

Single-Warehouse Multi-Retailer Inventory Systems with Full TruckLoad Shipments*

Yue Jin
Ana Muriel[†]

June 2005; revised March 2006 and March 2008

Abstract

We consider a multistage inventory system composed of a single warehouse that receives a single product from a single supplier and replenishes the inventory of n retailers through direct shipments. Fixed costs are incurred for each truck dispatched and all trucks have the same capacity limit. Costs are stationary, or more generally monotone as in Lippman (1969). Demands for the n retailers over a planning horizon of T periods are given. The objective is to find the shipment quantities over the planning horizon to satisfy all demands at minimum system-wide inventory and transportation costs without backlogging. Using the structural properties of optimal solutions, we develop (1) an $O(T^2)$ algorithm for the single-stage dynamic lot sizing problem; (2) an $O(T^3)$ algorithm for the case of a one-warehouse single-retailer system; and (3) a nested shortest-path algorithm for the one-warehouse multi-retailer problem that runs in polynomial time for a given number of retailers. To overcome the computational burden when the number of retailers is large, we propose aggregated and disaggregated Lagrangian decomposition methods that make use of the structural properties and the efficient single-stage algorithm. Computational experiments show the effectiveness of these algorithms and the gains associated with coordinated versus decentralized systems. Finally, we show that the decentralized solution is asymptotically optimal.

1 Introduction

The most common mode of transportation in industry applications is the full truckload mode. Large consumer product companies, such as Kimberly-Clark, Wal-Mart and Procter&Gamble, use 53 footers almost exclusively to move goods through their distribution systems. Some companies use their own fleet of vehicles, others contract out to outside

*Research supported by NSF Contract DMI-9732795.

[†]Department of Mechanical and Industrial Engineering, University of Massachusetts, Amherst, MA

providers. In either case, significant savings can be achieved by coordinating inventory and shipment decisions across the entire system to facilitate load consolidation.

The *Single-Warehouse Multi-Retailer Problem with Full Truckload (FTL) Shipments* can be stated as follows: A number of retail facilities faces known demands of a single product over a finite planning horizon. They order goods from a warehouse whose inventory is in turn replenished by an external supplier. All shipments from warehouse to retailers are direct; that is, trucks travel directly from the warehouse to a single retailer and back, see Gallego and Simchi-Levi (1990) [12]. There is no limit on the quantity ordered each period, but there are cargo constraints that require additional trucks to be dispatched when exceeded. There is a fixed cost per truck dispatched from supplier to warehouse and from warehouse to retailers, and linear holding costs at the warehouse and retailers. The objective is to decide when and how many units to ship from supplier to warehouse and from warehouse to retailers each period so as to minimize total transportation and holding costs over the finite horizon without any shortages. We consider the administrative ordering setup costs to be negligible relative to the fixed costs of dispatching a truck. In certain business environments, the full demand for a particular retailer and time period, d_t^i , may be required to travel together from supplier to warehouse and then from warehouse to retailer. We refer to this more restrictive version of the problem as the Non-Splitting Single-Warehouse Multi-Retailer Problem. In the general Single-Warehouse Multi-Retailer Problem, however, demands are allowed to be split over two or more shipments to avoid sending additional trucks.

In this paper, we consider the Single-Warehouse Multi-Retailer Problem with FTL Shipments and stationary costs; i.e., the fixed and variable transportation charges and the linear holding costs do not change over time. As a result, linear transportation costs do not affect the optimal shipping strategy and are thus ignored. We also assume, as is the case in much of the multistage inventory literature, that holding costs are no higher at the warehouse than at any of the retail locations. All the properties and dynamic programming algorithms presented are also valid in the case of monotone shipping cost functions introduced by Lippman (1969)[21].

The Single-Warehouse Multi-Retailer Problem with fixed charge costs (i.e., unlimited truck capacity) and stationary costs (*SWMRP*) has been shown to be NP-complete by Arkin,

Joneja and Roundy (1989) [5]. The *Joint Replenishment Problem (JRP)*, see also Joneja (1990) [18], can be modeled as a special case of the *SWMRP* by making the holding costs at the warehouse identical to those at the retailers. In the *JRP*, a single facility replenishes a set of items over a finite horizon. Whenever the facility places an order for a subset of the items, two types of costs are incurred: A joint set-up cost and an item-dependent set-up cost. These costs are stationary. The objective in the joint replenishment problem is to decide when and how many units to order for each item so as to minimize inventory holding and ordering costs over the planning horizon. Since the joint replenishment problem is NP-hard, the Single-Warehouse Multi-Retailer problem under consideration in this paper is also NP-hard even if all costs are stationary, inventory costs at the warehouse are required to be no higher than those at the retailers, and all transportation cost functions are fixed charge cost functions.

The complexity of optimizing the discontinuous step functions, also referred to as staircase or multiple setup cost functions, associated with Full TruckLoad (FTL) transportation has slowed research in this area. Other cost functions, such as fixed-charge¹ (see Gendron, Crainic and Frangioni (1999) [13] for a review), incremental discount (e.g. Muriel and Simchi-Levi (2004) [24], Balakrishnan and Graves (1989) [6], Amiri and Pirkul (1997) [2]) or modified all unit discount (Chan et al. (2002a, 2002b) [8] [9]), have been more widely studied.

A number of papers do explicitly consider staircase functions. In what follows, we first review the literature tackling single-stage and then multiple-stage systems.

For the basic single-stage dynamic lot sizing problem with multiple setups and stationary (or more generally monotone) costs, Lippman (1969) [21] shows that there is an optimal solution such that (1) no partially filled trucks are shipped in periods with positive initial inventory, and (2) the inventory in each period is less than the truck capacity. These two properties have been the cornerstone for much of the posterior research, including the present work. Lippman develops an $O(T^3)$ dynamic programming algorithm. In Section 3, we propose a backwards-recursion dynamic programming algorithm that reduces the computational complexity to $O(T^2)$. For general time-varying costs, Pochet and Wolsey (1993) use extreme

¹Observe that the capacitated fixed-charge network design problem would generalize the single-warehouse multi-retailer problem if parallel arcs with capacity equal to truck capacity are considered.

flow arguments to show that there is at most one partially filled truck between two regeneration points (points with zero inventory), and develop a forward dynamic programming algorithm that runs in $O(T^2 \min(T, W))$, where T is the number of periods and W the batch size. When applied to the stationary case, its complexity remains the same.

Lipmann (1969) [21] also considers a more general problem where multiple truck sizes to choose from are available in a single period, and an infinite horizon stationary model. Alp, Erkip and Güllü (2003) [1] characterize optimal policies and develop a dynamic programming algorithm for the problem with stochastic lead times. Li, Hsu and Xiao (2004) [20] consider a very general model that allows for demand backlogging and includes time-varying fixed and practical transportation-related ordering costs representing both fixed costs per full truck dispatched and linear costs associated with Less-than-Truckload (LTL) shipments. They develop a $O(T^3 \log T)$ algorithm based on a related dynamic lot-sizing model with batch ordering. Anily and Tzur (2004, 2005) [3] [4] consider multiple products to be delivered from warehouse to retailer in capacitated vehicles, each incurring a fixed cost per trip, and propose both a dynamic programming algorithm [3] and a search algorithm [4] to solve the problem optimally.

In recent years, the focus on integrated logistics management has led to the increased study of systems with multiple stages. The model recently studied by Lee, Çetinkaya and Jaruphongsa (2003) [23], focuses on the coordination of inventory replenishments and dispatch schedules at a warehouse that serves a single retailer. The warehouse orders incur a fixed cost and the outbound transportation cost function consists of a fixed cost per delivery plus a cost per vehicle dispatched. Jaruphongsa, Çetinkaya and Lee (2005) [16] consider a similar model with two available outbound shipment modes: one with a fixed setup cost structure and the other with a multiple setup cost structure. Diaby and Martel (1993) [11] develop a Lagrangian relaxation based procedure to solve a more general problem for multi-echelon distribution systems (each facility has a single predecessor) with general piece-wise linear ordering and transportation cost functions.

The Non-Splitting Single-Warehouse Multi-Retailer Problem has been addressed in Levi, Roundy and Shmoys (2005) [19] who develop constant approximation algorithms for the problem with fixed-charge ordering costs and later extend them to accommodate the multiple

setup cost structure. This results in a 4.796-approximation algorithm for the problem under consideration in this paper. Their LP-rounding approach constructs non-splitting solutions where the entire demand of a retailer at a particular time period (d_t^i) must travel together (i.e., in the same time period) from supplier to warehouse and then again together from warehouse to retailer. More general piece-wise linear transportation costs, which include both FTL and LTL (Less than TruckLoad) realistic cost functions, have been considered in Croxton, Gendron and Magnanti (2003) [10] to model the selection of different transportation modes and shipment routes in merge-in-transit operations, also under the assumption of non-splitting shipments. In this case, a set of warehouses coordinates the flow of goods from a number of suppliers to multiple retailers with the objective of reducing costs through consolidation.

