected routing costs while preserving a given service level.

The use of telemetry is common. For example, industrial gas suppliers, soil monitoring systems, and vending machines use telemetry for real time information tracking ([6],[3],[8]). The costs of telemetry units in all of these industries can be high. For some small and medium businesses, installing telemetry at all customers may not be financially viable. However, our research has shown that significant savings in routing costs can be realized by using a limited number of telemetry units. Knowledge of customer usage correlation can be useful in further reducing expected routing costs via strategic placement of telemetry units in all of these applications. This is the issue addressed in this paper.

In this article, we present the related literature in Section 2.2. In Section 2.3, we formally define the problem we solve. In Section 2.4, we discuss how the given service level can be preserved for non-telemetry, telemetry and correlated customers.

The details of our solution method are presented in Section 2.5. This involves a local search among different assignments of customers to telemetry as well as heuristics for solving the resulting routing problems in each of the realizations. We examine the performance of our approach as well as the impact of different problem characteristics in Section 2.6. We provide our final conclusions in Section 2.7.

2.2 Literature Review

The related literature can be divided into three major categories: (1) studies of information sharing in inventory management involving demand correlation among multiple products, multiple customers, or time (2) studies concerned with use of correlation specifically for the Vehicle Routing Problem (VRP) (3) studies on how to generate correlation coefficients.

Demand correlation among different inventory items is studied in the literature (see [37], [31] and [24]). For example, Liu and Yuan [31] consider a two-item inventory system with correlated Poisson demands. They model the two-item inventory system involving correlation using a bivariate Markov process. The coordinated replenishment of the two correlated products leads to savings in ordering or setup costs and transportation costs. The analytical results show that as demand correlation increases, the overall demand uncertainty decreases, and the total holding cost reduces.

Correlation among demands of different customers is also investigated in the literature (see [42], [47] and [36]). For example, Zhu and Thonemann [47] consider a supply chain with a single retailer and multiple customers, where customer demands are normally distributed and correlated. Each pair of customer demands are correlated with the same correlation coefficient. A customer's demand forecast for the following period is available to the retailer if the customer and the retailer share Future Demand Information (FDI). Sharing FDI incurs a cost. The degree of imperfection in FDI is captured by a parameter which ranges between 0 and 1. The authors consider two decisions for the retailer: (1) the optimal number of customers sharing FDI with the retailer (2) the optimal order quantity for the retailer. They develop an optimal solution through a two-stage dynamic program. The objective is to minimize the expected sum of information, ordering, shortage penalty, and inventory holding costs for the retailer. The authors derive equations for the expected total cost and the optimal reordering policy that involve the correlation coefficient. They infer that FDI is crucial when the information cost is low, the demand correlation is high, or the demand uncertainty is high.

Erkip et al. [25] consider a depot-warehouse system where demands are correlated both between warehouses and over time. They create a newsvendor model to make ordering decisions for the warehouses to minimize the total expected inventory holding and backorder penalty costs across the network. Experiments suggest that a high level of correlation among both warehouses and time increases the standard deviation in total demand resulting in larger amounts of safety stock compared with the uncorrelated case.

The VRP has been examined in conjunction with demand correlation among customers in Chiang [18]. The author studies a variant of the VRP where the objectives are to minimize the total travel distance and the number of vehicles while the demands are correlated. Also, the total demand is satisfied with a given service level. The authors develop a genetic algorithm to solve practical instances. In the solutions, customers with negatively correlated demands are assigned to the same route, and customers with positively correlated demands are assigned to different routes.

Another related research area deals with the generation of correlation coefficients. A correlation coefficient matrix is a symmetric positive semidefinite matrix with a unit diagonal. Xu and Evers [46] study the generation of valid correlation matrices where the correlations between all pairs of variables can be different. The authors postulate that, given correlation coefficients between $X1$ and $X2$ and between $X2$ and $X3$, the correlation coefficient between $X1$ and $X3$ is bounded by a function of the known correlation coefficients. Bucci et al. [13] solve this problem for a practical instance by creating a statistically valid correlation matrix whose terms are proportional to the distance between customer locations. They first group customers based on geographical location. A customer pair is assigned a correlation coefficient depending on the assigned groups. The degree of correlation between the customer pair decreases as the distance between a customer pair increases.

2.3 Problem Definition

The problem we want to solve is where to install M telemetry units to minimize the average routing costs across a set of R customer usage realizations over a given time horizon T while maintaining a given customer service level for all n customers. We assume that the daily usage rates of product for each customer follow a normal distribution. We use μ_i to represent the mean daily usage rate and σ_i for the standard deviation in daily usage rate for customer i ($i = 1$ to n). For each realization, the daily usage value is based on the μ_i and σ_i values.

A given number u of these n customers are divided into K groups of size $g_1, g_2, ..., g_K$. Every customer pair in each of these groups has demand correlated with a uniform correlation coefficient $\rho_k, k = 1, 2, ..., K$. We model correlation in this way to capture the correlation among stores of the same chain. We allow ρ_k to potentially vary across k to consider the effect that there may be more or less correlation among different kinds of stores.

We assume there is a fleet of V capacitated vehicles available each day, and each customer can be visited by at most one vehicle on a day (no split deliveries). We also assume that a vehicle fills up the customer's tank when a delivery occurs. The set of arcs A includes directed edges between all customers, as well as edges between each customer and the depot (indexed by 0 and $n + 1$). For each edge $(i, j) \in A$, we assign a distance t_{ij} and use a speed s to convert distance into travel time. The duration of a vehicle's tour has a given time limit l, and stops at customers require both a fixed stop time and variable stop time depending on the quantity delivered. We refer to this problem hereafter as the Correlated Telemetry Location Problem (CTLP).

2.4 Model

We will first look at how deliveries can be planned for different categories of customers in the CTLP. The first category consists of customers who do not have telemetry and are not correlated with any of the customers currently selected for telemetry installation. We will refer to these as non-correlated non-telemetry customers. The second category comprises the customers that are currently chosen for the installation of telemetry units. We refer to these as telemetry customers. The third category consists of non-telemetry customers that are correlated with a telemetry customer. We refer to these as correlated customers. The level of correlation will impact the delivery planning for only these customers.

To model the stochasticity of customer usage, we create a set of realizations of daily usage at telemetry customers over the entire planning horizon. The realizations reflect the type of information we would get from telemetry units on a daily basis. These simulated usage values help us create realistic delivery dates for telemetry and correlated customers for each realization, which helps us create a good assessment of the routing costs associated with a particular telemetry assignment. In each realization, the inventory level is reduced each night by the generated usage rate for that day and set to tank capacity when a delivery occurs. The following subsections discuss how we can use these updated inventory levels as well as other customer information to determine when deliveries are needed by each type of customer.

The service level is maintained by holding sufficient safety stock. In traditional inventory theory ([44]), safety stock at a customer location is given by:

$$z\sqrt{L}\sigma_i \tag{2.1}$$

where z is the z-statistic that corresponds to a particular in-stock probability (service level) , σ_i is the standard deviation of daily demand for a customer i, and L represents the lead time in days for a delivery once an order is placed. A longer lead time increases the size of safety stock. Here, lead time will be replaced by the number of days until the customer demand level is known since the customer will be vulnerable during those days.

2.4.1 Non-correlated Non-telemetry Customers

For customers without telemetry, we do not have the daily tank readings but only know the mean and standard deviation of usage. We can compute a value f_i for each customer that is the longest a customer can go between deliveries while preserving the service level. The value f_i replaces L in Equation 2.1. A customer i has a f_i (in days) when:

$$\mu_i f_i + z\sqrt{f_i}\sigma_i = I_i^{max} \tag{2.2}$$

where μ_i, σ_i and I_i^{max} denote the mean daily usage rate, standard deviation in daily usage and the tank capacity. Equation 2.2 says that a delivery every f_i days would allow a service level corresponding to z. The challenge with using a delivery frequency of f_i in practice is that deliveries may not always occur at the same time. For example, consider a customer with $f_i = 5$. If delivery on day 1 is scheduled at 9 a.m, then the next delivery on day 6 must be scheduled by 9 a.m. to guarantee the service level is maintained. To allow for flexibility in delivery time to promote good routing options, we suggest a adjusted maximum time between deliveries $f_i^{'} = f_i - 1$ for customers without telemetry as in Equation 2.3:

$$\mu(f_i^{'} + 1) + z\sqrt{f_i^{'} + 1}\sigma = I_i^{max}. \tag{2.3}$$

We can solve for $f_i^{'}$:

$$f_i^{'} = \frac{z^2\sigma_i^2 + 2\mu(I_i^{max} - \mu_i) - z\sigma\sqrt{(z^2\sigma_i^2 + 4\mu I_i^{max})}}{2\mu_i^2}. \tag{2.4}$$

Next, we will discuss how we will use $f_i^{'}$ to decide which day a customer i should receive their next delivery. For customer i with d_i days since the last delivery, a delivery must occur the next day if:

$$f_i^{'} - d_i = 1. \tag{2.5}$$

This is postponing the delivery as much as possible while preserving the service level. A better option from a routing perspective is to allow more flexibility in the choice of delivery day. For this purpose, we compute the delivery deadline D_i for customer i. This represents the number of days until the customer will need a delivery in order to preserve the service level. For these non-telemetry customers, D_i is given as:

$$D_i = f_i^{'} - d_i. \tag{2.6}$$

For example, if the deadline is 1, the delivery must happen by the next day. In this case, Equation 2.6 reduces to Equation 2.5. If we want to consider customers who require deliveries within the next 5 days in our routing plan, we consider those customers with $D_i \leq 5$.

