

**DEVELOPING VEHICLE ROUTING AND
OUTBOUND FULFILLMENT SYSTEMS FOR AN E-GROCERY COMPANY**

By

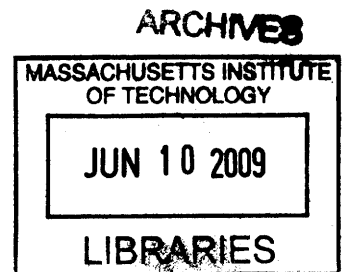
Nicholas Barker

B.S./B.A. 2004, Industrial & Systems Engineering
University of San Diego

Submitted to the MIT Sloan School of Management and the Department of Mechanical Engineering
in Partial Fulfillment of the Requirements for the Degrees of

**Master of Business Administration
AND
Master of Science in Mechanical Engineering**

In conjunction with the Leaders for Manufacturing Program at the
Massachusetts Institute of Technology
June 2009



© 2009 Nicholas Barker. All rights reserved.
The author hereby grants to MIT permission to reproduce
and to distribute publicly paper and electronic copies of this thesis document in whole or in part.

ARCHIVE

Signature of Author _____

MIT Sloan School of Management
MIT Department of Mechanical Engineering
May 8, 2009

Certified by _____

Daniel Whitney, Thesis Supervisor
Senior Research Scientist, CTPID

Certified by _____

Donald Rosenfield, Thesis Supervisor
Senior Lecturer, MIT Sloan School of Management
Director, Leaders for Manufacturing Fellows Program

Accepted by _____

David E. Hardt, Chairman
Committee on Graduate Students, Department of Mechanical Engineering

Accepted by _____

Debbie Berechman, Executive Director of MBA Program
MIT Sloan School of Management

DEVELOPING VEHICLE ROUTING AND
OUTBOUND FULFILLMENT SYSTEMS FOR AN E-GROCERY COMPANY

By

Nicholas Barker

Submitted to the MIT Sloan School of Management and the Department of Mechanical Engineering
on May 8, 2009 in Partial Fulfillment of the Requirements for the Degrees of
Master of Business Administration and Master of Science in Mechanical Engineering

ABSTRACT

This paper outlines areas for improvement within the outbound fulfillment network of an emerging online grocery (“e-grocery”) company offering home delivery to the customer. In particular, the research focuses on developing an efficient, scalable home delivery network, as a result of the known challenges and relatively high fulfillment costs associated with this business model.

Last-mile home delivery accounts for a substantial portion of total e-grocery fulfillment costs. The Vehicle Routing Problem (VRP), a well-known *NP-hard* combinatorial optimization problem, is examined in the context of e-grocery and its impact on last-mile delivery costs. The paper emphasizes an integration of scalable vehicle routing systems with efficient order fulfillment operations.

Practical analytical approaches, as well as new case experiments, serve as a framework of recommendations for an emerging e-grocer or similar last-mile delivery provider. The paper presents analysis using a real case study, serving as a basis for example, as well as more broad recommendations in the field. Moreover, it directs the reader to a wealth of literature in the fields of logistics, grocery fulfillment operations and the VRP class.

Thesis Supervisor: Donald Rosenfield
Title: Senior Lecturer, MIT Sloan School of Management; Director, LFM Fellows Program

Thesis Supervisor: Daniel Whitney
Title: Senior Research Scientist, CTPID

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

ACKNOWLEDGMENTS

I would like to thank the following people:

The faculty and staff of MIT and the Leaders for Manufacturing (LFM) Program
for the opportunity to be part of an incredible learning experience.

My thesis supervisors, Don Rosenfield and Daniel Whitney
for their guidance and support throughout this project.

Our gracious partners at Amazon.com
for sponsoring this project in conjunction with an LFM internship.

My internship supervisor Steven Harman, LFM Class of 1997,
for his mentorship and collaboration on the project.

The AmazonFresh management team,
especially Jim Sharkey, Peter Ham, George Donegan, and Juan Garcia
for sharing their time, resources, and thoughtful ideas.

Finally, a special thanks to my family
for their constant love, support, and guidance over the years.

Nicholas Barker

Cambridge, MA May 2009

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

Table of Contents

Abstract	3
Acknowledgments	5
PART I: Introduction	11
1.1 Purpose	11
1.2 AmazonFresh Background	12
1.3 History of E-Grocery	13
1.4 Outbound Fulfillment and the Vehicle Routing Problem	14
1.5 Thesis Overview	15
PART II: Improving Operations at AmazonFresh	17
2.1 AmazonFresh Background	17
2.2 Project Selection	18
2.2.1 Improvement Opportunities	18
2.2.2 Main Cost Drivers	19
2.2.3 Outbound Fulfillment Operations and Last-Mile Delivery	20
2.2.4 Project Justification	20
2.3 Literature Research	21
2.3.1 The Vehicle Routing Problem (VRP)	21
2.3.2 Narrowing the Problem Class to AmazonFresh	22
2.3.3 Framework for Evaluating Vehicle Routing Solutions	26
2.3.4 Benchmarking VRP Algorithms to Identify State-of-the-Art	26
PART III: Theory to Practice	31
3.1 VRP Systems: To Buy or Develop In-House?	31
3.1.1 Commercially-Available Solutions	31
3.1.2 Advantages/Disadvantages of In-House Development	32
3.1.3 Cost-Benefit Analysis	32
3.2 Implementation Considerations	39
3.2.1 VRP Software Deployment	39
3.2.2 Fulfillment Center	41
3.3 Integration with Fulfillment Center Operations	41
3.3.1 Order Picking and Packing Implications	41
3.3.2 Sorting and Truck Loading	43
3.3.3 Fleet Management	47
3.4 Roadmap for VRP Early-Stage Development	47
3.5 Scalability Considerations, Methods and Techniques	49

3.5.1 Fulfillment Center Capacity	49
3.5.2 Demand Shaping	49
3.5.3 Delivery Process	49
3.5.4 Distribution Network and Fulfillment Center Location	53
3.5.5 Strategic Expansion	56
PART IV: Conclusion	58
4.1 Summary of Findings	58
4.2 Recommendations and Future Steps	59
REFERENCES	60

LIST OF FIGURES

Figure 1: VRPSolver Euclidean Distance Solution.....	36
Figure 2: VRPSolver Road Network Distance Solution	37
Figure 3: Histogram of E-Grocery Orders by Day of Week.....	42
Figure 4: Histogram of E-Grocery Order Patterns by Time of Day	42
Figure 5: E-Grocery Demand Pattern Variability	43
Figure 6: Illustration of VRP Zonal Systems	45
Figure 7: VRP Solutions by Zonal System.....	46
Figure 8: High-Level VRP Process Flow Diagram	48
Figure 9: Truck Capacity vs Theoretical Improvement in Total Distance and # Routes.....	52
Figure 10: Truck Capacity vs Theoretical Fixed and Variable Cost Improvement.....	52
Figure 11: VRP Solution Based on 100% Density Increase	54
Figure 12: VRP-Solved Driving Distance/Stop vs % Density Increase.....	55
Figure 13: VRP-Solved Driving Distance/Stop vs % Density Increase (Capacity Adjusted).....	56

LIST OF TABLES

Table 1: Comparison of AmazonFresh to Peapod	12
Table 2: Common VRP Subtypes.....	15
Table 3: Best Known Solutions for 200-customer VRPTW Benchmark Instances	28
Table 4: Best Known Solutions for 400-customer VRPTW Benchmark Instances	29
Table 5: Commercially-Available VRP Software.....	31
Table 6: Correlation of Possible VRP Cost Metrics	40
Table 7: Impact of Scaling Delivery Capacity on VRP	51

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

PART I: Introduction

1.1 Purpose

The primary objective of this paper is to identify areas for improvement within the outbound fulfillment network of an online grocery (“e-grocery”) company offering home delivery to the customer. In particular, the paper focuses on developing an efficient, scalable home delivery network. It presents a case study of the emerging online grocery company AmazonFresh, along with useful analytical techniques for optimizing its outbound fulfillment and vehicle routing systems.

The potential benefits of home grocery delivery, relative to the traditional model of each individual customer driving to shop at their nearest grocery store, are significant both in terms of greater convenience to the customer and reduced impact on the environment. For example, Siikavirta et al [2003] illustrate the significant potential for reducing greenhouse gas (GHG) emissions through the implementation of e-grocery home delivery strategies. The research indicates that GHG emissions generated by grocery shopping are reduced by 18% to 87% through e-grocery home delivery strategies, compared with the situation in which household members go to the store themselves.

A critical aspect of this increased distribution efficiency, and key driver of total fulfillment costs, is the last-mile delivery problem. Logistics as a whole represent more than an estimated \$700 billion of the US economy annually, or 11% of GNP. De Backer et al. [1997] and Golden and Wasil [1987] estimate that distribution costs account for nearly half of the total logistics costs, and in some industries such as food and beverage, may account for up to 70% of the value added costs of goods. Transportation costs in e-grocery, which is based upon home delivery to the customer, can be particularly high. The paper focuses on vehicle routing based on the significant economic impact associated with the last-mile home delivery model. Moreover, the project that serves as the subject of this research was initially selected by AmazonFresh based on the organization’s need for integrating a scalable vehicle routing system with order fulfillment operations.

The paper aims to provide insight through both analytical techniques, and also the extensive outside literature, applicable to such a project. It presents analysis pertaining to the AmazonFresh case study, as well as more general recommendations in the field. Furthermore, the paper directs the reader to a selection of applicable literature in the fields of logistics, grocery fulfillment operations and the Vehicle Routing Problem (VRP) class.

The basis of this paper primarily lies on the experiential learning from a six-month internship with the e-grocery startup AmazonFresh, in their test market of greater Seattle, Washington. The extended internship experience, made possible through the MIT Leaders for Manufacturing (LFM) program’s unique partnership with companies such as Amazon, is used as a case study in both developing a framework for investigating and optimizing a home delivery network, as well as demonstrating practical approaches at reducing e-grocery fulfillment costs.

1.2 AmazonFresh Background

AmazonFresh (www.amazonfresh.com) is an innovative online grocery delivery service that started its operation in July 2007 as a limited pilot in the Seattle, Washington area. As this operation is still in its early stages, many aspects of its operation have not yet been optimized for efficiency. The scope of this paper centers on ways to optimize various parts of the AmazonFresh fulfillment network, in particular the outbound fulfillment and delivery systems. It also seeks to develop the groundwork for a network of future improvements. AmazonFresh has expressed interest in optimizing its home delivery network vis-à-vis developing a proprietary, scalable, and integrated solution to its particular Vehicle Routing Problem (VRP).

AmazonFresh faces the fundamental challenges of the e-grocery industry yet strives for a particularly high level of customer service. These challenges include delivering relatively low-margin, often perishable and low shipping-density items within narrow time windows.

At the time of this writing, the company provides prompt, free delivery of a wide variety of grocery products to area residences within narrow customer-defined time windows. Table 1 below illustrates a basic comparison of service offerings between AmazonFresh and Peapod, an established e-grocery company.

	AmazonFresh	Peapod
Available Delivery Windows	3hr Unattended, 1hr Attended*	Overlapping 2 & 3.5 hr Attended
Minimum Order Lead Time	4hr	10hr
Delivery Cost	Free**	\$6.95-\$9.95***
Use Existing Grocery Stores	No	Yes + 2 Fulfillment Centers

* “Unattended” orders are delivered to the customer’s doorstep in sealed totes. “Attended” delivery requires customer to be present to accept delivery

** At the time the internship took place, there was no delivery charge for orders above \$30

*** Orders over \$100.00: \$6.95, less than \$100.00: \$9.95, Minimum order amount: \$60.00. Fuel surcharge in some areas

Table 1: Comparison of AmazonFresh to Peapod

Additional challenges in e-grocery include traditionally non-uniform demand patterns¹, consumer price sensitivity, greater quality consciousness, variability of fulfillment process times (e.g. order picking, packing, driving, and delivery service), and escalating fuel prices². Finally, many e-grocery offerings consist of relatively low-margin, perishable, bulky, and/or generally inefficient to transport products.

1.3 History of E-Grocery

The term e-grocery often brings about comparison with the dot-com bubble, as several firms during this period collectively brought the idea of ordering groceries online into the mainstream. Certainly one of the most spectacular failures of this period was Webvan, which spent approximately US\$1.2 billion in its two-year lifespan, ultimately going bankrupt in 2001. Webvan adopted a unique approach of building expensive, highly-automated warehouses to fulfill customer orders. As it turned out, customer demand was simply insufficient to sustain Webvan's rapid growth strategy and high capital investment in the automated infrastructure.

Other e-grocery companies have been more successful, the most notable of which are Peapod, FreshDirect (New York City), and Tesco (United Kingdom). Note that these businesses have generally been constrained to urban centers having relatively high population density. Of course, higher customer density is generally advantageous to last-mile delivery from an operations standpoint, which may explain this trend.

Historically, various e-grocers have employed both "pure play" and "brick-and-mortar" e-grocery business models. The former model relies on proprietary fulfillment centers, while the latter leverages existing grocery store infrastructure and warehousing to some extent. In some cases, e-grocers implement a combination of the two strategies. For example, Peapod leverages a partnership with existing grocery stores in most locations, but also uses two dedicated 75,000 square foot warehouses in its Chicago and Washington, DC regions.

A recent paper [Tong, 2008] characterizes numerous factors influencing the commercial viability of e-grocery, through meta-analysis of both successful and unsuccessful firms, and also provides a detailed historical account of the industry for further reference. The findings suggest key success factors in e-grocery include (1) having knowledge of and experience in the grocery business,

¹ E-grocery demand patterns indicate notable fluctuation by time-of-day and day-of-week. Reference section 3.5.2

² Comparing weekly average price of diesel fuel between 7/14/2007 and 7/14/2008, regional fuel prices increased 62% over the year (<http://www.eia.doe.gov/>)

(2) using a cautious and slow expansion strategy, and (3) leveraging a store-pick model in most markets, with possibly a warehouse-pick model for markets with high customer demand.

1.4 Outbound Fulfillment and the Vehicle Routing Problem

At the core of the last-mile delivery challenge is the Vehicle Routing Problem (VRP), a combinatorial optimization and nonlinear programming problem with the objective of minimizing total transportation cost, subject to serving a number of customers with a fleet of vehicles. Since Dantzig [1959] first proposed the optimization problem, the VRP has remained critical in the field of transportation logistics. Inherent to the problem is the goal of minimizing the cost of distributing goods. Researchers have developed a number of *exact* and *heuristic* (i.e. approximate) solution methods over the years, but for all but the smallest problems, finding the global minimum for the cost function remains computationally complex. Lenstra and Rinnooy Kan [1981] show the underlying combinatorial optimization problem is nondeterministic polynomial time hard (*NP-hard*), implying that there is no known polynomial time exact solution algorithm [Garey and Johnson, 1979]. In essence, as the number of home delivery customers increases, finding optimal delivery routes quickly becomes computationally difficult.

The VRP is of course encountered frequently in industry. In fact the problem class has been widely studied for nearly half a century. Yet the VRP and its many variants remain notoriously difficult to solve in practice. Table 2 below illustrates common subtypes within the problem class.

Abbreviation	VRP Variant
*CVRP	Capacitated VRP. * “C” Generally omitted
VRPTW	VRP with Time Windows
VRPPD	VRP with Pick-up and Deliveries
VRPPDTW	VRPPD with Time Windows
MDVRP	Multiple Depot VRP
MDVRPTW	MDVRP with Time Windows
PVRP	Periodic VRP
PVRPTW	PVRP with Time Windows
SDVRP	Split Delivery VRP
SDVRPTW	SDVRP with Time Windows
TDVRP	Time-Dependent VRP
TDVRPTW	TDVRP with Time Windows
DVRP	Dynamic VRP
SVRP	Stochastic VRP
VRPB	VRP with Backhauls
MVVRPB	Mixed Vehicle VRP with Backhauls
FSMVRP	Fleet Size and Mix VRP
VRPSF	VRP with Satellite Facilities

Table 2: Common VRP Subtypes

As with many e-grocers, AmazonFresh most directly confronts a Vehicle Routing Problem with Time Windows (VRPTW). Section 2.3.2 expands on the VRPTW in greater detail.

