Vendor Managed Inventory vs. Order Based Fulfillment in a Specialty Chemical Company

by

Dimitrios Andritsos
M.S., Industrial Engineering, 2004
Purdue University

B.Sc., Mining Engineering and Metallurgy, 2002
National Technical University of Athens, Greece

and

Anthony Craig
B.Sc. Computer Engineering, 1999
Iowa State University

Submitted to the Engineering Systems Division in Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Logistics

at the

Massachusetts Institute of Technology

June 2006

Signatures of Authors..
Engineering Systems Division
May 12th, 2006

Certified by ....
Stephen C. Graves
Abraham J. Siegel Professor of Management Science and Engineering Systems
Thesis Supervisor

Accepted by ....
Yossi Sheffi
Professor of Civil and Environmental Engineering
Professor of Engineering Systems
Director, MIT Center for Transportation and Logistics

Vendor Managed Inventory vs. Order Based Fulfillment in a Specialty Chemical Company

by

Dimitrios Andritsos

and

Anthony Craig

Submitted to the Engineering Systems Division
on May 12th, 2006 in Partial Fulfillment of the
Requirements for the Degree of Master of Engineering in
Logistics

Abstract

In this thesis, an analysis of the existing order based fulfillment process for one product line of a specialty chemicals manufacturer is made and the potential benefits from the implementation of a vendor managed inventory (VMI) system are quantified. A single facility is considered and our focus is on the possible reductions in transportation effort.

Initially, a set of criteria are defined for classifying which storage tanks will be served under the VMI system and which under the existing order based process. Subsequently, a cluster first – route second approach is implemented, where customer locations are first separated into clusters based on geographical proximity and routes are then designed for each of the clusters. A mathematical model is constructed that aids in the design of delivery routes that minimize the total number of delivery trips. Finally, the total transportation effort that would be required for replenishing the VMI and non-VMI tanks is estimated and a comparison is made with the current system. Key performance indicators are compared between the existing order based fulfillment process and the potential VMI implementation. Limitations of the proposed approach are discussed and directions for future research are highlighted.

Thesis Supervisor: Stephen C. Graves
Title: Abraham J. Siegel Professor of Management Science and Engineering Systems

Acknowledgements

I would like to thank our advisor, Prof. Stephen Graves, for his guidance throughout this work and for his valuable help in developing the mathematical model that is presented in this thesis. In addition, I would like to thank my colleague, Anthony Craig, with whom I have worked together in carrying out this research. Finally, I am thankful to Jose Hernandez and Eric Simeonidis for their help and insightful comments.

Dimitrios Andritsos

I wish to thank Professor Stephen Graves for contributing his time, experience, and advice to the development of this thesis; Dr. Christopher Caplice for his advice and tireless efforts on behalf of the MLOG program; Jose Hernandez for his support and whose previous research provided a basis for much of the work seen here; and Eric Simeonidis for his advice and help. I would like to acknowledge the support of my classmates for sharing their collective wisdom and experiences during the past year. This thesis would not be possible without the help of my partner, Dimitrios Andritsos, and I thank him for his effort in this collaboration. Finally, I would like to thank Emily Sheldon for providing encouragement and a sympathetic ear when I needed it the most.

Anthony Craig

Dedication

I would like to dedicate this thesis to my family. Without their continued support none of this would have been possible. To them, I owe everything that I have achieved in my life. In addition, I wish to dedicate this work to my friends Petros, Simos and Dimitris with whom I have spent endless hours here at MIT and together we have been able to go through most of the difficulties that we faced. I am sure that our friendship will last even though each of us will take a different path from here.

Dimitrios Andritsos

I wish to dedicate this thesis to my parents, John and Mary Lynn. They have supported their children in all our endeavors, including my decision to attend graduate school at this stage of my life. Without their support I would not be in a position to follow my interests and continue my academic career. For their help, and all that they have done for me, I thank them.

Anthony Craig

Table of Contents

List of Tables

List of Figures

1 Introduction

Recent advances in technology have made it possible for Company A to consider Vendor Managed Inventory (VMI) in its specialty chemical business. As customers request this service and Company A looks for ways to reduce costs, it is necessary to determine how VMI could work in this context. This thesis explores a method for determining which customers to serve by VMI and measuring the potential benefits.

1.1 Company profile

Company A manufactures and distributes specialty chemicals in bulk. These products are used as admixtures to their customers' production processes and their role is to enhance product performance and quality. Currently, Company A produces products that are categorized in three product lines, entailing more than 40 distinct product families. In this thesis we consider one product line that consists of 26 product families.

Company A is an international company that serves approximately 30 sales districts throughout the world. Within North America, it currently operates a network of 48 facilities that serve more than 11,000 distinct customer sites. Sales districts within the United States of America have very distinct characteristics. Typically, in the east, sales districts cover medium to small service areas and exhibit high customer density and medium sales volume. In the central area of the United States, sales districts cover large service areas and exhibit both high customer density and sales volume. Finally, in the west and southwest sales districts cover large service areas that exhibit high customer density and sales volume in metropolitan areas, and low customer density and medium sales volume in suburban areas. For the product line under study, 5 dedicated plants and 17 distribution centers exist in North America and serve more than 4,000 customers.

Company A operates a private fleet of compartmentalized tanker trucks for serving the distribution needs of its plants and distribution centers. Currently, approximately 70 tanker trucks are used throughout the United States with each truck having a capacity of 4,000 gallons and the ability to carry up to 5 distinct products in its compartments. Compartments are unequally sized and their capacities range from 40 to 1,500 gallons.

1.2 Current inventory fulfillment practice

Replenishment of products at the customer sites is initiated by customer orders. Customers maintain at their sites tanks of different capacities that are dedicated to specific products. Customers have the ability to order products at their discretion and without any limitation regarding order sizes. It is company A's responsibility to replenish the product tanks at the customers' sites. Once an order is placed, the company usually has approximately a week to fulfill it. In some instances, an order can be placed on an emergency basis, in which case the company attempts to satisfy it within the next working day.

Design of delivery routes is performed manually by a dedicated planning team at company A that is in charge of designing routes for the entire network. On a weekly basis, orders are collected and subsequently routes are designed so that truck utilization is maximized and total miles driven are kept as low as possible. Routes usually include stops at multiple customer locations with the possibility of delivering multiple products at each stop. Movement of product from one customer location to another is very rare and trucks usually return empty at their starting locations.

1.3 Vendor Managed Inventory

Vendor Managed Inventory is the practice under which the vendor is responsible for monitoring product inventories at the customer's sites and for determining when to replenish a product as well as what quantities to deliver. VMI has been implemented in various industries such as the petrochemical industry for the replenishment of gas stations, the grocery industry for the replenishment of products at supermarkets, the soft drink industry for the replenishment of vending machines as well as the automotive industry for the replenishment of service parts (Campbell & Savelsbergh, 2004). Essentially, VMI provides the setting for a partnership between the vendor and its customers that is based on information sharing. The end goal is to match supply and demand as closely as possible (Angulo, Nachtmann & Waller, 2004).

With recent advances in technology, remote monitoring of customers' inventories has become easier and VMI has the potential of being implemented to a wide range of settings. Within the specialty chemicals industry, VMI is becoming increasingly popular and certain customers have expressed a willingness to enter into such agreements with their vendors. Towards this direction, company A is considering the possibility of offering VMI to its customers. For this reason, it has initiated a pilot project of installing tank monitoring telemetry devices at selected customer sites. These devices allow the company to obtain daily readings of tank levels and get an accurate picture of the true consumption rates for each product.

1.3.1 VMI opportunities

Within the context of company A's operations, VMI will improve demand visibility and can potentially eliminate the uncertainty and variability that are induced through the ordering system that currently exists. Large, unexpected orders will no longer appear and company A will have the ability to better serve its customers by anticipating their actual inventory needs based on their true consumption rates. From the customer side, VMI offers the potential for improved satisfaction, reduction in the number of stock – outs that are experienced and possibly higher inventory turns.

VMI offers tremendous opportunities for more effective use of the company's production and distribution resources. Through the improved demand visibility and the subsequent reduction in demand uncertainty, company A will have the opportunity to reduce the amounts of safety stock that are required at the company's plants and distribution centers under the existing ordering system in anticipation of orders with highly variable sizes. In addition, since company A will control the timing of the replenishments, significant opportunities will also exist for grouping customer locations on delivery routes and achieving reductions in transportation costs via the higher utilization of trucks and the reduction in total miles driven via the design of more effective delivery routes.

Furthermore, under a VMI setting, since company A will be able to better control the inventories located at customer sites, it will be able to effectively offer the option of consignment to its customers. Although, various consignment schemes have been implemented in practice, consignment in general can reduce and potentially eliminate inventory holding costs for company A's customers since inventory at the customer's site will still be owned by company A. Therefore, the option of consignment can act as a powerful negotiation tool for reducing the large payment cycles that currently exist. By offering the option of consignment, company A can negotiate with its customers for shorter payment cycles, where a payment cycle is defined as the amount of time that elapses between actual delivery of the product to the customer and receipt of payment by company A. The increases in inventory holding costs for company A will be offset by the improved cash flows to the company since the amount of capital that will be tied up will be reduced.

1.3.2 VMI requirements

In order to offer the option of managing its customers' inventories, company A will need to install telemetry devices that will allow it to monitor product consumption rates on a daily basis. Therefore, significant setup costs are associated with this requirement. In addition, investments in information technology infrastructure will be necessary for effectively collecting and processing the data that will be generated under this setting.

If consignment is offered as part of an arrangement with certain customers, company A will also experience increases in its inventory holding costs since transfer of product ownership will generally happen at a later stage compared to the existing practices. Since company A will maintain ownership of the product for a larger time period, possibly until actual product consumption depending on the consignment deal, it will be charged the associated holding costs.

VMI introduces a new paradigm of operation. Consequently, relationships with customers will need to be set on a new basis and planning processes will need to be redesigned. Furthermore, the costs of implementing and operating VMI are large. Therefore, VMI will not be offered to the entirety of company A's customer base. As a result, the additional administrative burden of dealing with two distinct inventory fulfillment practices will have to be dealt with.

1.4 Objective

The objectives of this thesis are a) to develop a policy for segmenting customer and products in to those served by VMI and those served by order based fulfillment, b) to create a model for determining regular delivery schedules for customers served by VMI, and c) to identify the potential savings by serving these customers through VMI net of any additional costs to serve non-VMI customers.

1.5 Scope

This research is focused on the transportation savings available through VMI. Issues related to the implementation of VMI, including the cost of holding inventory, consignment, and installation of tank monitoring equipment, are considered outside the scope of this thesis. Rather, we assume that it is possible to implement VMI in the manner proposed, and we measure the savings of such an implementation. Previous research with Company A has shown that inventory savings possible through VMI are small compared to the transportation costs (Shen, 2004), and therefore this thesis focuses only on the transportation elements. Further, the scope is limited to a single facility; however, the process developed in this thesis can be applied to each of Company A's locations to determine total company-wide savings.

1.6 Structure of the Thesis

Chapter 2 analyzes the customer, product, order, and delivery data examined in this thesis. Chapter 3 discusses the methodology developed for segmenting customers and products, placing customers into clusters, creating specific delivery routes within clusters, and measuring the total delivery effort. Chapter 4 presents the results of applying this methodology and compares it with the actual delivery effort. Chapter 5 provides the conclusion and presents areas for extension of this research.

