ABSTRACT

SARANGI, BIBHUTI BHUSAN. Public Consolidation Mechanisms for Freight Transport. (Under the direction of Michael G. Kay).

Consolidating multiple shipments onto a single carrier can increase the efficiency of freight transport. A shipper using only for-hire trucking for one-time shipments to different destinations is limited to using a less-than-truckload (LTL) service or a third-party private consolidation service that matches carriers to shipments using private information regarding local carrier supply and shipment demand. The shipper is given a take-it-or-leave-it offer from the consolidator because it is not feasible to negotiate for a one-time shipment. A public consolidation mechanism is proposed as an alternative to private freight consolidation. Since shippers willing to transfer their loads have to depend on a private carrier, such carriers get access to critical shipment information from the shippers which they subsequently use to form consolidated loads and thus generate a large amount of revenue. Since shippers are thus at a disadvantage and the market so formed is inefficient, the consolidation mechanism proposed in this thesis will improve market efficiency by posting public information regarding the shipments available for consolidation and the carriers available for transport. By using public information and standardized route procedures, the mechanism allows shippers to form consolidated loads that will then be posted to trucks for acceptance. The mechanism also develops protocols for trucks to accept consolidated loads, including reneging and accepting loads between pairs of trucks. The mechanism will also be extended to include time preferences and their impacts on shipments. While an agent-based approach has not been undertaken in this thesis to simulate the mechanism, a number of measures, such as softmax functions and linear assignment, have been used to evaluate the mechanism. Comparing shipment consolidation done by the public consolidation mechanism with that done by a private firm reveals that the overall costs are lower for a private firm, but when the shipment market becomes thick so that the overall reach of one private firm becomes limited, the public mechanism provides the lower cost. Experiments with various cost allocation methods in the mechanism show that the proportional cost allocation gives a lower overall transport cost for the public consolidation mechanism. Greedy assignment was used to assign trucks to loads with the objective of maximizing the total profits of trucks assigned to loads. With this approach, trucks find opportunities to improve their profits by reneging and swapping

their loads with other trucks. A model for time preferences has been developed using a simplified home delivery network that considers customers to be having high and low time priorities. The same home delivery model has also been used, in conjunction with a public logistics network, to develop a basic agent-based model consisting of shipment and truck agents that use load consolidation services and allocation methods to propose loads to truck agents that deliver items to customers.

Public Consolidation Mechanisms for Freight Transport

by
Bibhuti Bhusan Sarangi

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Industrial Engineering

Raleigh, North Carolina
2025

APPROVED BY:

Russell King	Brandon McConnell
Donald Warsing	Michael G. Kay Chair of Advisory Committee

BIOGRAPHY

The author was born in Sambalpur in the state of Odisha, India, where he completed his schooling and his engineering graduation. He completed his masters from Indian Institute of Technology, Roorkee. He then joined the Industrial and Systems Engineering department of North Carolina State University as a PhD candidate.

ACKNOWLEDGEMENTS

I thank Dr. Michael Kay for his guidance during my PhD journey. I am grateful to him for agreeing to become my advisor. Given his busy schedule, Dr. Kay has been very generous with his time. I admire his patience in answering my questions, clearing my doubts, and helping me progress with the thesis work. While I stumbled through completing the various chapters, he always had a long-term vision of the thesis. Dr. Kay not only has superb knowledge in the field of logistics, his grasp of the fundamentals of economics is also excellent. I thank him for introducing me to authors like Peyton Young, Ken Binmore, and others during our many discussions. I also appreciate the opportunity Dr. Kay has given me by arranging access to Coupa's Llamasoft software that wonderfully blends theoretical knowledge of supply chain with modeling practical problems. Dr. Kay is a great advisor. I am fortunate to be his student.

I am also grateful to Dr. Brandon McConnell, Dr. Russell King, and Dr. Donald Warsing for agreeing to be on my thesis committee. I appreciate the insightful comments they made and the advice they offered, which were crucial in completing this thesis.

I thank the faculty in the ISE department and the administration staff for developing an environment that fosters research. The course professors under whom I studied were excellent. Jasmine Petway and Kendall Walker were very prompt in responding to my queries and clearing any administrative doubts.

I am also grateful to my wife Anita, my daughter Shreen, and my family in India for their patience and support during my PhD journey. I thank them for being with me on this long journey and for providing constant encouragement during my doctoral studies.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER

1

INTRODUCTION

1.1 Overview

Consolidating multiple shipments onto a single carrier can increase the efficiency of freight transport. In trucking, a firm that operates its own private fleet for its shipping needs can look for consolidation opportunities as part of its normal transportation planning process. A shipper using only for-hire trucking for small one-time shipments to different destinations is limited to using a less-than-truckload (LTL) service or a third-party private consolidation service that matches carriers to shipments using private information regarding local carrier supply and shipment demand. The shipper is given a take-it-or-leave-it offer from the consolidator because it is not feasible to negotiate for a one-time shipment. A public consolidation mechanism is proposed as an alternative to private freight consolidation. Such a mechanism posts public information regarding the shipments available for consolidation and the carriers available for transport. By making the information public and standardizing the route formation and allocation procedures, each shipment and carrier can determine which loads to join or transport, respectively.

The primary motivation for developing the public consolidation mechanism is to improve economic efficiency. Truck capabilities provided by carriers are generic in the sense that the location of trucks and the type of service that they provide are known, while it is the shippers that want to ship items that have the relevant information (for example, their time sensitivities to pick up and delivery of their shipments). This is in contrast to most market transactions where the efficiency of the market supports the entity with the more specialized information. For example, a manufacturer not only knows the various types of items that it can produce but also understands the relative market demand of each category of the items that it produces with respect to other similar manufacturers. It is such tacit knowledge among various manufacturers that will ultimately drive demand and supply of their production leading to an efficient market. On the other hand, in logistics it is the shippers that have the more private and critical knowledge that does not get reflected in the market. Carriers generate a significant amount of revenue from consolidation since the shippers have to rely on carriers and their trucks to ship their commodities. This leads to market inefficiency since the carriers have private knowledge about their shipments while an individual shipper using the trucks of a carrier will neither be aware of the presence of other shippers nor whether an efficient route was determined by the carrier for transportation. The proposed public consolidation mechanism will make this market more efficient by allowing shippers to come together to form consolidated loads and propose those loads to trucks thereby keeping critical shipment information private and making public only those pieces of information that will be required to operate the mechanism.

The proposed public coordination mechanism (Figure 1.1) is also motivated by the same desire to reduce friction that gives ridesharing apps like Uber their popularity. Attempts have been made to apply this same approach to transport, but companies like Uber Freight have had limited success because they assume that people and freight are the same (Helguera Sanchez and Hendra Mukti 2018). This research recognizes the fundamental difference between people and freight with respect to the variability of their time sensitivity. No person would be willing to arrive 24 hours early at an airport to save half the cost of a rideshare to get there an hour before their flight. In contrast, such savings could significantly reduce the cost of providing a shipment that is not needed at its destination for at least a day, while another shipment may be a critical spare part that needs to reach its destination as soon as possible. The proposed consolidation mechanism flips the ridesharing approach and makes the

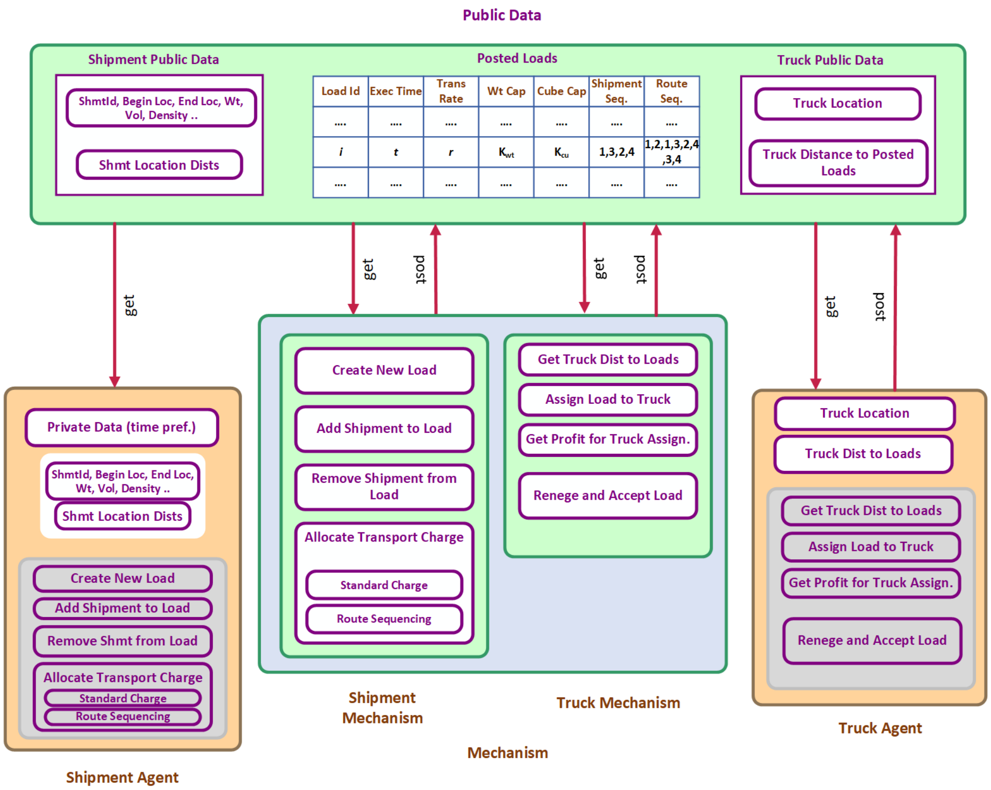

Figure 1.1: Public consolidation mechanism.

carrier, instead of the passenger, the price taker. For instance, a shipment can use its time sensitivity to adjust its offer for transport subject to its desired pickup and delivery times.

The proposed work also addresses a very significant inefficiency associated with trucking operations, namely, the difficulty in finding shipments to fill a truck. Trucks are estimated to travel empty for 20% of their miles, increasing costs, greenhouse emissions, and traffic congestion (Kearney 2022). Total trucking costs in the U.S. were $830.4 billion in 2022, representing 45% of total logistics costs and 3.26% of the U.S. GDP. The proposed work, if successful, will demonstrate that it is possible to find shipments to fill otherwise empty miles and, even when not traveling empty, it is possible to increase the utilization of a truck's capacity by finding additional shipments for consolidation. This increase in efficiency will reduce overall freight transport costs because almost every ocean, rail, or air cargo shipment starts and ends its journey via truck transport, which will have widespread beneficial impacts throughout the supply chain. The U.S. trucking industry is very fragmented, with 350,000 independent single-vehicle owner-operators and, of the 110,000 carriers, 90% own six or fewer trucks (Trucks 2023). The public consolidation mechanism developed in this work will allow a single owner-operator to access the same information regarding shipment availability as the largest carrier.

As the proposed mechanism will be used by shippers willing to form consolidated loads, there will be transaction costs associated with running the mechanism. Since there will be many shippers and carriers spread across a wide geographical area vying for the services of the mechanism, services that will include online route formation, cost allocation, and assigning loads to trucks, a centralized server running the mechanism will become easily overwhelmed with all the requests that it will receive to form and assign loads. Instead of a centralized server doing all the real-time computations, it is expected that each shipper will be able to run the computation locally, which will include route formation when new shippers join or leave the load. Information regarding the route formation and publicly available information about the shippers will then be used by the mechanism to allocate costs to the shippers or to assign consolidated loads to trucks. This decentralized nature of the public consolidation mechanism will mean that shippers will be able to see other shippers joining or leaving their loads and the formation of the route when these activities happen before the load is actually proposed to carriers and their trucks. While locally there could be computational overhead when a shipper initiates a load, the overhead will be for that particular load and will not impact the formation of other loads. Given this context, the effect of transaction cost for running the mechanism locally will be a small percentage of the allocated cost to the shipper and we expect its overall impact on the efficiency of the mechanism to be insignificant.

1.2 Organization of the Thesis

This thesis consists of the following main chapters, besides this introductory chapter.

Chapter 2 describes important components of the shipper side of the mechanism. It develops procedures to form consolidated loads, online route formation, and cost allocation to the shippers when the load is formed. Various cost allocation methods are explored to understand which method(s) are the most economically efficient for the mechanism.

Chapter 3 extends the work done in Chapter 2 to build components for the carrier side of the mechanism. Once consolidated loads are proposed by shippers, procedures are developed to efficiently assign these loads to trucks. Factors that will drive economic efficiency from a carrier perspective are explored especially in the context of allowing trucks the freedom to renege and exchange already assigned loads among themselves.

The work done in both Chapters 1 and 2 does not simulate the mechanism by using agents for the shipments and for the trucks. Instead, the mechanism is evaluated using various approaches like using the softmax function to reflect shippers' acceptance probabilities for a load in Chapter 2 or using the linear assignment as a means to arrive at the lower bound for the total deadhead distance for trucks in Chapter 3.

The analysis in the aforementioned chapters does not consider the time preferences of shippers. It is assumed that when a truck picks up a consolidated load, it will deliver all constituent shipments of the load within a single day or within an eight-hour period. The assumption thus made is that all pickup and delivery time windows of all member shippers in a load are met, the truck gets the full price of the load, and no penalties are incurred by or imposed on the carriers. While the absence of time preferences of shippers forming loads may seem to be a strong assumption, it is not entirely irrelevant if all shipments originate and end within a particular state or are located in two nearby states and a truck has an average speed of 50 miles per hour.

Chapter 4 extends this analysis by considering the time preferences of shippers as a part of the mechanism. Time preferences of shippers will become important because the mechanism will be available to shippers across the country, and will not be just limited to a few states. Time preferences are also important for local delivery of goods in which the truck speeds will be within $25 - 30$ miles per hour. Given the public nature of the mechanism, the problem will be quite different from the traditional vehicle routing problem with time windows (VRPTW). The VRPTW assumes all time windows of shippers are information that is available to the firm providing this service; that will not be the case with a public mechanism where shippers will wish to keep their time preferences private. Which aspects of a shipper's time preferences will be made public for others to see will be an important decision while extending the design of the mechanism. As an initial model, a home delivery logistics network with time preferences has been designed where a certain batch of customers wants to ship items from one distribution center (DC) to another that is close to their homes. While more details and initial findings have been summarized in Chapter 4, a modeling assumption that has been made is that only the lead shipper (the shipper that initiates the load) makes its time window public.

The assumption that the items for the customers will ship directly from one DC to another destination DC is not realistic since there could be DCs in between where the shipments may have to be transferred from one vehicle to another. But one of the other alternative for direct delivery is using drones. Drones will not have the overheads associated with road transport, like traffic rules and refueling, and can be used to transfer items from one DC to another DC without any intervention in between. The assumption made in this home delivery model is that all customers are willing to transmit items as truckload (TL). But drones can also be used to transfer items for customers who want to go LTL. The actual ramifications on pricing and efficiency of using drones can be an interesting research topic, but the work done in this research is general enough that it can be also conducive to using alternate means of transport like drones, and not just carriers and their trucks.

Chapter 5 provides protocols for a public logistics network for home delivery that will help to reduce car trips to the stores by allowing package deliveries to homes in a matter of hours. The design will enable scale economies to be realized in performing each logistics function since each element of the network has access to potentially all of the network's demand. A simple agent based model is developed that uses package and truck protocols to illustrate how loads will be formed in distribution centres, costs allocated, and loads assigned to trucks. The chapter concludes with a discussion that argues that with the development of autonomous vehicles and drone technology a public logistics network can be a viable alternative for home delivery as opposed to services provided by Amazon, USPS, etc.

Chapter 6 provides a summary of the work done in the thesis with a discussion about several extensions that can be made to the models developed in the aforementioned chapters.

1.3 Open Questions

The work done in this thesis will not only develop mechanisms for consolidating loads proposed by shippers, but also develop truck protocols to accept the loads once they have been made final by the shippers participating in the load formation. The mechanism will further be extended to include the time preferences of shippers and their impacts on load formation. The objective while developing the mechanisms, both for load formation and for truck acceptance protocol, will be to achieve economic efficiency.

The following are some of the major questions that work in this thesis will attempt to answer :

1. Allocation of costs to individual shippers forming a consolidation will be a very important aspect of the research in Chapter 2. What are the features of various cost allocation methods? Which method will be economically efficient?

2. How can trucks be assigned consolidated loads in an economically efficient manner? What factor will drive economic efficiency?

3. How will the time preferences of shippers affect load formation? How much of their time preferences should shippers reveal? Should all shippers in a load coalition reveal their time preferences, or only the shipper initiating the load should be one to reveal its time preferences?

CHAPTER

2

SHIPMENT PROTOCOL FOR LOAD FORMATION

2.1 Introduction

The purpose of this chapter is to develop a mechanism that uses publicly available shipper information to consolidate multiple shipments into a single load. Consolidating multiple shipments can lead to significant economies of scale in transportation and can reduce transportation costs, a major portion of total logistics costs. An online route construction procedure that uses two-optimal improvement is used to construct routes as and when shippers join a coalition. Thus, as new shippers join the coalition, the route is reconstructed and the new transport charge is reallocated to each shipper. This allows shippers to independently determine their allocation of the transport charge when they want to be part of more than one consolidated load. In contrast to a private consolidation service that provides consolidation of shipments and keeps a significant portion of the cost savings with itself, the decentralized nature of the mechanism allows shippers to divide almost all of the savings among themselves.

We make some simplified assumptions regarding shipments in order to limit the scope of the mechanism. These assumptions were made to focus on the shipper side of the mechanism. For example, we assume that any shipment posted by a shipper will be delivered from its origin to its destination within an eight-hour period. We assume a transport rate r of \$1 per ton-mile to compute the transport charge and that the actual allocated charge is collected from the shipments at the time of route execution. Shipments interested in forming consolidations are encouraged to post their availability so that other shipments are encouraged to join the load and allow the carrier to plan the transfer of the load in the best way possible. With every new shipment that joins, the total weight and volume of the putative load are compared with the truck trailer weight and volume so that the load is within the trailer weight and volume limits. We assume that there are sufficient numbers of trucks to pick up loads and that a truck is located at the origin of the consolidated route of a shipment coalition. We also assume that shipments are not time-sensitive and that delays in pickup or delivery of shipments do not incur penalties. This chapter discusses the elements of the mechanism needed for load formation prior to a truck accepting a load for transfer.

In order to have a stable consolidation, it is important that the total costs or savings are allocated among shippers in a way that reduces their independent costs so that there is a significant incentive for shippers to form coalitions. Since shippers can be a part of multiple coalitions, it becomes imperative to understand how the costs accrued from consolidation are divided among the shippers. This motivates the use of cost allocation methods, and an important aspect of this research is to compare and contrast the features of multiple models of cost allocation. The consolidation mechanism is designed to apply a cost allocation method to decide which shipper participates in a coalition and how much cost it will incur. Shippers do not have to enter into a binding agreement to coordinate among themselves; instead, any coordination and cooperation among shippers is self-enforced through the mechanism.

An important goal of this research is to adopt a cost allocation method that promotes economic efficiency so that the welfare of shippers participating in consolidations is maximized. Since it is difficult to claim the efficiency of the mechanism by using just one method of allocating costs, in this chapter, we develop and analyze the proportional and the Shapley models of cost allocation. The cost of a single load is just the total length of the consolidated route multiplied by the transport charge. The total cost of all the coalitions that are formed when the mechanism allows shippers to form consolidations in both allocation methods is used as a measure of economic efficiency. The total cost under different allocation schemes is also compared with the cost realized by a private consolidation service that consolidates shipments using an offline procedure, and with the total cost when all shipments travel as independent shipments. Numerical experiments run for this research show that a private consolidation service provides consolidation at the lowest transport cost, followed by the proportional cost allocation and the Shapley cost allocation. Note that since we are not doing a detailed agent-based simulation of the mechanism, we devise a number of methods in this chapter that aid us in evaluating the mechanism (for example, using set covering and the softmax function).

Since shipments will likely join multiple loads, an equally important aspect of this research is to understand which coalitions will be accepted by its shipment members, and which coalitions will be rejected. A softmax distribution was used to model the individual decision-making process of each shipment and to evaluate the mechanism. Each shipment is assigned a probability of acceptance of a consolidated load based on the cost allocated to it in the load. Given that the work in this chapter does not explicitly model the behavior of trucks, we assume that a truck randomly selects a load for transport. The decision problem that a shipper faces is either to accept (or reject) the load if the probability of its accepting the load is higher (lower) than a threshold value. If the shipper rejects the load, it leaves the coalition, and the remaining shippers in the load are then allowed to re-form coalitions among themselves. This process is continued until all coalitions are acceptable to their respective members.

Some of the important assumptions made in this chapter are listed below:

1. A consolidated load, after it is posted for transport, will be delivered within an eight-hour period. This time window is enforced, and shippers do not make their time preferences explicit.

2. Shippers prefer to be in loads in which they receive a higher savings when compared to their independent costs, if they are a part of more than one load.

3. Shippers are not time-sensitive. They do not have personal time preferences for their shipments, nor do they impose any penalty cost on trucks if the pick-up or delivery target is not met.