1.5 Thesis Overview

The research and development of this paper mainly took place from February 2008 to August 2008, and is in large part a product of the cooperation and collaboration of MIT faculty and Amazon employees.

The thesis is broken down into three basic parts. Part I introduces the e-grocery industry and the VRP in the context of a background for this paper. This section further includes a brief outline of the research behind the thesis.

Part II illustrates approaches and techniques for improving outbound fulfillment operations at e-grocery businesses such as AmazonFresh. This section describes the VRP class in greater detail,

as it pertains to the specific problem confronting AmazonFresh, and further directs the reader to extensive research in the literature.

Part III draws from real experiences during the course of the project and provides the reader with a generalized framework for identifying, selecting and improving upon an e-grocery firm's distribution network. Part III also presents several unique experiments, based on the AmazonFresh case study, which provide insights into real-world vehicle routing and e-grocery fulfillment. These experiments are presented as follows:

- 1) Validating VRP (i.e. computer optimized) versus Manual Routing (Section 3.1.3)
- 2) Determining the Effects of VRP Sub-Problems or Zonal Systems (Section 3.3.2)
- 3) Determining the Effects of Scaling Capacity (Section 3.5.3)
- 4) Determining the Effects of Scaling Density (Section 3.5.4)
- 5) Determining Strategic Expansion Zones by VRP Simulation (Section 3.5.5)

Part IV comprises a summary, conclusion and recommendations for further study.

PART II: Improving Operations at AmazonFresh

2.1 AmazonFresh Background

Amazon.com launched AmazonFresh, an independently-operating subsidiary, in the summer of 2007 as a limited pilot project. At the time of this internship, the e-grocery startup served approximately 1/5th of area neighborhoods in greater Seattle, Washington.

The core value proposition is one of persuading customers, who would traditionally drive their private vehicle to a brick-and-mortar grocery store, to instead shop online and have those same grocery products delivered to their home. Central to this business are three interrelated, key elements of competitive advantage: price, selection, and convenience. An illustration of each of these key dimensions is given below:

1) Price

- The added convenience of home delivery must be high enough to justify the total as-delivered price of groceries. This value proposition may differ among customers, so the e-grocery company may be well advised to examine price elasticity and, more generally, to offer the lowest pricing that is economically feasible.

2) Selection

- Customers prefer a broad selection of products. To the extent that the variety of products offered online meets or exceeds a customer's expectation, he or she is more likely to "convert" from traditional grocery shopping to e-grocery.

3) Convenience:

- The typical experience of shopping at a brick-and-mortar grocery store is time consuming and relatively inefficient in terms of transportation logistics. While some individuals enjoy the experience of physically browsing the aisles of their favorite grocery store and hand-selecting perishable goods, others value the convenience of online shopping. For example, certain professionals and busy parents may find the grocery shopping excursion burdensome. Key to this element of convenience is offering prompt delivery within customer-defined time windows, while still ensuring that perishable goods will be of high quality and "picked" to their specification.

Interestingly, e-grocery businesses such as AmazonFresh may offer two additional competitive elements, which are less feasible among traditional brick-and-mortar stores:

4) Information

- Online shopping has changed the face of retail, but not in the way we once expected. Only about 3% of retail is purchased online, but the effect has been much more informed shoppers and lower prices everywhere³. Examples of improved information access through e-grocery include “smart” or tailored shopping lists, customer product reviews, extended product information, greater traceability to perishables, and various health/RDA data.

5) Discovery

- Through online services and intelligent data mining, it is possible to delight the customer with something they did not expect. For example, e-commerce companies such as Amazon.com strategically use online marketing tactics such as “Have you seen..., Did you forget..., People also like...” to promote sales. Of particular note with such techniques is the significant potential for cross-selling various products, making it easier for the customer to purchase something they otherwise would not have found.

Note that the competitive criteria above are not independent, but rather interrelated. In particular, the key areas of price, selection, and convenience may be viewed as a triangle of linked criteria. At least in theory, added convenience may be offered with less selection, or a higher price, and so on. For additional reference, see a detailed study into the competitive aspects of various e-grocers in Tong [2008].

2.2 Project Selection

2.2.1 Improvement Opportunities

Despite the attractive qualities of e-grocery to the consumer, the business itself presents a number of operational challenges. Providing home delivery of perishable and non-perishable grocery products is notoriously difficult to execute in a cost-effective manner. Evidence of this can be seen in the failure of such e-grocery firms as Webvan, a company which adopted an approach of

³ Online purchases comprise 3.4% of all retail purchases. US Census Data, Q4 2008: (<http://www.census.gov/mrts/www/data/pdf/08Q4.pdf>)

elaborate automation in order to minimize fulfillment costs. The grocery industry itself is challenging, historically being characterized by razor-thin profit margins in the low single-digit percent range.

Moreover, the improvement opportunities at a startup company are generally more numerous, relative to opportunities at more established companies. E-grocery operations may utilize existing brick-and-mortar grocery store infrastructure⁴ or a dedicated fulfillment center (FC). AmazonFresh uses a proprietary warehouse, divided into three temperature zones: ambient, chilled, and frozen. The fulfillment operations can be categorized by tracing the flow of products chronologically through the FC as follows:

1) Inbound Operations

- Example: Purchasing, receiving, stocking, managing first-line Quality Assurance.

2) Warehouse Fulfillment Operations

- Example: FC capacity planning, order picking, packing, sorting, managing Quality Assurance and product shrink (i.e. loss, theft, or expiration of products).

3) Outbound Fulfillment Operations

- Example: Delivery capacity planning, vehicle routing, final customer order sorting, truck loading, navigation, delivery, empty tote pickup.

In order to maintain a reasonable scope, this paper emphasizes outbound fulfillment operations, with particular attention devoted to developing vehicle routing systems. The project selection process ultimately converged on developing the home delivery network because of the overwhelming extent to which home delivery impacts total fulfillment costs.

2.2.2 Main Cost Drivers

The main cost drivers of outbound fulfillment are the fixed and semi-fixed costs associated with the fleet of vehicles (e.g. lease or purchase cost of each vehicle plus insurance), as well as the following variable costs:

1) FC associate labor

- Includes the final sorting of customer orders, and loading of delivery trucks.

⁴ Peapod utilizes both dedicated FC's and Stop-and-Shop grocery stores, depending on location

2) Delivery driver labor

- Includes preparation for driving the route, performing vehicle safety checks, navigating the vehicle to customer addresses, and serving the customer at their doorstep.

3) Delivery truck expenses

- Variable costs attributable to operating a delivery vehicle, including consumables such as fuel, oil, tires, and maintenance items.

Naturally we seek to minimize the number of delivery vehicles in the fleet, subject to an overarching constraint of having enough vehicles and drivers to satisfy peak demand. We also seek to minimize the truck loading time, delivery time and driving distance associated with each delivery route. The routine problem constraints that must be met are: each customer being served once, with an on-time delivery according to a customer-specified time window, and by a delivery vehicle that cannot be loaded beyond its given capacity.

2.2.3 Outbound Fulfillment Operations and Last-Mile Delivery

In addition to characterizing AmazonFresh's particular delivery problem, it is important to understand the interaction between the FC fulfillment processes (e.g. inbound receiving, stocking, inventory management, capacity planning, order picking, packing, tote sorting) and outbound processes (i.e. vehicle route planning, truck loading, home delivery). Of particular importance are daily patterns of order checkout, FC picking, sorting, and outbound truck loading processes. For instance, if the final FC sorting step and subsequent truck loading cannot be postponed until after the customer order placement deadline (i.e. as a result of FC process cycle times) then the underlying VRP would become, to some degree, stochastic rather than deterministic. The resulting problem formulation would have significant implications in terms of problem complexity and approach. Fortunately it may be possible to avoid a Stochastic VRP (SVRP), by adding a preliminary order sorting step⁵. Practitioners generally prefer deterministic formulations where possible because they tend to be relatively more straightforward and robust than the stochastic variant.

2.2.4 Project Justification

The selection of vehicle routing and outbound fulfillment as a research topic is based on this component's relatively high contribution to total fulfillment costs in e-grocery. Refer to section 3.1.3

⁵ Based on preliminary manual sorting time study data

for experimental evidence suggesting that the implementation of vehicle routing systems is indeed justified.

2.3 Literature Research

2.3.1 The Vehicle Routing Problem (VRP)

The Vehicle Routing Problem (VRP), one of the most studied combinatorial optimization problems, aims to determine the best routes for pickup and/or delivery of goods in a distribution system. In the classical VRP a number of capacity-constrained vehicles located at a central depot must serve a set of geographically-dispersed customers. Each customer has a given demand and each vehicle has a given capacity. The objective is to minimize the total cost (i.e. distance or time) of travel. First proposed by Dantzig [1959] the VRP has been the subject of extensive research for approximately half a century. Interest in the VRP has been fueled by its inherent complexity, as well as the frequent occurrence of the problem in industry and the extent to which efficient transportation impacts the bottom line of businesses.

Not only is the VRP a common and important problem, but it is also notoriously difficult to solve in practice. Recall that the VRP is well known to be an *NP-hard* combinatorial optimization problem [Lenstra and Rinnooy Kan, 1981]. Problems of a size encountered in real world situations are generally approached *heuristically*, as it is prohibitive to solve *exactly* in cases where the problem size is larger than approximately $n=100$ nodes. The VRP is generally formulated as a mixed integer programming (MIP) model, with integer variables associated with each arc between locations, termed the Vehicle Flow Model [Bodin et al., 1983].

In modern practice, the most efficient approach to larger VRP's generally uses one of the more recently developed metaheuristics⁶, selected according to the attributes and constraints associated with a unique problem. Laporte (2007) provides a survey of literature outlining the state-of-the-art in the classical VRP.

In the case of unattended delivery at AmazonFresh, the classical capacity-constrained VRP is sufficient as a basic model. AmazonFresh must serve these n customer orders, in no particular order of precedence, within a 3-hour period. Note that in the unattended delivery scenario, the vehicle

⁶ Metaheuristics are high-level algorithmic frameworks or approaches for optimization problems, which often combine other heuristics in the search for feasible solutions. Some common approaches include Genetic Algorithms, Simulated Annealing, Tabu Search, Local Search, and Ant Colony Optimization

capacity is sometimes constrained by time (i.e. the number of grocery deliveries a driver can make within a 3-hour window) and sometimes the physical capacity of the vehicle. AmazonFresh sets a nominal truck capacity such that each driver may readily deliver the truck's manifest of orders within the 3-hour period (adjusting for variables in driver performance and route difficulty), subject to the number of totes assigned to a route being within the truck's physical capacity. The truck capacity may equal the number of deliveries that a driver can make within a 3-hour span, allowing for a margin of safety, or otherwise the actual physical capacity of the vehicle. Therefore, the degree to which the drivers can reduce their average driving and service times (i.e. increase stops per hour) may improve the truck capacity in some instances. Section 3.5.3 presents an experiment that illustrates significant cost reducing effects following an improvement in truck capacity.

Also note that the 3-hour unattended delivery window is not optimized in terms of truck capacity utilization, but rather is driven by other factors such as sales, marketing, and ensuring the integrity of temperature-sensitive grocery items. To the extent that the 3-hour window can be relaxed, a higher capacity utilization per truck is possible, up to the point where the driver's nominal stops-per-hour performance equates to the physical capacity constraint of the truck. For example, assuming that a single truck could hold 24 orders and could theoretically conduct 15 stops in a 3-hour window, as given in Section 3.5.3, this implies an optimal unattended delivery "time window" of 4.8 hours in order to maximize the physical capacity utilization of the delivery truck. Toth and Vigo [2001] present a thorough review of the VRP for additional reference.

2.3.2 Narrowing the Problem Class to AmazonFresh

In practice, the unattended deliveries having a large delivery window may be mixed with narrower 1-hour attended deliveries on the same route. The core problem then becomes a Vehicle Routing Problem with Time Windows (VRPTW). In this problem variant, the solution must fulfill each delivery node's demand within a certain time constraint. The overarching decisions within this VRPTW are (1) assigning orders to vehicles (2) routing vehicles to customer addresses, and (3) scheduling to satisfy demand within the time constraints promised with each order.

Moreover, the VRPTW is similar the classical capacity-constrained VRP, except the delivery addresses have time window constraints within which the deliveries must be made. Refer to Bräysy and Gendreau [2005a, 2005b] and Toth and Vigo [2001] for an in-depth treatment of the VRPTW.

Note there is often a tradeoff between customer service level and cost. The relative optimality of a set of deliveries is generally, from a cost standpoint (i.e. distance or time), negatively affected by the additional complexity of time windows. To the extent that a customer is promised a narrow delivery window, the distance-only VRP may be subverted to meet the time window promise. Managing this tradeoff is one of the keys to last mile delivery businesses such as AmazonFresh. Effective means to mitigate this tradeoff may include sales and marketing tactics to facilitate demand shaping. For example, one may offer customers an incentive to accept a particular or simply broader delivery window through rebates, promotions, or even appeals to environmental responsibility. Peapod employs a rebate strategy targeting customer acceptance of more broad delivery windows⁷. Beyond sales and/or demand-side techniques, an efficient routing system is critical to facilitating the delivery network.

In addition to the basic VRPTW, there are other extensions of the VRP that may be relevant to e-grocers such as AmazonFresh. First, the Dynamic VRP (DVRP) is applicable in the event that information needed to design a set of routes is revealed dynamically to the system. Typically these dynamic inputs comprise new customers, demand levels, vehicle status updates or traffic delays. Note that each of these inputs is part of an urban e-grocery business to some extent, and will likely need to be addressed as the business develops. Fortunately, recent technological advancements in GPS, mobile devices, and real-time traffic data have made this more feasible. Unfortunately, there are relatively unavoidable obstacles to the DVRP as well, such as delivery truck design. For instance, if delivery trucks utilize a standard last in, first out (LIFO) packing configuration, then subsequent dynamic VRP solutions may be infeasible in terms of either physical loading or route modification. Examples of dynamic inputs found in the e-grocery model are illustrated as follows:

- New Orders
 - Example: Customer orders arrive throughout the day. It is beneficial from an operations standpoint to level the fulfillment tasks in a reasonable manner. To the extent that AmazonFresh can accelerate operations at the FC, that is by not having to wait until all customer orders have arrived (i.e. “cut-off”), the fulfillment workload is less confined to a short turnaround time. Associates can accomplish tasks such as preliminary sorting of orders and even truck loading earlier, thereby increasing the

⁷ www.peapod.com

likelihood of an on-time truck departure and facilitating prompt service to customers.

- New Customers

- Example: Dispatchers may send delivery trucks already mid-route to new customers. E-grocery differs from other last-mile providers such as courier services in that it generally entails *delivery only* rather than *pickup-and-delivery*. However, e-grocers may deliver in reusable “totes” which must be picked up. According to AmazonFresh, customers frequently request prompt pickup of the empty totes⁸. As long as the empty tote pickup service is offered to customers it is essentially a non-value added “legacy cost” of the initial order. Therefore, to the extent that vehicle routing can include tote pickups dynamically (on a selective basis), these “legacy costs” may be reduced. Inserting a tote pickup into an existing vehicle route may not be the optimal way to handle impromptu tote pickups, but in the case where such a stop is especially close or “on the way” to next stop, it could make sense. Note the slight distinction between this and the standard PDVRP, which has a physical capacity constraint embedded in the pickup. The empty tote pickup constitutes a route cost and service time cost, but because the empty totes stack in a nested fashion, the physical capacity component is negligible.