Staircase or step cost functions have also been considered in facility location applications; see Holmberg (1994) [14] and Holmberg and Ling (1997) [15].

One-Warehouse Multi-Retailer systems with stationary fixed charge costs and constant demand over an infinite horizon have been extensively studied. The seminal work of Schwarz (1973) [28] and Roundy (1985) [27] analyzes the problem with fixed ordering costs at both the warehouse and retailer locations. Schwarz (1973) [28] characterizes the properties of optimal solutions: retailers only order when their inventory is down to zero and the warehouse only orders when both its inventory and that of one of the retailers is down to zero. For the more complex dynamic problem with cargo constraints, these results hold for partial shipments; see Property 2.4 and Property 4.1. Roundy (1985) shows that Power-of-Two policies are highly effective (within 2%). Lu and Posner (1994) [22] present approximation algorithms that further improve the quality of the solutions.

The paper is organized as follows. Section 2 describes the model under consideration and presents the main structural properties of optimal solutions. In Sections 3 and 4, we develop exact algorithms for the one-warehouse multi-retailer system under decentralized and centralized management, respectively. Under decentralized management of the system, each member makes their own self-optimizing decisions and thus solves a single-stage problem. For that purpose, we develop an algorithm for the single-stage dynamic lot sizing problem with stationary costs with complexity $O(T^2)$. In Section 4.1, we develop a $O(T^3)$ algorithm

for the two-stage, single-retailer case. This algorithm is then generalized to any number of retailers in Section 4.2. Due to the exponential growth of the complexity of the algorithm as the number of retailers increases, Section 5 introduces alternative algorithms based on Lagrangian decomposition that make use of the structural properties of optimal solutions and the efficient single-stage algorithm to solve large-scale problems effectively. Finally, we demonstrate the effectiveness of the heuristic Lagrangian-based algorithms and compare the performance of centralized versus decentralized management of the system through computational experiments. The outstanding performance of the decentralized solution as the number of retailers grows then lead us to show, in Section 7, that it is indeed asymptotically optimal.

2 Model

Consider the single-warehouse multi-retailer system with full truckload shipments described above. Let T be the time horizon over which demands from n retailers are known, and let the demand of retailer i at time t be d_t^i , $i = 1, 2, \dots, n$, $t = 1, 2, \dots, T$. All demand must be satisfied without backorders at the end of each period. We assume that the transportation and inventory cost parameters are stationary, with A^0 denoting the fixed cost of dispatching a truck from supplier to warehouse, A^i the cost of dispatching a truck from the warehouse to retailer i , and h^i the inventory cost per unit left over in inventory at retailer i at the end of each period. Inventory can be carried at the warehouse as well, at a rate h^0 , $h^0 \leq h^i$ for all i . All trucks are identical with capacity of W units.

The optimal solution will be determined by the quantities x_t^0 and x_t^i , for $t = 1, 2, \dots, T$ and $i = 1, \dots, n$, to ship from supplier to warehouse and warehouse to retailer i , respectively. For simplicity we will denote a solution vector by x , $x = (x^0, x^1, \dots, x^n)$ and $x^i = (x_1^i, x_2^i, \dots, x_T^i)$. We denote the resulting inventory at the beginning of period t at the warehouse by I_t^0 and at retailer i by I_t^i , $t = 1, 2, \dots, T + 1$. To simplify the exposition of the algorithms, we assume w.l.o.g. that the initial inventory at warehouse and retailers is zero; i.e., $I_1^i = 0$ for $i = 0, 1, \dots, n$. The extension to positive initial inventories at the retailers is straightforward, by reducing the retailer demand in the initial period(s). As we

shall see, the structural properties and thus the resulting algorithms are easily extended for positive initial inventories at the warehouse.

Let a *warehouse (retailer) regeneration point* be a period where initial inventory at the warehouse (retailer) is zero. A *warehouse (retailer) LTL period* is a period in which a partial, less than full, truckload is shipped from supplier to warehouse (warehouse to retailer). In the remainder of this section, we characterize the relationship between regeneration points and LTL periods in optimal solutions to the Single-Warehouse Multi-Retailer Problem. We first present the basic properties of the optimal solutions to the single-stage problem that have been presented in the literature and then we generalize them to the one-warehouse multi-retailer setting. These properties are the foundation for the algorithms developed in the paper under both centralized and decentralized management of the system.

2.1 Single-Stage Problem: Known Results

For the general single-stage economic lot sizing problem with multiple setups and non-stationary (fixed plus linear) costs, Pochet and Wolsey (1993) show the following property using extreme flow arguments.

Proposition 2.1 *Between two consecutive regeneration points, there is at most one LTL period.*

Lippman (1969) had earlier showed this and further properties, Property 2.2 and Property 2.3 below, under the assumption of a “monotone cost model,” which generalizes our stationary cost assumptions. A monotone cost model is characterized by shipping cost functions, $c_t(x)$, for $t = 1, 2, \dots, T$, that satisfy the following three conditions. Let $W_t > 0$ be the capacity of the truck at time t , $t = 1, 2, \dots, T$, and $c_t^*(\cdot)$ a non-negative, non-decreasing, concave function on $[0, W_t]$ with $c_t^*(0) = 0$. Then $c_t(x) = \frac{x - x \bmod W_t}{W_t} c_t^*(W_t) + c_t^*(x \bmod W_t)$, where $x \bmod W_t$ is the unique number r such that $0 \leq r \leq W_t$ and $x = r + kW_t$ for some integer k . Observe that $c_t^*(W_t)$ represents the cost associated with dispatching a full truck. Let A_t be the magnitude of the jump at zero, that is, $A_t = c_t^*(0+)$.

1. $c_t^*(u+\epsilon) - c_t^*(u) \geq c_{t+1}^*(v+\epsilon) - c_{t+1}^*(v)$, for all $0 < u \leq u+\epsilon \leq W_t$, $0 < v \leq v+\epsilon \leq W_{t+1}$, and $t = 1, 2, \dots, T-1$;

2. $A_t \geq A_{t+1}$, for all $t = 1, 2, \dots, T-1$;
3. $M_t \leq M_{t+1}$, for all $t = 1, 2, \dots, T-1$.

Inventory costs, $h_t(\cdot)$ are only required to be nondecreasing, left continuous functions of the inventory at the end of period t . In what follows, we use the same notation presented in the previous section, but omit the superindex indicating the facility and add the time subindex when needed.

Property 2.2 *Inventory in each period is less than one cargo capacity. That is $I_t < W_t$ for $t = 1, \dots, T$.*

Property 2.3 *A partial shipment, $0 < x_t < W_t$ is dispatched only in periods where initial inventory is zero. That is, $I_t(x_t \bmod W_t) = 0$.*

Consequently, between two consecutive regeneration points there is at most one LTL period, as later shown for the more general² case by Pochet and Wolsey. Furthermore, if there is one, it must be the first of the regeneration points.

2.2 Single-Warehouse Multi-Retailer Problem: Extended Properties

Pochet and Wolsey discuss the extension of the concept of regeneration points and their extreme flow arguments to directed subtrees: “based on the structure of the extreme flows there is at most one node in each regeneration subtree in which production is below capacity, and thus it is possible to find the minimum cost solution for each subtree in polynomial time.” Consequently, Property 2.1 above holds for the Single-warehouse Multi-retailer Problem with Full Truckload non-stationary costs as stated in the following property.

Property 2.4 *Between two consecutive warehouse or retailer regeneration points there is at most one LTL period.*

²Pochet and Wolsey (1993) consider non-stationary fixed plus linear shipping costs and linear inventory costs, and stationary capacity. The monotone cost functions (Lippman (1969)) place more restrictions on how costs change over time, but consider more general fixed plus concave shipping costs and general non-decreasing inventory costs, and allow the capacity to possibly increase from period to period.

Now, for the more restrictive case of stationary, or more generally monotone costs as defined above, we can easily extend the properties in Lippman (1969) to the single-warehouse multi-retailer scenario. In using the concept of monotone costs in a multi-facility environment, we require that the truck capacity be the same for all shipping links in the same period, while it can still increase from period to period, and that holding goods at the warehouse in a particular period be always no higher than holding them at the retailer.

Property 2.5 *Inventory at the warehouse and each of the retailers in each period is less than one cargo capacity. That is $I_t^i < W_t^i$ for all $i = 0, 1, 2, \dots, n$, and $t = 1, \dots, T$.*

This is true since otherwise a full truckload shipment could be delayed without incurring any additional shipping costs while saving holding costs.

Property 2.6 *A partial shipment, $0 < x_t^i < W_t^i$, $i = 0, 1, \dots, n$ is dispatched only in periods where initial inventory is zero. That is, $I_t^i(x_t^i \bmod W_t^i) = 0$.*

This is true since otherwise the shipment in that period should be increased and the previous shipment reduced to reduce inventory holding costs and possibly save shipping costs (given the monotone cost structure). Therefore, between two consecutive warehouse (retailer) regeneration points there is at most one LTL period, as shown for the non-stationary case by Pochet and Wolsey. Furthermore, if there is one, it must be the first of the regeneration points. We must note that if initial inventory is positive, $I_0^i > 0$, the first period in the horizon behaves as a first regeneration point and can have a partial shipment.