All deliveries can be planned in this way except the first delivery. Assuming that we start with the tank at I_i^0 at the beginning of the planning horizon, the first delivery deadline D_i^0 is determined by replacing I_i^{max} with I_i^0 in Equation 2.3:

$$\mu_i(D_i^0 + 1) + z\sqrt{D_i^0 + 1}\sigma_i = I_i^0. \tag{2.7}$$

If $D_i^0 \leq 1$, the delivery must happen by the next day. If we want to consider customers with delivery flexibility of 5 days in our routing, we will consider those customers with $D_i^0 \leq 5$. The value of D_i^0 is reduced by one each day until the first delivery occurs.

2.4.2 Telemetry Customers

With telemetry, this situation changes. A customer is not vulnerable in the same way since the vendor knows the status of each customer's inventory level every night. Every night, the vendor can see if a customer's tank level is very low and can decide to go there the next day assuming capacity is available. Thus, a customer is only vulnerable for runout for a single day, since each day their inventory is viewed. This allows their safety stock to be greatly reduced while preserving the same service level. Now a delivery must occur the next day if :

$$\mu_i + z\sigma_i \geq I_i^t. \tag{2.8}$$

where I_i^t represents the tank level of customer i at the end of day t in that realization. Equation 2.8 says that if tomorrow's expected usage plus a safety stock buffer appears to more than what is remaining in the tank, it is time to make a delivery. If the tank level is higher than that, this formula implies the delivery can be postponed while preserving the service level. Like with non-telemetry customers, this may create challenges in delivery timing. For example, for a particular customer, this may indicate a delivery is not needed on day 3, but will require a delivery early on day 4 to ensure the service level is maintained. Thus, we slightly modify Equation 2.8 to require a delivery the next day, i.e. $D_i = 1$, when:

$$2\mu_i + z\sqrt{2}\sigma_i \geq I_i^t. \tag{2.9}$$

Thus, we are essentially forcing the delivery to occur one day earlier than necessary, so that we have the whole day to make the delivery.

We want to note that Equation 2.9 is not the only way a delivery can be made to a telemetry customer. We again can make a delivery earlier if it is more cost effective. For telemetry customers, the value for $D_i = w$, where $w \geq 1$, if w is the lowest integer where the following holds:

$$(w+1)\mu_i + z\sqrt{w+1}\sigma_i \geq I_i^t. \tag{2.10}$$

The value for $D_i^0 = w$ is obtained using Equation 2.10 by replacing I_i^t with I_i^0, assuming that we start with tank at I_i^0 at the beginning of the planning horizon.

The use of telemetry and the reduction in safety stock translates to an increase in the average time between deliveries relative to customers without telemetry. The average value for $\hat{f}_i$, the longest a telemetry customer can go between deliveries while preserving the service level, is computed as follows:

$$\hat{f}_i = \frac{I_i^{max} - 2\mu_i - z\sqrt{2}\sigma_i}{\mu_i}. \tag{2.11}$$

The value $\hat{f}_i$ value helps us understand what average cost savings may result over a time horizon as a result of using telemetry and is used in evaluating which customers make good telemetry candidates in Section 2.6.

2.4.3 Correlated Customers

If a non-telemetry customer is correlated with a customer that has telemetry, the telemetry reading will yield an updated mean and variance in usage for the correlated customer. In this way, the vendor can better estimate the daily usage at a customer correlated with a telemetry customer. This new estimate also depends on the degree of correlation between the two customers. However, if a telemetry customer is correlated with another telemetry customer, the best usage information is still received daily by the telemetry units. Thus, the correlation will not be considered in this case.

We will consider a customer y correlated with a set X of customers, all belonging to group j. The set X represents those customers in group j that have telemetry units. The customer y does not have telemetry. Assuming g_j^T represents the number of customers in group j that have telemetry units, according to the properties of conditional multivariate normal distributions in ([22]), if mean μ and covariance Σ are partitioned as follows:

$$\mu = \begin{bmatrix} \mu_y \\ \mu_X \end{bmatrix} \text{ with sizes } \begin{bmatrix} 1 \times 1 \\ g_j^T \times 1 \end{bmatrix} \tag{2.12}$$

$$\Sigma = \begin{bmatrix} \Sigma_{yy} & \Sigma_{yX} \\ \Sigma_{Xy} & \Sigma_{XX} \end{bmatrix} \text{ with sizes } \begin{bmatrix} 1 \times 1 & 1 \times g_j^T \\ g_j^T \times 1 & g_j^T \times g_j^T \end{bmatrix} \tag{2.13}$$

then, the distribution of $\mathbf{y}$ conditional on $\mathbf{X} = \mathbf{a}$ is multivariate normal $N(\bar{\mu_y}, \bar{\sigma_y^2})$ with mean

$$\bar{\mu_y} = \mu_y + \Sigma_{yX}\Sigma_{XX}^{-1}(\mathbf{a} - \mu_X) \tag{2.14}$$

and variance

$$\bar{\sigma}_y^2 = \sigma_y^2 - \Sigma_{yX}\Sigma_{XX}^{-1}\Sigma_{Xy}. \tag{2.15}$$

Here $\mathbf{a}$ is a vector of observed values for customers in $\mathbf{X}$ representing the daily observed usage values at the correlated telemetry customers. Note that $\Sigma_{xy} = \rho_{xy}\sigma_x\sigma_y$ for a customer $x \in X$. The dimension of the matrix $\Sigma_{yX}\Sigma_{XX}^{-1}\Sigma_{Xy}$ will be 1 x 1, so, the computation for $\bar{\sigma_y^2}$ will yield a scalar.

Let μ_y^d and $(\sigma_y^d)^2$ represent the updated forecasts for mean usage values and variance in usage rates for the customer y on day d based on the observed values for telemetry customers. The values for μ_y^d and $(\sigma_y^d)^2$ are computed daily from the expressions for $\bar{\mu_y}$ and $\bar{\sigma_y^2}$ provided in Equations 2.14 and 2.15. Since the updated usage mean and variance for the correlated customer are dependent on the observed values, the updates will happen at the end of each day d when daily telemetry readings are received.

A correlated customer should receive a delivery when the updated estimate of the tank level indicates that a customer needs a delivery to ensure the service level is maintained. A delivery must occur the next day, when:

$$\sum_{d=1}^{t} \mu_y^d + \mu_y + z\sqrt{\sum_{d=1}^{t}(\sigma_y^d)^2 + \sigma_y^2} \geq I_y^t \tag{2.16}$$

where I_y^t denotes the tank level at end of day t for the customer y. Like with other customers, this may create challenges in delivery timing. Hence we will add an extra day for delivery flexibility. Now we should deliver by day $t+1$ when:

$$\sum_{d=1}^{t} \mu_y^d + 2\mu_y + z\sqrt{\sum_{d=1}^{t}(\sigma_y^d)^2 + 2\sigma_y^2} \geq I_y^t. \tag{2.17}$$

Equation 2.17 corresponds to the scenario where delivery deadline $D_y = 1$. Similar to telemetry customers, the delivery deadline D_y for the customer y will be equal to

w, $\forall w \geq 1$ if w is the lowest integer where the following holds:

$$\sum_{i=1}^{t} \mu_y^i + (w+1)\mu_y + z\sqrt{\sum_{i=1}^{t} (\sigma_y^i)^2 + (w+1)\sigma_y^2} \geq I_y^t. \tag{2.18}$$

As with telemetry customers, the first delivery deadline D_i^0 is determined based on Equations 2.18 by replacing I_y^t with I_y^0, assuming that we start with tank at I_y^0 at the beginning of the planning horizon.

The expected maximum time between deliveries $\bar{f}$ can be computed based on Equation 2.17 by substituting I_y^{max} for I_y^t. It will help us understand the potential impact of correlation on average savings in routing costs. Because the observations from telemetry units are independent across days, the sum $\sum_{d=1}^{t} \mu_y^d$ reduces to addition of t independent μ_y^d terms. The law of large numbers states that the sample mean of n observations i.e. $\frac{\sum_{i=1}^{n} a_i^n}{n}$ approaches the mean μ_x as n approaches ∞. Hence for sufficiently large t measured in days, $\sum_{d=1}^{t}(a_i - \mu_x) \rightarrow 0$. Similarly $\sum_{d=1}^{y}(\sigma_y^d)^2$ also reduces to sum of t independent $(\sigma_y^d)^2$ terms. Thus, the value for $\bar{f}$ is determined by:

$$(\bar{f}+2)\mu_y + z\sqrt{\bar{f}(\sigma_y^i)^2 + 2\sigma_y^2} = I_y^{max}. \tag{2.19}$$

2.5 Solution Methodology

The VRP is a subproblem of the CTLP. Since the VRP is an NP-Hard problem ([29]), the CTLP is also NP-Hard, so finding an exact solution to the proposed model is computationally intractable. We propose a heuristic to address this new problem. The heuristic starts with an initial assignment of customers to telemetry, estimates the cost of routes corresponding to this assignment across all realizations, and iterates to a new assignment. We propose a way to create the initial assignment of customers to telemetry and different choices for iterating through the assignments of customers to telemetry. We will examine the impact of these choices not only on solution quality but also on run time in Section 2.6. Next, we describe how we evaluate the routing costs associated with a particular assignment across the planning horizon. Note that every assignment of customers to telemetry units will influence the set of customers correlated with telemetry customers as well as the set of customers who are not correlated with telemetry customers.

2.5.1 Rolling Horizon Framework

First, we will discuss the rolling horizon approach used to estimate the routing cost associated with a particular assignment of telemetry units for each realization. We use a rolling horizon since we want to allow customers to receive deliveries slightly earlier than necessary if it creates good routes. Thus, we make routing decisions for several days (p) at once and repeat this process over the planning horizon. If we use a strict p period plan, the customers having a deadline of $p+1$ days will be served on the first day of the next p days and have no delivery flexibility. We would want to avoid such instances for good route planning, so that the deliveries are spread evenly across the planning horizon. A customer with a delivery deadline of $p+1$ days might also combine well with customers having deadlines on an earlier day which helps to reduce the overall routing costs. Thus, we use a planning period of $p+f$ days. Even though our algorithm makes a routing plan for $p+f$ days, it only uses the first p days of the plan before creating a new plan. Here f is a parameter that specifies the additional days of delivery flexibility in the rolling horizon.