- Demand levels

- Example: Customers may place an initial order, and then subsequently amend the order prior to delivery. In essence the subsequent order merges with the initial order, assuming the customer address and scheduled delivery time are the same. Therefore, the delivery address is already known in the VRP, but the demand has increased to the extent of the secondary order. The somewhat longer service time associated with the larger order may also be taken into account by the model, in the case where time window constraints are imposed (i.e. VRPTW). A more simple solution is to re-run the standard VRP at a few key intervals, in such case where it is unimportant to react immediately to a change in demand.

⁸ Based on interviews with AmazonFresh managers

- Vehicle status updates
 - Example 1: The e-grocer maintains its fleet of delivery vehicles in order to ensure minimal downtime. However, occasionally the vehicles may “break down” as a result of mechanical issues, a flat tire, or dead battery. The set of orders allocated to any such vehicle, assuming the vehicle is observed to have become disabled prior to departing the depot, should then be re-routed among the fleet as expediently and efficiently as possible. Because a spare vehicle may not be available, the fleet manager may set capacity utilization to account for the possibility of a disabled vehicle within the fleet.
 - Example 2: As in the case of tote pickups, such customer requests may occur at any time including when the delivery vehicles are mid-route. Traditionally, a transportation dispatcher would communicate with drivers on the road via radio or phone, requesting the driver in closest proximity to make the pickup. Note this method is not necessarily optimal, and requires a human dispatcher to be employed. In a delivery network employing modern GPS-based fleet management technology, each truck sends location data wirelessly, in real time, to a central dispatching system. The location data for each truck feeds into the DVRP model to re-route a particular truck automatically.
- Traffic delays
 - Traffic congestion represents a major threat to last-mile delivery in terms of ensuring low cost and a high level of customer service. Delays in service not only affect the immediate customer, but also the FC operators, who rely on the delivery truck returning on time so that it can be loaded for the next route. Moreover, traffic delays are of particular concern to e-grocery businesses, which tend to be centered in urban areas, and are therefore subject to heavier traffic densities. Real-time traffic data is available for many urban areas via the Web and radio broadcasts. This data can be used to update the VRP dynamically, as unexpected traffic delays occur. Such unexpected delays may occur as a result of traffic accidents, disabled vehicles, or construction projects. Similarly, delays recurring somewhat routinely could also be built into the model.

In addition to the DVRP case, an extension of the *deterministic* VRP/VRPTW to the *probabilistic* case may be appropriate for an e-grocery company. The term Stochastic Vehicle Routing Problem (SVRP) describes a number of cases where at least one of the inputs is variable. Also, while the classical VRP generally assumes a *homogeneous* fleet of vehicles, and indeed historically this is the case of many e-grocery companies, it is reasonable to assume a *heterogeneous* case is possible going forward. As the urban-centered e-grocery business develops into new markets, it stands to reason that a hub-and-spoke distribution system with same-size vehicles may not necessarily be ideal. Golden [2008] provides a more comprehensive treatment of the literature and recent developments in the VRP class.

2.3.3 Framework for Evaluating Vehicle Routing Solutions

Clearly the VRP variant most directly relevant to the AmazonFresh business model is the VRPTW. As long as customers are able to specify a delivery time window, within which they expect to receive their order, the time window constraints are a key component of the model.

Assuming VRPTW as the base problem, the e-grocer must then decide on a solution approach comprising either (1) purchasing commercial vehicle routing software, or (2) developing a proprietary software solution. Section 3.1 illustrates some tradeoffs between these two approaches. The micro-level criteria for evaluating any solution comprise the performance or solution quality (i.e. relative to best known solution, least distance/time cost, fewest vehicles, and computation time) afforded by the system. The macro-level criteria for evaluating a VRP system include the return on investment (ROI), ease of implementation, robustness/risks, and forward compatibility as the business develops.

If the e-grocer chooses to develop a propriety solution, careful consideration should be given to the extensive body of research already available, so as to leverage this valuable research and avoid unnecessary effort. The following section compares the VRPTW unique to AmazonFresh with analogous benchmark instances. Through this process, e-grocers may identify the best-known solution approaches to VRPTW's which most closely resemble the actual real-world problems they face.

2.3.4 Benchmarking VRP Algorithms to Identify State-of-the-Art

As we have seen the VRPTW is an *NP-hard* combinatorial optimization problem. So how do we identify the best methodology to solve VRPTW? Fortunately, the VRP and its variants have been

widely researched for nearly fifty years. Many algorithms and heuristics developed to solve these problems have been tested on standardized benchmark problems, as an effective means for comparison.

Literature suggests a strong trend toward metaheuristic methods as a result of their generally superior performance. Fortunately researchers have pitted the many solution methods developed over time against a set of benchmark problem instances⁹. The specific problem type at an e-grocer can be related closely to specific benchmark problem instances¹⁰. For example, if we assume that the problem confronting AmazonFresh is of a size approximated by $n=200$ or $n=400$ nodes, served by roughly 10-15 delivery trucks, this directs us to the extended VRPTW benchmark instances by Homberger. The problem size represents a set of customers receiving home deliveries from one FC on given shift. Within this problem size, we can further approximate specific benchmark problems by observing that the geographical distribution of customers in greater Seattle is neither solely R-type (“Random” or uniformly-distributed) nor C-type (“Clustered”), but rather RC-type (a combination of both)¹¹. The C1, R1, and RC1 problem class is based on a shorter scheduling horizon, characteristic of VRP instances having relatively many vehicles with small capacities. In contrast, the C2, R2 and RC2 problem class is based on a longer scheduling horizon, characteristic of longer delivery routes with fewer vehicles. AmazonFresh serves customers on a planning horizon somewhere in between the short and long extremes depicted in the benchmark instances, although most near the RC1 class¹². Thus, we can deduce what may be the best performing metaheuristic methods for a real-world vehicle routing scenario based on the documented performance these algorithms on analogous benchmark problems. The best solution approaches, consistent with the VRPTW parameters described above, are illustrated in Tables 3-4.¹³

⁹ In this case, benchmark VRPTW instances from Solomon ($n=100$) and extended by Homberger ($n=200, 400, 600, 800, 1000$). Available 4/1/2009 <http://www.top.sintef.no/vrp/benchmarks.html>

¹⁰ Analogous instances generally within Solomon Random-Clustered class, having short horizon (i.e. “RC2” class problems)

¹¹ Designated “RC” class in the benchmark instances to characterize networks of partial uniformity and partial clustering

¹² Designated “RC1” and “RC2” class, respectively, in the benchmark instances

¹³ <http://www.top.sintef.no/vrp/bknown2.html> and <http://www.top.sintef.no/vrp/bknown4.html>. Available 4/1/2009

Benchmark Case	# Vehicles	Best Known Solution (Distance)	Authors	Date of Best Solution
rc1_2_1	18	3602.80	PGDR	17-oct-07
rc1_2_2	18	3249.50	PGDR	17-oct-07
rc1_2_3	18	3008.33	BSJ2	20-sep-07
rc1_2_4	18	2852.62	MB	10-may-05
rc1_2_5	18	3385.88	PGDR	17-oct-07
rc1_2_6	18	3324.80	PGDR	17-oct-07
rc1_2_7	18	3189.32	BSJ2	20-sep-07
rc1_2_8	18	3083.93	PGDR	17-oct-07
rc1_2_9	18	3083.41	MB	16-sep-03
rc1_210	18	3008.53	PGDR	17-oct-07
rc2_2_1	6	3099.53	BSJ2	20-sep-07
rc2_2_2	5	2825.24	PGDR	17-oct-07
rc2_2_3	4	2603.83	PGDR	17-oct-07
rc2_2_4	4	2043.05	MB	10-may-05
rc2_2_5	4	2911.46	BSJ2	20-sep-07
rc2_2_6	4	2975.13	RP	25-feb-05
rc2_2_7	4	2525.83	BSJ2	20-sep-07
rc2_2_8	4	2293.35	PGDR	17-oct-07
rc2_2_9	4	2175.04	BSJ2	20-sep-07
rc2_210	4	2015.60	MB	16-sep-03

Table 3: Best Known Solutions for 200-customer VRPTW Benchmark Instances

Benchmark Case	# Vehicles	Best Known Solution (Distance)	Authors	Date of Best Solution
rc1_4_1	36	8630.94	PGDR	17-oct-07
rc1_4_2	36	7958.67	PGDR	17-oct-07
rc1_4_3	36	7562.60	PGDR	17-oct-07
rc1_4_4	36	7332.59	PGDR	17-oct-07
rc1_4_5	36	8249.63	PGDR	17-oct-07
rc1_4_6	36	8223.12	PGDR	17-oct-07
rc1_4_7	36	8001.12	PGDR	17-oct-07
rc1_4_8	36	7836.29	PGDR	17-oct-07
rc1_4_9	36	7811.55	PGDR	17-oct-07
rc1_410	36	7668.77	PGDR	17-oct-07
rc2_4_1	11	6688.31	PGDR	17-oct-07
rc2_4_2	9	6355.59	RP	25-feb-05
rc2_4_3	8	4958.74	PGDR	17-oct-07
rc2_4_4	8	3635.04	MB	30-jun-05
rc2_4_5	9	5923.95	PGDR	17-oct-07
rc2_4_6	8	5863.56	PGDR	17-oct-07
rc2_4_7	8	5466.70	PGDR	17-oct-07
rc2_4_8	8	4848.87	PGDR	17-oct-07
rc2_4_9	8	4599.57	MB	30-jun-05
rc2_410	8	4311.59	PGDR	17-oct-07

Table 4: Best Known Solutions for 400-customer VRPTW Benchmark Instances

Clearly PGDR¹⁴, MB¹⁵, RP¹⁶, and BSJ2¹⁷ are high-performing metaheuristic approaches that should be considered by the e-grocery company. Yet another promising metaheuristic approach for the VRPTW confronting the e-grocer is the Multiple Ant Colony System (MACS-VRPTW) [Gambardella, 1999] based on Ant Colony Optimization (ACO). Broadly speaking, ACO is a class of metaheuristics falling within the concept of "swarm intelligence," inspired the collective the behavior

¹⁴ **PGDR** - Eric Prescott-Gagnon, Guy Desaulniers and Louis-Martin Rousseau. A Branch-and-Price-Based Large Neighborhood Search Algorithm for the Vehicle Routing Problem with Time Windows. (2007)

¹⁵ **MB** - Mester, D. and O. Bräysy (2005), "Active Guided Evolution Strategies for Large Scale Vehicle Routing Problems with Time Windows". *Computers & Operations Research* 32, 1593-1614.

¹⁶ **RP** S. Ropke & D.Pisinger. "A general heuristic for vehicle routing problems", technical report, Department of Computer Science, University of Copenhagen.

¹⁷ **BSJ2** - Bjørn Sigurd Johansen, DSolver version2 05-2005.

of social insects seeking self-organization in biological systems. See Dorigo and Stutzle [2004] for more background on ACO, notably chapter 5, which provides a thorough treatment of the application of ACO to VRP including the MACS-VRPTW algorithm recommended above.

PART III: Theory to Practice

3.1 VRP Systems: To Buy or Develop In-House?

The literature provides a wealth of research into state-of-the-art VRP approaches. For a company such as Amazon, reputed for achieving competitive advantage through its information technology and software development capabilities, the option of developing a proprietary, flexible, state-of-the-art solution is feasible. In many cases the average last-mile delivery company would be better served by purchasing one of several commercial VRP solutions. This section illustrates criteria for making a decision between purchasing a commercial vehicle routing system and developing a proprietary system, as well as highlights some available options.

3.1.1 Commercially-Available Solutions

Several commercial VRP software solutions are available. Table 5 below depicts these products as of 2008. An excellent survey of these commercial solutions is also available¹⁸.

Product	Vendor	Year Introduced
Descartes Routing & Scheduling	The Descartes SystemCGroup	-
Direct Route	Applan Logistics Software, Inc.	1996
DISC	MJC2 Limited	1990
ILOG Dispatcher	ILOG, Inc.	1997
JOpt.SDK	DNA Evolutions	-
Optrak4 Vehicle Routing & Scheduling	Optrak Distribution Software	2002
ORTEC Routing and Scheduling	ORTEC	1981
Paragon Routing and Scheduling System	Paragon Software Systems, Inc.	1997
Prophesy Total Transportation System	Prophesy Transportation Solutions	1999
PTM Pro Online, Pupil Transportation Manager	Spatial Decisions Support Systems	2003
REACT	MJC2 Limited	1990
Roadnet Anywhere	UPS Logistics Technologies	2006
Roadnet Transportation Suite	UPS Logistics Technologies	1983
STARS 5.0	SAITECH, Inc.	1995
StreetSync Desktop	RouteSolutions, Inc.	2005
The LogiX Suite	Distribution Planning Software Limited	1985
TourSolver for MapPoint / MapInfo Pro	Magellan Ingenierie	2002
TruckStops Routing and Scheduling Software	MicroAnalytics	1984

Table 5: Commercially-Available VRP Software

¹⁸ Commercial products on the market (2008): http://www.lionhrtpub.com/orms/surveys/Vehicle_Routing/vrss.html

Note that generally the commercial software providers do not disclose detailed information about the algorithms used in their VRP applications. The following section outlines the advantages and disadvantages of commercial VRP software, relative to developing a proprietary solution.

3.1.2 Advantages/Disadvantages of In-House Development

Not surprisingly, because several relatively versatile software solutions exist, most companies facing VRP's choose to implement an off-the-shelf vehicle routing software package. The rationale behind implementing a commercial solution often includes: relative ease of implementation, faster deployment, availability of suitable solution, lack of in-house software development or IT resources, or otherwise a lack of justifiable need to build a customized VRP solution. The disadvantages to deploying a commercial package typically include the high initial cost¹⁹, dependency on third-party support, relinquishment of potentially sensitive data to a third party, limitation of basic algorithms, and relative inflexibility to adapt to dynamic business needs.

An e-grocery company such as AmazonFresh, having substantial experience in software development, may have interest in developing a proprietary VRP solution in order to mitigate the aforementioned risks, to provide a foundation upon which to continuously improve, and to afford the greatest overall flexibility and scalability potential. To that end, such an undertaking requires a substantial commitment of resources, and should be weighed heavily from a cost-benefit standpoint on an individual project basis.

3.1.3 Cost-Benefit Analysis

First consider the benefit of a vehicle routing system relative to traditional manual routing practices. Traditionally, manual routing processes have constituted a dispatcher or driver visualizing customer addresses, and subsequently allocating and/or sequencing the routes according to their own intuition. Such a manual process is perhaps suitable for infrequent, small problem instances. But with a few hundred customers per delivery cycle, the problem size soon becomes unmanageable. These individuals must have intimate knowledge of the road networks. Even still, the manual process is time consuming, risk prone, almost certainly suboptimal to some degree, and not particularly scalable. Note that given the absence of VRP software for comparison, it was not possible to quantify the relative sub-optimality associated with manual routing. However, research indicates that computer-optimized transportation routing may yield savings ranging from 5% to

¹⁹ Software costs generally range from \$10K to over \$100K (OR/MS Today, Vehicle Routing Software Survey, February 2008)

20%²⁰. An interesting validation experiment would be to compare a sample of the manually-determined routes with the VRP (globally-optimized) equivalent, on the dimension of total time or distance. Then the ROI of a VRP implementation may be calculated through the theoretical reduction of transportation costs and observed reduction in manual routing labor.