2 Data Analysis

2.1 Dataset description

Company A has provided an extensive dataset that was obtained from its Enterprise Resource Planning system. The information that was provided includes the following:

- Classification of service areas into sales organizations, sales regions, sales districts and sales territories.
- Geographical locations of the company's facilities defined in terms of latitude and longitude.
- Geographical locations of the customer sites being served (defined in terms of latitude and longitude) in addition to their classification into regions, districts and territories for sales purposes.
- Individual product characteristics (e.g. product gross and net weights) as well as their classification into distinct product families.
- Number and capacity of storage tanks installed at the various customer sites as well as the products to which these tanks are dedicated.
- Delivery history: Delivery information for three full years has been provided. Each delivery record, identified by a unique delivery number, includes: a) the

sourcing location; b) the customer site being visited; c) and the products and the respective quantities that were delivered on that visit.

- Trip information: Trip information for one full year has been provided. Each trip record, identified by a unique schedule number, includes: a) the total mileage of the trip; b) the total duration of the trip separated into the time spent loading the truck at the facility, the time spent driving and the time spent at the customer sites that were visited on the given trip; c) the total quantity that was delivered as well as any quantity of product that was transferred between customer sites or returned to company A's facility; d) and the number of stops that were made on the trip.

- Delivery history and delivery trips linking information: Based on the above mentioned sources of information, it was not possible to determine which locations were actually visited on a given trip. Therefore, information was provided that enabled us to find out the specific customer sites that were visited on a given trip as well as the specific products and quantities that were delivered at each stop. However, it was not possible to determine the actual sequence with which the individual customer sites were visited on a given trip.

- Tank monitoring telemetry information: As was mentioned previously, company A has initiated a pilot project of installing remote tank monitoring devices at selected tanks at certain customer locations. Readings obtained during the period of the first eight months of operation were provided.

Since delivery history and delivery trip information overlap for a single year, delivery history information for that year only will be used for carrying out the analysis.

2.2 Location description

In order to limit the scope of the thesis, we decided to focus on a single facility located in Southern California. The facility serves a large metropolitan area that exhibits high customer density as well as high sales volumes. Seasonality of demand for the specific facility under study is considered to be mild. In total, 289 customer sites were visited at least once during the year under study and 25 distinct products, belonging to 11 product families, were distributed. The total number of distinct tanks that were replenished during the course of the year was 948.

Using the geographical information available, a map and a plot were created that exhibit the 289 customer locations that were visited during the year. The map was constructed using a tool available in http://www.batchgeocode.com. Complete address information for each customer location was pulled from the available dataset and was used as input to the tool that subsequently processed the information and created the map shown in Figure 1. For constructing the plot that is shown in Figure 2, company A's facility was used as the origin and the customer sites were plotted according to their Euclidean distances from the facility. For computing the Euclidean distances, the relative location of each customer site, indexed as location *i*, to the facility, indexed as location *o*, had to be determined. This was performed by using the following approximate relationship, which assigns an average length of 69 miles per degree of latitude (Simchi – Levi, Kaminsky & Simchi – Levi, 2003):

$(x_{io}, y_{io}) = (69*(longitude_i - longitude_o), 69*(latitude_i - latitude_o))$, *where*

x_{io} = *horizontal distance, in miles, from location o to location i*

y_{io} = *vertical distance, in miles, from location o to location i*

The Euclidean distance from location o to location i would then be given by $\sqrt{x_{io}^2 + y_{io}^2}$.

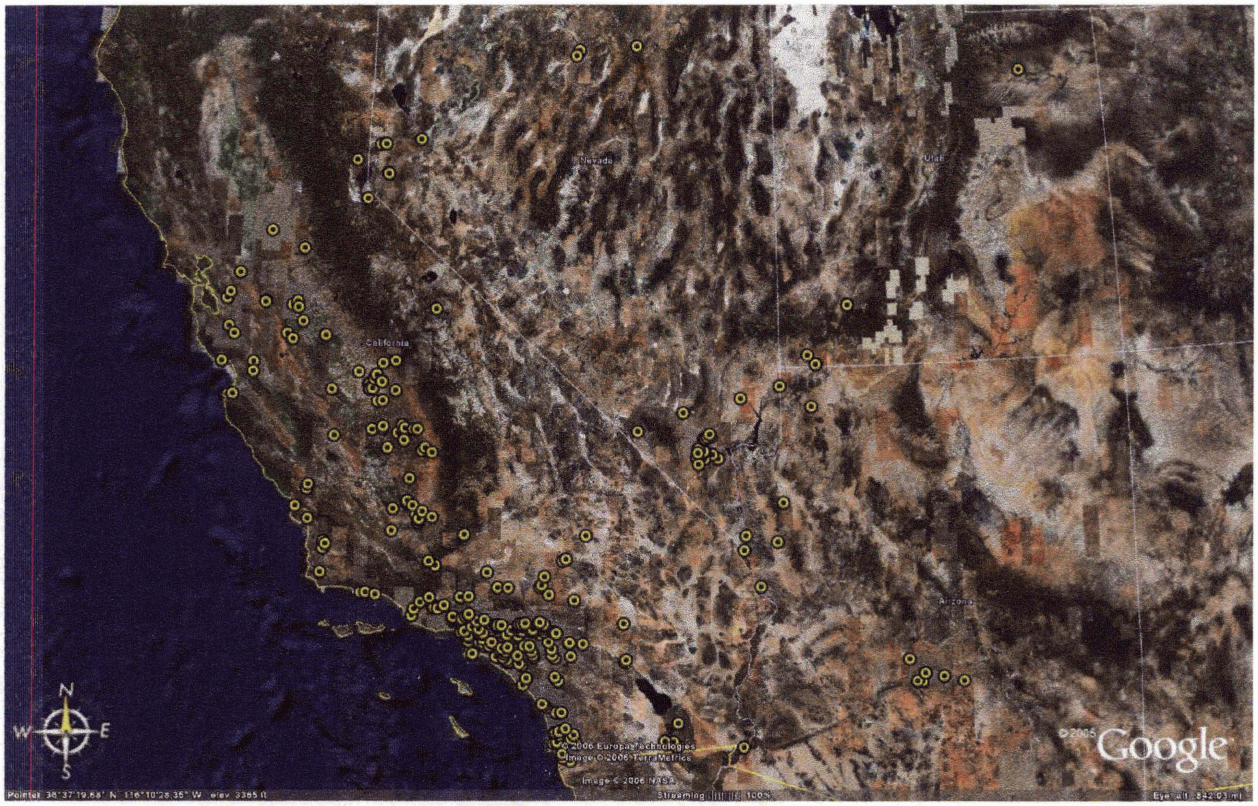

Figure 1: Map of all customer sites visited

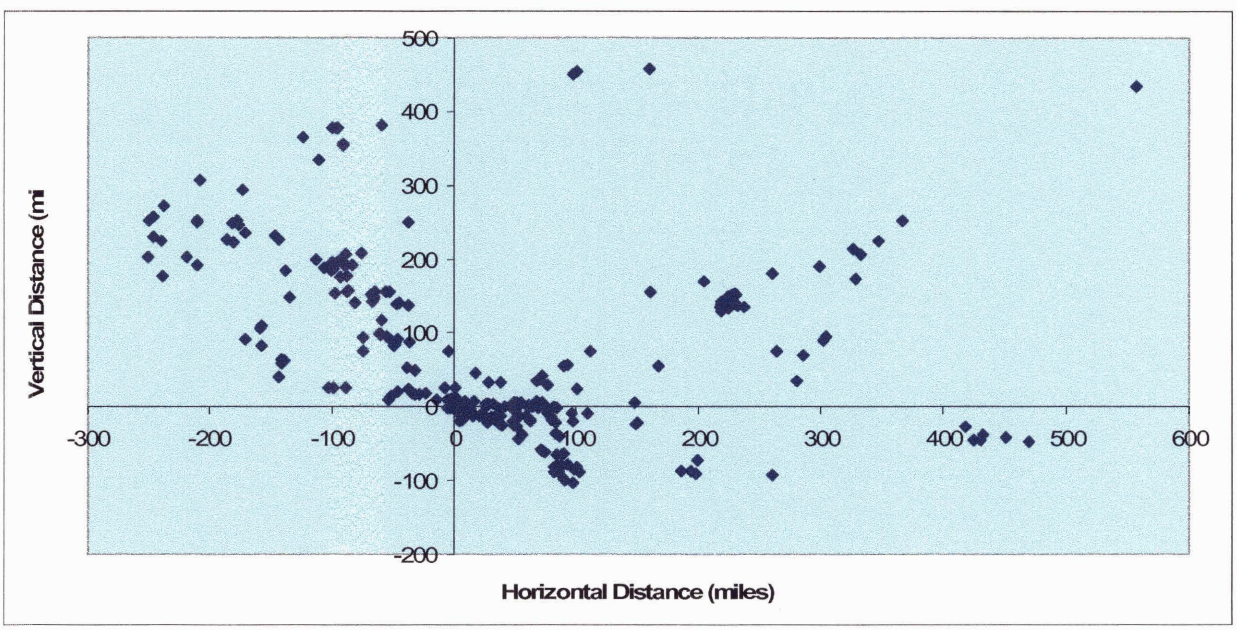

Figure 2: Plot of all customer sites visited

2.3 *Order history analysis*

2.3.1 Cumulative sales by customer location

In order to gain a better understanding of the distribution system under study, our analysis initially focused on the order history. In order to identify each location's contribution to the facility's sales, a plot of cumulative sales volume percentages by customer location was created. Sales volumes, measured in gallons sold, were considered across all products. The plot is provided in Figure 3.

The customer site with the highest total volume ordered, accounts for 5.16% of total sales. The top 39 customer sites (13.49% of the total number of customer sites that were served during the year under study) account for 50.74% of total sales and the top 90 customer sites (31.14% of the total number of sites) account for 80.41% of total sales. Identifying the top selling customer locations is very important for guiding a large scale implementation of VMI since the highest selling locations are the primary candidates. A map that exhibits the top 90 customer sites is provided in Figure 4 and a plot of the same locations, with the origin at company A's facility, is provided in Figure 5. In Figure 5, each location is displayed as a bubble, where the size of the bubble is indicative of the quantity of product ordered.

