

Analysis of Demand Variability and Robustness in Strategic Transportation Planning

by

Ahmedali Lokhandwala

B.S. Marine Engineering
Birla Institute of Technology & Science, Pilani (2007)

Submitted to the Department of Civil and Environmental Engineering
and Engineering Systems Division
in partial fulfillment of the requirements for the degrees of

Master of Science in Transportation

and

Master of Engineering in Logistics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

© Massachusetts Institute of Technology 2010. All rights reserved.

Author
Department of Civil and Environmental Engineering
and Engineering Systems Division
May 17, 2010

Certified by.....
Chris Caplice
Executive Director, Center for Transportation & Logistics
Thesis Supervisor

Accepted by.....
Daniele Veneziano
Chairman, Departmental Committee for Graduate Students

Accepted by.....
Yossi Sheffi
Professor, Engineering Systems Division
Professor, Civil and Environmental Engineering Department
Director, Center for Transportation & Logistics
Director, Engineering Systems Division

Analysis of Demand Variability and Robustness in Strategic Transportation Planning

by

Ahmedali Lokhandwala

Submitted to the Department of Civil and Environmental Engineering
and Engineering Systems Division
on May 17, 2010, in partial fulfillment of the
requirements for the degrees of
Master of Science in Transportation
and
Master of Engineering in Logistics

Abstract

Creation of a long-term strategic transportation plan is critical for companies in order to make informed decisions about fleet capacity, number of drivers needed, fleet allocation to domiciles, etc. However, the inherent demand variability present on a transportation network, in terms of weekly occurrences of lane volume, results in emergency weekly shipments that deviate from the long-term plan. This leads to a sub-optimal weekly execution, resulting in higher overall costs, compared to initial projections. Hence, it is important to address this variability while creating a strategic plan, such that it is robust enough to handle these variations, and is easy to execute at the same time. The purpose of this thesis is to create a stochastic annual plan using linear programming techniques for addressing demand variability, and prove its robustness using simple heuristics, so that it is easy to execute at an operational level. Through the use of simulations, it is shown that the proposed planning methodology is within 6% of the optimal solution costs and handles 71% of the demand variability occurring on a weekly basis, making it easy for operational managers to execute. Thus, the proposed plan reduces the optimality gap between long-term planning and weekly operations, creating a tighter bound over the projected versus actual costs incurred, which helps develop a better transportation strategy.

Thesis Supervisor: Chris Caplice

Title: Executive Director, Center for Transportation & Logistics

Acknowledgments

I would like to express a deep sense of gratitude and acknowledge the following people for a wonderful journey at MIT for the past two years. I am sure that the list is incomplete, and there are definitely many more folks who have helped me directly or indirectly during my time here.

My mentor and advisor, Dr. Chris Caplice for his constant advice, support and belief in my capabilities over the past two years. His boundless energy, quick thinking and infallible insights continue to amaze me! He has constantly pushed me to challenge the limits of my capabilities and has been instrumental in my personal and professional growth at MIT.

Dr. Francisco Jauffred, the Principal Investigator on this project (and my office-mate) for his amazing insights and help during my tenure at MIT. Whenever I was stuck with a problem, I knew the solution was only a “question to Francisco” away! No wonder, we lovingly call him the processor that runs the Wal-Mart computer (*Francisco Inside!*). In addition, I would also like to thank Jeff for his help with the MATLAB codes for my planning scenarios, which have been an instrumental part of my thesis.

My MLOG classmates from the Classes of 2009 and 2010, who have been a constant source of inspiration and advice on worldly as well as academic matters. Words cannot express the amount of learnings I have gotten from them - be it during the long hours at the MLOG Lab or the unbelievable amounts of fun that we had outside of it! The good times with them is undoubtedly the part about MIT that I will cherish and miss the most! I would also like to mention my friends outside of MLOG, in particular Kushal, Sarvee, Kashi and countless others whose unconditional friendship and help, as well as the long discussions about random topics have made the past two years a memorable experience.

The Administrative Staff from both my programs - in particular, Patty Glidden, Kris Kipp, Ineke Dyer and Mark Colvin who have been of immense help in crossing any administrative hurdles; and Jon Pratt for always being there as a friend and a

supporter, in addition to his help with the recruiting process. The CTL staff on the 2nd floor of E40, particularly Mary, Nancy, Karen & Eric for counting me as a part of the CTL community as well as for all the monthly lunches!

And last, but not the least, I would like to thank my family for their unconditional love and support.

Dedication

I would like to dedicate this thesis to my Mom and Dad for believing in me, and for placing unconditional trust in my abilities and judgement. This thesis has been brought to fruition through your blessings. The values that you have instilled in me are a big part of the person that I am today, and the person that I someday hope to be!

I would also like to share this moment of happiness with my sister Tasneem, brother-in-law Murtuza, my grandparents and in particular, my darling niece Mariya! I cannot thank you'll enough for always being there for me - through good times and bad.

Contents

1	Introduction	15
2	Trucking Industry Overview	19
2.1	U.S. Transportation System	19
2.2	Transportation Modes and Industry Trends	21
2.3	Trucking Industry	22
2.4	Types of Trucking Fleets	27
2.4.1	Private Fleet	27
2.4.2	For-Hire Carriers	29
2.4.3	Dedicated Fleet	32
2.5	Summary	34
3	Thesis Motivation: Demand Variability and Transportation Planning	35
3.1	Variability	35
3.2	Criticality of Transportation Planning	37
3.2.1	Long-term Planning	38
3.2.2	Short-term planning	40
4	Large-Scale Linear Optimization Model for Transportation Planning	43
4.1	Freight Network Optimization Tool (FNOT)	44
4.2	Understanding FNOT	45

5	The Data and Initial Analysis	57
5.1	Data Consolidation	58
5.2	Lane Demand Distributions & Goodness of Fit	59
5.3	Seasonality in Demand	61
5.4	Correlations in Demand	62
5.5	Sweet Spot on Optimality Curve	64
5.5.1	Metrics	65
6	Planning Methodologies & Weekly Operations	73
6.1	Long Term Annual Planning	73
6.1.1	Stochastic Plans	74
6.1.2	Deterministic Plans	74
6.2	Sub-network used for analysis	75
6.3	Simulation of Weekly Demands	75
6.4	Operational Flexibility Scenarios	79
6.4.1	Complete Operational Flexibility	79
6.4.2	Zero Operational Flexibility	80
7	Simulation Results Analysis	87
7.1	Network Level Metrics	87
7.2	Facility Level Metrics	95
7.3	Tour Level Metrics	99
7.4	Lane Level Metrics	108
8	Conclusions	117
8.1	Opportunities for Future Research	119
	Bibliography	122

List of Figures

2-1	(a) Evolution of the US Transportation Sector in terms of the Value of freight carried; and (b) Evolution of the US Transportation Sector in terms of the Tonnage of freight carried; and (c) Evolution of the US Transportation Sector in terms of the Ton-Miles of freight carried . . .	20
2-2	(a) Break-up and evolution of the modes of transportation within the US in terms of value of freight carried; and (b) Break-up and evolution of the modes of transportation within the US in terms of tonnage of freight carried; and (c) Break-up and evolution of the modes of transportation within the US in terms of ton-miles of freight carried .	23
2-3	(a) Breakup of Revenues generated by the different transportation sectors within the U.S.; and (b) Breakup of Tonnage carried by the different transportation sectors within the U.S.; and (c) Breakup of Ton-miles traveled by the different transportation sectors within the U.S. .	25
2-4	Evolution of the types of trucks used for freight carriage	26
4-1	Example of the possible ways in which a load can be moved from Point A to Point B, at a distance ‘d’ from each other	46
4-2	Example of the distribution of weekly lane demand and the break-down of demand allocation to fleet versus for-hire	47
4-3	Comparison of a generic Fleet and For-Hire Cost Structure for a single out-and-back shipment	50
4-4	Optimal Routing in the absence of continuous moves	52
4-5	Optimal Routing when continuous moves are allowed	53

5-1	A Generic Representation of Wal-Mart's Supply Chain	58
5-2	Aggregated Distribution of all Inbound Lanes from Vendors to a DC per week	60
5-3	Aggregated Distribution of all Outbound Lanes from a DC to Stores per week	61
5-4	Change in the No. of Potential Tours generated by the model	66
5-5	Change in the No. of Tours assigned by the model	66
5-6	Behavior of Total Transportation Costs	67
5-7	Behavior of Tour Efficiency in terms of Percentage Full Loads carried	67
5-8	Behavior of Tour Efficiency in terms of Percentage Loaded Miles	68
5-9	Behavior of Optimality Gap in terms of Total Solution Costs	69
5-10	Behavior of Optimality Gap in comparison to the number of potential tours generated	70
6-1	Hierarchy of transportation planning - from long-term planning, to development of weekly operational plans	74
6-2	Example of weekly demand calculation from randomly generated prob- abilities	77
6-3	An example of demand occurrence during a random week over a tour	82
7-1	Example of a tour with lane details in order to explain metrics	101
7-2	The graph gives a visual indication of the inverse correlation between the Actual Probability of Tour Occurrence and the Tour Fickleness factor calculated	102
7-3	Histogram and CDF of Occurrence of All Tours whilst running the model with Full Flexibility	104
7-4	Histogram and CDF of Occurrence of All Tours occurring within the Stochastic Annual Plan	105
7-5	The graph indicates that about 71% of tours generated on a weekly ba- sis within the Full Flexibility operational plan that are already present in the Stochastic Annual Plan	106

7-6	(a) Distribution of the number of freight lanes present in all tours generated in the Full Flexibility Operational Plan; (b) Distribution of the number of freight lanes present in the Stochastic Annual Plan; and (c) Distribution of the number of freight lanes present in all tours generated in the Full Flexibility Operational Plan which are not present in the Stochastic Annual Plan	107
7-7	The graph depicts the distribution of volume occurrence on the lanes within the network over the 52-week simulation period	109
7-8	The graph depicts the distribution of the lane fleet fickleness factor over the 52-week simulation period	110
7-9	The graph depicts the distribution of the lane fleet fickleness factor over the 52-week simulation period	111
7-10	The graph depicts the distribution of the lane fleet fickleness factor over the 52-week simulation period	112
7-11	Depiction of the use of lane demand statistics as a criterion for lane robustness	113

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

2.1	Top ten Private Fleet Owners in the U.S. during 2009	28
2.2	Top ten TL Carriers in the U.S. during 2009	30
2.3	Top ten LTL Carriers in the U.S. during 2009	31
2.4	Differences between Private Fleet, Dedicated Fleet and For-Hire carrier capacity	34
5.1	Correlation Matrix for a Sample DC	63
5.2	Correlation Statistics for a sample DC	64
6.1	Characteristics of the Mini-Network	76
7.1	Summary of Network-level statistics for key metrics	91
7.2	Facility level metrics for DC6006	97
7.3	Facility level metrics for DC6062	98

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

Trucking plays a pivotal role in the U.S. economy and nearly every good consumed in the country has traveled on a truck at some point. As a result, the trucking industry hauled 68.9% of all freight (in tons) transported in the United States, equating to 9.1 billion tons in 2008 (IBISWorld, 2010). The trucking industry amassed an astounding \$660 billion in revenue during 2008, representing about 83% of the U.S. commercial freight transportation market (Kirkeby, 2010). Put another way, on average, trucking collected 83 cents of every dollar spent on freight transportation.

Trucks transport the tangible goods portion of the economy, which is nearly everything consumed by households and businesses. However, trucking also plays a critical role in keeping costs down throughout the business community. Specifically, for businesses that produce high-value, low-weight goods, inventory carrying costs can be considerable. But, many of these producers now count on trucks to deliver products efficiently and timely so that they can keep stocks as low as possible. In fact, inventory-to-sales ratios continue to fall, indicating that motor carriers and their customers are working well together in this area, saving the economy billions of dollars in costs.

The trucking industry is made up of three major types of players:

- **Private fleet owners** - move cargo solely on privately owned vehicles
- **For-Hire Carriers** - lease trucks to shippers on a short or long-term basis

- **Combination shippers** - own a private fleet, and supplement them with for-hire capacity

All of these players face the fundamental task of anticipating shipping requirements and creating a strategic transportation policy to fulfill the demand occurrences at an operational level. Because of the competitive nature of the industry, carriers are under constant pressure to meet service level targets at the lowest possible costs. However, occurrence of demand over a carrier's network is uncertain in terms of the number of weekly loads to be carried from one point to another. The traditional tools and methodologies commonly used to create the overarching transportation policy do not adequately address this key decision criterion. As a result, this inherent demand variability creates a gap between the forecasted plan for the network and actual weekly execution.

In order to address this gap, there has been ongoing research at the Center for Transportation & Logistics to incorporate demand variability in the transportation planning process, using stochastic planning methodologies. Led by Dr. Chris Caplice and Dr. Francisco Jauffred, the Freight Lab team has created a Freight Network Optimization Tool (FNOT), for deciding the allocation of volume occurrences on a lane to an optimal mix of private fleet and for-hire carriers, as well as finding optimal routes for moving volume over the private fleet. FNOT is a Large-Scale Linear Optimization Model - a detailed discussion about the model and its capabilities is carried out in later chapters. Using the planning scenario created by FNOT, the key questions that this thesis addresses are:

- *How do long-term stochastic planning methodologies perform in the presence of demand variability?*
- *Can simple heuristics be used to reduce the complexity associated with executing a long-term annual plan on a weekly basis?*
- *What are the key metrics for analyzing the performance of the stochastic plan in comparison to the plans created on the basis of heuristics?*

- *What are the trade-offs associated with minimizing costs versus simplicity of execution using heuristics?*

Thesis Summary

For the purpose of this thesis, we shall analyze Wal-Mart's transportation network across the continental US. At later stages, we shall pare the network to a smaller sub-section. This serves the purpose of maintaining the overall integrity and characteristics of a real-world transportation network, while providing us with a workable data set that can be thoroughly analyzed to draw insights.

We start the thesis by providing an overview of the transportation industry in the US, highlighting the importance of trucking in order to get a better understanding of the major stakeholders within the industry and the issues they face. Having identified the key questions in the previous paragraph, we lay down the motivation behind the thesis by discussing issues related to demand variability and the criticality of planning in transportation. After that, we move into an initial data analysis by sifting through and creating a coherent database, as well as looking for patterns that could potentially explain the demand distributions, seasonality and correlations. This lays down the foundation for creating a long-term annual planning model and testing its performance by simulating weekly demand using Monte-Carlo simulations. The next part of the thesis deals with creating heuristics with varying levels of operational flexibility to figure out ways in which the annual plan can be used to satisfy weekly demand in real-time. The scenarios essentially highlight the difference between the best-case-assignment of demand over the network, versus assigning loads using simple heuristics based on the stochastic annual plan. The analysis compares the weekly execution scenarios using various metrics and throws light on the trade-off between dispatching weekly loads with complete flexibility in the most optimal manner possible, versus executing weekly demand using heuristics based on the annual plan. The analysis is carried out at four levels, in decreasing order of transportation network aggregation. In the process of analyzing these scenarios, we create new metrics that provide a

deeper understanding about the weekly operations and help us get a comparison of the robustness of the annual plan and reasons behind it. Finally, we summarize our findings, discuss the implications and make recommendations for creation of a robust transportation plan that addresses demand variability over the network.

As a slight tangent to the demand variability topic, we also analyze the sweet spot for the trade-off between transportation cost savings obtained by relaxing the constraints on certain run parameters and the computational requirements for running the model.

Chapter 2

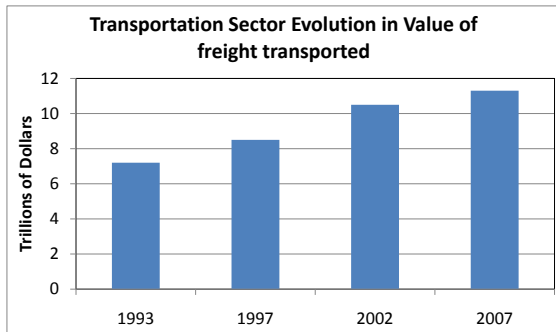
Trucking Industry Overview

This section provides a brief overview of the transportation industry within the US. Using statistical data and reports generated over the years, the aim of this section is to provide an insight into the transportation sector and the role that the trucking industry, in particular, plays in the overall US economy. In addition, we apprise the reader about the main players within the trucking sector to build a basic understanding of the types of fleet and considerations that go into transportation planning.

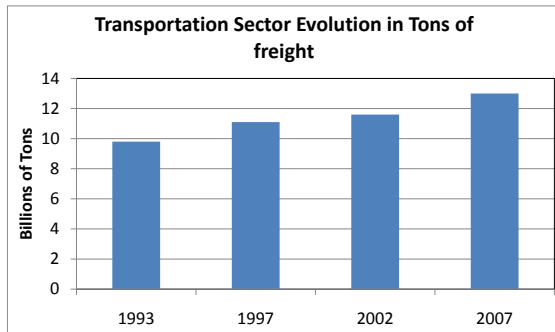
2.1 U.S. Transportation System

More than 13 billion tons of freight, valued at \$11.8 trillion, were transported nearly 3.5 trillion ton-miles in the United States during 2007, according to preliminary estimates from the 2007 Commodity Flow Survey (CFS)(M. Margreta, 2009). The tonnage, value, and ton-miles of 2007 freight shipments all increased over 2002 totals. Tonnage was up 12 percent, inflation-adjusted value up 13 percent, and ton-miles up 11 percent.

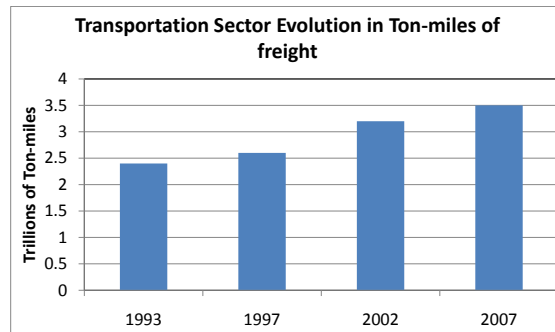
This steady growth in freight movements was possible because of growth in the U.S. economy, an increase in U.S. international merchandise trade, improvements in freight sector productivity, and the availability of an extensive multimodal transportation network in the United States. The statistics for historical market trends within the transportation sector are visible in Fig. 2-1.



(a)



(b)



(c)

Figure 2-1: (a) Evolution of the US Transportation Sector in terms of the Value of freight carried; and (b) Evolution of the US Transportation Sector in terms of the Tonnage of freight carried; and (c) Evolution of the US Transportation Sector in terms of the Ton-Miles of freight carried

According to the US Department of Transportation(USDOT, 2006), the transportation sector moves large volumes of freight to support economic activities in the nation. More than \$1 out of every \$10 produced in the U.S. Gross Domestic Product is related to transportation activity. Transportation employs nearly 20 million people in America - 11 million in direct transportation and transportation-related industries (e.g., pilots, train operators, autoworkers, and highway construction workers) and another 9 million in non-transportation industries (e.g., truck drivers for retail and grocery stores, wholesale shipping clerks, and distribution managers for manufacturing firms).

On a typical day in 2007, over 35.7 million tons of goods, valued at \$32.4 billion, moved nearly 9.6 billion ton-miles on the nation’s transportation network. Nearly 93 percent of the total tonnage and 81 percent of the total value of freight were shipped by means of a single transportation mode, while the remainder was shipped using two or more modes(M. Margreta, 2009).

2.2 Transportation Modes and Industry Trends

Each mode plays an important role in the US freight transportation system - railroads and barges haul bulk commodities and perishable goods over long distances, trucks carry smaller packages to the main streets and back roads of America, and airplanes fly expensive goods overnight across the country. The following discussion is based on the Commodity Flow Survey (CFS) data published by the U.S. Department of Transportation every five years¹. Between 2002 and 2007, shipments by trucks grew the most, measured by value, while rail shipments experienced the highest increase in terms of tons or ton-miles of freight carried. The value of multi-modal shipments increased by 54 percent during this time, followed by increases in pipeline shipments of 51 percent and trucking of 11 percent. By tonnage, multi-modal shipments experienced an increase of 257 percent, followed by rail with 24 percent trucking with

¹Source: U.S. Department of Transportation, Bureau of Transportation Statistics, based on 1993, 1997, 2002 and preliminary 2007 Commodity Flow Survey data plus additional estimates from Bureau of Transportation Statistics

18 percent. And by ton-miles, air cargo grew by 63 percent, followed by truck with 56 percent and multimodal combinations by 37 percent. Water transportation experienced a severe downturn during this period, dropping by 78 percent in terms of tonnage and 69 percent in terms of ton-miles of freight carried. A detailed comparison of the modes of transportation in terms of value, tons and ton-miles carried over the years can be seen in Fig. 2-2.

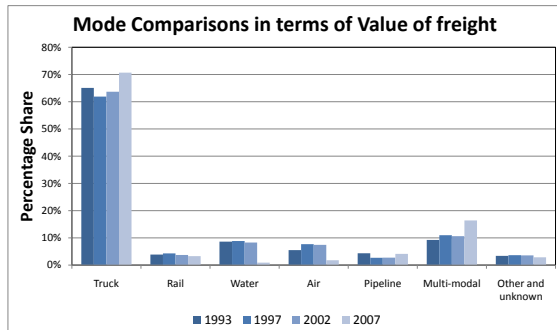
Trucking continued its dominance of the US freight transportation system, as visible in the Fig. 2-3. In 2007, trucks hauled about 71 percent of the value, 69 percent of the tonnage, and 40 percent of the ton-miles of total shipments, exhibiting a positive trend in comparison to 2002 numbers. Measured by ton-miles, trucking was followed by rail at 37 percent and multi-modal at 14 percent. In general, trucking dominated shipment distances of less than 500 miles while rail dominated the longer distance shipments. Multimodal transportation, i.e., shipments moved by more than one transportation mode, grew substantially in value (54 percent) during this period. Of these shipments, parcel, postal, or courier services (typically involving more higher value and smaller size shipments) grew the most rapidly and accounted for over 83 percent of the value of multimodal shipments in 2007. A comparison of the modes of transportation across a variety of metrics is visible in Fig. 2-2.

Thus, it is quite clear that the trucking industry is quite predominant and by far, the biggest and most important mode of transportation for commercial freight within the US.

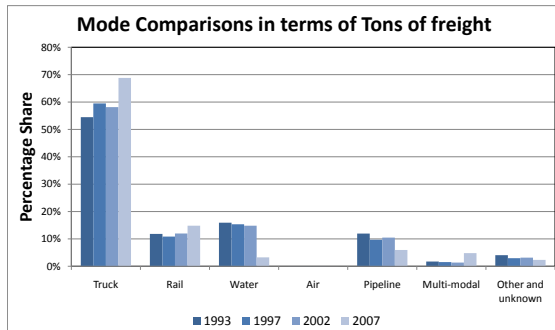
2.3 Trucking Industry

Trucking is the largest mode in both value of shipments handled and tonnage. According to 2007 CFS preliminary data(M. Margreta, 2009), truck shipments accounted for:

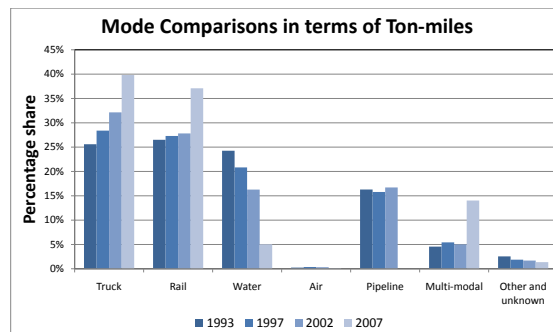
- about \$8.4 trillion worth of goods, an inflation-adjusted gain of 9.1 percent from 2002, and 71 percent of the total value of all shipments



(a)



(b)



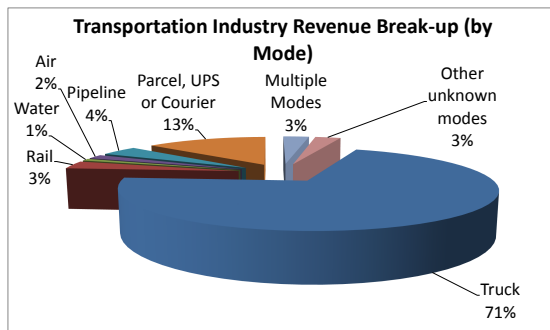
(c)

Figure 2-2: (a) Break-up and evolution of the modes of transportation within the US in terms of value of freight carried; and (b) Break-up and evolution of the modes of transportation within the US in terms of tonnage of freight carried; and (c) Break-up and evolution of the modes of transportation within the US in terms of ton-miles of freight carried

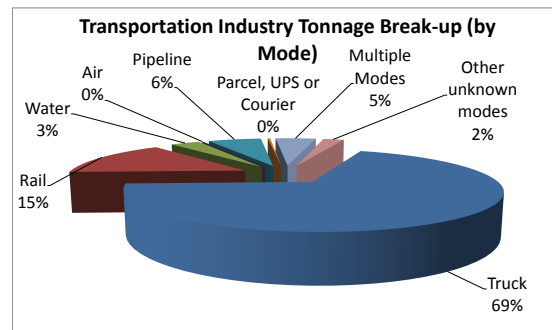
- about 9.0 billion tons of goods, an increase of 14.2 percent from 2002, and 69 percent of all tonnage;
- about 1.4 trillion ton-miles, representing 40 percent of all ton-miles; and
- an average distance of 187 miles per shipment

In 2007, trucking (both for-hire and private) continued its dominance of the freight industry, moving 71 percent of the nation's commercial freight, measured by value, and 69 percent of the tonnage. However, by ton-miles, trucks moved just slightly more than rail, 40 percent compared to 37 percent, followed by multi-modal shipments at 14 percent. These numbers show a faster growth in shipments by truck, compared with rail, and the decline in water transportation since 2002. Truck ton-miles grew by 24 percent, rail by 33 percent, and water declined by about 69 percent. A decade ago, trucks moved almost 28 percent of ton-miles and rail moved about 27 percent, followed by water with 21 percent and pipeline with 16 percent. Fig. 2-3(M. Margreta, 2009) indicates the dominance of trucking in comparison to other modes of transportation, as far as revenues, tonnage and ton-miles are concerned.