In order to carry out the validation experiment, we utilize a standalone implementation of Microsoft MapPoint²¹ in combination with an academic VRP application in Java available for non-commercial use. This MapPoint software is well suited for the experiment based on its performance characteristics, low cost, and ease of use. A simple address report tool²² that generates a source file of specified delivery addresses is then integrated with the system. MapPoint does not provide VRP capability, but does provide some notable capabilities. First, it provides the requisite geocoding functionality. Geocoding is the process of finding geographic coordinates (i.e. latitude and longitude) from other geographic data, such as customer street addresses. These geographical location coordinates are important because the associated origin-destination (O-D) cost matrix data (e.g. point-to-point distances between all customers) constitutes the foundation of the VRP. In addition to geocoding, MapPoint provides access to high quality commercial map data, and a built-in shortest-path solver for predetermined route sequences along an actual road network. It further includes a built-in optimization function that solves the Travelling Salesman Problem (TSP). The TSP may be viewed as similar to VRP except that it seeks to minimize the total cost of a single route rather than globally for multiple routes. Thus, the TSP is relatively more simplistic than the VRP.

In combination with the mapping package, we leverage an academic VRP system available for non-commercial use. The Java-based VRP solver implementation by Snyder²³ is used for solving the VRP as part of this experiment. The software is based on the well known Clark-Wright savings algorithm [Clarke and Wright, 1964] with various improvement heuristics. Illustrated below is the experimental process. The following experiment illustrates the value of VRP software generically, whether a commercial package or developed in house, relative to manual routing processes.

²⁰ P. Toth, D. Vigo: "The Vehicle Routing Problem". Monographs on Discrete Mathematics and Applications. SIAM, Philadelphia. 2001

²¹ MS MapPoint software was installed on dedicated laptop for experiments

²² Address Report Tool generates a user-defined .csv source file from Data Warehouse based on parameters date, time, and delivery type

²³ The author thanks Professor Lawrence Snyder, Department of ISyE, Lehigh University.
<http://www.lehigh.edu/~lvs2/software.html>

EXPERIMENT: Validating VRP (i.e. computer optimized) versus Manual Routing

Summary:

First, the manually-routed theoretical driving distance (i.e. manually allocated and sequenced routes, but computer-solved for shortest path along road infrastructure) for a typical unattended route is **25%** greater than that of “optimal” VRP solution. It is important to note that this experiment compares theoretical VRP-solved total route cost to that of the manually-allocated and sequenced (but still theoretical shortest path-solved along the road network, using MapPoint) total route cost. Therefore the route is not actually driven, but instead the experiment represents a fair comparison of VRP technology to human intuition. A skilled AmazonFresh dispatcher or “lead driver” generates the manual routing data as is typical for the given route. In this case, the dispatcher takes approximately one hour to complete the manual routing. In contrast the VRP takes approximately nine seconds to solve. As a caveat, the 212 X 212 O-D matrix takes nearly three hours to compute using actual road network data(in MapPoint), likely because of the significant computational cost associated with the interpolating accurate road distance for the large matrix. We address this matter further in subsequent sections. The following illustrates the VRP validation experiment in detail.

Assumptions:

- 1) VRP uses only sample unattended orders (i.e. no Time Window constraints)
- 2) Capacity constraint was set at a nominal 15 orders/truck.

Steps:

- 1) Download sample unattended customer orders from data warehouse via SQL query.
- 2) Import .csv file of these addresses into MS MapPoint. Generate geocodes (latitude, longitude) and Origin-Destination (O-D) cost matrices using third-party add-on utility MileCharter²⁴, with the MapPoint software.
- 3) For testing purposes, repeat step 2 to generate O-D matrices based on Euclidean (straight line) distance, actual road distance, and actual road driving time. Including the central depot, these matrices are 212 x 212, comprising 44944 O-D “costs”. Note the actual road network distances and times are provided asymmetrically (i.e. distance A-B is not necessarily equal to B-A). Computation runtime for generating the actual road distance O-D matrix is significant, in this case 2 hours and 50 minutes. The lengthy computation time is

²⁴ Ref Winwaed Software Technology <http://www.mapping-tools.com/milecharter/index.shtml>

attributable to the extensive interpolation required along the road network for each of the thousands of O-D data points. Note that once an O-D data point is calculated it may be stored, thereby mitigating computational expense in the case of repeat customers. Euclidean (straight line) distances are calculated much more quickly, and naturally symmetric. Output to an .xls spreadsheet.

- 4) Convert latitude-longitude geocodes into Cartesian coordinates relative to the origin depot. Units are distance given in miles. Output a tab-delimited .txt file with fields X, Y, and Demand for each of $n=212$ addresses. Set demand to 1 for each stop, and assign a 15 order/truck nominal capacity constraint. Set demand for depot to 0.
- 5) Import .txt file per step 5 into VRP software. This allows visualization of the addresses. The software interprets cost matrices internally by Euclidean (straight line) distance, Great Circle (on a sphere) distance, or via uploading a .txt file with an auxiliary O-D cost matrix (e.g. actual road distance/time as per steps 3-4).
- 6) Set capacity constraint to 15 orders/truck and maximum route distance constraint²⁵ arbitrarily high (e.g. 1000 miles).
- 7) First solve the VRP using Euclidean distance metric. Figure 1 depicts an optimal solution with objective value (min total distance) of 333.61 miles, using 15 routes, each having a max capacity of 15 orders. This sample solution takes roughly nine seconds to compute including construction (Clark-Wright) and improvement heuristics. See graphical result below:

²⁵ Note the max route cost constraint is more useful when using Actual Road Driving Time matrix (i.e. 3.5 hours).

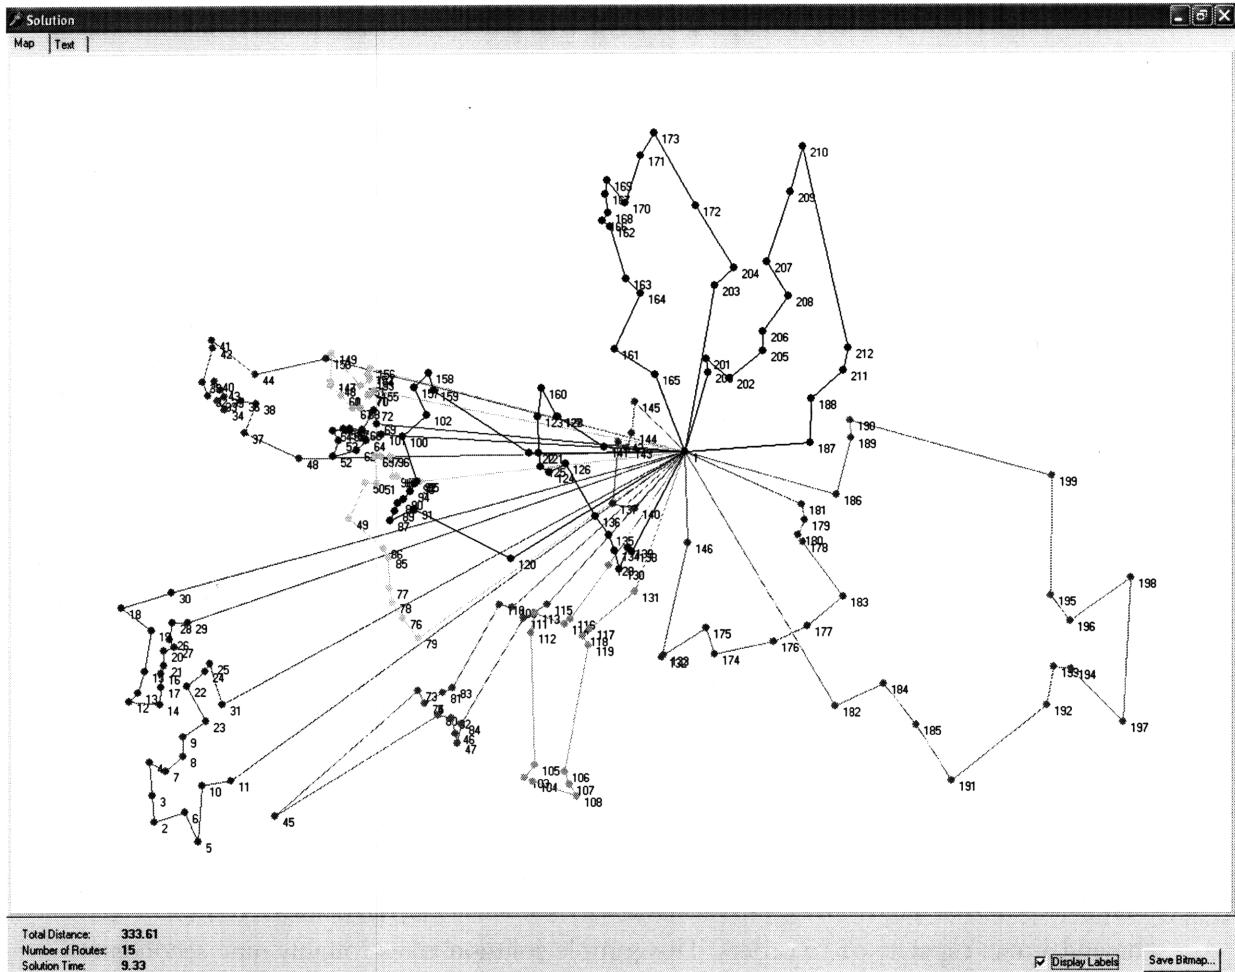


Figure 1: VRPSolver Euclidean Distance Solution

8) In order to solve the same problem using actual road distance, import the auxiliary cost matrix .txt file generated per steps 3-4. Note the road distances were converted from asymmetric to symmetric in order to interface properly with VRPSolver. This is relatively insignificant, as an additional study in Section 3.2.1 shows a high 0.996 correlation between actual road distance from A-B and B-A. Solving the VRP, note the optimal solution using actual road network data has an objective value (min total distance) of **415.70** miles, using 15 routes, each having a max capacity of **15** orders. This computation time is **9.22** seconds. Also observe the greater visual overlap of routes, attributable to using actual road network cost data. Figure 2 illustrates the VRP solution below.

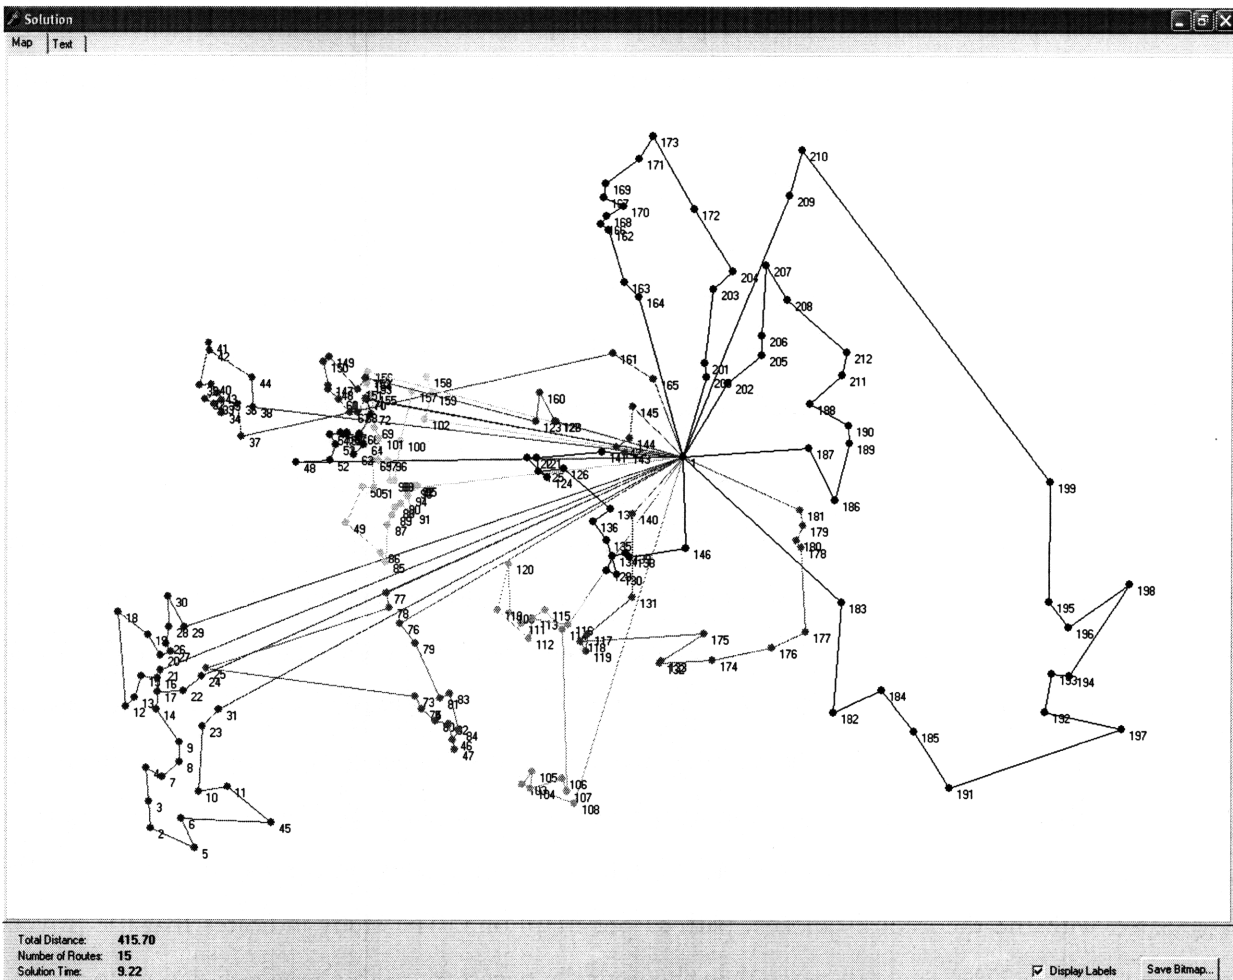


Figure 2: VRPSolver Road Network Distance Solution

9) Extract the total theoretical distance travelled based on legacy manual routing process, using manual route allocations and sequences as the input into MapPoint. For each manually-allocated route, input the sequence of addresses and MapPoint will calculate the shortest path distance subject to maintaining the predetermined sequence. The sum of these theoretical distances for each manually-allocated route is the basis for comparison with VRP solution. Total theoretical distance for a manually-based process is in this case **518.73** mi. Compared with the VRPSolver-optimized solution of **415.70** mi., based on the same actual road network data, this implies a **25%** longer total route distance with the manual process than is theoretically necessary. This data suggests that even an experienced human dispatcher cannot rival a simple vehicle routing algorithm in solution quality, for a problem of practical size.

To clarify, there are actually two important comparisons to be made: (1) the theoretical VRP-solved total route cost versus the manually-allocated and sequenced (but theoretical shortest path-solved) total route cost, and (2) the theoretical versus actual as-driven total route cost. As we have just observed, the manually-planned vehicle routing is, theoretically, about 25% suboptimal to that of the VRP solution. But what about the route as it is actually driven? In actuality the as-run route may be relatively better or even worse than the 25% suboptimal. This is true to the extent that a delivery driver does not actually follow the planned route. The driver may deviate from a planned route as a result of traffic, various road obstacles/detours, or based on their degree of knowledge in navigating an area.

With the deployment of a TSP solver based on quality road network data (i.e. MapPoint) we perform a second test, aiming to characterize the extent of driver deviation from a planned route, based on actual as-run route data over the course of one week. The underlying question is: how well can delivery drivers route themselves optimally through local neighborhoods? To capture this we provide each driver a delivery manifest with a predetermined sequencing. We then instruct the driver to navigate along the quickest path according to their own intuition, using onboard GPS as required. After the driver returns to the central depot, the odometer reading is compared to the distance associated with the theoretical quickest path as per MapPoint. The study indicates that the drivers travel on average 27% longer distance than theoretically necessary, with a range of -1% to over 40% longer distance than necessary. Interestingly, one driver did beat the theoretical optimum slightly, most likely by using more obscure side roads than were permitted in the MapPoint parameter settings.