The research in this chapter will have the following main deliverables:

1. Which cost allocation method, the proportional or the Shapley, is more economically efficient?

2. Which set of consolidation loads are acceptable, and thereby feasible, to all shippers?

2.2 Related Research

It has been accepted by the supply chain community that combining many shipments into large loads can reduce total logistics costs to the shipper, reduce the number of trucks used by the carrier, and be environmentally beneficial. Higginson and Bookbinder (1994) provide a number of recommendations by suggesting three consolidation policies: Time, quantity, and time and quantity. They provide simulation results for the three policies and conclude that the policy to choose will be a managerial decision depending on cost and time savings. Consolidation of shipments that is fair to various stakeholders necessitates the design of good mechanisms, and so a variety of supply chain coordination mechanisms use mechanism design. Key ingredients for good mechanism designs are described in Maskin (2008). Like many supply chain applications, the public consolidation mechanism uses a decentralized information system (Fan et al. 2003). Our research comes under the purview of distributed coordination mechanism design because, while a trusted center will have publicly available information, stakeholders make decisions that compete with each other. Shipment consolidation has also been investigated for global 3PL providers that must make consolidation decisions at their terminals and how they choose cargo flights after consolidation (Tyan et al. 2003). An important issue, elaborated in Cheong et al. (2007), that 3PL firms face is getting efficient performance from their consolidation services based on how they arrange their consolidation flows. Besides addressing hub locations, they also study the impact of inventory holding and replenishment policies for consolidation. Consolidation services can also be extended to perishable items (for example, cut flowers), as shown by Nguyen et al. (2014), especially since severe time constraints exist for such items. The authors use stochastic dynamic programming models to compare consolidation strategies with independent shipment alternatives. Since consolidation involves creating efficient routes, heuristics for solving vehicle routing problems are developed in Chu (2005). Most of the heuristics developed are modifications of the Clark and Wright algorithm and include a fixed-cost component for operating various kinds of trucks. A savings procedure is used by Kay et al. (2022) to construct periodic long-distance trucking routes based on minimizing the sum of transport and cycle inventory costs. This procedure is only suitable for use by a private fleet since it uses the inventory carrying cost of each shipment in its calculation.

Some of the issues involved in this research are described in Kay and Jain (2002) in the context of a public logistics network. An initial design of the protocols needed for a distributed coordination mechanism is presented in Kay (2004). This work is limited to point-to-point transport of unit-size loads. Some of the basic features of the mechanism, especially constructing consolidated shipment loads using publicly available information, are described in Kay and Warsing (2008). Economic effects of freight consolidation have been studied by Kay et al. (2023) by using datasets from cities in Turkey. A study done by Liu et al. (2019) for the U.S. Department of Transportation, using data from a private carrier, found that consolidated LTL shipments reduced the number of trucks and saved money for the carrier. The study did a ten-run simulation for unconsolidated costs for ten medium-sized trucks and found that under consolidation, the number of trucks, the total costs, and the total distance were reduced. So while private consolidation services benefit the carrier, they will accumulate deadhead (i.e., empty travel) costs since their route structures might ignore shippers that are not using their services. The carrier will invariably pass such deadhead costs to the shipper(s) (Bartolini et al. 2013). According to Aronsson and Huge Brodin (2006), a public consolidation mechanism will help reduce deadweight loss and improve consumer surplus as more and more shippers use its services. Note that the deadweight loss to the trucking industry was estimated to be about $32.1 billion as shown in Nehiba (2020). The need for consolidation of item quantities by different agents has also been studied in the broad area of economic product quantity (EPQ) in inventory. Agents facing EPQ shortages can agree to form coalitions to order jointly, giving rise to inventory games. Such games have been studied in detail in Meca et al. (2003, 2004). A fundamental difference between such games and a public consolidation mechanism is that in the former, agents do not need to make their orders public, while in a public consolidation mechanism, this information is publicly available. Further, inventory games allow the computation of the optimal order quantity, so costs are fixed apriori when agents decide to order jointly. This is not the case in public consolidation mechanisms since the most effective route changes with each new shipment, and so do the cost allocations.

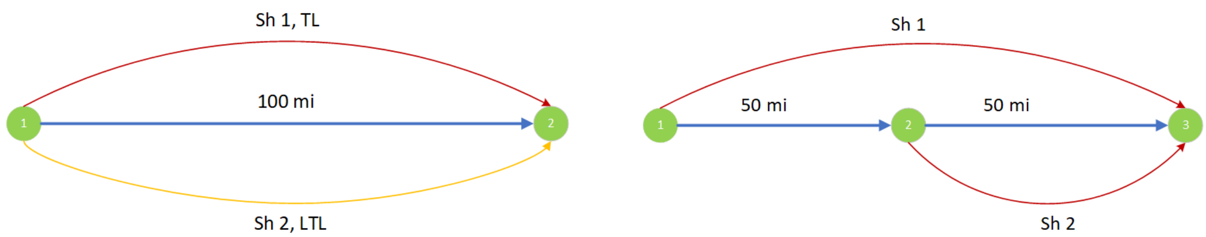

(a) TL and LTL Shipments.

(b) Two TL Shipments.

Figure 2.1: Two Shipment examples.

2.3 Cost Allocation Mechanisms

2.3.1 Shapley Cost Allocation

The Shapley value remunerates each participant (in this research, a shipper) by taking a simple average of all the marginal contributions containing that participant.

The complete Shapley value for shipment i can be defined as follows ((Moulin 1991)):

$$\alpha_i = \sum_{0 \leq m \leq n-1} \frac{m!(n-m-1)!}{n!} \sum_{M \subset N \backslash i, |M|=m} (\sigma_{M \cup \{i\}} - \sigma_M) \tag{2.1}$$

Here, N refers to the unordered set of shipments in the coalition, M denotes the subsets of shipments, σ_M is the sequence-independent cost associated with the subset M of the shipments, and $\sigma_{M \cup \{i\}}$ is the cost of the combined shipments when shipment i joins the coalition. The computation of the exact value for shipment i will necessitate computing marginal costs of shipment i for all $n!$ arrangements, which can become prohibitively time-intensive for higher values of N.

In Figure 2.1a, two shipments travel from origin 1 to destination 2 and cover a distance of 100 miles. Shipment 1 prefers to go truckload (TL) and pays a transport charge of \$1 per mile for an independent cost of \$100. Shipment 2 prefers to go less-than-truckload (LTL) and is willing to pay \$0.50 per mile for a total cost of \$50 to go from point 1 to point 2. In Figure 2.1b, both shipments 1 and 2 are TL, but shipment 1 has to cover a distance of 100 miles with its origin at point 1 and destination at point 3 while shipment 2 has to cover a distance of 50 miles with its origin at point 2 and destination at point 3. Both the shipments pay a transport charge of 1 per mile so shipment 1 is willing to pay \$100 as its independent cost, while shipment 2 is willing to pay \$50.

We observe that in Figure 2.1a the entire 100 miles is shared by the two shipments, while in Figure 2.1b, only a portion of the route, 50 miles, is shared by the two shipments. In both these cases, since routes overlap, the shipments can form a coalition to reduce their total transport cost to \$100 instead of the \$150 that they would pay if they both go as independent shipments. To specify the marginal contributions to calculate the Shapley value, we consider the two shippers to be randomly ordered with every ordering equally likely. That is, shipper 1 can initiate the coalition, and shipper 2 can join it as the second member, or shipper 2 can initiate the coalition, and shipper 1 can be the second member. The marginal costs in either of the cases are shown in Table 2.1. When shipper 1 is the first member of the coalition, its marginal cost is \$100 since it has to bear the entire cost of the route while shipper 2 has a marginal cost of \$0. When shipper 2 is the first member of the coalition, its marginal cost is \$50 for its share of the route so that shipper 1 incurs a marginal cost of \$50. The average of these marginal costs for the shippers is the Shapley value cost allocation, and this is shown in the last row of Table 2.1. Using the Shapley value, when shippers 1 and 2 form a coalition, they pay \$75 and \$25 respectively in both the cases in Figure 2.1.

Table 2.1: Marginal Costs(\$) using Shapley Cost Allocation

Coalition Sequence	MC(1)	MC(2)
1,2	100	0
2,1	50	50
Average	75	25

Table 2.2: Allocated Costs using Proportional, Shapley and Hybrid

Shipments	Proportional	Shapley	Hybrid	Independent
1	66.67	75	72.73	100
2	33.33	25	27.27	50
Total	100	100	100	150

2.3.2 Proportional Cost Allocation

The proportional cost allocation imputes costs in proportion to the total savings (or costs) that shippers receive. More generally, the proportional cost allocation can be written as

$$c_{y_i} \quad = \quad c_L \frac{c'_{y_i}}{\sum_{j=1}^{N'} c'_{y_j}}, \tag{2.2}$$

where c'_{y_i} is the independent cost of shipment y_i, N' is the number of shippers, c_L is the total route cost for the consolidated load and c_{y_i} is the allocated cost to shipper i under the proportional cost allocation scheme. We now look at an example in Figure 2.1.

In both the cases of Figure 2.1, if shippers 1 and 2 decide to come together, they get a total savings of $150 - 100 = 50$. Shippers split this saving in proportion to their independent costs. The cost to shipper 1 is then $100 - (50)(100)/150 = \$66.67$, while for shipper 2 it is $50 - (50)(50)/150 = \$33.33$. Note that this is also equivalent to allocating the combined route cost of \$100 in proportion to their independent costs. That is, for example, for shipper 1 this is $(100)100/150 = \$66.67$ (and similarly for shipper 2). The proportional allocation thus allocates a higher cost to the shipper that has a higher independent cost.

2.3.3 Hybrid Cost Allocation

Table 2.2 summarizes the costs allocated under various allocation schemes. Table 2.2 also shows the costs allocated to shippers 1 and 2 under a hybrid scheme. If C_1^{Sh}, C_1^{Pr} are the costs allocated to shipper 1 under the Shapley and the proportional schemes and C_2^{Sh}, C_2^{Pr} are the corresponding costs to shipper 2, then the Hybrid allocation scheme takes the minimum of the two costs for each shipper, normalized to the actual route cost. That is if $C_1 = \min(C_1^{Sh}, C_1^{Pr})$ and $C_2 = \min(C_2^{Sh}, C_2^{Pr})$, then

$$C_1^{Hy} = (C_1)(TC_R)/(C_1 + C_2), \tag{2.3}$$

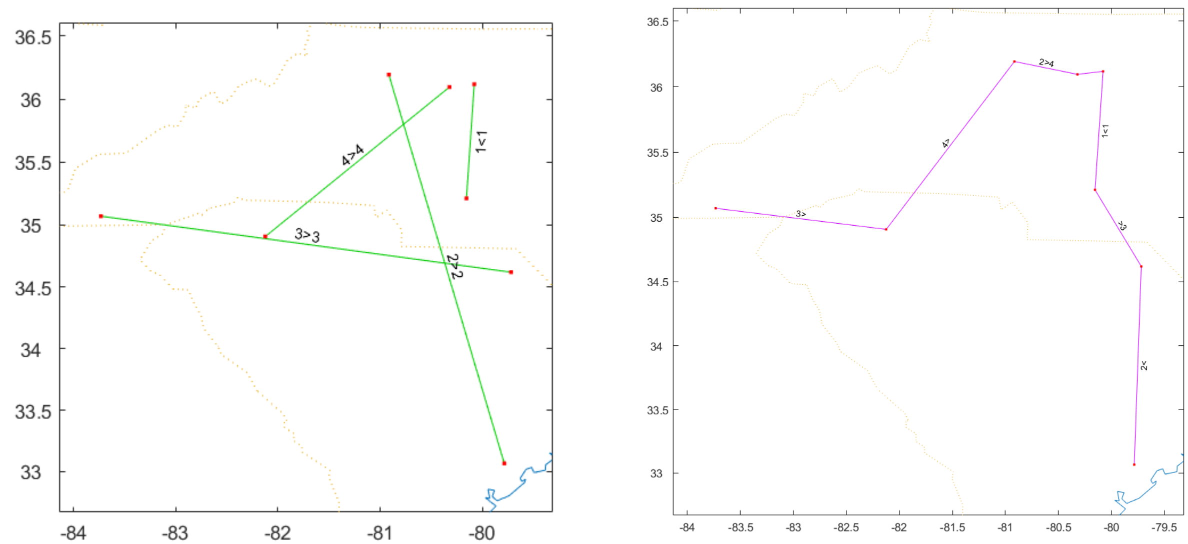

(a) As Independent Shipments.

(b) As Consolidated Shipments.

Figure 2.2: Five Shipment examples.

where TC_R is the total cost of the combined route. For both the cases in Figure 2.1, $TC_R = \$100$, $(C_1, C_2) = (66.67, 25)$ and the normalized costs $(C_1^{Hy}, C_2^{Hy}) = (72.73, 27.27)$.

Even though the two examples in Figure 2.1 have the same cost allocations under the various allocation schemes (Shapley, Proportional, and Hybrid), there is an important difference between the Shapley and the Proportional allocation besides the respective ways in which they arrive at the cost shares. The proportional cost allocation assumes that the entire route of 100 miles is in contention in both the examples in Figure 2.1. It does not distinguish between the fact that in one case, the entire route is shared by both the shippers, while in the other case, only a part of the route is shared by the shippers. The Shapley value does in fact make this distinction; an alternate way of arriving at the Shapley cost allocations in Figure 2.1b is by noting that shipper 1 travels the distance of 50 miles from location 1 to location 2 by itself, while the remaining 50 miles from 2 to location 3 is shared by both the shippers. Since shipper 1 does not share the leg $1 - 2$ with shipper 2, it bears the full cost of that leg of \$50, while the cost of the leg $2 - 3$ is shared equally by both the shippers for a cost of \$25 each. Shipper 1 thus pays $50 + 25 = \$75$ while shipper 2 pays \$25 which are the Shapley values that were already computed.

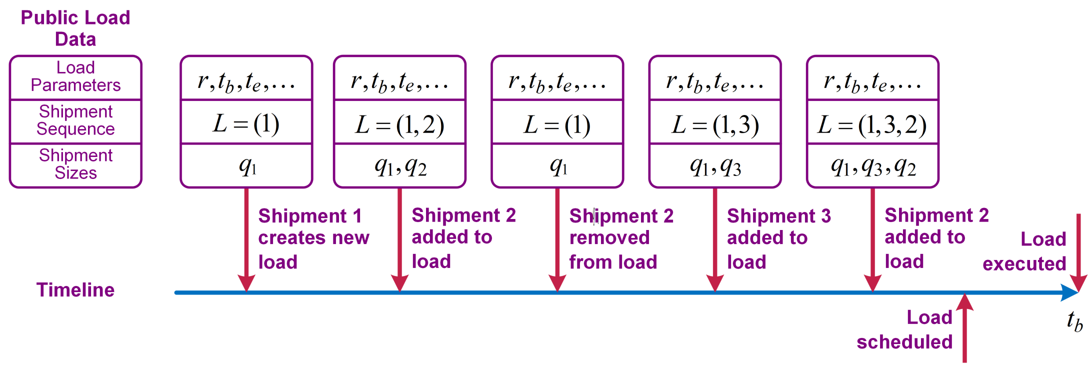

Figure 2.3: Load formation (adapted from Kay and Warsing, A Distributed Coordination Mechanism for Shipment Consolidation).

2.4 Designing the Public Consolidation Mechanism

An example of a public consolidation mechanism used to form consolidated loads by shippers is shown in Figure 2.4. Shipments will post public data that will be used to form consolidated loads. Examples of public data could be the beginning and end location of a shipment, the distance matrix for shipment locations, the shipment densities, and the like. They will also have critical data that they will keep private and use them to decide if they want to join a load or not (for example, their time preferences).

We now describe the important design features of the public consolidation mechanism, which will include a description of route construction procedures and the coalition formation process.

2.4.1 Route Sequencing

The standard procedure to determine the minimum distance route sequence takes the sequence $L = (y_1, ..., y_n)$ of n shipments in a load as input and outputs a $2n$-element route sequence $R = (z_1, ..., z_{2n})$ as shown in 1, representing the sequence in which the loading and unloading of each shipment occur to minimize the total distance travel from the loading of the first shipment to the unloading of the last shipment in the route.

The first occurrence of $z_j = y_i$ in R indicates the loading of shipment y_i, and the second occurrence of $z_k = y_i$ indicates its unloading, where $j < k$. Each element z_i of route sequence R has a corresponding location x_i, $X = (x_1, ..., x_n)$, representing either

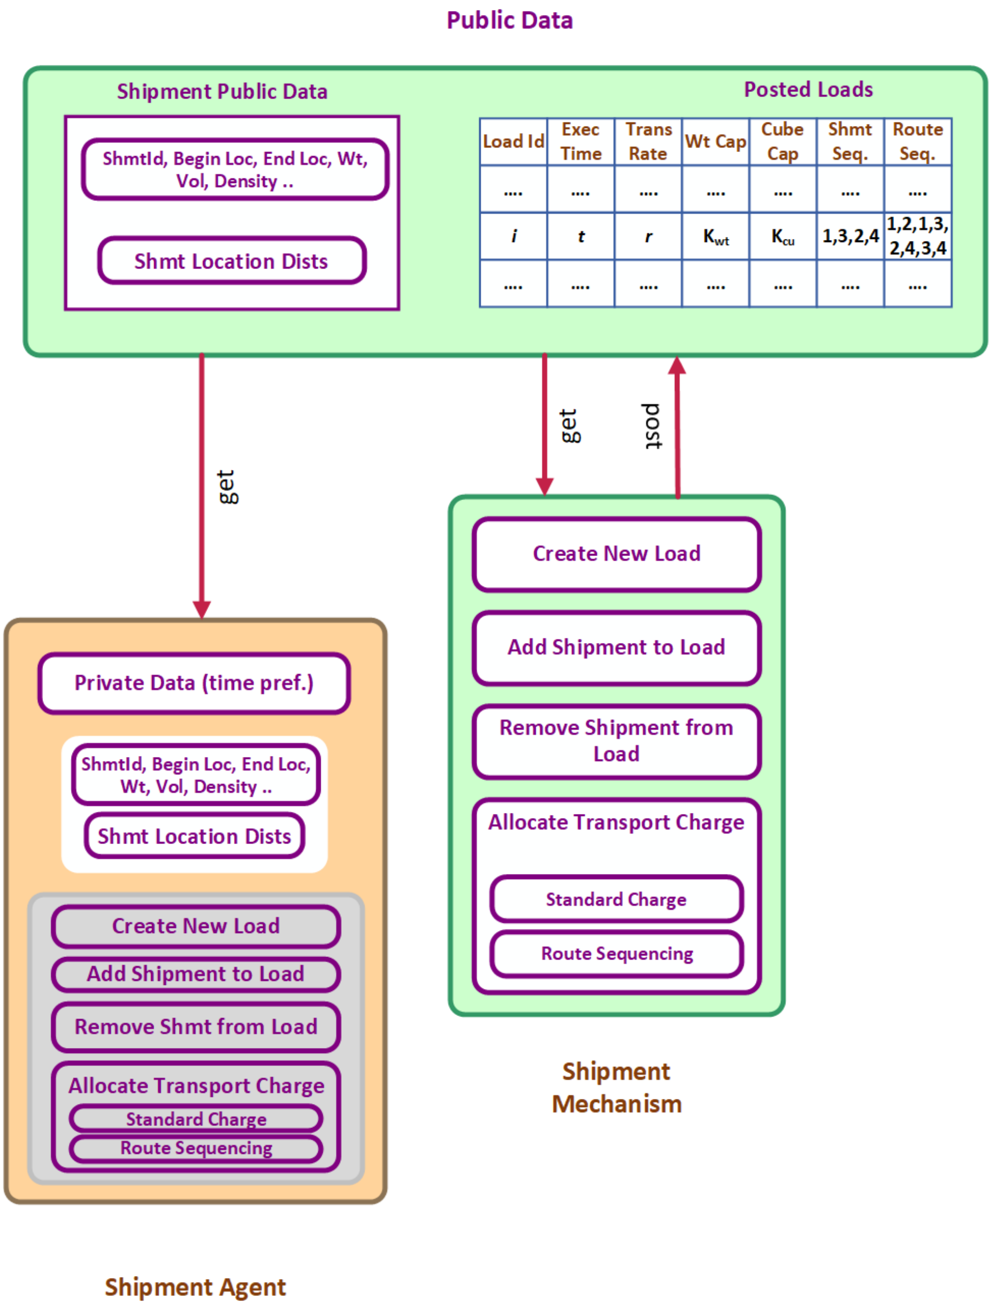

Figure 2.4: Public consolidation mechanism for shippers.

Algorithm 1 sequenceInsert: Construct a Route R

Input $y_i \in R$
Set $R = (y_1, y_1)$
for each $i \in \{2, ..., |L|\}$ **do**
    Set $R = minDistanceInsert(y_i, R)$
    Set $R = twoOptImprove(R)$
end for
Return R

Algorithm 2 minDistanceInsert: Insert Shipment With Minimum Increase in Distance

Input $y, z_i \in R$
Set $d_R = d(R)$
for each $i \in \{1, ..., |R| + 1\}$ **do**
    for each $i \in \{1, ..., |R| + 1\}$ **do**
        Set $R^{'} = (z_1, ..., z_{i-1}, y, z_i, ..., z_{j-1}, y, z_j, ...z_{|R|})$
        if $d(R^{'}) < d_R$ **then**
            $d_R = d(R^{'})$
            $R = R^{'}$
        end if
    end for
end for
Return R

Algorithm 3 twoOptImprove: Pairwise Exchange of Route Segments

Input $z_i \in R$
Repeat (Continue until no segments can be exchanged)
done = *true*, $i = 1, j = 2$
while *done* is true and $i < |R|$ **do**
    while *done* is true and $j < |R| + 1$ **do**
        Set $R^{'} = (z_1, ..., z_{i-1}, reverseSequence(z_i, ..., z_j), z_{j+1}, ...z_{|R|})$
        if $d(R^{'}) < d_R$ **then**
            $d_R = d(R^{'})$
            $R = R^{'}$ (Replace R with improved route)
            done = *false*
        end if
        $j = j + 1$
    end while
    $i = i + 1, j = i + 1$
end while
until *done* = *true*
Return R

the origin or destination location of the shipment. A $2n$-element route R has $2n-1$ segments, each consisting of a consecutive pair of elements in R, (z_i, z_{i+1}). The total distance of R is the sum of the distance of each of its segments. Assuming that the distances d_{ij} between all pairs of locations i and j are publicly known,

$$d(R) = \sum_{i=1}^{2n-1} d_{x_i, x_{i+1}} \tag{2.4}$$

represents the distance of route R. For example, referring to Figure 2.5a, the route sequence $R = (3,1,2,2,1,3)$ would correspond to travel from location 5 after loading shipment 3 to location 1 to load shipment 1, to location 3 to load 2, to 4 to unload 2, to 2 to unload 1, and to 6 to unload 3. The total distance of the route is 430, where $d_{5,1} = 60$, etc.