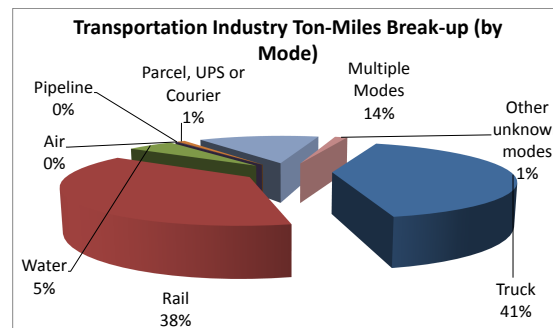
In recent years, as trucking maintained its dominance, the number of trucks traveling on the nation's highways steadily increased and the truck fleet mix changed(USDOT, 2006). While two-axle single-unit trucks are the most common commercial trucks on the nation's roads, the number of larger combination trucks grew at a much faster rate, increasing about 59 percent over this period, compared to 30 percent for single-unit trucks. In 2003, combination trucks accounted for 28 percent of the commercial truck fleet, up from 24 percent in 1980. These larger trucks also travel more miles per vehicle than the single-unit trucks. Combination trucks generated a total of 138 billion vehicle-miles of travel (VMT) in 2003, compared to 78 billion miles by single-unit trucks. Since 1980, overall truck vehicle-miles have doubled from 108 billion to 216 billion in 2003. Despite this growth in truck VMT, commercial truck's share of total highway vehicle-miles remained steady, hovering between 7.1 and 7.5 percent over this period. This was primarily because travel by all highway vehicles, including passenger cars, buses, and light trucks (e.g., pickup trucks, sport utility vehicles, and



(a)



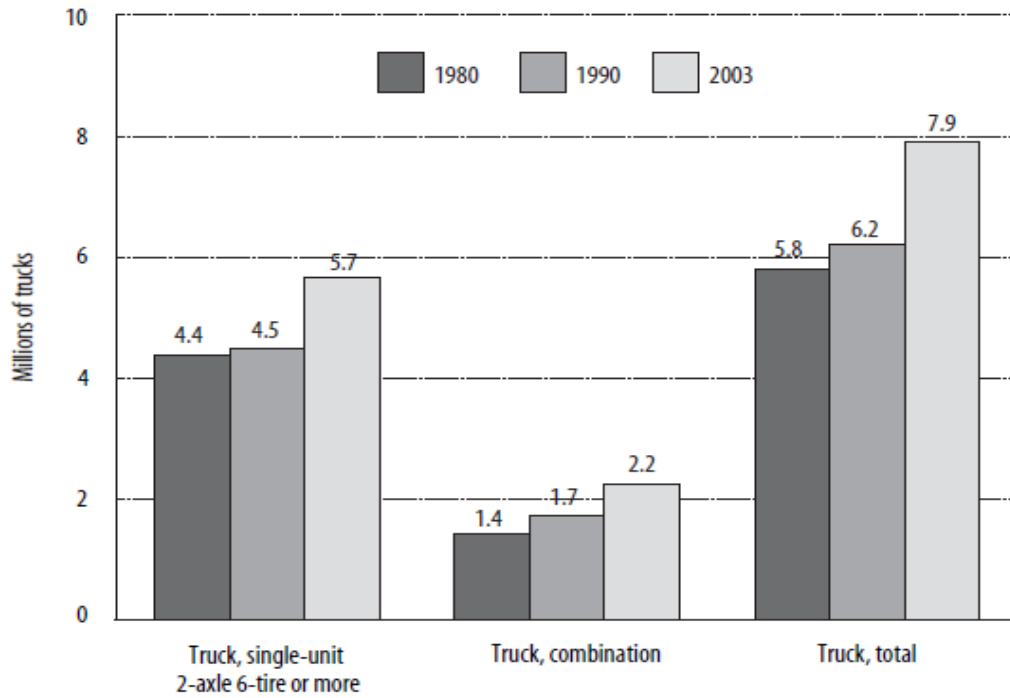
(b)



(c)

Figure 2-3: (a) Breakup of Revenues generated by the different transportation sectors within the U.S.; and (b) Breakup of Tonnage carried by the different transportation sectors within the U.S.; and (c) Breakup of Ton-miles traveled by the different transportation sectors within the U.S.

minivans) also grew at a similar pace. The evolution of the various types of trucks and their growth is indicated in the Fig. 2-4(USDOT, 2006).



NOTE: Total trucks exclude light trucks such as pickup trucks, sport utility vehicles, and minivans.

SOURCE: U.S. Department of Transportation, Research and Innovative Technology Administration, Bureau of Transportation Statistics, *National Transportation Statistics 2005*, available at http://www.bts.gov/publications/national_transportation_statistics/ as of August 2005.

Figure 2-4: Evolution of the types of trucks used for freight carriage

2.4 Types of Trucking Fleets

In order to understand how to create a well conceived transportation policy, it is imperative that the reader have an adequate understanding of how private, dedicated and for-hire carriers differ in terms of cost structures and services provided. This section will provide a high level overview of each of these different types of transportation.

2.4.1 Private Fleet

A private fleet is owned and operated by the shipping entity, whose primary business is something other than transportation, such as Wal-Mart, Target, etc. The principle objective of the fleet is to support the shipper's internal distribution requirements. The shipper leases or owns the physical assets such as tractors, trailers and/or straight trucks. The drivers are normally employees of the company. Private carriers are a major part of motor carriage operations. Although little financial information is available on private carriage, the American Trucking Associations estimates that companies running their own shipping operations provided services valued at some \$288 billion in 2008, or about 44% of the motor carriage market(Kirkeby, 2010). According to estimates from the National Private Truck Council, a trade group, private fleets operate more than two million trucks, make up about 82% of the medium- and heavy-duty trucks registered in the United States, and account for around 56% of all freight tonnage carried by medium- and heavy-duty trucks.

The Table 2.1 lists the ten largest private fleets owned and operated within the US, during 2009, as per FleetOwner.com (2010).

Table 2.1: Top ten Private Fleet Owners in the U.S. during 2009

Private Fleet Owners	Total Vehicles	Straight Trucks	Tractors	Trailers
AT&T	78,070	78,000	70	22,000
Verizon Communications	44,973	44,858	115	116
Pepsi Bottling Group	38,500	-	-	-
Republic Services	22,582	21,960	622	1,378
Waste Management	22,000	20,895	1,105	4,305
Time Warner Cable	18,010	18,000	10	440
Coca-Cola Enterprises	17,400	9,500	7,900	10,000
PepsiCo.'s Frito-Lay	17,109	17,109	-	-
Tyco International	15,600	15,600	-	-
The ServiceMaster Co.	15,450	15,450	-	1,890

According to Mulqueen (2006), customer service is the preeminent driver of these types of fleets, and is given priority over optimal routing sequences and vehicle utilization. Because of the nature of this business, routes are typically constrained more by delivery time than physical space on the vehicle. Optimization is typically used during tactical planning in order to create driver territories and establish route sequences, but due to the expectation of high service levels, static route planning is often used for daily execution.

For-Hire carriers are usually not considered in these operations due to the need for extremely high level of customer service. Fleets are also much more cost effective than for-hire carriers in these operations due to the high density of low volume stops. This allows a single fleet truck to support the delivery requirements of many of the wholesale distributor's customers.

Our research is more applicable to shippers that use private fleets for longer haul, full truckload transportation. This is common in many industries including grocery retail (Albertsons, HEB), big-box retail (Wal-Mart, Target) and food/consumer packaged goods (CPG) manufacturers (P&G, Kimberley Clark). In this segment, for-hire carriers are often times used to supplement fleet capacity. Typically, for-hire carriers are more efficient due to advantages they hold in terms of economies of scale and scope, however, many shippers still view fleets as an important component in their overall transportation strategy.

As pointed out by Mulqueen (2006), the most common reasons cited for having a fleet include:

- Better perceived service to their customers. Fleet drivers are viewed as important assets in maintaining a strong shipper/customer relationship.
- Fleet drivers can be requested to perform special services during the delivery that for-hire carriers will not do or would do only for an additional charge.
- More leverage with contract carriers during rate negotiations by sending a message to the carrier that it can be replaced by an internally managed fleet.
- Marketing advantages of having the shipper's name on the trailer, thereby acting as a rolling billboard for the company. Provides assurance of freight capacity times of tight capacity, such as exist in the current environment.
- More control over transportation operations.

2.4.2 For-Hire Carriers

The for-hire carrier industry is extremely fragmented, with 240,000 for-hire truckload carriers in the United States in 2010, of which the top five carriers contribute only 13.5% to the overall industry revenue (IBISWorld, 2010). The activities of this industry can be broadly segmented into consignments weighing greater than 10,000 pounds known as 'truckload' and less than 10,000 pounds (less-than-truckload).

Truckload carriers dedicate full trucks to one customer and make deliveries of goods from start to finish. Less-than-truckload (LTL) carriers take partial loads from multiple customers on a single truck and then route the goods through a series of terminals where freight is transferred to other trucks with similar destinations. LTL transportation providers consolidate numerous orders generally ranging from 100 to 10,000 pounds from varying businesses at individual service centers in close proximity to where those shipments originated. Utilizing expansive networks of pickup and delivery operations around these local service centers, shipments are moved between

origin and destination utilizing distribution centers when necessary, where consolidation and de consolidation of loads occurs.

The industry is dominated by truckload carriers. IBISWorld (2010) estimates that more than 80.0% of all establishments are truckload carriers contributing 71.1% of total revenue. LTL is more labor intensive, employing 35.0% of all employees in the industry, costing 40.0% of total wages. This disproportional cost reflects the higher labor intensity of LTL load transport. Revenue is generated per ton hauled; It takes the same labor to ship a 5,000 ton load as 10,000 ton load, but two trips (double the labor) is often required to generate the same revenue.

The Tables 2.2 and 2.3 provide the rankings for the top ten TL and LTL Carriers in the US respectively, along with their revenues during 2009 (Schulz, 2010).

Table 2.2: Top ten TL Carriers in the U.S. during 2009

Top TL Carriers	Revenue (in \$ millions)
Swift Transportation	2489
Schneider National	2380
Werner Enterprises	1433
U.S. Xpress Enterprises	1333
J.B. Hunt Transport Services	1204
Prime Inc.	992
C.R. England	866
Crete Carrier Corp.	849
CRST International	610
Knight Transportation	585

A For-Hire TL carrier is contracted by outside organizations to move freight at a pre-determined rate and operate in environments where loads are greater than 10,000 lbs, which is the approximate breakpoint where the variable nature of LTL costs begin to exceed the fixed nature of TL costs. These carriers pick up freight at the origin point and move it to the final destination without any intermediate loading and unloading of the shipment, although the shipper can contract the TL carrier to perform multiple pickups and/or deliveries under the same bill of lading. This is markedly different from the LTL and parcel transportation network models, which utilize hub and spoke systems that require multiple transfer points to move product

Table 2.3: Top ten LTL Carriers in the U.S. during 2009

Top LTL Carriers	Revenue (in \$ millions)
FedEx Freight	3618
YRC National	3177
Con-way Freight	2574
UPS Freight	1807
ABF Freight System	1260
YRC Regional	1226
Estes Express Lines	1174
Old Dominion Freight Line	1158
R+L Carriers	862
Saia Motor Freight Line	794

from the origin to the ultimate destination.

The primary benefit of a hub and spoke network is that it enables consolidation of shipments going between terminals. This benefit is not recognizable in a TL environment, since the vehicle is, theoretically, already fully utilized, and injecting a full TL into a hub and spoke network would simply add additional transit time and handling expense to the process.

Additionally, capacity commitments are often specified within TL contracts. From the shipper's perspective, capacity commitments require the carrier to cover a certain number of loads on a given lane over a specified period of time. This is often done by requesting that the carrier agree to haul a set percentage of total load volume on each lane. This, in theory, provides the shipper the capacity needed to manage the weekly variations in load volumes that a fixed volume commitment would not support.

One important facet of TL transportation that needs to be recognized is that unlike LTL or parcel carriers, For-Hire TL carriers will often reject undesirable loads; even those loads under contract. This occurs if the carrier does not have available capacity or, as carriers get more technologically savvy the load is deemed operationally unprofitable given the current location and status of the carrier's assets. The frequency of carriers turning down loads has increased in recent years as US domestic TL capacity has tightened and the carriers have begun to exert their new found power in the buyer/seller relationship. Harding (2005) showed that the cost of a turndown

was estimated to be between 2% and 7% of the freight spend. In this analysis, over 25% of tendered loads under study were rejected by the primary carrier. The effect of declined freight is discussed extensively by Harding (2005). This tendency of For-Hire carriers increases the uncertainty related to available capacity for the shipper, and can potentially have dire consequences for the latter's transportation plan. However, for the purpose of this thesis, we shall take the relationships and contractual agreements between the shippers and carriers as a given, and assume that the loads assigned to carriers are never rejected, or are carried by an alternate carrier at the same price.

2.4.3 Dedicated Fleet

Unlike private fleets, dedicated fleets are not owned by the shipper, but are provided on an exclusive basis to the shipper for a contractually specified period of time. Most large Truckload (TL) carriers like Schneider National, JB Hunt, Swift and Werner have active and growing dedicated fleet businesses. The advantages of a dedicated fleet over a private fleet is that a dedicated fleet does not require a large capital expenditures outlay as is required when expanding the capacity of a shipper's private fleet. Dedicated fleets provide the advantages of guaranteed capacity in constrained markets and increased control over the asset and its usage. This has become an important advantage as shippers compete with each other for the available TL capacity on the market (Bradley, 2005).

Like private fleets, shippers will incur dedicated fleet costs regardless of whether the assets are used, since a large component of the cost is fixed versus load based, as in the case of contract carriage. Idle fleet assets and excessive dwell time are especially costly for both dedicated and private fleets. Additionally, the variable component of dedicated contracts typically includes per mile charges that are incurred regardless of whether the vehicle is loaded, thereby penalizing inefficient shippers, as defined by the percentage of empty miles in their networks.

Dedicated and private fleets are both most effective when the shipper has the ability to maximize equipment utilization by minimizing deadhead distance and dwell time as well as by utilizing the assets on shorter-distance runs that incur high mini-

mum fees when executed by for-hire carriers. While there are key differences between private and dedicated fleets, this thesis views them as fundamentally the same. Both modes have significant sunk costs and are perishable in nature in the sense that if the capacity is not used, it is lost. This contrasts with for-hire carriers, which are paid only when used to execute the movement of a load.

Private Fleet and Third Party Freight

A benefit that some companies take advantage of with regard to their private or dedicated fleets is the ability to generate revenue by moving freight for other shipping entities. As the TL capacity in the United States continues to become scarce and rates are driven up, it has become economically compelling for private fleets that have significant empty miles built into their network to acquire common carrier authority and move other shipper's freight. This industry trend is cyclical, and moves with the macro-economic conditions of supply and demand within the industry. The ability of private fleets to carry third party loads not only helps to generate extra revenue for the company, but also serves to reduce overall empty miles traveled by the fleet and in turn, increases asset utilization. This enables the fleet to generate revenue and turn into a profit center for the shipping organization. This facet of private fleet operation is critical to our thesis. The ability to carry third-party freight helps private fleet owners reposition their trucks back to the domicile, incurring minimal empty miles along the way. This is done by creating multi-legged tours for the fleet. However, there is a cost versus service trade-off involved in using the private fleet as third party carriers, since overdoing it might affect overall company service levels. This trade-off needs to be kept in mind at all times and freight carriage needs to be prioritized accordingly. The thesis uses a limited application of this ability of private fleet carriers, by allowing them to haul their own freight from vendors to the distribution centers. This serves a dual purpose. On one hand, it helps bring in additional revenue for the company and act as a profit center. On the other hand, it can also be justified from a service stand-point since the fleet is carrying the company's freight, leading to a more reliable service without compromising on the uncertainty associated with hauling third-party

freight, which might bring in additional demand variability into the network.

Fifty-six percent of private fleets operate with common carrier authority today (Terrerri, 2006), although it is not known what percentage of these fleets are actually moving third party freight since there are benefits aside from generating revenue that drive a fleet to attain common carrier authority. For the purpose of this thesis, we do not consider the fringe benefits associated with common carrier authority, and concentrate on it solely as a revenue generating and operational efficiency mechanism.

2.5 Summary

In summary, we can see that the trucking industry is fragmented with a large number of players, because of minimal barriers to entry and exit. We also need to remember that, in practice, all private/dedicated fleet owners utilize for-hire carrier capacity to some extent in order to meet emergency requirements. For the purposes of this thesis, it is critical to note that the private and dedicated fleet pay carriage charges for all the miles traveled by the fleet, while the for-hire carriers charge customers according to the origin-destination distance. In addition, the available fleet capacity for a private fleet at domiciles remains static for a private fleet, change on a periodic basis (based upon contractual agreements) for a dedicated fleet, while it is completely flexible for for-hire carriers. These differences can have significant consequences while creating a transportation plan, and must be taken into account. However, for this thesis, we are assuming that the dedicated fleet and private fleet is equivalent. The Table 2.4 highlights the important differences between private, dedicated and for-hire fleet for the purposes of this thesis.

Table 2.4: Differences between Private Fleet, Dedicated Fleet and For-Hire carrier capacity

Characteristics	Private	Dedicated	For-hire
<i>Capacity Flexibility</i>	Static	Changes per contract (monthly)	Completely flexible
<i>Capacity Availability</i>	Always available	Based on contract	Not always available
<i>Payments based on</i>	All miles traveled	All miles traveled	Direct distance and per mile rate
<i>Min. charge for load carriage</i>	Low	Low	High
<i>Control over operations</i>	Full control	Moderate control	Low control (for shippers)

Chapter 3

Thesis Motivation: Demand Variability and Transportation Planning

The following chapter provides a synopsis of the types of variability inherent in a transportation network and the importance of addressing demand variability in order to create a robust transportation plan. The need for incorporating these key decision criteria for creating a robust transportation plan is the main motivation behind this thesis.

3.1 Variability

One of the critical aspects that this thesis tries to address is variability. There are many connotations and meanings to the term “variability” - It is derived from the word variable which is from the Old French *variable*, from the Latin *variabilis* “changeable” from *variare* “to change”. As far as transportation planning is concerned, variability can occur in three major forms: demand, supply and geography. However, for the purpose of our thesis, we shall only be concentrating on the demand variability, since the remaining aspects of variability do not have a significant impact on the network under consideration.

- **Supply variability**

This is the type of variability wherein there are constant fluctuations in the availability of fleet or for-hire trucks in order to carry the shipments from one point to the other. This aspect of variability is not a major concern for Wal-Mart because of the large contingent of their private fleet and contractual relationships with for-hire carriers, but could be a crucial consideration for other, smaller shippers.

- **Geographic variability**

Changes in the distribution network with regard to vendor, distribution center and store locations can potentially create geographic variability over the transportation network. This would mean that, over time, transportation lanes may be added or removed from the existing distribution network, as demand patterns change. However, this variability has a relatively long time horizon, and hence, is not critical to the planning time frame that we are addressing in our thesis. Thus, we assume a static distribution network that does not change over the time frame for all practical purposes.

- **Demand variability**

This is the type of variability wherein there are constant changes in the volume of shipments moving on a given transportation lane over a pre-defined period of time, usually assumed as a week. This aspect of variability can have tremendous impacts associated with creating a strategic plan for long-term fleet capacity as well as developing contracts with third party carriers.

For the purpose of this thesis, demand variability is the most critical aspect, and shall be dealt with in detail. Most contemporary off-the-shelf software tools neglect this aspect of variability and instead use a single value of ‘average loads per week’ for planning purposes. However, the underlying variability on the lanes leads to executional strategies that are quite different in reality.

The following thesis particularly addresses planning problems faced by combination carriers, who own a private fleet, and use for-hire carriers to supplement their capacity. Thus, they basically incur a fixed cost associated with the private fleet, and variable for-hire costs for having flexible capacity for meeting demand fluctuations over the year. The two classic cases that could be the result of demand variability for these shippers are discussed below:

- **Actual demand is higher than planned**

This would lead to short-term dependence on for-hire carriers to carry the freight, which would have a dual impact. Firstly, for-hire carrier charges would be a lot more for hauling freight on an ad hoc basis. Secondly, there would be uncertainty associated with the availability of for-hire carriers on a short-term basis, which would directly affect service levels.

- **Actual demand is lower than planned**

This would lead to under-utilization of the private fleet and would affect the shipper in the form of wasted sunk costs related to idle fleet and driver capacity.

Typically, over the course of the year, both of these cases will occur multiple times because of demand variability creating volume fluctuations with respect to the planned capacity. Thus, in order to mitigate these losses, unexpected, short-term decision making takes place at an operational level, which leads to system sub-optimality. These decisions create a wide gap between predicted costs during long-term strategic planning and actual costs incurred. Additionally, this short-term sub-optimal decision making might also affect overall company service levels, in cases where demand cannot be met in time. One of the major points that this thesis is trying to address is to minimize this gap between planning and execution.

3.2 Criticality of Transportation Planning

Because of the inherent variability in transportation demand, creating a transportation policy that at least closely simulates real-time operations is quite critical. The gap

between planning and execution can be addressed by creating a robust transportation policy that gives due consideration to variability over the distribution network and proposing a plan that is as close to real-life execution as possible. This plan should be capable of incorporating the operational uncertainties associated with demand, as a criterion for decision making. The tools and methodologies currently available do not adequately address this key decision criterion for creating an overarching transportation policy, which is the main motivation behind the research carried out in this thesis.

Transportation planning must be capable of satisfying long-term as well as short-term decisions effectively. Let us consider the decisions that companies hauling freight by using a mix of privately owned fleet and for-hire carriers need to make, in the long as well as the short-term.

3.2.1 Long-term Planning

This level of planning is required for deciding the overall transportation strategy of a company. The time horizon associated with these decisions could be anywhere between one to five years. Long term decisions usually have large amounts of capital expenditure and sunk costs associated with them. Hence, having a robust long term plan that would be operationally viable would help a company in the following ways:

- **Making decisions for expenditure related to fleet acquisition**

Depending upon the demand variability over the transportation network and policies related to satisfying demand using the private fleet, a company needs to decide upon their fleet size. Also, since the average life span of a truck is around five years, this decision needs to be robust, in order to utilize the fleet efficiently in the long run.

- **Allocation of fleet to distribution centers based on tour plans**

Depending upon the long-term projected demand flowing into or out of a distribution center, a decision needs to be made for allocating fleet capacity to the DC, such that it is capable of handling demand requirements adequately.

- **Fixed/ Sunk costs associated**

There could be fixed or sunk costs related to construction of shipping docks and other miscellaneous expenses that need to be taken depending upon the transportation requirements.

- **Hiring of Drivers**

Since driver salaries are a fixed cost, the hiring decision heavily depends upon the projected utilization that a company hopes to get out of their drivers. Idle drivers lead to a lot of wasted resources, and hence, this decision needs to be made carefully. In order to mitigate these losses, most firms have a mix of full-time employees based on long-term demand patterns and contract-based employees to satisfy demand fluctuations in the short run.

- **Long Term contracts with for-hire carriers to get better rates for specific lanes**

In order to ensure that a company gets the best for-hire lane rates, it is important to build long term relationships and contracts with the carriers. If a load is assigned to a for-hire carrier on an ad-hoc basis, they might qualify the load as an emergency shipment, and charge a high amount to move the load. If the contracts for load carriage by for-hire carriers are made well in advance, it would not only ensure that a for-hire truck is available whenever needed, but would also enable the company to get lower rates.

3.2.2 Short-term planning

After figuring out the long-term planning decisions for the transportation network, a short-term plan is needed to make the daily or weekly decisions for operational managers, in order to ensure optimal utilization of available resources. The two main decisions that need to be made over a short-term are:

- **Allocation of available fleet capacity to optimal routes, based on weekly demand**

Thus, depending upon the demand, the operational managers need to make a decision on how to route the private fleet in the most optimal manner such that the loads are moved in the most cost-effective way and meet the requisite service level targets set by the company.

- **Short term allocation of capacity to for-hire in order to meet unmet weekly demand**

In cases where demand fluctuations are so heavy that the load occurrences cannot be satisfied by the private fleet, the operational managers need to make alternate arrangements for moving the load, using short-term for-hire contracts.

Demand variability has a huge role to play in the short-term as well as long-term decision making process. If the variability is not handled effectively by the transportation plan, there tends to be a significant gap between the expected plans and the actual operations. This divide leads to extra costs and low service levels at the operational level.