For simplicity in what follows we consider the case of stationary fixed plus linear shipping costs, stationary and linear inventory costs, and identical truck capacities throughout the network and over time, as described in the model introduction. The dynamic programming algorithms can be easily extended to Lippman's monotone cost model, since all the properties hold in that more general case.

Without loss of generality, we assume that the demand at each retailer in each period is less than a full truckload. Otherwise, an optimal solution would send the full truckload(s) directly from supplier to warehouse to retailer in that period and coordinate the remaining less than truckload demands.

Finally, we let $\lfloor a \rfloor$ be the maximum integer less than or equal to a and $\lceil a \rceil$ be the minimum integer greater than or equal to a .

3 Decentralized System

The single-warehouse multi-retailer system could be managed in a decentralized fashion, where each of the individual members makes its own decisions based on its local demands and costs. In this setting, retailer i observes its demands d_t^i , for $t = 1, 2, \dots, T$, and minimizes its total transportation and inventory costs, which are composed of a fixed cost A_i per vehicle dispatched and a linear holding cost of h_i per unit left over at the end of each period. The retailer then places its cost-minimizing orders, x_t^i , $t = 1, 2, \dots, T$, to the warehouse. As a result, the warehouse faces demands $d_t^0 = \sum_{i=1}^n x_t^i$. Given transportation costs of A_0 per truck dispatched and a linear holding cost h_0 , the warehouse finds its corresponding cost-minimizing ordering quantities. Observe that we can assume that $d_t^0 < W$ in solving the problem, as we did for retailer demands, since otherwise, if $d_t^0 \geq W$, it is optimal to send a full truckload in period t and consider an equivalent problem with $d_t^0 \leftarrow d_t^0 - W$.

Thus, the *Decentralized Problem* at facility i , $i = 0, 1, \dots, n$, can be written as:

$$\begin{aligned}
 \text{Problem } DP_i : \quad & \text{Min} \quad \sum_{t=1}^T (A_i y_t^i + h_i I_{t+1}^i) \\
 & \text{s.t.} \\
 & x_t^i \leq W y_t^i, \quad \forall t = 1, 2, \dots, T, \\
 & x_t^i + I_t^i = d_t^i + I_{t+1}^i, \quad \forall t = 1, 2, \dots, T, \\
 & I_1^i = 0, \\
 & x_t^i \geq 0, \quad \forall t = 1, 2, \dots, T, \\
 & y_t^i \in \{0, 1\}, \quad \forall t = 1, 2, \dots, T, \\
 & I_t^i \geq 0, \quad \forall t = 1, 2, \dots, T,
 \end{aligned} \tag{1}$$

As in Lippman (1969) and Pochet and Wolsey (1993), the single-stage Decentralized Problem at facility i , $i = 0, 1, \dots, n$, can be modeled as a shortest path problem from node

1 to node $T + 1$ on a network with nodes $1, 2, \dots, T + 1$ and arcs (u, v) , $1 \leq u < v \leq T + 1$, representing two consecutive regeneration points. The length of arc (u, v) , which we denote by L_{uv}^i , is the minimum transportation and holding cost at facility i associated with covering all of its demands between periods u and $v - 1$ without shortages, given that u and v are consecutive regeneration points (that is, given that $I_u^i = I_v^i = 0$ and $I_t^i > 0$ for all $u < t < v$). We denote the problem of calculating the length of arc (u, v) as the *Decentralized Subproblem* DS_{uv}^i .

The following proposition and the subsequent algorithm show that all the arc lengths in the network can be calculated in time $O(T^2)$. Consequently, the single-stage Decentralized Problem at facility i , $i = 0, 1, \dots, n$, can be solved in $O(T^2)$ and the system-wide solution under decentralized management can be found in $O(nT^2)$. This is in contrast to the algorithms of Lippman (1969) and Pochet and Wolsey (1993), which run in time $O(T^3)$ and $O(T^2 \min(T, W))$, respectively. Note that our algorithm is valid under the more general monotone cost assumption in Lippman (1969) and that the complexity of the forward dynamic programming algorithm by Pochet and Wolsey (1993) remains the same under stationary costs.

Proposition 3.1 *Given two consecutive regeneration points u and v , the optimal transportation and inventory quantities for any period t between them, $u < t < v$, can be determined independent of the exact timing of u . That is, the optimal quantities and costs in period t are identical for problems DS_{uv}^i and $DS_{u-k,v}^i$ for any $0 < k < u$.*

Proof. As long as there are no regeneration points between t and v , the optimal action is to send full trucks to the retailer as late as possible while ensuring that (1) the final inventory is $I_v^i = 0$, and (2) the less-than-truckload quantity required to meet demand without backlogging must be carried in inventory from previous periods. ■

This proposition allows us to solve the Decentralized Subproblems efficiently. In particular, for each period v , $1 < v \leq T + 1$, the following algorithm calculates the costs on arcs (u, v) for all $1 \leq u < v$ in time $O(T)$.

Algorithm to solve Decentralized Subproblems DS_{uv}^i :

Let \tilde{L}_{tv}^i denote the transportation and holding cost between period t ($u < t < v$) and $v - 1$ in the optimal solution for arc (u, v) , given that u and v are consecutive regeneration points. Using a backwards recursion for $t = v - 1, \dots, 1$, we can calculate the costs between periods t and $v - 1$ simultaneously for both the case where t is not a regeneration point, i.e., \tilde{L}_{tv}^i , and the case where $t = u$ is a regeneration point, i.e., L_{uv}^i , as follows:

1. Initialize $\tilde{L}_{vv}^i = 0$, $I_v^i = 0$ and $L_{v-1v}^i = A_i \lceil \frac{d_{v-1}^i}{W} \rceil$.
2. Recursively, backwards in time for periods $u = v - 2, v - 3, \dots, 1$ and $t = u + 1$:
 - (a) Calculate \tilde{L}_{tv}^i and I_t^i .
 - If $d_t^i + I_{t+1}^i > W$, then $x_t^i = W$.
A full truck has to be dispatched to retailer i since t is not an LTL period and the inventory carried in any period is below cargo capacity (Observation 2.5).
 - Otherwise, $x_t^i = 0$.
That is, $d_t^i + I_{t+1}^i$ must be entirely covered by the initial inventory I_t^i and no shipment is needed.

Thus,

$$x_t^i = \lfloor \frac{d_t^i + I_{t+1}^i}{W} \rfloor W, \quad I_t^i = I_{t+1}^i + d_t^i - x_t^i, \quad \text{and} \quad \tilde{L}_{tv}^i = A_i \frac{x_t^i}{W} + h_i I_t^i + \tilde{L}_{t+1,v}^i$$

- (b) Calculate L_{uv}^i . Since u is a regeneration point and $t = u + 1$,

$$I_u^i = 0, \quad x_u^i = d_u^i + I_t^i, \quad \text{and} \quad L_{uv}^i = A_i \lceil \frac{x_u^i}{W} \rceil + \tilde{L}_{tv}^i$$

If the sum $d_u^i + I_{u+1}^i$ exceeds a full truckload W , periods u and v cannot be consecutive regeneration points in the overall optimal solution: a lower cost solution can be constructed by sending $x_u^i = d_u^i$ and shipping a partial truckload with I_{u+1}^i when needed. This saves inventory costs without dispatching any more trucks, but adds an intermediate regeneration point $u + 1$, contradicting the initial assumption that u and v are consecutive regeneration points. The associated arc (u, v) can thus be removed from the shortest path network.

4 Centralized System

We now consider the case where the single-warehouse multi-retailer system is managed by a centralized decision maker whose objective is to minimize system-wide transportation and inventory costs over the planning horizon. The *Centralized Single-Warehouse Multi-Retailer Problem*, Problem CP, can be written as follows.

$$\begin{aligned}
 \text{Problem CP :} \quad & \text{Min} \quad \sum_{t=1}^T \sum_{i=0}^n (A_i y_t^i + h_i I_{t+1}^i) \\
 & \text{s.t.} \\
 & x_t^i \leq W y_t^i, \quad \forall t = 1, 2, \dots, T, i = 0, 1, \dots, n, \\
 & x_t^i + I_t^i = d_t^i + I_{t+1}^i, \quad \forall t = 1, 2, \dots, T, i = 1, 2, \dots, n, \\
 & x_t^0 + I_t^0 = \sum_{i=1}^n x_t^i + I_{t+1}^0, \quad \forall t = 1, 2, \dots, T, \\
 & I_1^i = 0, i = 0, 1, \dots, n, \\
 & x_t^i \geq 0, \quad \forall t = 1, 2, \dots, T, i = 0, 1, \dots, n, \\
 & y_t^i \in \{0, 1\}, \quad \forall t = 1, 2, \dots, T, i = 1, \dots, n, \\
 & 0 \leq y_t^0 \leq n, \quad \text{integer} \quad \forall t = 1, 2, \dots, T, \\
 & I_t^i \geq 0, \quad \forall t = 1, 2, \dots, T, i = 0, 1, \dots, n,
 \end{aligned} \tag{2}$$

Observe that the number of shipments to the warehouse in any period t is bounded by the number of retailers, n , since in period t at most one truckload shipment will be sent to each retailer (recall that we have assumed w.l.o.g. that the demand at each retailer in each time period is less than a truckload).