The route sequencing problem is equivalent to the traveling salesman problem, and simple heuristics can be used to provide reasonably good solutions. To make the route sequencing procedure especially simple and fast, an online procedure is used that takes as input only the current route sequence and the shipment to be added to the sequence. This is opposed to a more time-consuming offline procedure that would use all the shipments to construct the route sequence. While an offline procedure can sometimes produce better results and would make the resulting route independent of the order in which shipments are added to the route, the extra time required is not deemed worth the effort since the goal of the consolidation mechanism is to allow many different potential route sequences to be considered with a reasonable amount of effort. Also, shipments are added to a route to minimize its total distance. This criterion is used, as opposed to minimizing the total transport charge of the route, because the transport charge can be a function of the shipment size and could require that the route be re-determined each time there is a change in size.

Given a load sequence L, the procedure *sequenceRoute* shown in Algorithm 1 is used to construct a route R. Each shipment y in L is added to R using *minDistanceInsert* in Algorithm 2, which determines where the loading and unloading of the shipment can be added with the minimum increase in distance. In *minDistanceInsert*, y is inserted between subsequences of R, where any subsequence $(z_i, ..., z_j)$ where $i > j$ corresponds to the empty sequence $()$; thus, for example, when $i, j = 1$, $(z_1, z_0, y, z_1, z_0, y, z_1, ..., z_{|R|})$ reduces to $(y, y, z_1, ..., z_{|R|})$, a route where the loading and unloading of the new shipment occur before any of the other shipments in the sequence.

After each shipment has been inserted into R, the pairwise exchange subprocedure

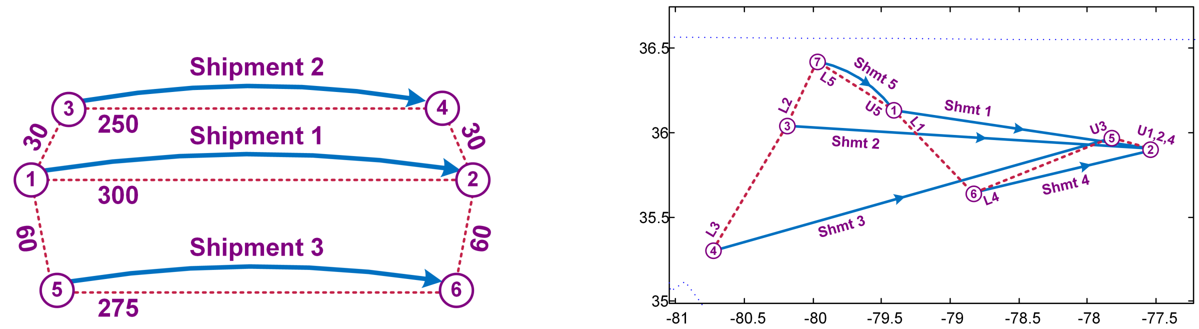

(a) Three-shipment example.

(b) Five-shipment example.

Figure 2.5: Shipment examples (adapted from Kay and Warsing, A Distributed Coordination Mechanism for Shipment Consolidation).

twoOptImprove shown in Algorithm 3 is used to possibly improve the route by replacing nonconsecutive pairs of segments (z_{i-1}, z_i) and (z_j, z_{j+1}) in R with segments (z_{i-1}, z_j) and (z_i, z_{j+1}), and reversing the subsequence $(z_i, ..., z_j)$ until a reduction in distance is found; if so, the R is replaced with the improved route and the process restarts, continuing until no improvement is found for all pairs of segments. The total distance of each route created by the procedure depends on the order of the shipment sequence L and may not correspond to the minimum possible distance.

For example, referring to the three shipments shown in Figure 2.5a and the final load sequence $L = (1, 3, 2)$ formed in 2.3, the result of applying *sequenceRoute* is the route $(2, 1, 3, 3, 1, 2)$ with a total distance of 455, which is greater than the minimum possible distance of 430 corresponding to route $(3, 1, 2, 2, 1, 3)$. The extra effort required to find a better route would place a significant additional computation burden on the shipments and trucks using the mechanism because each might be continuously comparing hundreds of potential loads. In another example shown in Figure 2.5b, five shipments are consolidated. Given the load sequence $(1, 2, 3, 4, 5)$, and the resulting route and location sequences are $(3, 2, 5, 5, 1, 4, 3, 1, 2, 4)$ and $(4, 3, 7, 1, 1, 6, 5, 2, 2, 2)$, respectively.

2.4.2 Consolidated Load Formation

Shippers form consolidated loads by forming coalitions among themselves. The first shipper that wants to invite other shippers to join its load begins by posting a random number seed, along with its origin and destination, the transport charge r and the weight and volume of its shipment. The rate r represents the cost per mile that will be paid by the consolidated load to the truck for its transport, subject to the truck providing a minimum capacity of K_{wt} and K_{cu} for the load. Without loss of generality, this research assumes an r of \$1 per mile. While any alpha-numeric key can be used to uniquely identify a consolidated load, the purpose of using a random number seed is to be able to replicate the approximation that is done for the Shapley value cost allocation using random sampling.

When the random number seed is posted for load formation, other shippers join the load by posting their weight and volume requirements and their origin and destinations. Once a new shipper joins a load, the total weight and volume of the combined load of the existing shippers and this new shipper are checked against the truck's weight and volume requirements. If the combined weight or volume exceeds the truck's limits, the new shipper is removed from the coalition, and the next available shipper is invited to join the load. Assuming that a shipper meets the load weight and volume requirements, the consolidated route is formed using the origins and destinations of all the shippers in the current load using the route sequencing procedures already mentioned, and the total route cost is calculated using the transport charge. This route cost is then allocated among the shippers using an allocation method like the proportional or the Shapley value.

Pareto Improving Cost Allocation

The mechanism is designed so that, irrespective of the allocation method used, each new shipper that joins the load gets a lower cost as compared to its independent transport charge. The newly allocated costs are also lower than the previously allocated costs of the existing shippers. This ensures that shippers get a nonincreasing allocated transport cost whenever a new shipper joins the load. In Table 2.3, when shipper 2 joins shipper 1, a lower transport cost is achieved in both the proportional and the Shapley cost allocation methods, as compared to their independent costs, so shipper 2 stays in the load. When shipper 3 joins the load, the proportional allocation further lowers the previously allocated costs of shippers 1 and 2, while shipper 3 gets an allocated cost that is lower than its independent cost. As a result, the mechanism allows shipper 3 to stay in the load in the proportional cost allocation method, and the next shipper available is invited to join this load.

But using the Shapley value, shipper 3 by joining the load leads to an increment

Table 2.3: Coalitions formed by Proportional but not by Shapley Mechanism

Shipments	Proportional	Shapley	Independent
1, 2	209.7, 108.03	219.25, 98.48	244.37, 123.61
1, 2, 3	207.64, 101.29, 197.51	231.16, 66.32, 208.96	244.37, 123.61, 235.01

Table 2.4: Individual, Pairwise and Three-way Costs($) for Shipments in Table 2.3

Coalition	Cost
1	244.37
2	123.61
3	235.01
1,2	317.73
1,3	491.61
2,3	282.68
1,2,3	506.43

Table 2.5: Marginal Costs($) using Shapley Allocation for Shipments in Table 2.3

Coalition Sequence	MC(1)	MC(2)	MC(3)
1,2,3	244.37	73.36	188.70
1,3,2	244.37	14.83	247.23
2,1,3	194.12	123.61	188.70
2,3,1	223.75	123.61	159.07
3,1,2	256.59	18.83	235.01
3,2,1	223.75	47.67	235.01
Average	231.16	66.32	208.96

in the already allocated cost of shipper 1 from \$219.25 to \$231.16. The reason for this is shown in Tables 2.4 and 2.5. The individual (or independent), pairwise, and three-way costs for each of the shippers in the putative load are shown in Table 2.4. The pairwise cost for $\{1, 3\}$ of \$491.61 is higher than the sum of the independent costs for both shippers 1 and 3. This means when the marginal costs are computed for the Shapley value in Table 2.5, the coalition sequences $\{1, 3, 2\}$ and $\{3, 1, 2\}$ lead to marginal costs for shippers 3 and 1 respectively, that are higher than their independent costs. These higher costs are more than enough to offset the lower marginal costs that shipper 1 achieves in other sequences so that its final allocated cost is higher than its previous allocated cost when shipper 3 had not joined the load. Since shipper 3 leads to an increment in the already existing allocated cost of shipper 1, it is removed from the load, and a new shipper is given the opportunity to join the load. Note that a shipper removed from the load because it was not able to provide a Pareto improving allocation for the existing shippers has the opportunity to join the load at a later period of time, as long as its load weight and volume requirements along with the weight and volume requirements of the existing coalition members meet the truck weight and volume limits.

Approximation of the Shapley value

One of the allocation methods to allocate the total transport charge associated with each consolidated load to each shipment in the load is determined by using the Shapley value (Moulin 1991). An approximation is used because determining the complete Shapley value would involve determining the expected costs associated with all of the possible subsets of shipments, which requires exponential complexity, and, using the route sequencing procedure as specified in Algorithm 1, the value of the total allocated costs of each subset of shipments is sequence dependent. Instead, only a subset of the total number of sequences associated with the number of shipments is used to approximate the Shapley value.

Since the determination of the actual Shapley cost values will involve the computation of expected savings associated with all possible subsets of the coalition, this computation quickly becomes exponential in the number of shippers as more shippers join the coalition. As an alternative, an approximate Shapley value is used that only looks at a subset of the permutations of the number of shippers forming the coalition (Fatima et al. 2008; van Campen et al. 2018). Stratified random sampling is used to pick the number of subsets of permutations. Part of the information posted for each load in the mechanism is the random number seed used to generate the permutations used in the allocation. Since in $N!$ permutations of N, there will be exactly $(N-1)!$ permutations that begin with 1, $(N-1)!$ permutations that begin with 2, and so on up to N, ρ subsets are randomly picked from each batch of permutations that begin with 1, 2 and so on, for a total number of permutations $N' = \rho N$. For each permutation π_j, $j = 1, 2, .., N'$, the marginal cost of the ith shipment is the difference between the cost of the subsequences $a_1 \cdots a_i$ and $a_1 \cdots a_{i-1}$, where a_1 can range from $1, 2, .., N$ and a_i is the position of shipper i in the entire permutation π_j. A matrix with the marginal costs computed for each shipper i, for each sample permutation in the N' permutations is created, and the columns are then summed and averaged over N' to get the approximated Shapley costs for each shipment which is

$$c'_{y_i} = \frac{\sum_{j=1}^{N'} \left[c_j(a_1, \ldots, a_i) - c_j(a_1, \ldots, a_{i-1}) \right]}{N'} \tag{2.5}$$

$$c_{y_i} = c_L \frac{c'_{y_i}}{\sum_{j=1}^{N'} c'_{y_j}}, \tag{2.6}$$

where $c_j(a_1, \ldots, a_i)$ is the cost of the subsequence $a_1, \ldots, a_i$ in the permutation π_j.

Once the approximated Shapley costs have been determined for each shipment, the total cost of the route is just the sum of the Shapley allocated costs of all shipments, $c_L = \sum_{i=1}^{N} c_{y_i}$. The savings can then be computed as

$$c_L^{sav} = \sum_{y_i \in L} c_{y_i}^0 - c_L, \tag{2.7}$$

associated with the consolidation where $c_{y_i}^0$ is the independent transport charge associated with shipment y_i.

Selecting Coalitions

This research develops the shipper side of the mechanism in the absence of carriers and their trucks. In this section, we describe a method by which coalitions formed can be accepted by shippers. A shipper initiates a coalition by posting one or more random number seeds along with the origin and destination of its shipment. A new shipper wishing to join a consolidation sees a number of seeds posted, each seed associated with a putative load. Shippers are not restricted from joining a coalition; they attempt to join all existing loads. Due to the load and volume restriction and the Pareto improving behavioral assumption, they will be able to join some or all of the existing loads. Given a pool of shippers, each shipper initiates a load, and every other shipper is given the opportunity to join these loads. When all possible coalitions are formed, a shipper may find itself in a number of loads with varying allocated costs. It is also possible that some shippers are not able to join any consolidated load since their shipment has been rejected by these loads, and so such shippers have to transmit their shipments as independent shipments. It is assumed that a shipper will prefer to be in a coalition in which it receives a higher saving as compared to its independent costs. The situation for shippers after all coalitions have been formed can be something that is shown in Table 2.6.

Under the coalition column of the table, we see that shippers 43 and 876 are a part of both coalitions. A vector consisting of cost allocations along with the independent cost is formed for each shipper. We assume that the shipper will attribute a positive probability of acceptance to each allocated cost value, and since the lower the cost, the higher the savings, and the higher its acceptance probability, we attribute a softmax distribution to the cost vector after normalizing the costs according to their mean and standard deviation. With this process, each allocated cost for a shipper gets an acceptance probability, which is then remapped to the original coalition of which the shipper was a part. Column "Acceptance Probabilities" shows that shipper 43 has an acceptance probability of 0.89 for the first coalition and 0.93 for the second coalition. The research assumes a threshold probability of 0.8 for accepting a coalition. That is, a shipper will accept a coalition if its probability of allocated cost acceptance is 0.8 or higher. Once probabilities have been attributed to all coalitions, coalitions in which every shipper has a threshold of 0.8 or higher are considered to be acceptable to all coalition members, and a disjoint set of such coalitions is randomly picked as acceptable coalitions. Any shippers not members of the selected coalitions are again allowed to form coalitions, and the above process of coalition formation, probability attribution, and acceptance is followed until no more coalitions are possible to form. Note that if a shipper is not a member of any coalition and decides to go as an independent shipment, its acceptance probability is considered to be one. Based on this analysis, both the coalitions in Table 2.6 are acceptable to all shippers as the acceptance probabilities are higher than the threshold in both rows, and so both are candidate coalitions to be selected, but since the coalitions are not disjoint, only one of the coalitions will be randomly selected; remaining shippers, if not a part of any selected coalitions, will be allowed to reform coalitions and go through the process of acceptance. Note that the calculated probability of acceptance is only used to illustrate the behavior of shipment agents and is not part of the mechanism itself.

Table 2.6: Acceptance Probabilities of Coalitions

Coalition	Allocated Costs	Acceptance Probabilities
353, 43, 876	8.99, 97.44, 180.43	0.8, 0.89, 0.82
876, 43, 448, 504	137.73, 79.42, 53.85, 184.83	0.9, 0.93, 0.93, 0.94

2.5 Computational Experiments and Analysis

In this section, the various allocation methods in the mechanism introduced before are tested using North Carolina (NC) and South Carolina (SC) shipment data sets. Experiments were conducted with shipment data generated using reallocation and population density information from North and South Carolina ZIP codes using the methods mentioned in Kay and Warsing (2009) and using the Matlog toolbox (Kay 2023). Since the total transport cost using the allocation methods was also compared with the total transport cost in a centralized mechanism, the Matlog toolbox was also used for the experimental studies. Assuming that the availability of shipments represents a thick market, a total of 1000 shipments were generated in NC and SC. This allows more shipments to be candidates for consolidation; a small data set with geographically dispersed shipments will make it difficult for the shipments to be consolidated both in the centralized and the decentralized mechanisms.

Since it is not obvious which allocation method will give a lower total transport cost for the mechanism, as a first step, using shipment sizes of 10, lower bounds for the mechanism were generated for the proportional and the Shapley allocation methods using a set covering approach. The set covering method is a way of quantifying the least transport cost that the mechanism can provide under a particular cost allocation method. It is not necessary that the coalitions selected by the set covering will be the coalitions that the shippers themselves will select when they are a part of multiple coalitions.

Given a batch of ten shipments, all possible coalitions were generated using all the non-empty subsets of the batch. For a shipment size of 10, the number of non-empty subsets is 1023, which includes the ten shipments as independent shipments. Each subset was then passed to the mechanism which decided if the subset was a valid coalition using the Pareto-improving method described above. If the subset was a valid coalition, it was included in the list of candidate sets to be considered for set covering; otherwise, the subset was discarded. Note that since subsets overlap and it

Table 2.7: Transport Costs($) for Three 10 Shipment Batches using Set covering

	Shapley($)	Hybrid($)	Proportional($)	Independent($)
1	1170.78	1172.80	1159.59	1875.70
2	1133.89	1092.27	1034.11	1376.00
3	1280.88	1264.83	1235.40	1753.00
Average	**1195.18**	**1176.63**	**1143.03**	**1668.23**

Table 2.8: Transport Costs($) for Three 50 Shipment Batches using Set Covering

	Shapley($)	Hybrid($)	Proportional($)
1	4561.39	4483.48	4607.58
2	4780.54	4612.28	4512.66
3	5412.26	4774.93	4582.06
Average	**4918.06**	**4623.56**	**4567.43**

is not necessary that all members of a subset will form a coalition, it was possible for coalitions of a smaller size to be formed from the subset passed to the mechanism. Thus, any duplicate coalitions formed were also removed. Set covering was then used to get the least cost subsets that cover the entire set of 10 shipments. The total cost under various allocation schemes of three samples of 10 shipment batches with this approach is summarized in Table 2.7. The same experiment done on three samples of 50 shipment batches is shown in Table 2.8. The average of these costs for each allocation scheme, as well as when the shipments travel independently, is also shown in the last row. The table reveals that the average lower bound for the proportional cost allocation scheme is lower than the Shapley allocation. The set covering method thus gives us a metric to compare various allocation methods when the mechanism does the best it can to lower overall transport cost. Since the coalitions selected by set covering may not be the actual allocations that shippers will accept, set covering is not part of the mechanism.

The second case study that was done was to consider shipments at the firm level. To make this analysis, we consider two edge cases. At one end of the spectrum is the case where all shipments are owned by one single firm; at the other end is the case where all shipments are owned by distinct firms (Figure 2.6). In the first case, all shipments have a common origin but distinct destinations. (The case where all origins are distinct, but the shipments have a single destination is symmetrical.) In the second edge case,

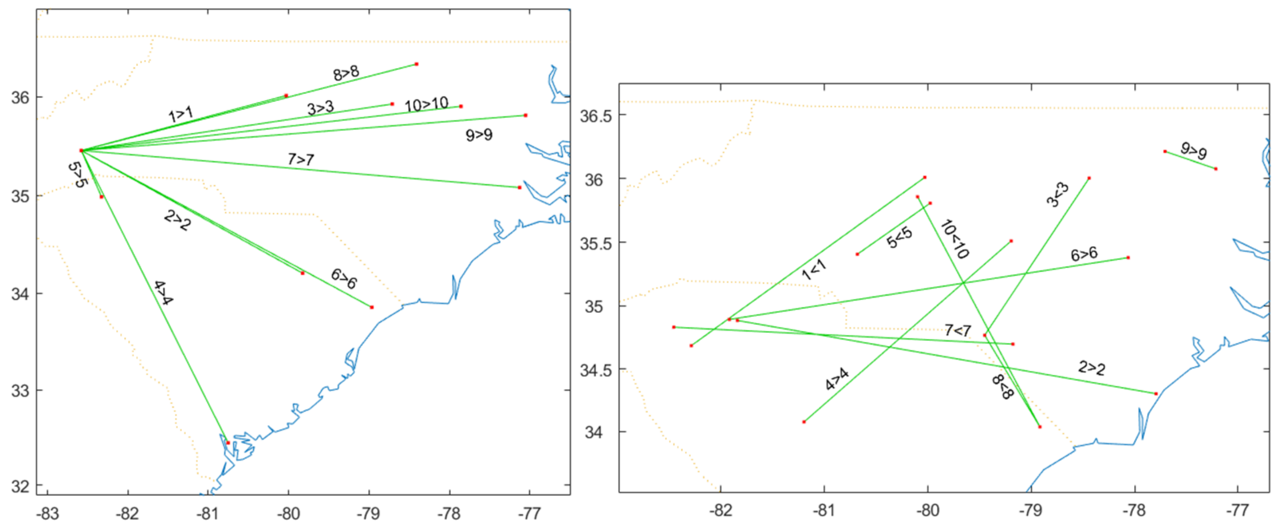

(a) Shipments with a single origin.

(b) Shipments with multiple ODs.

Figure 2.6: Single versus Multiple Firms.

Table 2.9: Cost Comparison for Single Firm versus multiple Firms

		Centralized		Shapley		Proportional	
	Independent	WoS	WS	NPI	PI	NPI	PI
Single Firm	2608.6	899.9	1137.7	899.9	1113.58	899.9	1036.18
Distinct Firms	1861.43	1261.8	1284.62	1387.02	1353.19	1360.29	1269.17

all shipments have distinct origins and destinations. Table 2.9 and Figures 2.7 and 2.8 summarize the results of this test. The costs for each cell in Table 2.9 were arrived at by taking the average of the total transport cost that the mechanism generated for five samples of ten shipment sizes. The same sets of shipments were used for the various mechanisms for a particular case. The total independent costs in both cases are higher than in all other schemes. The centralized mechanism runs an offline procedure to calculate the total transport costs for the shipments. The charges under the centralized are when the process computes costs using a savings function (WS) versus when it does not use the savings function (WoS). The centralized uses the savings function to sort pairs of shipments with savings in descending order (Kay et al. 2022). The higher cost when the centralized runs the savings procedure is due to the fact that some shipments do not give pairwise savings with any of the other shipments, and so they travel independently, which raises the total cost for the centralized. When the

(a) NPI: Non Pareto Improving Allocation.

(b) PI: Pareto Improving Allocation.

Figure 2.7: Shipments With Single Origin (WoS: Without Savings, WS: With Savings).