Hence, there is an inherent trade-off involved between -

- Optimizing a complex transportation plan to generate maximum theoretical savings by creating best possible scenarios on an everyday operational basis
- Adequately addressing the uncertainties, and creating a simplified plan that is easy to execute

It is our hope that the proposed techniques in this thesis shall balance this trade-off and help create a robust long-term transportation plan that actually materialize in practice, by taking the demand uncertainties into account.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Large-Scale Linear Optimization Model for Transportation Planning

Mulqueen (2006) laid down the groundwork regarding the importance of addressing demand variability while creating a transportation policy. He created a two-step strategy for transportation planning - the first part involved creating a deterministic linear programming model to assign volumes to tours based on user-defined confidence levels; the second part of the strategy involved testing the optimization results with random demand simulations in order to understand the behavior of the deterministic optimization policy in a real-life scenario. This comparison indicated that the savings predicted by a deterministic transportation policy were highly dependent upon user-defined confidence intervals and hence, not robust enough to emulate real life scenarios. These results were crucial in reiterating the true stochastic nature of transportation lane volumes. The shortcomings pointed out by Mulqueen (2006) led to further research into the area and paved the path for creation of the Freight Network Optimization Tool (FNOT), which forms the base of this thesis. The following chapter provides a high-level summary of the methodology used by FNOT for creating a transportation policy that addresses demand variability. In the first part, we discuss the allocation of lane volume to an optimal mix of fleet and for-hire carriers. In the next part, we discuss the formation of tours for volume to be carried over the private fleet.

4.1 Freight Network Optimization Tool (FNOT)

Realizing the importance of addressing demand variability in transportation networks, there has been ongoing research at the Center for Transportation & Logistics at the Massachusetts Institute of Technology for developing a tool that addresses this uncertainty. Led by Dr. Chris Caplice and Dr. Francisco Jauffred, the Freight Lab team has come up with a Linear Programming Optimization model called the Freight Network Optimization Tool (FNOT)(F. Jauffred, 2010). FNOT is basically is a steady state model that conducts large scale linear optimization. The objective of the model is to assign shipments over the distribution network at an optimal mix of private fleet and for-hire carriers while minimizing the overall costs for the network. Additionally, in order to reduce the number of empty miles travelled by the private fleet within the network, the model also looks for tours (a set of continuous moves for the fleet) that can satisfy lane demand in the most cost-effective way, and assigns the lane volume accordingly. In order to address the demand variability aspect, FNOT has the capability of modeling the demand distributions over the transportation lanes using stochastic distributions like Normal, Poisson, etc. as well as by using histograms of demand formulated from empirical data. Furthermore, FNOT also has the capability of breaking down a freight lane into a set of two or more moves, called relays, if it is cost-effective to do so. Unlike Mulqueen’s multi-stage optimization and simulation approach, FNOT integrates stochastic conditions within the optimization methodology to find the optimal transportation policy that takes demand uncertainty into account.

4.2 Understanding FNOT

The objective function behind the FNOT engine is to minimize total transportation costs over the distribution network while optimally assigning volumes to the private fleet and for-hire carriers, based on some supply-demand constraints.

Objective function: Minimize (Total Fleet Costs + Total For-Hire Costs + Relay Costs)

Subject to the constraints:

- All lane volume must be satisfied by either fleet or for-hire
- Fleet assignments must not exceed available fleet or driver capacity
- Actual lane volume assigned by the model to private fleet must be between the overall minimum and maximum values of lane demand

FNOT assigns lane volume to a mix of fleet and for-hire carriers in a unique manner. This calculation is carried out using a modification of the Newsvendor formulation for transportation planning purposes. The essence of this calculation is the use of loss functions for for-hire volume assignment. In statistics, a loss function represents the loss (cost in money or loss in utility in some other sense) associated with an estimate being different from either a desired or a true value, as a function of a measure of the degree of deviation (generally the difference between the estimated value and the true or desired value). In our case, the loss function represents the degree with which our estimate about volume planned on a lane to be covered by private fleet is lower than the actual volume occurring on the lane, i.e. the expected volume not covered by the private fleet. As such we have to incur a penalty by fulfilling the remaining demand using for-hire carriers, and thereby paying an extra cost to them over and above what was originally planned for.

Let X be the optimal fleet demand allocation point on the empirical histogram. That is, whenever demand occurs on this freight lane which is less than or equal to X , it makes economical sense to assign the loads to private fleet, and all demand above it should go to for-hire carriers. The value of X would be our final decision in creating

a volume assignment plan for fleet and for-hire carriers for this transportation lane. The underlying question is, “How do we find X ?”

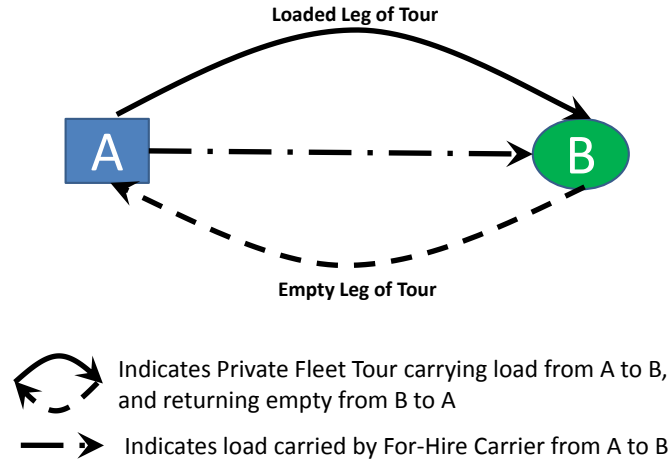


Figure 4-1: Example of the possible ways in which a load can be moved from Point A to Point B, at a distance ‘d’ from each other

For the sake of simplicity, let us consider the example shown in Fig. 4-1, consisting of a transportation lane from point A to B. Demand occurring on this lane over a particular week can be transported by using either the private fleet or a for-hire carrier option. The private fleet option would be to send a truck from A to B and bring it back to A empty. The second option would be to use a for-hire carrier to ship the load.

This lane has an empirical histogram of demand information, as shown in the Fig. 4-2 (Note, that FNOT is capable of producing the same analysis using probability distributions as well). The histogram presents a pictorial representation of the portion of lane volume that shall be carried over the private fleet, with the excess demand (beyond X) assigned to for-hire. This allocation is obtained by comparing the costs for carrying the load using private fleet with for-hire carrying costs, in conjunction with the demand information available through the histogram.

Let

c = Cost per mile for carrying the load using the private fleet (\$ /mile)

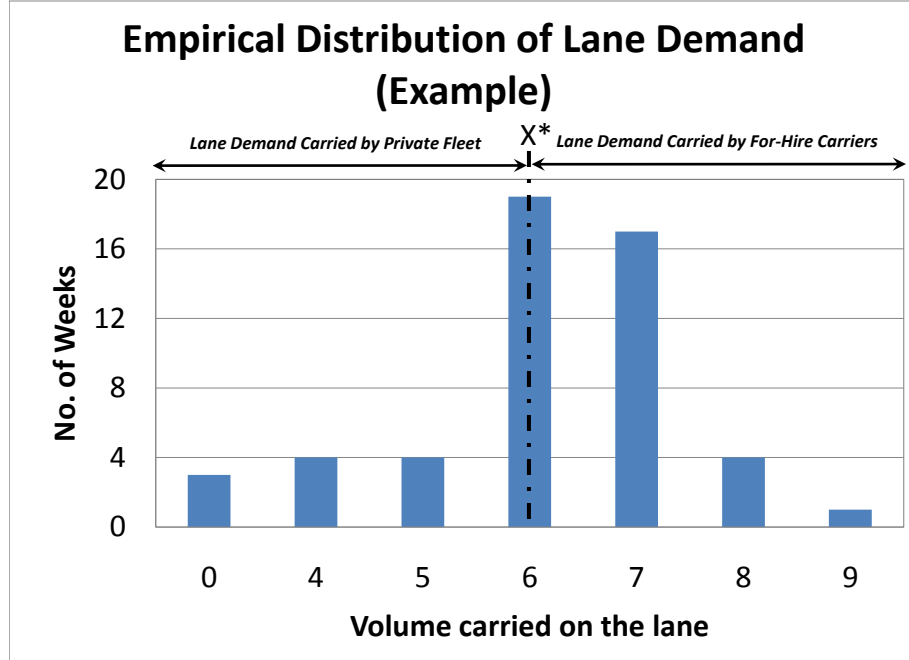


Figure 4-2: Example of the distribution of weekly lane demand and the break-down of demand allocation to fleet versus for-hire

d = Distance between points A and B (miles)

p = For - Hire cost per mile for carrying a load from point A to B (\$ /mile)

X = Planned optimal fleet demand allocation point

D = Actual value of demand on a lane

$P(D > X)$ = Probability that the actual demand occurring on the lane is greater than X

In the event that demand for the lane is less or equal to the fleet planned volume, the entire demand can be satisfied by the private fleet.

Hence,

$$\text{Total Transportation Cost} = \text{Fleet Cost} = 2cdX$$

In the event that demand for the lane is greater than the fleet planned volume,

the excess demand needs to be satisfied using for-hire carriers.

Hence,

$$\text{Fleet cost} = 2cdX$$

$$\text{For-Hire cost} = pd \int_X^\infty P(D > \psi) d\psi$$

$$\text{Total Transportation Cost (TTC)} = \text{Fleet Cost} + \text{For - Hire Cost}$$

$$= 2cdX + pd \int_X^\infty P(D > \psi) d\psi$$

As per the definition of Loss function discussed previously, the extra costs related to meeting the unplanned excess demand on the transportation lane corresponds to for - hire carriage costs. This loss would only be incurred if the planned optimal fleet demand allocation is smaller than the actual demand occurrence.

Hence, the loss function formula can be generically stated as,

$$\text{Loss function} = E[\max(0, D - X)] = pd \int_X^\infty P(D > \psi) d\psi$$

As per the Newsboy Equation, the optimal point would occur when the Total Transportation Cost function reaches its minimum; i.e.,

$$\text{Marginal Cost of Private Fleet} - \text{Marginal Cost of For - hire carriers} = 0$$

$$\frac{\partial}{\partial X}(2cdX) - \frac{\partial}{\partial X}pd \int_X^\infty P(D > \psi) d\psi = 0$$

$$2cd - pdP(D > X) = 0$$

$$P(D > X) = \frac{2c}{p}$$

Using this probability value, the optimal X can be easily figured out from a cumulative empirical distribution of lane demand that we possess, as shown in Fig. 4-2. The Loss function is continuous and convex for all positive values of X thus the total

transportation cost, $TTC(X)$, is continuous and convex in the same domain. In consequence $TTC(X)$ always has at least one optimal minimum value.

Intuitively, this formulation makes sense as well. Marginal costs are defined as the change in total cost incurred when the quantity produced increases by one unit. In our case, the marginal costs for private and for-hire fleet can be considered as the effect that one extra capacity over the private fleet or for-hire carriers has, on the total costs of the network. When demand is below optimal (X), we have idle fleet capacity that adds to the total costs, since it is not fully utilized. When demand is above optimal, we have to incur costs related to allocating the excess demand to for-hire carriers. Hence, the only way we can have minimal impact on the total cost equation is to obtain the condition where we are indifferent between having an extra capacity of fleet or for-hire carriers available to us, since our demand is being met optimally by the private fleet, i.e. marginal costs of private fleet equals marginal costs of for-hire carriers. Thus, generalizing the discussion, we can say that the allocation of freight to for-hire carriers depends on the loss function of the demand distribution.

We know from the calculations above that demand allocation to fleet or for-hire carriers is essentially dependent upon the costs associated. Hence, let us consider the extreme cases for fleet versus for-hire costs, and see how it affects our assignments.

Case 1: *For-Hire costs (p) are substantially higher than Total Fleet Costs ($2c$)*

Hence,

$$\frac{2c}{p} \approx 0 \Rightarrow P(D > X) \approx 1$$

Therefore, all demand occurring on the lane would be assigned to the private fleet.

Case 2: *For-Hire costs (p) are substantially lower than Total Fleet Costs ($2c$)*

Hence,

$$\frac{2c}{p} \approx 1 \Rightarrow P(D > X) \approx 0$$

Therefore, all demand occurring on the lane would be assigned to the for-hire carriers.

After deciding the optimal loads on all freight lanes within the network, the next part of the model deals with trying to find continuous moves for the private fleet over the network, to further minimize the costs.

If we were to consider using a private fleet solely for out-and-back moves (as illustrated in the example shown in Fig. 4-1), we place a great limitation on the effective use of the private fleet to carry loads over longer distances. This is because of the cost structure generally used in the transportation industry for private fleet and for-hire carriers. The behavior of a generic fleet and for-hire carriage cost structure with respect to distance traveled over a lane is indicated by the Fig. 4-3.



Figure 4-3: Comparison of a generic Fleet and For-Hire Cost Structure for a single out-and-back shipment

This dictates that for-hire carriers would make more economic sense if we were to compare them to any out-and-back tours that are longer than 95 miles (one way). The reason is that, for out-and-back tours, 50% of the tour cost is incurred because of empty miles travelled by the fleet in order to return to its origin. However, if we

were to somehow find other possible lane volumes that the fleet could satisfy through continuous moves, we could minimize the overall empty percentage miles travelled by the truck. This would make a strong case for allocating more loads to the private fleet and increase asset utilization simultaneously. For the example shown in Fig. 4-4, if continuous moves are disallowed over the network, the optimal allocation for volume occurring over the network is broken down into a for-hire move (A-B) and an out and back tour (B-C-B). The associated costs are:

A-B: Carrier costs, which would be the maximum of \$250 (minimum amount charged by the for-hire carriers for carrying this load) and distance times the for-hire rate per mile based on the carrier bid (\$2.21 per mile for this lane). The distance traveled over this lane is 368 miles, which results in for-hire costs of \$813.28 per load.

$$\text{Carrier Cost} = \text{Max}[250, \text{Distance} \times \text{For - Hire Cost per mile}]$$

$$\text{Carrier Cost} = \text{Max}[250, 544 \times \$ 2.21] = \$ 813.28$$

B-C-B: Private fleet costs are calculated as the total distance traveled by the fleet times the per-mile fleet costs incurred by the company. The distance from B-C is 272 miles. Hence, the total fleet costs for the tour B-C-B would be the maximum of out-and-back distance times the per-mile fleet costs and the minimum charge for moving a load (\$50, as decided by the private fleet owners). This leads to the overall tour costs equating to \$652.80 per load.

$$\text{Fleet Cost} = \text{Max}[50, \text{Distance} \times \text{Fleet Cost per mile}]$$

$$\text{Fleet Cost} = \text{Max}[50, 272 \times 2 \times \$ 1.2] = \$ 652.80$$

Thus, for this example, the overall costs to the company for moving the two loads over the network would be $\$813.28 + \$652.80 = \$1466.08$.

However, when continuous moves are allowed, both loads can be carried by a private fleet truck on a tour A-B-C-A, as shown in Fig. 4-5, wherein the truck has to travel empty from C-A. At the fleet rate of \$1.2 per mile, the overall costs for carrying

both loads over the network come up to $1.2 \times 786 \text{ miles} = \943.20 . This results in overall savings of \$523. Thus, by actively looking for all possibilities of linked tours that can potentially reduce overall network costs, we can create a lot of savings and at the same time, improve asset utilization.

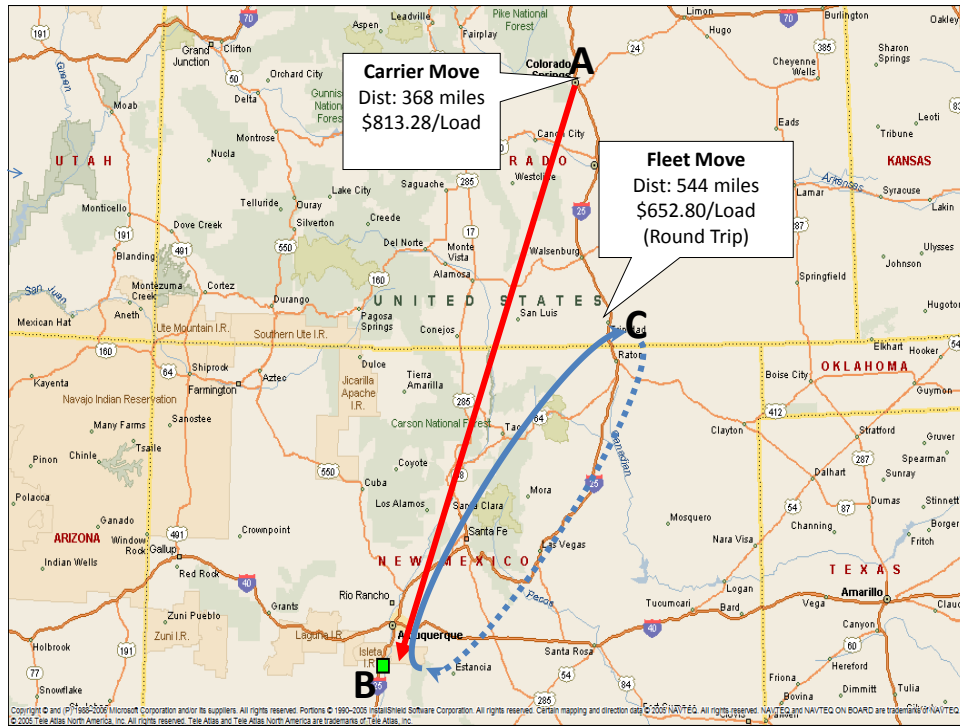


Figure 4-4: Optimal Routing in the absence of continuous moves

In summary, the FNOT model is a tool created for meeting the following goals in the process of producing a transportation plan:

- Allocate loads occurring over lanes within the network to an optimal mix of fleet and for-hire carriers, while taking demand uncertainty into account
- Create a chain of continuous moves (tours) to satisfy the demand carried by the private fleet in a cost-effective manner
- Minimize the overall empty miles travelled by the private fleet, and in the process, also increase asset utilization
- Look for optimal relay points within the network, which again serve the purpose of increasing asset utilization while minimizing transportation costs



Figure 4-5: Optimal Routing when continuous moves are allowed

The following is the actual large scale optimization linear programming model, which is part of the working paper (F. Jauffred, 2010).

$$\begin{aligned}
 & \text{MIN} \sum_{t \in \{\text{Tours}\}} c_t x_t + \sum_{\alpha \in \{\text{Freight Lanes}\}} p_\alpha S_\alpha + \sum_{\alpha \in \{\text{Freight Lanes}\}} \sum_{l_i \in \{\text{Truck Lanes}\}} \sum_{l_o \in \{\text{Truck Lanes}, l_i \neq l_o\}} q w_\alpha(l_i, l_o) \\
 & \text{s.t.} \\
 & S_\alpha \geq \text{Loss}_\alpha [y_{\alpha,i}] - P_\alpha (U > y_{\alpha,i}) \left(v_\alpha + \sum_{l_i \in \{\text{Truck Lanes}\}} \sum_{l_o \in \{\text{Truck Lanes}, l_o \neq l_i\}} w_\alpha(l_i, l_o) - y_{\alpha,i} \right) \\
 & \forall \alpha \in \{\text{Freight Lanes}\}, \forall i \in \{\text{Samples of } \alpha\} \\
 & \sum_{t \in \left\{ \begin{array}{l} \text{Tours adjacent} \\ \text{to lane } l \end{array} \right\}} x_t = v_l + \sum_{\alpha \in \{\text{Freight Lanes}\}} \left(\sum_{l_i \in \{\text{Truck Lanes}\}} w_\alpha(l_i, l) + \sum_{l_o \in \{\text{Truck Lanes}\}} w_\alpha(l, l_o) \right) \\
 & \forall l \in \{\text{Truck Lanes}\} \\
 & \sum_{t \in \left\{ \begin{array}{l} \text{Tours starting at} \\ \text{Domicile } d \end{array} \right\}} H_t x_t \leq L_d \quad \forall d \in \{\text{Domiciles}\} \\
 & \text{FleetMin}_\alpha \leq v_\alpha + \sum_{l_i \in \{\text{Truck Lanes}\}} \sum_{l_o \in \{\text{Truck Lanes}, l_o \neq l_i\}} w_\alpha(l_i, l_o) \leq \text{FleetMax}_\alpha \quad \forall \alpha \in \{\text{Freight Lanes}\}
 \end{aligned}$$

where:

x_t is the number of loads per week in tour. It is always equal or greater than zero

S_α is the Expected number of carrier loads in lane α

v_α is the total number of fleet loads per week in freight lane α moved without relays.

It is always equal or greater than zero

$w_\alpha(l_i, l_o)$ is the total number of fleet loads per week in freight lane α relayed using truck lane i and o as inbound and outbound relay lanes respectively. It is always equal or greater than zero

c_t is the cost of tour

p_α It is the penalty associated with the loss function in lane α . In general it will be the cost per load of a for-hire carrier or the allowance per load in a prepaid lane

H_t is the number of driver-hours required to complete tour t

L_d is the Total number of driver hours per week available at domicile d

$FleetMin_\alpha$ and $FleetMax_\alpha$ are the Minimum and Maximum fleet loads per week in lane α

S_α is an approximation of the loss function in lane α

y_α is the total number of fleet loads per week in lane α

P_α is the cumulative demand probability distribution for lane α

f_α is the demand density function for lane α

D_α is the average demand in lane α

The optimization model is coded in C++ using the CPLEX callable library; and data handling is carried out using INFORMIX, with Microsoft Access as the interface. At a high level, the following are the inputs and outputs generated by the model, which shall be analyzed further in the rest of this thesis.

The annual plans generated by FNOT form the crux of this thesis. Although FNOT addresses demand variability over the network before deciding the optimal allocation of volume, we need to be aware that FNOT is a steady state model. This means that in the presence of available weekly demand characteristics, FNOT creates the best possible allocation of volume for that particular week. Hence, it is important to test the robustness of the plans generated, in terms of its capability of handling real-world demand variability on a week-to-week basis. Keeping the annual plans generated as a basis, we create various weekly planning scenarios and test their

performance at an operational level using Monte-Carlo simulations.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

The Data and Initial Analysis

The following chapter begins by building an understanding about Wal-Mart's supply chain and the original data provided by them for analysis. Additionally, it discusses the motivation behind consolidating the data to create a workable dataset. Using this dataset, we carry out an initial analysis of the data to develop an understanding about the demand occurrences with respect to correlations, seasonality and distributions.

The raw, shipment level national data for Wal-Mart consisted of approximately 5.2 million shipments annually over their distribution network. Each shipment is representative of a full truckload (TL) movement from an origin location to a destination within the continental US. The basic supply chain for Wal-Mart can be represented by the flow diagram indicated in Fig. 5-1:

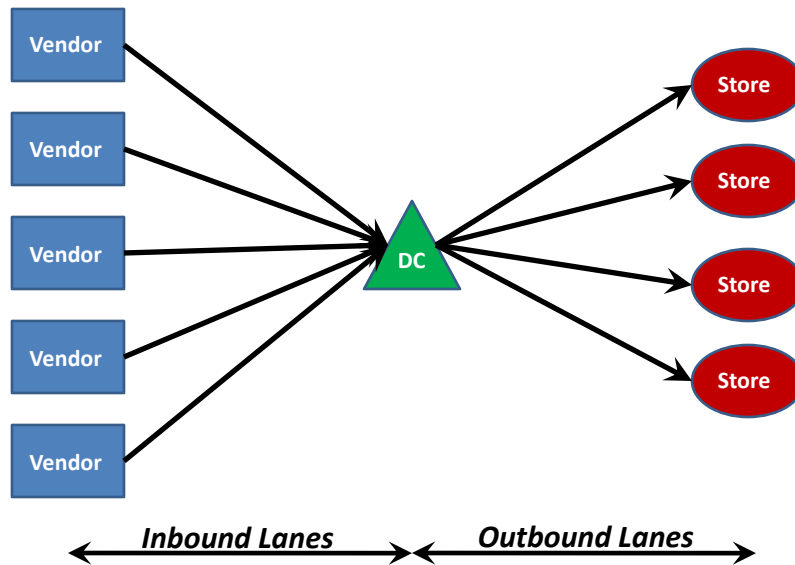


Figure 5-1: A Generic Representation of Wal-Mart's Supply Chain

Fig. 5-1 is a pictorial depiction of freight movement over Wal-Mart's network, as it is forwarded by the vendors, consolidated at the distribution centers and finally sent out to stores. Thus, the overall supply chain can be broken down into an inbound lane, signified by a Vendor - DC move, and an outbound lane, denoted by a DC - Store move.

5.1 Data Consolidation

Because of the large number of shipments and the highly dispersed nature of locations, the shipment level data was consolidated on the basis of zip codes. This reasoning behind it was to achieve a more manageable dataset size without losing its geographic essence. The demand data was also aggregated into weekly buckets of volume occurrence.

The precision level of zip codes within the continental US increase with the number of digits added to the right. To give the reader an example let us consider the 3 digit

zip code 021 representing the area encompassing Boston, Cambridge as well as the neighboring towns of Braintree, Waltham, etc in the state of Massachusetts. Thus, a 3-digit zipcode encompasses an area of much larger magnitude and is relatively vague in nature. On the other hand, the zip code 02139 indicates the smaller subset of Cambridge within this large area. Using the same principle, the vendors are consolidated at the level of a 3-digit zip code, whereas the stores were aggregated at a 5-digit zip code level.

After consolidation, the overall distribution network consists of 124 Distribution Centers and about 7000 vendors and stores all over the continental US, which combine together to form a total of approximately 21,000 transportation lanes.

The movement of cargo over the distribution network is quite heavy, with approximately 120,000 full truckload movements per week. This translates into approximately 30,000,000 weekly loaded miles traveled by trucks all over the continental US. In order to give you a perspective about the magnanimity of the miles traveled, consider this - the distance from the center of the earth to the moon is just 238,000 miles!!

As such, the network was pared down to a workable size at different times in the project, in order to dig deeper and draw generalizations from the analysis. The smaller networks selected shall be discussed at appropriate places in the thesis.

5.2 Lane Demand Distributions & Goodness of Fit

It is clear from the discussion about variability in demand that the assumption about averaging out weekly loads carried on transportation lanes every week, results in sub-optimal transportation plans in actual practice. In order to adequately address this variability, annual freight lane volumes need to be modeled stochastically.