As mentioned in the introduction, this problem is NP-hard. However, the following section shows that the single-retailer problem can be solved in polynomial time $O(T^3)$. The exact algorithms for the centralized system developed in the next two sections rely on the fact that only regeneration points can be LTL periods and thus they are the only ones that need to be coordinated, since full truckloads are shipped on the same period from supplier to warehouse to retailer.

4.1 Single-Retailer System

The single-retailer problem has additional properties that we can exploit in the development of an exact algorithm.

Property 4.1 *In the optimal solution, if one period is a warehouse LTL period, it must be a retailer LTL period. That is, if $0 < x_t^0 < W$, then $0 < x_t^1 < W$.*

This is easy to show by contradiction; since the LTL period must be a warehouse regeneration period, if the retailer shipment is either 0 or a full truckload, then the partial shipment to the warehouse can and must be postponed to reduce inventory costs.

Property 4.2 *In the optimal solution, if one period is a warehouse LTL period, it must be a system regeneration point, i.e. a period in which the initial inventories at both warehouse and retailer are zero, or the first period in the horizon.*

Property 4.3 *Between two consecutive system regeneration points, there is at most one warehouse LTL period. If there is one, it must be the first period between the two consecutive system regeneration points.*

Again, by the same argument, there is at most one warehouse LTL period between the first period in the horizon and the following system regeneration point. Furthermore, the LTL period, if there is one, must be period 1.

The Single-Warehouse Single-Retailer problem can be modelled as a shortest path problem in a network with $T + 1$ nodes, indexed $1, 2, \dots, T + 1$, and arcs (s, l) for each $1 \leq s < l \leq T + 1$. The cost of an arc from period s to l , C_{sl} , for all $1 \leq s < l \leq T + 1$, is the optimal cost to cover the demands from periods s to $l - 1$ assuming both s and l are *system* regeneration points. The shortest path from node 1 to node $T + 1$ provides the optimal solution to the Single-Warehouse Single-Retailer problem.

Given the lengths of all arcs, the shortest path can be found in time $O(T^2)$. The only issue remaining is how to calculate the cost of each arc, C_{sl} .

Since s and l are two consecutive system regeneration points, we know that $I_s^0 = I_l^0 = 0$ and the quantity shipped to the warehouse in period s , the only possible LTL period, must

be,

$$\Delta_{sl}^0 = \sum_{t=s}^{l-1} d_t^1 - \lfloor \frac{\sum_{t=s+1}^{l-1} d_t^1}{W} \rfloor W.$$

Between the two consecutive system regeneration points there may be several retailer regeneration points. Let u be a retailer regeneration point between system regeneration points s and l . The inventory at the warehouse at the beginning of period u , $s < u < l$, is

$$I_u^0 = \sum_{t=u}^{l-1} d_t^1 - \lfloor \frac{\sum_{t=u}^{l-1} d_t^1}{W} \rfloor W.$$

Therefore, the initial warehouse inventory at a retailer regeneration point u within consecutive system regeneration points does not depend on the timing of the initial system regeneration point, s . To reflect the dependence on the second system regeneration point, l , we will denote it by $I_u^0(l)$.

We calculate the cost, C_{sl} , associated with each pair of consecutive system regeneration points s and l , as a shortest path on a network with nodes $s, s+1, \dots, l$. An arc from node u to node v , $s \leq u < v \leq l$, represents the optimal system ordering policy to cover the demands from period u to period $v-1$, given that s and l are two consecutive system regeneration points, and u and v are two consecutive retailer regeneration points. The length of the arc is the minimum cost, which we denote by $F_{sl}(u, v)$. Since the initial warehouse inventory $I_u^0(l)$ does not depend on s for $s < u < v \leq l$, the value of $F_{sl}(u, v)$ remains the same for all $s < u$. Using this property, we develop an exact algorithm for the Single-Warehouse Single-Retailer problem that runs in time $O(T^3)$.

4.1.1 Single-Retailer Algorithm

Step 1: For all u and v , $1 \leq u < v \leq T+1$, solve a Decentralized Subproblem DS_{uv}^1 (see Section 3) and let $x_t^1(u, v)$, for $u \leq t < v$, be the optimal replenishment quantities and L_{uv}^1 be the optimal cost. Compute also the quantities $Y_{uv} \equiv \sum_{t=u+1}^{v-1} \frac{x_t^1(u, v)}{W}$.

Step 2: For all u and l , $1 \leq u < l \leq T+1$, calculate $I_u^0(l)$.

Step 3: For each l , u and v , $1 \leq u < v \leq l \leq T+1$, calculate the following quantities, assuming that u and v are consecutive retailer regeneration points, l is a system

regeneration point and there are no other system regeneration points between u and l .

1. The inventory costs, $H_{uv}^0(l)$, at the warehouse between retailer regeneration points u and v : $H_{uv}^0(l) = h_0(v - u)I_v^0(l)$.
2. The total supplier-warehouse transportation cost from $u + 1$ to $v - 1$: $A_0 Y_{u,v}$.
3. The supplier-warehouse shipment quantities in period u .
 - (a) Assuming u is not a system regeneration point,
 - i. If $x_u^1(u, v) + I_v^0(l) \geq W$, then $x_u^0 = W$.
 - ii. Otherwise, $x_u^0 = 0$.
 - (b) Assuming u is a system regeneration point, a (possibly) partially loaded truck with $x_u^0 = x_u^1(u, v) + I_v^0(l)$ units is dispatched to the warehouse.
4. The total resulting supplier-warehouse transportation cost in periods u through $v - 1$, which we refer to as $T_{uv}^0(l)$ when u is not a warehouse regeneration point, and as $\tilde{T}_{uv}^0(l)$ when u is a warehouse regeneration point.
5. $F_{.l}(u, v) \equiv L_{uv}^1 + H_{uv}^0(l) + T_{uv}^0(l)$. (Observe that $F_{sl}(u, v) = F_{.l}(u, v)$ for all $s < u$.)
6. $F_{ul}(u, v) = L_{uv}^1 + H_{uv}^0(l) + \tilde{T}_{uv}^0(l)$.

Step 4: Calculate the arc cost C_{sl} for each $l = 2, 3, \dots, T + 1$, and $1 \leq s < l$ in $O(T^2)$ as follows. Let $R_{t,l}$ be the cost associated with periods t through l , given that t is a *retailer* regeneration point, l is a *system* regeneration point, and there are no other system regeneration points in between t and l .

1. Initialize $R_{ll} = 0$.
2. For each $s = l - 1, l - 2, \dots, 1$,

$$R_{sl} = \min_{k, s < k \leq l} \{F_{.l}(s, k) + R_{kl}\}$$

$$C_{sl} = \min_{k, s < k \leq l} \{F_{sl}(s, k) + R_{kl}\}$$

Step 5: Calculate the shortest path between 1 and $T + 1$ in a network with nodes $1, 2, \dots, T + 1$ and arcs (s, l) for each $1 \leq s < l \leq T + 1$ with length C_{sl} .

4.2 Multi-Retailer System

In the general case of $n > 1$ retailers, a warehouse regeneration point is not necessarily a system regeneration point. In this section, we show that we can still use a network (shortest path) approach to solve the Centralized Single-Warehouse Multi-Retailer Problem. However, the network is far more complex since details on the status of each retailer at each warehouse regeneration point need to be specified in order to calculate the costs associated with two consecutive warehouse regeneration points.

Construct an acyclic graph $G = (V, A)$, where

$$V = \{\bar{u} = \langle u_0, u_1, \dots, u_n \rangle \mid 1 \leq u_0 \leq u_i \leq T+1, i = 0, 1, \dots, n\} = \underbrace{\mathcal{T} \times \mathcal{T} \times \dots \times \mathcal{T}}_{n+1 \text{ times}},$$

$$A = \{\langle u_0, u_1, \dots, u_n \rangle \rightarrow \langle v_0, v_1, \dots, v_n \rangle \mid u_0 < v_0, u_i \leq v_i \text{ for } i = 1, 2, \dots, n\}$$

Each node $\langle u_0, u_1, \dots, u_n \rangle$ represents a warehouse regeneration point, u_0 , along with the earliest regeneration points for each retailer on or after that point, $u_i \geq u_0, i = 1, 2, \dots, n$.