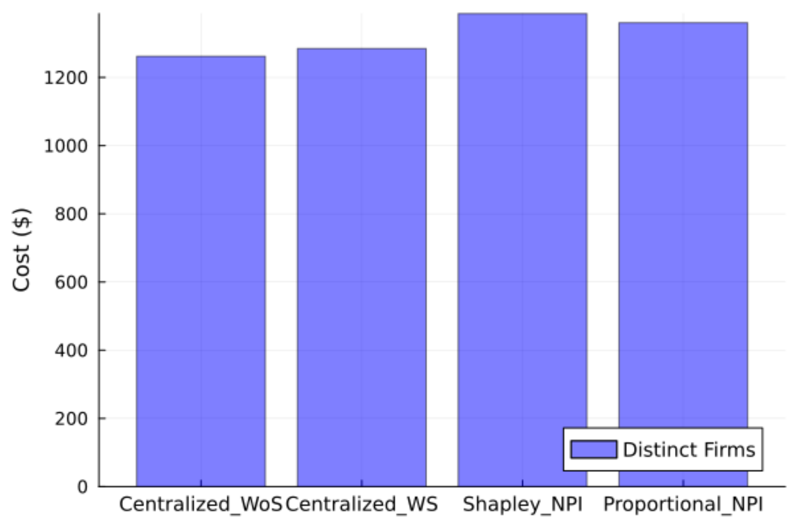

(a) NPI: Non Pareto Improving Allocation.

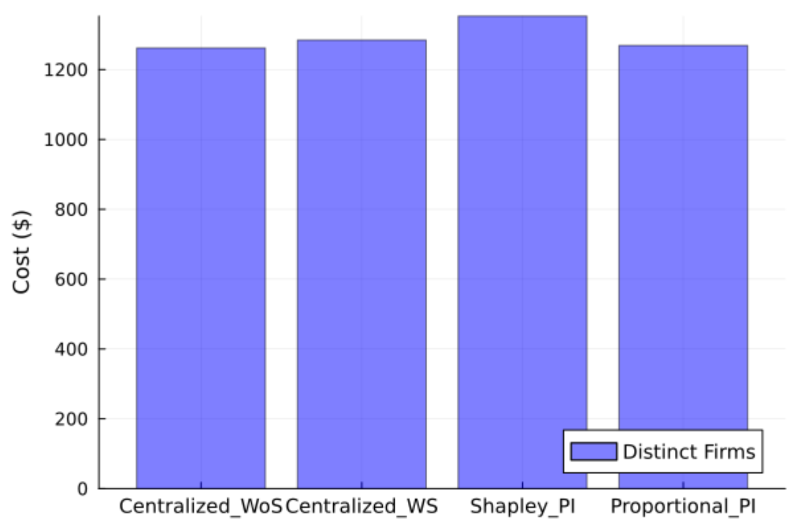

(b) PI: Pareto Improving Allocation.

Figure 2.8: Shipments With Distinct Origins and Destinations.

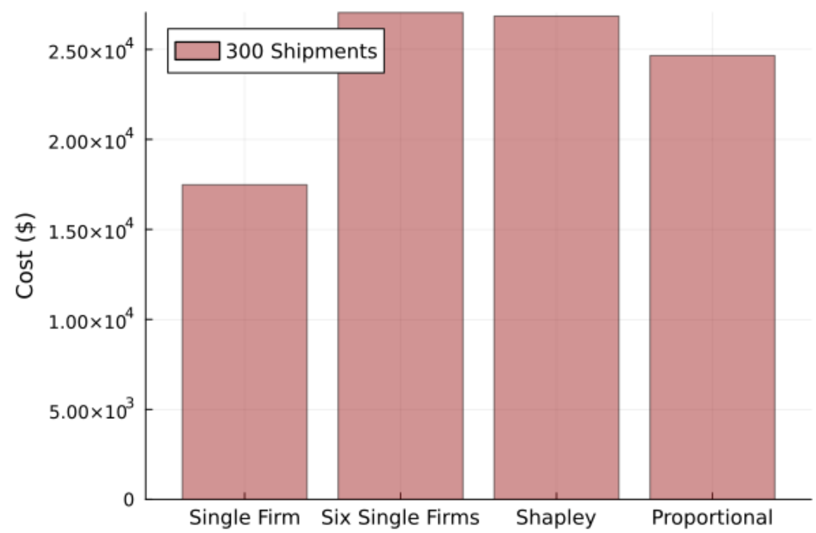

Figure 2.9: Consolidation of 300 Shipments.

centralized computes total transport charge without using the savings procedure, this is similar to the mechanism running without Pareto improving cost allocations. For the single firm, the transport charges are the same when the centralized runs without the savings function and the mechanism is non Pareto improving. For shipments with distinct firms, the centralized mechanism without the savings procedure gives a lower cost compared to the non-Pareto improving mechanism under any allocation scheme. Overall the centralized has the lower transportation cost in both the single firm case as well as the multiple firm case. Note that the firm running a centralized mechanism will run whichever process, with or without savings, that gives it the lower cost. In both cases, the proportional allocation has a lower transport charge compared to the Shapley cost allocation.

Table 2.10: Total Cost($): Consolidation of 300 shipments

Independent	Single Firm	Six Firms(50 shmts each)	Shapley	Proportional
49347.49	17480.12	27046.82	26849.58	24653.08

A salient benefit of the public consolidation mechanism designed in this work is its capability to use the thickness of the shipment market. Unlike a centralized consolidation service that will have access to certain segments of the shipment market, the public consolidation mechanism will have a much wider geographical reach. In order to illustrate this advantage, we considered a sample of 300 shipments from our pool of 1000 NC and SC shipments. The total cost with a centralized service with savings, and with the public consolidation mechanism under the various allocation schemes using Pareto improvement are shown in Table 2.10. The entry in the table "Single Firm" implies that a centralized service is being run by a single firm that has access to all 300 shipments at the same time. While this is highly unlikely, it serves to highlight the fact that the single firm trying to consolidate 300 shipments does much better in terms of total transport cost. The more likely case is where firms will have access to a lower number of shipments at any time. The heading "Six Firms(50 shmts each)" implies that there are six firms with access to 50 shipments each for a total of 300 shipments. Each firm now runs its own individual consolidation mechanism on its 50 shipments. The cost of $27046.82 is the total transport cost after each of these firms has run its centralized consolidation service. Figure 2.9 illustrates that the single firm consolidating all 300 shipments has the least cost while a centralized mechanism being used by six firms gives a much higher cost as compared to the public consolidation mechanism (using either Shapley or proportional cost allocation). Here too, the proportional cost allocation mechanism gives a lower total transport cost compared to the Shapley cost allocation method.

2.6 Conclusion

In this chapter, a public consolidation mechanism for shippers was proposed that uses only publicly available data of shippers to allow shippers to form consolidated loads. For the purpose of experiments and analysis, the only publicly available data used were the pickup and drop points of the shipments to construct routes. One of the main features of the mechanism is that it allows shippers to take control of the load formation. Each shipper can start its own consolidated load or attempt to join any existing loads. The mechanism is designed so that shippers join a load if they see nonincreasing costs (that is Pareto improving) allocated to them as new shippers join a load. The main purpose of this chapter was to investigate various cost allocation schemes for consolidated loads and to choose the most efficient scheme. Experiments conducted with the Shapley and the proportional cost allocations on various shipment batch sizes revealed that the proportional allocation is more efficient than the Shapley with respect to the total cost of the mechanism. An exception was found in only one case of a fifty-shipment batch where the total cost under the Shapley allocation mechanism ($4917.08) was lower than that with the proportional cost allocation ($5161.83). This comparison between the two schemes also introduces the efficiency versus equity argument. The Shapley is more equitable than the proportional since shippers share costs if their routes overlap. In the Shapley allocation, a shipper bears the full cost of the route segment if its load is the only one to use that segment. The proportional allocation, on the other hand, depends only on the initial share of the total independent cost for shippers. After the consolidated route is formed, it imputes costs to the coalition members based on their proportional shares; the fact that some segments of the routes are shared my multiple shippers is not considered by the proportional cost allocation.

As shown in Table 2.3 and Table 2.5, coalitions formed under the proportional allocation scheme may not be valid under the Shapley scheme since Shapley uses pairwise, three-way and so on, to arrive at the marginal costs and these interim costs could be very high (as compared to when a new shipper did not join the existing load) for some shipments so that they will see their overall costs going up when a new shipper joins the load. As a result, the proportional cost allocation, on average, was able to form larger coalition sizes compared to the Shapley cost allocation scheme.

The computational effort required for the two allocation schemes also differs. The proportional is relatively simpler to implement compared to the Shapley and does not require any approximation schemes. It is well known that the Shapley allocation does not scale well as the number of members in the coalition rises. Since it computes marginal contributions of the members over all possible permutations, even for a dozen members in a coalition, the number of permutations becomes huge (it is about half a billion). This inevitably necessitates the use of approximation methods like Monte Carlo to keep the computation time under control.

The consolidation mechanism was also compared with a centralized firm that provides consolidation services on its own. The difference with the centralized firm is that shippers forming their own consolidation groups know which shippers are a part of their group and they are aware of the online route updates that occur when a new shipper joins them. (A centralized firm is more likely to perform offline route construction.) For small shipment sizes, a firm using a centralized mechanism for consolidation does better than the public consolidation mechanism, irrespective of the allocation method used. However, the important observation to be made is that the public consolidation mechanism has a much wider market reach than a firm providing centralized consolidation services. The contrast in total cost is revealed when the shipment market is increased to 300 shipments. While a single firm does do better, it is highly unlikely that it will have access to all 300 shipments. It is more likely that similar firms at different locations will have access to subsets of these shipments. In other words, as shown in Figure 2.9 and in Table 2.10, the mechanism proposed here will be able to better utilize the thicker nature of the market since it is implicitly designed without any geographic boundaries.

Since the overall goal of the thesis is to design an efficient market where consolidated loads proposed by shippers are allocated to trucks provided by carriers, work for this chapter ends with the selection of consolidated loads from the set of loads formed by shippers. A shipper may want to be a part of multiple loads, but disjoint sets of loads were selected so that each shipper was a part of one load. In order to evaluate the mechanism, the softmax function with a predefined threshold was used to select disjoint consolidations. These loads will then be proposed to trucks and based on efficiency criteria defined therein, loads will be allocated to the trucks.

CHAPTER

3

TRUCK PROTOCOL FOR LOAD ACCEPTANCE

3.1 Introduction

As mentioned in Chapter 1, the logistics market has inherent market inefficiencies since shippers have the more critical information about their loads, but they have to reveal this information to a private firm that provides transportation. The firm then uses this information to consolidate loads in a way that benefits the firm more than the shippers. The inefficiency stems from the fact that the market is not in favor of the counterparty (the shippers) with the more relevant information but favors the counterparty (carriers with trucks) whose services are highly generic. Chapter 2 described how consolidated loads can be formed by shippers forming coalitions among themselves, without having to reveal valuable private shipment information. It ended with a method to choose among consolidated loads that are valued the highest by its constituent members.

In this chapter, the mechanism is further extended to include carriers (trucks) that will accept the loads based on certain protocols that will be enforced by the mechanism. Since separate firms can own each consolidated load and truck, the coordination of package transport is more difficult than in a private logistics network where a single firm can provide centralized control of the entire network. The goal of the decentralized coordination mechanism is to make it possible for shippers to provide a consolidated load to minimize their cost of transport and for trucks to maximize their revenue based on the consolidated loads the trucks are assigned. The objective is to assign loads to trucks so as to maximize truck profits. By maximizing truck profits, the mechanism will also minimize deadhead distance for the trucks since any price posted by a load will have to cover a truck's deadhead distance to pick up that load along with the route distance for the load. The alternative, to minimize deadhead distance as a way of assigning trucks to consolidated loads, is not viable from the carrier's perspective since the carrier would prefer to maximize its profits given a deadhead distance for its truck. For example, if two trucks are at the same location and they see an opportunity to swap their loads if it increases their profits, then they would prefer to renege on their current loads irrespective of the impact on the deadhead distance.

Some important outcomes of the aforementioned market inefficiency are the formation of routes that favor the carriers more than the shippers and the rise in deadhead distances for trucks since they have to complete less-than-optimal routes to pick and drop the components of their consolidated loads. As reiterated earlier, since the primary objective of the mechanism proposed in this thesis is to make this market more efficient, this efficiency must be achieved by allowing shippers to keep critical information private and reveal only those pieces of information about their loads that are important for consolidation and at the same time keeping the functions of trucks simple and generic to the extent that they accept the loads that are proposed to them in a way that minimizes the total deadhead distance. The mechanism is designed so that while a truck will be able to see all available loads, the only information about a consolidated load that a truck sees are those that are made public by the shippers in the load; an important part of this information is the route that the truck will follow to complete the load (route construction was described in Chapter 2). Similarly, any load that enters the market sees all available trucks, along with their locations, that are willing to pick up loads.

With this context in place, the proposed mechanism (Figure 3.1) is now extended to design methods that will aid trucks in accepting loads. Consolidated loads available

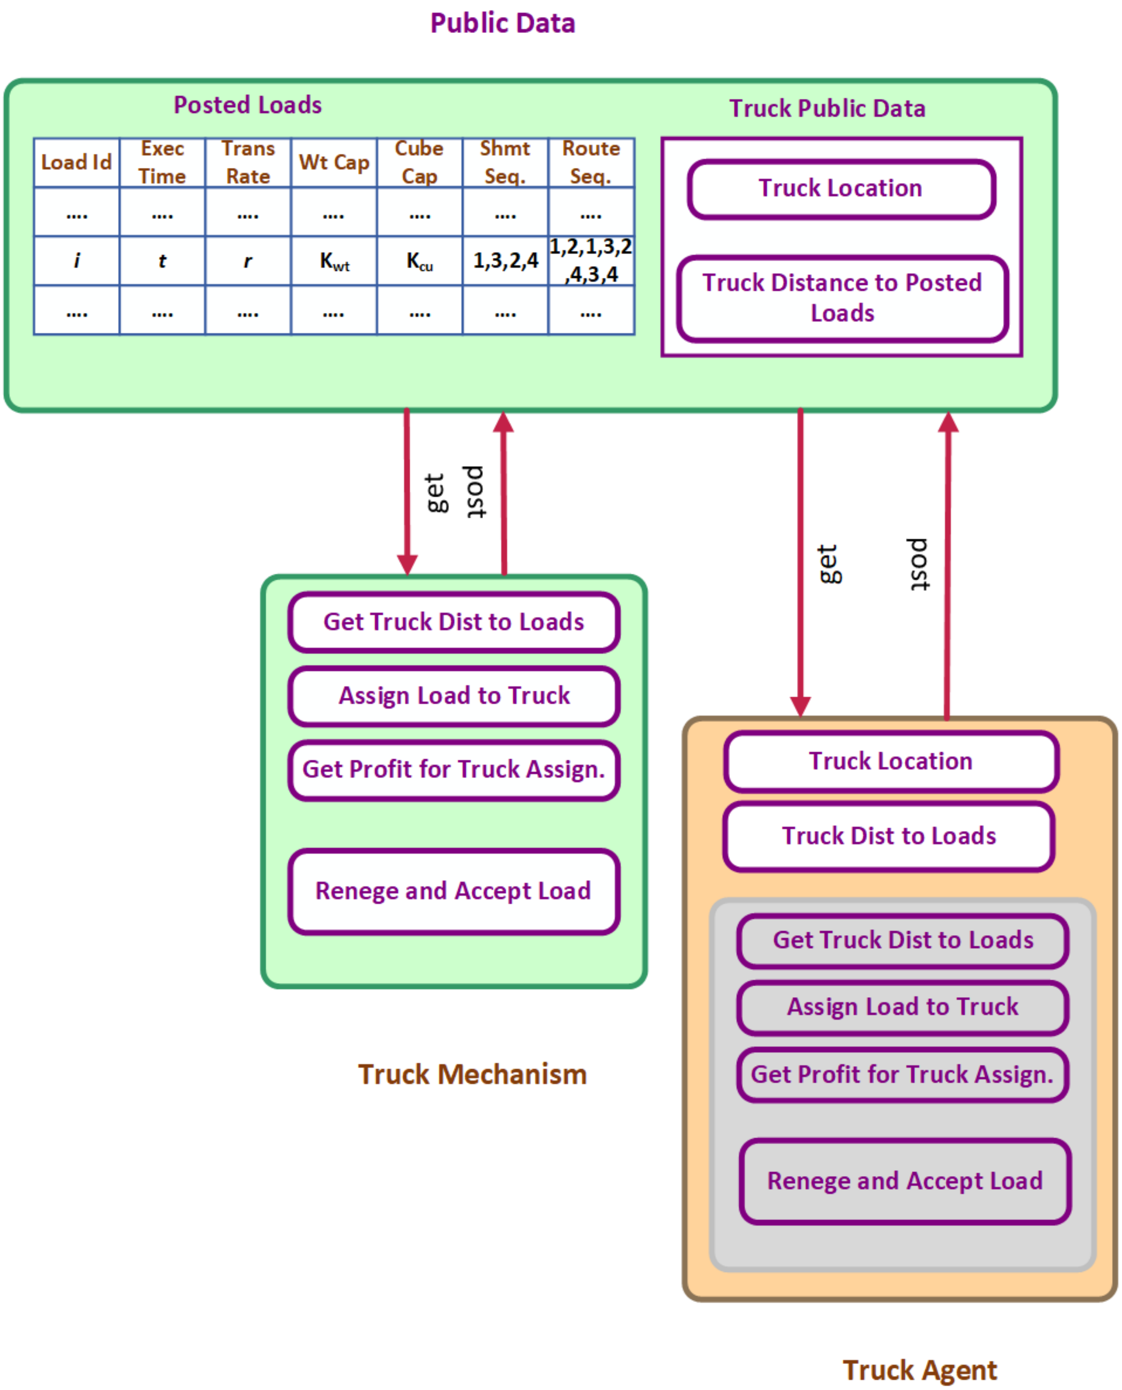

Figure 3.1: Public consolidation mechanism for carriers.

in the market will have to propose a price that is above their reservation price. The reservation price for a load will be the total cost of the route. Since a truck has to cover a certain amount of distance to reach the pickup point of the load (the deadhead distance), a truck would expect that any price proposed by a load must be more than enough to cover both the deadhead distance as well as the route distance.

After loads propose prices to trucks, various methods are explored in this chapter to assign loads to trucks. These include linear assignment, greedy assignment, and stable matching. The linear assignment is expected to be the preferred method that will be used by a private firm providing consolidation and having trucks at its disposal for the shipment of loads. On the other hand, the stable matching method is used to demonstrate that a Pareto-optimal solution to the consolidated load-truck matching problem exists where the two sides to be matched are deadhead distances of trucks to load origins and the truck profits proposed by each load to each available truck. Since the mechanism neither reflects the behavior of a private firm nor aims for Pareto-optimality, it uses greedy assignment to assign loads to trucks. As mentioned in Chapter 2, we are not performing a (truck) agent based simulation of the mechanism but use methods like linear assignment and stable matching approaches to evaluate the mechanism.

Some of the important assumptions that we make in this chapter include:

1. All available trucks and loads can mutually see each others' locations.

2. A consolidated load posts a price that is profitable to all available trucks.

3. Shippers do not have any time preferences for their loads. Any time windows within which loads are delivered are enforced by the mechanism.

4. Trucks pick up and deliver all component shipments of consolidation on time.

5. Trucks with assigned loads are willing to renege and swap their loads with other trucks if reneging improves their profit.

This chapter answers the following research questions:

1. What factor determines economic efficiency when trucks are assigned consolidated loads?

2. What is the (economic) impact of trucks reneging on their loads and swapping loads among themselves? How can reneging be included in the mechanism?

3.2 Related Research

The literature on greedy assignment techniques and the applications of stable matching is vast and extensive. A small subset of recent papers in these fields, as applied to supply chain and logistics, will be reviewed here.

Many heuristics have been developed to implement the greedy assignment problem. Hochbaum and Pathria (1998) study the NP-hard problem of covering a maximum set of elements with a fixed number of subsets and utilizing a greedy-like approximation algorithm for such problems. From a packing perspective, their k-stage approach packs a single bin at a time using subsets not already picked. In logistics, this can be considered to be the packing of k identical vehicles with a set of items with maximum weight to be delivered to a common location. Since capacity planning becomes more challenging due to rising demands in all sectors, Han and Cheng (2020) study crowd-sourced logistics as an alternative in which delivery tasks are completed by part-time vehicle owners. Beginning with an LP model, their heuristic iteratively constructs a route in a greedy manner based on the fractional flows provided by the LP model until the route forms a cycle or some constraints are violated.

Ozdamar and Yi (2008) build a greedy l-neighborhood technique that selects partial paths to append to vehicles' itineraries. The rationale behind this method is that vehicle assignment to locations in emergency logistics is different from normal conditions. A greedy search-based method based on a genetic algorithm is proposed by Chang et al. (2014) as opposed to a mixed-integer formulation to implement solutions for emergency scheduling problems. Several feasible solutions are given by the method instead of one optimal solution so that planners have a set of efficient routing schedules to enable emergency relief. Mishra et al. (2017) apply round-robin greedy search based on distance matrix and demand status to find the nearest demand region from a supply point. The round-robin method places each demand region in a queue so that each region gets a chance to be served by a supply point.

The problem of assigning the members of two disjoint sets to one another was first studied, and an algorithm was proposed by Gale and Shapley (1962). The individual preferences of one set for the members of the other set can be represented as weak orderings (for example, such orderings exist in the processing of college admissions and kidney transplant assignments). Sotomayor (1999) showed that there always exists a pairwise stable matching in many-to-many matching. Few studies have been done to use the stable matching method in supply chain and logistics.

Wei et al. (2019) apply the stable matching method to the logistics selection problem. They set an objective to maximize the satisfaction of shippers and clients based on the condition of stable matching between logistics and clients. Mofidi and Pazour (2019) study how platforms can coordinate decentralized resources to fulfill demand requests using personalized recommendations. They propose a hierarchical approach to create recommendations for shippers and contrast their method with centralized, decentralized, and many-to-many stable matching approaches.

A study of related literature reveals that the benefits and losses associated with reneging on an assignment and accepting another assignment have not been studied well. One of the earliest studies on the impact of reneging was done by Niederle and Roth (2009). They consider college admission markets that can be open or exploding (that is, assignment is binding). Applicants can renege on their offers if the offer is open, at a cost of one point. Note that in this case, firms may not be keen to make early offers since such offers can be used to renege and get better offers. Kay (2004) introduces basic ideas of reneging by trucks when they accept a load to transfer from one DC to another. For example, when a truck reneges on a load, it can accept that load again after all other trucks have rejected the load. Further, two trucks can agree with each other to simultaneously renege and accept the other's loads.