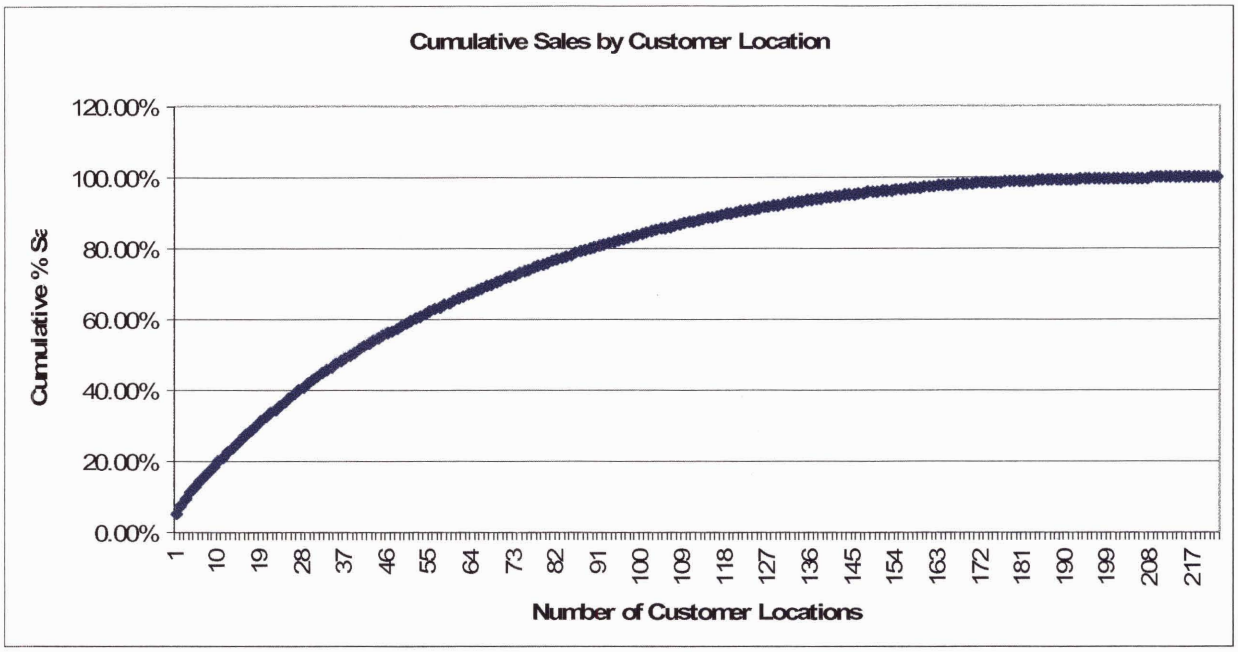

Figure 3: Cumulative sales by customer location

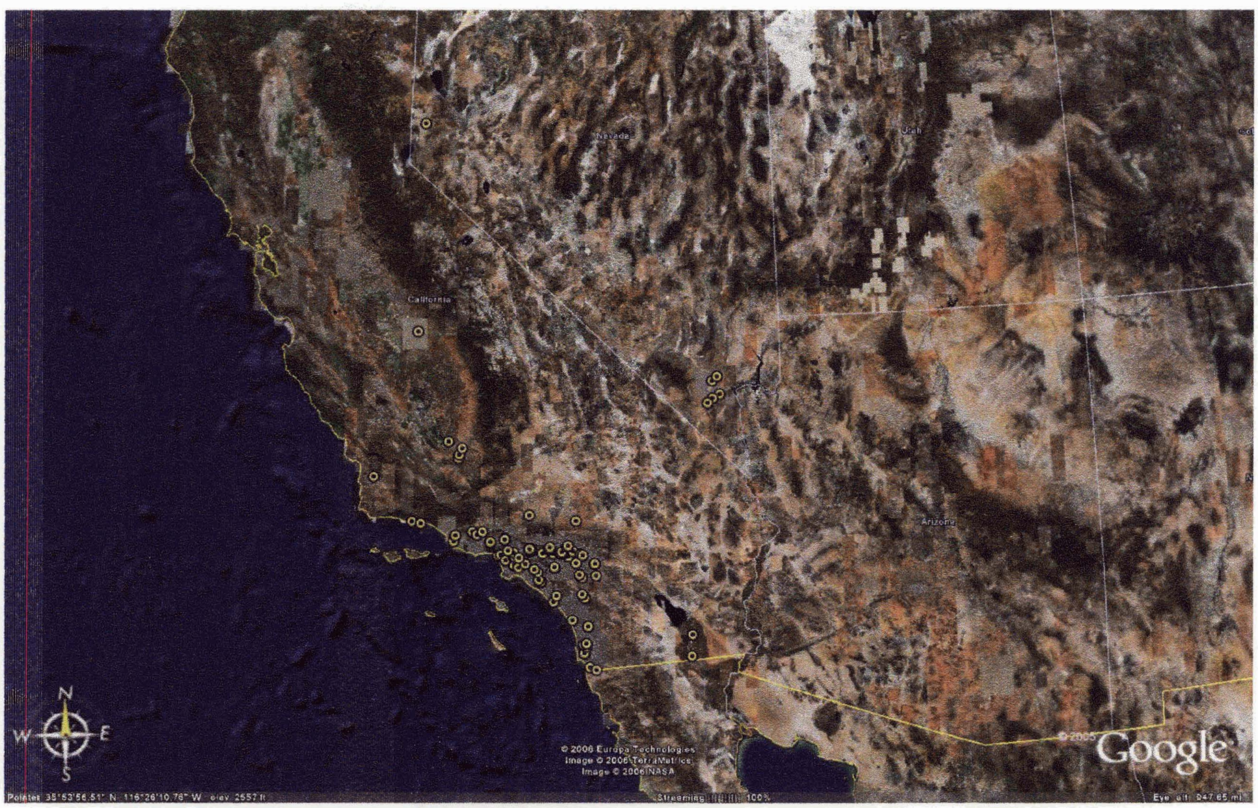

Figure 4: Top customer sites in terms of total volume ordered (Map)

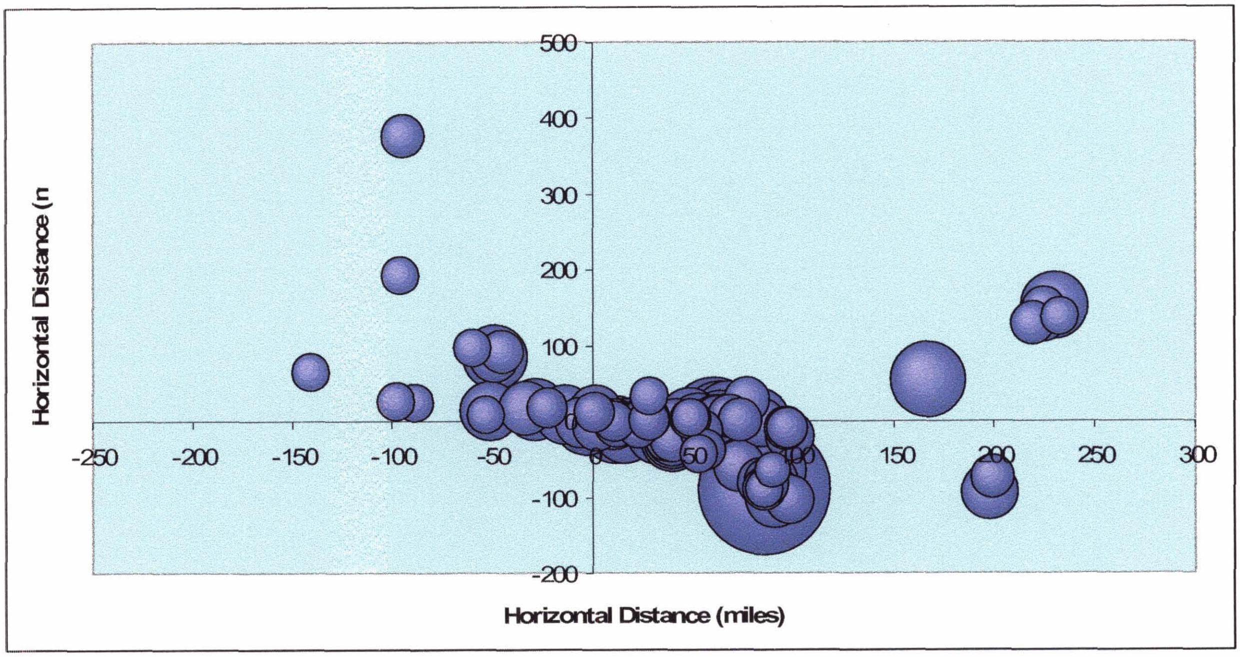

Figure 5: Top customer sites in terms of total volume ordered (Plot)

It can be seen from the plot of the top customer sites that the majority of them are located close to the facility. Although certain sites are in excess of 400 miles from company A's facility, the majority of the sites that account for 80% of the total sales volume are located within a radius of 100 miles from the facility.

2.3.2 Cumulative sales by product

Subsequently an analysis was made to determine which products account for the majority of the facility's sales. High volume products are the primary candidates for being monitored and replenished under a VMI setting since they drive a large percentage of transportation costs. Better coordination of replenishments made for such products would lead to transportation costs savings. In addition, this would allow for better production planning at the company's production facilities. A plot of cumulative sales volume percentages by product was constructed and is presented in Figure 6. Sales were measured in gallons of product.

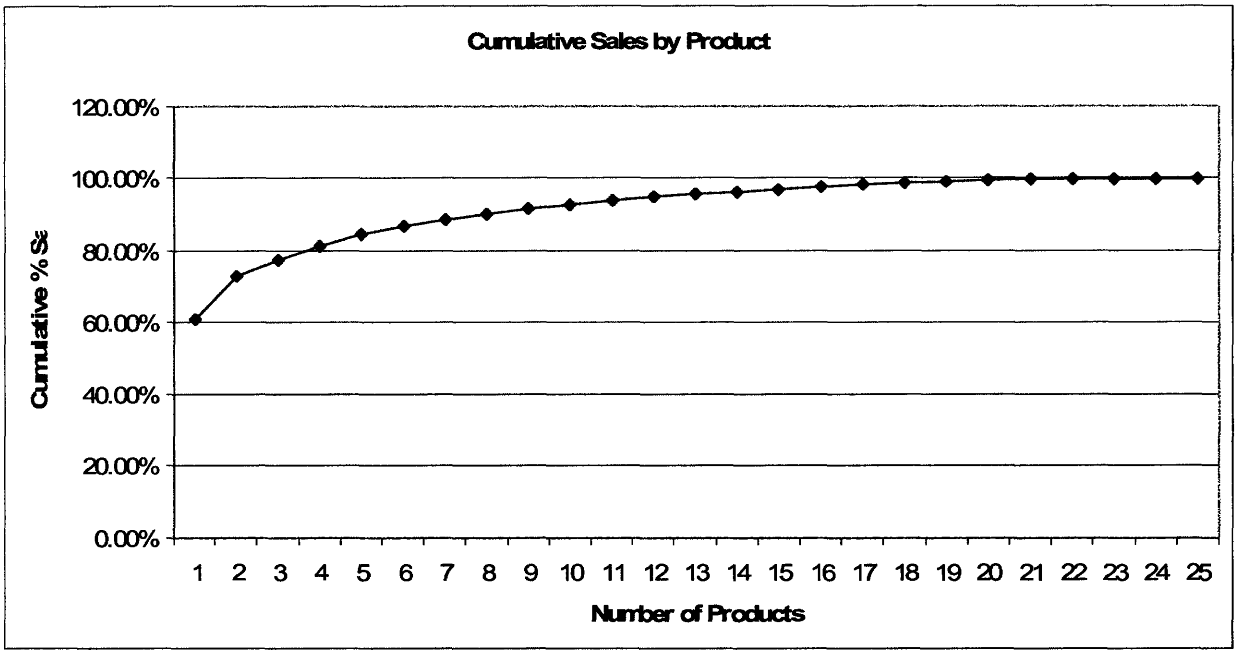

Figure 6: Cumulative sales by product

The most important finding from this analysis is that one product accounts for 60% of the total sales volume throughout the year. In addition, although a total of 25 distinct products were delivered to customer sites throughout the year, four of them account for 81.11% of the total sales volume and the other 21 products account for the remaining 18.89%. Therefore, the top four products become primary candidates for being offered under a VMI agreement.

2.3.3 Order size analysis

After identifying which customer sites and which products account for the majority of the sales volume, an analysis was conducted by product and customer location (at the tank level) to quantify the variability of order sizes, identify the frequency with which storage tanks are replenished, determine the relationship between order sizes and tank capacities and finally determine the relationship between order sizes and truck capacity. Initially, summary statistics were computed at the tank level and these were subsequently aggregated across all tanks in order to compute the aggregate summary statistics that are presented in Table 1.