For each realization, we use equations from Section 2.4 to identify the customers whose delivery deadlines fall within the upcoming $p + f$ days. For route planning, we will only consider these customers and allow them to receive a delivery up to p days early. For example, on Sunday a business may make a route plan for $5(p) + 3(f)$ delivery days. This means customers whose $D_i \leq 8$ will be routed on day $max(1, D_i - 4)$ to day D_i of the 8 day plan.

In executing the p day plan, each night the feasibility of the next day's plan is checked based on simulated readings for the realization. If this dictates that a change needs to be made in order to satisfy the service level, a change is made. After p days, the process is repeated, until the end of horizon T is reached.

2.5.2 Routing Cost Algorithm

We use the well known Variable Neighborhood Search (VNS) algorithm presented in [28] to generate the routes for each $p+f$ day period. The order for insertion of customers for route planning is determined by the delivery deadlines. We first insert the customer with the earliest deadline. This customer is most important to serve well. Each customer is inserted in the cheapest point in the routes over the $p+f$ days. Once the routes are constructed, two improvement routines are run. These routines are focused on improvements within a single day or a single vehicle's tour. We do not allow the transfer of any customers across different days of the planning horizon. These limitations are made to reduce run time since we need to compute a routing cost across the whole planning horizon for each realization for each assignment of customers to telemetry. These improvement ideas were carefully chosen, after a series of computational experiments, to offer maximum improvement in a limited amount of run time.

First, we run a one-shift improvement routine that improves the routes on each day of the planning horizon as in [30]. Customers are potentially swapped to different locations on the same tour or to other tours on the same day. All possible moves are considered for each customer being served on a given day, and the move that creates the best improvement is selected. This process is iterated until no more improving moves are found on a given day, and then the improvement process moves to the next day.

Second, we run the standard 2-OPT improvement heuristic ([30]) to improve individual routes on each day. The best improving move is chosen among all the candidates on each route, and the process is repeated until no more improvements are found. This process is repeated for each route on each day of the planning period. Once the improvement process is completed, we now have the total routing cost associated with a potential assignment of customers to telemetry over the $p + f$ days for each realization. For each realization, we keep a running sum of the values for the first p days over the whole horizon. We report the average routing cost across all realizations over the whole horizon as the cost associated with a particular assignment of customers to telemetry.

2.5.3 Initialization of Telemetry Set

Next, we will present a method to initialize the set of customers that have telemetry that has a big impact on the run time. We use the f_i' and $\hat{f}_i$ values discussed in Section 2.4 to estimate the impact of telemetry on each customer. We first use the VNS heuristic from Section 2.5.1 and Section 2.5.2 with no customers assigned to telemetry. From this solution, we compute C_i^r values for each customer for each realization r, where C_i^r reflects the total cost to serve customer i without telemetry over the planning horizon. The total cost C_i^r for a customer is computed as follows. Each time that customer i is served, we compute the insertion cost based on its location relative to preceding and succeeding customers. The insertion cost represents the extra cost required to add i to an existing route. These insertion costs are summed over the planning horizon to create the total insertion cost C_i^r which is a reflection of the cost to serve customer i without telemetry in realization r. We average the C_i^r values over the realizations to create a value $C_i^{'}$. This value represents the current cost to serve customer i on average over the R realizations.

The expected reduction E_i in routing costs by telemetry installation for a customer i is useful for determining the value of placing telemetry at customer i. The following expression captures the estimated reduction in routing costs:

$$E_i = \frac{C_i^{'}}{\frac{T}{f_i^{'}}}(\frac{T}{f_i^{'}} - \frac{T}{\hat{f}_i}). \tag{2.20}$$

The average cost per visit is captured by the first term in Equation 2.20, and the expected reduction in number of visits is given by the second term. Thus, Equation 2.20 captures the estimated reduction in cost associated with adopting telemetry at customer i. We can evaluate Equation 2.20 for each customer and choose the M customers that have the maximum value as our initial set of telemetry customers. This approach has been shown to greatly reduce run times as compared with starting with a random initial allocation.

2.5.4 Telemetry Set Improvement

Next, we consider strategic ways to iteratively improve the telemetry set. In our local search, we replace one telemetry customer with one non-telemetry customer at a time, evaluating the cost of this new assignment using the procedure described in Section 2.5.2 and keeping the new assignment if an improvement in cost is found. Computational experiments show that a first-improving search is much faster than a best improving search while achieving similar quality solutions.

Once we have an initial set of telemetry customers, the expected gain by adding a telemetry unit to a customer varies depending on one of the three possible scenarios. First, the customer i is not correlated with any telemetry customers. This expected gain is similar to the initialization phase discussed in Section 2.5.3. Hence, Equation 2.20 is used to compute the expected gain from the adoption of a telemetry unit.

The second case is when customer i is the first representative of a correlated group of customers g_k in the telemetry set. In such a case, we should also consider the contribution from the expected savings in routing costs from the correlated customers. Therefore, we modify Equation 2.20 to:

$$E_i = \frac{C_i'}{\frac{T}{f_i'}}\left(\frac{T}{f_i'} - \frac{T}{\hat{f}_i}\right) + \sum_{j \in g_k} \frac{C_j'}{\frac{T}{f_j'}}\left(\frac{T}{f_j'} - \frac{T}{\bar{f}_j}\right). \tag{2.21}$$

The second term in the right hand side captures the contribution in savings due to the correlated customers.

Third, customer i is correlated with a customer in the telemetry set. Here, customer i already has a representative from its own correlated group g_k in the telemetry set. In such a case, we should not consider the additional contribution in expected savings in routing costs from the correlated customers. So, we modify Equation 2.20 as:

$$E_i = \frac{C'_i}{\frac{T}{f'_i}}\left(\frac{T}{\bar{f}_i} - \frac{T}{\hat{f}_i}\right). \tag{2.22}$$

This equation says that the reduction in routing costs is because of the transition of customer i from the category of a correlated customer to a telemetry customer.

We use these estimates to identify an ordering of how we consider alternative solutions. We want to order non-telemetry and correlated customers so that we consider adding customers who are likely to create more cost improvements with the addition of telemetry. For non-telemetry and correlated customers, the higher values of expected gain correspond to higher cost improvement from adding telemetry. Thus, we consider non-telemetry and correlated customers from highest to lowest values of expected gain in our one-exchange. For the telemetry customers, we want to remove customers that are not creating much actual savings as a result of telemetry. For telemetry customers, we can easily adapt the above calculations to yield the expected cost of removing telemetry units. For telemetry customers, the higher the value of the expected gain is, the higher cost increases are expected from removing telemetry. Thus, we sort this list from low to high, so the customers with the lowest value are considered first for removal of telemetry. This ordering is used in our local search algorithm. Specifically, we remove a customer from the telemetry set one at a time in increasing order of expected cost. Each time a customer is removed from the telemetry set, we consider swapping in a non-telemetry or correlated customer in decreasing order of expected gain. Each time an improvement is found, the telemetry set is changed, and the new solution is used to recompute the E_i values.

2.6 Computational Experiments

Computational experiments were carried out to assess the impact of the proposed solution strategy. We also wanted to evaluate the impact of different problem characteristics on the solutions, including the level of correlation, the number of telemetry units, and size of correlated groups.

The algorithms were implemented and executed in MATLAB Version 7.8.0.347. The experiments were performed on a 2.40 GHz Intel Core 2 Quad processor running the LINUX operating system [4].

Our computational analysis uses eight datasets for the symmetric traveling salesman problem from TSPLIB [7]. These include dantzig42, st70, eil101, tsp225, gil262, rat783, dsj1000 and pr1002. The number associated with the name is the number of customers in the dataset. For example, eil101 has 101 customers. We chose these datasets because they represent a nice variety in terms of both the numbers of customers and how the customers are distributed across the geographical area. The distances between customers for four datasets (dantzig42, st70, and gil262) were computed directly from the location data in miles with no adjustments, while for four others (tsp225, rat783, dsj1000 and pr1002), the distances were scaled so that no two customers were more than eight hours travel time apart with the use of a travel speed of 45 miles per hour. The depot was fixed at the mean latitude and mean longitude position for each dataset.

In our experiments, we examine the average costs across 100 usage realizations over a period of 12 weeks. This time period is chosen so that every customer will have at least one delivery given our choice of usage rates. The number of vehicles for each experiment is the minimum number required to create a feasible routing solution with no telemetry units. This minimum is largely driven by the restriction that routes are no more than eight hours long. We set the fixed stop time to 2.67 minutes and a variable delivery time to .04 minutes per unit. We assume usage and deliveries occur five days per week. We assume all trucks have a capacity of 6000 units and use a service level of 95%. For the base case, we assumed a tank capacity of 350 units at every customer. The mean usage rate μ_i for each customer is randomly assigned a value between 10 units per day and 30 units per day. We fix the variance in usage rate as a fixed proportion of the mean usage rate such that $\sigma = 0.25\mu$. These values are based on real values at $\mathrm{NuCO_2}$. We use a random initial tank level for each telemetry customer for each realization like [14] and [10]. Half of the customers are chosen at random to be divided into three groups of equal size. Each pair of customers belonging to the same group is correlated with a uniform level of correlation ρ. The rest of the customers are uncorrelated. In the base case, $\rho = 0.5$, $p = 5$, and $f = 3$. For all datasets, we solve for the best solution with 5, 10, and 20 telemetry units.