3.3 Assigning Loads to Trucks

The locations of trucks and loads can be as shown in Figure 3.2. The location of loads is assumed to be the beginning point of the consolidated route and not the location of the shipper that initiates the consolidation. For this example, the trucks are assumed to be randomly located. The loads shown are the available loads and will be visible to all trucks. Note that some trucks will be positioned quite close to a load (for example, the trucks near loads 5 and 6), while others could be at locations such that they have to decide which load to pick up. As shown in the figure, there are fewer trucks than loads, so the figure represents a thin market for the trucks and a relatively thick market for the loads, but this may not necessarily be the case.

A load will not be assigned to a truck simply because it offers the highest price. Figure 3.3 shows four loads vying for the services of a truck. Since the objective of the mechanism is to maximize profits for trucks, a truck will pick up a load with a price that covers its deadhead distance and the total distance of the route. Assuming

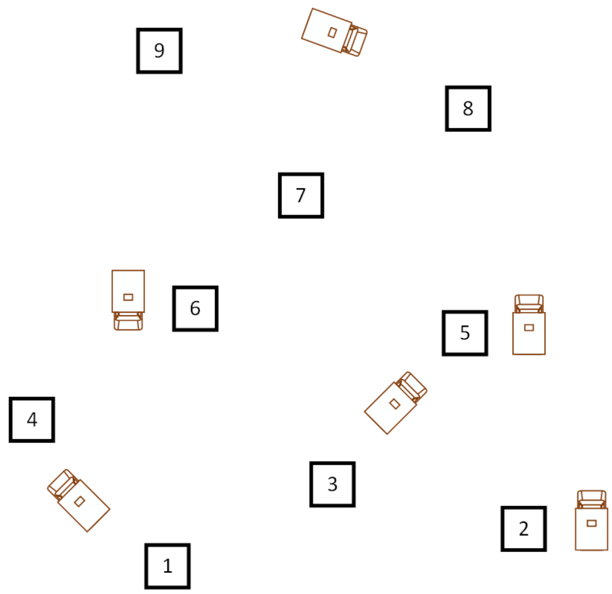

Figure 3.2: Consolidated Loads and Truck Locations.

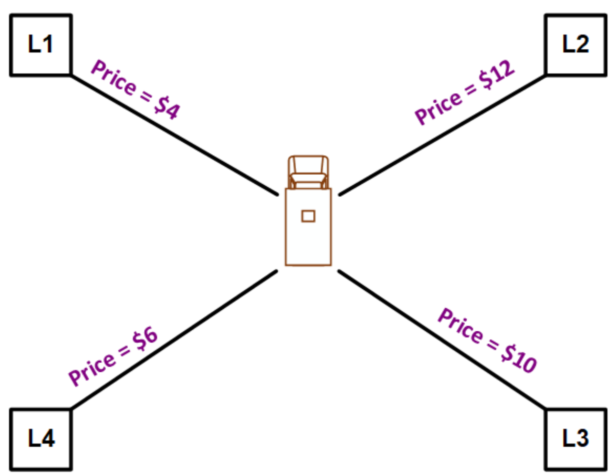

Figure 3.3: Loads With Offer Prices.

that each load offers such a price, the truck will pick the load that gives the maximum profit.

Since the greedy assignment is used to assign trucks to loads, it works by assigning the truck to the load that provides the highest profit. Once this assignment is done, the truck-load pair is removed from the assignment pool. The next truck is assigned the load with the highest profit. This pair is also removed from the pool and this process is carried out for all trucks. Once all trucks have been assigned, the total deadhead distance and the total truck profit is computed. Since in reality, the number of trucks available may be lower than the number of loads, that is, trucks may show availability as and when they become free to pick loads, the advantage of the greedy assignment lies in the fact that a load may be assigned to a truck as it is available. For example, in Table 3.1, truck T4 will be assigned load L4, then truck T3 will be assigned L2, and so on. This is unlike a private carrier whose trucks will be available to transfer whatever number of loads it consolidates.

Table 3.1: Truck to Load Deadhead Distances (mi)

	L1	L2	L3	L4	L5
T1	36.33	205.67	18.08	207.02	243.52
T2	43.66	232.16	43.22	265.83	300.41
T3	197.06	175.24	182.87	43.02	80.6
T4	167.5	266.79	165.66	127.05	173.88
T5	302.81	63.8	280.64	235.86	218.21

3.3.1 Computing Truck Profits

Consider a consolidated load $L1$ that has to travel a total distance of b miles from pickup point O to the drop point D (O and D may not belong to the same shipment). Note that since this is a consolidated load, there will be other pickup and drop points in between O and D. It is assumed that when a load is posted, all available trucks are visible to the load. Thus when $L1$ is posted, three trucks $T1, T2, T3$ are visible and ready to accept the load. In order for the truck to pick up a load, it has to travel a certain distance a to reach the pickup point O, which, from the perspective of the

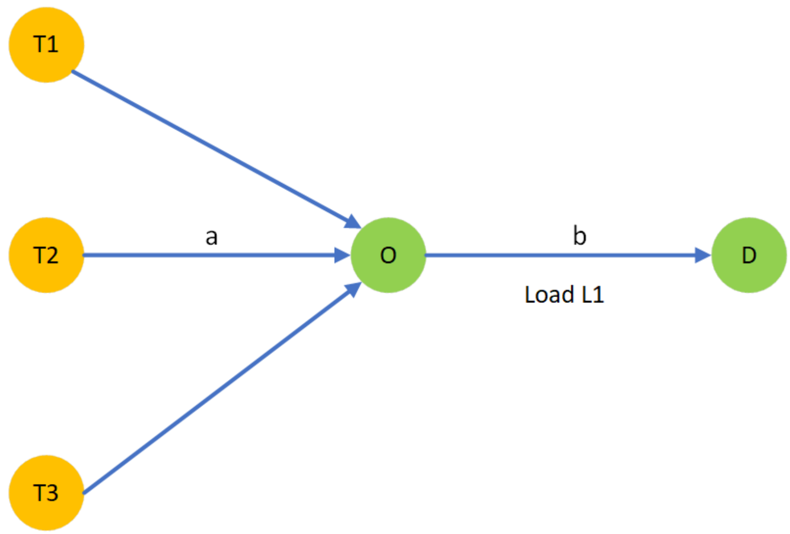

Figure 3.4: Truck Profit depends on deadhead distance a and load distance b.

truck, is a deadhead distance. The load then has to post a price that is good enough for the truck to make a profit since it has to travel a total distance of $a + b$ miles. For example, if truck $T2$ accepts load L1, its total cost of delivering the load is $(a + b)y$ where y is the fixed overall per mile expenditure that any truck incurs when it is on the road. since we have taken a TL rate of $r = 1$ as the route cost, let x_2 be the rate that $L1$ offers above this rate. Then, the truck makes a profit if,

$$b(1 + x_2) > (a + b)y \tag{3.1}$$

$$x_2 > \frac{(a + b)y - b}{b}, \tag{3.2}$$

x_2 can now be written as

$$x_2 = \frac{(a + b)y - b}{b} + \epsilon, \tag{3.3}$$

where a, b, y are known for all trucks and $\epsilon > 0$ is empirically chosen to be uniformly distributed between $(1, 1.25)$. In this way, x_1, x_2, x_3 can be chosen for all three trucks $T1, T2, T3$. Since a load will offer the same price to all three trucks, x for a particular load can be chosen as $x = \max(x_1, x_2, x_3)$. This ensures that a load offers the same price to all trucks and all trucks make a positive profit.

Table 3.2: Truck Profits($) from Five Loads

	L1	L2	L3	L4	L5
T1	376.87	132.08	320.34	77.2	155.1
T2	368.22	100.82	290.68	7.8	87.97
T3	187.21	167.99	125.89	270.72	347.34
T4	222.09	59.96	146.2	171.57	237.27
T5	62.43	299.48	10.52	43.17	184.96

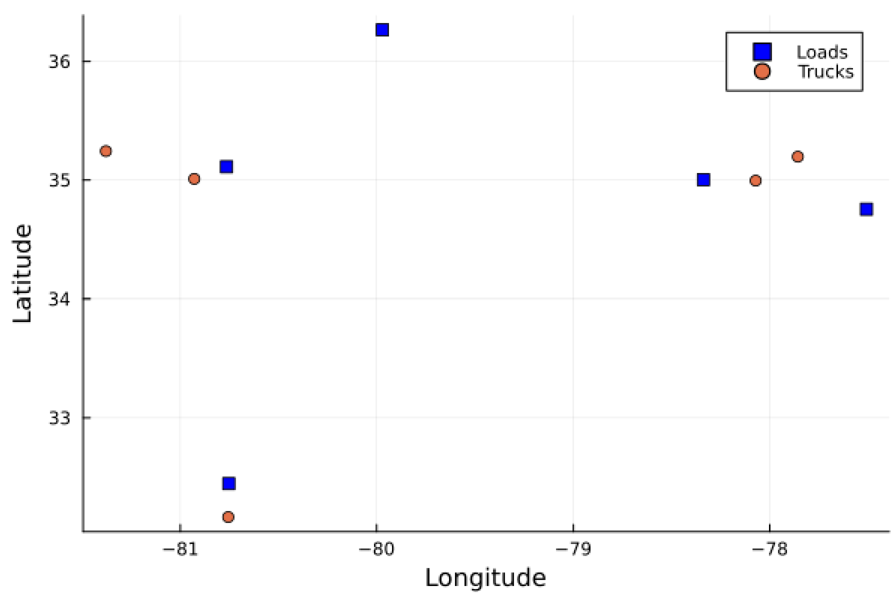

Figure 3.5: Locations of Five Loads and Five Trucks.

Tables 3.1 and 3.2 show the deadhead distances in miles between five trucks and five loads shown in Figure 3.5 and the profits earned by the trucks if they will pick up one of these loads. Assuming that all these loads are available for all the trucks, a greedy assignment of the loads to the trucks that maximizes the profit for each truck gives the assignment shown in the first row of Table 3.3. The total deadhead is 649.40 miles, and the total truck profit is $896.60.

A linear assignment model that maximizes truck profits as the objective can also be used to assign trucks to loads. This model can be visualized as the case when all trucks are owned by a single carrier. The model gives the upper bound on truck profits (and the lower bound on the total deadhead distance) for the truck load assignment problem. For the five trucks and loads example, the total deadhead is now 525.00

miles and the total truck profit is $1043.39. The linear assignment of trucks to loads is shown in the second row of Table 3.3.

Table 3.3: Assignment of Trucks to Loads

	L1	L2	L3	L4	L5
Greedy Assignment (GA)	T1	T5	T2	T4	T3
Linear Assignment (LA)	T2	T5	T1	T4	T3
Stable Matching (SM)	T1	T4	T2	T3	T5

3.3.2 Stable Matching of Loads and Trucks

The aim of the mechanism is to achieve economic efficiency by reducing total deadhead distance for trucks when they are assigned to loads. However, the assignment of loads to trucks can also be viewed from a matching perspective. The purpose of stable matching is to achieve Pareto optimality. That is, once trucks and loads have been matched to one another using the stable matching algorithm, trucks do not have the incentive to renege on their loads to reduce their deadhead or increase their profits without making some other truck worse off. In order to use stable matching, two sets of ranking can be devised. One is where each load ranks the deadhead distance for each truck from least to greatest; the other is when each truck ranks each load in terms of increasing profits. Note that the loads themselves do not rank the trucks but it is the mechanism itself that does this ranking, since the loads do not care which truck is assigned to them.

After ranking the five loads and trucks in terms of profit and deadhead distances, the assignment of loads to trucks using stable matching is shown in Table 3.3 (SM refers to Stable Matching). With this method, the total deadhead is 542.76 miles and the total profit is $1022.43.

3.3.3 Reneging Assigned Loads

Since the linear assignment model gives the lower bound on the total deadhead distance that trucks have to travel to pick up their loads, there is room for improvement

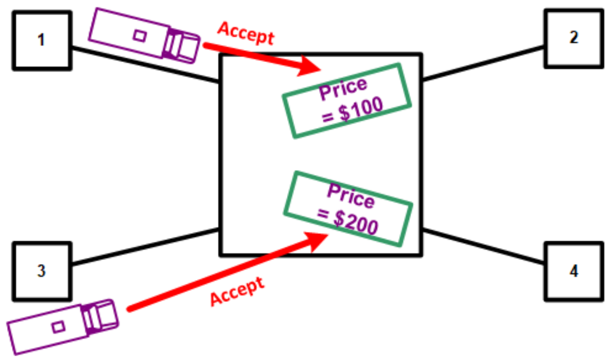

Figure 3.6: Two Trucks Accepting Loads (adapted from Kay (2004))

if any other method is used to make the assignment. The mechanism allows reneging such that if two trucks can mutually renege on their assigned loads and accept the other load, they can do so if it reduces their total deadhead distance and improves their total profit. Note that there will be costs that a truck may have to pay if it decides to renege on an already accepted load. This can be implemented by keeping the entire cost of the load in an escrow account when a load is accepted by a truck. The money can be released to the truck when it actually arrives at the pickup point of the first load to be picked up. In case a truck reneges, the truck must pay a certain percentage of the offered price of the load as a penalty to the escrow account. This amount can then be redistributed among the shippers using an allocation method.

Figure 3.6 shows two trucks that have accepted loads offering \$100 and \$200. Load 1 begins at location 1 and ends at location 2 and load 2 begins at location 3 and ends at location 4. The two trucks renege on their loads (Figure 3.7) and accept the other load when they realize that reneging allows them to reduce their total deadhead distance and gives them an increase in profit.

Table 3.4 is a summary of the total deadhead distance and the total profit for trucks when loads were formed by 10 shipments. In order to simulate reneging, loads were available to trucks in a random sequence one at a time. Greedy assignment was used to assign a load to a truck. Once a truck was assigned a load, it was removed from the truck pool and the remaining trucks competed for the next available load. Once all loads were assigned, the mechanism enforced pairwise reneging for trucks by considering all pairs of trucks.

Two cases of reneging were considered: allow each truck to renege once. Each pair was allowed to renege. If a pair switched their corresponding loads to save on

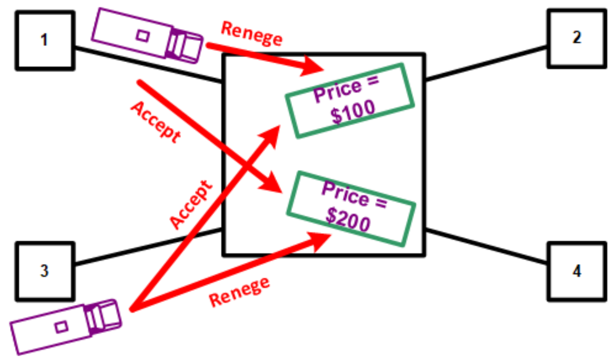

Figure 3.7: Trucks Mutually Reneging and Accepting Loads (adapted from Kay (2004)).

their deadhead distance and improve their profits, then the trucks in the pair were no longer considered for further reneging. While this reduced the total deadhead distance, the total deadhead was still far from the lower bound achieved with the linear assignment. The second case was to allow multiple reneging of trucks, again by considering all pairs of trucks but this time allowing trucks to renege more than once. Rows 4 and 5 of Table 3.4 show the results of the two cases. For a small number of loads, allowing trucks to renege more than once allowed the total deadhead distance and total profit to converge to the linear assignment values.

Table 3.4: Truck Load Assignment: 10 shipments

	Type	Total Deadhead(mi)	Total Profit($)
1	GA	351	1486
2	LA	333	1507
3	SM	608	1183
4	Renege (Once)	333	1507
5	Renege (Multiple)	333	1507

Table 3.5 is a summary of the results for loads formed by 50 shipments. Note that total deadhead in multiple reneging does not converge to the linear assignment value.

Table 3.5: Truck Load Assignment: 50 shipments

	Type	Total Deadhead(mi)	Total Profit($)
1	GA	1416	7299
2	LA	1114	7656
3	SM	3790	4498
4	Renege (Once)	1323	7410
5	Renege (Multiple)	1182	7575

3.4 Conclusion

The work done in this chapter continued with building the carrier side of the consolidation mechanism market. The work done in the second chapter proposed consolidated loads that shippers formed and that are available for acceptance by trucks. A truck acceptance protocol was developed in this chapter that allowed trucks to accept as well as renege on their acceptances of loads. The main objective when trucks accepted loads was to maximize truck profits. A number of cases were considered with this objective in mind. For example, in small experiments, it is possible that all trucks were owned by a single carrier. In this case, the carrier will make an assignment as optimal as possible so that the total profit is maximized. The carrier, in effect, thus solves a linear assignment problem that gives the optimal assignment of trucks to loads (this also includes the case where the number of trucks is not equal to the number of loads so that dummy trucks or loads are added to the model). The total profit thus found provides an upper bound for the truck load assignment problem. Relaxing the condition that all trucks are owned by a single carrier and that trucks have independent ownership, trucks were then assigned loads using the greedy assignment technique that gave a lower truck profit (and a higher value for the deadhead distance).

Since the overall goal of the thesis is to design an efficient market, the two individual pieces of the market must also be efficient. During consolidated load formation, choosing an efficient cost allocation method was the driving motive. While assigning trucks to loads, maximizing truck profits has been the main aim. Higher deadhead distances at the expense of lower truck profits mean more traffic congestion, more greenhouse gasses, and increasing overhead costs for the truck owners. The impact of these issues can be minimized by maximizing truck profits as it also lowers deadhead distances.

As an alternative to the greedy assignment or the linear assignment methods, the stable matching algorithm was also implemented with the ranking of truck profits by trucks and the ranking of deadhead distance from the perspective of loads. The stable matching gives the Pareto optimal solution for the truck load assignment where trucks do not have the incentive to swap their loads with other trucks to increase profits or reduce deadhead distances. The important observation here is that there is no guarantee that the stable matching method will reach the upper bound for truck profits achieved by the linear assignment model.

The mechanism also allowed the provision for trucks reneging on their loads and swapping with the load of another truck. While reneging will normally be triggered by some change in the environment for the trucks and the loads, for example, when a new shipper joins an existing load so that its route and offer price change, the mechanism allows reneging between pairs of trucks if there is room to reduce the deadhead distance and improve the combined profits for the truck pair. By allowing reneging among all pairs of the set of trucks, it was found that the total profit improved as compared to the original assignment. Again, to actually achieve the upper bound of the truck profits, multiple rounds of reneging may have to be done among pairs. The incentive for trucks to renege also depends on how thick the truck market is. That is, if compared to the number of consolidated loads offered in the market, the number of trucks available to pick up the loads is higher, then loads will not have the incentive to proffer higher prices to trucks, which will lower the incentive to renege assigned loads on the part of trucks. If the truck market is thin and fewer trucks are available to transfer loads than the number of loads in the market, then the incentive for loads to proffer higher prices becomes lucrative since there is no guarantee for the availability of trucks to increase. This also provides more opportunities for trucks to renege and exchange their loads if they see an improvement in their profits.

Now that both sides of the consolidation market have been developed, the next chapter adds an extra layer of complexity by introducing and analyzing time penalties. A major assumption while developing the mechanism until now is that all load pickups and drops will be made within a day. Shipments did not have any time preferences, and their choice of loads was based on an acceptance probability (described in Chapter 2). We relax this assumption in the next chapter such that individual load deliveries can take more than a day, and shippers will be able to apply penalties to carriers and their trucks if pick-up and drop times are not met. Shipments will have different time preferences which will drive their becoming a part of a consolidation and their load acceptance probabilities.

CHAPTER

4

IMPACT OF TIME PREFERENCES ON SHIPMENT CONSOLIDATION

4.1 Introduction

The previous two chapters focused on the creation of consolidated loads using the most efficient cost allocation mechanism, and then assigning those loads to trucks with the objective of reducing deadhead distances. In those chapters, time preferences and related penalties were not a factor when loads were formed and assigned to trucks. While cost allocation and load assignment are important features of the mechanism, there are two important aspects of load formation and assignment that were not considered in those chapters. The first aspect is the time preferences of shippers. When a shipper initiates a load, and another shipper joins the load, because of the online route formation, the new shipper can easily see the tentative pickup time of its load. The shipper can make a decision, based on its preferred pickup time and the tentative pickup time, whether it wants to join the load. If it is willing to accept the deviation, if any, from its preferred pickup time and joins the load, then the shipper coalition so formed can drive the load formation by allowing other shippers to join or remove from the load by observing deviations from their preferred pickup times.

Implicit in this aspect is the fact that the value of time can be significantly different for two shippers willing to transmit identical loads. For example, one shipper might wish to urgently send a part to its destination, while another shipper would send the same part to be kept in storage. The first shipper obviously will be willing to pay much more for transport while the other shipper would be willing to wait if it leads to a lower transport cost. From a time preference perspective, the first shipper has a time priority while the second shipper has a lower time priority. This necessitates the use of modeling techniques that will incorporate such distinction among shippers.

The second aspect is the time at which loads are picked up by the trucks and delivered since they improve the reliability of the trucks and the carriers. This is especially relevant when a shipment might request a one-day delivery, while another shipment is comfortable if the delivery is made within a week. Freight transportation usually involves a number of interested parties, and unreliable transport times can cause delays and lead to more complicated problems that can be difficult to quantify. If there are any delays in the pickup and drop of shipments along the consolidated route, then the shipper will apply a penalty to the truck and deduct that amount from its payment to the truck. This would be a public act and will be part of the mechanism. The challenge in this case will be to arrive at reasonable values to use for the time penalties associated with missing pick-up and delivery targets. Note that the penalty applied by a shipper for missed time targets is quite different from the penalty that each shipment uses internally (and which is private information) to decide to join a load.