While we perform the two experiments above independently, not in combination, we may nonetheless estimate the worst-case combined effect of these empirical results. For instance, it is conceivable that the dispatcher performs a manual vehicle routing that is 25% longer than theoretically optimal, which is then driven an additional 27% longer than necessary because of driver error, etc. The resulting total route cost = $1.25 \times 1.27 = 1.5875$, or 58.75% worse than the VRP solution. Conversely, it is also possible that the total degree of suboptimality may be better than that of the manually-planned route, to the extent that the planned route is 25% suboptimal, but the driver deviates in a manner that is actually beneficial to the route. The extent of extra distance travelled translates into lost capacity, affecting not only variable transportation costs but also fixed costs (e.g. fleet size). Refer to the capacity scaling experiment in Section 3.5.3 for an indication of

these fixed and variable costs. The experiments above suggest there is a justifiable need for automated vehicle routing and likely driver training.

3.2 Implementation Considerations

3.2.1 VRP Software Deployment

First, developers should consider the underlying O-D cost matrix data associated with the VRP. Commercial map data providers such as Tele Atlas and NAVTEQ are generally thought to provide the most accurate and current data. However, the cost of this data service may be prohibitive to VRP system development in the early stages. As we have seen, calculating large O-D matrices based on actual road network data is also computationally expensive. This begs the question of whether geometric distances (e.g. calculations based on coordinates) afford “good enough” solution quality.

The following correlation study aims to help answer this question. Consider a set of 212 customer addresses to which a fleet of vehicles must deliver goods on a particular day. The addresses are geographically dispersed across the greater Seattle area, in a rather typical pattern that embodies a mix of randomization and clustering. We calculate an O-D matrix for the set of customer addresses based on three common metrics: (1) Great Circle Distance (2) Actual Road Network Distance (3) Actual Road Network Driving Time.

In addition to the TSP solver, MapPoint utilizes high-quality commercial map data²⁶ and calculates cost matrices (i.e. point-to-point driving distances or times) based on actual road distances. In its present form we leverage this capability to run another valuable experiment. Table 6 illustrates the correlation among various cost metrics based on typical demand. Note a very strong correlation between Great Circle distance, actual road distance and driving time. This study suggests that simple Great Circle distance may be used as a suitable proxy for calculating VRP cost matrices. The simple distance metrics would further benefit from incorporating constraints around physical obstacles such as bodies of water. Calculating actual road distances is computationally expensive²⁷ and adds a layer of complexity we may wish to avoid during the early implementation phase.

²⁶ Commercial map data from Tele Atlas and NAVTEQ are widely used and accurate for this purpose

²⁷ 2:50:00 calculation time for asymmetrical n=212 distance matrix calculation (i.e. 44,944 point-to-point distances)

	<i>Road Distance A- B (mi)</i>	<i>Road Distance B- A (mi)</i>	<i>Road Driving Time B-A (min)</i>	<i>Road Driving Time A-B (min)</i>	<i>Great Circle Distance (mi)</i>
Road Distance A-B (mi)	1				
Road Distance B-A (mi)	0.996	1			
Road Driving Time B-A (min)	0.960	0.964	1		
Road Driving Time A-B (min)	0.949	0.957	0.993	1	
Great Circle Distance (mi)	0.969	0.972	0.934	0.929	1

Source: Sample Unattended Delivery Data, 212 Orders

Note that each of these correlation coefficients is very close to 1, signifying a strong positive correlation between the variables. For instance, if we elect to use a Great Circle (i.e. along a sphere) distance as a proxy for actual road instance, note the strong positive correlation of 0.969. $R^2 = 0.939$, implying that 93.9% of the variance in actual road distance is explained by Great Circle distance.

Table 6: Correlation of Possible VRP Cost Metrics

See Love, Morris, and Wesolowsky [1988, ch.10], and also Alberta [2004] for further reference on the various distance metrics.

One distinct advantage to utilizing actual road network data is its intrinsic handling of unique geographical constraints. For instance, road infrastructure is built around obstacles such as mountains, lakes, and rivers. Therefore it is likely that the O-D matrix generated from actual road network data (assuming the data is accurate and up-to-date) reflects the true cost or distance associated with travelling from origin to destination.

In contrast, a VRP having an O-D cost matrix based on simple geometric distances must include artificial constraints to prevent a physically-impossible solution. A simple example, but one rather common in Puget Sound, is having two customer addresses physically close to one another but separated by a body of water.

Another manner in which to handle geographical constraints, involves splitting the VRP into logical sub-problems beforehand. Consider as an example the case of two populous neighborhoods divided by a large body of water. Assuming the neighborhoods have sufficient demand to justify dedicated routes, it may not make sense to include both neighborhoods in the same VRP at all. Taken a step further, a company may split a large service territory in to several zonal systems.

But what effect does splitting the primary vehicle routing problem into geographical sub-zones have on implementation and solution quality? This concept is tested in Section 3.3.2. Advantages of splitting the VRP into sub-problems may include:

- Simplification of customer order sorting and outbound fulfillment processes FC, since there is no need to postpone initial zone sorting step until the VRP is run.
- Simplification for delivery drivers. A delivery route that is coarsely sorted by zone or neighborhood lends itself to familiarity and more standardized work, comprising a “routine” among local drivers. As an example, market leaders in last-mile delivery such as Fedex, UPS, and USPS depend on locally-knowledgeable drivers to deliver efficiently within typically small neighborhoods.
- Reduction in solution time and computing time required.

However, there are some disadvantages to splitting the VRP into sub-problems as well. Possible disadvantages include:

- Reduction in “optimality” of solution. That is, the total cost of the objective function (in time, distance, number of routes, etc.) for the original VRP is likely to be lower (i.e. better) than that of the sub-problem divided VRP (reference Section 3.3.2 for a comparative study).
- Need to determine logical “dividing line” between sub-zones, which may shift over time as demand patterns fluctuate.

3.2.2 Fulfillment Center

A key consideration with any VRP implementation is its integration with fulfillment processes at the FC. The implications of vehicle routing on FC processes such as picking, packing, sorting, and loading are treated in section 3.3, including an experiment to determine the effect of splitting the VRP into sub-problems.

3.3 Integration with Fulfillment Center Operations

3.3.1 Order Picking and Packing Implications

Ideally a VRP system should integrate well with the warehouse fulfillment processes. As a starting point, the FC’s demand and fulfillment patterns provide valuable insights. Data mining within the e-grocer’s order database via SQL queries yields valuable statistical analysis. Figure 3 below shows a typical distribution of AmazonFresh orders by day of week, from a peak on Mondays to a low on Thursdays.

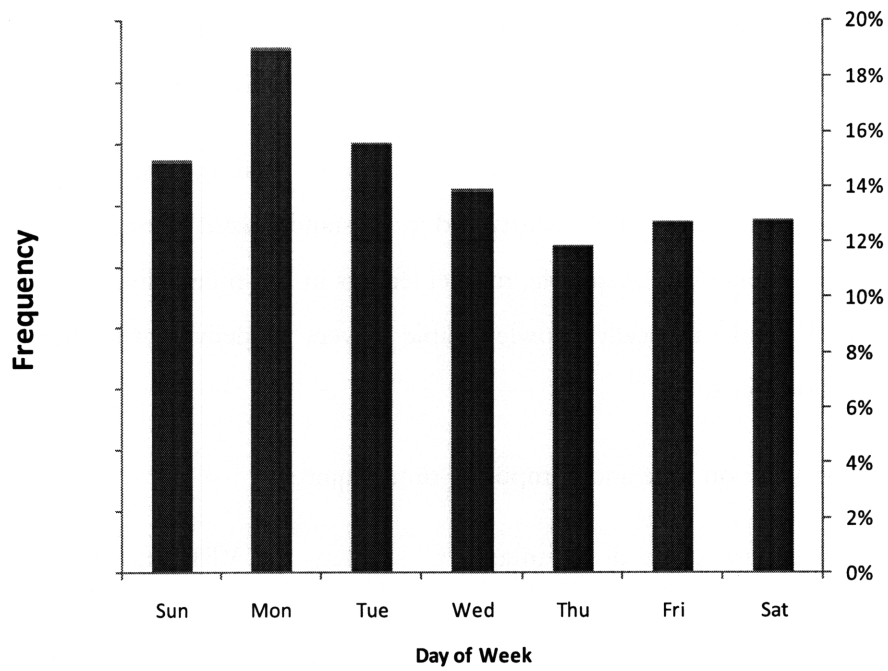
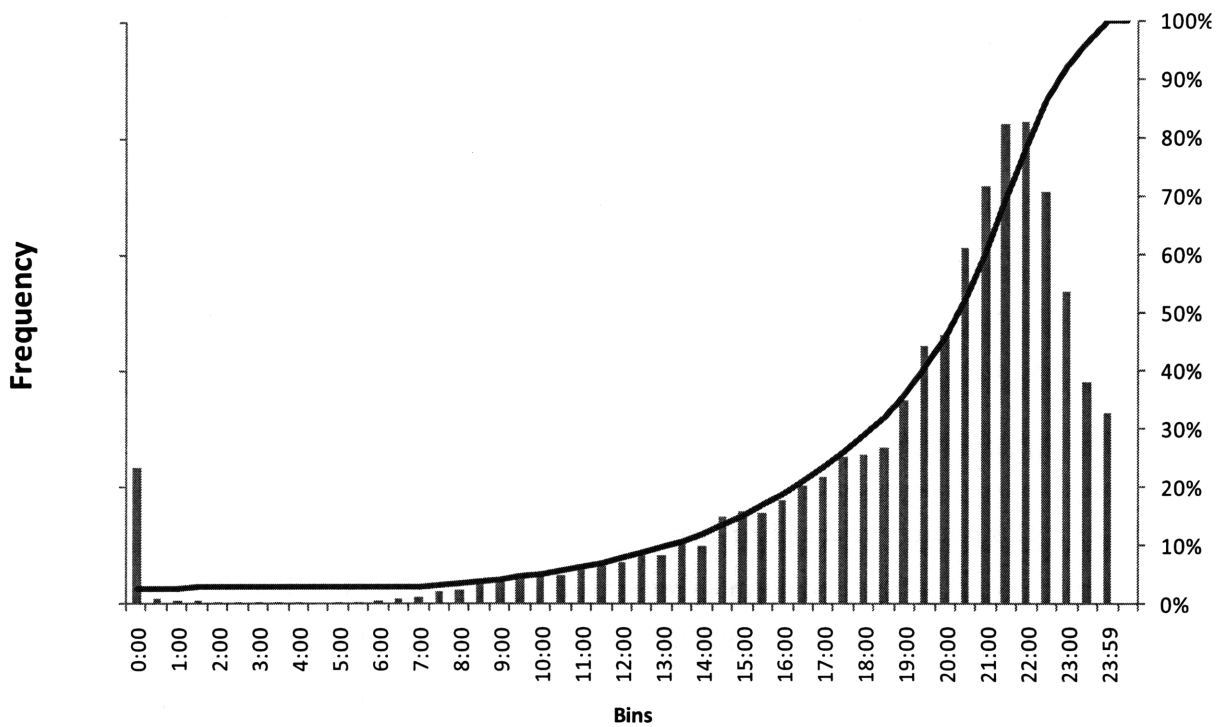


Figure 3: Histogram of E-Grocery Orders by Day of Week

Figure 4 depicts typical customer order patterns by time of day. Observe that on average approximately 92% of orders are received by 11PM, based on an order deadline of midnight.



■ Frequency — Cumulative %

Figure 4: Histogram of E-Grocery Order Patterns by Time of Day

Finally, Figure 5 characterizes demand pattern variability. The two bell shaped histograms illustrate sample distributions of the time at which 80% and 90%, respectively, of a day's cumulative orders have been received by the e-grocer.

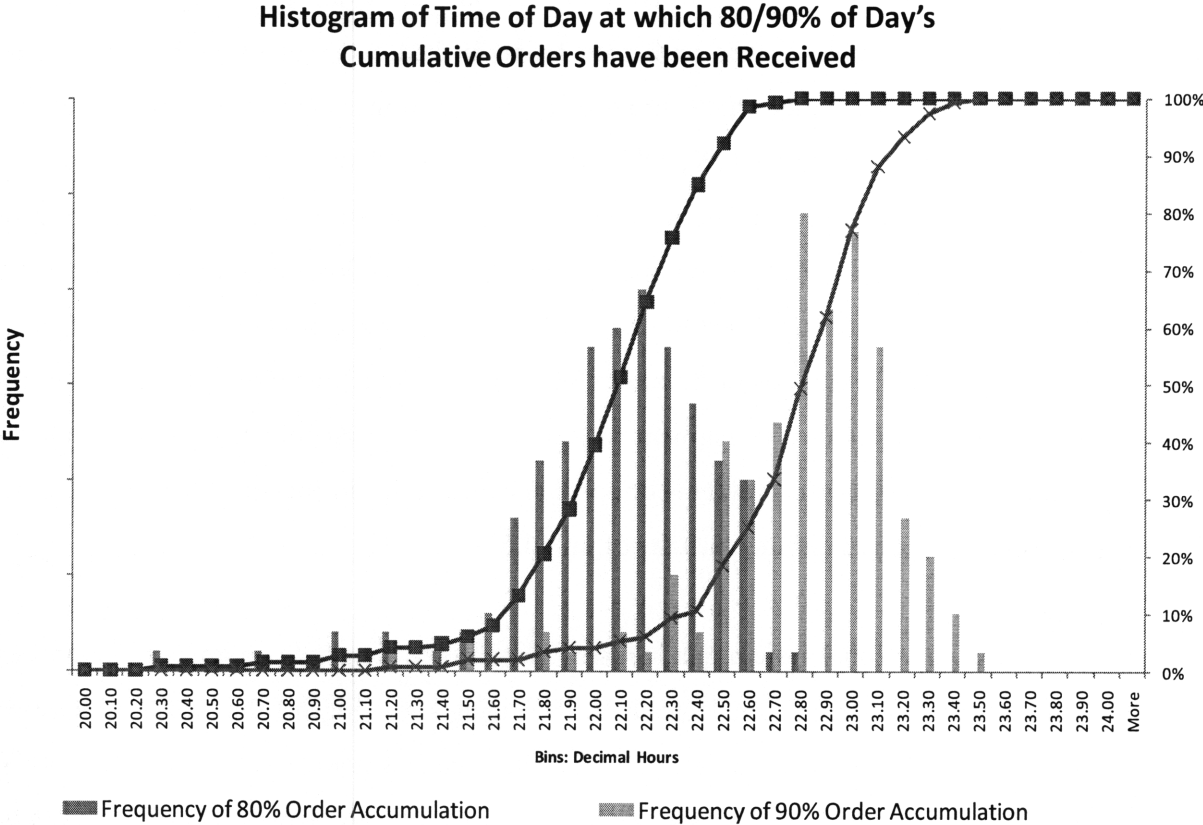


Figure 5: E-Grocery Demand Pattern Variability

The standard deviation for both distributions is approximately .37 hours, which implies 95% confidence that these cumulative order thresholds will be met within +/- .74 hours of the mean (expected) time. The value of this data is that it demonstrates the possibility of an early (i.e. before 12AM customer order placement deadline) VRP run, and thus an early start to final order sorting and truck loading if needed.