2.6.1 Evaluation of Estimates

The estimates described in Section 2.5.4 are used for the initialization and improvement of the telemetry set. The reason for developing these estimates is to yield a faster convergence than a random ordering of customers both in and out of the telemetry set. In order to evaluate the efficacy of the estimate list using a first-improving routine, we compare the solution quality and solution time with a best improving routine. The best improving routine chooses the best exchange move in the neighborhood and thus ordering does not matter. We also compare with first improving routine based on an arbitrary customer ordering. Table 2.1 summarizes these results for all datasets. The columns "Dataset", "Tele", "Veh", "Soln 1", "Time 1", "Soln 2", "Time 2", "Soln 3", "Time 3" and "% Diff" represent the name of the dataset, number of telemetry units used, number of vehicles used, solution and run time using best improvement, solution and run time improvement using first improvement with arbitrary customer ordering, solution and run time improvement using first improvement with our estimates, and percent improvement of run time with first improvement using estimates relative to best improvement.

Looking at Table 2.1, we notice that, with the choice of a first improving routine, use of our estimates outperforms the arbitrary customer ordering both in terms of solution quality and run time. We also observe that our estimates lead to the same solution quality as the best improving routine in all of our tests. However, the solution time decreases. The percentage reduction in run time ranges from 2.93% to 14.16%. This is because unnecessary one exchange swaps are assigned lower priority using the estimates. Since the estimate list provides the same solution quality with a significant decrease in computational time, we use it in the rest of the tests in this section.

In terms of other observable trends, we notice that the run time increases as the size of the dataset increases. This happens because the complexity of the VRP subproblem increases. This is depicted by solution times for dantzig42 compared to dsj1000. The solution times can also depend on the distribution of customers. However, large improvements found in run time do not seem to follow a pattern in terms of the different datasets or the number of telemetry units.

To reduce the run time for larger sets of customers, we reduce the candidate set of customers that are considered for possible telemetry placement to those with the top 50% (for 200-500 customers) or 25% (for more than 500 customers) values for expected gain from the adoption of telemetry.

2.6.2 Solution Characteristics

Next, we will examine the impact of considering correlation on routing costs. For this purpose, we report the routing cost with and without capitalizing on correlation in the generation of delivery deadlines. We will compare it with the routing cost corresponding to the case where we consider all three categories of customers. The results are shown in Table 2.2. The columns "Dataset" and "Tele" represent the dataset name and the number of telemetry units used. The columns "Soln 1-1", "Soln 2-1" and "Soln 3-1" represent the routing cost without capitalizing on correlation with $\rho = 0.3$, 0.5 and 0.7. The columns "Soln 1-2", "Soln 2-2" and "Soln 3-2" represent the routing cost for different levels of correlation. The columns "% Imp 1","% Imp 2" and "% Imp 3" represent the percent improvement in "Soln 1-2" relative to "Soln 1-1", "Soln 2-2" relative to "Soln 2-1", and "Soln 3-2" relative to "Soln 3-1".

Table 2.1: Solution Comparison

Dataset	Tele	Veh	Soln 1	Time 1	Soln 2	Time 2	Soln 3	Time 3	% Diff
dantzig42	0	1	9011.36		9011.36		9011.36		
	5		8865.27	3279.27	8865.27	3211.99	8865.27	3183.27	2.93
	10		8639.68	6512.93	8639.68	6002.21	8639.68	5852.11	10.15
	20		8401.96	16573.18	8401.96	15273.36	8401.96	14683.28	11.40
st70	0	1	8129.88		8129.88		8129.88		
	5		8007.59	4936.83	8007.59	4518.27	8007.59	4237.90	14.16
	10		7913.90	11092.48	7913.90	10255.02	7913.9	9919.02	10.58
	20		7746.23	24153.59	7746.23	22485.75	7746.23	21183.75	12.30
eil101	0	1	6812.29		6812.29		6812.29		
	5		6519.85	7089.93	6692.15	6403.28	6519.85	6245.82	11.91
	10		6337.93	15892.12	6558.03	15007.38	6337.93	14158.28	10.91
	20		6125.29	26828.75	6329.98	25752.85	6125.29	23112.63	13.85
tsp225	0	6	68834.70		68834.70		68834.70		
	5		67915.38	13875.84	68382.63	12996.84	67915.38	12586.18	9.29
	10		66795.11	20542.11	66997.11	19564.61	66795.11	18726.17	8.84
	20		65516.24	35729.27	65972.12	33836.75	65516.24	31676.93	11.34
gil262	0	4	43003.62		43003.62		43003.62		
	5		41778.65	9286.38	42383.02	8837.57	41778.65	8598.23	7.41
	10		41003.56	26156.49	41723.82	24387.86	41003.56	23652.28	9.57
	20		39102.48	54937.54	39767.53	52866.15	39102.48	51380.01	6.48
rat783	0	6	55819.02		55819.02		55819.02		
	5		54756.16	23739.65	54832.03	23086.96	54756.16	21726.18	8.48
	10		53582.71	38155.27	53997.18	37296.28	53582.71	35893.28	5.93
	20		50163.29	66286.11	50825.92	64982.37	50163.29	60155.24	9.25
dsj1000	0	8	105517.6		105517.56		105517.56		
	5		104973.40	28344.10	105087.16	26980.56	104973.39	25820.12	8.90
	10		102238.20	39268.57	104727.91	38265.05	102238.15	37211.77	5.24
	20		100919.20	77231.03	101863.26	72783.29	100919.17	69827.64	9.59
pr1002	0	8	86734.38		86734.38		86734.38		
	5		85147.12	25968.50	85982.25	25726.18	85147.12	23773.21	8.45
	10		84685.27	40626.12	85395.65	38317.82	84685.27	36817.36	9.38
	20		82957.21	70255.11	83107.16	65892.55	82957.21	63927.54	9.01

We can observe from Table 2.2 that the impact of capitalizing on correlation translates to savings in routing costs. The percentage savings range from 0.42 % to 10.95 %. Also, the savings increase as the level of correlation increases. However, there is no observable trend as we increase the number of telemetry units. For bigger datasets and higher level of correlation, the savings of dsj1000 is bigger compared to pr1002. This is because the distribution of customers in dsj1000 is non-uniform with clusters situated far from the depot. Here the placement of telemetry units translates to savings of bigger magnitude. On the other hand, the customers in pr1002 are distributed uniformly.

Next, we will further examine the impact of different levels of correlation over different number of telemetry units. Savings in routing costs are measured relative to the case without telemetry units. The units of percentage savings provide a fair comparison between varying levels of correlation for the same dataset. Table 2.3 summarizes the results. The columns “Dataset”, “Tele”, “Soln 1” and“% Imp 1” represent the name of the dataset, number of telemetry units used, solution using ρ = 0.3, and percent improvement in solution relative to the correlated case without telemetry units. The columns with suffix 2 and 3 represent the similar quantities for $\rho = 0.5$ and $\rho = 0.7$ respectively. The column “Rel 2vs1” represents the difference between “% Imp 2” and “% Imp 1”. Similarly, the column “Rel 3vs2” represents the difference between “% Imp 3” and “% Imp 2”.

Table 2.2: Impact of Capitalizing on Correlation on Routing Costs

Dataset	Tele	Soln 1-1	Soln 1-2	% Imp 1	Soln 2-1	Soln 2-2	% Imp 2	Soln 3-1	Soln 3-2	% Imp 3
dantzig42	0	9267.26			9195.25			9035.27		
	5	9078.55	8917.73	1.77	8932.17	8765.27	1.87	8856.39	8590.27	3.00
	10	8852.17	8786.22	0.75	8715.21	8639.68	0.87	8692.14	8215.18	5.49
	20	8625.93	8537.12	1.03	8509.23	8401.96	1.26	8524.04	8169.39	4.16
st70	0	8286.04			8203.89			8100.62		
	5	8230.63	8135.19	1.16	8140.57	8007.59	1.63	7957.19	7857.92	1.25
	10	8145.17	8056.91	1.08	8064.82	7913.90	1.87	7901.87	7525.19	4.77
	20	7970.73	7896.02	0.94	7902.98	7746.23	1.98	7873.15	7312.03	7.13
eil101	0	7246.64			7159.03			6902.28		
	5	7052.15	6982.29	0.99	6953.24	6519.85	6.23	6855.40	6215.55	9.33
	10	6759.17	6641.10	1.75	6698.27	6237.93	6.87	6639.27	5912.03	10.95
	20	6499.75	6383.05	1.80	6395.16	6025.29	5.78	6392.31	5752.64	10.01
tsp225	0	70578.81			69884.39			69014.29		
	5	70114.17	69002.75	1.59	69501.75	67915.38	2.28	68682.38	66326.38	3.43
	10	69572.83	68135.38	2.07	69118.27	66795.11	3.36	68152.65	65007.34	4.62
	20	67984.36	67283.87	1.03	68376.38	65516.24	4.18	67732.35	63855.14	5.72
gil262	0	44656.03			43957.27			43484.95		
	5	43119.05	42588.64	1.23	42991.88	41778.65	2.82	42736.27	40009.37	6.38
	10	42581.89	41835.30	1.75	42602.93	41003.56	3.75	42395.92	38956.46	8.11
	20	40816.92	39937.06	2.16	41387.32	39102.48	5.52	41557.16	37412.11	9.97
rat783	0	57296.73			56976.34			56102.86		
	5	56850.60	55107.29	3.07	55672.19	54256.16	2.54	55006.28	52036.82	5.40
	10	55125.23	54682.02	0.80	54603.17	53582.71	1.87	54267.54	51261.73	5.54
	20	52873.83	51940.30	1.77	52736.28	50163.29	4.88	52577.13	48149.46	8.42
dsj1000	0	108582.30			106862.36			105927.73		
	5	106196.30	105702.40	0.47	104936.42	103503.40	1.37	104586.47	97378.98	6.89
	10	104972.70	104289.40	0.65	103623.21	102238.20	1.34	103515.04	96106.90	7.16
	20	101365.00	100936.60	0.42	102411.09	98900.17	3.43	102507.52	92137.94	10.12
pr1002	0	88636.42			87635.11			87010.25		
	5	86983.76	86189.23	0.91	86384.35	85147.12	1.43	86002.37	83595.26	2.80
	10	86110.25	85719.06	0.45	85725.24	84685.27	1.21	85438.92	83028.02	2.82
	20	85472.14	84722.15	0.88	85296.79	82957.21	2.74	83884.05	81292.11	3.09