During the initial part of the project, much work was done in order to try to understand demand patterns occurring over the lanes. This would help in drawing generalizations for demand distributions of lanes and using them in conjunction with FNOT's ability to model stochastic demand. In order to do that, a small sample of

data for inbound and outbound lanes for a particular DC was collected and analyzed for demand distributions over the annual period. The first step in this analysis was to characterize lanes based on the average weekly volumes and standard deviations over the lanes. The Figs. 5-2 and 5-3 indicate the distribution of inbound and outbound volume at a sample Distribution Center (DC) within the network.

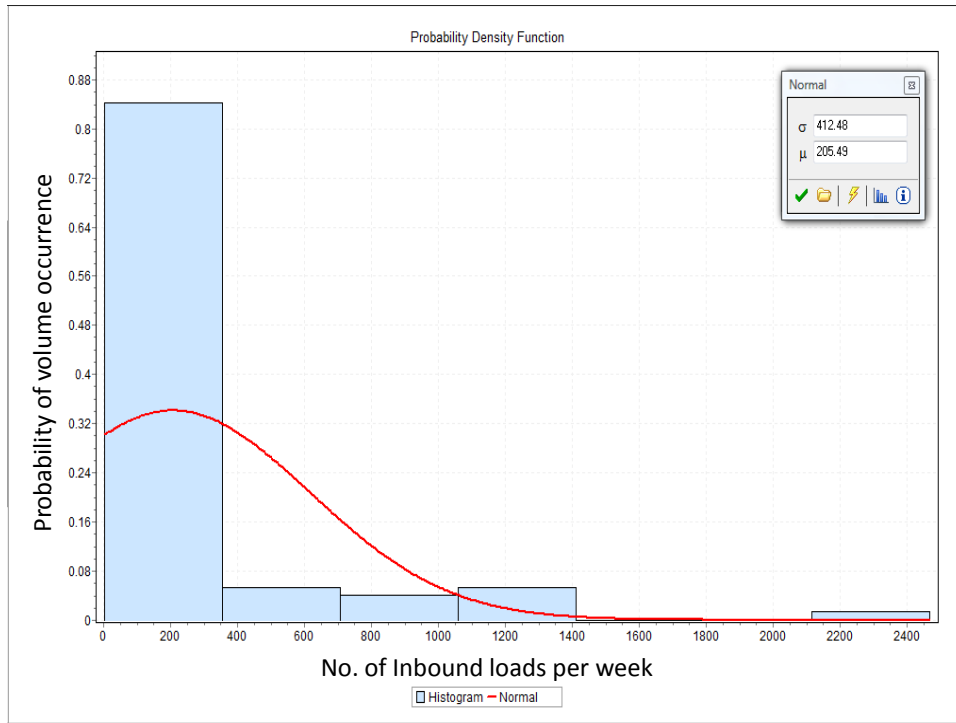


Figure 5-2: Aggregated Distribution of all Inbound Lanes from Vendors to a DC per week

A general insight formed from this analysis was that the volume on the inbound lanes is sparser as compared to the outbound lane volume, which shows a much more consistent behavior. Qualitatively, this behavior makes sense as well, since vendor volumes are expected to be sporadic, based upon distribution center requirements. On the other hand, loads from multiple vendors get consolidated at the DC and are sent out to stores at regular intervals. Hence, the outbound lanes are relatively more consistent in volume occurrence in comparison to inbound lanes.

The next step in the analysis was to compare the lane volume histograms with standard distributions in order to check for goodness of fit. This would provide insights useful in modeling the lane volumes using probability distributions at a gen-

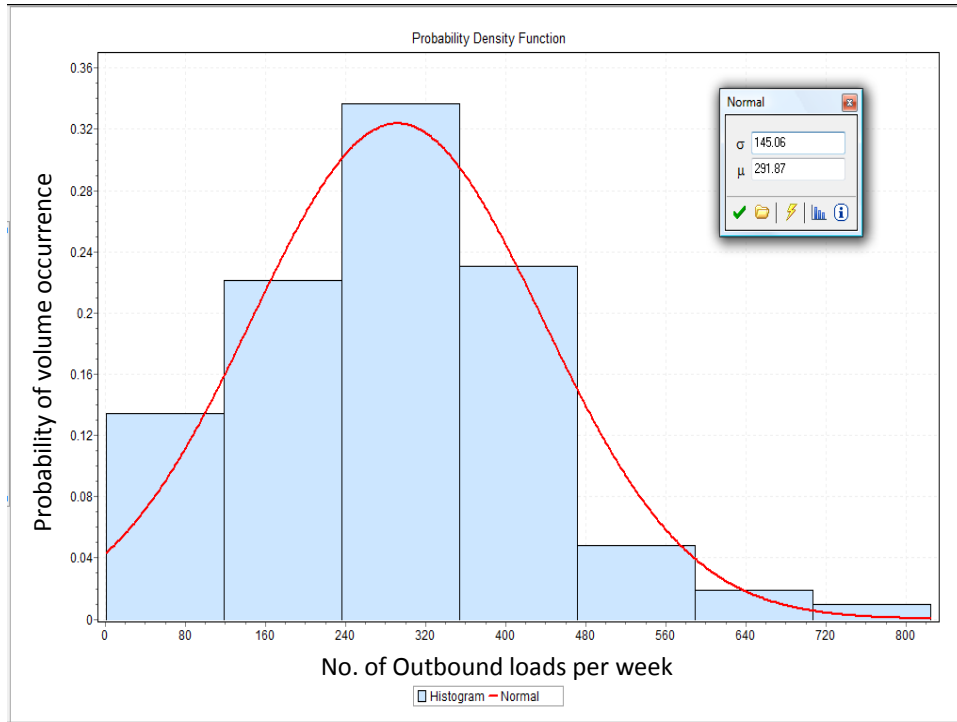


Figure 5-3: Aggregated Distribution of all Outbound Lanes from a DC to Stores per week

eralized level. Using the chi-squared goodness of fit test, we were able to examine whether a particular probability distribution could be used to model the volume occurrence over a freight lane, with a certain confidence level.

It is clear from the dispersed nature of volume occurrences over the lanes that all freight lanes cannot be modeled using a single generic distribution. The conclusions from our research, combined with the findings by Iliadis (2009) suggest that using empirical histogram information would be the best way to model freight lane demand. This is because all lanes have their own peculiarities about volume occurrence that cannot be modeled using a single, generalized probability distribution.

5.3 Seasonality in Demand

One of the possible reasons for variability within the demand data could be seasonality of demand, which is the result of external factors like promotions, holiday seasons, etc. This would create a spike or dip in demand which could be attributed to known

factors. Capturing this information would help reduce the unexplainable portion of demand uncertainty. Thus, we analyzed lane demand to find out if the transportation network exhibited certain peculiar characteristics that could be attributed to seasonality. Comparing lane volumes over the year, it was noted that there was approximately a 10% demand spike in the weeks before Thanksgiving. However, since multi-year data was not available, it was relatively difficult to provide conclusive evidence about this demand spike. Also, the spikes and dips in volumes were not consistent enough across lanes to conclusively prove any definite seasonality in the data set. Hence, we neglected seasonality as a possible cause for demand uncertainty and used the data as-is.

5.4 Correlations in Demand

Correlations play an important role in allocating fleet capacity to a DC. Correlations are indicators of a predictive relationship between demand occurrences that can be exploited in actual practice. Thus, if two freight lanes are statistically proven to be perfectly negatively correlated, it would be an indicator that in the event of volume occurring on one lane, there would be zero volume on the other. Thus, while planning fleet capacity, we can assign a smaller fleet to a Distribution Center with negatively correlated lanes. This phenomenon is conceptually similar to risk pooling.

In order to further analyze and figure out effect of correlations over our network, we selected a sub-set of DC's having an equal mix of heavy volume, medium volume and low volume lanes, and tested the correlations on the inbound and outbound lanes. The correlations were calculated using the Pearson's coefficient, shown below:

$$\rho_{XY} = \frac{\text{cov}(X, Y)}{\sigma_X \sigma_Y} = \frac{E[(X - \mu_X)(Y - \mu_Y)]}{\sigma_X \sigma_Y}$$

The Table 5.1 gives an example of a correlation matrix created for a low volume DC for its outbound lanes.

The next step was to segment the correlations of lanes into 8 main categories:

Table 5.1: Correlation Matrix for a Sample DC

D6022	<i>S44720</i>	<i>S46140</i>	<i>S46214</i>	<i>S48193</i>	<i>S55438</i>	<i>S60466</i>
<i>S44720</i>	1.00					
<i>S46140</i>	0.04	1.00				
<i>S46214</i>	0.00	0.08	1.00			
<i>S48193</i>	0.04	0.22	(0.15)	1.00		
<i>S55438</i>	(0.02)	0.43	(0.12)	0.22	1.00	
<i>S60466</i>	0.01	(0.01)	(0.06)	0.47	0.00	1.00

- Low positive correlation (0 to 0.3)
- Medium positive correlation (0.3 to 0.5)
- High positive correlation (0.5 to 0.7)
- Very high positive correlation (0.7 to 1)
- Low negative correlation (0 to -0.3)
- Medium negative correlation (-0.3 to -0.5)
- High negative correlation (-0.5 to -0.7)
- Very high negative correlation (-0.7 to -1)

These correlation calculations were carried out for approximately 50 different distribution centers, and statistics were computed to understand the behavior of these correlations.

The results shown in Tables 5.1 and 5.2 are representative of the results observed in all the DC's that were tested. They indicate that demand on a high percentage of lanes was either independent or had a very low correlation. Hence, we decided to go ahead with the assumption that the lane demands within the network are independent of each other, without loss of generality.

Table 5.2: Correlation Statistics for a sample DC

Correlation types		Count	Percentage
Positive Correlations	Low(0 to 0.3)	5	33%
	Medium (0.3 to 0.5)	2	13%
	High (0.5 to 0.7)	0	0%
	Very High (0.7 to 1.0)	0	0%
Negative Correlations	Low (-0.3 to 0)	8	53%
	Medium (-0.5 to -0.3)	0	0%
	High (-0.7 to -0.5)	0	0%
	Very High (-1.0 to -0.7)	0	0%
Total		15	100%

5.5 Sweet Spot on Optimality Curve

The following section discusses the relationship between solution optimality for a large-scale linear program and the computational requirements associated with it. When working with routing problems, computing resources and time taken to obtain a solution is exponentially proportional to network size, since routing problems are considered as NP-Hard(J.K. Lenstra, 2006). Additionally, as the solution nears optimality, there is ever decreasing marginal benefit obtained from higher computing resources beyond a certain point. This point is generically called the “Sweet Spot” on optimality curves. It has been shown through research (M. Arnold, 2001) that selective optimization strategies work just as well in obtaining close-to-optimal solutions in much less time. Hence, for our thesis as well, we pursued the idea of selectively constraining the model to look for close-to-optimal solutions.

It has been explained previously in Section 4.2 that FNOT looks for all possibilities of tours that can route volume over the selected network, and includes only tours that reduce the overall objective function costs in the actual results. As network size increases, the number of tours that FNOT needs to look for increase exponentially as well. Constraining the possibilities of tours that FNOT has to look for can have a direct positive impact on computation time. This is carried out by manipulating the run parameter of empty percentage miles allowed on a tour, to find out what would

be the sweet spot of the run parameter. By finding the sweet spot, we can find a close-to-optimal solution in much less time, as compared to looking at all possible tour options.

In order to do this, we selected a smaller sub-section of the overall US network to carry out a deeper analysis of the effect of empty percentage miles on a tour on the solution generated by FNOT. We re-ran the model multiple times, by changing the empty percentage miles factor in steps of 5%, and computed the following outputs from the solution of each run. The impact of each output is discussed subsequently.

5.5.1 Metrics

The following section lays down the definitions of the metrics used to analyze the results, along with the graphs showing the outputs. The next section analyzes these metrics and their impact on our understanding of the sweet spot.

- **Potential Tours generated by FNOT** - is the total number of all possible tour routes generated by FNOT. This includes tours that may be profitable enough to be included in the final solution, as well as those that are eventually discarded as being unprofitable.
- **Tours assigned by FNOT** - is the total number of tours that FNOT recognized as profitable, and included in the final solution.
- **Total Transportation Costs for the Network** - are the combined fleet and for-hire costs for moving loads over the network.
- **Tour Efficiency in terms of percentage full loads carried** - is an indicator of the efficiency of a tour. Whenever a truck has to travel empty from a destination point to an origin point to pick up another load and continue the tour, it results in empty miles being logged, which, in turn, reduces the overall efficiency of the private fleet.

$$\text{Tour Efficiency in \%age full loads carried} = \frac{\text{No. of Full Loads carried by Fleet}}{\text{Total No. of Origin to Destination moves made by fleet}}$$

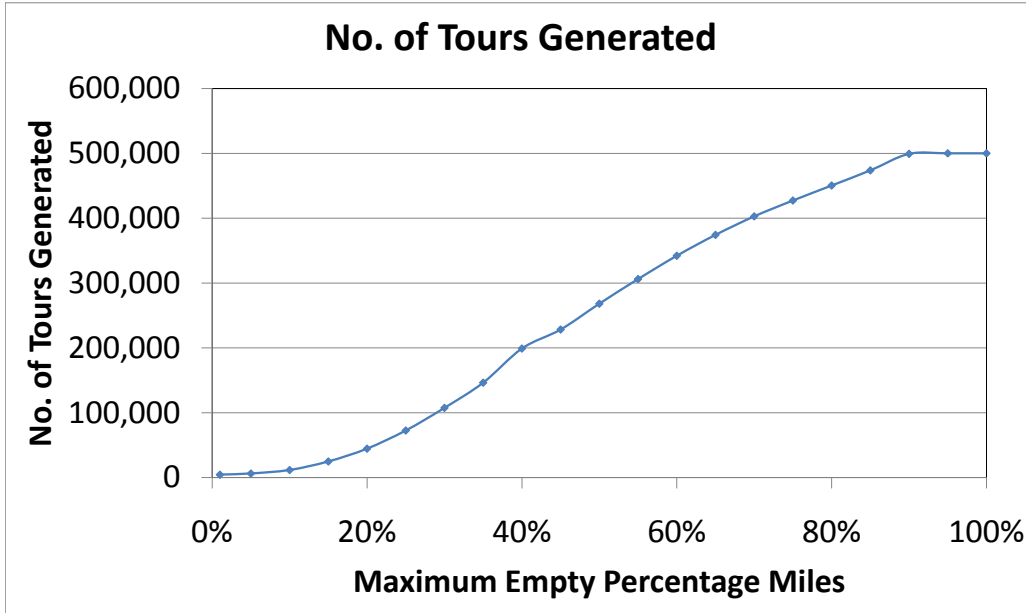


Figure 5-4: Change in the No. of Potential Tours generated by the model

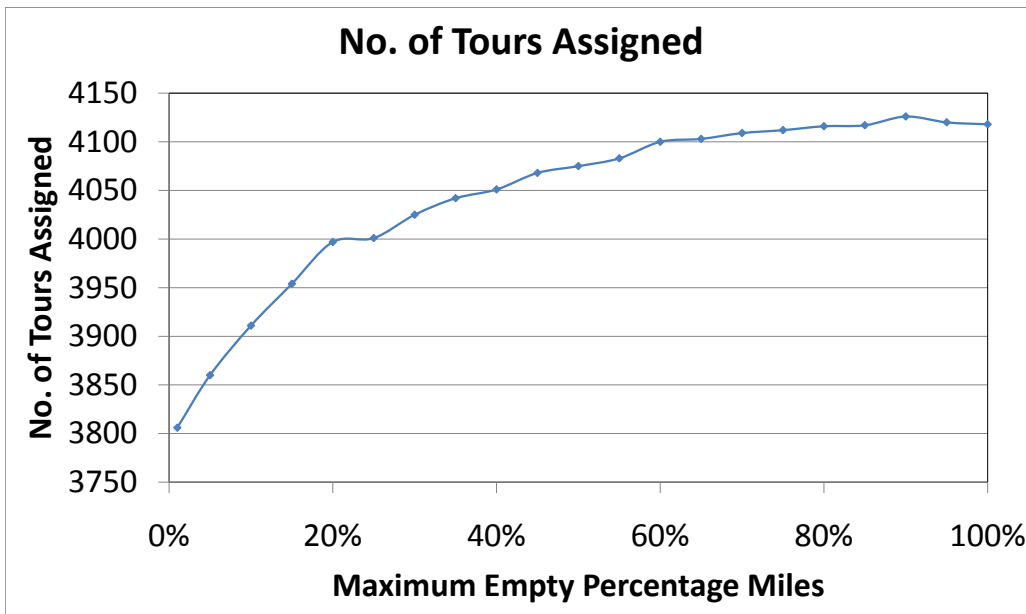


Figure 5-5: Change in the No. of Tours assigned by the model

- **Tour Efficiency in terms of percentage loaded miles** - is another indicator of the efficiency of the private fleet. It is measured as a percentage of miles traveled by the fleet while carrying a load, to the overall miles traveled by the

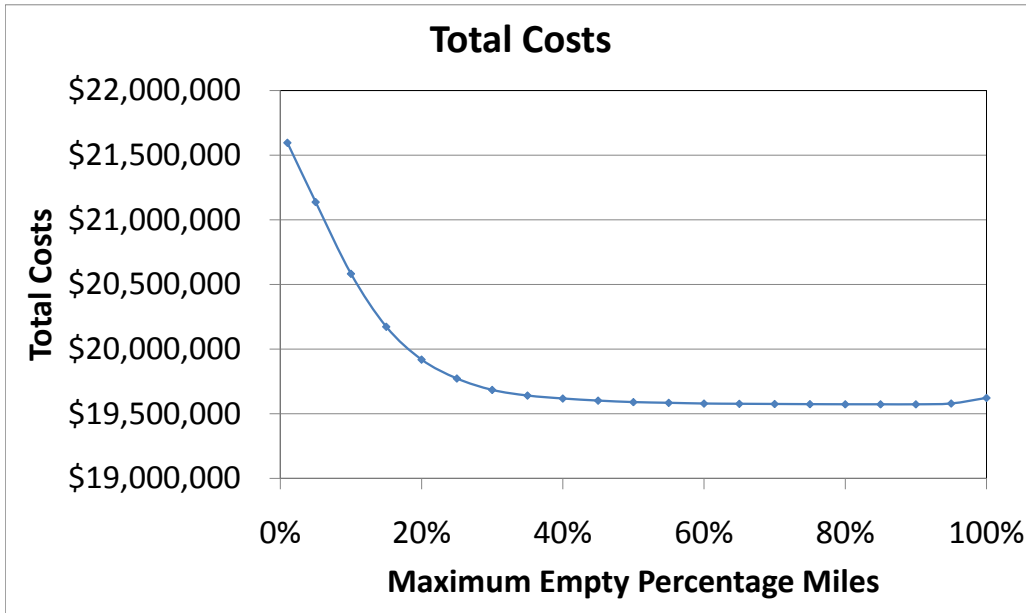


Figure 5-6: Behavior of Total Transportation Costs

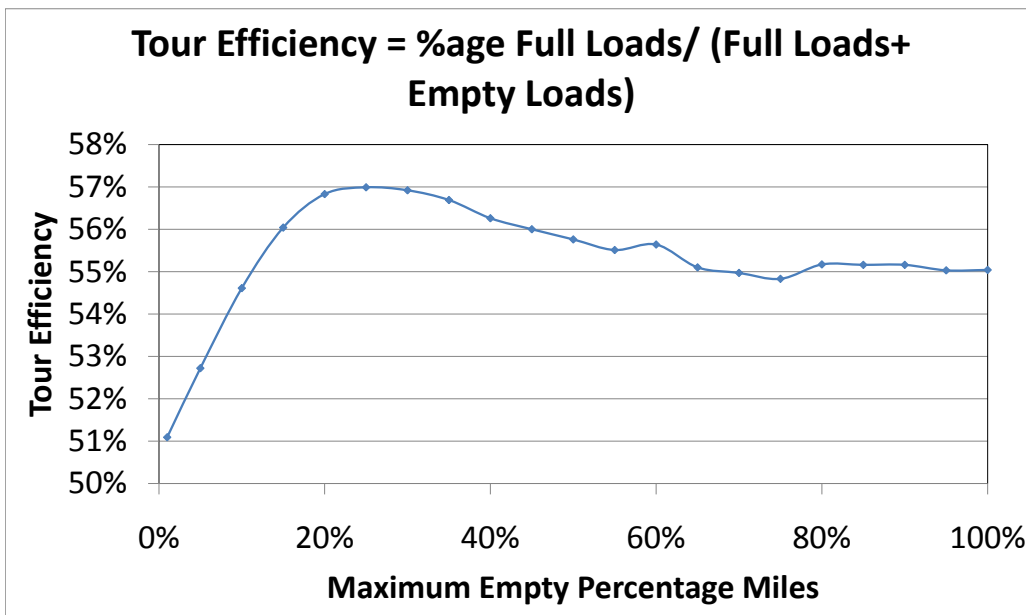


Figure 5-7: Behavior of Tour Efficiency in terms of Percentage Full Loads carried

fleet.

$$\text{Tour Efficiency in \%age miles traveled} = \frac{\text{No. of Loaded miles traveled by Fleet}}{\text{Total No. of miles traveled by Fleet}}$$

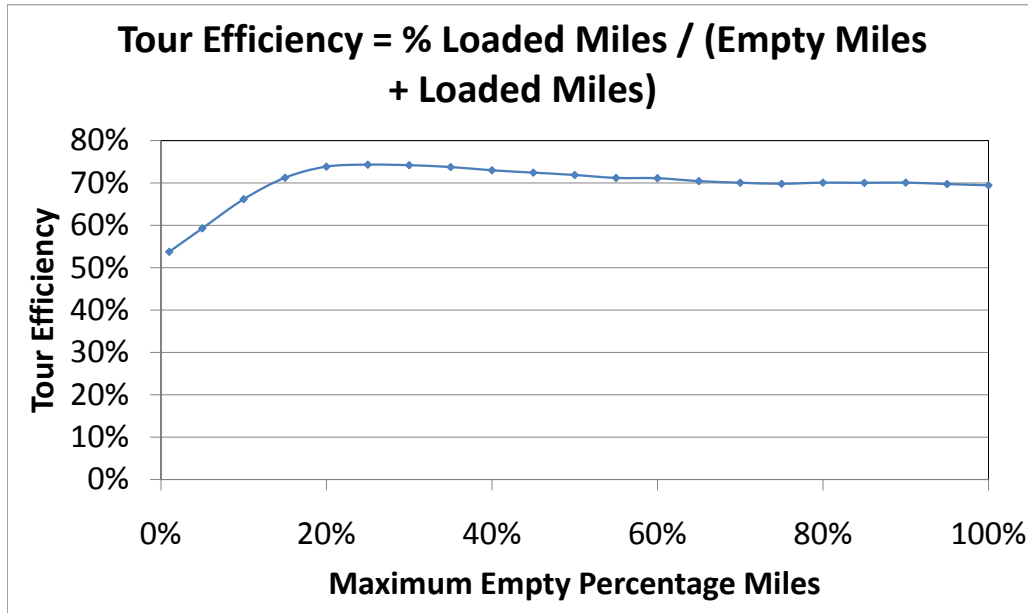


Figure 5-8: Behavior of Tour Efficiency in terms of Percentage Loaded Miles

- **Total Cost Optimality Gap** - is the percentage difference between the lowest possible solution cost and the actual solution cost. In this analysis, the lowest possible solution cost would be obtained when the fleet is not constrained at all in terms of the percentage empty miles traveled. This cost is considered as the baseline case for comparison with the cases where FNOT restrains the solution in terms of percentage empty miles that are allowed for a tour.

$$\text{Optimality Gap} = \frac{\text{Baseline Cost} - \text{Constrained Solution Cost}}{\text{Baseline Cost}}$$

Analysis of Metrics

As can be seen in the graph presented in Fig. 5-4, the number of potential tours generated by FNOT increases as we relax the constraint of allowable percentage empty miles in tours. However, as observed in the graph shown in Fig. 5-5, the marginal benefit generated by these extra potential tours provides increasingly reduced returns. This is because, as the percentage empty miles allowed in a tour increase, we are allowing FNOT to look at a larger sub-set of potential tours with the expectation

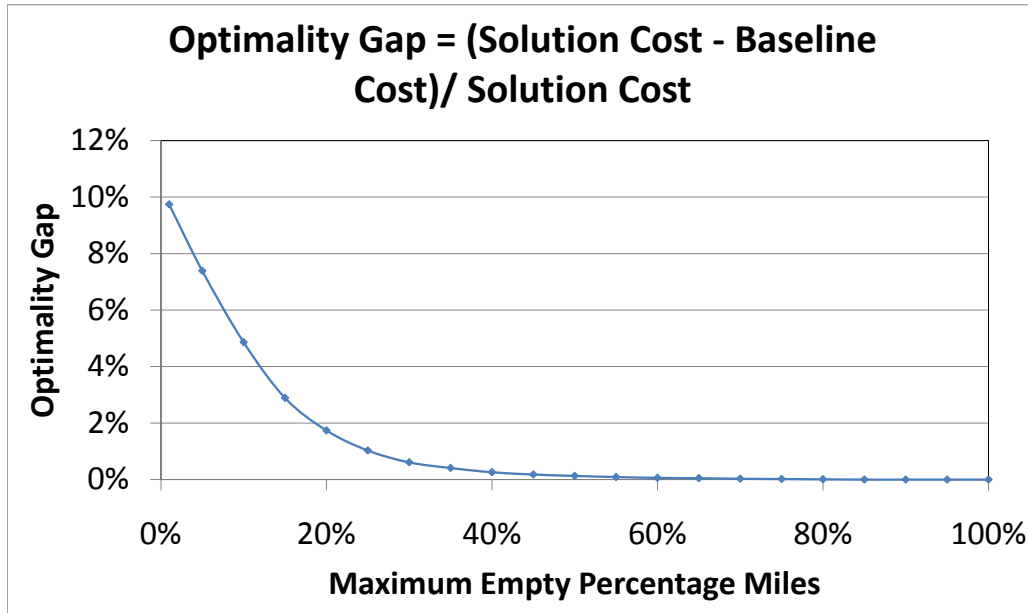


Figure 5-9: Behavior of Optimality Gap in terms of Total Solution Costs

that some of them might have a positive impact on the objective function in terms of decreasing the total network costs. However, these savings are obtained at the cost of increased computational requirements needed in searching these tours.