	Average Order Size (gallons)	Order Size CV	Delivery Frequency (days between orders)	Order Size/Tank Capacity	Order Size/Truck Capacity
Max	2,608.50	1.24	344.00	100.00%	65.21%
Average	627.40	0.30	52.69	48.09%	15.68%
Median	510.00	0.29	32.20	46.25%	12.75%
Min	100.00	0.00	0.00	6.72%	2.50%
Std. Dev.	394.96	0.16	66.12	21.98%	9.87%

Table 1: Order size summary statistics across all tanks

The average coefficient of variation (CV) of order sizes is relatively low (0.30), however CVs range from 0 for products that have been ordered only once to 1.24 for products with highly variable order sizes. Delivery frequency, defined as the number of days between successive orders, ranges from 0 to 344 days with an average value of 52.69 days and a median value of 32.20 days indicating the existence of products that are ordered very infrequently at certain locations. As far as the order size is concerned, both the average and median values are close to each other and indicate that products are ordered in quantities close to half the size of their dedicated tank. Finally, a certain order occupies on average 15.68% of a truck's capacity.

In Table 2, similar summary information is provided as in Table 1 for the top product only, which accounts for 60% of the total sales volume. The order sizes exhibit similar variability on average although the range of the values of the CVs for the various locations is narrower. As expected, delivery frequency is significantly lower on average (from 52.69 to 37.40 days between orders) and with significantly lower variability (standard deviation drops from 66.12 to 37.60 days).

	Average Order Size (gallons)	Order Size CV	Delivery Frequency (days btw orders)	Order Size/Tank Capacity	Order Size/Truck Capacity
Max	2,283.92	0.74	210.00	100.37%	57.10%
Average	914.59	0.29	37.40	47.37%	22.86%
Median	872.39	0.28	28.59	46.33%	21.81%
Min	201.00	0.00	0.00	14.08%	5.03%
Std. Dev.	425.90	0.11	37.60	18.00%	10.65%

Table 2: Order Size sumarry statistics for the top selling product at all customer locations

2.4 Delivery trips analysis

2.4.1 Trip metrics

After analyzing the order history we turned to the delivery trips. Since one of the major goals of instituting a VMI system is to improve the performance of a company's distribution resources, we initially tried to understand how effective the current dispatching practices are and identify potential areas for improvement. Certain metrics were computed using the delivery history that was available and are presented in Table 3. A total of 1040 trips were available to analyze in the delivery history for the year under study.

	Quantity Delivered (gallons)	Distance Traveled (miles)	Num Stops	Truck Util.	Gallons / Mile	Gallons /Stop	Miles /Stop
Average	3,411.59	260	3	84.17%	21.70	1204.19	86.52
Std. Dev.	795.03	201.23	1.17	18.67%	19.91	544.45	71.63
CV	0.23	0.77	0.37	0.22	0.92	0.45	0.83
Median	3,618.50	210	3	90.46%	15.98	1091.17	68.42

Table 3: Trip metrics

Truck utilization is a very important metric for assessing the efficiency of a distribution system. From the delivery history it appears that on average trucks are 84.17% utilized with relatively low variability in the total quantity delivered on a trip (CV of truck utilization is 0.22). The number of gallons delivered per mile driven is used frequently for assessing the effectiveness of different distribution strategies. However, this metric is affected by the physical location of the customer sites being served. For the facility under study, the average number of gallons per mile equals 21.70 but this metric exhibits high variability and has a CV equal to 0.92. On the other hand, the number of gallons delivered per stop exhibits lower variability (CV equals 0.45) and equals 1204.19 gallons/stop on average.

2.4.2 Trip characteristics

In this section we present two histograms, which depict the distribution of the 1040 trips that were performed throughout the year, according to the number of stops per trip and the number of products delivered per trip, as well as a histogram depicting the distribution of the total number of stops according to the number of products delivered per stop.

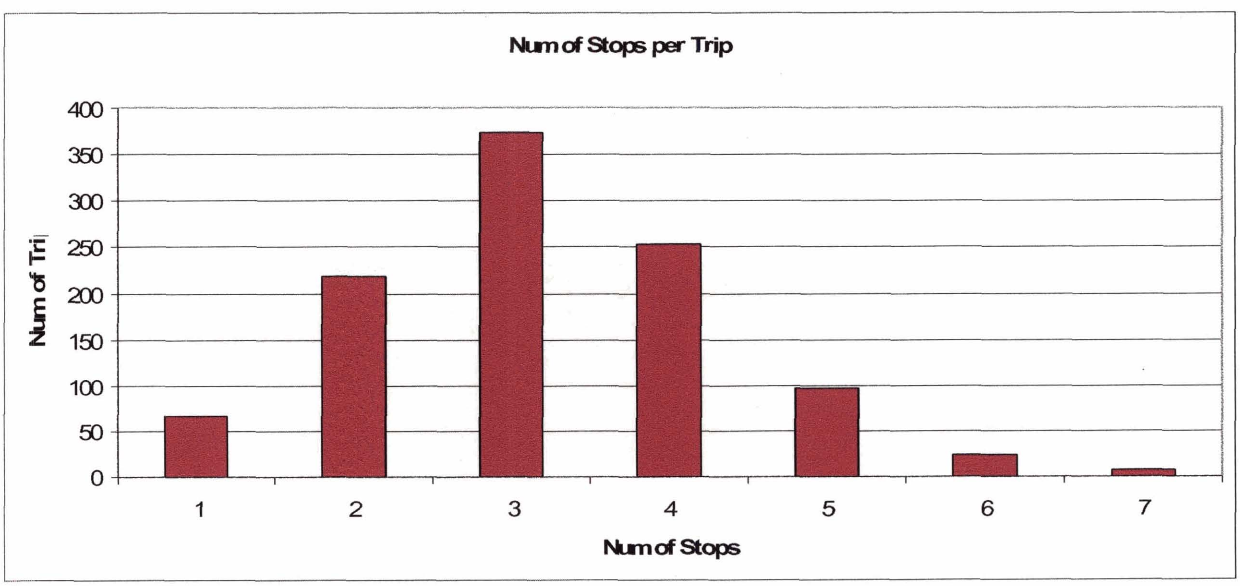

Figure 7: Number of stops per trip

Figure 7 presents the histogram depicting the number of stops per trip. It appears that most (81.25%) are two – three – or four – stop trips with three – stop trips accounting for the largest share of the total number of trips (35.96%).

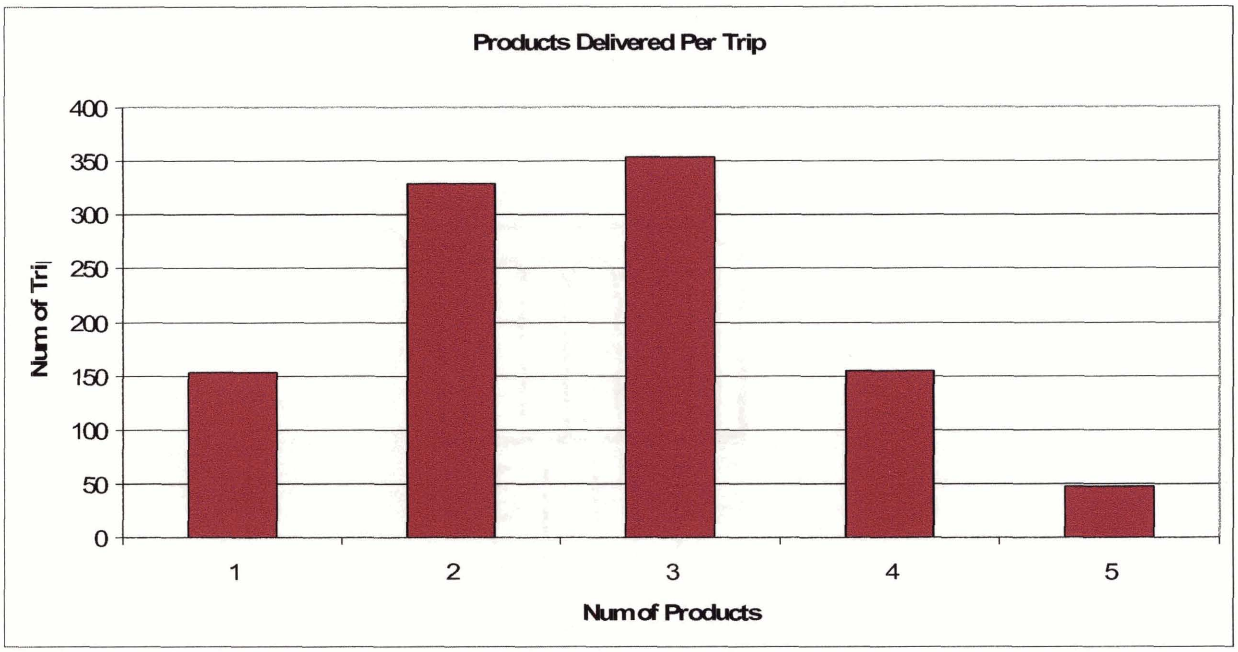

Figure 8: Products delivered per trip

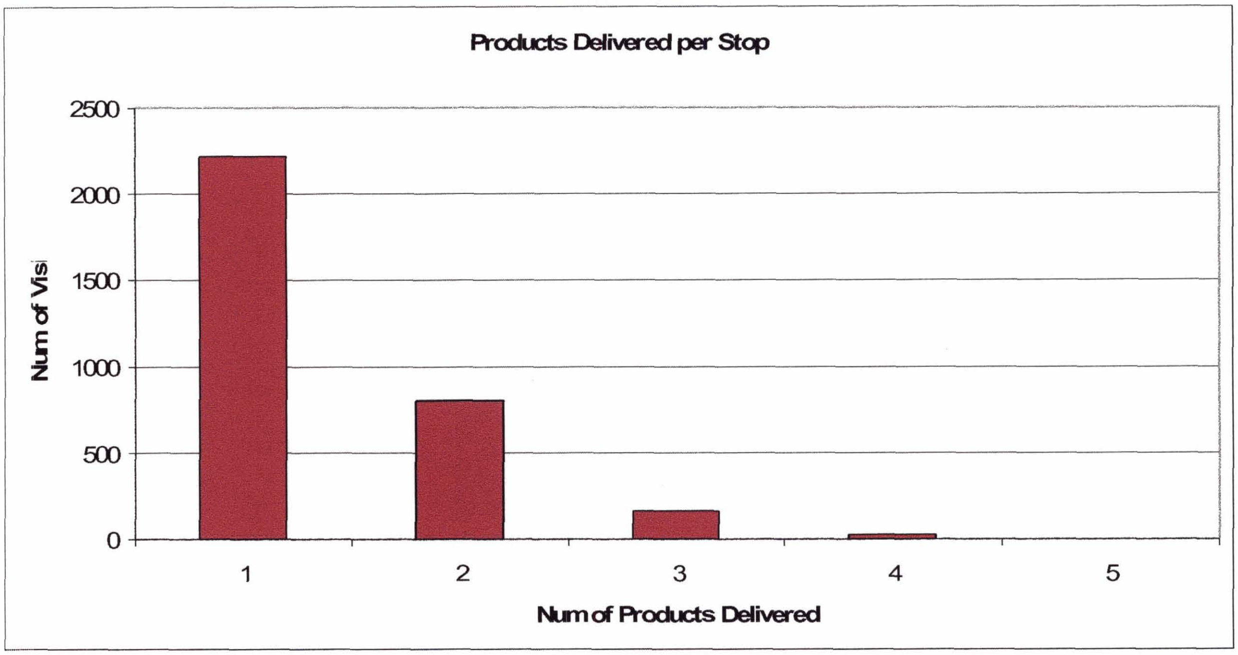

Figure 9: Products delivered per stop

Figure 8 presents the histogram of the distribution of the number of products delivered per trip. It appears that in 80.54% of the total number of trips only one, two or three products were delivered. Five – product trips account for only 4.53% of the total number of trips indicating that only rarely are all five compartments of the trucks used for different products.