Table 2.3: Effect of Correlation on Routing Costs

Dataset	Tele	Soln 1	% Imp 1	Soln 2	% Imp 2	Soln 3	% Imp 3	Rel 2vs1	Rel 3vs2
dantzig42	0	9156.24		9011.36		8837.21			
	5	8917.73	2.60	8765.27	2.73	8590.27	2.79	0.13	0.06
	10	8786.22	4.04	8639.68	4.12	8215.18	7.04	0.08	2.91
	20	8537.12	6.76	8401.96	6.76	8169.39	7.56	0.00	0.79
st70	0	8199.06		8129.88		8005.78			
	5	8135.19	0.78	8007.59	1.50	7857.92	1.85	0.73	0.34
	10	8056.91	1.73	7913.90	2.66	7525.19	6.00	0.92	3.35
	20	7896.02	3.70	7746.23	4.72	7312.03	8.67	1.02	3.95
eil101	0	7157.82		6812.29		6528.36			
	5	6982.29	2.45	6519.85	4.29	6215.55	4.79	1.84	0.50
	10	6641.10	7.22	6237.93	8.43	5912.03	9.44	1.21	1.01
	20	6383.05	10.82	6025.29	11.55	5752.64	11.88	0.73	0.33
tsp225	0	69752.15		68834.70		67278.92			
	5	69002.75	1.07	67915.38	1.34	66326.38	1.42	0.26	0.08
	10	68135.38	2.32	66795.11	2.96	65007.34	3.38	0.65	0.41
	20	67283.87	3.54	65516.24	4.82	63855.14	5.09	1.28	0.27
gil262	0	43782.29		43003.62		41186.28			
	5	42588.64	2.73	41778.65	2.85	40009.37	2.86	0.12	0.01
	10	41835.30	4.45	41003.56	4.65	38956.46	5.41	0.20	0.76
	20	39937.06	8.78	39102.48	9.07	37412.11	9.16	0.29	0.09
rat783	0	56372.05		55819.02		53738.26			
	5	55107.29	2.24	54256.16	2.80	52036.82	3.17	0.56	0.37
	10	54682.02	3.00	53582.71	4.01	51261.73	4.61	1.01	0.60
	20	51940.30	7.86	50163.29	10.13	48149.46	10.40	2.27	0.27
dsj1000	0	107571.7		105517.6		99378.39			
	5	105702.4	1.74	103503.40	1.91	98556.18	2.01	0.17	0.10
	10	104289.4	3.05	102238.20	3.11	96227.29	3.29	0.06	0.18
	20	100936.6	6.17	98900.17	6.27	94501.35	7.29	0.10	1.02
pr1002	0	87756.81		86734.38		85184.19			
	5	86189.23	1.79	85147.12	1.83	83595.26	1.87	0.04	0.04
	10	85719.06	2.32	84685.27	2.36	83028.02	2.53	0.04	0.17
	20	84722.15	3.46	82957.21	4.35	81292.11	4.57	0.90	0.21

One might expect that for a given number of telemetry units we would see a larger percentage savings on the smaller datasets since the number of customers impacted by correlation would make up a larger proportion of the dataset. This is somewhat true. For instance, with 20 telemetry units, we see the largest percentage improvement, 11.88 %, in one of the smaller datasets, eil101. However, we also achieve larger savings with bigger datasets, for example, 10.13 % with rat783 dataset. This shows that it is hard to quantify the impact of both telemetry units and correlation on routing cost without using a complete model, such as the one presented here.

For all datasets and all levels of correlation, we observe a steady improvement in the percentage savings with increasing number of telemetry units. For most datasets, the savings with 10 telemetry units is roughly double the savings with 5 telemetry units. The notable exception are the smaller datasets, where the effect is more pronounced. For example, the increase in savings of dantzig42 from 2.79 % with 5 units to 7.04 % with 10 units and $\rho = 0.7$ is remarkable. The higher savings are attributed to the larger proportion of customers impacted by the high correlation.

We also see a steady improvement in the percentage savings from telemetry with increasing level of correlation. The changes are especially significant in smaller datasets, for example, the increase from 3.70 % to 8.67 % with an increase of ρ from 0.3 to 0.7 for st70 and 20 telemetry units. The effect of correlation becomes more prominent as the level of correlation increases. The percentage savings in routing cost show a definite increase with an increase of ρ but there is definitely less improvement with the largest two datasets.

One more observation that is interesting is that in some datasets, the benefits of a higher correlation level outweigh the benefits attained by installing additional telemetry units. For example, in rat783, 3.00 % savings in routing costs is attained with 10 telemetry units for $\rho = 0.3$. However, for the same dataset, we can achieve slightly more savings of 3.17 % with only 5 telemetry units for $\rho = 0.7$. This indicates that knowledge of correlation can be as valuable as 5 telemetry units in this case. This insight is beneficial for small companies with a fixed budget who want to extract the maximum benefits from a limited number of telemetry units.

2.6.3 In-depth Example

We observe from Table 2.3 that all datasets show considerable savings from the adoption of telemetry. We wanted to examine what happens as the number of telemetry units increases further with $\rho = 0.7$ for dsj1000. We compare the percentage savings compared with the case of zero telemetry units. We ran additional experiments using 30 and 50 telemetry units for dsj1000. We can see from Figure 2.1 that in presence of correlation most of the benefit is from 30 telemetry units with only marginal increments after that. There is little value in installing additional telemetry units after the threshold of 30 units, and we conjecture such thresholds exist for all datasets.

We also observe from Table 2.3 that the savings from adoption of telemetry units increases as the level of correlation increases. We wanted to examine this behavior for additional levels of correlation for a particular dataset. We ran additional experiments for dsj1000 for 20 telemetry units with ρ as 0.1, 0.2 , 0.4, 0.6 and 0.8 to

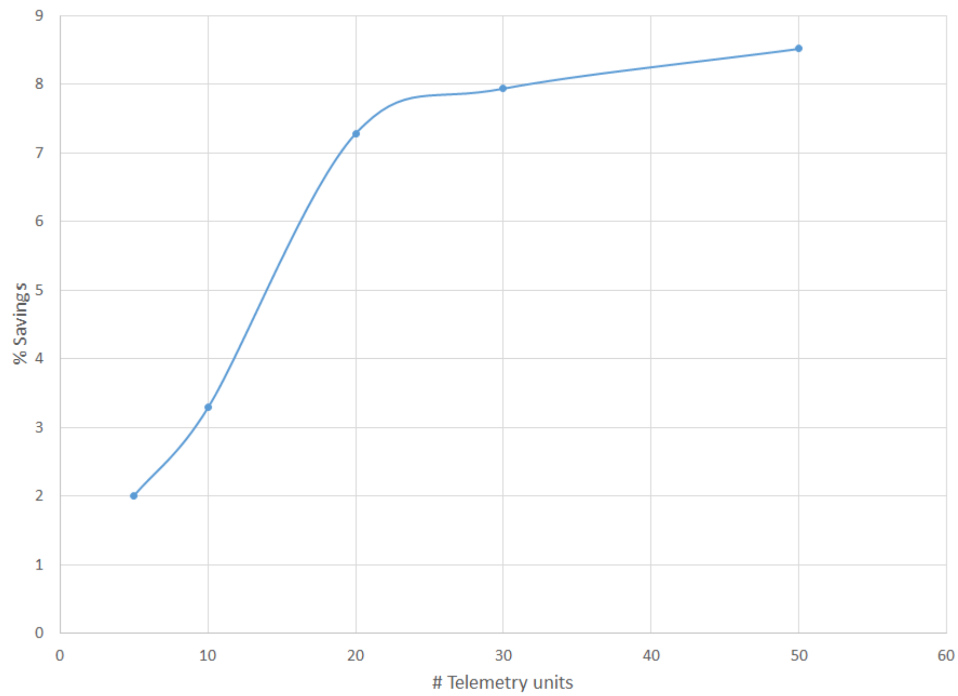

Figure 2.1: Savings vs. Number of Telemetry Units

examine this behavior. As we can observe from Figure 2.2, low levels of correlation do not add much value in the route planning. However, for higher levels of correlation the savings in routing costs are considerable with a careful choice of telemetry units.