Hence, the purpose of this analysis is to find the trade-off spot, where the increasing computation time and resource requirement is not worth the possibility of finding a marginally profitable tour from an exponentially increasing pie of potential tours. This trade-off can be clearly seen in Fig. 5-6, which shows that the marginal improvement in Total Network Costs have decreasingly small improvements after a certain point of time. Upon cursory observation, it is clear that the Total Costs curve essentially flattens beyond the maximum percentage miles factor of 40%. This result is further substantiated by the Figs. 5-7 and 5-8, wherein we can see that the Tour Efficiency in terms of percentage loaded legs and percentage loaded miles is approximately at its peak around the empty percentage miles factor of 40%. Finally, we check for the sweet spot for the empty percentage miles factor by comparing it to the Optimality Gap. The gap between the best possible solution obtainable (by allowing upto a 100% empty percentage miles) and the Total Costs observed at 40% empty percentage miles was as low as 0.12%! It is clear from the analysis that the

trade-off point for an incremental improvement in the overall costs of the network by 0.12% needs computational efforts (in terms of extra potential tours to be generated) equating to almost 60% of the overall requirements!! This trade-off can also be seen in the graph comparing the Optimality Gap to the Potential Tours generated in the Fig. 5-10, which indicates that the number of tours required to be generated increase exponentially, as we approach the optimal solution.

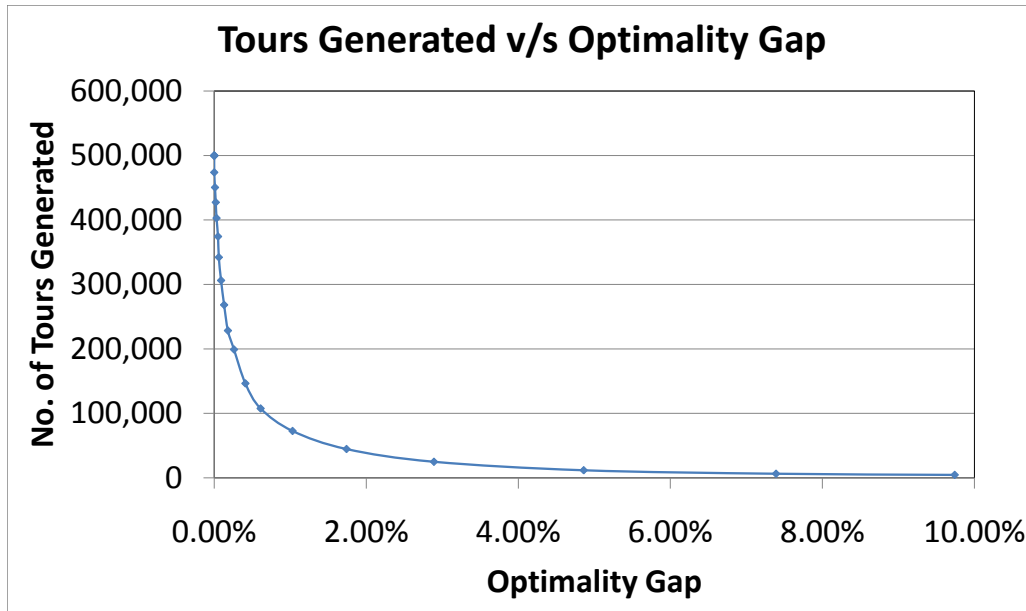


Figure 5-10: Behavior of Optimality Gap in comparison to the number of potential tours generated

Thus, it can be conclusively said that the sweet spot for the run parameter of empty percentage miles is approximately 40% (as a conservative estimate). Adding this constraint into the model provides us with a solution that is extremely close to optimal in terms of total costs, but utilizes only a fraction of resources and time. This analysis was repeated on three other sub-networks with identical results. Thus, we can say with relative certainty that the results obtained from this analysis can be used generically for similar retail networks, where the inbound (Vendor-DC) lanes have a low mean and high standard deviation; and the outbound (DC-Store) lanes are more consistent, with a high mean and low standard deviation.

However, we must qualify this statement by mentioning that the 40% empty per-

centage miles factor that was obtained for the network selected might not be an overall standard for all networks. This is because there might be certain esoteric network qualities like demand density, variability, geography, fleet and for-hire costs, etc. that might skew the results one way or the other. The main purpose of the analysis was to highlight the fact that using smart metrics to constrain an optimization model can be significantly beneficial in saving resources without compromising on the actual solutions obtained.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Planning Methodologies & Weekly Operations

In the previous chapters, we have gone through the need for using probabilistic distributions for modeling the uncertainty in demand patterns and discussed the capabilities of FNOT with regard to creating transportation plans. The following chapter builds upon these capabilities in order to actually create long term transportation plans and analyze their effectiveness at an operational level. The hierarchy of transportation planning, from Long Term Annual Planning methods to weekly operational scenario generation (based on operational flexibility) used in this thesis is illustrated in the Fig. 6-1. A detailed discussion of each node within the hierarchy follows.

6.1 Long Term Annual Planning

Annual Plans are long-term plans, which provide the company with a strategic direction to plan volume on transportation lanes using a mix of private fleet and for-hire carriers. Using the capabilities of FNOT, we can create these plans using stochastic distributions (probabilistic or empirical) or deterministic assignments of lane demands. A detailed discussion of these two types of plans is carried out in this section.

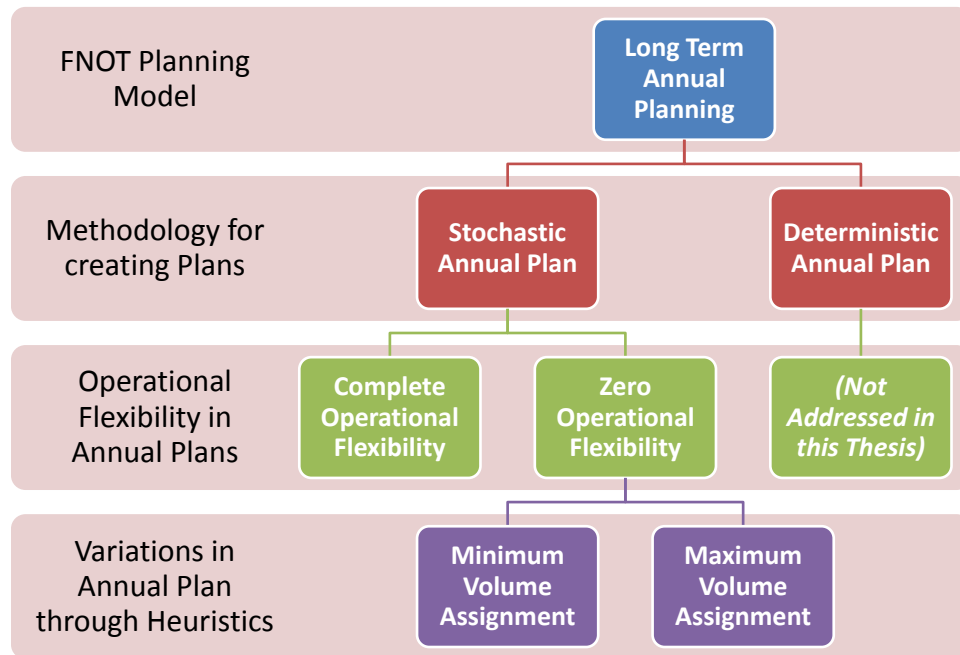


Figure 6-1: Hierarchy of transportation planning - from long-term planning, to development of weekly operational plans

6.1.1 Stochastic Plans

Stochastic Plans are Annual Plans that are created by modeling transportation lane demand using probability distributions or based on empirical lane demands. The key advantage of these plans is that they account for demand variability on lanes while planning for volume to be carried by the private fleet. Based upon the final assignments of lane demand by these plans, we can decide the fleet capacity needed to satisfy this demand and the tours that can best satisfy demand at minimal costs. At the same time, the plan also gives us an idea about which lanes need to be bid out to for-hire carriers.

6.1.2 Deterministic Plans

Deterministic Plans are created by modeling transportation lane demand using average lane demand, or by using heuristics like average demand plus a certain standard deviation. These plans assume that the calculated demand occurs every week, and do

not account for the possible variations in demand. Such planning methodologies are prevalent in the industry today. However, as per our discussion in previous chapters, the goal of this thesis is to move away from this paradigm and account for demand variability at an operational level. Hence, we shall not pursue these plans further in the thesis.

6.2 Sub-network used for analysis

In order to perform an in-depth analysis about the effectiveness of the annual plans generated by FNOT, we had to pare down the overall network into a small, but representative dataset that would be easy to work with. This would also help in analyzing the outputs provided by the model and comparing them across scenarios to generate insights.

The Table 6.1 indicates the key characteristics about the network that was selected for use in the remaining part of the thesis.

Utilizing the capabilities of FNOT, we are able to generate a stochastic annual plan for optimal fleet and for-hire assignments for this network. This plan also contains a set of tours that optimally route the volume carried by the private fleet over the network. Based on our conclusions from discussion about demand distributions and goodness of fit in Section 5.2, we know that empirical distributions provide the best results for planning. Hence, the annual plan for this network is created using empirical lane data.

6.3 Simulation of Weekly Demands

FNOT is a steady state model. This means that, given the inputs about lane volumes along with standard deviations (for probabilistic distributions) or lane-level demand histograms (for empirical distributions), the model calculates the best possible volume assignment over the network. However, this is not to say that there might not be any problems in executing the plans at an operational level, . Because of the inherent

Table 6.1: Characteristics of the Mini-Network

Network Characteristics	Statistics	
No. of domiciles (DC's)	2	
No. of Vendors	141	
No. of Stores	183	
No. of freight lanes	324	
No. of Inbound Lanes	141	
No. of Outbound Lanes	183	
Minimum Lane distance	50	miles
Maximum Lane distance	1028	miles
Average Lane distance	429	miles
Maximum Lane volume	103	loads
Minimum Lane volume	0	loads
Average Lane volume	4.83	loads
Private fleet minimum charge (per load)	\$ 60	
Private fleet rate (per mile)	\$ 1.20	
For-hire minimum charge (per load)	\$ 250	
For-hire rate (per mile)	\$ 0.66	
Fuel Price	\$ 1.95	
DOT Max. Hours of Work reqmt. for drivers per day	11	hours
DOT Min. Hours of Rest reqmt. for drivers per day	8	hours
Maximum allowable tour distance	1500	miles

variability in lane demand, there could be some weeks where lane demand might go over or under the planned volume, and contingencies have to be developed in the form of for-hire assignments. This would lead to a gap in planned volume that was assigned to fleet and the actual volume observed in practice.

In order to understand the behavior of the model at an operational level, it is necessary to simulate weekly demand conditions and analyze the performance of the Annual Plan on a week-to-week basis. The demand for lanes is generated using a Monte-Carlo simulation consisting of a random number matrix representing demand occurrence probabilities for all the lanes within the network, over the 52 simulation weeks. Next, we create a cumulative probability distribution table for the lane demands from the historical demand information available to us. Using the random numbers from the matrix, we can calculate the demand occurrence on a lane from the cumulative probability distribution of its demand. In order to understand this demand calculation, an example is shown below.

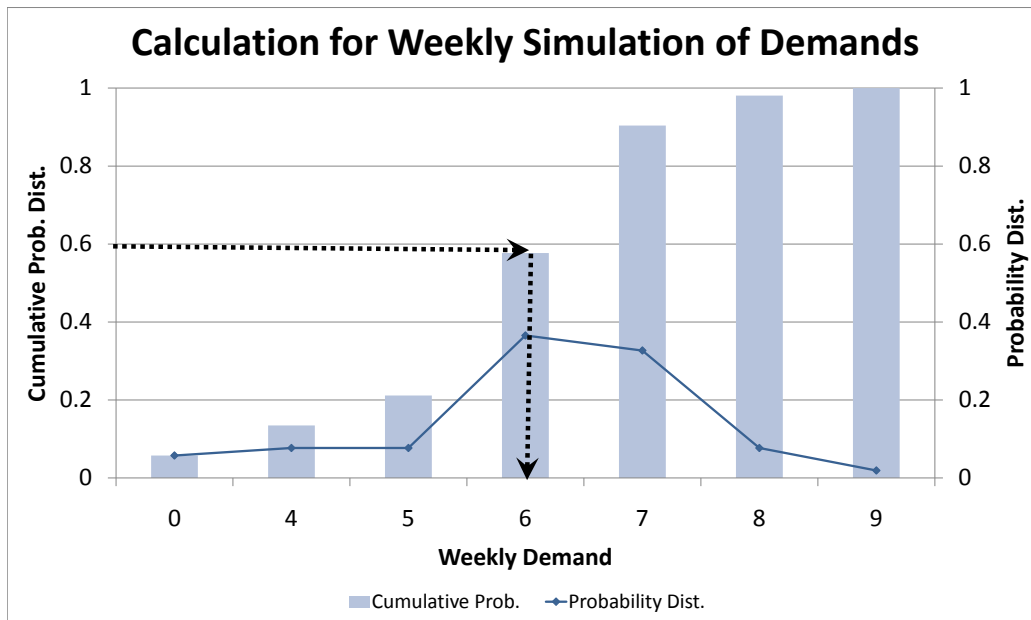


Figure 6-2: Example of weekly demand calculation from randomly generated probabilities

The graphs shown in 6-2 depict the probability distribution function (PDF) and cumulative probability function (CDF) of a sample lane within the network. Suppose

the Monte-Carlo simulation generated a random probability of 0.5 during a particular week for this lane. In order to calculate the corresponding demand for the lane during that week, we developed a code in MATLAB. The algorithm basically looks for the demand interval within which the randomly generated probability lies. It then assigns the higher value of the interval as observed demand for the lane for that particular week. Thus, for the example indicated, the demand assigned to the lane for the week under consideration would be 6 truckloads. Note that this calculation leads to a slight approximation error, since demand on a lane occurs in discrete intervals (0,1,2,3....) and cannot be a continuous variable (for e.g., demand cannot be 5.536), since we are considering full truckload movements within our network. However, since the probabilities generated are completely random in nature, this approximation gets evenly distributed over the length of the simulation.

Repeating the same process for all the random values in the matrix, we end up with a demand matrix for all the lanes over the 52 weeks. In order to confirm the validity of the simulated demand, we performed a check to confirm that the data is not skewed and is completely random. This can be done by comparing the average of lane demands generated over the 52 weeks with the actual average of the lane demand, as observed from the histogram. After completing this check, we can be sure that the demand generation is successful, and provides an honest picture what the company might face on a week-to-week basis.

Weekly Operations

The main question that we are trying to answer in this thesis is, "*How robust is the Annual Plan generated by FNOT?*" In other words, can the plan handle the difficulties created within the network because of demand variations on a weekly basis and still produce low cost results? In order to test this hypothesis, we need to understand the performance of the Annual Plan in the presence of random demands simulated in the previous chapter. The following section discusses various scenarios through which the annual plan can be used at an operational level to allocate weekly demand to the

private fleet and for-hire carriers.

6.4 Operational Flexibility Scenarios

In reality, it is quite difficult and complex to have complete operational freedom to decide the optimal allocation of lane demand to a mix of private fleet tours and for-hire carriers on a weekly basis. Hence, the robustness of an Annual Plan in being able to handle the weekly demand fluctuations with minimal operational flexibility is of paramount importance to any company. In a zero flexibility environment, only the tours produced by the annual plan are allowed to be operated in practice. This is the opposite extreme of having complete flexibility with regards to allocating loads based on weekly demand dynamics, which result in optimal assignments. Hence, the performance of an annual plan in terms of its robustness can truly be tested by comparing its performance in a zero flexibility environment with the complete flexibility scenario. If the annual plan allocation is robust, it should be capable of matching (or at least coming close to) the results obtained with complete flexibility.

The following sections provide a summary of how these operational scenarios with varying levels of flexibility are created and tested.

6.4.1 Complete Operational Flexibility

This scenario can be considered as every operational planner's dream! In this case, we use the Annual Plan solely for the purpose of deciding fleet and driver capacity; and the operational weekly plans are allowed to occur on an ad-hoc basis, after actually visualizing the demand. Thus, in this case, the operational planner knows exactly what the demand is going to be each week, and, keeping that in mind, they decide the optimal allocation of demand to private fleet tours or for-hire carriers. The operational planner also has the freedom of deciding what the optimal routes for tours served by the private fleet would be every week and these tours can change from week-to-week as deemed fit by the planner. Additionally, it can be assumed that they have long-term contracts set up with for-hire carriers for all the lanes within the

network, and if we had to assign a load to them, they would carry the load at the pre-arranged prices and cannot reject it for any reason.

This scenario can be easily simulated in FNOT. The weekly lane volumes obtained from Monte-Carlo simulations are assigned to the lanes under the assumption of a deterministic (fixed volume) distribution. The driver and fleet capacity is constrained using the numbers generated by the Annual Plan. Under this set of conditions, the FNOT model is re-run 52 times, using the 52 different weekly volumes generated by the simulation.

The resulting output obtained from these runs would be the optimal allocation of lane volume to a mix of fleet and for-hire carriers week-after-week. Also, the model would produce the best possible tours that the private fleet should operate upon, keeping the weekly lane demand variances in mind. Hence, this operational scenario simulates the best case environment against which the annual plan needs to match up, in order to truly prove its robustness.

6.4.2 Zero Operational Flexibility

This scenario is the opposite extreme of the full flexibility environment. In this case, we assume that the Annual Plan created at the beginning of the year needs to be followed exactly as it is, in terms of tour routing. Thus, the private fleet tours planned by the Annual Plan are the only ones that can be executed in practice. This means that once the Annual Plan is generated, the operational planners are restricted to assigning weekly loads occurring on lanes to these tours only, whenever the tours may be feasible. All the remaining loads go to for-hire carriers. Hence, this plan allows for zero flexibility in deviating away from the Annual Plan, as far as tour formulations are concerned. However, the tour volumes can be flexible, depending upon the actual volume occurrences on lanes, on a week-to-week basis.

At first glance, this kind of planning, by itself, seems to be too constricting in terms of handling operational demand variability effectively. Hence, a simple heuristic is used to further refine the allocation of volume to the tours generated by the Annual Plan. This heuristic gives due importance to the metrics of overall asset

utilization and minimization of empty percentage miles, and ranks the Annual Plan tours accordingly. It basically gives more preference to tours that have lower empty percentage miles, since they are generally assumed to be more cost effective and also tend to improve asset utilization by reducing the empty miles travelled. This heuristic is most effective where the fixed and variable costs per mile for the private fleet may be different for distribution centers across the network, i.e. one distribution center might value its fleet in a different way than another. In cases, where these costs are common across the network (as is the case for our scenario), we can also use the tour cost per loaded mile traveled, as a metric for ranking the tours.

Once the tours are ranked as per the aforementioned heuristic at a DC level, the next step would be to test the Annual Plan's performance under this scenario, at an operational level. That is, in the presence of weekly lane demands generated by the Monte-Carlo simulations, how does this stringent planning methodology perform? The annual plan can be operationalized under the zero flexibility scenario using two different variations of volume assignment to tours. In one case, the volume assigned to a tour for a particular week equals the minimum volume on any of the lanes within the tour during that week. This is called the Minimum Volume Assignment Plan. In the second case, the volume assigned to a tour for a particular week equals the maximum volume on any of the lanes within the tour during that week. This is called the Maximum Volume Assignment Plan. A detailed discussion of these plans, and the algorithms used to operationalize them are discussed in next sections.

Volume assignment calculations

As discussed before, the zero flexibility planning methodology constrains the operational planners to using the private fleet only for tours generated by the Annual Plan. However, the planners still have to make the decision about the number of times these tours shall occur during the week. This is dependent upon the volumes occurring on the lanes comprising the tours on a weekly basis.

There are two schools of thought with regards to assigning lane volume to the tours generated by the Annual Plan, which can be illustrated by the example shown

in Fig. 6-3:

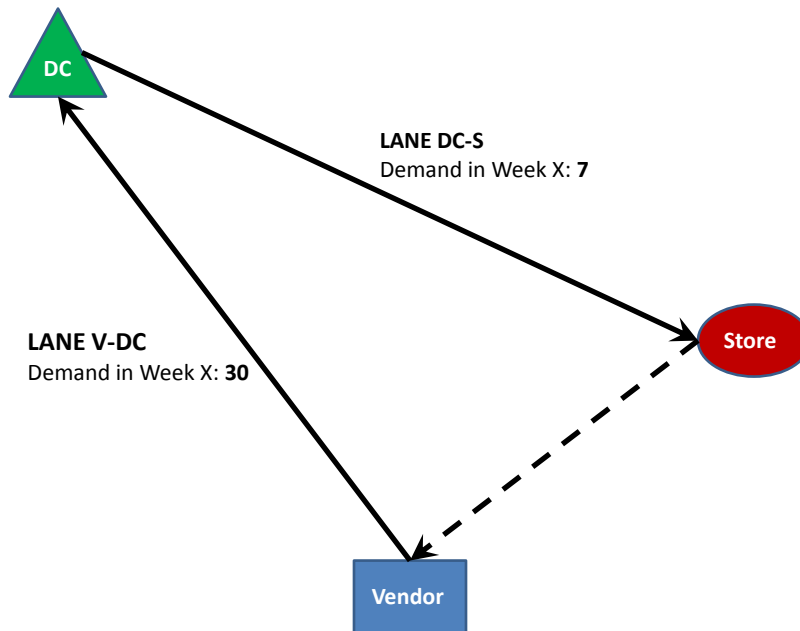


Figure 6-3: An example of demand occurrence during a random week over a tour

One school of thought dictates that we must assign the minimum volume occurred over all the lanes within a tour to the fleet, while the excess volume is assigned to For-Hire Carriers. Thus, for the example shown in Fig.6-3, the volume assigned to the tour would be equal to 7, which is the minimum lane volume over the tour occurring on Lane DC-S; the remaining 23 loads on the V-DC lane is assigned to for-hire carriers.

However, at the other end of the spectrum, we might have risk-averse operational managers who prefer moving all volume that occurs on lanes within the annual plan using the private fleet. This behavior might be based upon company rules which dictate that whenever volume occurs on any lane that is a part of the annual plan, that volume must be satisfied by the private fleet. Although this might seem a bit too conservative, service or other qualitative considerations might necessitate such a move. Also, concerns related to fleet relocation back to its domicile for future load requirements might be another constraint that may drive operational managers to make such a decision. In this case, the volume assigned to the tour shown in Fig. 6-3

would be equal to 30.

Thus, based upon the assignment policies discussed above, we have created algorithms for allocating the randomly generated weekly volume over our network to annual plan tours. Analyzing the performance of these execution schemes would provide insights into the robustness of using the annual plan for weekly operations, in the presence of demand variability.

Minimum Volume Assignment Case

In this case, the number of times a tour is operated during a week is equal to the minimum volume during the week over all the lanes comprising that tour. This assignment makes intuitive sense, since a tour would not realistically materialize if any of the lanes within that tour has no volume to be carried for a particular week.

Thus, in order to assign lane volumes generated by the Monte-Carlo simulation to tours on a weekly basis, we developed a code in MATLAB, which basically follows the algorithm listed below, for creating the operational plan:

1. Create a table consisting of all the tours generated by the Annual Plan, along with the lanes comprising the tour - we can call this the 'tour table'.
2. Sort the tour table at the DC level in the order of decreasing preference for volume allocation, based upon the empty percentage miles heuristic discussed previously
3. Create another table indicating the randomly generated lane volumes for the week under consideration, with an additional field called 'lane volume assigned to tours'. This field is initialized with a value of zero, for all lanes within the network. We can call this table as the 'lane table'.
4. Starting from the top of the tour table, calculate the difference between lane volume and lane volume assigned to tours for the week under consideration, for all lanes comprising the tour

5. The minimum value of the difference calculated in the previous step is assigned as the ‘tour volume’
6. Update the ‘lane volume assigned to tours’ field in the lane table for all the lanes within this tour with the ‘tour volume’ value calculated above
7. Move to the next tour within the tour table, and repeat steps 4 to 6
8. Calculate the difference between the lane volume and lane volume assigned to tours fields in the lane table. This would be the volume assigned to for-hire carriers for the week under consideration
9. Repeat steps 3 to 7, for the randomly generated demands for all 52 weeks

Thus, using the algorithm above, we can generate an operational plan for the entire year with the Minimum Volume Assignment case. An important characteristic about the Annual Plan is that it allows lane volume to be assigned to multiple tours, as deemed most profitable and/or feasible. This adds additional complexity into the minimum volume assignment algorithm, since we now need to maintain an account of lane volume that has already been assigned to tours, and remaining lane volume yet to be assigned - in order to avoid double counting. This is done using the lane volume assigned to tours field.