Finally, in Figure 9 a histogram is presented that depicts the distribution of the number of products delivered per stop. At 69.03% of the total number of stops only one product was delivered. In an additional 25.05%, two products were delivered and finally, the percentage of stops at which three or more products were delivered only equals 5.91%.

2.4.3 Distance between stops in multi – stop trips

In order to identify inefficiencies in the existing route design process, we grouped the delivery trips that are available in the delivery history according to the number of stops that were made on each trip. Subsequently we computed the average distance between stops for each group of trips. As it was presented in the previous section, the majority are two –, three – and four – stop trips.

The actual sequence in which customer sites were visited on every given trip was necessary in order to compute the average distances between stops. However, since the structure of the available information in the dataset did not make it possible to determine what that sequence was, a method described in (Hernandez, 2004) was used for this purpose. Essentially, for each trip, the mid – angle of the trip, which is measured counter – clockwise from the north, is used in order to determine whether a given customer location was visited when the truck was going out of the facility or when it was returning to the facility. This is achieved by comparing each location's angle to the mid – angle of the trip and by classifying each location with an angle smaller than the mid – angle as going out of the facility and each location with an angle greater than the mid – angle as returning to the facility. Subsequently, the visit sequence is determined by rank ordering all locations according to their distance from the facility. The results of the analysis are presented in Table 4.

It appears that the average distances between stops do not differ widely across the various route groups with the exception of seven – stop routes. The lowest average distance is for two – stop routes and equals 33.477 miles. The average distance increases as the number of stops on the route increases and becomes maximum for the seven – stop routes when it equals 66.100 miles. However, it should be mentioned that limited data for seven – stop trips were available and therefore this result might not be statistically reliable.

Average Distance Between Stops (miles)		Num of Trips
2-Stop Routes	33.477	221
3-Stop Routes	33.516	350
4-Stop Routes	38.447	229
5-Stop Routes	39.574	84
6-Stop Routes	38.916	17
7-Stop Routes	66.100	7

Table 4: Average distance between stops

In order to assess the efficiency of the current dispatching process we need some metric against which the average distances can be compared. A natural choice would be the average distance between the customer locations that are served by company A's facility. If we consider all 289 customer sites that were visited at least once, order them clockwise and compute the average distance between adjacent locations, then this average distance is found to be equal to 80.994 miles. Therefore, it appears that the current route design process is effective in that only customer sites that are located close to each other are assigned on the same trip.

In computing the average distance between adjacent locations, each customer site was given equal weight regardless of the number of times that it was visited throughout the year. For certain customer sites located in distances in excess of 300 miles from the facility, we observed that they had placed only a very small number of orders. Hence, the inclusion of these sites in the above computations greatly increases the computed average distance between stops.

Analysis of the number of orders placed by each customer site, indicates that 154 sites account for 90% of the total number of visits made throughout the year. These sites, which have generated the majority of the company's transportation expenses, are plotted in Figure 10. We see that almost all of these sites are within a radius of 200 miles from the facility. After ordering them clockwise, it turns out that the average distance between stops is 41.768 miles. Therefore, it appears that even after considering only the customer sites that are mostly visited, the computed average distance between stops indicates that the current route design process is effective since only closely located sites are included on the same delivery trips.

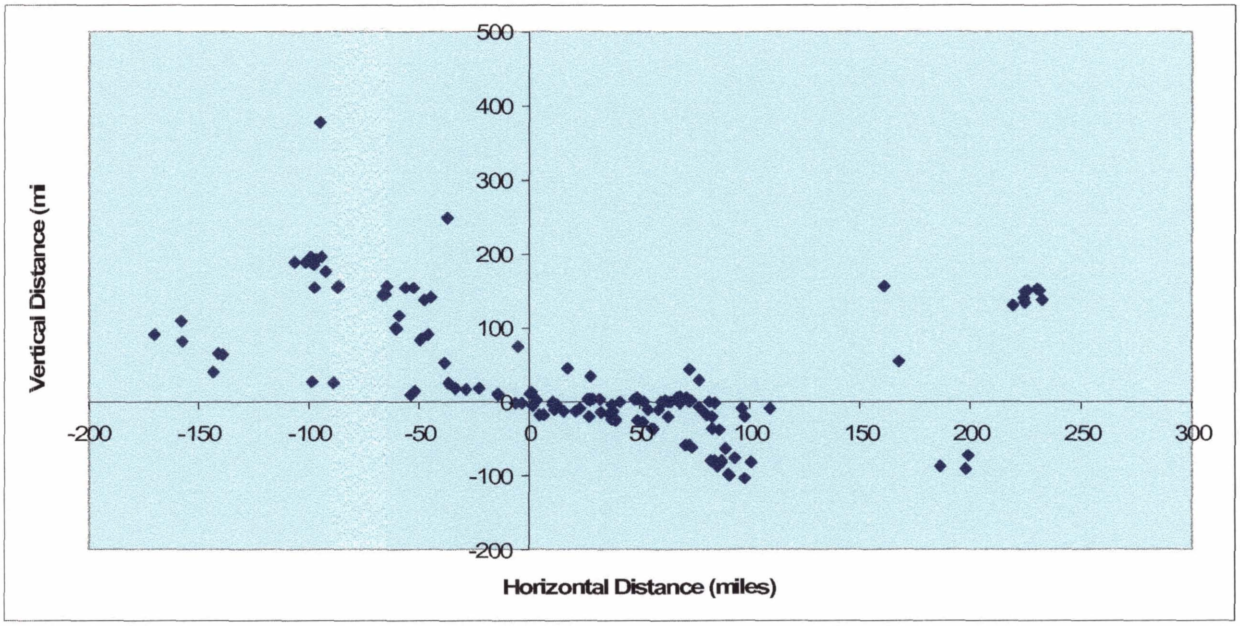

Figure 10: Top customer sites in terms of customer orders placed

2.4.4 Site replenishment using dedicated trucks

From the order history, we observed that several locations were being visited multiple times within given weeks. In most cases, the multiple visits were for the replenishment of different products although certain cases had been noted where the same product had been replenished multiple times in a given week. In addition, as shown above, in 69.03% of the total stops that were analyzed only one product was replenished. Under a VMI setting, when the company will have control over when replenishments will be made as well as over what quantities will be delivered, it might be possible to deliver larger quantities of product and replenish multiple products when a certain location is visited.

In this section we consider the possibility of sending dedicated trucks to customer sites. The reasoning behind this type of replenishment versus the multi – stop replenishment trips that are currently used, is that by sending dedicated trucks to certain locations, higher truck utilization can be achieved and the total number of visits to each customer site will be reduced. This would result in a reduction of the total mileage driven.

In order to test this proposal, we developed a mixed – integer program to determine the optimal number of trips that would be necessary to fulfill the requirements of a customer location. The mixed – integer program was implemented in the OPL Studio modeling environment and was solved separately for each selected customer location.

The decision variables for the program are the following:

x_i is a binary variable indicating whether a truck will be sent to the customer site on week i $(1 \leq i \leq 52)$

y_{ji} is a binary variable indicating whether product j will be delivered on week i (j is customer location specific)

q_{ji} is the quantity of product j delivered on week i

I_{ji} is the inventory of product j on week i

The inputs to the mixed – integer program are the following:

D_{ji} is the demand for product j on week i. Since only delivered quantities are available from the database, we assumed that the replenishment quantities for a given product in a given week represent the demand for that product since the last replenishment.

C_j is the tank capacity dedicated to product j

The objective of the program is to minimize the total number of visits to the customer site. This can be expressed in mathematical terms as follows:

$$\min \sum_i x_i$$

The constraints of the program are defined as follows:

$$\sum_j y_{ji} \leq 5, \forall i$$

$$\sum_j q_{ji} \leq 4{,}000, \forall i$$

$$I_{ji} = I_{ji-1} + q_{ji} - D_{ji}, \forall i, j$$

$$I_{ji} \geq 0, \forall i, j$$

$$I_{ji} \leq C_j, \forall i, j$$

$$x_i \geq \frac{1}{5} \sum_j y_{ji}, \forall i$$

$$y_{ji} \geq \frac{1}{4{,}000} q_{ji}, \forall i, j$$

$I_{ji} \geq 0, q_{ji} \geq 0$, x_i binary, y_{ji} binary, $\forall i, j$

The first constraint ensures that the total number of products delivered will not exceed the total number of compartments available at the truck. The second constraint forces the total quantity delivered not to exceed the total capacity of the truck. The third constraint is used for computing the inventory of every product at the end of each week. Initial inventories are set to be equal to half the tank sizes. The fourth constraint ensures that demand for a product on a given week is met. The quantity available for use on every week is equal to the inventory at the end of the previous week plus the quantity that was delivered at the beginning of the current week. The sum of these two quantities has to be greater than or equal to the demand for that product during the week. The fifth constraint ensures that the inventory at the end of the week does not exceed the available tank capacity. The sixth constraint ensures that x_i is always equal to one when a truck is sent on a given week. Finally, the seventh and eighth constraints ensure that y_{ji} is equal to one when product j is delivered on week i and is zero otherwise.

We implemented the model on customer sites that were among the top both in terms of total volume ordered and in terms of total number of orders placed. In Table 5 the computational results from the implementation of the model are presented. The minimum required number of visits is presented as well as the average truck utilization that was achieved. For comparison purposes we include in the table the actual number of visits that were made. No direct comparisons can be made with the available data in terms of truck utilization since the locations were being visited on multi – stop trips.

Location	Actual Num. Visits	Optimal Num. Visits	Average Truck Util.
1	85	34	97.47%
2	55	20	59.05%
3	53	12	96.34%
4	42	14	47.53%
5	42	15	56.39%

Table 5: Optimal number of stops for individual customer locations

From the results it appears that large reductions are possible in the total number of visits that would be necessary to replenish a customer site, assuming that a dedicated truck is assigned to perform the replenishment for each of the 5 sites. However, the truck utilization that can be achieved depends on the characteristics of each location.

Location 1 is the top location in terms of total volume ordered as well as total number of orders placed. During the year for which the analysis is made, frequent orders were being placed for nine different products. Therefore, this allowed the construction of replenishment trips with very high truck utilization. Location 3 also exhibits very high truck utilization. However, only one product is actually being ordered throughout the year and this has permitted the construction of full truckloads.

As far as the other locations are concerned, it seems that the achievement of high truck utilization is not possible. This is due to the fact that although each of these locations orders multiple products throughout the year, only one product is ordered frequently with the others being ordered on a very rare basis. For these locations, the tank capacity of the high volume product becomes a binding constraint and leads to the generation of trips that only replenish that product, which results in very low truck utilization.

As a result, it appears that use of trips dedicated to replenishing multiple products to a single location is a viable alternative only if several fast moving products are in use at that location. However, the vast majority of customer sites that are served by company A's facility do not possess this characteristic. Therefore, the use of multi – stop trips is necessary in order to achieve high utilization of the fleet.

3 Methodology

This section describes the process used to determine candidates for VMI, place them into delivery clusters, create routes within those clusters, and measure the effort required to service VMI and non-VMI customers. First, we develop a policy to segment tanks into those that will be served by VMI and those that will continue with order based fulfillment. Next, we group all customers into geographic clusters that are used to develop delivery routes. We then develop a program was then developed that creates regular delivery routes within a cluster for VMI tanks. Finally, a method for measuring the delivery effort for non-VMI tanks is described.