Last, we examine the customers that are selected for telemetry in more detail for one dataset. We again analyze the results of dsj1000. Table 2.4 summarizes the distribution of telemetry across the different correlated groups for $\rho = 0.3$ and $\rho = 0.7$. The columns "Tele","Group1", "Group2", "Group3", represent the number of telemetry units, the number of customers in the telemetry set belonging to the first, second and third correlated group. The column "Uncorr" represents the number of customers picked in the telemetry set which are not correlated with any of the

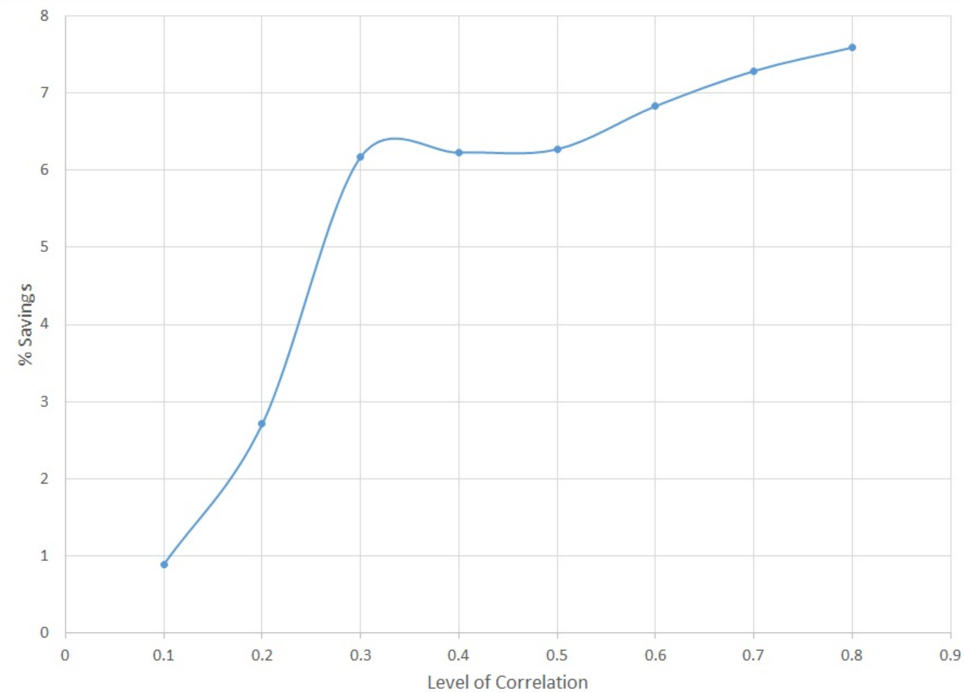

Figure 2.2: Savings vs. Level of Correlation

customers.

Table 2.4: Distribution of Telemetry Customers among Groups

	ρ=0.3				ρ=0.7			
Tele	Group1	Group2	Group3	Uncorr	Group1	Group2	Group3	Uncorr
5	1	0	0	4	2	2	0	1
10	3	2	1	4	5	3	1	1
20	7	3	3	7	9	5	4	2

As we can observe from Table 2.4, the presence of Group 1 members are prominent in the telemetry set. Group 1 consists of more customers with the high estimated gains from telemetry and thus there are higher chances that customers from this group will be picked for telemetry installation. We also notice that as the level of correlation increases the number of uncorrelated customers picked for telemetry installation decreases. In other words, the impact of correlation becomes more pronounced as the level of correlation increases. For example, the number of uncorrelated customers who are picked reduce significantly from 7 to 2 for 20 telemetry units as we increase the level of correlation from 0.3 to 0.7. The two uncorrelated customers that were present in the telemetry set of 20 units for $\rho = 0.7$ were far from the depot and were crucial for forming good routes. In total, 18 out of 20 customers belonged to one of the three correlated groups for $\rho = 0.7$.

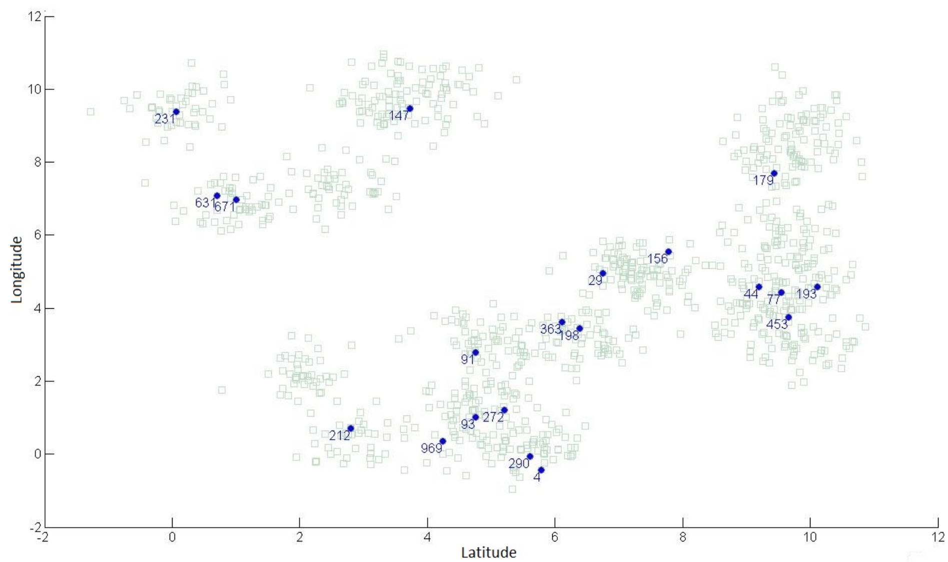

Figure 2.3: Distribution of Telemetry Units, ρ=0.3

Figures 2.3 and 2.4 depict the location of telemetry units for $\rho = 0.3$ and $\rho = 0.7$ respectively. Almost all parts of the graph include customers assigned to

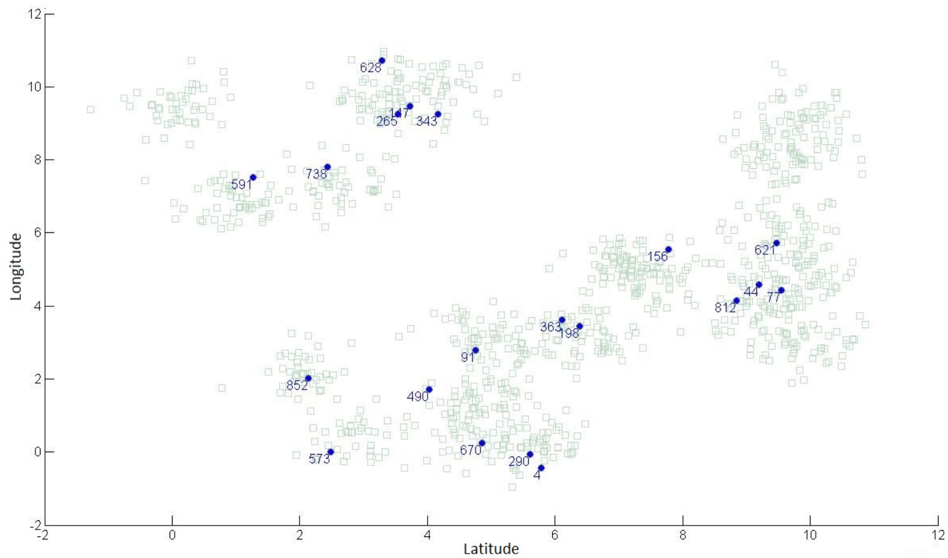

Figure 2.4: Distribution of Telemetry Units, ρ=0.7

telemetry. The customers selected to receive 20 telemetry units with $\rho = 0.3$ are not necessarily included within the set picked with 20 telemetry units with $\rho = 0.7$. Here, 10 customers among the 20 customers selected for telemetry are not identical. As the level of correlation increases, different customers are picked. This indicates that the level of correlation should be considered carefully while route planning with telemetry units, especially for higher levels of correlation.

2.7 Conclusions

Incorporating the impact of correlation can increase the value of a limited number of telemetry units. We have shown savings in routing costs ranging from 0.42 % to 10.95 % over different levels of correlation. We have seen that capitalizing on correlation is especially important when the degree of correlation is high. The choice of initialization and ordering routines for the local search algorithm impacts both solution quality and run time. We present a method which achieves improvement in both these aspects.

CHAPTER 3
LOCOMOTIVE FUEL PLANNING

3.1 Problem definition

The Railway Applications Section (RAS) of INFORMS presented the following problem for their 2010 Problem Solving Competition [2]: "Fuel expenses are a significant part of any railroad's operating costs. Fuel delivery costs differ from location to location because of the differences in distribution, marketing costs and other factors; e.g., as of August 1, 2006, one gallon of diesel costs one of the Class-I railroad company $ 2.2057 in Atlanta, GA, but $ 2.2823 at Augusta, GA. A railroad faces the problem of identifying a cost effective plan to fuel the locomotives that power its trains.

A train schedule defines, among several other attributes, the sequence of yards in which a train stops on its route from origin to destination. It also defines its train-starts or days of operation. If a train operates 3 days per week, say Monday, Wednesday and Friday, then the train-starts for that train at its origination yard are Monday, Wednesday and Friday. This is usually represented by a string notation (M-W-F OR YNYNYNN). The sequence of yards in which a train stops is identical on any train-start. Further, a train-start may run over one or more days, for example, a Monday train-start may run for two days, starting on Monday from its origin and arriving at its destination yard on Tuesday. For this competition it is assumed that all trains are started everyday (train start-MTWTFSS OR YYYYYYY) and that the schedule of each train repeats every week.

Trains are powered by locomotives, and although the number of locomotives might differ from train to train, for this competition it is assumed that each train requires exactly one locomotive and that all locomotives are identical. A locomotive plan assigns engines to train start sequences. This plan dictates the train-starts that each locomotive will power. The plan is setup in such a way that for each locomotive the sequence of train-starts repeats over one or more weeks constituting what is called a locomotive train assignment cycle. The duration of these cycles determine the planning horizon. For this competition we assume that this duration is the same for all locomotives and therefore the planning horizon under consideration corresponds to the duration of any locomotive train assignment cycle.

Locomotives consume fuel at a per mile rate and there is a limit to the maximum amount that can be dispensed into a locomotive determined by the locomotive tank-capacity. Locomotives can only be refueled at yards. For this competition it is assumed that there is only one source of fuel: fueling trucks. Fueling truck contracts are annual, and it is further assumed that currently there are no active contracts. Fueling trucks have a weekly operating cost and a fuel price per gallon that can change from yard to yard. Each truck has a capacity that determines the maximum amount of fuel that it can dispense in one day. A locomotive may be refueled at the train's origin or intermediate yards where the train stops; however, a locomotive may not be refueled at the train's destination yard. Assume instantaneous refueling time, and that a train incurs a fixed cost if it is refueled. Additionally, there is a restriction on the maximum number of times a train can stop to be refueled (excluding the origin).