3.3.2 Sorting and Truck Loading

In order to characterize the impact of a VRP system on order sorting and truck loading processes, it may be valuable to understand the extent to which the problem can be decomposed into logical sub-problems. As referenced in section 3.2.2, an experiment into these effects is conducted as follows:

EXPERIMENT: Determining the Effects of VRP Sub-Problems or Zonal Systems

Summary:

To the extent that it may be possible to delay a VRP solution and final order sorting, the e-grocer may offer the customer a shorter turnaround time on their order. By adding a preliminary, coarse sorting step, for example, that is by having orders picked and segregated into four zones initially, the final sorting of customer orders into specific delivery routes may occur in parallel at four sorting nodes (e.g. corresponding to NW, SW, NE, SE zones). Particularly if a manual sorting process is employed, the ability to distribute these sorting operations into parallel work may be advantageous from a time and efficiency standpoint. But what impact does segregating the VRP into sub-problems have on the global objective function of minimizing transportation cost? Note that we are utilizing one central FC, with customer orders being picked into four virtual zones within the FC, temporarily pending their final assignment to a particular truck. **Result:** Separating the global VRP for sample data into four sub-problems or zonal systems results in the same number of routes allocated (15) and a suboptimal distance penalty of only 1.6%.

Steps:

- 1) Determine zonal system quadrants by latitude and longitude coordinate boundaries. For the experiment, this is done logically by leveraging knowledge of the local road network and geographical boundaries, as well as balancing demand among the quadrants to some extent. A more sophisticated model might take care of this dynamically. The quadrant boundaries may be shifted to generate a slightly better solution. Graphical representation of the zonal systems in this experiment (i.e. NW, SW, NE, and SE quadrants) are given in Figure 6:

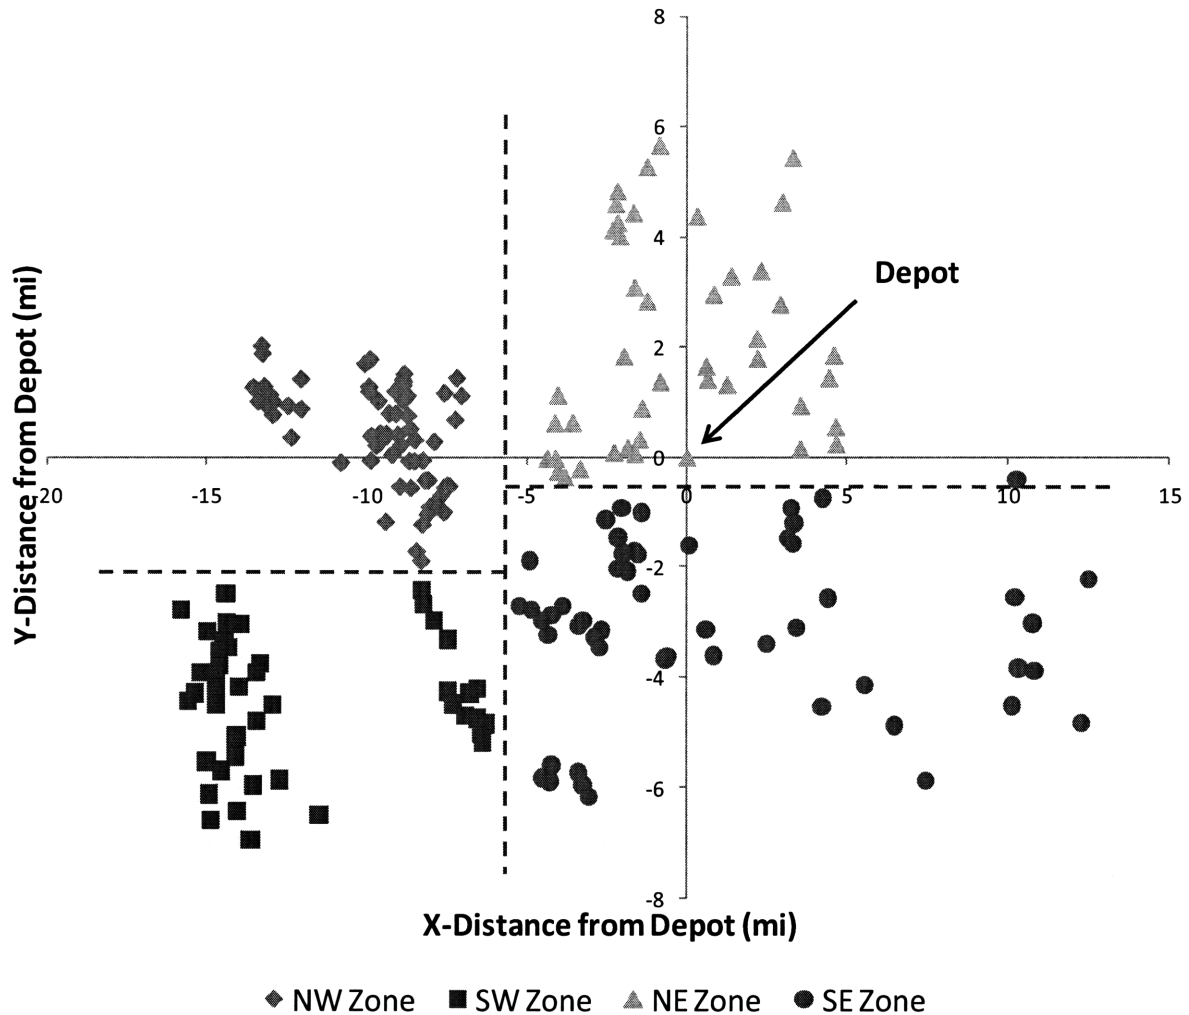


Figure 6: Illustration of VRP Zonal Systems

- 2) Segregate addresses into these four zonal systems by geocoded latitude-longitude coordinates. Output four tab-delimited .txt files, one for each zone, with Cartesian coordinates and Demand (=1 for each address) as per previous VRPSolver Experiment 1 step 5.
- 3) Run VRPSolver using Euclidean distance metric for each of the four zones. Maintaining the same parameter settings as before, we see the following results:

	# Routes	Route Distance (mi)
NW	5	111.77
SW	3	102.5
NE	3	45.39
SE	4	79.25
Total	15	338.91
w/o Megazones		333.61
% suboptimal		1.6%

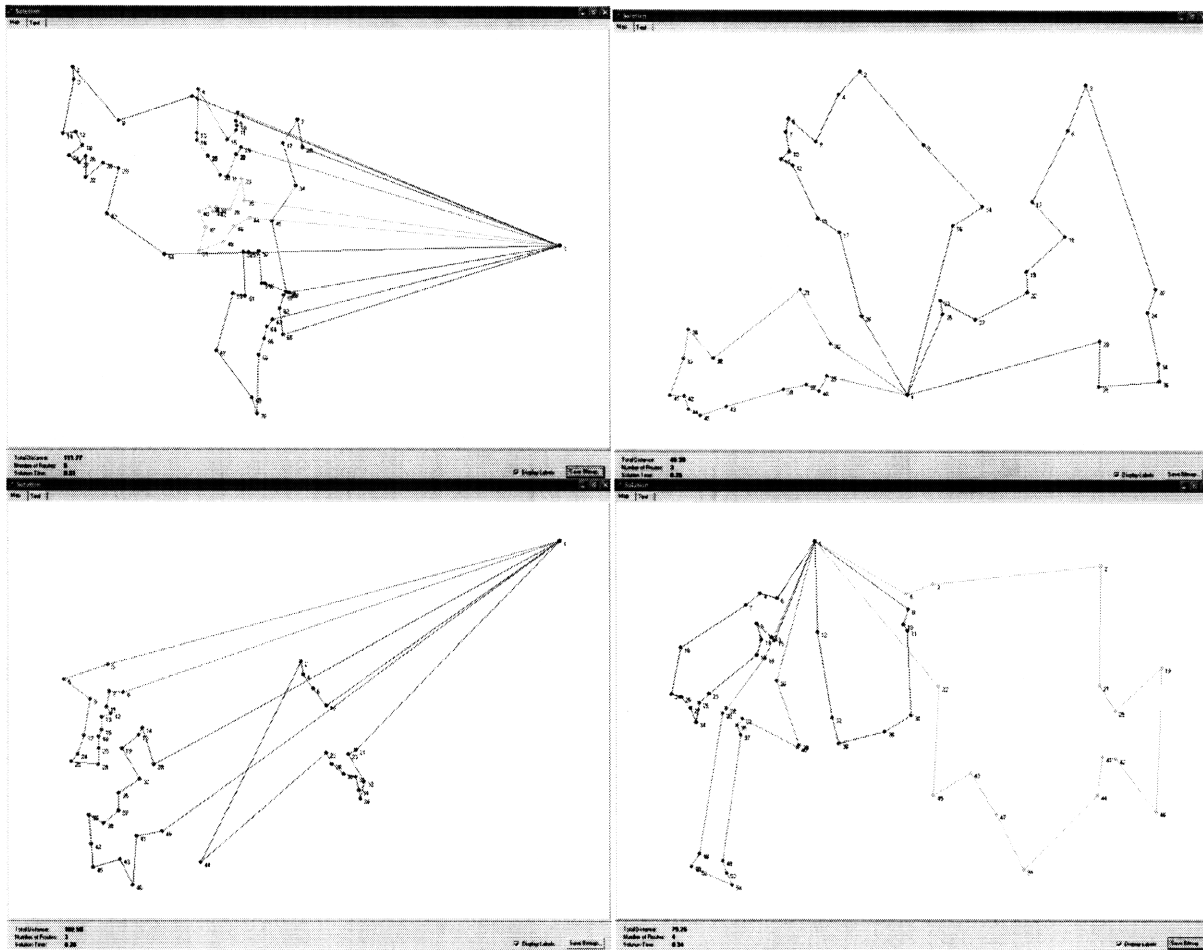


Figure 7: VRP Solutions by Zonal System

Therefore, the VRP solution for the problem divided into four quadrants or zonal systems utilizes the same number of delivery vehicles and is only 1.6% suboptimal in route distance. A logical extension of this experiment would be to vary the number and size of zones in order to model the impact on the objective function. Also note that the problem is very dependent on unique geography and how the zonal systems are divided. The practitioner should consider this data as a demonstration only, and repeat the experiment based on the unique problem they confront.

3.3.3 Fleet Management

Technological advances in mobile devices and GPS tracking have brought these tools within reach for many last-mile delivery providers. GPS-based fleet management technology reduces the need for a dedicated dispatcher in many cases, further ensuring driver accountability, safety and security. Coupled with real-time, wireless delivery confirmation (e.g. barcode scanning the customer order upon delivery), the e-grocer then has access to real-time driving and service time data at the customer level. This data may be used in a feedback loop to update the central dispatching system dynamically, provide higher-granularity input data into the VRP, and help continuously improve customer service.

3.4 Roadmap for VRP Early-Stage Development

The basic components of a real-world VRP implementation comprise (1) Geographic Information System (GIS) or map data (2) geocoding software (3) an optimization framework and (4) a Graphical User Interface (GUI). Fortunately there is free, readily available US TIGER/Line census data that may sufficient for experimentation and early-stage development. More accurate commercial map data (e.g. NAVTEQ, TeleAtlas) will likely be needed for real-world VRP implementation, as well as an effective Address Validation Service²⁸ (e.g. Uniserv). An open-source Geocoding component is also available for experimentation²⁹. The preceding sections propose metaheuristic optimization methods for the e-grocery specific routing problem. Finally, Figure 8 shows a high-level VRP Process Flow Diagram for reference.

²⁸ Address Validation Services are commonly used in postal and delivery networks to correct common errors, account for unique addresses and generally reduce the incidence of unidentifiable or undeliverable addresses.

²⁹ Open source, Perl-based geocoding resource: Geocoder.us

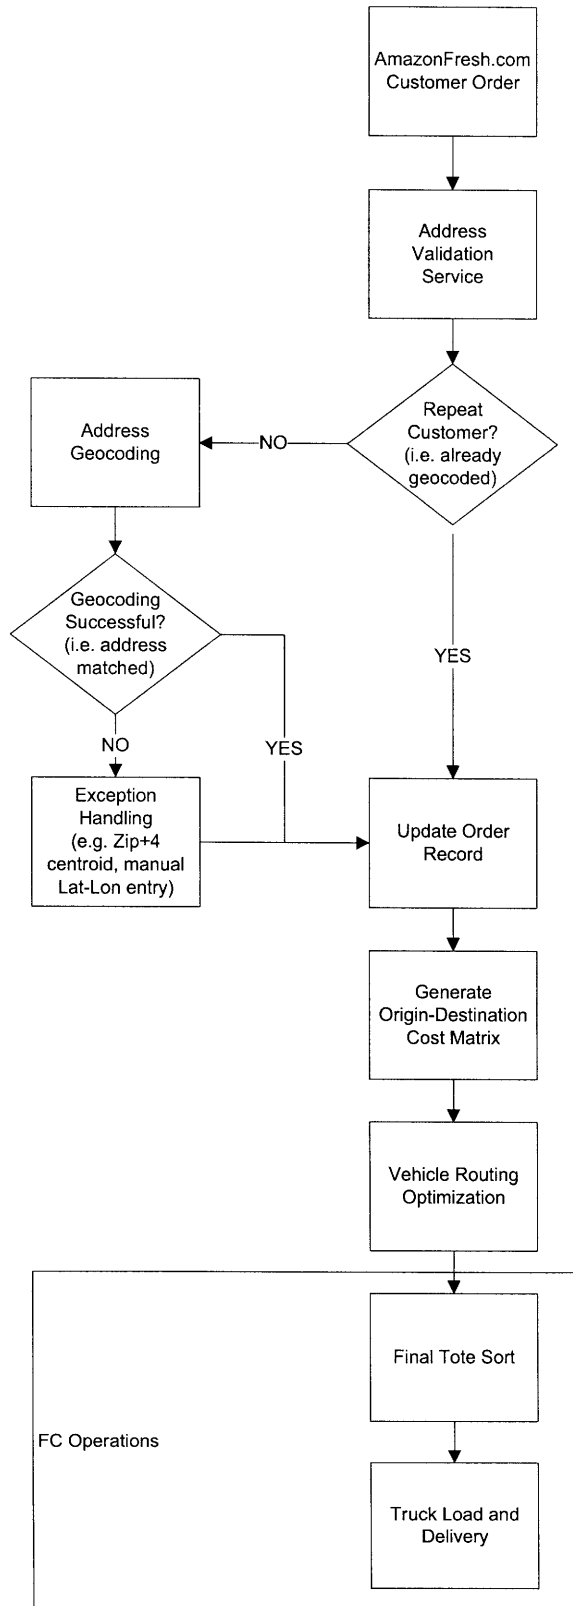


Figure 8: High-Level VRP Process Flow Diagram

3.5 Scalability Considerations, Methods and Techniques

3.5.1 Fulfillment Center Capacity

A key consideration of scalability at the FC, as it pertains to vehicle routing and outbound transportation, is the inherent limitation of loading dock space. Geometric limitations exist on the inbound and outbound docks of any FC such that the e-grocer may be capacity constrained as customer demand grows. For example, if we assume for simplicity a square warehouse with loading docks along one linear side of the building, floor area (a proxy for order fulfillment capacity in the warehouse) scales as the square of the warehouse's outbound dock capacity. This may be a significant risk in an emerging, high-turnover delivery business like e-grocery where dock space is at a premium. A common solution is to send trucks out at a variety of times, effectively bolstering throughput by sharing the dock capacity among a larger fleet.

3.5.2 Demand Shaping

To the extent that customer demand patterns can be understood and manipulated through various incentives, outbound fulfillment costs can be reduced through improved fleet utilization and more efficient vehicle routing. Note there is generally a tradeoff between customer service level and outbound fulfillment cost, and demand shaping through various incentives is an effective means to maintain customer satisfaction while reducing operating costs. Refer to section 3.3.1 for background pertaining to demand patterns. Peapod is an example of an e-grocer offering price incentives to customers who are willing to accept more broad or operationally more efficient delivery windows.

3.5.3 Delivery Process

The scalability of delivery networks is a well studied area of logistics as a result of its importance in industry. Rosenfield, Engelstein, and Feigenbaum [1989] present useful strategies for addressing the delivery territory sizing problem. The literature outlines key analytic relationships for determining the optimum number of service territories, discusses an application to the sizing of postal delivery territories, and addresses the issue of varying density.