Given that there can be two shipments with identical load parameters, including pickup and drop points, but different time sensitivities, that is, one shipment has a tighter time window for its pickup or drop, while the other shipper has a more relaxed time window. What would be acceptable ways to specify the preferences of both these shippers?

1. The shipper with the tighter time window may post a higher price for transferring its load compared with the one that has a more relaxed time window. The extra amount bid by the shipper will depend upon how sensitive it is to its time window.

2. Shippers can be ranked according to their time preferences, and the ones with

tighter time windows will be given first preference to be the lead shipper and start the load. This will allow shippers with relaxed time windows to join later.

As and when the route is formed with new shippers joining the load, the lead shipper that initiated the load can post a pickup time window and a drop time window for its load. This will then allow all shippers in the load to know beforehand, at least tentatively, based on their pick-up and drop-off positions in the route, if their personal time windows are being met (arrival times at a particular point can be determined by assuming a 50 mph truck speed). For example, if a second shipper joins the load, based on the posted route information and the lead shipper's time windows, it can decide if its personal time windows can be met by a truck traveling at an average speed. If it believes its window cannot be met, it drops from the consolidation, and some other load is given the opportunity to join. The advantage of this approach is that only one shipper posts its time window preferences allowing other shippers to keep their time preferences private.

The research issues raised so far can be handled by assuming a vehicle routing problem with time window preferences, but as mentioned, time preferences are private information for a shipper. If all time preferences are publicly available, then shippers will assume that someone with tight time windows will most likely bid higher to transfer its load. This then becomes an incentive for a shipper with a relaxed time window to actually bid low when it becomes part of a load with tighter time windows and take advantage of other shippers.

The actual computation of time penalties can be done by assuming that the transportation time is equal to the replenishment lead time for a retailer. Any delay in picking up or dropping off a load can thus be assigned a shortage penalty cost to the truck.

The work done in this chapter will address the following research questions:

1. In order to accommodate time preferences, what minimum information needs to be made public by shippers as part of the load formation protocol?

2. Given that shippers have time preferences, is a fixed time for pick up and delivery preferable to trucks, or time windows preferable to trucks?

4.2 Related Research

Vehicle routing with time window penalties requires efficient heuristics to obtain good solutions. Some of the heuristic techniques applied include genetic algorithms (Thanglah et al. 1991; Potvin and Bengio 1996) and simulated annealing (Kokubugata et al. 1997). Taniguchi and Thompson (2002) used genetic algorithms along with dynamic traffic simulation to simultaneously determine the departure time and the assignment of vehicles. Much of the work on vehicle routing with time preferences has been done to estimate the travel time of vehicles to distribution centers (DCs). van der Spoel et al. (2017) develop methods to predict the arrival times of trucks at DCs so that DCs can respond adequately. They use predictive analytics instead of causal-exploratory analysis to predict arrival times and estimate the power of the model.

Using supply chain data from Taiwan, Chu et al. (2017) use stochastic models for vehicle routing rather than deterministic models and find that the former reduces unfulfilled demands compared to the latter. They distinguish between a planned scenario in which all windows are met versus unplanned scenarios where some or all of the windows are unmet and impute an unanticipated cost to the unplanned scenario. Note that this unplanned scenario is equivalent to the penalty cost that a shipper will impose if its time preferences are not met.

Bolis and Maggi (2003) conduct experiments to determine the most important qualities in freight transport and their monetary values (in Swiss franc, CHF). Their experiments indicate that, per net ton, the reliability of freight services (that is, meeting time windows) has a higher monetary value than transporting goods in an hour or less. For example, for a full truckload of 28 tons, a firm was willing to pay on average 17.25 CHF more to save an hour, but 36.30 CHF to gain 1% more in reliability. This conclusion was further confirmed in a review by Shams et al. (2017) in which the authors found that the value of freight reliability was four to six times that of savings in travel time.

4.3 Simple Example

In order to understand the impact of time preferences on consolidated load formation, a home delivery logistics network is first considered (Figure 4.1). In a home delivery logistics network, orders originating from stores that are located far from a home

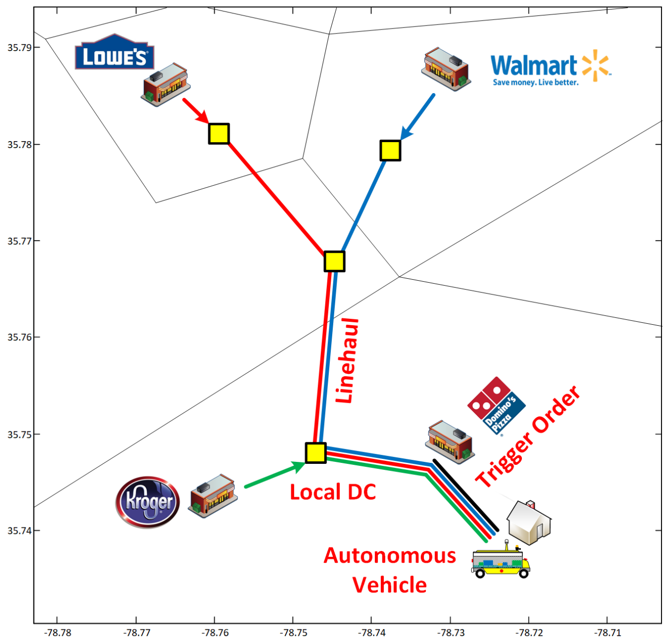

Figure 4.1: Design of a Home Delivery Network (adapted from Kay (2024)).

could be sent to a distribution center (DC) that is closest to the stores. The order could then be routed through a sequence of DCs before it reaches a DC closest to the home. If the home does not need immediate delivery of these orders, they can stay at the DC until an urgent order (like a restaurant takeout) arrives at the DC. Since the DC that will deliver to the home is located close to the home, the trigger order allows all accumulated orders till that point to be delivered to the home in one trip Kay (2024).

As an example, we propose a simple home delivery system where there is one source DC and one destination DC delivering to ten different customers. The data for the ten customers are listed in Table 4.1. The pickup times are in twentyfour hour format and are the times at which the order will be picked up by the truck at the source DC. The home delivery model considers two different types of customers. Customers 1 to 4 have a low time sensitivity for the items they want delivered and so the penalty they apply for each hour of delay is $1. Customers 6 to 10 are high time-sensitive customers and apply a penalty of $5 per hour of delay. The model does not differentiate between early and late pickups, so the same penalty is applied if the order is picked up before the scheduled time for a customer. The Deadline column lists the hours that the customers are willing to wait for the order to be delivered. The low time-sensitive customers are willing to wait up to 10 hours for their order to be delivered, while the high time-sensitive customers will wait up to 3 hours.

Table 4.1: Home Delivery: Time Preferences for 10 Customers

Customer	Pickup Time	Penalty($/hr)	Deadline(hrs)
1	10	1	10
2	2	1	10
3	16	1	10
4	15	1	10
5	18	5	3
6	6	5	3
7	4	5	3
8	20	5	3
9	14	5	3
10	7	5	3

With the above setup, the mechanism allows consolidated loads to be formed but now considers time deadlines also as a condition for joining the load. The customer that initiates the load makes its load pickup time public, while all other customers willing to join the load keep their pickup times private. As each new customer joins the load, it is able to see the online route formation and, assuming the lead customer's order will be picked up on time, can calculate its pickup time and the penalty if there are any deviations, based on an assumed truck speed (the assumed truck speed is the same for all customers). If the deviation from the pickup time exceeds the deadline set by the customer, it does not join the load. For the home delivery example, since the pickup and destination DCs are the same for all customers, the pickup time will be the same for all customers in a load and so the time deviation is just the deviation from the lead customer's pickup time. Executing the mechanism with these data gives the following load formations (Table 4.2), with a total cost of $532.21 for delivering the loads of all customers.

Note that the allocation method to allocate costs is not relevant for a home delivery system since all customers in a load will share the total cost equally, which in this case will be the simple average of the total cost for each load. The penalty column lists the individual penalties for each customer in the load. For example, customer 8 in Consolidation 2 has a pickup time at 20 hours. Customer 5 joins the load and

Table 4.2: Home Delivery: Load Formations for 10 Customers

Consolidation	Customers	Allocated Cost($)	Penalty ($)
1	[6, 10, 7, 4, 3, 2, 1]	25.34	[0.0, 5.0, 10.0, 9.0, 10.0, 4.0, 4.0]
2	[8, 5]	88.71	[0.0, 10.0]
3	[9]	177.41	0

is willing to incur a delay of 2 hours and charge a penalty of $10 (its pickup is at 18 hours) since the 2 hours delay is within its deadline of 3 hours. The table also shows that consolidation 3 is formed by just one load with customer 9 which goes independent and thus bears the full cost of the route.

4.4 Discussion and Conclusions

This example assumes a single pickup time with the same pickup and drop locations for all shipments. The only information that is assumed to be made public is the lead customer's pickup time. There can be numerous extensions and modifications to this model to understand which case(s) are more economically efficient. For example, early arrival of a truck could incur zero penalty, while late arrivals could incur a linearly increasing penalty. There can be separate penalties for drop times also with customers considering either a linearly increasing penalty or a flat penalty amount. We would also like to explore the amount of information that can be made public by customers. For example, the impact on load formation and total costs when shippers publish time windows instead of single pickup and delivery times. These cases will become more relevant when the model is extended to the more general condition with multiple and possibly distinct pickup and drop locations.

CHAPTER

5

PROTOCOLS FOR HOME DELIVERY USING A PUBLIC LOGISTICS NETWORK

5.1 Introduction

In this chapter, a public logistics network (PLN) is proposed as an alternative needed for home delivery (Kay and Parlikad 2002). A sequence of public distribution centers in metropolitan areas can be used to transport a package from a store to its destination, similar to the transportation of packages through an internet network (Varian and MacKie-Mason 1993). This will allow delivery of the package to homes in a matter of hours, thereby reducing car trips to the stores. Currently, it is common for a single logistics firm like UPS and FedEx to handle a package throughout its transport. In a PLN, the different functions of the network would be separated so that a single firm would not be required to coordinate. This would enable scale economies to be realized in performing each logistics function since each element of the network has access to potentially all of the network's demand. The increase in scale would make it economical to have many more transshipment points. Each transshipment point, or DC, could be an independently operated facility that serves as both a freight terminal and a public warehouse and could be established in small cities and towns that would never have such facilities if they were served as part of a proprietary, private logistics network (Park et al. 2023).

The proposed Public Logistics Network (PLN) will provide significant advantages to state municipal governments, as well as shippers and carriers. Since a number of micro DCs will be located within $1-2$ miles of each urban household, a carrier can provide a truck full of items at the same cost that a company like Amazon takes to deliver about two packages to a single household. The DCs will use a consolidation mechanism to run the network and the time preferences of the shippers will drive the consolidation. Currently, Amazon's delivery model is beneficial to the drivers and the vehicles, not the customers. With the help of micro DCs, time-sensitive items can be delivered to customers at a cheaper cost compared to what customers currently pay to companies like Amazon. The PLN will have a significant impact on environmental sustainability since each DC can use reusable containers. For example, part of the demand at a DC could be items that have to go to a restaurant in these containers. The PLN can be provided as a planning tool to a municipal government, with demographic information and location of DCs that will be available for download, so that the government can do a pilot test of the PLN.

One significant benefit of providing consolidated deliveries to the home is that, after delivery, the vehicle is available for backhaul (Figure 5.1), making it possible for things like take-out food to be delivered in reusable containers. Being able to return reusable containers conveniently would provide a foundation for a sustainable environment. It would make it possible to eliminate the need for most recycling and the one-time use of plastics.

5.2 Related Work

Pricing issues related to the transfer of packages via routers have been discussed in Varian and MacKie-Mason (1993). Basics of a public logistics network for home delivery were developed in Kay and Parlikad (2002) and in Kay (2018) and Kay (2022). These papers highlight the fact that by using a PLN, the need for people to travel to

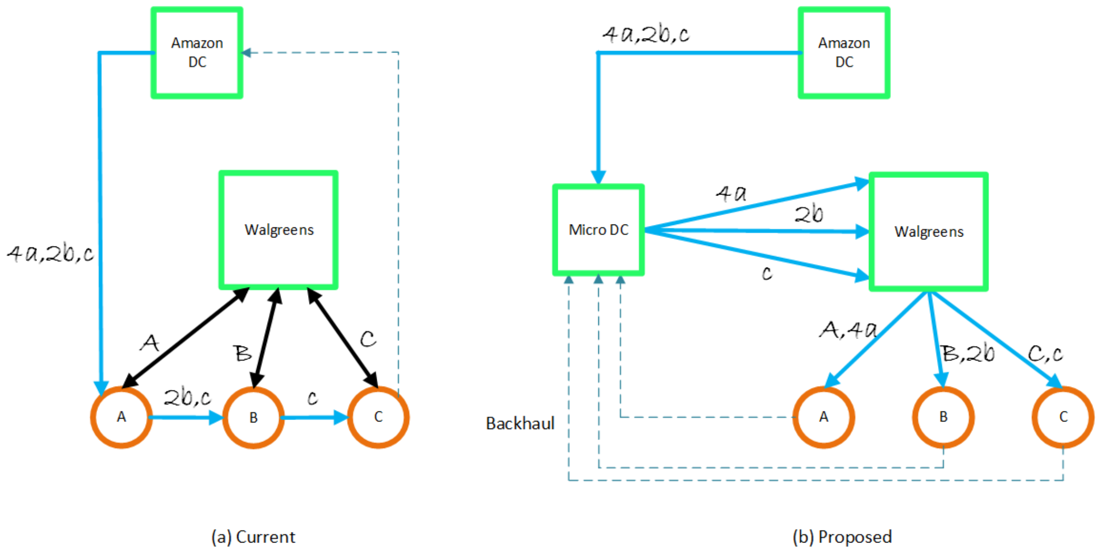

Figure 5.1: Current logistics network and proposed PLN.

and shop in a store can be significantly reduced. Since the current logistics providers do not effectively provide high-cost human labor, the use of autonomous vehicles and drones has become economically feasible. More recently, home-based logistics and its impact on last-mile delivery have been discussed in Park et al. (2023). Five strategies are proposed, including augmenting the transportation capacity of freight trains, expanding or establishing hubs and terminals, an advanced delivery system for the last mile, training workers in this domain, and establishing a technical, financial, and legal framework that supports this infrastructure.

5.3 Delivery Costs

Compared to the current logistics network for home delivery, using a PLN for home delivery can provide significant savings to customers, as shown in Figure 5.1. Consider three customers, A, B, and C, that use Amazon to get four, two, and one orders of non-priority items a, b, c, respectively. The delivery of these items takes over distinct periods of time; that is, the items are not ordered by the customer in one request. The Amazon delivery van makes these deliveries from its DC by dropping an item at customer A, then traveling to customers B and C to make their deliveries. The customers also make a personal trip to Walgreens to buy an important item that they need within a short period of time (for example, within an hour). These 'trigger' orders are labeled A, B, C in Figure 5.1.

Table 5.1: Vehicles used for home delivery (summarized from Appendix A, Kay (2022)).

	Vehicle Type	Cost per delivery($)
1	Delivery Van	1.58
2	Nuro delivery vehicle	N/A
3	Autonomous vehicle	2.27
4	Starship delivery vehicle	1.89
5	Drone	0.80

Table 5.1 lists various home delivery vehicle alternatives with their per delivery costs. With the current logistics network and using the per delivery costs associated with a delivery van and the cost associated with one customer round trip to a store, the total cost to deliver four items of a and one item A is $4(1.58) + 11.94 = \$18.26$. ($11.94 is the total cost estimate per customer pickup trip. Kay (2020))

In the proposed PLN for a home delivery network using autonomous vehicles, the non-priority items can be transferred from an Amazon DC and stored in a micro DC that would be located within $1.5 - 2$ miles from customer A. When the customer requests a trigger order from a Walgreens store, the non-priority items, along with the trigger order (picked up from Walgreens), can be delivered to the customer for a delivery cost of $2.27. The savings for customer A using a PLN is thus about 8 times that of the current logistics network. After making the delivery, the delivery van will be available to backhaul reusable containers. Additional savings can also be realized since the van can also be used to return items by customers, a service that is not currently provided by companies like Amazon, where customers have to make a trip to the nearest USPS or some other store to return items they do not want.

5.4 Agent Based Coordination Mechanism

Software agents will control packages and trucks. The DC provides services that will allow the agents to interact with each other (see Figure 5.2). For example, each package agent will be provided with all the active load bids and the expected arrival times of

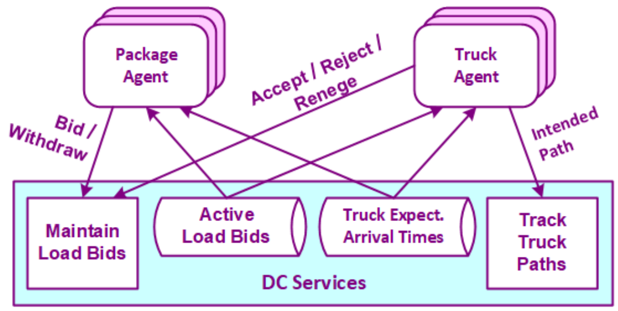

Figure 5.2: Framework for agent-based coordination (adapted from Kay (2018)).

trucks, along with the location of trucks as part of the service provided by the DC. The load bids and truck locations, arrival times and their paths will be maintained by the DC. The DC will also determine the order in which trucks will be assigned the active loads.

5.4.1 Truck Protocol

The truck protocol consists of two rules and is used to determine which truck is used to transport what load at what DC. The goal for truck operation is to try to match the load that values transport the highest with the truck that can provide that transport service at the least cost (e.g., the truck closest to the DC).

1. **Priority for Accepting Loads**: *Opportunity to accept or reject load will be based on truck's arrival time at DC*. Priority for accepting or rejecting an available load at a DC is first given to trucks already at the DC based on their order of arrival, earliest first, and then to trucks not at the DC based on their expected time of arrival at the DC along their intended path to the DC.

    - If all trucks reject a load, the load is posted at the DC and is then available for any truck to accept.
    - The expected time of arrival of each truck is posted at the DC.
    - Although the intended path chosen by a truck does not have to be the quickest path to the DC, late expected arrival at the DC can affect the truck's priority for being able to accept the load.
    - A truck's portion of the load bid is fixed after acceptance.
    - Each truck's current location is assumed to be known at all times.

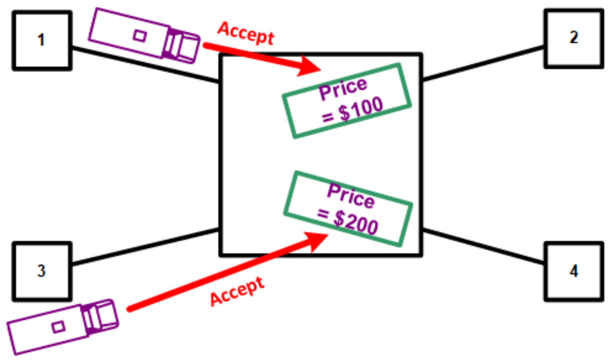

(a) Two trucks accepting loads

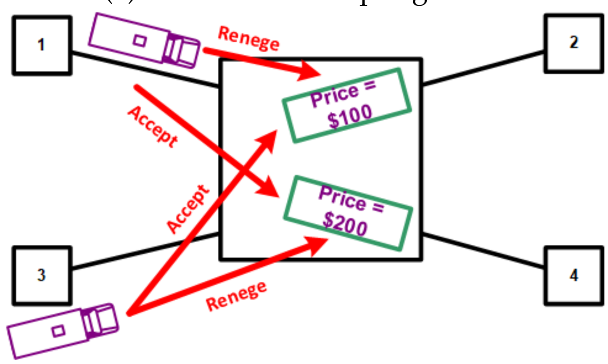

(b) Trucks reneging and accepting Loads.

Figure 5.3: Truck reneging loads example (adapted from Kay (2004)).

2. **Reneging**: *After reneging, a truck cannot again accept the same load until all other trucks have rejected it.* A truck can renege on its acceptance of a load at any time, but it will not be able to accept the same load until all other trucks have been offered the load and have rejected it. Figure 5.3 presents an example of how two trucks can both renege on their loads in order to capture all of the increase in a load bid.

5.4.2 Package Protocol

The package protocol will consist of rules that will be used to determine which packages are selected to join a load. The goal for package selection is to encourage a package to submit a bid that represents its true value for transport as soon as possible, thereby allowing trucks to be more responsive and discouraging multiple-bid auction-like behavior.

1. **Load Formation**: A load will be initiated by a package agent managing packages. Other package agents can then allow their packages to become a part of the putative load. Each agent bids a price for its package, and the sum of the bids before a truck accepts the load becomes the price that will be offered to a truck that finally accepts the load. Packages can leave the load at no penalty before a truck accepts the load, but after a truck accepts the load, a package will be charged the full bid amount if it wants to leave the load.

   - All loads are assumed to have the same capacity (i.e., maximum size).
   - Both the bid amount and the size of a package determine whether it will be selected for a load.
   - When necessary, equivalent package bids are selected based on the elapsed time a package has spent in the network.

2. **Allocation of Load Bid**: *Truck's portion of a load bid does not increase after acceptance*. The portion of a load bid given to the truck that accepted the load is equal to the amount of the load bid at the time of acceptance; all subsequent increases in the load bid are given to the packages that were in the load at the time of acceptance (and remain in the load) in proportion to their bid amounts.

   - Once a load has been accepted, no further increase in the amount given to the truck is needed to change its behavior, while giving the increase to the original packages encourages them to submit their bids early, thus increasing the load bid while it can still influence truck behavior.
   - Once a truck has accepted a load at a DC, it has no incentive to delay its departure to wait for additional packages, thus resulting in predictable departure times.
   - A truck can renege on an accepted load in order to try to capture all of the increase in a load bid, but it must take into account that all other trucks will first be given the opportunity to accept the load.