The railroad needs to identify a fueling plan that minimizes total fuel-related expenses over the planning horizon. This includes truck operating costs as well as fuel purchasing costs. This plan must specify the number of trucks that should be contracted at each yard, and the quantity of fuel that must be dispensed into each locomotive at every yard on every train-start in the locomotive train assignment cycle. The plan must ensure that all locomotives have enough fuel to run all trains according to the schedule (i.e., no locomotive can run out of fuel on its route between yards)."

3.2 Our Approach

The fueling decisions associated with operating a railroad are complex and, as in our test problem, can also be of a large scale. This makes it challenging to find good quality solutions in a relatively short amount of time. After careful thought and experimentation, we decided to pursue a mathematical programming approach because of the solution quality guarantees and the ease of implementation it offers the user.

3.2.1 Notation

We will first describe the notation used throughout this section. Many of the parameters and variables are based on the notation used in [38].

Parameters

N	set of yards
J	set of locomotives
n_j	number of stops for locomotive j,$j \in J$
q_{ijs}	1 if yard i is the s^{th} stop for locomotive j; 0 otherwise, $i \in N, j \in J, s = 1 \ldots n_j$
d	number of days in planning horizon
H_k	set of locomotive and stop combinations falling on horizon day k, $k = 1..d$
C_{jk}	set of stops at intermediate yards in the train sequence falling on day k for locomotive j, $j \in J, k = 1 \ldots d$
U_j	set of unique yards visited by locomotive j, $j \in J$
Z	set of yards that should not be visited
B_{js}	set of stops beginning with stop s on locomotive j and going forward that it is feasible to travel to within a locomotive tank's capacity, $j \in J, s = 1 \ldots n_j$
c_i	fuel price at yard i, $i \in N$
α	fixed refueling cost per stop
F	weekly operating cost per truck
b	tank capacity of a locomotive
r	fuel consumption rate per mile
l_{js}	distance between the stops s and $s+1$ for the locomotive j, $j \in J, s = 1 \ldots n_j$
p	permitted number of stops per train run
t	truck capacity
M_i	total number of visits at yard i by all of the locomotives, $i \in N$

e_{js} the last possible stop belonging to the same yard as s such that the mileage required to traverse from s is within the tank capacity for locomotive j, $j \in J, s = 1 \ldots n_j$

Decision Variables

g_j initial fuel estimate for locomotive j, $j \in J$

w_{js} amount of fuel acquired at stop s for locomotive j, $j \in J, s \in 1 \ldots n_j$

z_i integer number of trucks contracted at yard i, $i \in N$

$$x_{js} = \begin{cases} 1 & \text{if locomotive j uses stop s for refueling, } j \in J, s = 1 \ldots n_j \\ 0 & \text{otherwise} \end{cases}$$

3.2.2 Initial Model

Because of the similarity in problem definition and the quality of the model in [38], our initial integer programming (IP) model contains many of the same constraints as [38]. Key differences are the fact that [38] allows the use of emergency fuel, where our model does not. We also do not use the frequency f_j values found in [38]. In terms of additional constraints, our model must reflect a daily maximum number of stops per locomotive and refueling truck capacity. These additional constraints are noted below.

$$\min 2F \sum_{i \in N} z_i + \sum_{j \in J} \sum_{s=1}^{n_j} \left(\sum_{i \in N} (c_i q_{ijs} w_{js}) + \alpha x_{js}\right) \tag{3.1}$$

subject to

$$w_{js} \leq bx_{js} \; \forall j \in J, \forall s = 1 \ldots n_j \tag{3.2}$$

$$g_j + \sum_{s=1}^{k-1}(w_{js} - rl_{js}) + w_{jk} \leq b \; \forall j \in J \;, \forall k = 1 \ldots n_j \tag{3.3}$$

$$g_j + \sum_{s=1}^{k}(w_{js} - rl_{js}) \geq 0 \; \forall j \in J \;, \forall k = 1 \ldots (n_j - 1) \tag{3.4}$$

$$\sum_{s=1}^{n_j}(w_{js} - rl_{js}) = 0 \; \forall j \in J \tag{3.5}$$

$$\sum_{j \in J} \sum_{s=1}^{n_j}(q_{ijs}x_{js}) \leq M_i z_i \; \forall i \in N \tag{3.6}$$

$$\sum_{(j,s) \in H_k} (q_{ijs}w_{js}) \leq tz_i \; \forall i \in N \;, \forall k = 1 \ldots d \tag{3.7}$$

$$\sum_{s \in C_{jk}} (x_{js}) \leq p \; \forall j \in J \;, \forall k = 1 \ldots d \tag{3.8}$$

$$g_j \geq 0 \; \forall j \in J \tag{3.9}$$

$$w_{js} \geq 0 \; \forall j \in J, \forall s = 1 \ldots n_j \tag{3.10}$$

$$z_i \in \mathbb{Z}^+ \; \forall i \in N \tag{3.11}$$

$$x_{js} \in \{0, 1\} \; \forall j \in J, \forall s = 1 \ldots n_j \tag{3.12}$$

The objective function (3.1) minimizes the sum of truck contracting costs, fuel purchasing costs and stop costs. Note that unlike [38], emergency stop costs and frequency parameters are removed, and the truck contracting cost is modified to reflect a two week contracting period. Constraint (3.2) ensures that a locomotive must stop at a yard to be able to purchase fuel there, like constraint (2b) in [38]. Constraint (3.3) stipulates that the total fuel in the locomotive can never exceed the truck capacity, as constraint (2d) does in [38]. Constraint (3.4) ensures that no locomotive runs out of fuel before arriving at the next stop, like constraint (2e) in [38]. Constraint (3.5) guarantees that the total fuel purchased is equal to the total fuel consumed. The constraint (2f) in [38] uses an inequality for this constraint, but we have modified this constraint to be an equality constraint to help speed up convergence. Constraint (3.6) ensures that a truck must be contracted before any fueling stops can occur at a yard, as in constraint (2g) in [38].

Constraint (3.7) ensures that the total fuel purchased at a yard does not exceed a contracted truck's capacity. Constraint (3.8) limits the total number of stops, excluding the source yard, for a train run. Note that both of these constraints must be enforced for each day in the planning horizon. Finally, constraints (3.9-3.12) define the variables.

3.2.3 Improved Model

Next, we will introduce the constraints that we have found that tighten the above formulation. These constraints increase the value of the linear programming (LP) relaxation of the problem, enabling better bounds to be used in the branch and bound algorithm used to solve the IP. As noted below, some of the constraints here replace constraints in Section 3.2.2, where others are added to the formulation. Experiments in Section 3.3 will demonstrate the improvements that occur as a result of these new constraints.

$$q_{ijs}x_{js} \leq z_i \ \forall i \in N \ , \forall j \in J, \forall s = 1 \ldots n_j \tag{3.13}$$

$$\sum_{s=1}^{n_j} x_{js} \geq \lceil (r/b) \sum_{s=1}^{n_j} (l_{js}) \rceil \ \forall j \in J \tag{3.14}$$

$$\sum_{(j',s')\in B_{js}} x_{j's'} \geq 1 \; \forall j \in J, \forall s = 1 \dots n_j \tag{3.15}$$

$$\sum_{u \in U_j} z_u \geq 1 \; \forall j \in J \tag{3.16}$$

$$z_i = 0 \; \forall i \in Z \tag{3.17}$$

$$x_{j1} + x_{j2} = 1 \; \forall j \in J : \mid U_j \mid = 2 \tag{3.18}$$

$$q_{ijs'} x_{js'} \leq 1 - q_{ijs} x_{js} \forall i \in N, \forall j \in J, \forall s = 1..n_j, \forall s' = s + 1..e_{js} - 1 \tag{3.19}$$

Constraint (3.13) performs the same job as (3.6) but does so in a way that is stronger in terms of the LP because the sum in (3.6) is separated into individual constraints. Thus, (3.13) will replace (3.6) in the revised formulation of the problem. Constraint (3.14) provides a lower bound on the number of stops a locomotive will make. The next type of strengthening constraints are based on the idea that a locomotive must make at least one stop during each block of distance that equates to a full locomotive tank capacity. Constraint (3.15) states that there should be at least one stop in each of the blocks proceeding from each stop. For each locomotive, the first block reflects that the largest amount of fuel at the first stop before any refueling is equal to $b - rl_{jn_j}$. For the later stops in each locomotive's sequence, the B_{js} set will contain stops from the beginning of the train sequence to reflect the repeating nature of a sequence. Constraint (3.16) ensures that a truck should be contracted at least one of the unique yards visited by a locomotive. We can eliminate some yards from consideration and set the associated z_i values to zero. We can eliminate those yards that do not provide the minimum cost fueling option to any of the locomotives, as these will never be selected in the optimal solution. This idea is reflected in (3.17), but in the model file, we modify the set N to be $N - Z$.

In our preliminary computational experiments, we found that many alternate optima exist for this problem, as with many location problems. Here, many different combination of stops for refueling a locomotive can yield the same total cost which makes it difficult for an IP solver to converge in a reasonable amount of time. Thus, the next set of constraints were created to reduce the symmetry and enforce a structure on the solution to reduce the number of nodes in the branch and bound tree. Some of these constraints also strengthen the model and increase the LP value. In the example problem, many locomotives visit only two yards, and many different stop combinations yield the same cost solution. We restrict the first stop to occur during the first visit to one of these two yards in constraint (3.18). This will not create additional stops, as it will simply prevent stops later in the locomotive's sequence. The second idea is to force that if a locomotive stops at a particular yard for refuelling, the locomotive cannot stop at the same yard again if there is another stop at that yard before a locomotive runs out of fuel. Constraint (3.19) eliminates such alternate solutions in which intermediate stops are allowed. The final model consists of constraints (3.1)-(3.5), (3.7)-(3.19).