The following section presents another empirical experiment based on the case of AmazonFresh, which aims to characterize the scalability of delivery capacity within an existing fleet of vehicles.

EXPERIMENT: Determining the Effects of Scaling Capacity

Delivery capacity (i.e. customer orders fulfilled per vehicle within a given route) is varied to determine the effect on optimal (VRP-solved) total route distance as well as the number of routes. From these solutions we can infer a rough cost impact. For the purpose of this example, let us assume a hypothetical delivery truck capacity of 120 totes (or 24 orders assuming five totes per order), and a driver capacity of 75 totes (or 15 orders assuming five totes per order). That is, assume that the driver can only deliver some fraction (e.g. $5/8^{\text{th}}$) of physical capacity of the truck because the driver is actually constrained by an overriding time window, within which the deliveries must be made. Note that delivery stops per hour will reach a theoretical limit at the point where the problem-limiting constraint shifts from time (i.e. the time needed to complete driving and delivery stops) to volumetric capacity. This would occur as the density of stops increases and time between stops decreases. Further efficiency gains would be dependent on an improvement in tote packaging density (i.e. more items packaged per unit volume). In this experiment, the nominal capacity is assumed to be 15 orders per truck, based on a 3-hour unattended window constraint. The volumetric limit of a delivery truck is set to 120 totes, or approximately 24 orders based on an assumed average grocery order size of five totes. Table 7 below presents the results of the experiment:

VRP SOLUTION								
Capacity (Orders/Truck)	# Routes	Total Distance (mi)	Avg Distance/Route (mi)	% Improvement in Total Distance	% Improvement in # Routes	Nominal Stops/Hr		
10	22	445.74	20.26	-33.6%	-46.7%	2.9	Lowest Performing Driver	
11	20	411.4	20.57	-23.3%	-33.3%	3.1		
12	18	385.65	21.43	-15.6%	-20.0%	3.4		
13	17	364.43	21.44	-9.2%	-13.3%	3.7		
14	16	349.62	21.85	-4.8%	-6.7%	4.0		
15	15	333.61	22.24	0.0%	0.0%	4.3	Baseline Avg	
16	14	316.58	22.61	5.1%	6.7%	4.6	Improvement Region w/ Current Tote Configuration	
17	13	303.09	23.31	9.1%	13.3%	4.9		
18	12	293.55	24.46	12.0%	20.0%	5.1		
19	12	287.65	23.97	13.8%	20.0%	5.4		
20	11	280.71	25.52	15.9%	26.7%	5.7		
21	11	275.59	25.05	17.4%	26.7%	6.0		
22	10	266.81	26.68	20.0%	33.3%	6.3		
23	10	260.02	26.00	22.1%	33.3%	6.6		
24	9	256.09	28.45	23.2%	40.0%	6.9		Truck Capacity Constraint (i.e. 120 Totes/Truck @ 5 Totes/Order Avg)
25	9	254.07	28.23	23.8%	40.0%	7.1		
26	9	248.21	27.58	25.6%	40.0%	7.4		
27	8	236.83	29.60	29.0%	46.7%	7.7		
28	8	236.49	29.56	29.1%	46.7%	8.0		
29	8	231.99	29.00	30.5%	46.7%	8.3		
30	8	231.24	28.91	30.7%	46.7%	8.6		
31	7	215.58	30.80	35.4%	53.3%	8.9		
32	7	215.47	30.78	35.4%	53.3%	9.1		
33	7	214.84	30.69	35.6%	53.3%	9.4		
34	7	214.19	30.60	35.8%	53.3%	9.7		
35	7	212.42	30.35	36.3%	53.3%	10.0		

Table 7: Impact of Scaling Delivery Capacity on VRP

Note the improvement region within which a driver may dramatically improve the VRP solution simply by being more efficient in their nominal delivery rate (e.g. by driving more efficiently and/or reducing service time per stop). The capacity constraint is actually a function of time (i.e. meeting a customer promise time) up to the physical volumetric constraint of 24 orders per truck on route. The data in Table 7 illustrates that if drivers improved their pace so as to maintain a capacity of 16 customer orders on a route, a 5.1% reduction in total route distance and 6.7% reduction in fleet size would result. Moreover, if drivers improved so as to maintain a capacity of 24 customer orders on a route (the volumetric capacity of the truck), a 23.2% reduction in total route distance and 40% reduction in fleet size would result.

Figures 9 and 10 illustrate the data in graphical form. Figure 9 depicts Truck Capacity versus Theoretical Improvement in Total Distance and # Routes. Figure 10 shows Truck Capacity versus Theoretical Fixed and Variable Cost Improvement.

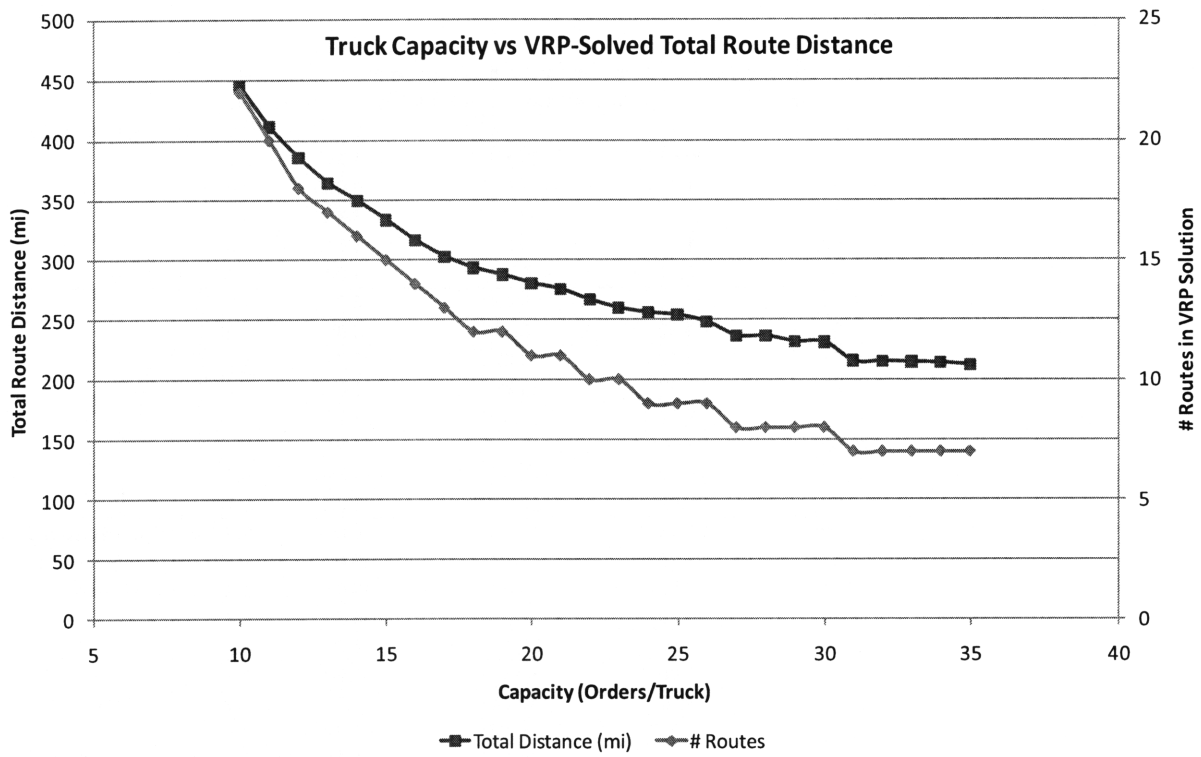


Figure 9: Truck Capacity vs Theoretical Improvement in Total Distance and # Routes

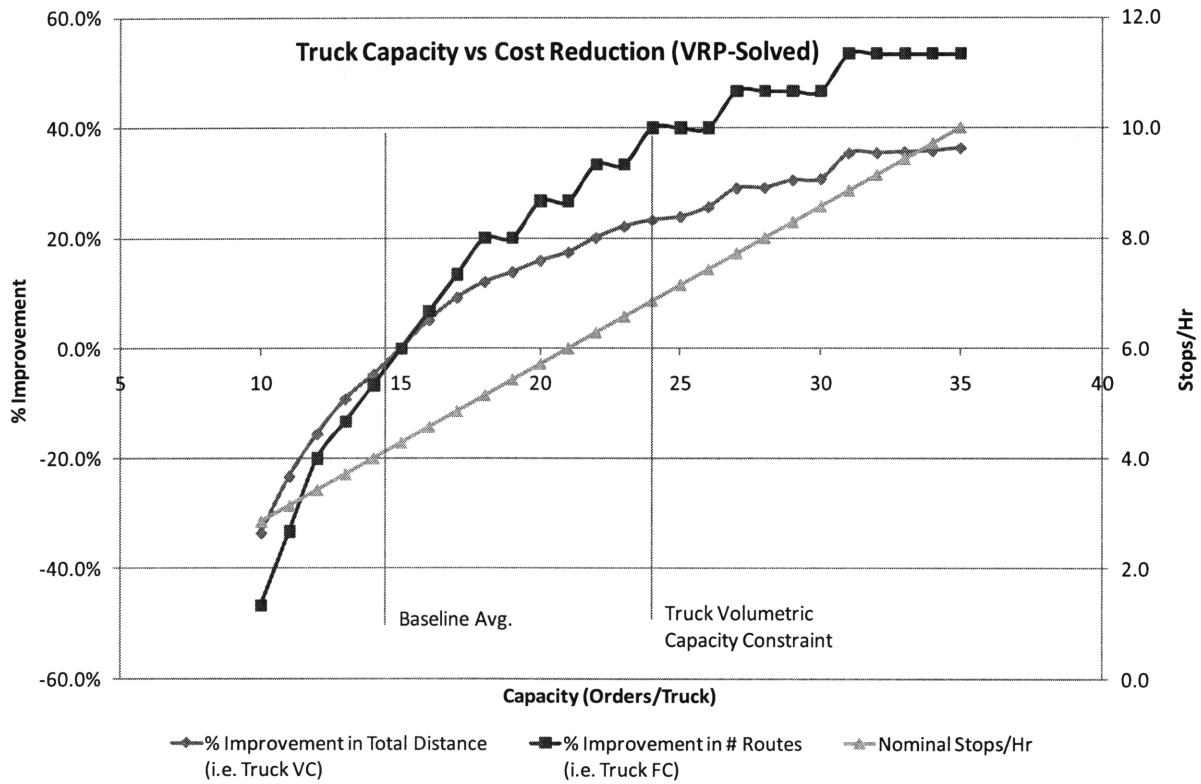


Figure 10: Truck Capacity vs Theoretical Fixed and Variable Cost Improvement

3.5.4 Distribution Network and Fulfillment Center Location

The following experiment presents a simulation to determine the effect of scaling customer density within the service territory, utilizing VRP software results as the theoretical output.

EXPERIMENT: Determining the Effects of Scaling Density

Steps:

- 1) Simulate increased density by augmenting a sample route from the pool of known customers. For the purposes of this experiment we instead generate “fictional” addresses at a midpoint between adjacent known addresses. Initial Euclidian distance VRP solutions are used to identify clusters and the set of adjacent addresses. To simulate a 100% increase in density, we generate 211 fictional addresses within predetermined clusters of known addresses. To simulate less than 100% density increase, a random number generator selects a given percentage of fictional addresses within the set (e.g. 25%, 50%, 75%). Note this is not entirely realistic. In reality the new customers would be more widely scattered. However this does allow for simple geographical bounding that is consistent with an experiment in increasing customer density within a region.
- 2) Solve the VRP for each set of addresses, including a number of fictional addresses in proportion to the degree of customer density simulation. Note that for the VRP solution based on 100% density increase, in Figure 11, while the number of customer addresses doubles, the number of routes and total distance is less than double as a result of the efficiency gain.

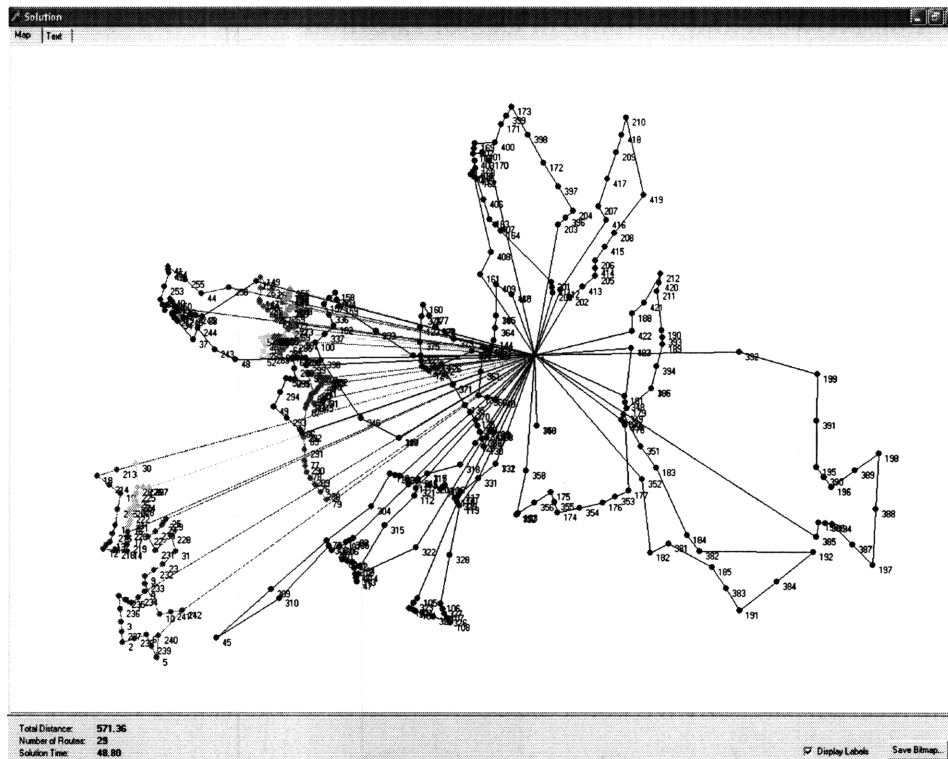


Figure 11: VRP Solution Based on 100% Density Increase

3) Plot the VRP solutions in distance/stop (correlated with time) as a function of % increase in density. Note that these solutions initially assume truck capacity is held constant at 15 orders/truck. The effects of the density simulation are depicted in Figure 12:

# Addresses	Nominal % Density Increase	Actual % Density Increase (Random Generated)	VRP Solved Distance (mi)	# Routes	Distance/Stop (mi)	% Distance Reduction	Theoretical Driving Time		% Driver Labor Saved incl Service Time	Implied addl. Capacity (Orders/Truck)	Adjusted Capacity (Orders/Truck)
							Saved (hr/Predawn)	% Driving Time Saved			
211	0%	0%	333.61	15	1.58	0.0%	0.00	0.0%	0.0%	0.00	15
251	25%	19%	374.73	17	1.49	5.6%	1.88	5.6%	3.6%	0.54	15
303	50%	44%	424.4	21	1.40	11.4%	3.85	11.4%	7.3%	1.10	16
377	75%	79%	525.65	26	1.39	11.8%	3.99	11.8%	7.6%	1.14	16
422	100%	100%	571.36	29	1.35	14.4%	4.85	14.4%	9.2%	1.39	16