5.4.3 Simple Example 1

This example and Table 5.2 depict how loads with bids will be created, and how bids will be allocated among existing packages. Figure 5.4a shows loads available at DC 1 whose destinations are to DCs $2, 3, 4, 5$. For example, for the load whose total current bid is \$12, it is assumed that a package agent with a package that wished to transport the load to DC 2 initiated this load with a bid price of \$5. This load was then posted so that it was visible to all package agents. Two package agents then joined this load with package bids of \$4 and \$3. Note that agents can join multiple loads. For example, the package agent with a bid of \$4 could have also joined the load with a current bid of \$10. If the final destination of the package is DC 2, the agent can join a load that goes directly to that DC or can join a load that goes to DC 3. It can then join another load at DC 3 that goes to DC 2. If a truck immediately accepts the \$12 load, it can drop its package from the other load. If both loads have been accepted, a penalty will be incurred by dropping since it can only be in one load.

Loads are offered to trucks based on their location from the DC. Trucks located at the DC will be given first preference for acceptance at all loads posted by the DC. Once the nearest truck rejects a load, the next nearest truck will be offered the load. Figure 5.4a shows one truck located at or close to the DC and another truck that is en route to the DC but is still farther away from the first truck. The truck nearest to DC 1 will be offered all current, unaccepted loads. The truck agent decides to accept a load based on the reservation price for its truck, which will be private information. For example, for the truck nearest to the DC, if the reservation price for this truck along all routes (from DC 1 to the other DCs) is higher than \$12, it will not accept any of the loads. But if its reservation price for route DC 1 to DC 2 is \$10, it will accept the load with bid \$12. The second truck will then get an opportunity to accept the other loads.

The process with which loads will be accepted and charges allocated among packages is shown in Table 5.2. A package agent initiates a load with a package bid of \$8. The Bid/Cost column shows that the cost if this load is accepted, will be \$8 to the package agent. A second package agent joins the load with a package bid of \$2. The load is now accepted by a truck whose reservation price must be lower than \$10, the total current bid of the load. A third package agent then joins the load with a package bid of \$5. While the total bid is now \$15, the truck will only get \$10. The extra \$5 of the bid is used to reduce the cost of the packages that existed before the load was accepted. The extra bid is allocated to these packages in proportion to their bid costs. In this case, \$5 will be allocated to the \$8 and the \$2 in the ratio 4 : 1, and this allocation will then be used to reduce the cost of the packages. The cost of the \$8 package reduces to \$4, and the cost of the \$2 reduces to \$1.

The package agent with the \$5 bid then decides to leave the load. Since the load

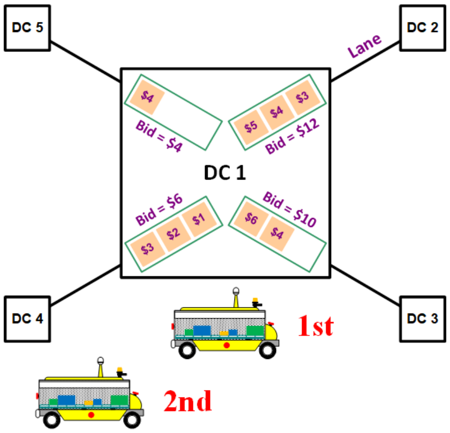

(a) Loads with load bids at a DC.

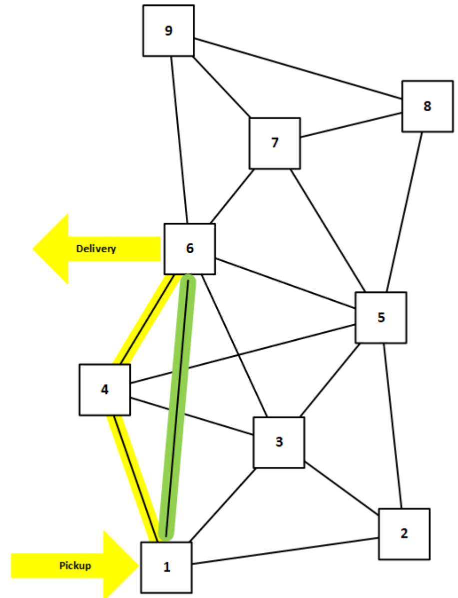

(b) Package pick-up and delivery from DCs in a PLN.

Figure 5.4: PLN network with load bids, pickup and delivery (adapted from Kay (2004)).

has already been accepted and the truck is guaranteed its \$10, the agent is charged the full cost of \$5 (If a package agent drops its package from an unaccepted load, it will not incur any penalty.). This amount is placed in a Balance account which will be used by the DC to pay the truck. A new package agent then joins the load with a bid of \$3. Since the bid is now \$13, this new package has to pay \$3 and with a balance of \$5, the extra $10 - 8 = \$2$ is allocated among the \$8 and the \$2 to further reduce their costs to \$1.6 and \$0.4 respectively.

5.4.4 Simple Example 2

Consider the DC network shown in Figure 5.4b. Seven packages P_1 through P_7 are located at DC1. Some of these packages have to be delivered at DC4, while others have to be delivered at DC6. Each package is managed by a package agent that sets a bid range for its package. For example, as shown in Table 5.3, P_1 has a low bid of \$30 and a high bid of \$40. When a package agent initiates a load, it wants its load to be accepted as soon as possible, so it starts with a high bid price. High bid prices will also ensure that truck reservation prices are met within short time frames.

Table 5.2: Allocation of Load Bid

Package Event	Truck Response	Load Bid	Truck Portion	Allocated Portion	Load Bid/Cost	Balance
Bid & Join	**Reject**	8	8	0	**8/8**	0
Bid & Join	**Accept**	10	10	0	**8/8, 2/2**	0
Bid & Join	-	15	10	5	**8/4, 2/1, 5/5**	0
Drop	-	10	10	5	**8/4, 2/1**	5
Bid & Join	-	13	10	3	**8/1.60, 2/0.40, 3/3**	5

For example, a package agent initiates a load L_1 with P_1 with a bid price of \$40. Packages P_2 and P_3 then join the load with bid prices of \$25 and \$50 respectively. The total bid price is now \$115. The DC (here DC1) maintains a priority queue of trucks available for load acceptance, with preference given to trucks nearer to the DC. In this case, truck T_1 (Table 5.4) is first offered the load as it is located at DC1. Since its reservation price is \$120, it rejects the load. The load is then offered to truck T_2 that is located at DC3 and is en route to DC1. T_2 accepts the load since its reservation price of \$100 is lower than the current bid, \$115, of the load. Once the load is accepted, T_2 gets a fixed price of \$115. This gives an incentive for other packages to join the load at low bid prices which reduces the cost of the existing packages in the load. For example, packages P_4 and P_5 join the load after acceptance with low bids of \$12 and \$47 respectively. The costs allocated to each package after P_4 and P_5 join the load is shown in Table 5.5.

A new load L_2 is then initiated at DC1 by a package agent to transfer package P_6 to DC6. The route in yellow in Figure 5.4b is the route that P_4 would have taken by joining the load initiated by the agent for P_1 (pick-up at DC1), by going to DC4 and then joining another load at DC4 to reach DC6 (delivery at DC6). The green route is the route that P_4 wants to take to directly reach DC6. Since the agent for P_6 initiates this new load, packages P_7 and P_4 join the load for a total bid of \$125 which is now accepted by T_1 (assuming it is the first truck to be offered the load) whose reservation price is \$120. The agent managing P_4 then decides to drop P_4 from L_1 and incurs a penalty of \$12 since this load was already accepted. As discussed in the earlier

Table 5.3: Packages with bid prices($)

Package Id	Begin	End	Bid Range($)
P_1	DC1	DC4	$[30, 40]$
P_2	DC1	DC4	$[5, 25]$
P_3	DC1	DC4	$[30, 50]$
P_4	DC1	DC6	$[10, 15]$
P_5	DC1	DC4	$[45, 52]$
P_6	DC1	DC6	$[40, 60]$
P_7	DC1	DC6	$[30, 50]$

example, the $12 is placed in a balance account, and none of the costs allocated to the existing packages in L_1 are impacted.

Table 5.4: Trucks with reservation prices($)

Truck Id	Location	Volume(cu. ft)	Weight(ton)	Reservation price($)
T_1	DC1	2750	25	120.0
T_2	DC3	2750	25	100.0

Table 5.5: Packages with allocated costs. Allocated cost for existing packages does not change after P_4 drops from the load.

Packages	Bid($)	Cost($)
P_1	40	19.48
P_2	25	12.17
P_3	50	24.35
P_4	12	12.00
P_5	47	47.00
Total	174	115.00

5.5 Agent Based Model Development

Based on the protocols discussed above, an agent based model consisting of package agents and truck agents was developed in Julia using the API of Agents.jl. As shown in Table 5.6, the package agent now consists of a number of fields. In particular, package agents bid prices based on a triangular distribution, consisting of a low bid, a high bid and a bid peak point. The package agent structure also consists of boolean values like committed (whether a package is committed to be in the load) and expired (whether the maximum pickup time of the package has passed). Package agents also have various bidding strategies that will be stored in the field bidding_strategy. For example, package agents can be aggressive or conservative in their bidding strategies.

Table 5.6: Package agent and Truck agent fields.

Package Agent	**Truck Agent**	**Load**	**Distribution Center**
id	id	id	location
pos	pos	packages	queued_trucks
size	current_load_id	total_size	available_loads
high_bid	arrival_time	total_bid	
low_bid	expected_arrival_times	accepted_by_truck	
current_bid	intended_path	accepting_truck_id	
bid_peak_point	has_rejected	max_capacity	
penalties	reneged_loads	dc_location	
min_pickup_time	current_dc	rejection_count	
max_pickup_time	dc_arrival_order	posted_at_dc	
entry_time	reservation_price	trucks_rejected	
in_load		trucks_considered	
committed		fixed_truck_payment	
expired			
bidding_strategy			

The *Truck Agent* consists of fields that provide information about its arrival time at a distribution center (DC), the current load it has been assigned to at a DC, the loads it has reneged and its reservation price, among others. Since the location of trucks will be known at all times, its intended path to a DC is also stored in the intended_path field.

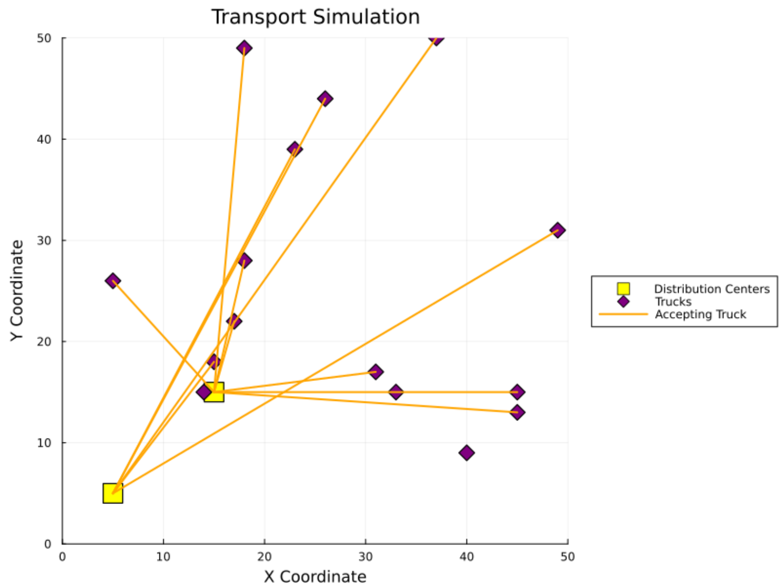

Figure 5.5: 2 DCs, 200 Packages, 15 Trucks.

The *Load* column is not an agent based structure and is just defined as a mutable structure in Julia. It consists of the packages it is made up of, whether it has been accepted by a truck, and the identity of the truck if it has been accepted, the trucks that have considered or rejected it, and the truck payment it will make. Similarly, *Distribution Center* is a mutable structure in Julia, and consists of its location, the trucks that are en-route to it to pickup loads and the available loads.

In order to run a simulation using Agents.jl, a $50x50$ continuous space grid was created and two DCs were located at $(5,5)$ and $(15,15)$. Packages were then located at either one of these two DCs, with their destination to the other DC. Trucks were randomly located throughout the grid space. The simulation of the agent based model was then done for 200 packages and 15 trucks. Loads were created i the two DCs and trucks were allowed to accept or reject loads based on their location from the DCs. The result is shown graphically in Figure 5.6. Some trucks were not assigned loads because when the load was offered, the price was below their reservation price, but when the load price finally exceeded their reservation price some other truck had accepted the load.

The behavior of bid prices of package agents was also investigated with respect to the number of loads at the DC. With 800 packages, a negative correlation between bid price and number of loads at a DC is observed for most packages because many

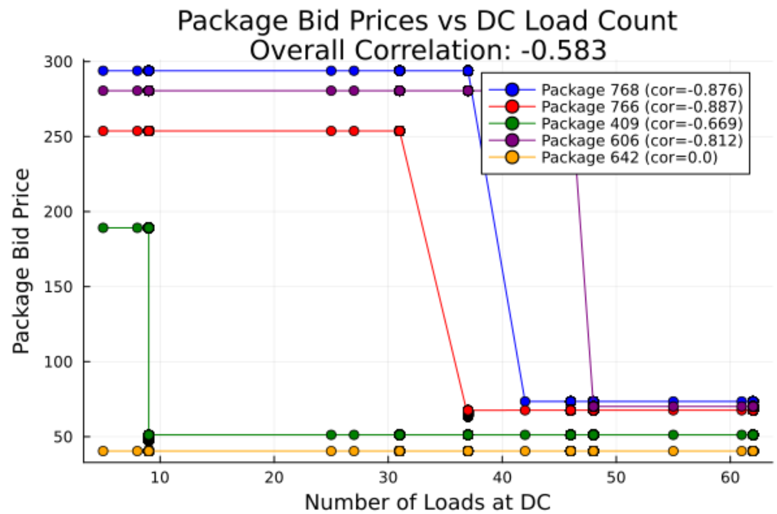

Figure 5.6: Bid Prices vs DC Loads: 800 packages.

packages do not have the incentive to bid higher when the number of loads increases, since the probability of loads being accepted by a truck increases. This is also true if there are a high number of trucks present in and around the DC. With more number of trucks, package agents do not have to bid higher since, with more number of packages forming loads, the reservation prices of trucks can be easily met or exceeded.

5.6 Discussion and Conclusions

A public logistics network for home delivery was developed in this chapter using a simple example consisting of truck agents and package agents. Package agents are encouraged to initiate a load or join a load as early as possible since their portion of the allocated cost does not change if packages join or drop the load after the load has been accepted. Package agents begin or join loads based on protocols that define which packages will be able to become part of loads. Similarly, truck agents are assigned loads based on the truck protocols defined. For example, the DC will maintain a priority queue of trucks that it will use to assign loads to trucks. If the closest truck to the DC declines a load, it will be available for the next truck to accept. The simple example developed in this chapter showed how various components of the PLN will work to create loads and transfer them. The example used point-to-point routing, but the model can be extended to include multi-stop routing. The total round-trip distance traveled by vehicles for point-to-point deliveries is about twice the distance required to make the same deliveries as part of a multi-stop route, but with point-to-point, after unloading a delivery at home, the vehicle is available to travel back to the DC. Using what would otherwise be an empty backhaul for transport from home would make it possible, for only a little extra cost, to send empty containers back to the DC so that they can then be reused. This would be a low-cost way of providing logistical support for encouraging sustainability through the reuse of containers instead of just recycling or disposing of all containers and packaging (currently, only 10% of plastics are recycled Sullivan (2020)).

The development of autonomous vehicles and drone technology means that a PLN can be a viable alternative for home delivery, as opposed to services provided by Amazon, USPS, etc. Interior villages to which road communication is delayed or sometimes impossible due to inclement weather can benefit from drone services that can deliver packages without the hassle of road transport. The availability of backhaul after the packages are dropped at their locations means that containers can be used to return packages that households do not need, send dirty laundry for washing, or distribute goods manufactured at homes.

CHAPTER

6

CONCLUSIONS

The work done in this research demonstrates the feasibility of shipper-driven consolidation mechanisms under various allocation schemes and the assignment of consolidated loads to carriers motivated by profit maximization.

6.1 Summary

A shipper-driven consolidation mechanism is designed so that it can serve as an almost cost-free means to enable all of the benefits of consolidation to be divided among the shippers, as opposed to other means of consolidation that share a significant portion of benefits with a third party as the cost of facilitating the consolidation. The benefit, or "savings," associated with consolidation is the difference between the sum of the charges to transport each shipment independently and the charge to transport the shipments as a consolidated load. How the savings are allocated among the shippers, the carrier, and any other third party that helps facilitate the consolidation is based on knowledge of the number of shipments available for consolidation and the variability with respect to the location and frequency of the shipments. A key aspect of the research was the choice of cost allocation mechanism that gave the most incentive to shippers to join a load. The proportional, as compared to the Shapley, was found to be more efficient in the long term and incentivized more shippers to form consolidated loads.

The carrier side of the mechanism was then developed to assign trucks to consolidated loads. The incentive for trucks to choose loads was based on profit maximization, so various methods, such as linear assignment and greedy assignment, were used to show how this objective was attained. Two extensions of the mechanism were also developed. The first two chapters, in which the shipper and carrier sides were developed, did not explicitly consider the time preferences of shippers. Chapter 4 shows the impact of time preferences in the context of point-to-point delivery using a simple home delivery model where customers were grouped into two segments, high priority, and low priority. Chapter 5 provides a framework for a home delivery system using a public logistics network. Consolidated loads in distribution centers are formed by shipper agents using a proportional cost allocation method and proposed to truck agents that accept or reject loads based on certain protocols. Two simple examples show how the loads will be formed and then accepted by trucks, with trucks nearest to the DC given first priority assuming it has not been assigned a load.

Computational experiments using shipments in North Carolina and South Carolina reveal important insights about the mechanism. In the short term, when a fixed set of shipments has to be consolidated, the Shapley cost allocation is the most equitable for allocating cost to each shipment that closely matches its impact on the overall cost of the consolidated load. In the longer term, the proportional scheme incentivizes more shipments to form consolidated loads and thus promotes efficiency due to the savings associated with increased consolidation; this is because when many potential shipments could be consolidated, it is not known apriori which shipment groups will lead to lower costs.

From the context of a centralized firm providing consolidation, with the same number of shipments for consolidation, a centralized mechanism is as good as or better than the public mechanism since it can give a lower total cost due to available tacit shipment information, like its willingness to pay. Coalitions formed by a centralized firm may not be incentive compatible for individual shipments since they do not require cost allocation, and thus, larger coalitions can form. But, there is a limit to the number of shipments that a centralized firm can control, beyond which the firm would dominate the market and risk government regulation. For a particular example consisting of 300 shipments, when a single firm controls the entire market, its total consolidation cost is 30% lower than that of a public mechanism. Since the more realistic scenario is that multiple firms will control portions of the market, when a firm controls one-sixth of the market, the total consolidation cost with six firms is about 10% higher than the cost under a public consolidation mechanism.

From the carrier's perspective, the assignment of trucks to loads has the highest profit when the carrier owns all trucks. With this upper bound, the greedy approach to assigning trucks to loads based on profit maximization and allowing trucks to renege loads among themselves causes the total profit to approach the upper bound when not all trucks are owned by the same carrier.

The natural way to extend the work on time preferences in Chapter 4 was to analyze their impact on shippers and carriers. But as an alternative, a model for home delivery using a public logistics network was developed in Chapter 5 which can be carried over nicely into an NSF proposal.

6.2 Future Work

The work done in this research can be extended in a number of ways. As mentioned in Chapter 2, the softmax distribution with a built-in threshold was used to propose consolidated loads to trucks since shippers can be a part of multiple loads. The softmax approach by itself is not part of the mechanism and was a way to demonstrate how loads can be proposed to trucks so that all proposed loads are acceptable to the participating shippers.

Instead of the softmax distribution, an agent-based approach can be used to develop the mechanism which can include shipment and truck agents. Different types of agents can be developed for shipments and carriers to allow a variety of different types of behaviors to utilize the mechanism. In particular, different strategies for forecasting will be available. Some shipment agents may try to predict load charges into the future while others may operate in a greedy manner; some carrier agents may be better at predicting future shipment demand at different locations and, as a result, use this information in selecting which shipments to accept. Another issue will be the possibility of a shipment specifying using different carrier capacities. Specifying less than a carrier's full capacity would allow the carrier to accept other loads utilizing the remaining capacity. This adds a level of complexity to the carrier agents. In order to evaluate these agents, multiple rounds of shipment demand will be simulated, and those shipment and carrier agents that perform the best will be used preferentially. Each shipment agent will be given funds proportional to its transport and time penalty, and those shipments not able to reach their destination before exhausting their funds will be eliminated. In a similar manner, each carrier will be assumed to have the same investment and operating cost, and those carrier agents not able to make enough revenue to cover their costs will be limited. In this manner, the best agents can be identified. The actual internal behavior of an agent will not evolve as part of this process. Also, the suitability of applying the consolidation mechanism to other modes of freight transport will be investigated. In most cases, the current application of the mechanism to trucking is the most complex and has the greatest potential, so its extension to other modes would likely not need to use the entire mechanism.

The time preferences of shippers and carriers will also play an important role in extending the mechanism. The models in Chapter 4 provide just one aspect of how time preferences can be modeled. However, extensions to the model can include analyzing the impact on overall cost if shippers relate their time preferences to inventory carrying costs. For example, a perishable or high-demand item will be highly time-sensitive and have high inventory costs since its shelf life will vanish quickly. Inventory costs can then be used to model deviations from single pick-up and delivery times for shipments.

The home delivery example can also be extended to include multiple pickup and drop points with pick-up and drop time preferences. Possible penalty measures can then be implemented and analyzed when a truck misses its pick-up and drop targets, and how those penalties affect the overall cost of the load.

REFERENCES

Aronsson, H. and Huge Brodin, M. (2006). The environmental impact of changing logistics structures. *The International Journal of Logistics Management*, 17(3):394–415.

Bartolini, E., Cordeau, J.-F., and Laporte, G. (2013). An exact algorithm for the capacitated arc routing problem with deadheading demand. *Operations Research*, 61(2):315–327.