We define the length of arc $\bar{u} \rightarrow \bar{v}$, where $\bar{u} = \langle u_0, \dots, u_n \rangle$ and $\bar{v} = \langle v_0, \dots, v_n \rangle$, as the minimum system-wide transportation and holding costs between periods u_0 and $v_0 - 1$ given that they are consecutive warehouse regeneration points. Observe that the pairs (u_i, v_i) are needed so that we can calculate the LTL quantities required by retailer i and subsequently the LTL quantity Δ that should be carried to the warehouse in period u_0 . Specifically,

$$\Delta = \sum_{i=1}^n \sum_{t=u_i}^{v_i-1} d_t^i - \lfloor \frac{\sum_{i=1}^n \sum_{t=u_i}^{v_i-1} d_t^i}{W} \rfloor.$$

It is easy to see that the shortest path from $\langle 1, 1, \dots, 1 \rangle$ to $\langle T+1, T+1, \dots, T+1 \rangle$ in $G = (V, A)$ corresponds to finding the optimal system ordering policy. Unfortunately, the network grows exponentially as the number of retailers increases.

In what follows we focus on calculating the cost of arc $\bar{u} \rightarrow \bar{v}$. For that purpose, we break time up in smaller increments such that there are no retailer regeneration points in between. We construct a new network $G_{(\bar{u} \rightarrow \bar{v})} = (V_{(\bar{u} \rightarrow \bar{v})}, A_{(\bar{u} \rightarrow \bar{v})})$, where

$$V_{(\bar{u} \rightarrow \bar{v})} = \{\bar{p} = \langle p_1, \dots, p_n \rangle \mid u_i \leq p_i \leq v_i\},$$

$$A_{(\bar{u} \rightarrow \bar{v})} = \{ \langle p_1, \dots, p_n \rangle \rightarrow \langle q_1, \dots, q_n \rangle \mid p_i < q_i \text{ if } p_i = \min_k p_k, \text{ and } p_i = q_i \text{ otherwise.} \}$$

The nodes represent successive regeneration points for each retailer, i.e. if we let $p_{min} = \min_{k=1,2,\dots,n} p_k$, then for each retailer, say i , p_i is the earliest regeneration point on or after p_{min} . The cost on arc $\langle p_1, \dots, p_n \rangle \rightarrow \langle q_1, \dots, q_n \rangle$ is the minimum system-wide cost between periods p_{min} and $q_{min}-1$ (where $q_{min} = \min_i q_i$), under the assumption that there are no regeneration points at any intermediate time in any facility in the system. Consequently, the associated costs can be calculated as follows.

1. For each retailer i , $i = 1, 2, \dots, n$:

- If $p_i = p_{min}$, solve the Decentralized Subproblem $DS_{p_i q_i}^i$ (see Section 3) and consider only the optimal cost and replenishment quantities between p_i and $q_{min} \leq q_i$.
- If $p_i > p_{min}$, solve the Decentralized Subproblem $DS_{p_{min}-1, p_i}^i$ and consider only the optimal cost and replenishment quantities between p_{min} and q_{min} .

2. Given the retailer replenishment quantities, calculate the coordinated warehouse shipping quantities and cost between p_{min} and q_{min} as in Step 3 in Section 4.1.1.

The cost associated with arc $\bar{u} \rightarrow \bar{v}$ is the shortest path between nodes $\langle u_1, u_2, \dots, u_n \rangle$ and $\langle v_1, v_2, \dots, v_n \rangle$ in network $G_{(\bar{u} \rightarrow \bar{v})}$. We refer to this type of algorithms, which consist of solving a shortest path on a network where the cost of each arc is calculated as the shortest path on a related network, as *nested* shortest-path algorithms.

5 Lagrangian Decomposition

The dynamic programming algorithm for Problem CP becomes computationally expensive as the number of retailers increases. Solving the associated single retailer problems using the properties derived above, however, is relatively fast. As a result, the problem appears well suited for a Lagrangian Decomposition approach that would allow us to break the problem down into a subproblem for each retailer while maintaining the coordination between them through Lagrangian multipliers. The solutions to these subproblems will provide both a

lower bound on the cost of an optimal solution and a starting point to construct good feasible solutions to the problem effectively.

Observe that the only constraints that link all the retailer facilities together are $x_t^0 + I_t^0 = \sum_{i=1}^n x_t^i + I_{t+1}^0 \forall t = 1, 2, \dots, T$. We develop two algorithms, which we denote by Aggregated and Disaggregated Lagrangian Decomposition, respectively, by adding the following variables and constraints:

- *Aggregated*: For each t , $t = 1, 2, \dots, T$, we add a new variable z_t^0 and a new constraint $z_t^0 = \sum_{i=1}^n x_t^i$. The linking constraint is then written as $x_t^0 + I_t^0 = z_t^0 + I_{t+1}^0$.
- *Disaggregated*: For each t and i , $t = 1, 2, \dots, T$, and $i = 1, 2, \dots, n$, we add a new variable z_t^i and a new constraint $z_t^i = x_t^i$. The linking constraint is then written as $x_t^0 + I_t^0 = \sum_{i=1}^n z_t^i + I_{t+1}^0$.

These new constraints will be relaxed so that the problem can be decomposed into one warehouse and n retailer subproblems. As we shall see in the computational section, the disaggregated method provides stronger lower bounds and better feasible solutions. However, it increases the computational time required to generate solutions. This trade-off needs to be considered when deciding which one to use in each particular case.

Let λ_t denote the Lagrangian Multipliers for the aggregated decomposition and λ_t^i those for the disaggregated counterpart. The objective functions of the resulting Lagrangian problems are:

Aggregated Lagrangian Decomposition Objective:

$$\text{Min} \quad \sum_{t=1}^T \sum_{i=0}^n (A_i y_t^i + h_i I_{t+1}^i) + \sum_{t=1}^T \lambda_t (z_t^0 - \sum_{i=1}^n x_t^i)$$

Disaggregated Lagrangian Decomposition Objective:

$$\text{Min} \quad \sum_{t=1}^T \sum_{i=0}^n (A_i y_t^i + h_i I_{t+1}^i) + \sum_{t=1}^T \sum_{i=1}^n \lambda_t^i (z_t^i - x_t^i)$$

The following sections study the subproblems associated with the two decompositions.

5.1 Retailer Subproblem

For each retailer i :

$$\begin{aligned}
 \text{Problem } RSP : \quad & \text{Min} \quad \sum_{t=1}^T (A_i y_t^i + h_i I_{t+1}^i - \lambda_t x_t^i) \\
 & \text{s.t.} \\
 & x_t^i \leq W y_t^i, \quad \forall t = 1, 2, \dots, T, \\
 & x_t^i + I_t^i = d_t^i + I_{t+1}^i, \quad \forall t = 1, 2, \dots, T, \\
 & I_1^i = 0, \\
 & x_t^i \geq 0, \quad \forall t = 1, 2, \dots, T, \\
 & y_t^i \in \{0, 1\}, \quad \forall t = 1, 2, \dots, T, \\
 & I_t^i \geq 0, \quad \forall t = 1, 2, \dots, T.
 \end{aligned} \tag{3}$$

The corresponding subproblem in the disaggregated Lagrangian decomposition technique is identical except that the multipliers are denoted by λ_t^i .

For this subproblem, the first part of Theorem 2.4 still holds. That is, between two consecutive retailer regeneration points, there is at most one LTL period in the optimal solution. If there were two LTL periods s and t ($s < t$), we can move a unit forward from s to t with an additional cost of $-(t-s)h_i + \lambda_s - \lambda_t$ or backward from t to s with an additional cost of $(t-s)h_i - \lambda_s + \lambda_t$. At least one of them should be non-positive. Thus, there is always an optimal solution with a single LTL period. The second part of the theorem, however, does not necessarily hold. The period to send the LTL shipment will depend on the new value of “transporting” it, λ_t , which now varies from period to period. It may thus be suboptimal for the subproblem to have the first period between regeneration points as the LTL period.

Observe, however, that the optimal solution to the original problem does satisfy both properties and thus we can restrict the feasible set to solutions that satisfy them. This strengthens the lower bound obtained through the Lagrangian decomposition and allows us to solve the subproblem using the shortest path algorithm presented in Section 3. The only difference is that the arc cost, L_{uv}^i , is now composed of the transportation and holding costs plus the costs associated with λ_t .

5.2 Warehouse Subproblem

The warehouse subproblem is where the aggregated and disaggregated methods differ, with the latter becoming stronger.