3.1 Segmentation

Based on the analysis of customer and product order history a segmentation policy was developed that identifies candidates for VMI at the tank level. A tank is defined as a single product at a single customer location. In reality a location may have multiple tanks containing the same product, but for the purpose of this thesis they are considered a single tank having a capacity equal to the total capacity of those tanks.

The total volume delivered to each tank was calculated. The tanks were then sorted in decreasing order based on total volume. We added tanks to the VMI group in order of volume until 80% of total delivered volume was in the VMI group. This produces a set of 248 tanks at 163 different customer locations. These tanks represent 67% of total orders and 80% of total delivery volume, yet comprise only 26% of the total number of tanks. A summary of the tanks selected for VMI is shown in Table 6.

	VMI	Total	% of Total
Tanks	248	946	26%
Deliveries	3020	4535	67%
Volume(gallons)	2854754	3568104	80%

Table 6: Tank Segmentation Statistics

The remaining tanks are not considered candidates for VMI and will continue to be served through order based fulfillment. This includes tanks that are located at customer locations where other tanks are served through VMI, but do not meet the threshold for VMI candidacy. The low volume and infrequency of delivery means these tanks require relatively low effort to service, and therefore it is unlikely that any improvements could be made that would offset the vast increase in cost and complexity required to plan for these tanks in VMI.

3.2 Clustering

Clustering is defined as the unsupervised classification of patterns into groups. It can be a difficult problem, but has a wide range of applications including data mining, pattern recognition, and grouping (Jain et al, 1999). This thesis is concerned with its applications for grouping, specifically to take a collection of customer locations and place them in groups based on their geographic locations. Once the customers are clustered into locations, actual routes are devised within the clusters. Separating customers into clusters reduces the complexity of the routing problem.

3.2.1 K-means Clustering Algorithm

The clustering algorithm employed is the k-means algorithm. This algorithm is a simple and economical procedure for separating the population into k groups (MacQueen, 1967). The procedure works by choosing k initial customer locations as centers of clusters. Additional customer locations are added, and placed in the closest cluster. After all locations have been added to clusters a new center for each cluster is computed, and locations are shifted to the cluster with the closest center. This procedure continues until new centers are computed and no locations are shifted between clusters. Mathematically the algorithm attempts to optimize the following objective function:

$$\min \sum_{j=1}^{K} \sum_{i=1}^{n_j} \sqrt{(x_i^j - c_j)^2 + (y_i^j - d_i)^2}$$

where:

n_j = number of elements in cluster j

x_i^j, y_i^j = the x and y coordinates of the i^{th} element of cluster j

c_j, d_j = the x and y coordinates of the center of cluster j

The k-means algorithm does not guarantee a global optimal solution. That is, it may converge to a local optimal solution, and there may exist a different grouping of clusters that results in a lower total sum of distances to the centers of clusters. Despite the lack of guarantee of a global optimization and the sensitivity to the initial seed points, its intuitive behavior and ease of computation make it very popular (Jain et al).

3.2.2 Results of Clustering

All 289 customer locations were placed into clusters. For the VMI candidates the clusters will be used to design delivery routes within the cluster. Figure 11 shows the results of the clustering algorithm for the 163 customer locations served by this facility that are candidates for VMI. Each location is assigned an X and Y coordinate based on their position relative to the company facility, which is located at the origin. The value of k for this algorithm was chosen as 4, based on an estimation of the number of trucks required to serve the entire region. After performing the initial clustering Cluster 2 made up a significant amount of customer locations and total volume, and this cluster was then split into Cluster 2 and Cluster 5. We ran the clustering algorithm several times with different initial seeds and compared the results. In each instance the same general pattern of clusters appeared, and more than 80% of customer locations remained in the same cluster regardless of the initial seeds.

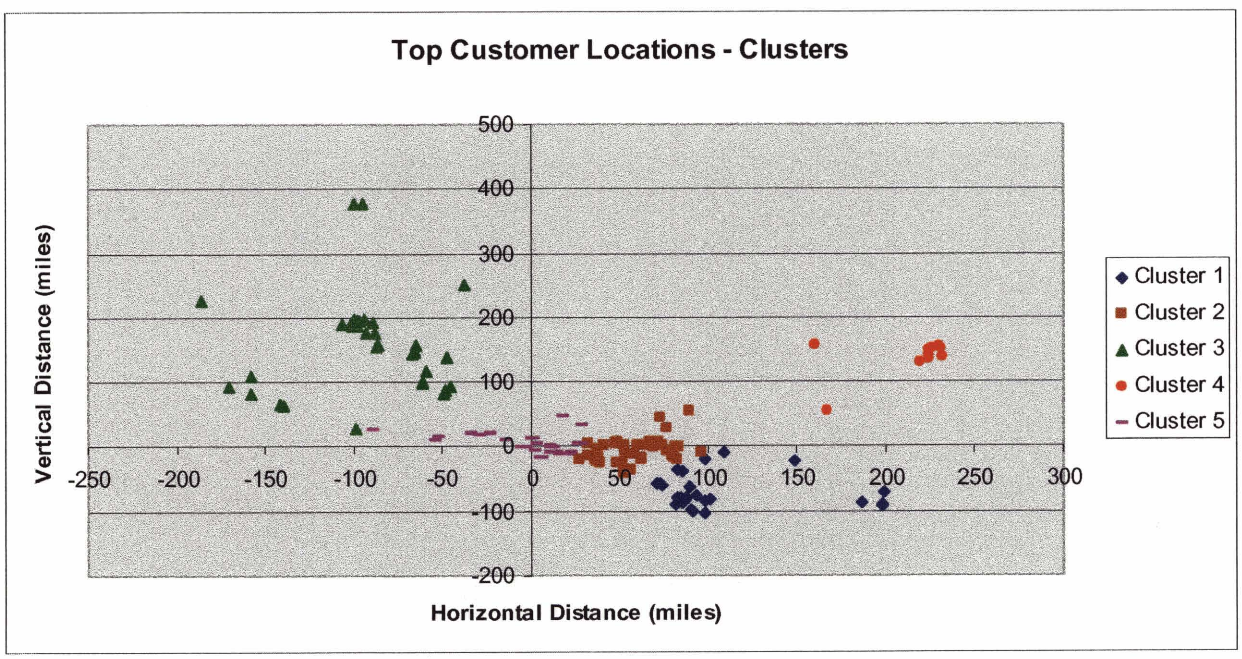

Figure 11: Results of Clustering

Table 7 shows the statistics for each of the 5 clusters.

Cluster	Customers	Tanks	Annual Volume (gallons)
1	30	55	751900
2	51	78	1090087
3	35	43	263605
4	11	22	182444
5	36	50	566748

Table 7: Cluster Statistics

3.3 Routing

Once the designated customers have been placed into clusters, it is possible to construct a series of routes to serve each cluster. A mixed integer program was constructed in order to determine the series of routes that meet the specified demand in the fewest number of trips subject to constraints on the capacity of the vehicles and customer tanks. The output of this program is a series of routes. For each route we specify the tanks that receive a delivery, the quantity delivered to each tank, and the frequency the route is run.

In order to determine daily demand at each tank the three year order history was examined. A long run average demand was computed by calculating the total quantity delivered and dividing that by the number of days elapsed. For the purposes of this model, demand is considered to be deterministic and non-seasonal.

Each route is subject to two constraints on the capacity of the vehicle and one constraint on the quantity delivered to each tank. The vehicle constraints consist of the volume capacity of the vehicle and the number of compartments. Vehicles are considered to be identical and have a capacity of 4,000 gallons. Each vehicle is also compartmentalized, allowing it to hold multiple products. For this model each vehicle has a maximum of five compartments. There are no constraints on the capacity of individual compartments; rather, they are considered adjustable to hold any volume as long as the total volume of the compartments does not exceed the 4,000 gallon constraint on total vehicle capacity. The capacity of the tank at the customer location provides another constraint. Based on the tank size and the daily demand rate of the tank, a minimum frequency with which the tank must be serviced is computed. Servicing the tank less frequently than this minimum value would require delivery of more product than the tank allows.

For example, consider a tank of capacity 1000 gallons and a demand rate of 10 gallons/day. This tank then has a maximum time between deliveries of 100 days, for a minimum frequency of .01 times per day. Attempting to service the tank less frequently, such as .005 times per day, would result in receiving a delivery every 200 days. In order to meet the demand of the tank, it would require a delivery of 2000 gallons, violating the constraint on the capacity of the tank.

The mixed – integer program was developed and implemented in the OPL Studio modeling environment. The decision variables that were used are described below:

F_k is the frequency in trips/day with which route k is executed

Q_{ijk} is the quantity of product i delivered to location j on route k in gallons/day

x_{ik} is a binary variable indicating whether product i is delivered on route k.

y_{jk} is a binary variable indicating whether location j is visited on route k.

z_{ijk} is a binary variable indicating whether product i is delivered to location j on route k. This variable is defined only when product i is indeed stored at location j.

Since the model will be run for each geographical cluster separately, thus ensuring the geographical proximity of the customer locations to be served by the routes that will be designed, we decided to create routes that would be executed less frequently. Therefore, the objective of the mixed – integer program is to minimize total route frequency, which equals the number of truck trips per day. In mathematical terms the objective is presented below:

$$\text{Min} \sum_{k} F_k$$

The inputs to the mixed – integer program are:

D_{ij} is the demand rate in gallons/day for product i at location j

$Days_{ij}$ is the maximum allowable number of days between successive replenishments of product i at location j. As mentioned above, this quantity is computed by dividing the available tank capacity for product i at location j with the demand rate D_{ij}. Alternatively, if a delivery quantity less than the tank capacity is desired, the number of days between successive replenishments can be computed by dividing this quantity with the demand rate.

numLocations is the number of distinct customer locations within the cluster

numProducts is the number of distinct products delivered at the various customer locations within the cluster

The constraints of the program are presented below:

$$\sum_{i} x_{ik} \leq 5, \ \forall k$$

$$\sum_{j} y_{jk} \leq 5, \ \forall k$$

$$Q_{ijk} = D_{ij} * z_{ijk}, \forall i, j, k$$

$$\sum_{ij} Q_{ijk} \leq 4{,}000 * F_k, \forall k$$

$$F_k \geq \frac{1}{Days_{ij}} * z_{ijk}, \forall i, j, k$$

$$\sum_{k} z_{ijk} = 1, \forall i, j$$

$$x_{ik} \geq \frac{1}{numLocations} \sum_{j} z_{ijk}, \forall i, k$$

$$y_{jk} \geq \frac{1}{num\Pr oducts} \sum_{i} z_{ijk}, \forall j, k$$

$F_k \geq 0, Q_{ijk} \geq 0$, x_{ik} binary, y_{jk} binary, z_{ijk} binary, $\forall i, j, k$

The first constraint enforces the fact that only five products can be carried simultaneously by a truck. The second constraint limits the total number of visits that can be made on a given trip to five. The third constraint ensures that the daily quantity delivered to a tank is equal to that tank's daily demand rate. The fourth constraint ensures that the total quantity delivered on a route cannot exceed the truck's total capacity. The fifth constraint forces each route to be executed as frequently as the tank with the maximum visitation frequency on the route requires.

The sixth constraint ensures that each tank is served by exactly one route. The remaining two constraints are used as linking constraints between the binary variables of the program. The results of the model are presented in section 4.1.