The for-hire assignment field indicates the volume for lanes which was not assigned to any tour, and is hence, given out to for-hire carriers. Additionally there may be certain lanes within the network that are not present in the Annual Plan. In such cases, all the volume occurring on Non-Annual Plan lanes is always assigned to for-hire carriers.

Thus, using the methodology described above, we can generate a comprehensive operational plan for the Minimum Volume Assignment case, with respect to tour volumes and lane volume assignments to private fleet and for-hire carriers. A potential disadvantage of this planning mechanism is that we would still need to maintain contracts with for-hire carriers for all the lanes within our network in order to satisfy demand for lanes within the Annual Plan tours that cannot be met using the private

fleet. This leads us to our second case, which considers the opposite extremity for operational planning in a zero flexibility environment.

Maximum Volume Assignment Case

The main assumption in creating the operational plan for this case is that the entire volume occurring for all the lanes within the Annual Plan must be satisfied using the private fleet. Hence, in this case, we would need to maintain contracts with for-hire carriers only for the lanes which are not present in the Annual Plan. Intuitively speaking, such a planning model would make sense if the company's service level requirements are so high, that they cannot rely on for-hire carriers. Another possible explanation of such an assignment could be derived from the Aircraft Planning model, wherein the aircrafts follow a pre-decided planning schedule precisely, irrespective of the number of passengers onboard - since asset relocation back to the domicile holds priority over network inefficiencies.

In order to assign lane volumes generated by the Monte-Carlo simulation to tours on a weekly basis, we again developed a code in MATLAB, which basically follows the following algorithm for creating the operational plan:

1. Create a table consisting of all the tours generated by the Annual Plan, along with the lanes comprising the tour - we can call this the 'tour table'.
2. Sort the tour table at the DC level in the order of decreasing preference for volume allocation, based upon the empty percentage miles heuristic discussed previously
3. Create another table indicating the randomly generated lane volumes for the week under consideration, with an additional field called 'lane volume assigned to tours'. This field is initialized with a value of zero, for all lanes within the network. We can call this table as the 'lane table'.
4. Starting from the top of the tour table, calculate the difference between lane volume and lane volume assigned to tours for the week under consideration, for

all lanes comprising the tour

5. Assign the maximum of difference calculated above as the tour volume. We must constrain this value to be above zero. If not, the tour volume is indicated as zero.
6. Update the 'lane volume assigned to tours' field in the lane table for all the lanes comprising this tour with the 'tour volume' value calculated above
7. Move to the next tour within the tour table, and repeat steps 4 to 6
8. Repeat steps 3 to 7, for the randomly generated demands for all 52 weeks

The Maximum Volume Assignment algorithm used above will allocate any volume occurring on the lanes which are included in the Annual Plan to the private fleet. On the other hand, volume on lanes that are not a part of the Annual Plan can automatically be assigned to for-hire carriers. A downside of this planning mechanism is that the empty miles travelled by the fleet would be very high since even if volume occurs on a single lane within a tour, it must be satisfied by the private fleet. Also, the overall fleet size required for this executing plan would be quite high.

Thus, using the algorithms for Minimum and Maximum Volume Assignment, we can create operational plans which have zero flexibility with regard to the Annual plan tour routes. These plans do not have the option of creating alternate tours for private fleet assignment. In the next chapter, we shall dig deeper into the weekly operations of these plans, and compare them to the full flexibility scenario, to test our hypothesis about annual plan robustness.

Chapter 7

Simulation Results Analysis

After obtaining the results for the weekly execution of the Annual Plan using the methodologies discussed in the previous chapter, we assimilate the resulting facts together, in order to compare them with each other. The performance of these operational plans can essentially be compared at four main levels, in decreasing order of aggregation, as follows:

- Network Level
- Facility (Distribution Center) Level
- Tour Level
- Lane Level

Comparing the plans at varying levels of aggregation shall help us analyze the similarities and differences in weekly operations, and dig deeper into the reasoning behind it.

7.1 Network Level Metrics

At the Network level, we intend to compare the overall efficacy of the operational plans with each other in order to draw generalizations on their performance with certain key metrics discussed below. We believe that these metrics are indicators of

the Annual Plan behavior at a week-to-week execution level in comparison to the best-case, scenario with complete flexibility. It will also help us understand the associated advantages and disadvantages of using one execution strategy over the other.

- **Total Cost** - is the key objective function that is minimized by the FNOT model, as discussed in the section 4.1. These costs indicate the total amount spent on the private fleet and for-hire carriers.
- **Total Distance** - is the total number of miles travelled by the private fleet and for hire carriers. This number is an indicator of how efficiently the volume assignments over the lanes were made, but may mask the fact that a longer tour might have been cheaper than assigning the corresponding loads to for-hire.
- **Total Volume** - is an indicator of the total volume carried by the private fleet and for-hire carriers over the network. Since the volumes simulated by the Monte-Carlo simulations were frozen to get a good cross comparison of the behavior of the Annual plan in the different cases, these numbers would essentially be the same.
- **Fleet Cost** - is the portion of total costs incurred by the private fleet in dispatching volume over the network in the form of tours. These costs, when used in conjunction with the percentage volume transported over the fleet are a good indicator of fleet usage efficiency.
- **For Hire Cost** - is the portion of total costs incurred due to volume being carried by for-hire carriers over the network. Most companies would prefer this cost to be as low as possible, because for-hire carriage is generally conceived to be more expensive than fleet and is considered as an emergency shipment in the presence of demand variability, which results in lower perceived service quality.
- **Fleet Volume** - is the portion of total volume carried over the network by the private fleet, in the form of tours.

- **For Hire Volume** - is the portion of total volume carried over the network by for-hire carriers. This number can be used along with the For-Hire Cost as an indicator of how expensive emergency shipments have been, during the weekly execution of the Annual Plan.
- **Fleet Distance** - is the portion of total distance travelled by the private fleet over the network.
- **For Hire Distance** - is the portion of total distance travelled by for-hire carriers over the network.
- **Average Cost per Load (ACPL)** - As the name suggests, is the average cost per load incurred in carrying the total volume across the network using fleet and for-hire carriers. This number is an indicator of the efficiency of the operational plan, keeping lane costs in mind.
- **Average Cost per Mile (ACPM)** - is the average cost incurred per total miles travelled over the network by the private fleet and for-hire carrier trucks. This number is also an indicator of the cost efficiency of the operational plan, as a function of the distance travelled across the network.
- **Average Length of Haul (ALOH)** - is the average distance travelled by loads carried over the private fleet and for-hire trucks across the network. Along with ACPM, ALOH is also an indicator of the efficiency with which loads are transported over the network, as a function of distance travelled by the loads.
- **Fleet ACPL** - is the average cost per load incurred in carrying the total volume across the network using private fleet only. It is an indicator of the cost efficiency of loads carried in tours across the network.
- **Fleet ACPM** - is the average cost incurred per mile travelled over the network by the private fleet only.
- **Fleet ALOH** - is the average distance travelled by loads carried using the private fleet only. When used in conjunction with the other metrics for private

fleet, it is an indicator of the operational efficiency of the tours generated.

- **For Hire ACPL** - is the average cost per load for the volume carried solely with for-hire carriers.
- **For Hire ACPM** - is the average cost per mile for the volume carried by for-hire carriers.
- **For Hire ALOH** - is the average distance travelled by volume carried by for-hire carriers over the entire network.

The following section goes through the results obtained from the simulation runs and discusses the significance of the values obtained for each of the metrics explained previously. The numbers shown in Table 7.1 are average values of the metrics obtained over the 52-week simulation runs. These numbers include only the weekly operations that were simulated, and exclude the annual plan numbers.

Table 7.1: Summary of Network-level statistics for key metrics

Metrics		Mean	Std. Error	Min.	Median	Max.	Std. Dev.	Range
Total Cost	Full Flex.	\$268,176	\$ 1,798	\$245,383	\$265,191	\$308,505	\$ 12,964	\$ 63,123
	Min. Vol.	\$283,754	\$ 1,669	\$262,608	\$282,006	\$325,295	\$ 12,034	\$ 62,687
	Max. Vol.	\$321,639	\$ 1,855	\$295,078	\$319,464	\$364,894	\$ 13,378	\$ 69,816
Total Distance	Full Flex.	223,721	1,570	194,789	223,322	262,127	11,321	67,338
	Min. Vol.	215,645	1,628	198,126	214,093	251,760	11,737	53,634
	Max. Vol.	270,150	1,751	251,714	267,408	312,060	12,627	60,346
Total Volume	Full Flex.	1,597	9	1,475	1,590	1,785	62	310
	Min. Vol.	1,597	9	1,475	1,591	1,785	62	310
	Max. Vol.	1,597	9	1,475	1,591	1,785	62	310
Fleet Cost	Full Flex.	\$159,460	\$ 2,856	\$123,908	\$170,200	\$181,829	\$ 20,597	\$ 57,920
	Min. Vol.	\$132,732	\$ 805	\$122,606	\$132,623	\$146,044	\$ 5,808	\$ 23,438
	Max. Vol.	\$215,482	\$ 1,078	\$198,057	\$216,086	\$229,598	\$ 7,773	\$ 31,541
For Hire Cost	Full Flex.	\$108,716	\$ 4,004	\$ 78,068	\$ 92,284	\$184,597	\$ 28,876	\$106,529
	Min. Vol.	\$151,022	\$ 1,692	\$124,215	\$149,844	\$192,889	\$ 12,202	\$ 68,674
	Max. Vol.	\$106,157	\$ 1,451	\$ 87,803	\$104,559	\$141,270	\$ 10,462	\$ 53,467
Fleet Volume	Full Flex.	1,187	13	1,026	1,227	1,292	91	267
	Min. Vol.	1,018	9	897	1,009	1,141	62	244
	Max. Vol.	1,197	6	1,122	1,194	1,299	45	177
For Hire Volume	Full Flex.	410	16	291	348	717	115	427
	Min. Vol.	579	7	475	574	742	48	267
	Max. Vol.	400	6	327	395	544	41	217
Fleet Distance	Full Flex.	132,883	2,380	103,257	141,833	151,524	17,164	48,267
	Min. Vol.	110,610	671	102,172	110,519	121,703	4,840	19,532
	Max. Vol.	179,568	898	165,047	180,072	191,331	6,477	26,284
For Hire Distance	Full Flex.	90,838	2,325	68,476	84,556	143,210	16,765	74,734
	Min. Vol.	105,035	1,478	84,759	103,586	138,884	10,655	54,125
	Max. Vol.	90,582	1,463	71,143	88,981	125,707	10,547	54,564
ACPL	Full Flex.	\$ 168	\$ 0.56	\$ 160	\$ 168	\$ 178	\$ 4	\$ 19
	Min. Vol.	\$ 178	\$ 0.71	\$ 167	\$ 178	\$ 194	\$ 5	\$ 27
	Max. Vol.	\$ 201	\$ 0.40	\$ 196	\$ 201	\$ 208	\$ 3	\$ 13
ACPM	Full Flex.	\$ 1.20	\$ 0.01	\$ 1.13	\$ 1.18	\$ 1.36	\$ 0.06	\$ 0.23
	Min. Vol.	\$ 1.32	\$ 0.00	\$ 1.22	\$ 1.32	\$ 1.38	\$ 0.04	\$ 0.16
	Max. Vol.	\$ 1.19	\$ 0.00	\$ 1.15	\$ 1.19	\$ 1.24	\$ 0.02	\$ 0.09
ALOH	Full Flex.	140	1	122	141	156	6	34
	Min. Vol.	135	1	124	135	146	4	22
	Max. Vol.	169	0	160	169	179	3	19
Fleet ACPL	Full Flex.	\$ 134	\$ 1.24	\$ 116	\$ 136	\$ 150	\$ 9	\$ 34
	Min. Vol.	\$ 131	\$ 0.51	\$ 124	\$ 130	\$ 141	\$ 4	\$ 17
	Max. Vol.	\$ 180	\$ 0.46	\$ 174	\$ 180	\$ 188	\$ 3	\$ 14
Fleet ACPM	Full Flex.	\$ 1.20	\$ 0.00	\$ 1.20	\$ 1.20	\$ 1.20	\$ 0.00	\$ 0.00
	Min. Vol.	\$ 1.20	\$ 0.00	\$ 1.20	\$ 1.20	\$ 1.20	\$ 0.00	\$ 0.00
	Max. Vol.	\$ 1.20	\$ 0.00	\$ 1.20	\$ 1.20	\$ 1.20	\$ 0.00	\$ 0.00
Fleet ALOH	Full Flex.	112	1	97	114	125	7	29
	Min. Vol.	109	0	103	108	118	3	14
	Max. Vol.	150	0	145	150	157	3	11
For Hire ACPL	Full Flex.	\$ 266	\$ 0.72	\$ 255	\$ 266	\$ 279	\$ 5	\$ 24
	Min. Vol.	\$ 261	\$ 0.37	\$ 254	\$ 261	\$ 268	\$ 3	\$ 14
	Max. Vol.	\$ 266	\$ 0.54	\$ 256	\$ 266	\$ 277	\$ 4	\$ 21
For Hire ACPM	Full Flex.	\$ 1.18	\$ 0.02	\$ 1.00	\$ 1.14	\$ 1.55	\$ 0.13	\$ 0.55
	Min. Vol.	\$ 1.44	\$ 0.01	\$ 1.24	\$ 1.44	\$ 1.61	\$ 0.08	\$ 0.37
	Max. Vol.	\$ 1.18	\$ 0.01	\$ 1.05	\$ 1.17	\$ 1.36	\$ 0.06	\$ 0.31
For Hire ALOH	Full Flex.	227	4	164	235	269	25	105
	Min. Vol.	181	1	158	182	210	11	52
	Max. Vol.	227	2	190	227	252	12	62

The following section analyzes the network level results observed in Table 7.1 for the key metrics defined previously.

- **Total Cost** - As expected, the cost values for the full flexibility stochastic plan would be the least. This is because we are providing the model with the most amount of freedom to route the volume occurring over the network optimally. However, an interesting observation is that the total costs for Minimum Volume Assignment case are not much farther off as compared to the optimal scenario. In comparison to the full-flexibility option, the total costs for Minimum Volume Assignment case are only 6% higher! Also, the standard deviation for the total cost values of Minimum Volume Assignment case is lower, indicating a tighter bound on the deviation from average values. As expected, the Maximum Volume Assignment case total costs are much higher, since the plan is quite wasteful in terms of a high number of empty miles travelled by the private fleet for covering the required volume.
- **Total Distance** - As observed, the total distance for Minimum Volume Assignment case is the lowest. This makes sense because in comparison to the full flexibility option, there is more volume being assigned to For-Hire carriers in the Minimum Volume Assignment case. Thus, the overall miles travelled would reduce, since for-hire load distances are considered from pick-up to drop-off location only. In the case of private fleet, we need to relocate the trucks back to the origin DC, which results in more miles being covered.
- **Total Volume** - Since we had frozen the simulated volumes across the different operational scenarios in order to obtain a direct and true comparison of the weekly execution, the volume numbers across all the three cases are identical.
- **Fleet Cost** - As expected, the fleet costs for the Maximum Volume Assignment case are the highest since the nature of operational execution is such that it forces the fleet to cover a certain requisite amount of volume over the lanes. The Full Flexibility option has higher fleet costs as compared to the Minimum

Volume Assignment option. This is because when the model is given the flexibility of choosing the optimal weekly allocation of volume over the network, it tries its best to assign lane volume to the private fleet whenever it is economically feasible to do so, thus increasing overall fleet costs. Since the Minimum Volume Assignment case is restricted to using only the tours generated in the Annual Plan, it has limited options in assigning volume to tours, resulting in reduced fleet costs.

- **For-Hire cost**- The For-hire costs for the Minimum Volume Assignment case are the highest, because the constraints associated with allocating volume to private fleet are the greatest in this case. As a result, all the left-over demand needs to be met by for-hire carriage. An interesting thing to note is that the difference between for-hire costs for the Maximum Volume Assignment case and the Full Flexibility option is not much. This is because lanes that are not present in the Annual Plan need to be covered by For-Hire carriers, regardless. Hence, there isn't much savings resulting from conservatively assigning all the volume occurring on Annual Plan lanes exclusively to the private fleet.
- **Fleet Volume** - The behavior of fleet volume allocations in the three different operational plans follows the same behavior as fleet costs. Again, it can be seen that the Maximum Volume Assignment case does not have a substantially higher volume assignment in comparison to the Full Flexibility plan. Thus, it can be seen that the Maximum Volume Assignment plan reaches close to optimality conditions as far as fleet volume carriage is concerned, but at the expense of a lot of extra empty miles traveled in the process.
- **For-Hire Volume** - The behavior of for-hire volume assignment is exactly opposite to the discussion for fleet volume, as expected.
- **Fleet distance** - The fleet distance travelled in the Maximum Volume Assignment case is the highest, since the operational plan constrains all the volume occurring over the lanes included in the Annual Plan to be satisfied by fleet.

Hence, the private fleet accumulates a lot of empty miles in satisfying this constraint. On the other hand, the Minimum Volume Assignment case has the least miles travelled by the fleet, since it is conservative in terms of volume assignment to fleet, and covers more volume through for-hire carriers. However, an important thing to note here is that because of demand variability, the spread of private fleet usage is much larger in comparison to the other cases. This might add to the costs observed, if we penalize the Full Flexibility option for idle fleet capacity.

- **For-hire distance** - The for-hire distance travelled in the Minimum Volume Assignment case is the highest because of the relatively higher volume assigned to for-hire carriers in this plan. However, the interesting thing to note again is that the for-hire distance traveled in the Maximum Assignment case is really close to the Full Flexibility plan. This reiterates the fact that there is not much value-add in religiously covering all the volume occurring in the Annual Plan lanes by private fleet. The for-hire costs incurred in the Maximum Volume Assignment Plan are independent of the Annual Plan lanes, and due consideration should be given to this fact before routing volume over the network. Again, the spread for for-hire carrier usage is much larger in the Full Flexibility option, which highlight the operational difficulties for executing the plan in reality, because of demand variability.
- **Average Cost per Load** - The ACPL for the Full Flexibility option is the least, as expected, since the model has complete freedom to route weekly volume occurrences in the best possible manner using a combination of private fleet and for-hire carriers. On the other hand, the Maximum Volume Assignment case turns out to be most expensive because of the conservative nature of the routing.
- **Average Cost per Mile** - The ACPM for the Minimum Volume Assignment case is the highest because of the greater proportion of volume that is assigned to for-hire carriers, who generally have higher minimum costs for carrying a load. In comparison, the Maximum Volume Assignment option has the least

ACPM since it moves a larger proportion of volume using the private fleet.

- **Average Length of Haul** - As expected, the ALOH for the Maximum Volume Assignment Plan is the highest since all Annual Plan lane volume is covered by private fleet, leading to extra empty miles traveled. On the other hand, the Minimum Volume Assignment case has the least ALOH since the for-hire assignments are greater in this case. The Full Flexibility plan optimizes the lane volume carriage as per costs, and hence it has a slightly higher ALOH, because of economic viability considerations.

The Fleet and For-Hire ACPL, ACPM and ALOH also display the same characteristics as discussed previously, and help reinforce our understanding of how the operational plans work.

The main take-aways from this analysis is the closeness of the optimal Full Flexibility plan in weekly execution in comparison to the Minimum Volume Assignment case. Although the latter is slightly more expensive, the ease of operation in the form or reduced complexity of execution at a weekly level makes this a viable alternative to pursue. The extra costs incurred can be considered as a trade-off between ease of operations and achieving optimal routing in the presence of uncertainty.

7.2 Facility Level Metrics

At the Facility level, we intend to compare the overall efficacy of the operational plans with each other in order to draw generalizations on their performance with certain key metrics discussed below. By disaggregating the network by one level, we hope to understand whether the behavior of metrics is a function of the operational plan itself, or it is induced because of esoteric, external factors that are not being captured by the operational plan.

- **Number of Tours generated** - is the breakup of tours generated within each network that originated from a particular DC. This number, when used in conjunction with other metrics can give an idea about how much volume is

carried by fleet, and provides an indirect signal of volume consistency originating from the DC.

- **Total Volume** - is the breakup of volume handled by each DC within the overall network. This number indicates how busy a DC is, in comparison to others
- **No. of Drivers utilized** - is an indicator of how efficiently the private fleet has been utilized by the DC's in comparison to others
- **Average Number of Loaded Legs** - is an indicator of the number of loaded legs present in the tours originating from each DC. When used with the number of tours generated, this metric gives an idea about the percentage volume handled by the DC, which moves over the private fleet, versus for-hire carriers.
- **Tour Distance** - is the average distance travelled by tours originating from a particular DC. It provides an idea about the network size that the DC is serving, and is an indicator of network density.
- **Total loaded and empty distance travelled** - gives the breakup of empty versus loaded miles travelled by the private fleet originating from each DC, and is an indicator of operational efficiency.
- **Percentage of empty miles** - combines the above metric to come up with a single comparable number across DC's.
- **Total Cost** - is the breakup of total costs contributed by each DC towards the network.
- **ALOH** - is the average length of haul traversed by tours originating at each DC within the network.

The results for facility level metrics are shown in Tables 7.2 and 7.3. The results are quite self-explanatory and mirror the network level metrics in terms of behavior. Again, we can see that the Minimum Volume assignment case seems quite appealing

Table 7.2: Facility level metrics for DC6006

Metrics for DC 6006		Mean	Std. Error	Min.	Median	Max.	Std. Dev.	Range
No. of Tours	Full Flex.	89	1	72	92	102	8	30
	Min. Vol.	99	-	99	99	99	-	-
	Max. Vol.	99	-	99	99	99	-	-
Total Volume	Full Flex.	771	7	688	792	834	47	146
	Min. Vol.	674	6.58	570	670	779	47	209
	Max. Vol.	1,293	7.69	1,154	1,288	1,430	55	276
Avg. Loaded Legs	Full Flex.	1.65	0.01	1.54	1.65	1.78	0.05	0.24
	Min. Vol.	1.76	0.00	1.76	1.76	1.76	0.00	-
	Max. Vol.	1.76	0.00	1.76	1.76	1.76	0.00	-
Max. Loaded Legs	Full Flex.	3.00	-	3.00	3.00	3.00	-	-
	Min. Vol.	3.00	-	3.00	3.00	3.00	-	-
	Max. Vol.	3.00	-	3.00	3.00	3.00	-	-
Min. Tour Distance	Full Flex.	100	-	-	100	100	-	-
	Min. Vol.	100	-	100	100	100	-	-
	Max. Vol.	100	-	100	100	100	-	-
Max. Tour Distance	Full Flex.	499	7	363	517	558	52	195
	Min. Vol.	461	0	461	461	461	0	-
	Max. Vol.	461	0	461	461	461	0	-
Min. Loaded %age miles	Full Flex.	100%	0%	100%	100%	100%	0%	0%
	Min. Vol.	50%	0%	50%	50%	50%	0%	0%
	Max. Vol.	50%	0%	50%	50%	50%	0%	0%
Max. Loaded %age miles	Full Flex.	97%	0%	94%	96%	99%	2%	5%
	Min. Vol.	99%	0%	99%	99%	99%	0%	0%
	Max. Vol.	99%	0%	99%	99%	99%	0%	0%
Total Distance	Full Flex.	84,476	1,235	68,958	88,516	94,244	8,908	25,286
	Min. Vol.	70,975	444.02	64,582	70,126	78,628	3,202	14,046
	Max. Vol.	125,803	616.08	113,434	125,639	137,172	4,443	23,737
Total Empty Distance	Full Flex.	33,098	587	25,177	34,473	40,031	4,235	14,854
	Min. Vol.	27,610	141.09	25,850	27,606	29,588	1,017	3,738
	Max. Vol.	75,837	707.96	71,201	77,809	80,344	5,098	9,142
Total Costs	Full Flex.	\$101,372	\$ 1,482	\$ 82,750	\$106,219	\$113,093	\$ 10,689	\$ 30,343
	Min. Vol.	\$ 85,170	\$ 533	\$ 77,498	\$ 84,152	\$ 94,354	\$ 3,842	\$ 16,856
	Max. Vol.	\$150,964	\$ 739	\$136,121	\$150,767	\$164,606	\$ 5,331	\$ 28,485
Driver Hours per week	Full Flex.	4,206	53	3,567	4,427	4,531	380	963
	Min. Vol.	3,544	28.13	3,095	3,516	4,036	203	940
	Max. Vol.	6,966	37.36	6,254	6,955	7,653	269	1,399
ALOH	Full Flex.	109	1	97	110	120	7	23
	Min. Vol.	106	0.51	98	106	116	4	18
	Max. Vol.	97	0.21	94	97	101	2	7
%age Loaded Miles	Full Flex.	61%	0%	58%	61%	64%	2%	6%
	Min. Vol.	61%	0%	58%	61%	64%	2%	6%
	Max. Vol.	72%	0%	71%	72%	73%	1%	3%
Drivers used per week	Full Flex.	31	0	27	33	34	3	7
	Min. Vol.	26	0	23	26	30	2	7
	Max. Vol.	52	0	47	52	57	2	10
Available Drivers	Full Flex.	38	-	38	38	38	-	-
	Min. Vol.	38	-	38	38	38	-	-
	Max. Vol.	38	-	38	38	38	-	-

Table 7.3: Facility level metrics for DC6062

Metrics for DC 6062		Mean	Std. Error	Min.	Median	Max.	Std. Dev.	Range
No. of Tours	Full Flex.	68	1	50	72	80	8	30
	Min. Vol.	70	-	70	70	70	-	-
	Max. Vol.	70	-	70	70	70	-	-
Total Volume	Full Flex.	415	6	336	428	466	45	129
	Min. Vol.	344	4.25	289	348	401	31	112
	Max. Vol.	517	5.28	429	517	592	38	163
Avg. Loaded Legs	Full Flex.	1.47	0.01	1.35	1.47	1.63	0.08	0.28
	Min. Vol.	1.50	-	1.50	1.50	1.50	-	-
	Max. Vol.	1.50	-	1.50	1.50	1.50	-	-
Max. Loaded Legs	Full Flex.	3.00	-	3.00	3.00	3.00	-	-
	Min. Vol.	3.00	-	3.00	3.00	3.00	-	-
	Max. Vol.	3.00	-	3.00	3.00	3.00	-	-
Min. Tour Distance	Full Flex.	100	-	-	100	100	-	100
	Min. Vol.	100	-	100	100	100	-	-
	Max. Vol.	100	-	100	100	100	-	-
Max. Tour Distance	Full Flex.	449	4	374	455	534	28	160
	Min. Vol.	459	0	459	459	459	0	-
	Max. Vol.	459	0	459	459	459	0	-
Min. Loaded %age miles	Full Flex.	100%	0%	100%	100%	100%	0%	0%
	Min. Vol.	50%	0%	50%	50%	50%	0%	0%
	Max. Vol.	50%	0%	50%	50%	50%	0%	0%
Max. Loaded %age miles	Full Flex.	97%	0%	97%	97%	97%	0%	0%
	Min. Vol.	97%	0%	97%	97%	97%	0%	0%
	Max. Vol.	97%	0%	97%	97%	97%	0%	0%
Total Distance	Full Flex.	48,407	1,171	34,176	53,378	58,379	8,441	24,202
	Min. Vol.	39,635	406.74	34,071	39,540	46,689	2,933	12,618
	Max. Vol.	53,765	515.71	46,851	53,510	60,415	3,719	13,563
Total Empty Distance	Full Flex.	17,541	523	10,404	19,134	23,070	3,769	12,666
	Min. Vol.	14,952	138.96	13,263	14,794	18,036	1,002	4,773
	Max. Vol.	20,694	573.14	16,893	19,682	22,038	5,133	5,146
Total Costs	Full Flex.	\$ 58,088	\$ 1,405	\$ 41,012	\$ 64,053	\$ 70,054	\$ 10,130	\$ 29,043
	Min. Vol.	\$ 47,561	\$ 488	\$ 40,885	\$ 47,448	\$ 56,026	\$ 3,520	\$ 15,141
	Max. Vol.	\$ 64,518	\$ 619	\$ 56,222	\$ 64,212	\$ 72,498	\$ 4,463	\$ 16,276
Driver Hours per week	Full Flex.	2,411	49	1,846	2,593	2,722	354	875
	Min. Vol.	1,949	22.51	1,641	1,954	2,282	162	640
	Max. Vol.	2,853	28.67	2,441	2,833	3,261	207	820
ALOH	Full Flex.	116	1	96	118	138	10	43
	Min. Vol.	115	0.67	109	114	132	5	24
	Max. Vol.	104	0.50	97	104	115	4	18
%age Loaded Miles	Full Flex.	64%	0%	60%	64%	70%	2%	10%
	Min. Vol.	62%	0%	59%	63%	66%	1%	7%
	Max. Vol.	71%	0%	68%	71%	74%	1%	6%
Drivers used per week	Full Flex.	18	0	13	19	20	3	7
	Min. Vol.	14	0	12	14	17	1	5
	Max. Vol.	21	0	18	21	24	2	6
Available Drivers	Full Flex.	22	-	22	22	22	-	-
	Min. Vol.	22	-	22	22	22	-	-
	Max. Vol.	22	-	22	22	22	-	-

in terms of overall costs as well as the tightness of the bound for all the metrics generated during weekly execution.