The aggregated subproblem for the warehouse can be written as follows:

$$\begin{aligned}
 \text{Problem } WSPA : \quad & \text{Min} \quad \sum_{t=1}^T (A_0 y_t^0 + h_0 I_{t+1}^0 + \lambda_t z_t^0) \\
 & \text{s.t.} \\
 & x_t^0 \leq W y_t^0, \quad \forall t = 1, 2, \dots, T, \\
 & x_t^0 + I_t^0 = d_t^0 + I_{t+1}^0, \quad \forall t = 1, 2, \dots, T, \\
 & I_1^0 = 0, \\
 & 0 \leq z_t^0 \leq nW, \quad \forall t = 1, 2, \dots, T, \\
 & 0 \leq y_t^0 \leq n, \quad \text{integer}, \quad \forall t = 1, 2, \dots, T, \\
 & I_t^0 \geq 0, \quad \forall t = 1, 2, \dots, T. \tag{4}
 \end{aligned}$$

This warehouse subproblem can be strengthened by adding inequalities always satisfied in optimal solutions to the original problem. In particular, the total quantity sent to retailer i up to time k in an optimal solution must be greater than or equal to the total demand up to k . In fact, it will be equal to that demand plus a portion of the demand in a certain number of succeeding periods up to the next regeneration period, say l , that cannot be consolidated into full trucks. That is, $\sum_{t=1}^k x_t^i = \sum_{t=1}^l d_t^i - \lfloor \frac{\sum_{t=k+1}^l d_t^i}{W} \rfloor W$ for some l such that $k \leq l \leq T$. Hence, we must have that:

$$\sum_{t=1}^k d_t^i \leq \sum_{t=1}^k x_t^i \leq \max_{k \leq l \leq T} \left\{ \sum_{t=1}^l d_t^i - \lfloor \frac{\sum_{t=k+1}^l d_t^i}{W} \rfloor W \right\} \tag{5}$$

Consequently,

$$\sum_{i=1}^n \sum_{t=1}^k d_t^i \leq \sum_{t=1}^k z_t^0 \leq \sum_{i=1}^n \max_{k \leq l \leq T} \left\{ \sum_{t=1}^l d_t^i - \lfloor \frac{\sum_{t=k+1}^l d_t^i}{W} \rfloor W \right\}$$

Solving the subproblem directly and repeatedly with an MIP solver in the Lagrangian

routine becomes computationally expensive for relatively large instances. To improve the computational efficiency of the Lagrangian decomposition algorithm, we solve the LP relaxation of WSPA first. Let \bar{y}_t^0 denote the value of y_t^0 in the solution of the LP relaxation. We substitute y_t^0 with $\lfloor \bar{y}_t^0 \rfloor + b_t^0$ ($b_t^0 \in \{0, 1\}$) to convert the general integer variable y_t^0 into a binary variable b_t^0 . The modified problem can then be solved with an MIP solver in a much shorter time. The resulting lower bounds are the same as those obtained when solving the subproblem directly in most cases and the differences observed in the remaining cases are very small. The justification behind this manipulation is that the full truck shipments sent in a particular period, ($\lfloor \bar{y}_t^0 \rfloor$), in the relaxed problem will also be in the optimal integer solution for most instances because the costs of sending a full truck are the same in each period and correspond to the linear costs. This modification to gain computational speed, may result in intermediate Lagrangian iterations where the so-called lower bound is not so. In the last iteration of the Lagrangian technique we solve the original mixed integer program for the warehouse subproblem directly to ensure that we have a true lower bound and quantify the distance from optimality of the final solution obtained.

The disaggregated subproblem for the warehouse is

$$\begin{aligned}
 \text{Problem } WSPD : \quad & \text{Min} \quad \sum_{t=1}^T (A_0 y_t^0 + h_0 I_{t+1}^0 + \sum_{i=1}^n \lambda_t^i z_t^i) \\
 & \text{s.t.} \\
 & x_t^0 \leq W y_t^0, \quad \forall t = 1, 2, \dots, T, \\
 & x_t^0 + I_t^0 = \sum_{i=1}^n z_t^i + I_{t+1}^0, \quad \forall t = 1, 2, \dots, T, \\
 & I_0^0 = 0, \\
 & 0 \leq z_t^i \leq W, \quad \forall t = 1, 2, \dots, T, i = 0, 1, \dots, n, \\
 & 0 \leq x_t^0 \leq nW, \quad \forall t = 1, 2, \dots, T, \\
 & 0 \leq y_t^0 \leq n, \quad \text{integer}, \quad \forall t = 1, 2, \dots, T, \\
 & I_t^0 \geq 0, \quad \forall t = 1, 2, \dots, T.
 \end{aligned} \tag{6}$$

We again strengthen the subproblem by adding constraints (5). To improve computational efficiency, we resort once more to the linear-programming-based approximation to a binary

problem described above. That is, we first solve the LP relaxation of this problem, then substitute y_t^0 with $\lfloor \bar{y}_t^0 \rfloor + b_t^0$ (where \bar{y}_t^0 is the value of y_t^0 in the solution of the LP relaxation and $b_t^0 \in \{0, 1\}$), and finally solve the modified subproblem. The disaggregated subproblem has more variables and takes significantly longer to solve than its aggregated counterpart, but it provides a better lower bound.

5.3 Lagrangian-Based Upper Bound

For both decomposition forms, a feasible solution can be easily constructed by fixing the optimal shipments obtained through each of the retailer subproblems and solving the associated decentralized problem for the warehouse.

5.4 Subgradient Optimization

After solving the subproblems and constructing the feasible solution, the Lagrangian Multipliers are updated using the following subgradient method.

For the aggregated form, in iteration k , the multipliers are updated as

$$\lambda_t^{k+1} = \lambda_t^k + t_k(z_t^0 - \sum_{i=1}^n x_t^i)$$

where t_k is a scalar step size such that $t_k \rightarrow 0$ and $\sum_{l=1}^k t_k \rightarrow \infty$ for $k \rightarrow \infty$, in order to assure convergence (Polyak (1967) [26]). Generally, we set

$$t_k = \frac{\alpha_k(\overline{UB} - LB)}{\sum_{t=1}^T (z_t^0 - \sum_{i=1}^n x_t^i)^2}$$

where $\alpha_k \in (0, 2]$ is a scalar parameter, \overline{UB} is the best upper bound available and LB is the current lower bound.

For the disaggregated form, in iteration k , the Lagrangian multipliers are updated as $(\lambda_t^i)^{k+1} = (\lambda_t^i)^k + t_k(z_t^i - x_t^i)$, where $t_k = \frac{\alpha_k(\overline{UB} - LB)}{\sum_{t=1}^T \sum_{i=1}^n (z_t^i - x_t^i)^2}$.

6 Computational Results

The focus of our computational experiments is to assess the quality of the solutions obtained through the two Lagrangian decomposition methods, determine the time required, and evaluate the gains obtained through system coordination. We consider 3 different problem scales: Small with 8 periods and 7 retailers, Medium with 10 periods and 20 retailers, Large with 12 periods and 50 retailers. We solve a total of 768 problem instances generated as follows.

Let $U(a, b)$ denote the uniform distribution over the range (a, b) . We construct 192 test instances of the One-Warehouse Multi-Retailer Problem with Full Truckload Shipments as follows.

We fix the cargo capacity constraint to $W = 25$ and test the algorithms over 4 different demand patterns: Low with data generated from $U(0, 6)$, Medium with $U(6, 12)$, High with $U(12, 20)$, and Wide Range with $U(0, 25)$. The fourth pattern represents the scenarios where demands change wildly over the time horizon.

To consider the trade-off between the fixed transportation costs per vehicle and the linear inventory holding costs, we generate the cost of dispatching a truck to retailer i , A_i , $i = 1, 2, \dots, n$, from $U(10, 20)$, for $i = 1, \dots, n$, and consider 4 different holding cost rates, h_i : High with data generated from $U(1.2, 2)$, Medium with $U(0.8, 1.2)$, Low with $U(0.4, 0.8)$ and Very Low with $U(0.1, 0.4)$.

The supplier-warehouse fixed cost per truck dispatched, A_0 , is generated as the product of a random number drawn from $U(10, 20)$ and a factor f_1 that captures the relative difference between supplier-warehouse and warehouse-retailer transportation costs. Normally, the warehouse is much closer to the retailers than to the supplier. Thus the transportation cost per cargo from the supplier to the warehouse is considered to be higher than or equal to that from the warehouse to the retailers. We use two factors: $f_1 = 4$ for a relatively high supplier-warehouse transportation cost and $f_1 = 1$ for a relatively low one.

The warehouse holding cost rate is calculated as $h_0 = f_2 \min \{ \tilde{h}_0, \min_{i=1, \dots, n} h_i \}$, where \tilde{h}_0 is randomly generated using the same distribution as for the h_i , $i = 1, 2, \dots, n$, described above, and f_2 is a factor that accounts for the relative difference between warehouse and retailer inventory holding rates. In our computational tests we consider $f_2 = 1$ for a relatively

high warehouse holding cost and $f_2 = 1/2$ for a relatively low one.

The values of the parameters used in the computational study are summarized in Table 1. They result in a total of 192 combinations for the base case (Case 1). For each combination of the parameter values, we generate a single instance.