3.4 Determining Travel Distance

In order to determine the total distance traveled, different approaches are used for the VMI and non-VMI tanks. The tanks serviced by VMI are served by a defined set of routes determined by solving a mixed – integer program run at regular intervals. It is possible to calculate the distance required for these routes with much more precision than for tanks served by order based fulfillment. For those tanks not served by VMI an estimate of total distance required is created based on the quantity to be delivered and the characteristics of the delivery clusters.

3.4.1 Tanks Serviced by VMI

For tanks serviced by VMI we will determine a defined set of routes and frequencies with which they are run by solving a mixed – integer program. Given the low number of routes and the small number of customer locations visited on each route, we can solve a Traveling Salesman Problem (TSP) for each route. The TSP defines the shortest path that visits each stop on the path once and only once. For this problem the route begins and ends at the plant and visits each of the customer locations on the route. Given the computed distance of the route and its frequency it is possible to calculate the number of miles required to serve the route over a given period of time, in this case on an annual basis.

3.4.2 Tanks Serviced by Order Based Fulfillment

The large number of tanks, locations, and orders makes calculating the travel distance for non-VMI tanks a difficult problem. Instead, a method is used to estimate the travel distance required based on the estimated number of trips required, stops per trip, and predicted average distance per trip (Hernandez, 2003). This procedure is applied for each cluster of customers defined by the clustering algorithm. By placing the customers into separate delivery regions we can obtain a better estimate than if we considered all customers to be in the same delivery region. The estimated number of trips, l, to each cluster is based on historical truck utilization information:

$$l = \frac{V_{Total}}{V_{avg}}$$

where:

V_{Total} = Total volume delivered to cluster

V_{avg} = Historical average delivery volume per trip

The number of stops on an average tour, c, is calculated based on the estimated total number of customer visits, v, required and the number of trips.

$$c = v / l$$

where:

v = total number of customer visits

l = total number of trips

The number of visits, v, is estimated based on the total number of orders placed for all customer tanks, the minimum number of visits required per customer, and the percentage of orders for different tanks at the same customer location that are combined on the same delivery. The percentage of a customer's orders that are combined on a single delivery, p, is estimated based on historical delivery and order information.

$$v = \sum_i v_{\min,i} + p * (o_i - v_{\min,i})$$

where:

o_i = total orders placed by customer i

$v_{min,\ i}$ = minimum number of visits required by customer i, calculated as the maximum number of orders received for a single tank

p = percentage of additional orders that result in new trips

Next, the average length of each trip must be estimated. Each trip can be thought to consist of a line haul to the first customer location, a series of local trips to other customers, and finally a back haul from the last customer to the plant. Figure 12 illustrates a typical delivery tour.

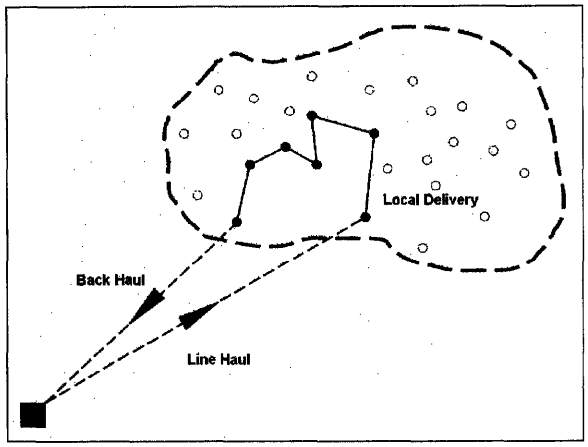

Figure 12: Illustration of a Vehicle Tour

The total distance, D_{tour}, is then given by:

$D_{tour} = D_{linehaul} + D_{local} + D_{backhaul}$

An estimation of this distance is provided by calculating the distance of the line haul and back haul as twice r, the distance from the plant to the center of the delivery region, plus the local delivery distance D_{local} (Daganzo, 1984).

$D_{tour} = 2r + D_{local}$

The local delivery distance can be calculated by first determining an average distance between stops in the delivery region (Bearwood et al, 1959). We estimate the length of a traveling salesman problem through n points in a delivery region of area A to be:

$$D_{TSP} = k\sqrt{nA}$$

where:

n = number of points

A = area of the region

k = constant estimated at .9 for rectangular delivery regions (Daganzo, 1984)

The average distance between stops is then given by:

$$D_{stop} = \frac{D^{TSP}}{n} = \frac{k\sqrt{nA}}{n} = k\sqrt{\frac{A}{n}}$$

If we then define the stop density, δ, to be:

$$\delta = \frac{n}{A}$$

The average distance per stop now becomes:

$$D_{stop} = k\delta^{-\frac{1}{2}}$$

The average local delivery distance is then found by multiplying that average distance between stops, D_{stop}, by the average number of stops in a tour, c (Daganzo). Now that an estimate for the number of trips required and the average distance of a trip in each cluster have been calculated an estimate for the total travel distance required can be calculated as:

$$D_{total} = l \cdot (2r + ck\delta^{-\frac{1}{2}})$$

where:

l = total number of trips

r = distance to the center of the delivery region

c = average number of stops per tour

k = constant value of .9

δ = stop density within the delivery region

By applying these methods we are able to estimate the total distance required to service all tanks, and compare this result with the current order based fulfillment system. The results of this methodology are presented in the next chapter.

4 Results

This chapter provides a summary of the results obtained by applying the methodology described in Chapter 3 to a single facility. A series of routes were generated by the model for the VMI tanks in each cluster. The total annual mileage traveled to service those tanks was calculated, and this was added to the estimated mileage required to serve the non-VMI tanks. This value was compared with the estimated mileage required to serve all tanks under the order based fulfillment strategy. Finally, the results of the mileage estimation procedure are compared against the calculated distances of the VMI tanks as predicted by solving the TSP problem for each route, as well as against actual data collected by Company A for its deliveries.

4.1 VMI Routes

For each cluster a series of routes were generated that lists which tanks are visited on the route, the quantity of product to be delivered, and the frequency the route is to be run. Table 8 shows a summary of the route information for each of the clusters.

Cluster	Number of Routes	Trips Per Year	Annual Mileage	Annual Volume Delivered (gallons)	Avg. Stops Per Trip	Avg. Products Per Trip	Avg. Volume Per Trip
1	16	188	67265	751907	3.45	2.61	4000
2	23	273	48022	1090094	3.66	2.58	4000
3	15	66	32694	263603	3.34	2.09	4000
4	6	46	25754	182438	3.21	2.25	4000
5	14	142	16020	566754	3.52	2.10	4000
Total	74	714	189755	2854796	3.52	2.43	4000

Table 8: VMI Route Statistics

The key figure from the trip information is the 100% truck utilization. The VMI routes feature full trucks, reducing the total number of trips required.

Figure 13 shows a histogram of the number of stops on a trip for each of the four clusters.

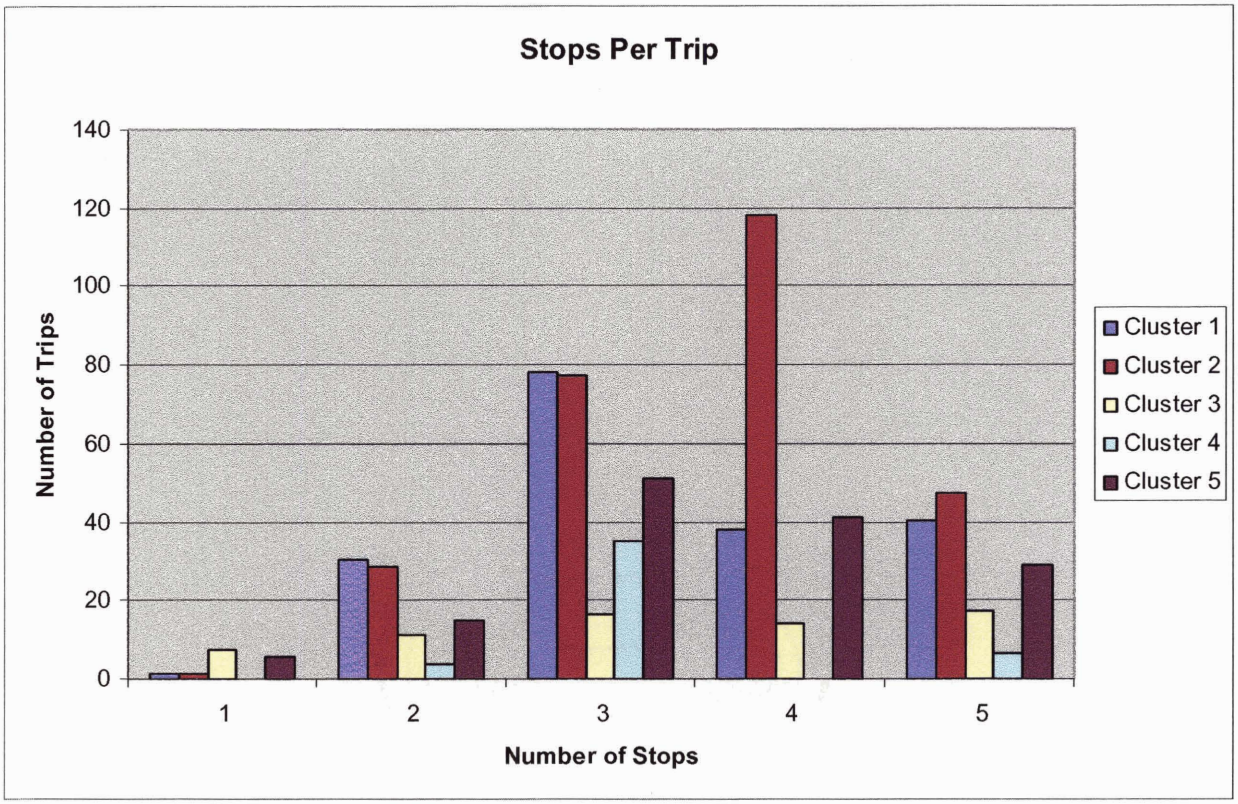

Figure 13: Stops Per Trip on VMI Routes

Approximately 78% of routes consist of 2, 3, or 4 stops; similar to the 82% of actual routes driven in 2005. In general, however, the VMI routes average a higher number of stops per trip.

Figure 14 shows a histogram of the number of products per trip for each of the four clusters.