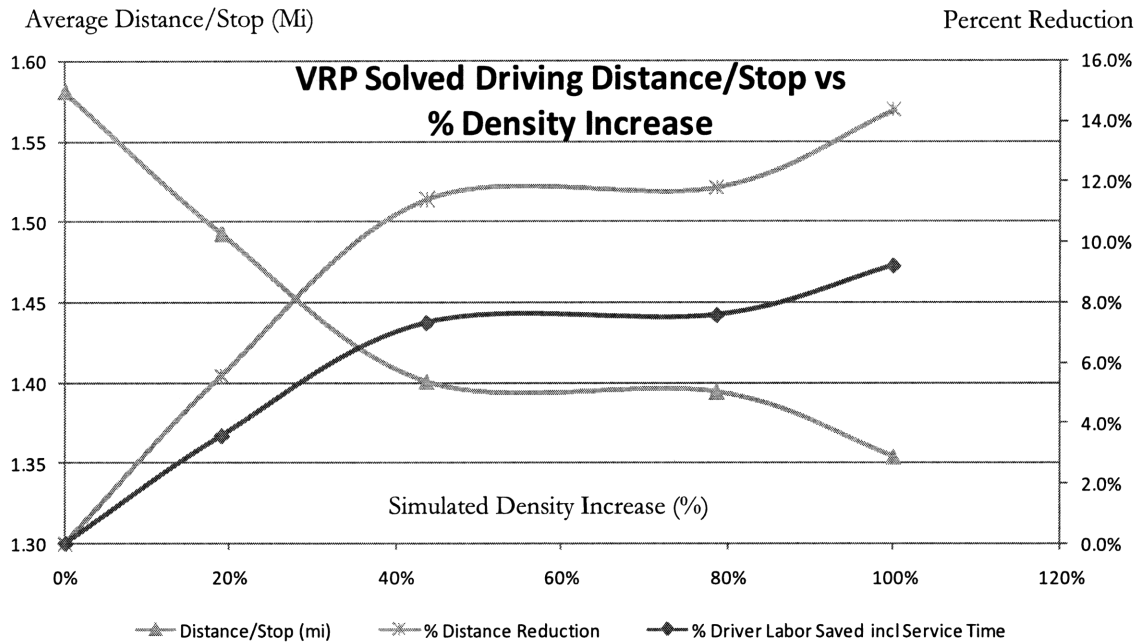


Figure 12: VRP-Solved Driving Distance/Stop vs % Density Increase

We take into account a second-order improvement in capacity and resolve the VRP. We base this capacity adjustment on the stops/hr efficiency improvement (attributable to higher customer density) yielding a higher effective capacity (since the real-world problem is time constrained not volumetric capacity constrained), which in turn yields an incrementally-better VRP solution once we account for this second-order capacity improvement. Figure 13 below illustrates the capacity-adjusted results:

# Addresses	Nominal % Density Increase	Actual % Density Increase (Random Generated)	Re-Solved VRP w/ Capacity Adjustment (mi)	Re-Solved VRP w/ Capacity		% Distance Reduction	Theoretical Driving Time Saved		% Driver Labor Saved incl Service Time
				(# Routes)	Distance/Stop (mi)		(hr/Predawn)	% Driving Time Saved	
211	0%	0%	333.61	15	1.58	0.0%	0.00	0.0%	0.0%
251	25%	19%	374.73	17	1.49	5.6%	1.88	5.6%	3.6%
303	50%	44%	405.71	19	1.34	15.3%	5.17	15.3%	9.5%
377	75%	79%	496.55	24	1.32	16.7%	5.63	16.7%	10.3%
422	100%	100%	544.78	27	1.29	18.4%	6.19	18.4%	11.4%

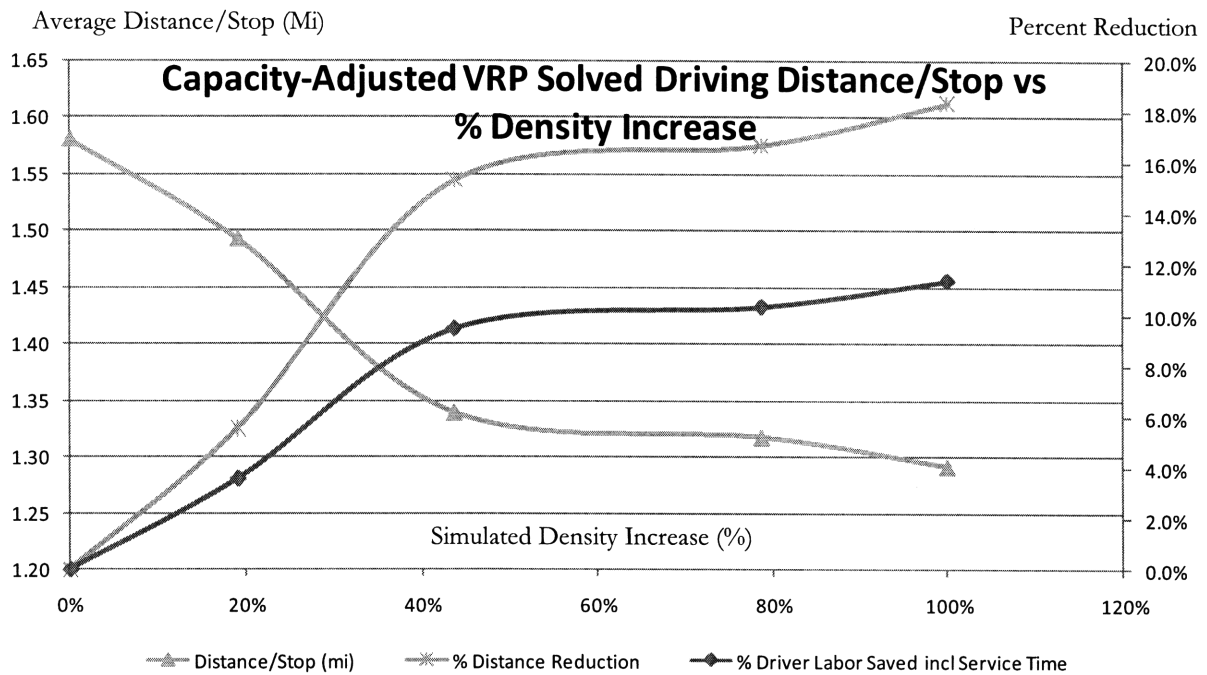


Figure 13: VRP-Solved Driving Distance/Stop vs % Density Increase (Capacity Adjusted)

Therefore, a simulation of doubling the customer density impacts vehicle routing by reducing the average driving distance between stops by more than 18%.

3.5.5 Strategic Expansion

Interestingly it is also possible to leverage VRP systems to determine an optimal expansion strategy. For example, consider the VRP to select the optimal expansion zone from a list of candidate zones. Assuming aggregate demand will be impacted similarly regardless of which new market is brought online, then the proposed distribution network can be simulated and optimized from the perspective of minimizing route fulfillment costs.

EXPERIMENT: Determining Strategic Expansion Zones by VRP Simulation

We again use VRPSolver to solve for a typical problem at AmazonFresh using real data, but augmented with simulated customers from various proposed expansion zones in each case. We then compare the total VRP-solved simulated delivery costs across all proposed new zones. We also compare the new zone simulations to a base case of simulated organic growth (distributed among the existing service territory, demand-weighted by individual zone). One might expect the transportation cost impact of opening a new zone to be somewhat worse than that of equivalent

organic growth (e.g. primarily an increase in customer density within current delivery territory), but this not necessarily the case. Based on a simulation of eight candidate zip codes (proposed new zones) at AmazonFresh, we determine not only the most favorable in terms of transportation cost impact, but also that two of the eight zip codes offer cost advantages over simple organic growth. While opening a new zone generally worsens the objective function relative to organic growth, two zip codes on the fringe of the existing service territory actually improve upon it by 1.5% and 1.1%, respectively, in the AmazonFresh case. This outcome is likely attributable to the fleet making more efficient use of stem distances and adjacent demand geographies. Thus, opening a new geographical zone can actually help transportation costs relative to relying on organic growth across all zones.

PART IV: Conclusion

4.1 Summary of Findings

The primary objective of this project is to identify areas for improvement within the outbound fulfillment network of an online grocery company, with emphasis on developing an efficient, scalable home delivery network. Presented in the paper is a prototypical case study of an emerging online grocery company, together with a number of analytical techniques useful for optimizing its outbound fulfillment and vehicle routing network. We review an extensive literature and offer practical advice, tailored to suit the interests of an emerging e-grocery company. The paper directs the reader to commercially-available solutions, and provides guidance into the development of a customized proprietary vehicle routing system.

Although academics and practitioners have studied VRP problem class, the *NP-hard* combinatorial optimization is still challenging to implement in practice. The work in this paper leverages the academic VRP software implementation VRPSolver, which proves useful for such experimentation. Generating an O-D matrix using actual road networks proves to be computationally expensive. We provide insights through new case experiments and analysis, which are intended to serve as an easily-accessible template for those interested to explore using their own data. First, a study of VRP-solved versus manually routed network costs is conducted. The experimental analysis suggests validation and justification of vehicle routing optimization relative to manual systems. A study of the correlation among various O-D cost metrics illustrates the sufficiency of basic geometric distance metrics, and perhaps a lack of need for road network data in more simple VRP cases. The segregation of a real-world VRP into four logical zonal systems indicates benefits in problem simplicity and customer order sorting, at only a miniscule cost to the aggregate VRP's optimality. Next, an experiment in scaling delivery truck capacity maps out a relationship between improved delivery driver performance and theoretical network costs. Even a slight improvement in delivery capacity yields not only reductions in variable transportation costs but also fleet size. Similarly, the paper presents an experiment characterizing the effect of scaling customer density, and mapping its relationship with fulfillment costs. Finally, we describe a scenario to utilize VRP systems as a means to optimize a company's expansion strategy. The following section presents recommendations and possible future steps.

4.2 Recommendations and Future Steps

An e-grocery company should explore the commercially-available VRP systems, weighing the advantages and disadvantages of each in the context of the firm's needs. For an organization interested in developing a proprietary VRP system, and having the necessary development resources, a vast amount of research is available in the literature. Benchmark problem instances such as Solomon and Homburger may be used to narrow down the academic research and state-of-the-art in metaheuristic approaches. In the case of either a commercial or proprietary approach, the vehicle routing system should be tightly integrated with warehouse fulfillment processes insofar as possible.

Future steps following the implementation of a VRPTW system within an e-grocery company may include additional experiments in the area of this paper. One extension would be comparing VRPTW (Attended 1hr windows) to the relaxed VRP case, in order to characterize the cost of narrow time windows. Research suggests that the cost of offering narrow service windows is particularly high³⁰. In addition, it may be valuable to compare the deterministic VRP to various real-world stochastic components such as urban traffic, service time, and demand variability. A Pareto analysis into these various factor would serve as interesting research. Based on the inherent proximity of e-grocery to urban centers, and the substantial influence that traffic delays will have on fulfillment cost and customer service, this factor in particular seems critical.

To characterize the effects of driving and service time variability on overall route time, a Monte Carlo simulation experiment may also be conducted. This simulation would help to characterize the effects of urban traffic within the fulfillment network, and confirm whether the VRP solution should leverage address-level driving/service times, so as to minimize the risk of using a simple average.

Deployment of mobile scanning and computing devices will allow actual customer-level driving and service times to be inferred electronically. However, for an e-grocery startup company, the upfront cost may not be immediately justified. In the long term, GPS fleet management in combination with a mobile device (e.g. latest-generation Symbol models) would facilitate streamlined delivery, fewer delivery errors, real-time order tracking, and embedded service data that would serve as a basis for continuous improvement.

³⁰ Punakivi and Saranen (2001) found that fully-flexible unattended windows reduced costs by up to 33%, relative to 2hr delivery windows.

REFERENCES

- Bodin, L., Golden, B., Assad, A., and Ball, M. (1983). "The state of the art in the routing and scheduling of vehicles and crews." *Computers and Operations Research* 10, 63–212.
- Bräysy O. and Gendreau M., (2005a). "Vehicle Routing Problem with Time Windows, Part I: Route Construction and Local Search Algorithms." *Transportation Science* 39, 104–118.
- Bräysy O. and Gendreau M., (2005b). "Vehicle Routing Problem with Time Windows, Part II: Metaheuristics." *Transportation Science* 39, 119–139.
- Calculating Distances in Alberta. (2004) Geographic Methodology Series No. 4. Health Surveillance Branch, Alberta Health and Wellness.
- Clarke, G. and J. Wright (1964). "Scheduling of vehicles from a central depot to a number of delivery points." *Operations Research* 12 (4), 568–581.
- Dantzig G.B and Ramser J.H. (1959). "The truck dispatching problem." *Management Science* 6, 80–91.
- De Backer B., V. Furnon, P. Prosser, P. Kilby and P. Shaw (1997). "Local Search in Constraint Programming: Application to the Vehicle Routing Problem," in *Proceedings of CP-97 Workshop on Industrial Constraint-Directed Scheduling*, 1–15, Schloss Hagenberg, Austria.
- Dorigo and Stutzle (2004). *Ant Colony Optimization*. MIT Press.
- Gambardella, L. M., Taillard, E., and Agazzi, G. (1999). "MACS-VRPTW: A Multiple Ant Colony System for Vehicle Routing Problems with Time Windows," in *New Ideas in Optimization*, D. Corne, M. Dorigo and F. Glover (eds), 63-76, McGraw-Hill, London.
- Garey, MR and Johnson, DS. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York. W.H Freeman
- Gehring, H. and Homberger, J. (2001). "A Parallel Two-phase Metaheuristic for Routing Problems with Time Windows." *Asia-Pacific Journal of Operational Research* 18, 35-47.
- Golden, B., Raghavan, S., and Wasil, E. (Eds.) (2008). *The Vehicle Routing Problem: Latest Advances and New Challenges. Operations Research/Computer Science Interfaces Series, Vol. 43*.
- Golden, B. and Wasil, E. (1987). "Computerized Vehicle Routing in the Soft Drink Industry." *Operations Research* 35, 6-17.
- Laporte G. (2007). "What you should know about the vehicle routing problem." *Naval Research Logistics* 54, 811-819.
- Lenstra, J. K. and A. H. G. Rinnooy Kan (1981). "Complexity of vehicle routing and scheduling problems." *Networks* 11, 221–227.
- Love, R.F., Morris, J.G., and Wesolowsky, G.O. (1988). *Facilities Location-Models and Methods*, ORSA Publications in Operation Research Series, North Holland.
- Mester, D. and Bräysy, O. (2007). "Active Guided Evolution Strategies for Large Scale Vehicle Routing Problems with Time Windows." *Computers & Operations Research* 32, 1593-1614.

- Prescott-Gagnon, E., Desaulniers, G., and Rousseau, L.M.. (2007). "A Branch-and-Price-Based Large Neighborhood Search Algorithm for the Vehicle Routing Problem with Time Windows."
- Punakivi, M. and Saranen, J., (2001). "Identifying the success factors in e-grocery home delivery." *International Journal of Retail and Distribution Management* 29 (4), 156-163.
- Rochat, Y. and Taillard, E.D. (1995). "Probabilistic Diversification and Intensification in Local Search for Vehicle Routing." *Journal of Heuristics* 1. 147-167.
- Rosenfield, D., Engelstein, I., and Feigenbaum, D. (1992). "An Application of Sizing Service Territories." *European Journal of Operational Research* 63, 164-172.
- Siikavirta, H., Punakivi, M., Kärkkäinen, M. and Linnainen, L. (2003). "Effects of e-commerce on greenhouse gas emissions: a case study of grocery home delivery in Finland.," *Journal of Industrial Ecology* 6 (2), 83-97.
- Taillard, E., Badeau, P., Gendreau, M., Geurtin, F., and Potvin, J.Y. (1997). "A Tabu Search Heuristic for the Vehicle Routing Problem with Time Windows." *Transportation Science* 31, 170-186.
- Tong, N. (2008). Graduate Thesis: "Strategies for Success in the e-Grocery Industry." Rochester Institute of Technology.
- Toth, P. and Vigo, D. (Eds). (2001). *The Vehicle Routing Problem*. Monographs on Discrete Mathematics and Applications. SIAM, Philadelphia.