3.3 Computational Experiments

For our experiments, the Excel file provided in the competition was converted to a file that could be read by a program in Matlab. The Matlab program creates the H_k, C_{jk}, U_j, Z, and B_{js} sets and computes the e_{js} value. It outputs these sets and the other parameters in one data file read by AMPL. The model and data file

Table 3.1: LP Results

Model	LP Value
(1)-(12)	11174178.51
(1)-(5), (7)-(12) + (13)	11214860.66
(1)-(12) + (14)	111204475.60
(1)-(12) + (15)	11220847.56
(1)-(12) + (16)	11258207.11
(1)-(12) + (17)	11174265.31
(1)-(12) + (18)	11187839.31
(1)-(12) + (19)	11174178.51
(1)-(5) + (7)-(19)	11339248.69

Table 3.2: IP Results

Model	Best IP Value	Solution Gap	LP Value
(1)-(12)	11442372.92	1.45%	11276000
(1)-(5) + (7)-(19)	11404774.56	0.15%	11388000

are solved by AMPL with a CPLEX solver (version 9.0). The experiments were run on a Intel Core 2 Quad CPU Q6600 @ 2.40GHz (x4).

Table 3.1 demonstrates the impact of each of the new constraints on the LP relaxation with a combined increase in the LP value of $165,070.18. Table 3.2 demonstrates how well we can solve the IP with our initial model and with modifications. It show that with a run time of 90 minutes, our model yields a solution with an optimality gap of only 0.15% and our improvements yield a cost savings of $37,598.36 from our initial model. Running this same model for 8 hours reduces this gap to 0.04%, representing a savings of $41,393.40 from the initial model.

3.4 Conclusions

Through the addition of constraints that strengthen the original model as well as reduce the symmetry in the solutions, we are able to save $37,598.36 (with an optimality gap of 0.15%) in 90 minutes and save $41,393.40 (with an optimality gap of only .04%) in 8 hours. The use of a math programming model allows our approach to be described in a very concise and readable way, and it avoids much of the overhead involved with implementing many types of approaches. A practitioner can easily change the parameters and/or the run time options and do repeated runs or "what if" types of analysis. The combination of AMPL and CPLEX used here offers a high quality means of solving our model, and if another solver is selected by the practitioner, the changes required are simple syntax changes in the data file generator and model file.

REFERENCES

[1] NuCo2: The only certified beverage grade co2 carbonation system for fountain drinks and beer carbonation, 2010. http://www.nuco2.com/about/index.html [Online; accessed 15-March-2010].

[2] 2010 RAS Competition - INFORMS, 2012. https://www.informs.org/Community/RAS/Problem-Solving-Competition/2010-RAS-Competition [Online; accessed 15-May-2010].

[3] Exemys: Oil and Gas, 2012. http://exemys.com/beta/english/solutions/oil.shtml [Online; accessed 15-March-2012].

[4] Main Page - ISOR Wiki, 2012. http://dollar.biz.uiowa.edu/~sburer/isorwiki/index.php/Main_Page [Online; accessed 15-March-2012].

[5] RCATS Telemetry Price Listing, 2012. http://www.rcatsystems.com/pdf/pricelist.pdf [Online; accessed 15-September-2012].

[6] Think Tank FAQs, 2012. http://www.calor.co.uk/customer-services/faqs/think-tank/ [Online; accessed 15-September-2012].

[7] TSPLIB, 2012. http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/ [Online; accessed 15-March-2012].

[8] Vending Operations: Reaping the Benefits of Telemetry Technology, 2012. http://www.telemetrytalk.com/resources/MEI%20Authored%20Telemetry%20Article%208.12.08%20FINAL.pdf [Online; accessed 15-September-2012].

[9] What is Telemetry?, 2012. http://www.sealtag.org/What%20Is%20Telemetry.html [Online; accessed 15-September-2012].

[10] J. Bard, L. Huang, P. Jaillet, and M. Dror. A decomposition approach to the inventory routing problem with satellite facilities. *Transportation science*, 32(2):189–203, 1998.

[11] J. Belenguer, E. Benavent, C. Prins, C. Prodhon, and C. Wolfler. A branch-and-cut method for the capacitated location-routing problem. *Computers & Operations Research*, 38(6):931–941, 2011.

[12] R. Bent and P. V. Hentenryck. Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52(6):977–987, 2004.

[13] M. J. Bucci, D. P. Warsing, M. G. Kay, R. Uzsoy, and J. R. Wilson. Modeling the effects of demand correlation in location problems that account for inventory pooling. 2006.

[14] A. M. Campbell, L. Clarke, and M. Savelsbergh. Inventory routing in practice. *The vehicle routing problem*, 9:309–330, 2002.

[15] A. M. Campbell and M. Savelsbergh. Inventory routing in practice. In P. Toth and D. Vigo, editors, *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications, 2002.

[16] A. M. Campbell and M. Savelsbergh. Delivery volume optimization. *Transportation Science*, 38:210–223, 2004.

[17] A. M. Campbell and J. H. Wilson. Forty years of periodic vehicle routing. *Networks*, 63(1):2–15, 2014.

[18] C. Chiang. The correlated vehicle routing problem. In *Wireless Communications, Networking and Mobile Computing, 2007. WiCom 2007. International Conference on*, pages 3824–3828. IEEE, 2007.

[19] N. Christofides and J. E. Beasley. The period routing problem. *Networks*, 14:237–256, 1984.

[20] L. C. Coelho, J. Cordeau, and G. Laporte. Thirty years of inventory-routing. Technical report, CIRRELT, 2012.

[21] J. Cordeau, M. Gendreau, and G. Laporte. A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, 30(2):105–119, 1997.

[22] M. Eaton. *Multivariate statistics: a vector space approach*. Wiley New York, 1983.

[23] R. Edelson and J. Krolik. The discrete correlation function-a new method for analyzing unevenly sampled variability data. *The Astrophysical Journal*, 333:646–659, 1988.

[24] G. Eppen. Note-effects of centralization on expected costs in a multi-location newsboy problem. *Management Science*, 25(5):498–501, 1979.

[25] N. Erkip, H. Hausman, and S. Nahmias. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Science*, 36(3):381–392, 1990.

[26] T. D. Gwiazda. *Crossover for single-objective numerical optimization problems*, volume 1. Tomasz Gwiazda, 2006.

[27] V. Hemmelmayr, K. F Doerner, R. F. Hartl, and M. Savelsbergh. Vendor managed inventory for environments with stochastic product usage. *European Journal of Operational Research*, 1:685–695, 2010.

[28] J. Kytöjoki, T. Nuortio, O. Bräysy, and M. Gendreau. An efficient variable neighborhood search heuristic for very large scale vehicle routing problems. *Computers & Operations Research*, 34(9):2743–2757, 2007.

[29] J. K. Lenstra and A. H. G. Rinnooy Kan. Complexity of vehicle routing and scheduling problems. *Networks*, 11:221–227, 1981.

[30] S. Lin and B. W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21:498–516, 1973.

[31] L. Liu and X. Yuan. Coordinated replenishments in inventory systems with correlated demands. *European Journal of Operational Research*, 123(3):490–503, 2000.

[32] F. R. J. López and A. J. López. Temperature telemetry system for agricultural crops using free software. In *Proceedings of the 6th Euro American Conference on Telematics and Information Systems*, Valencia, Spain, 2012.

[33] G. Mbiydzenyuy. *Assessment of Telematic Systems for Road Freight Transport*. PhD thesis, School of Computing, Blekinge Institute of Technology, 2010.

[34] H. Min, V. Jayaraman, and R. Srivastava. Combined location-routing problems: A synthesis and future research directions. *European Journal of Operational Research*, 108(1):1–15, 1998.

[35] G. Nagy and S. Salhi. Location-routing: Issues, models and methods. *European Journal of Operational Research*, 177(2):649–672, 2007.

[36] S. Netessine, G. Dobson, and R. Shumsky. Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research*, 50(2):375–388, 2002.

[37] S. Netessine and N. Rudi. Centralized and competitive inventory models with demand substitution. *Operations Research*, 51(2):329–335, 2003.

[38] S. M. Nourbakhsh and Y. Ouyang. Optimal fueling strategies for locomotive fleets in railroad networks. *Transportation Research Part B: Methodological*, 44:1104–1114, 2010.

[39] J. Ohlmann, M. Fry, and B. Thomas. Route design for lean production systems. *Transportation Science*, 42(3):352–370, 2008.

[40] D. Pacciarelli, A. D'Ariano, and M. Scotto. Optimal location of RFID antennas in a courier express warehouse. In *Proceedings of the 2011 International Conference on Industrial Engineering and Operations Management*, Kuala Lumpur, Malaysia, 2011.

[41] H. N. Psaraftis. Dynamic vehicle routing: Status and prospects. *Annals of Operations Research*, 61:143–164, 1995.

[42] S. Raghunathan. Impact of demand correlation on the value of and incentives for information sharing in a supply chain. *European Journal of Operational Research*, 146(3):634–649, 2003.

[43] R. Shuttleworth, B. L. Golden, S. Smith, and E. Wasil. Advances in meter reading: Heuristic solution of the close enough traveling salesman problem over a street network. In B. Golden, R. Raghavan, and E. Wasil, editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*. Springer, 2008.

[44] E. Silver, D. Pyke, and R. Peterson. *Inventory management and production planning and scheduling*, volume 3. Wiley New York, 1998.

[45] A. Verma and A. M. Campbell. Strategic placement of telemetry to reduce routing costs. *Networks*, 63(3):260–275, 2014.

[46] K. Xu and P. T. Evers. Managing single echelon inventories through demand aggregation and the feasibility of a correlation matrix. *Computers & Operations Research*, 30(2):297–308, 2003.

[47] K. Zhu and U. W. Thonemann. Modeling the benefits of sharing future demand information. *Operations Research*, 52(1):136–147, 2004.