Table 1: Values of the parameters used in the Base Case ($W = 25$, $A_i \in U(10, 20)$)

T	N	d	h_i	f_1	f_2
8	7	U(0, 6)	U(1.2, 2)	4	1
10	20	U(6, 12)	U(0.8, 1.2)	1	2
12	50	U(12, 20)	U(0.4, 0.8)		
		U(0, 25)	U(0.1, 0.4)		

We also explore scenarios with different types of retailers in three additional cases of 192 instances each. For this purpose, we first randomly generate 192 instances of the problem, each with a different combination of the parameter values, using the same parameters as in Table 1. In Case 2, we halve the transportation cost per truck A_i for half of the retailers. In Case 3, we halve the holding cost per unit per unit time h_i for half of the retailers. The holding cost per unit per unit time at the warehouse h_0 is then set equal to the lowest of the holding costs at the retailers. In Case 4, we halve the demand in each period d_t^i for half of the retailers.

In the Lagrangian Decomposition procedure, we start with $\alpha = 2$. If there is no improvement in the upper or lower bounds in 10 iterations, we halve the value of α . We initialize all Lagrangian Multipliers as 0. We consider three termination criteria: 1) $\alpha < 0.0001$; 2) $(\overline{UB} - LB)/LB < 0.001$; 3) running time is more than 1800s.

To benchmark the quality and speed of the solutions, we also solve Problem FTL for each of the instances generated with CPLEX MIP solver with a time limit of 1800s. Furthermore, to investigate possibly stronger formulations, we also solve the MIP version of the formulation introduced in Levi et al. (2005).

In Table 6, we summarize the quality of the solutions obtained by CPLEX using both the FTL and Levi's formulations, the aggregated and disaggregated Lagrangian Decomposition methods and the decentralized system. For that purpose, we report the relative difference

Table 2: Average and maximum percent relative gap of the feasible solutions (upper bounds) and the lower bounds generated with each method.

	T	N	Upper Bound					Lower Bound			
			Levi	FTL	Lag	AggLag	Decomp	Levi	FTL	Lag	AggLag
Case 1											
Average	8	7	0.03%	0.01%	0.21%	0.36%	0.87%	0.26%	0.10%	0.45%	0.64%
	10	20	0.19%	0.15%	0.19%	0.26%	0.44%	0.46%	0.55%	0.12%	0.14%
	12	50	0.32%	0.34%	0.14%	0.17%	0.24%	0.50%	0.70%	0.01%	0.01%
Max	8	7	0.66%	0.12%	3.54%	2.60%	6.65%	3.86%	2.74%	1.99%	2.49%
	10	20	1.10%	0.60%	0.76%	0.90%	1.82%	1.81%	3.88%	0.51%	0.62%
	12	50	1.89%	2.43%	0.53%	0.65%	0.98%	2.26%	3.95%	0.20%	0.14%
Case 2											
Average	8	7	0.01%	0.01%	0.18%	0.28%	0.79%	0.10%	0.02%	0.57%	0.82%
	10	20	0.15%	0.10%	0.21%	0.31%	0.49%	0.30%	0.32%	0.17%	0.20%
	12	50	0.31%	0.27%	0.17%	0.19%	0.26%	0.45%	0.60%	0.03%	0.03%
Max	8	7	0.05%	0.05%	1.65%	2.41%	4.95%	2.05%	0.89%	2.13%	3.42%
	10	20	0.96%	0.37%	0.81%	1.39%	1.80%	1.63%	2.63%	0.53%	0.65%
	12	50	2.60%	1.26%	0.84%	0.91%	1.07%	2.40%	3.85%	0.23%	0.21%
Case 3											
Average	8	7	0.03%	0.02%	0.13%	0.25%	0.63%	0.35%	0.11%	0.38%	0.54%
	10	20	0.27%	0.19%	0.14%	0.18%	0.27%	0.67%	0.73%	0.05%	0.08%
	12	50	0.42%	0.39%	0.09%	0.11%	0.14%	0.71%	0.99%	0.00%	0.00%
Max	8	7	0.34%	0.15%	1.15%	1.89%	5.46%	3.43%	1.74%	1.68%	2.46%
	10	20	1.26%	1.00%	0.57%	0.63%	0.80%	2.66%	6.08%	0.61%	0.88%
	12	50	2.36%	1.53%	0.41%	0.56%	0.68%	3.15%	5.08%	0.02%	0.04%
Case 4											
Average	8	7	0.01%	0.02%	0.17%	0.32%	0.91%	0.16%	0.04%	0.45%	0.71%
	10	20	0.52%	0.53%	0.55%	0.57%	0.73%	0.93%	1.26%	0.57%	0.62%
	12	50	0.44%	0.49%	0.30%	0.41%	0.34%	0.72%	1.45%	0.27%	0.40%
Max	8	7	0.20%	0.20%	1.32%	2.69%	3.89%	2.14%	1.39%	1.53%	3.06%
	10	20	4.26%	4.21%	4.21%	3.74%	4.21%	4.98%	7.23%	4.28%	3.82%
	12	50	2.16%	2.29%	2.66%	2.70%	2.66%	3.45%	7.27%	2.68%	2.75%

of both the lower bounds and the feasible solutions (upper bounds on total cost) generated by each method to the largest of the lower bounds obtained. Observe that CPLEX and the Lagrangian Decomposition methods provide both the current best feasible solution found and the current lower bound when they are terminated. The decentralized approach, however, does not provide a lower bound. Let LB^X be the lower bound provided by method X , $X = Levi, FTL, Aggregated, Disaggregated$, and UB^X be the feasible solution generated by method X , $X = Levi, FTL, Aggregated, Disaggregated, Decentralized$. Let $MaxLB = \max(LB^X, X \in \{Levi, FTL, Aggregated, Disaggregated\})$. We report the relative percent difference $\frac{UB^X - MaxLB}{MaxLB} * 100$ (%) and $\frac{MaxLB - LB^X}{MaxLB} * 100$ (%).

While the performance of CPLEX deteriorates as the problem size increases, the quality

Table 3: Running time (s)

	T	N	Levi	FTL	Lag	AggLag
Case 1						
Average	8	7	343.03	183.75	21.39	10.60
	10	20	1262.59	1389.48	114.52	10.49
	12	50	1797.75	1765.71	917.54	11.80
Max	8	7	1801.03	1800.39	40.56	21.27
	10	20	1800.63	1800.44	257.84	29.33
	12	50	1806.36	1800.56	1811.92	29.44
Case 2						
Average	8	7	219.06	77.89	22.59	10.63
	10	20	1043.65	1048.02	126.83	11.17
	12	50	1686.92	1651.20	1078.87	12.76
Max	8	7	1800.70	1800.23	35.72	19.28
	10	20	1800.95	1800.84	246.45	20.50
	12	50	1824.62	1802.50	1811.81	31.39
Case 3						
Average	8	7	528.59	313.06	20.18	9.59
	10	20	1585.20	1654.34	98.08	9.46
	12	50	1800.47	1800.15	632.84	9.61
Max	8	7	1802.80	1815.53	41.59	19.83
	10	20	1802.66	1800.42	261.88	21.81
	12	50	1802.48	1801.45	1811.47	28.25
Case 4						
Average	8	7	281.52	134.39	21.14	10.19
	10	20	1228.62	1448.94	124.58	11.27
	12	50	1800.47	1799.76	1036.05	12.30
Max	8	7	1800.09	1800.28	36.98	19.49
	10	20	1803.78	1800.42	256.64	20.77
	12	50	1802.70	1800.66	1813.42	27.36

of the solutions generated by the Lagrangian methods improves. This suggests that the proposed Lagragian Decomposition algorithms are superior to tackle large instances. The comparison of their running times in Table 6 further supports this finding. Recall that all algorithms are restricted by a running-time limit of 1800 seconds. The running time of CPLEX quickly reaches the limit, while the Lagrangian methods are very rarely constrained by it. The computational complexity of the Aggregate Lagrangian method grows very slowly with problem size, making it especially attractive to solve large-scale instances.

The difference in cost associated with the decentralized versus the centralized management of the system in the instances tested is rather small. We would like to point out, however, that there are examples where the relative difference can be as large as $\frac{1}{3}$: Consider a single retailer with demands over a planning horizon of two periods of $d_1 = \alpha W$ and $d_2 = (1 - \alpha)W$, $0 < \alpha < 1$. Assume that $2A_1 < A_1 + h_1 d_2$ and $A_0 + h_0 d_2 < 2A_0$ (that is, $A_1 < h_1 d_2$ but $A_0 > h_0 d_2$). In the decentralized solution, the retailer will order d_1 and d_2 separately in their respective periods while the warehouse will order d_1 and d_2 together in period 1, with a total cost of $2A_1 + A_0 + h_0 d_2$. If $h_0 \simeq h_1$, the optimal solution under centralized management is to ship d_1 and d_2 directly from supplier to warehouse to retailer and keep the inventory at the retailer. The total cost is $A_1 + A_0 + h_1 d_2$. The gap is $\frac{A_1}{A_1 + A_0 + h_1 d_2}$ which can get as close to $\frac{1}{3}$ as desired by taking $A_1 = h_1 d_2 - \epsilon$, $A_2 = h_1 d_2 + \epsilon$ and $\epsilon \rightarrow 0$.