7.3 Tour Level Metrics

This part of the analysis is critical to proving the robustness of the Annual Plan. The robustness can be tested by comparing the weekly occurrence of annual plan tours to all possibilities of tours generated under weekly operations with full flexibility. When we provide the model with complete flexibility during weekly execution, it comes up with the best possible routing for volume over the network in the form of different tours for that week. Hence, the complete flexibility planning model can be considered optimal in nature. However, the discussion in Section 7.1 indicates that providing this level of flexibility at a weekly level comes with many operational challenges. Many of the tours created under complete flexibility are inconsistent in nature. That is, many of these tours occur on an ad-hoc basis, depending on the demand conditions for that particular week only. Hence, these tours would not be a value-add for long term planning. On the other hand, there might be a particular sub-set of tours that occur consistently, even in the presence of complete flexibility. Knowing this set of robust tours before-hand would be extremely valuable in long-term planning. Also, if a majority of these tours are present in our Annual Plan, it would be an important indicator that our planning is robust, and it considers the demand variability component within the network effectively. Hence, this part of the thesis compares the tours created under the complete flexibility scenario with the Annual Plan tours to analyze the performance of the Annual Plan in the presence of demand variability. The final aim of this analysis would be to check how close our long-term planning is, with respect to the optimal weekly execution plan.

The following section defines key metrics that shall be used to analyze the performance of tours generated during weekly operations with full flexibility, and compare these tours to Annual Plan tours.

- **Tour occurrence/ consistency of occurrence** - is the probability of a tour

occurring during the 52 weekly simulations carried out under complete flexibility. The probability is indicative of the consistency with which a particular tour occurs over the 52 week period; the higher the number, the more robust the tour is.

$$\text{Tour occurrence} = \frac{\text{No. of times tour has occurred}}{\text{Total no. of weeks in simulation}}$$

- **Tour fickleness** - is a number that is indicative of the properties of lanes that make up a tour. We know that for a tour to occur consistently, it is important to understand the behavior of the lanes that make up the tour. There are 3 important indicators of lane behavior that directly affect the chances of a tour occurring within the network:
 - Lanes must have consistent occurrence of demand over the weeks
 - Lanes that get assigned to a large number of different tours reduce the chances of it having enough volume to cover a particular tour consistently
 - Lanes that consistently get assigned to for-hire carriers reduce the chances of it having enough volume to cover a particular tour consistently

Using these lane characteristics, we map out all the tours with the maximum number of times any lane within a tour has been assigned to different tours and been assigned to for-hire carriers over the 52-week execution period. The sum of these two numbers for a particular tour is defined as its tour fickleness factor. In order to better understand the calculation, consider the example indicated in Fig. 7-1:

Thus, as per the metrics defined, the fickleness of the tour in the example shown in Fig. 7-1 is the sum of the maximum number of times any lane within the tour has been assigned to a Private Fleet (49 as per the DC-S Lane) and the maximum number of times any lane within the tour has been assigned to For-Hire Carriers (23 as per the V-DC lane). Hence, the tour fickleness of the tour DC-S-V-DC is 72.

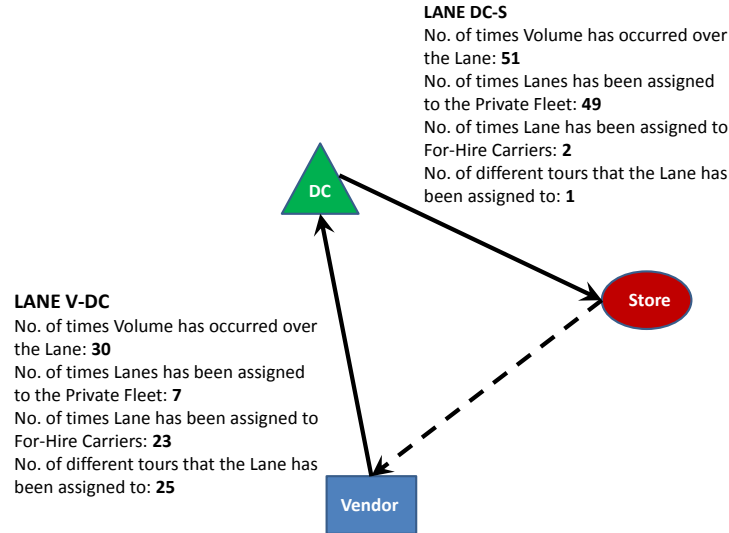


Figure 7-1: Example of a tour with lane details in order to explain metrics

Tour fickleness is an important characteristic that helps explain the actual tour occurrences within the weekly operational plans. It makes intuitive sense to say that if a tour is made up of fickle lanes, i.e., lanes that get assigned to a large number of different tours or lanes that consistently assigned to For-Hire carriers, the tour would find it difficult to occur in actual practice. Hence, we draw a correlation between the actual tour occurrence probabilities and their corresponding fickleness to understand how effectively this metric helps us in identifying fickle tours.

As shown in the Fig. 7-2, it can be seen that the Tour Fickleness is quite strongly negatively correlated with respect to the Tour Occurrence Probability. The correlation factor is -66%. For the Annual Plan tours, this factor is even higher at -76%. Hence, it can be said that the significance level of this factor in terms of detecting the actual tour occurrences is quite high.

- **Tour theoretical probability of occurrence** - relates to the demand occurrence over the lanes that comprise a particular tour. It can be intuitively said that a tour can actually occur in practice if and only if there is at least one load over all the lanes that make up the tour. Hence, the theoretical probability of

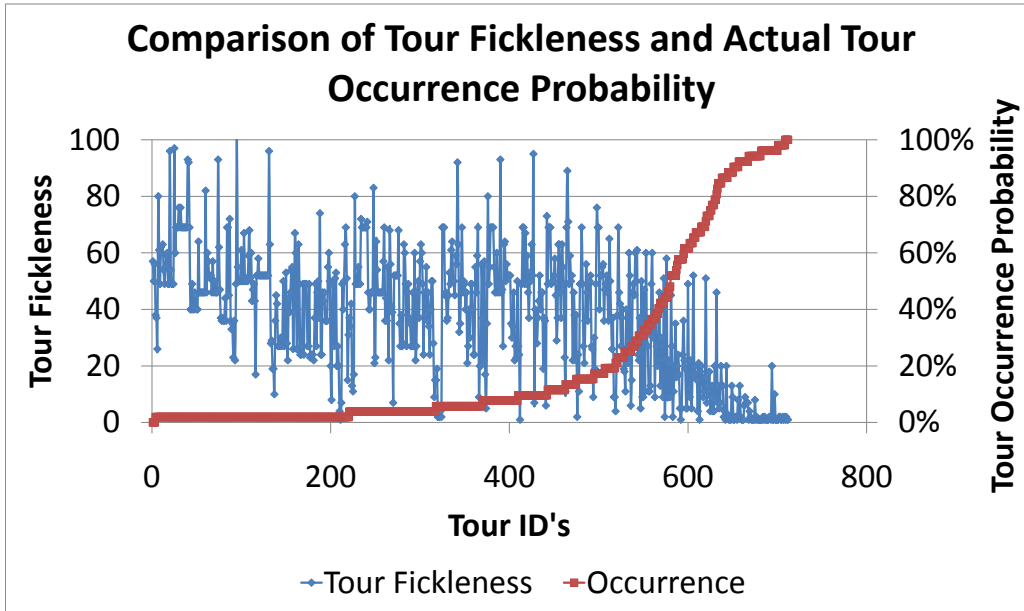


Figure 7-2: The graph gives a visual indication of the inverse correlation between the Actual Probability of Tour Occurrence and the Tour Fickleness factor calculated

a tour occurring can be defined as the product of the probabilities of non-zero demand occurring over all the lanes comprising the tour. This number can be calculated using the empirical histogram demand information for the lanes.

$$P_t = \prod_{l \in \{\text{lanes in tour}_t\}} P(D_l > 0)$$

$$P_t = 1 - \prod_{l \in \{\text{lanes in tour}_t\}} P(D_l = 0)$$

Using this formula, we can calculate the theoretical probability of occurrence for all the tours generated by the model, when provided with Full Flexibility. However, when we draw a comparison of this factor with the actual tour occurrences generated by the simulation of weekly operational plans, the correlation number is just 38%.

Hence, we raise the question, “*why is the actual tour occurrence probability obtained from simulations so different from the theoretical calculations?*” This

is because of network effects. The allocation of lane volume to a tour is not solely dependent on whether all the lanes within the tour have demand occurring every week. Because of the dynamic nature of weekly demand occurrence, there might be new tours that can be planned out which may be more profitable from one week to another. At the same time, there can be instances where allocation of demand to a tour might be more profitable in one week, and less in the other. This needs to be taken into account while considering the long term allocation of lane demand, and hence, just a theoretical calculation of demand occurrence over tours is not sufficient to explain the network effects that actually occur in practice. This metric highlights the importance of developing a simultaneous optimization strategy for network planning and optimization, instead of the two-step strategy proposed by Mulqueen.

- **Tour demand statistics** - the mean and standard deviations of the demands occurring over the lanes that comprise the tours can also help generate a lot of insight into its occurrence in practice. It can be intuitively said that tours having a high mean and lower standard deviation have a higher probability of occurring in practice, as compared to tours that have a low mean and a wider spread of demand. Thus, analyzing this information can help us identify tours that would be robust in actual occurrence, and can be pro-actively included in the long-term planning.

Comparison of Annual Plan Tours versus Full Flexibility Simulation Tours

In order to check for robustness of the Annual Plan, it is critical to compare the tours created in the planning run with the tours generated in the actual weekly operational runs to note the similarities and differences.

Under the full flexibility scenario, the simulations over the 52-week period resulted in a total of 709 different tours being created by the model, depending upon the lane demand characteristics for each week. However, as mentioned in the discussion about

tour occurrence probability, most of these tours occur on an ad-hoc basis. This is visible from the histogram of tour occurrence shown in Fig. 7-3.

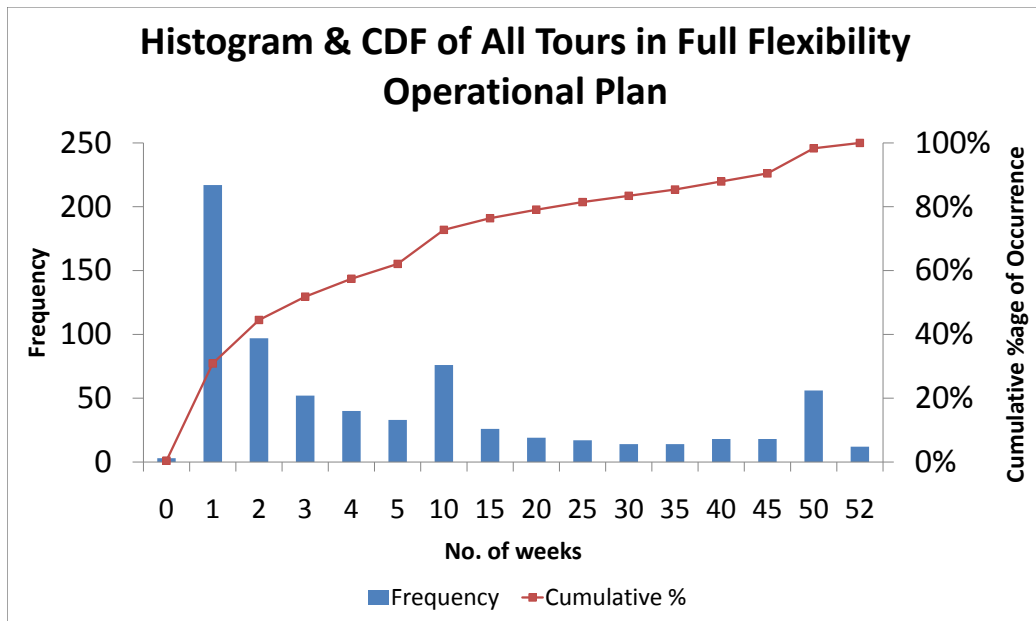


Figure 7-3: Histogram and CDF of Occurrence of All Tours whilst running the model with Full Flexibility

We can see from Fig. 7-3 that almost 62% of the all possible tours generated by the model occur 5 times or less in the 52 week period; also, 38% of the tours generated occur only once. These tours are extremely fickle and of no use for long-term planning. However, the peculiarity about the histogram is the small bump that can be seen at the far right end of the graph. These are the tours that have consistently occurred week after week in the simulation, and are robust in nature. These tours would be ideal for long-term planning, since the model is choosing them in consistently, even when it has the option of full flexibility. Hence, we need to dig a bit deeper and analyze the characteristics of these tours to get a better understanding of what makes them robust. If a large proportion of these tours are present within our Annual Plan, it would be a good indication that our planning is effectively taking demand variability into account, and has a good chance of handling operational uncertainty in the long run.

The graph shown in Fig. 7-4 shows the distribution of weekly occurrences of tours

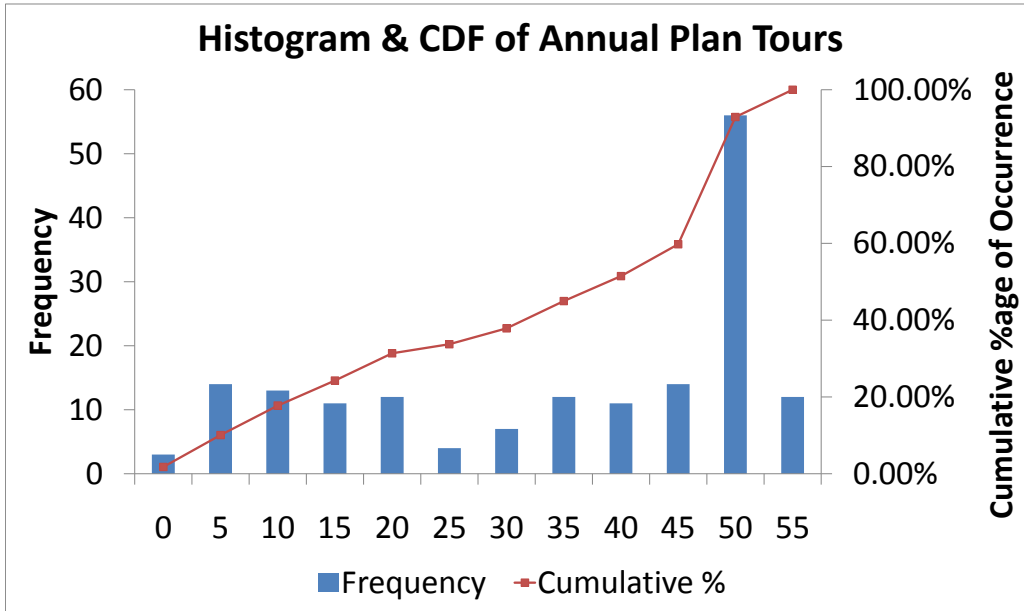


Figure 7-4: Histogram and CDF of Occurrence of All Tours occurring within the Stochastic Annual Plan

that are present in the Annual Plan. The Annual Plan comprises of just a 169 tours in comparison to the 709 that are generated by the full flexibility option. However, it is clearly visible, that a high proportion (50%) of tours present in the Annual Plan occurs in the weekly simulations consistently. This validates our hypothesis that the consistent portion of the tours at the far right of the graph in Figure 7-3 come from the Annual Plan tours!

The result discussed above has huge implications as far as long-term planning is concerned, and is a key milestone in our research work. Using simulations, it is shown that the Annual Plan is actually capable of handling the demand uncertainty that occurs in practice effectively. Another aspect highlighted by this result, and reinforced by the Minimum Volume Assignment case, is the ease of operation that can be achieved if we can trust the Annual Plan to handle the uncertainties in demand during weekly execution.

Another interesting aspect of the Annual Plan tours is highlighted when we compare the weekly tours generated in the simulations of the Full Flexibility plan with the Annual Plan tours. This comparison is termed as tour commonality, and it is

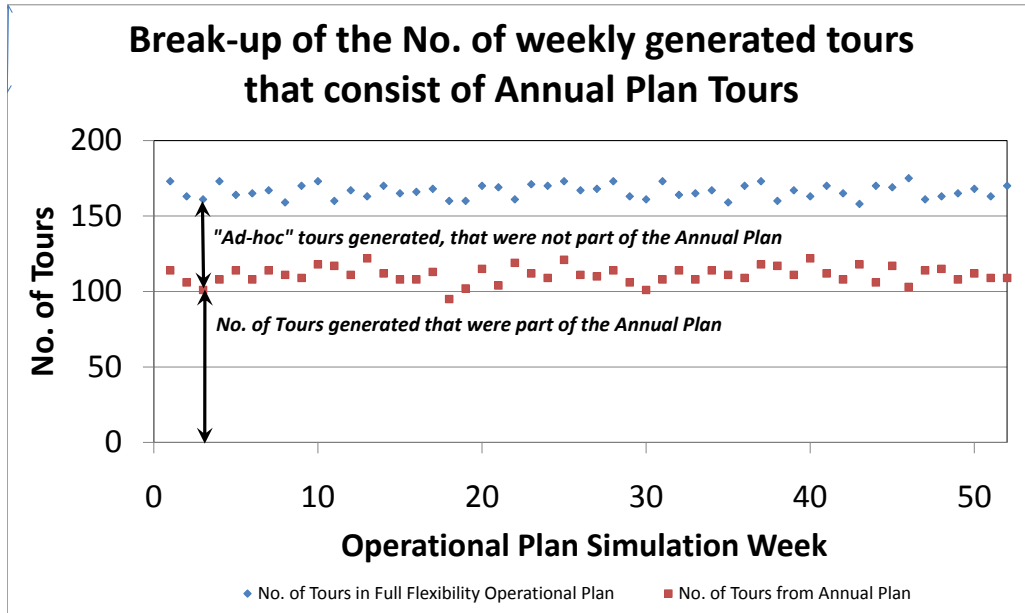
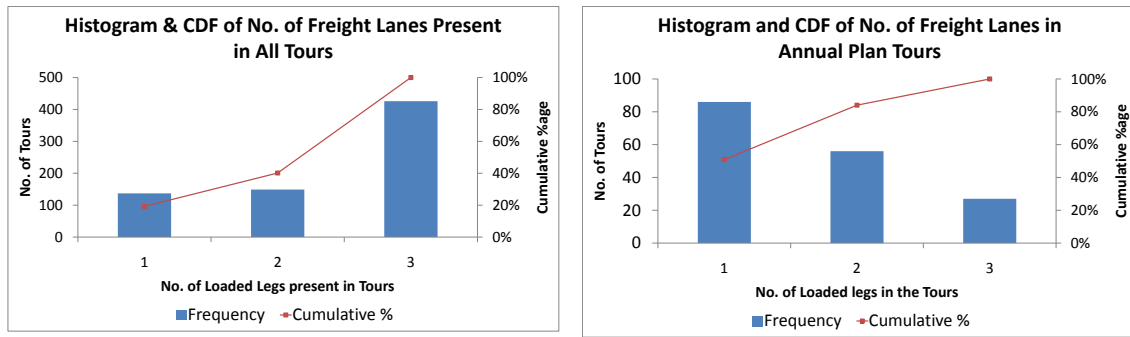


Figure 7-5: The graph indicates that about 71% of tours generated on a weekly basis within the Full Flexibility operational plan that are already present in the Stochastic Annual Plan

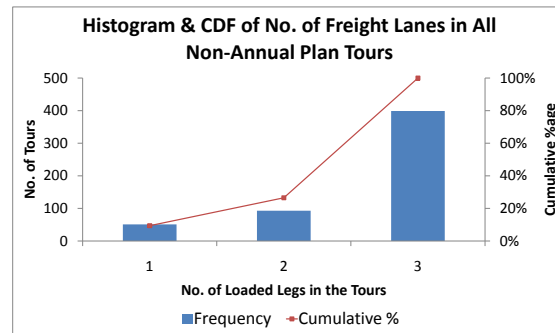
indicative of the proportion of tours generated on a weekly basis, that are a part of the Annual Plan. The graph in Fig. 7-5 shows that about 71% of the tours generated on a weekly basis for the Full Flexibility option were already a part of the Annual Plan, while the remaining 29% were ad-hoc, one-off tours that were generated by the model, as deemed fit, based upon weekly demand conditions. This consistency in weekly tours can give us a better understanding about the sweet spot of Annual Plan tours, and help us narrow down to the robust tours more effectively.

Now that we have proved that a large portion of Annual Plan tours are robust in nature, we need to figure out if there are any characteristics of these tours that assist them in being robust, as compared to the various combinations of tours created in the Full Flexibility plan. It can be intuitively said, that if a tour has a high number of loaded legs, its probability of occurring consistently decreases. This is because the chances of all the legs within a tour having enough volume to be carried everyday reduce. Hence, we analyze the tours by breaking them down into two categories. Firstly, we look at the histogram for the number of loaded legs in all tours created by



(a)

(b)



(c)

Figure 7-6: (a) Distribution of the number of freight lanes present in all tours generated in the Full Flexibility Operational Plan; (b) Distribution of the number of freight lanes present in the Stochastic Annual Plan; and (c) Distribution of the number of freight lanes present in all tours generated in the Full Flexibility Operational Plan which are not present in the Stochastic Annual Plan

the Full Flexibility plan (Fig. 7-6a), and then break them down into tours present in the Annual Plan (Fig. 7-6b) and those that are not (Fig. 7-6c).

As is visible in Fig. 7-6, the characteristics of tours that are present within the Annual Plan bear quite different statistics as compared to the tours that do not. This is because, as the model tries to optimize the network volume at a weekly level in the Full Flexibility plan, it creates a lot of long, one-off tours that do not consistently occur week-after-week. On the other hand, the Annual Plan tours are relatively shorter in length, which increasing their chances of reoccurring on a weekly basis. The long tours that are created during the weekly execution might be profitable to operate in terms of reduced empty miles, but there is a trade-off involved between the potential cost savings that can be obtained and their consistency. This is a trade-off

that the long-term planners must account for, when allocating lane volumes to tours.

7.4 Lane Level Metrics

This part of the analysis takes a closer look at the characteristics of the freight lanes present in the network, and tries to draw generalizations about its effect on long-term annual planning. This section follows up the tour level analysis and digs deeper into the performance of lanes in the complete flexibility scenario on a weekly basis.