Bolis, S. and Maggi, R. (2003). Logistics strategy and transport service choices: an adaptive stated preference experiment. *Growth and Change*, 34(4):490–504.

Chang, F.-S., Wu, J.-S., Lee, C.-N., and Shen, H.-C. (2014). Greedy-search-based multi-objective genetic algorithm for emergency logistics scheduling. *Expert Systems with Applications*, 41(6):2947–2956.

Cheong, M. L., Bhatnagar, R., and Graves, S. C. (2007). Logistics network design with supplier consolidation hubs and multiple shipment options. *Journal of Industrial and Management Optimization*, 3(1):51.

Chu, C.-W. (2005). A heuristic algorithm for the truckload and less-than-truckload problem. *European Journal of Operational Research*, 165(3):657–667.

Chu, J. C., Yan, S., and Huang, H.-J. (2017). A multi-trip split-delivery vehicle routing problem with time windows for inventory replenishment under stochastic travel times. *Networks and Spatial Economics*, 17:41–68.

Fan, M., Stallaert, J., and Whinston, A. B. (2003). Decentralized mechanism design for supply chain organizations using an auction market. *Information Systems Research*, 14(1):1–22.

Fatima, S. S., Wooldridge, M., and Jennings, N. R. (2008). A linear approximation method for the Shapley value. *Artificial Intelligence*, 172(14):1673–1699.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

Han, C.-K. and Cheng, S.-F. (2020). An exact single-agent task selection algorithm for the crowdsourced logistics. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence IJCAI 2020*.

Helguera Sanchez, I. and Hendra Mukti, L. P. (2018). Uberization effects on freight procurement. Master's thesis, Supply Chain Management, Massachusetts Institute of Technology, Massachusetts.

Higginson, J. and Bookbinder, J. H. (1994). Policy recommendations for a shipment-consolidation program. *Journal of Business Logistics*, 15(1).

Hochbaum, D. S. and Pathria, A. (1998). Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics (NRL)*, 45(6):615–627.

Kay, M. (2018). Public logistics networks. `https://mgkay.github.io/hdln/` [Accessed: 2024-11-01].

Kay, M. (2020). Challenges and opportunities associated with using autonomous vehicles and drones for home delivery. Personal Communication.

Kay, M. (2022). Public logistics networks for home delivery. `https://mgkay.github.io/hdln/` [Accessed: 2024-11-01].

Kay, M. (2023). Matlog: Logistics engineering matlab toolbox. `https://mgkay.github.io/matlog/` [Accessed: 2023-05-31].

Kay, M. (2024). Home delivery logistics networks using driverless delivery vehicles. `https://mgkay.github.io/hdln/` [Accessed: 2024-07-01].

Kay, M. and Parlikad, A. (2002). Material flow analysis of public logistics networks. *Meller R, Ogle MK, Peters BA, Taylor GD, Usher J, eds. Progress in Material Handling Research*, pages 205–218. Charlotte, NC: The Material Handling Institute.

Kay, M. G. (2004). Protocol design for a public logistic network. *Progress in Material Handling Research*, pages 181–188.

Kay, M. G. and Jain, A. (2002). Issues in agent-based coordination of public logistics networks. Technical report, Citeseer. `https://mgkay.github.io/pln/` [Accessed: 2023-05-31].

Kay, M. G., Karagul, K., Sahin, Y., and Aydemir, E. (2023). Shipment consolidation practice using matlog and large-scale data sets. In *Smart Applications with Advanced Machine Learning and Human-Centred Problem Design*, pages 701–717. Springer.

Kay, M. G., Karagul, K., Şahin, Y., and Gunduz, G. (2022). Minimizing total logistics cost for long-haul multi-stop truck transportation. *Transportation Research Record*, 2676(2):367–378.

Kay, M. G. and Warsing, D. P. (2008). Distributed shipment consolidation. *Progress in Material Handling Research*, pages 293–300.

Kay, M. G. and Warsing, D. P. (2009). Estimating LTL rates using publicly available empirical data. *International Journal of Logistics: Research and Applications*, 12(3):165–193.

Kearney (2022). State of logistics 2022: Executive summary. `https://www.kearney.com/documents/` [Accessed: 2023-05-31].

Kokubugata, H., Itoyama, H., and Kawashima, H. (1997). Vehicle routing methods for city logistics operations. *IFAC Proceedings Volumes*, 30(8):727–732.

Liu, T., Zhao, C., et al. (2019). Study the impacts of freight consolidation and truck sharing on freight mobility. Technical report, Transportation Consortium of South-Central States. `https://rosap.ntl.bts.gov/view/dot/42413/dot_42413_DS1.pdf` [Accessed: 2023-05-31].

Maskin, E. S. (2008). Mechanism design: How to implement social goals. *American Economic Review*, 98(3):567–576.

Meca, A., Garcia-Jurado, I., and Borm, P. (2003). Cooperation and competition in inventory games. *Mathematical Methods of Operations Research*, 57(3):481–493.

Meca, A., Timmer, J., Garcia-Jurado, I., and Borm, P. (2004). Inventory games. *European Journal of Operational Research*, 156(1):127–139.

Mishra, B. K., Dahal, K., and Pervez, Z. (2017). Dual-mode round-robin greedy search with fair factor algorithm for relief logistics scheduling. In *2017 4th International Conference on Information and Communication Technologies for Disaster Management (ICT-DM)*, pages 1–7. IEEE.

Mofidi, S. S. and Pazour, J. A. (2019). When is it beneficial to provide freelance suppliers with choice? A hierarchical approach for peer-to-peer logistics platforms. *Transportation Research Part B: Methodological*, 126:1–23.

Moulin, H. (1991). *Axioms of cooperative decision making*. Number 15. Cambridge university press.

Nehiba, C. (2020). Taxed to death? Freight truck collision externalities and diesel taxes. *Regional Science and Urban Economics*, 85:103577.

Nguyen, C., Dessouky, M., and Toriello, A. (2014). Consolidation strategies for the delivery of perishable products. *Transportation Research Part E: Logistics and Transportation Review*, 69:108–121.

Niederle, M. and Roth, A. E. (2009). Market culture: How rules governing exploding offers affect market performance. *American Economic Journal: Microeconomics*, 1(2):199–219.

Ozdamar, L. and Yi, W. (2008). Greedy neighborhood search for disaster relief and evacuation logistics. *IEEE Intelligent Systems*, 23(1):14–23.

Park, H., Jeong, S., and Lee, J.-S. (2023). Changes in patterns of home-based shopping and its last-mile delivery logistics services in Korea. `https://www.metrans.org/assets/research/` [Accessed: 2024-11-01].

Potvin, J.-Y. and Bengio, S. (1996). The vehicle routing problem with time windows Part II: Genetic Search. *INFORMS Journal on Computing*, 8(2):165–172.

Shams, K., Asgari, H., and Jin, X. (2017). Valuation of travel time reliability in freight transportation: A review and meta-analysis of stated preference studies. *Transportation Research Part A: Policy and Practice*, 102:228–243.

Sotomayor, M. (1999). Three remarks on the many-to-many stable matching problem. *Mathematical Social Sciences*, 38(1):55–70.

Sullivan, L. (2020). How big oil misled the public into believing plastic would be recycled. *National Public Radio (NPR)*. `https://www.wrongkindofgreen.org/tag/new-plastics-economy/` [Accessed: 2024-11-01].

Taniguchi, E. and Thompson, R. G. (2002). Modeling city logistics. *Transportation research record*, 1790(1):45–51.

Thanglah, S., Nygard, K., and Juell, P. (1991). A genetic algorithm system for vehicle routing with time windows. In *Miami. Proceedings of the Seventh Conference on Artificial Intelligence Applications*, page 323O325. Florida.

Trucks, C. (2023). Trucking industry trends, Statistics & Outlook for 2023. `https://www.cloudtrucks.com/blog-post/trucking-trends` [Accessed: 2023-05-31].

Tyan, J. C., Wang, F.-K., and Du, T. C. (2003). An evaluation of freight consolidation policies in global third party logistics. *Omega*, 31(1):55–62.

van Campen, T., Hamers, H., Husslage, B., and Lindelauf, R. (2018). A new approximation method for the Shapley value applied to the WTC 9/11 terrorist attack. *Social Network Analysis and Mining*, 8:1–12.

van der Spoel, S., Amrit, C., and van Hillegersberg, J. (2017). Predictive analytics for truck arrival time estimation: A field study at a European distribution centre. *International Journal of Production Research*, 55(17):5062–5078.

Varian, H. R. and MacKie-Mason, J. K. (1993). Pricing the internet. *Public Access to the Internet*.

Wei, H., Huang, M., and Chen, Y. (2019). One-to-one logistics-client stable matching. In *2019 Chinese Control And Decision Conference (CCDC)*, pages 965–970. IEEE.

APPENDIX

APPENDIX

A

SHIPMENTS USED IN TABLE 2.7 AND TABLE 2.8

The following figures show the shipments along with their longitudes (x-axis) and latitudes (y-axis) of their pick-up and drop locations used to compute the lower bound of the cost of the mechanism under different allocation schemes. All shipments are located in North Carolina and South Carolina. TL is truck-load and LTL is less-than-truckload.

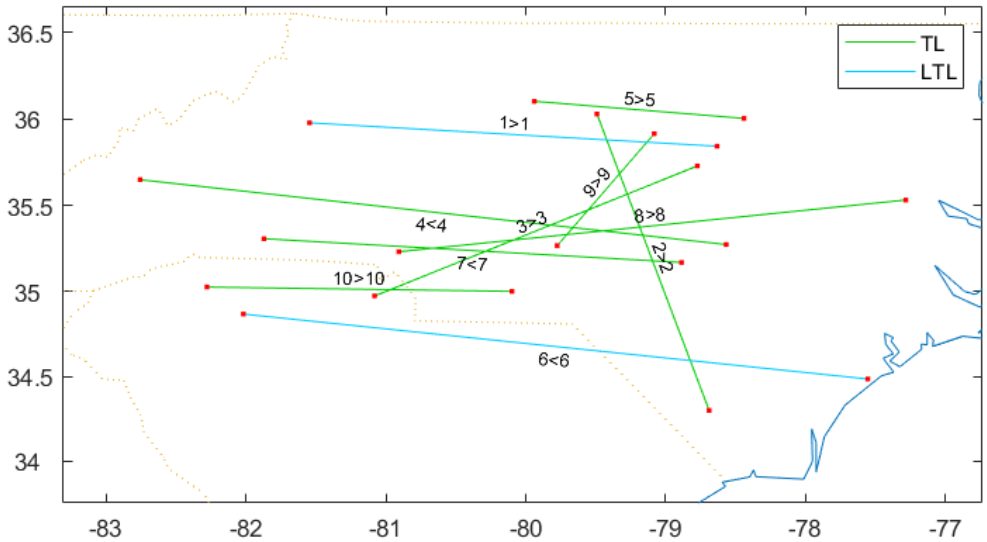

Figure A.1: Ten Shipment Batch 1.

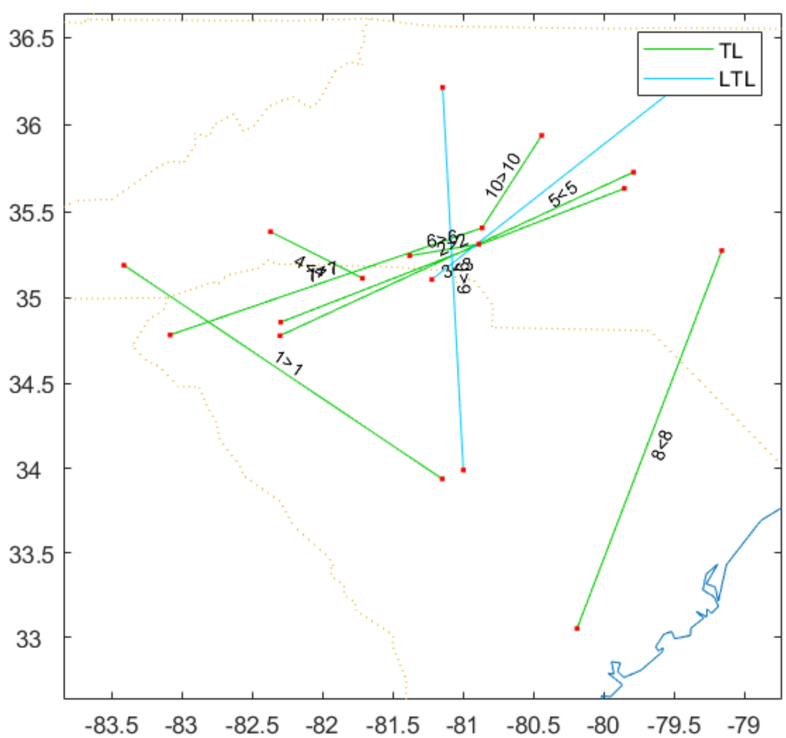

Figure A.2: Ten Shipment Batch 2.

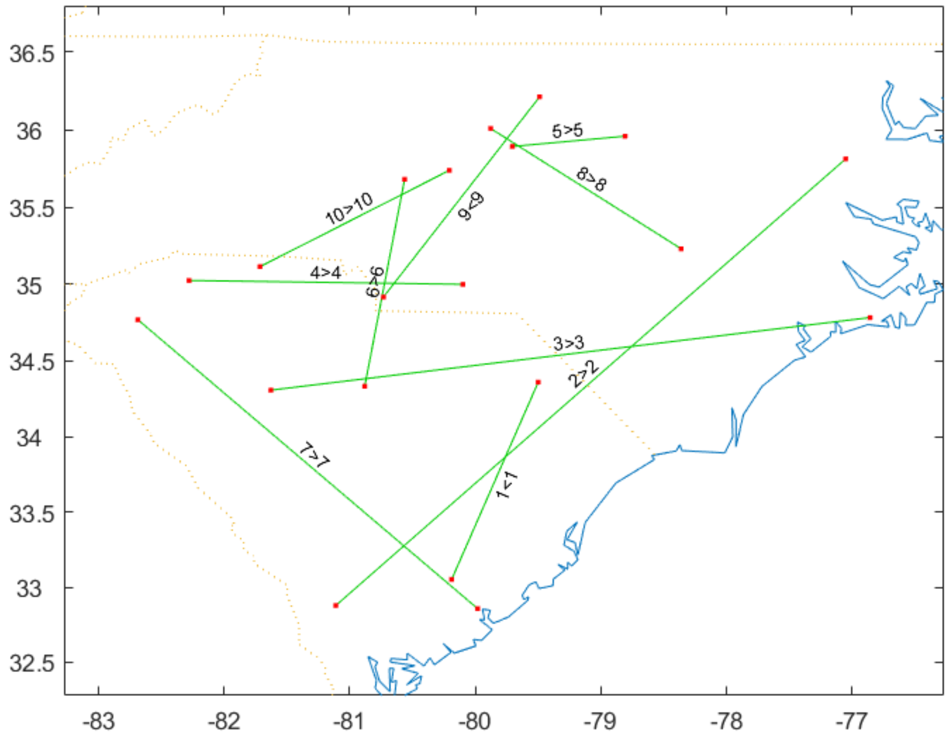

Figure A.3: Ten Shipment Batch 3.

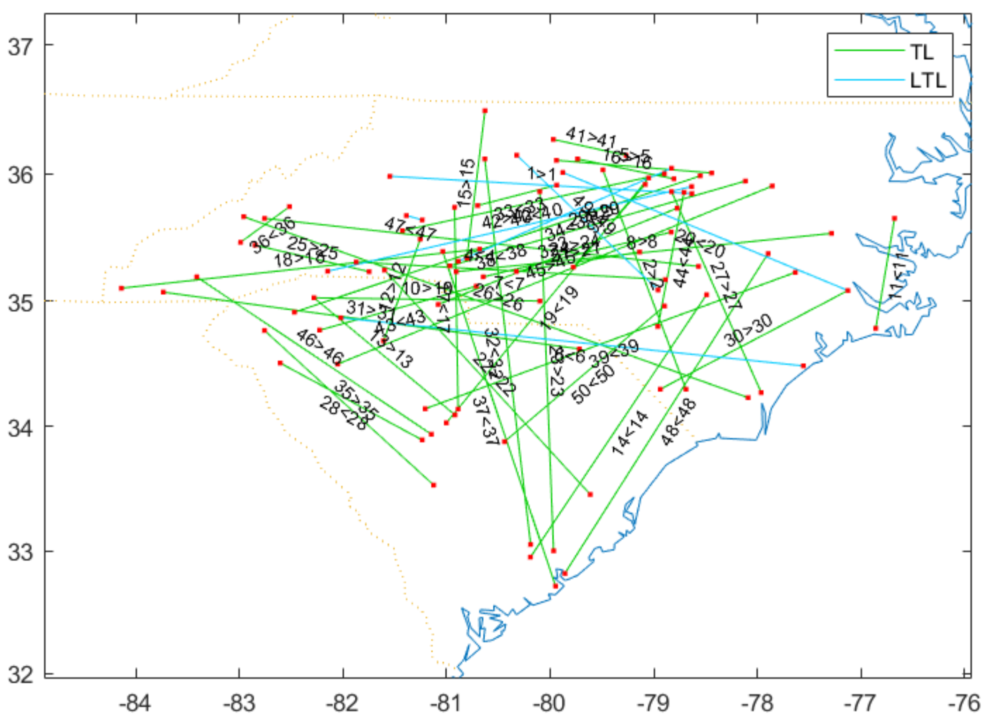

Figure A.4: Fifty Shipment Batch 1.

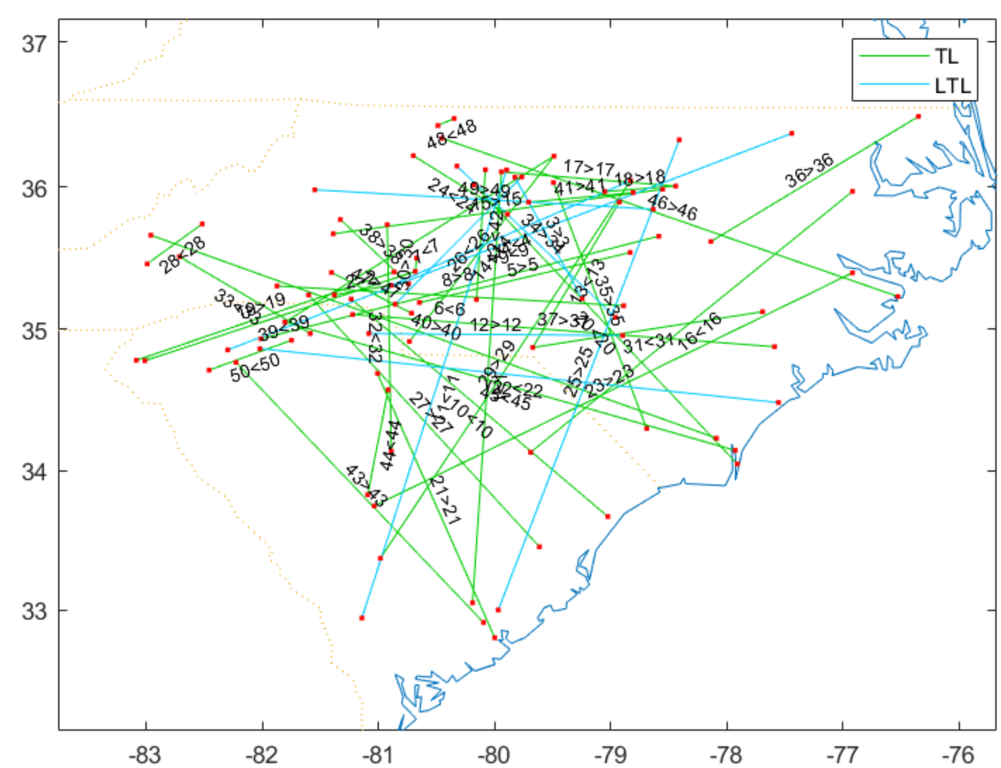

Figure A.5: Fifty Shipment Batch 2.

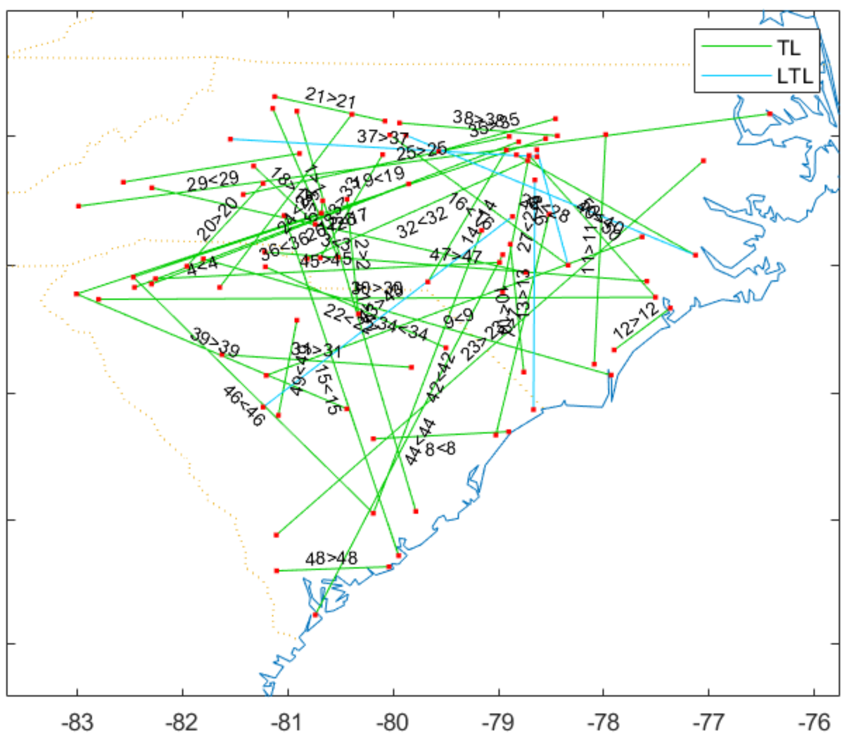

Figure A.6: Fifty Shipment Batch 3.