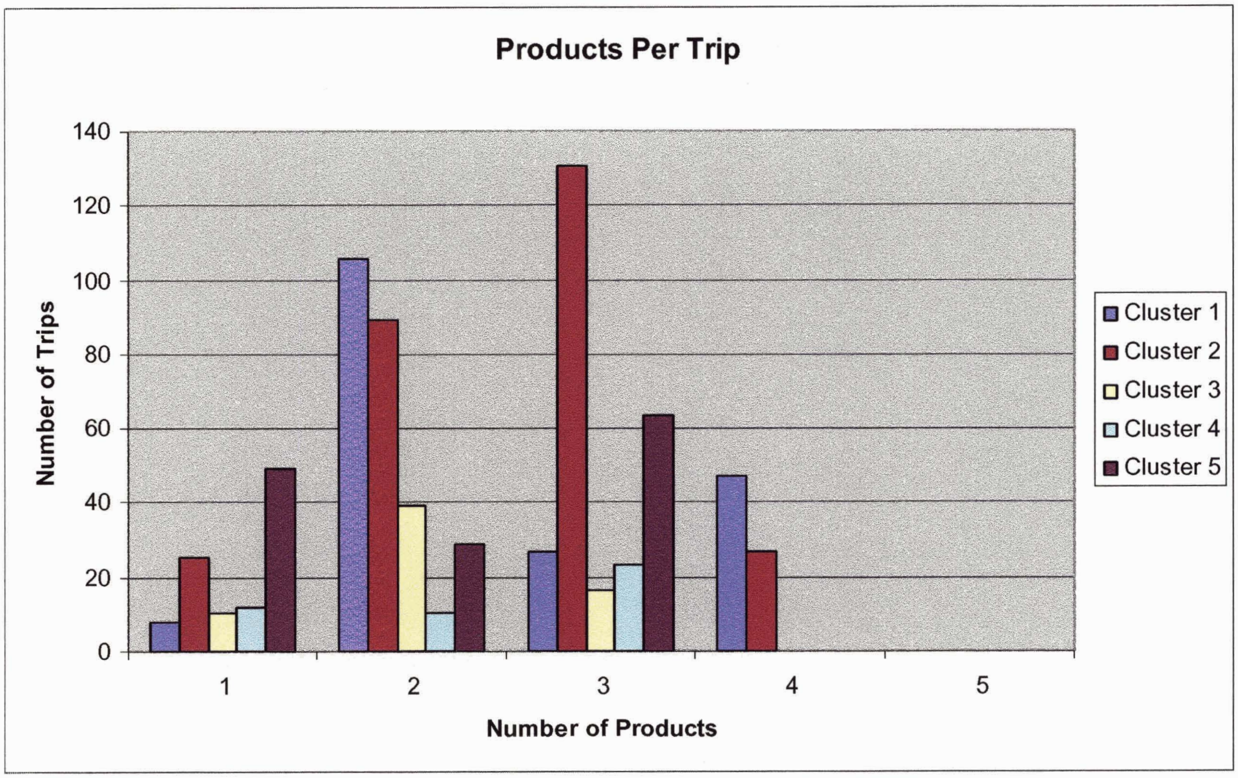

Figure 14: Products Per Trip on VMI Routes

None of the VMI routes use the full 5 product compartments available on the trucks. In general the number of products delivered on a VMI routes are similar to that of the actual delivery routes.

4.2 Order Based Fulfillment Tanks

For the tanks not served by VMI the method described in Chapter 3 was used to estimate the total mileage required to serve the tanks in each cluster. Table 9 summarizes the results of this procedure.

Cluster	Trips Per Year	Volume Delivered (gallons)	Avg Stops Per Trip	Mileage
1	35	119611	5.45	13216
2	44	148500	5.46	7553
3	59	199668	5.12	36012
4	26	89116	3.83	13664
5	46	156425	5.42	5804
Total	210	713320	5.15	76250

Table 9: Non-VMI Route Statistics

These calculations were done based on the following values for the average delivery volume and percentage of customer orders combined on a single delivery.

$V_{avg} = 3400$ gallons

$p = .4$

The average number of stops is significantly higher than either the VMI routes or actual delivery data. This is expected given the character of the tanks being served in this manner. The orders served here are smaller in volume and scattered across many different customer locations. This reduces opportunities for combining customer orders on a single delivery route and requires more tanks to be visited on a given route to achieve the average 3400 gallon utilization.

4.3 Comparison to Current Practice

Given the expected mileage required to serve both VMI and non-VMI tanks, it is now possible to calculate the total mileage and number of delivery trips and compare these values with the current order based fulfillment practice. Table 10 summarizes the performance of the new system and the current order based fulfillment system.

Method	Trips Per Year	Mileage	Total Volume (gallons)	Avg Stops Per Trip
VMI	714	189755	2854796	3.52
non-VMI	210	76250	713320	5.15
Total	924	266005	3568116	3.89
Current	1049	282146	3568104	3.17

Table 10: Comparison of Systems

Under the proposed system of selectively using VMI on high volume customer tanks this methodology predicts a savings of nearly 6% in total mileage driven and 12% fewer total delivery trips. This savings is achieved through superior truck utilization on routes serviced by VMI, where the total mileage required is estimated to be 10% lower than under the current system. Though the average route length is longer under the new system due to the increased stops per trip the reduction in number of trips achieved through better truck utilization outweighs this increase. This insight is illustrated by comparing the average line haul distance, estimated at twice the distance from the depot to the center of the delivery region, to the average distance between stops. As shown in Table 11 the average distance of a line haul is much greater than the distance between stops. Under these conditions the marginal cost of an extra stop on a delivery route is small when compared to the reduction in mileage obtained by elimination of an entire trip.

Cluster	Line Haul (miles)	Avg Distance Between Stops (miles)
1	292.34	15.29
2	124.61	8.84
3	462.14	29.54
4	445.91	19.69
5	66.00	11.10

Table 11: Line Haul vs. Average Stop Distance

4.4 Evaluation of Estimates

In order to evaluate how well the procedure outlined in Chapter 3 estimates delivery distances there are two comparisons that can be made. First, the total mileage estimated under the current order based fulfillment system can be compared to actual recorded delivery distances. Second, the procedure can be applied to estimate the delivery distance for the VMI generated routes, and those distances can be compared to the distances computed for those routes derived by solving a traveling salesman problem for each route.

4.4.1 Estimated vs. Actual Distance Under the Current System

Table 12 shows a comparison of the estimated distance traveled to the actual recorded distance. The estimated values were obtained by applying the methodology outlined in section 3.4.2 to the actual order history for a year where delivery information was available. Based on the orders within a cluster an estimate was made of the number of delivery trips required, the average number of stops per trip, and the average mileage per trip. The estimates for each cluster were then summed and compared with the actual recorded mileage and trip information from the delivery data. The total number of trips and average stops per trip are very similar, as expected since a study of the historical data was used to develop the methodology. The average distance per route was not influenced by historical data, however, as it was calculated by a method that used only the number of customer locations and total area in a delivery region. The results provided by the estimate are within about 1.2% of the actual recorded distance.

	Volume	Trips Per Year	Avg Stops Per Trip	Mileage
Estimated	3568104	1049	3.17	282146
Actual	3568104	1052	3.19	277765
Error	0.0%	-0.2%	-0.7%	1.6%

Table 12: Predicted vs. Actual Delivery Statistics

4.4.2 Estimated vs. Calculated Mileage Under VMI

For the VMI routes the actual distance traveled was calculated by solving a TSP for each route generated. The total annual mileage was determined by multiplying the calculated distance of each individual route by the number of times it was run on a yearly basis. In order to provide another test of the system used to estimate distances the results calculated by solving the TSP can be compared with the estimated distance. Table 13 shows a summary of this comparison. The method employed assumes a uniform random distribution of locations within each delivery region, but this assumption may not be valid for each region. An examination of Cluster 3 shows many locations tightly grouped near the center of the region, with others scattered along the edge. This suggests that assuming such a distribution of locations may not be valid for all delivery regions. However, though within a given cluster the error rate can be high, on an aggregate level the estimated mileage is within 0.6% of the calculated value.

Cluster	TSP Distance (miles)	Predicted Distance (miles)	Error
1	67265	66624	-0.95%
2	48022	44346	-7.66%
3	32694	38112	16.57%
4	25754	23729	-7.86%
5	16020	15862	-0.99%
Total	189755	188672	-0.57%

Table 13: Predicted vs. Calculated TSP Distance

5 Conclusion

The method outlined in this thesis shows that through the improved demand visibility provided by VMI it is possible to achieve reductions in the number of delivery trips and total mileage traveled to deliver Company A's products. It is difficult to judge the accuracy and applicability of these findings to actual practice given the assumptions required to reduce the complexity of the problem. Assuming deterministic demand and flexible compartment sizes within vehicles were necessary for construction of the model, but in practice the uncertain nature of demand and fixed compartment sizes reduce the opportunities to fully utilize truck capacity. The work contained in this thesis provides a basis for future research that can improve upon the results and accuracy, as well as provide several insights for helping Company A improve its vehicle routing. First, due to the high variations in demand rates of products located at the same customer site it does not appear that replenishment through dedicated trucks is a viable strategy. Second, by focusing on better truck utilization it is possible to reduce the total number of delivery trips. In a delivery area where the length of the line haul is much greater than the average distance between customers, the decrease in total number of tours more than offsets the additional distance of making more stops on a route. Finally, in many cases the limiting factor on improving efficiency was the capacity of the tank at the customer site. Tanks that were too small result in a higher frequency of deliveries and reduced opportunities for efficient utilization of truck capacity.

5.1 Future Research

Certain limitations are inherent in the structure of the model that was developed in this work. The most important is the use of deterministic demand rates that are constant throughout the year. As it has been mentioned in several places in the thesis, certain products at certain locations exhibit high seasonality and demand fluctuations. Therefore, a mathematical model that would capture these fluctuations through the use of stochastic demand rates would be a natural extension to the existing model. In addition, the use of the actual demand history, instead of the order history, for the estimation of demand rates will also lead to more accurate results.

Second, a more thorough model should also include an inventory component. After our discussions with people from company A, it was decided for this work to focus on the transportation component of the distribution system since this is the one that appeared to be most promising in terms of the savings that could be realized under a VMI setting. Therefore, the effects on inventory of implementing less frequent trips were not explicitly studied in this work.

In combination with the study of the inventory component, an analysis of the effects of varying tank sizes at the customer sites would also be insightful. It was shown in the analysis of the current system that for several fast moving products, it is the size of the dedicated tank capacity that becomes a constraining factor and leads to the generation of a large number of orders and visits at certain customer sites. Therefore, it would be of great interest to study how an increase in tank sizes for the fast moving products would lead to a reduction in the total number of visits necessary at the high volume customer sites. Such a study should be performed in combination with a study on the effects of tank sizes to the average amounts of inventory that would be held at the customer sites. This would be especially important to determine if consignment is to be offered to customers.

Finally, in designing truly optimal routes that could serve a VMI system, the location of each customer site as well as the demand characteristics of the products that are stored at that site should be considered jointly. In our approach, in order to simplify the mathematical model, we designed customer clusters based only on geographical proximity. Subsequently, the only criteria that were used to decide whether to include a certain location on a trip were the frequencies with which the tanks at each location should be replenished. Therefore, a joint consideration of replenishment frequencies and geographical locations could lead to the design of more effective routes.

Bibliography

Angulo, A., Nachtmann, H., & Waller, M. A. (2004). Supply Chain Information Sharing in a Vendor Managed Inventory Partnership. *Journal of Business Logistics, Vol. 25, No. 1,* 101 – 120

Bearwood J., Halton J. H. and Hammersley J. M. (1959). The shortest path through many points. *Proc. Cambridge Philosoph. Soc.* 55, 299-328.

Campbell, A. M., & Savelsbergh, M. (2004). Inventory Routing in Practice. *The Vehicle Routing Problem,* edited by P. Toth and D. Vigo, SIAM Monographs on Discrete Mathematics and Applications

Daganzo, C.F. (1984) The Distance Traveled to Visit N Points with a Maximum of C stops per Vehicle: An Analytical Model and Application. *Transportation Science, Vol. 18, No. 4,* 331-350.

Hernandez Lopez, J. J. (2003, May). Evaluation of Bulk and Packaged Distribution Strategies in a Specialty Chemical Company. Masters Thesis. MIT

Jain, A.K., Murty, M.N., & Flynn, P.J.(2000) Data Clustering: A Review. ACM Computing Surveys. Vol. 31, No. 3.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, 281-297.

Shen, Victoria (2005). VMI vs. Order Based Fulfillment. Masters Thesis. MIT

Simchi - Levi, D., Kaminsky, P., & Simchi – Levi, E. (2003). *Designing and Managing the Supply Chain. Concepts, Strategies and Case Studies.* Second Edition. McGraw – Hill. New York.