It can be intuitively said that information regarding lanes that have consistent volume, show consistency with regards to assignment over the private fleet or for-hire carriers and show loyalty of assignment to a small set of tours within the network. Knowledge of these robust lanes is quite useful for long-term planning, since it helps create a sub-segment of consistent lanes, that can be allocated to the annual plan. The following metrics have been developed to dig deeper into the lane performance characteristics and can be used in conjunction with the tour metrics to understand the robustness of the Annual Plan.

- **Lane volume occurrence** - is the factor that indicates the number of times volume has occurred on a lane over the 52-week simulation period. For example, referring to the Fig. 7-1, the lane volume occurrence for the V-DC lane is 30 and for the DC-S lane is 51.

$$\text{Lane Volume Occurrence} = \frac{\text{No. of weeks in which volume has occurred on lane}}{\text{Total No. of weeks in Simulation}}$$

As visible in Fig. 7-7, most of the lanes within the network have consistent volume occurrences over the simulation period. However, there are certain fickle lanes with very few occurrences of volume on the far left of the graph. Knowledge of these lanes is quite important for long-term planning purposes, since consistency of volume is crucial for planning routes over the private fleet, as well as for bidding out lanes to for-hire.

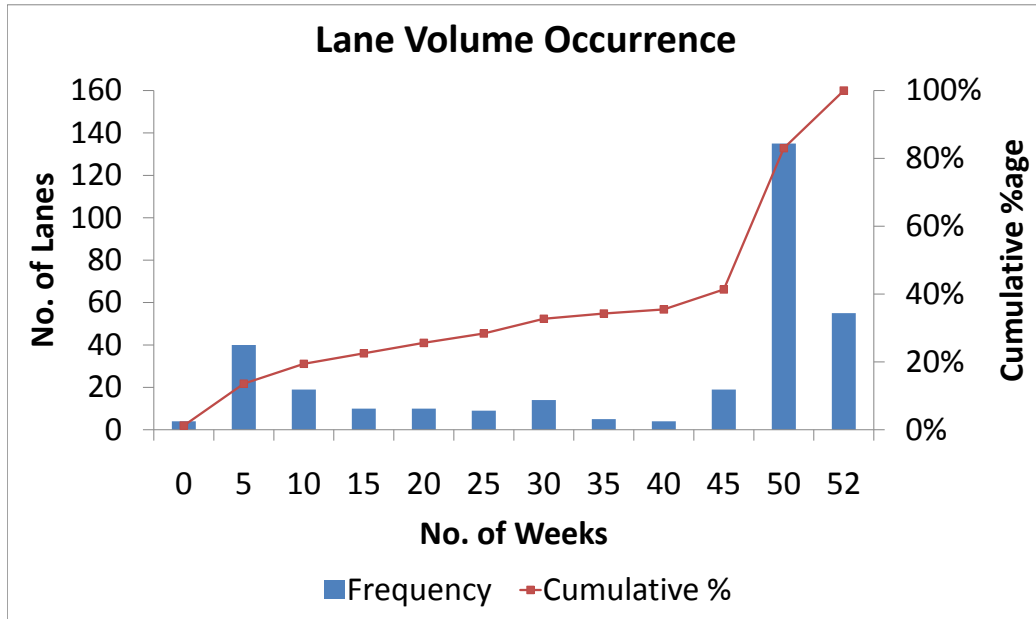


Figure 7-7: The graph depicts the distribution of volume occurrence on the lanes within the network over the 52-week simulation period

- Lane Fleet fickleness** - is a factor that indicates the number of times volume over a lane has been assigned to private fleet, whenever volume has occurred on the lane. Along with lane volume occurrence, this factor is a good indicator of the robustness of a lane in terms of its assignment to private fleet, which is valuable information for long-term planning. Thus, a lower lane fleet fickleness factor is indicative of the fact that whenever volume occurs on the lane, it has a high probability (robustness) of being assigned to the private fleet, and vice-versa. For instance, referring to the generic example shown in Fig. 7-1, the Lane Fleet Fickleness for the V-DC lane is $(1 - 7/30) \times 100 = 77\%$, and for the DC-S lane is $(1 - 49/51) \times 100 = 4\%$.

$$\text{Lane Fleet fickleness} = 1 - \frac{\text{No. of times volume on lane has been assigned to Fleet}}{\text{No. of weeks in which volume has occurred on lane}}$$

As visible in the Fig. 7-8, there are two distinct extremes for lane fleet fickleness values. This information is indicative of the fact that lanes show consistent behavior as far as assignment to fleet is concerned. There is a substantial

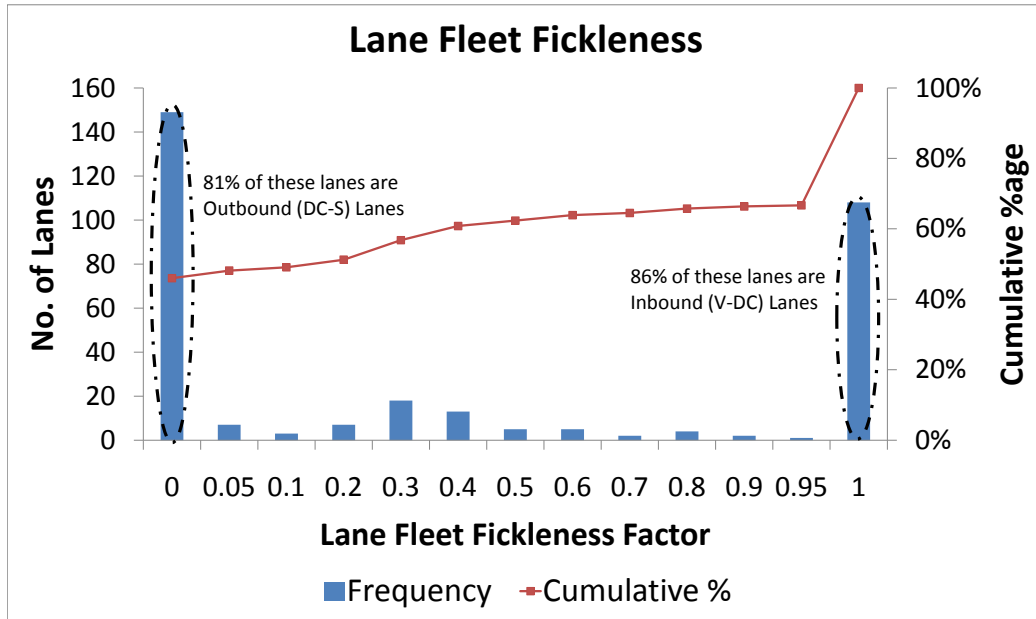


Figure 7-8: The graph depicts the distribution of the lane fleet fickleness factor over the 52-week simulation period

portion of lanes that are consistently assigned to fleet, whenever volume occurs on them, and most of the remaining lanes are always more economically viable for assignment to for-hire carriers. This information is useful for long-term planning purposes, and can be incorporated easily within the Annual Planning scenario.

- Lane Tour Fickleness** - is a factor that indicates the maximum number of different tours that the lane gets assigned to, whenever volume occurs on the lane. This factor is a good indicator of lane popularity within the private fleet. If Lane Tour Fickleness factor is low, it indicates that whenever volume occurs on the lane, it is consistently assigned to a certain small sub-set of tours, and hence, can be considered as robust for annual planning purposes. For instance, referring to the generic example shown in Fig. 7-1, the Lane Tour Fickleness for the V-DC lane is $25/30 \times 100 = 83\%$, and for the DC-S lane is $1/51 \times 100 = 2\%$.

$$\text{Lane Tour fickleness} = \frac{\text{Max. no. of diff. tours to which lane vol. has been assigned}}{\text{No. of weeks in which volume has occurred on lane}}$$

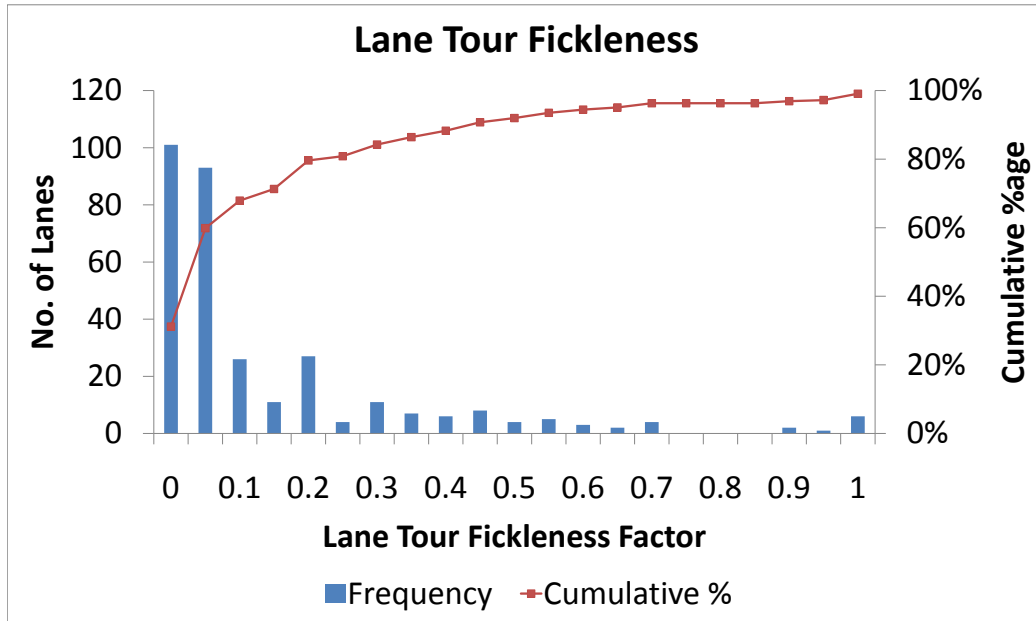


Figure 7-9: The graph depicts the distribution of the lane fleet fickleness factor over the 52-week simulation period

As can be seen in Fig. 7-9, there is a substantial portion of lanes that are consistently assigned to a small sub-set of tours whenever volume occurs on them. This information is particularly useful for long-term planning, since these lanes exhibit consistent behavior of assignment to tours. Hence, this information can be incorporated in the Annual Planning scenario effectively.

- Lane For-Hire Fickleness** - is a factor that indicates the number of times a particular lane has been assigned to for-hire carriers, whenever volume occurs on them. When used with lane volume occurrence factor, this number provides us with information regarding the robustness of the lane, and can be an important consideration during long-term planning. Thus, a low value of the lane for-hire fickleness factor indicates consistent (robust) assignment to for-hire carriers, whenever volume occurs over the lane. For instance, referring to the generic example shown in Fig. 7-1, the Lane For-Hire Fickleness for the V-DC lane is $23/30 \times 100 = 24\%$, and for the DC-S lane is $2/52 \times 100 = 96\%$.

$$\text{Lane For-Hire fickleness} = 1 - \frac{\text{No. of times volume on lane has been assigned to FH}}{\text{No. of weeks in which volume has occurred on lane}}$$

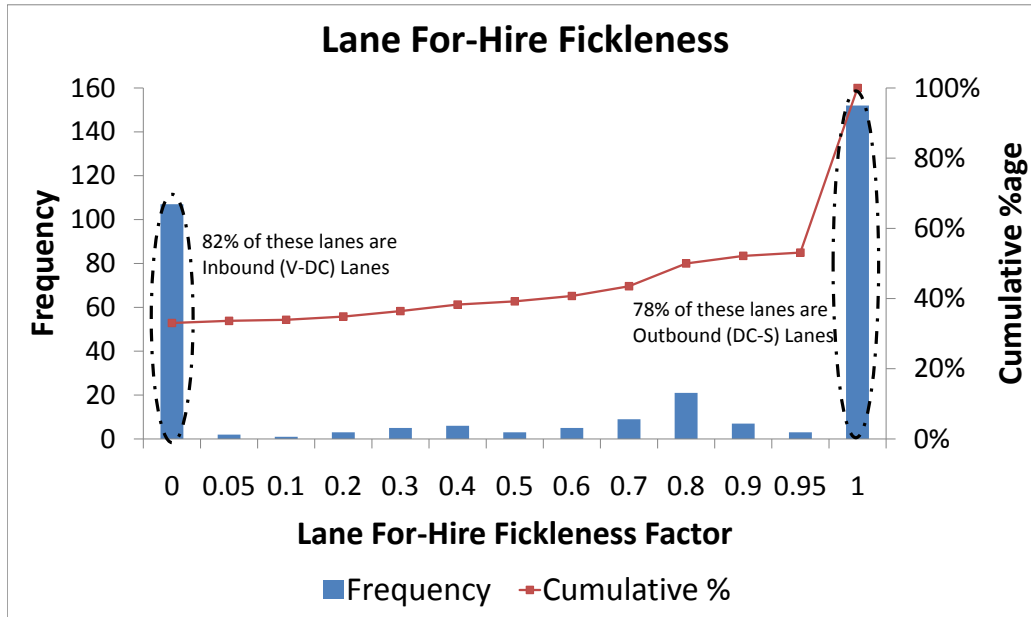


Figure 7-10: The graph depicts the distribution of the lane fleet fickleness factor over the 52-week simulation period

The behavior of the Lane For-Hire Fickleness factor, as seen in Fig. 7-10 is quite similar to the Fleet Fickleness factor. The lane behavior can be divided into two extremes, one which shows consistent allocation to for-hire carriers whenever volume occurs on them; and the other end, which is consistently assigned to private fleet upon volume occurrence.

- Lane demand statistics** - provide insight into the demand occurring over the lane. Using these statistics, insights can be drawn with regard to the effect of lane demand consistency on annual planning methods. Generically speaking, lanes with high mean demand and low standard deviations are expected to be more robust than lanes with low mean demand and high standard deviations. Using this information, we can further build upon the analysis to search for robust lanes on a pro-active basis. However, solely using lane demand statistics as the criterion for deciding whether a lane is robust or not would be insufficient because of the network effects, as discussed previously. An graphical example of lane demand statistics can be seen in Fig.7-11.

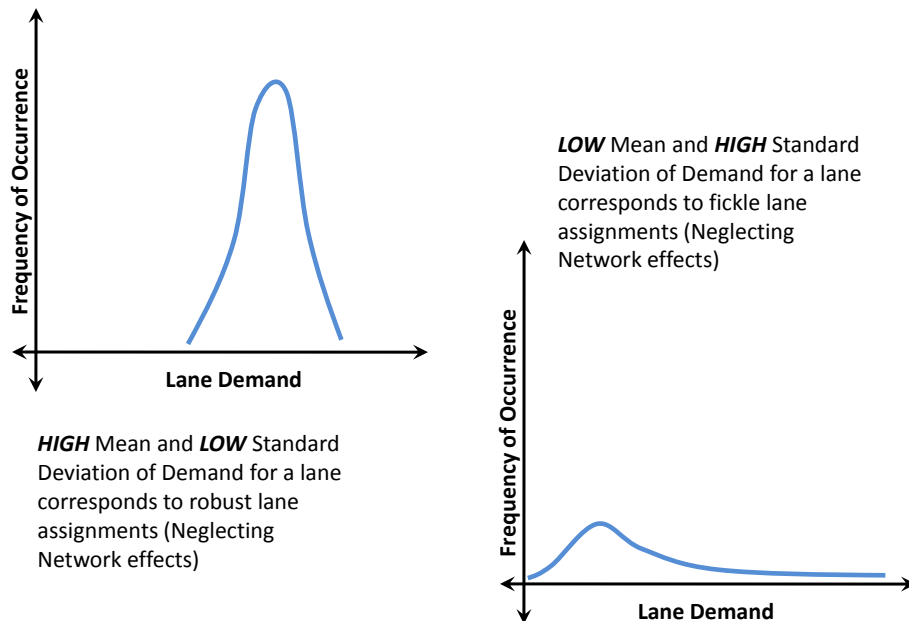


Figure 7-11: Depiction of the use of lane demand statistics as a criterion for lane robustness

- **Qualitative assessment of robustness** - the metrics discussed in this section help us in getting a better understanding of the behavior exhibited by lanes, and throw light on the corresponding implications for long-term planning purposes. However, an important thing to keep in mind is that these metrics are not independent of each other, i.e., in order to get a complete picture about lane robustness, multiple metrics need to be looked at simultaneously. Hence, further analysis needs to be done on this inter-dependence to draw definitive insights from this study.

A lane can definitively be termed as robust for private fleet only if it exhibits the following characteristics simultaneously:

- It has high volume occurrence
- It has a high lane fleet fickleness factor
- It has a low lane tour fickleness factor
- It has a low lane for-hire fickleness factor

Similarly, a lane can be definitively termed as robust for for-hire carriage only if it exhibits the following characteristics simultaneously:

- It has high volume occurrence
- It has a low lane fleet fickleness factor
- It has a low lane tour fickleness factor
- It has a high lane for-hire fickleness factor

Based upon these rules and by observing the results obtained for all the lane-level metrics within our network, we have set up a decision rule to classify lanes into three categories:

- Robust for tours
 - * Lane Volume Occurrence > 32
 - * Lane Fleet Fickleness Factor > 0.9
 - * Lane Tour Fickleness Factor < 0.2
 - * Lane For-Hire Fickleness Factor < 0.2
- Robust for for-hire carriers
 - * Lane Volume Occurrence > 32
 - * Lane Fleet Fickleness Factor < 0.1
 - * Lane Tour Fickleness Factor < 0.2
 - * Lane For-Hire Fickleness Factor > 0.9
- Fickle lanes All the remaining lanes that do not fit into the above two categories

Thus, we can incorporate this information of robust lanes while creating our long-term planning scenarios. Using this information, we can modify our Annual Plan further. As a result, we can generate a tighter bound on the demand uncertainty that we wish to address within our Annual Plan, depending upon our risk appetite.

However, it needs to be noted that the results obtained from the analysis are ex-post, i.e. the robustness of tours and lanes within the annual plan can be definitively measured only after the analysis is complete. Hence, it would be difficult to propose a robust annual plan on a pro-active basis, since each transportation network would have its own intricacies. In order to address this issue, the analysis in this thesis needs to be re-done for historical demand information for the network under consideration, so as to understand the robustness characteristics of the lanes and tours comprising the Annual Plan.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 8

Conclusions

The competitive market dynamics in the trucking industry have, and will continue to push private fleet owners and for-hire carriers towards achieving greater operational efficiencies without compromising on service levels. This trade-off can be addressed using robust transportation planning methodologies that address the demand variability inherent in the transportation networks.

An efficient transportation policy needs to bridge the gap between the expected long term planning and weekly operational occurrences of demand, and propose simple, yet effective ways of execution. The thesis addresses this problem by creating a transportation plan on an annual basis, keeping demand uncertainty in mind, through stochastic metrics. Additionally, it takes the analysis a step further by creation of heuristics with which the annual plan can be tested (through simulations) at a weekly, operational level at varying degrees of flexibility, to satisfy demand in a simplistic manner.

Comparisons of the weekly operations using heuristics like Minimum Volume Assignment and the volume allocation under Complete Flexibility showcase the robustness of the annual plan. In order to analyze the operational plans extensively with regards to its robustness in comparison to the Annual Plan, the simulation results were analyzed at four levels (in decreasing order of aggregation) at the Network Level, Facility (Distribution Center) Level, Tour Level and Lane Level.

The Network level results provide an overview of how the overall weekly operations

created under zero flexibility (through the use of heuristics) perform in comparison to the full flexibility option. Through the use of key metrics, the results of this analysis highlight that complete re-optimization of the network at a weekly level produces little incremental cost benefit (6% of cost savings), at the expense of greatly increasing the overall network complexity at an operational level. This is a trade-off that all companies must take into account, since it helps draw a comparison between increased network complexity for an optimal allocation and the potential cost savings that could be generated. Additionally, a comparison of the standard deviations indicate that the Minimum Volume Assignment numbers have a much tighter bound in comparison to the Complete Flexibility option, which is a direct indicator of the reliability of the results. The Facility level analysis also mirrors the results observed at the network level, and help strengthen the insights generated from the metrics.

The Tour Level analysis digs deeper into the occurrence of tours generated within the Annual Plan versus those generated weekly, on an ad-hoc basis, in order to understand their robustness in the presence of demand variability. Using the complete flexibility scenario as the the best-case, we are able to analyze what the ideal volume allocation on a weekly basis looks like, and compare that to the Annual Plan allocation. A key observation is that, on average, 71% of tours generated on a weekly basis under complete flexibility are already present within the Annual Plan. The remaining tours are ad-hoc assignments based on the nature of demand occurrence for a particular week. The consistency of this behavior leads us to believe that the long-term strategic plan generated by addressing demand variability can handle the weekly fluctuations effectively, and can be used in practice. The Annual Plan tour consistency can be attributed to the fact that the planning mechanism takes the demand variability on lanes into account, before allocating lane demand to the private fleet. This, in turn, results in the the Annual Plan tours being shorter (lower number of loaded legs in the tour) and more consistent, in comparison to the ad-hoc tours which are longer, and hence, difficult to materialize consistently in practice.

Lastly, the Lane Level metrics try to understand the behavior of lane volume assignments to private fleet and for-hire carriers in the complete flexibility scenario.

Using lane demand statistics and lane fickleness with respect to tours, fleet and for-hire as the key metrics, we are able segregate the network into robust versus fickle lanes. This segregation can provide valuable information about the network to operations planners during the development of a strategic transportation plan.

The simplicity of execution is a crucial aspect of transportation planning, especially when we keep in mind that the execution of plans shall be carried out by operational managers at a local level. Hence, by providing an easy-to-handle planning module that does not need much ad-hoc decision making would reduce the chances of sub-optimal localized decisions and maintain the network allocation at close to system optimal levels. This would directly translate into lowering the gap between expected transportation costs proposed by the strategic plan and actual costs incurred during execution.

8.1 Opportunities for Future Research

The thesis has made several assumptions while performing the aforementioned analysis, that could be valuable points to be addressed as future research opportunities.

- **Load Volume Pickup and Delivery dates/ time windows**

This analysis assumes that loads occur in weekly time buckets, and a more granular time window is not supported by the model. In order to use the analysis for actual day-to-day planning, the time window needs to be a lot smaller. Additionally, with the push of the Just-In-Time model where deliveries are synchronized within a specific time window during the day, this assumption constrains the applicability of the model in real life.

- **Network size and lane demand patterns**

The analysis within this thesis is performed on a small sub-network for Walmart. As such, it would be interesting to check for the applicability of the results observed in this thesis on a generalized transportation network which is larger and/or has a different demand pattern on the inbound and outbound lanes. It

is our hypothesis that the results obtained in the thesis can be generalized to a retailer network with sporadic inbound loads and consistent outbound loads.

- **Pro-active (Ex-ante) analysis of robustness**

The analysis carried out in this thesis proposes a robust Annual Planning mechanism after understanding the assignments of network demand over the lanes and tours. However, there could be opportunities for generalizing the results and proposing a planning strategy which pro-actively segregates the tours and lanes within a network as robust versus fickle. This can be done by analyzing the demand patterns on lanes and tours along with lane pricing strategies, to see which ones are more consistent than others.

- **Collaborative logistics**

Pooling logistics capabilities for private fleet and for-hire carriers could provide interesting opportunities for minimizing empty percentage miles and maximizing asset utilization, since pooled resources would reduce the supply variability within the network. However, this pooling effort could bring in a lot of complexities for assigning volume to loads on popular lanes, as well as issues related to transfer pricing. Additionally, having a dynamic fleet availability at the domiciles would increase the complexity for the optimization model.

Bibliography

- P. Bradley. Taking the company private. *DC Velocity*, 3(5):44–46, 2005.
- C. Caplice F. Jauffred. Working paper, 2010.
- FleetOwner.com. America’s top private fleets. Technical report, February 2010.
- M. Harding. Can shippers and carriers benefit from more robust transportation planning methodologies? Master’s thesis, Massachusetts Institute of Technology, Engineering Systems Division, <http://hdl.handle.net/1721.1/33322>, 2005.
- IBISWorld. Ibisworld insutry report: Long-distance freight trucking in the u.s. Technical report, January 2010.
- C. Iliadis. Can we improve the freight network solutions of a large retailer by using probability distributions to model demand? Master’s thesis, Massachusetts Institute of Technology, Dept. of Civil and Environmental Engineering, <http://hdl.handle.net/1721.1/51626>, 2009.
- A.H.G. Rinnooy Kan J.K. Lenstra. Complexity of vehicle routing and scheduling problems. *Networks*, 11(2):221–227, 2006.
- K. Kirkeby. Industry surveys transportation: Commercial. Technical report, Standard & Poor’s Industry Surveys, January 2010.
- B. Ryder M. Arnold, M. Hind. An empirical study of selective optimization, 2001. Rutgers University in collaboration with IBM Thomas J. Watson Research Center.
- M. Adhi Dipo M. Margreta, C. Ford. U.s. freight on the move: Highlights from the 2007 commodity flow survey preliminary data. Technical report, Research and Innovative Technology Administration Bureau of Transportation Statistics Special Report, September 2009.
- M. Mulqueen. Creating transportation policy in a network that utilizes both contract carriers and an internally managed fleet. Master’s thesis, Massachusetts Institute of Technology, Engineering Systems Division, <http://hdl.handle.net/1721.1/36141>, 2006.
- J. Schulz. Top 50 trucking companies trucking’s game-changing moment. *Logistics Management*, April 2010.

A. Terreri. Shippers are adding private fleets to their transportation mix. *World Trade Magazine*, 2006.

USDOT. Freight in america. Technical report, Research and Innovative Technology Administration Bureau of Transportation Statistics, January 2006.