Finally, we would like to compare the quality of our solutions with that of the solutions generated by algorithms that require the demand for each retailer in each particular time period to be satisfied by a single shipment, such as the 4.796-approximation algorithm of Levi, Roundy and Shmoys (2005) [19] and the cutting-plane procedure proposed by Croxton, Gendron and Magnanti (2003) [10] for more general settings. We refer to these solutions as non-splitting policies. To obtain a lower bound on the cost associated with non-splitting policies, we solved the integer programming formulation presented in [19] for small-size instances with 3 retailers and 5 time periods, randomly generated as described above. In a first group of 8 instances where demand is large, $U(12,20)$, relative to truck capacity of $W = 25$, the solutions are on average 3.3% over the optimal cost and reach a maximum deviation of 7%. In contrast, our algorithm provides the optimal policies in those cases and decentralized management of the system produces solutions that are on average 0.3% away

from optimality, with a maximum deviation of 1.7%. In a second group of 8 instances where demand is wide-ranging, $U(1,24)$, relative to the truck capacity of $W = 25$, the solutions are on average 1.65% over the optimal cost and reach a maximum deviation of 5%. In contrast, our algorithm provides the optimal policies in those cases and decentralized management of the system produces solutions that are on average 0.64% away from optimality, with a maximum deviation of 3.97%. For more details and examples where the best non-splitting policies lead to costs that are as much as $3/2$ times the optimal cost, we refer the reader to Jin (2006) [17].

7 Asymptotic Optimality of the Decentralized Approach

The computational results suggest that the performance of the decentralized solution increases as the problem size, in particular the number of retailers, increases. In fact, the following theorem shows that the decentralized solution is asymptotically optimal.

Theorem 7.1 *Let retailer demands d_t^i be i.i.d. with a mean of $E[d_t^i] = \bar{d} > 0$, and let Z^* and Z^{DC} denote the cost of the optimal and decentralized solutions, respectively. With probability one,*

$$\lim_{n \rightarrow \infty} \frac{Z^{DC} - Z^*}{Z^*} = 0,$$

where n is the number of retailers.

Proof. Observe that the total warehouse-retailer transportation and retailer inventory costs will be minimized by the decentralized solution, and can thus be ignored when upper-bounding the difference between decentralized and optimal solutions. We can then construct a lower bound for the optimal solution and an upper bound for the difference between decentralized and optimal solutions as follows.

The cost of the optimal solution is bounded below by the minimum possible cost to transport all demands from supplier to warehouse, associated with full truckload shipments at a cost per unit of A_0/W ; that is, $Z^* \geq (A_0/W) \sum_{i,t} d_t^i$.

The difference in cost between the decentralized and optimal solutions is bounded above by the difference between the costs associated with supplier-warehouse shipments and ware-

house inventory in the decentralized solution and the lower bound on those in the optimal solution, i.e., $(A_0/W) \sum_{i,t} d_t^i$. Observe that in the decentralized solution: (1) the inventory carried by the warehouse in any period will not exceed the truck capacity, W , since otherwise a full truckload could be postponed, and thus total warehouse inventory costs are bounded by Th_0W ; and (2) supplier-warehouse transportation costs will consist of a number of full truckloads and possibly one less than full truckload each period, at a cost that can thus be bounded by $(A_0/W) \sum_{i,t} d_t^i + TA_0$.

Therefore, we have that:

$$\frac{Z^{DC} - Z^*}{Z^*} \leq \frac{TA_0 + Th_0W}{(A_0/W) \sum_{i,t} d_t^i} = \frac{W(A_0 + h_0W)}{nA_0 \frac{\sum_{i,t} d_t^i}{nT}} \xrightarrow{n \rightarrow \infty} 0$$

This is true with probability one by the law of large numbers, since $E[d_t^i] = \bar{d} > 0$. ■

8 Conclusions

In this paper, we use structural properties of the optimal solutions to the one-warehouse multi-retailer problem with stationary full truckload costs to develop:

- an algorithm for the single-stage problem with complexity $O(T^2)$, which translates into solving the one-warehouse multi-retailer problem under decentralized management of the system in time $O(nT^2)$,
- an algorithm for the one-warehouse single-retailer problem with complexity $O(T^3)$, and its extension to the multi-retailer case with polynomial computation time for a fixed number of retailers, and
- heuristic algorithms based on two different Lagrangian decompositions of the problem: aggregated and disaggregated.

The structural properties and dynamic programming algorithms are valid as long as costs are nondecreasing over time so that it is always optimal to send the full truck as late as possible, i.e. under the monotonicity conditions in Lippman (1969). For more general non-stationary costs, Pochet and Wolsey (1993) showed that there is only one partial shipment

between consecutive regeneration points, but our algorithms fail because they also rely on the following two properties that do not hold in the general case: (1) the LTL shipment (if any) occurs in the first period between consecutive regeneration points; (2) inventory at any facility is always below truck capacity.

Our computational experiments show that the two Lagrangian Decomposition methods offer good solutions within reasonable time. For small and medium scale instances, the disaggregated Lagrangian Decomposition method offers better solutions. For large scale instances, however, the computational expense makes its aggregated counterpart more preferable. Finally, our computational experiments show that the gap between the centralized and decentralized solutions decreases as the problem scale increases. For large scale instances, managing the system in a decentralized fashion offers near-optimal solutions. In fact, we show analytically that the decentralized heuristic is asymptotically optimal.

References

- [1] Alp, O., N. K. Erkip and R. Güllü. 2003. Optimal Lot-Sizing/Vehicle-Dispatching Policies under Stochastic Lead Times and Stepwise Fixed Costs. *Operations Research* **51**(1), pp. 160–166.
- [2] Amiri, A. and H. Pirkul. 1997. New Formulation and Relaxation to Solve a Concave-Cost Network Flow Problem. *Journal of the Operational Research Society* **48**(3), pp. 278–287.
- [3] Anily, S. and M. Tzur. 2005. Shipping Multiple-Items by Capacitated Vehicles An Optimal Dynamic Programming Approach. *Transportation Science* **39**(2), pp. 233-248.
- [4] Anily, S. and M. Tzur. 2006. Algorithms for the Multi-Item Capacitated Dynamic Lot Sizing Problem. *Naval Research Logistics* **53**(2), pp. 157-169.
- [5] Arkin E., D. Joneja and R. Roundy. 1989. Computational Complexity of Uncapacitated Multi-Echelon Production Planning Problems. *Operations Research Letters* **8**, pp. 61–66. North-Holland.
- [6] Balakrishnan, A. and S. Graves. 1989. A Composite Algorithm for a Concave-Cost Network Flow Problem. *Networks* **19**, pp. 175–202.
- [7] Bramel, J. and D. Simchi-Levi. 1997. The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management, Springer Series in Operations Research, Springer, New York.
- [8] Chan, L. M. A., A. Muriel, Z. J. Shen, D. Simchi-Levi and C. P. Teo. 2002a. On the Effectiveness of Zero-Inventory-Ordering Policies for the One-Warehouse Multi-Retailer Problem with Piecewise Linear Cost Structures. *Management Science* **48**(11), pp. 1446–1460.
- [9] Chan, L. M. A., A. Muriel, Z. J. Shen and D. Simchi-Levi. 2002b. On the Effectiveness of Zero-Inventory-Ordering Policies for the Economic Lot-Sizing Model with a Class of Piecewise Linear Cost Structures . *Operations Research* **50**(6), pp. 1058–1067 .

- [10] Croxton K. L., B. Gendron and T. L. Magnanti. 2003. Models and Methods for Merge-In-Transit Operations. *Transportation Science* **37**, pp. 1–22.
- [11] Diaby, M. and A. Martel. 1993. Dynamic Lot Sizing for Multi-Echelon Distribution Systems with Purchasing and Transportation Price Discounts. *Operations Research* **41**(1), 48–59.
- [12] Gallego, G. and D. Simchi-Levi. 1990. On the Effectiveness of Direct Shipping Strategy for the One-warehouse Multi-retailer R-systems. *Management Science* **36**(2), 240–243.
- [13] Gendron, B., T. G. Crainic and A. Frangioni. 1999. Telecommunications Network Planning, in Multicommodity Capacitated Network Design, B. Sans and P. Soriano (eds.), Kluwer Academic Publishers, 1–19.
- [14] Holmberg, K. 1994. Solving the Staircase Cost Facility Location Problem with Decomposition and Piecewise Linearization. *European Journal of Operational Research* **75**, pp. 41–61.
- [15] Holmberg, K., and J. Ling. 1997, A Lagrangean Heuristic for the Facility Location Problem with Staircase Costs. *European Journal of Operational Research* **75**, pp. 41–61.
- [16] Jaruphongsa, W., S. Çetinkaya and C. Y. Lee. 2005. A Dynamic Lot-Sizing Model with Multi-Mode Replenishments: Polynomial Algorithms for Special Cases with Dual and Multiple Modes. *IIE Transactions* **37**(5), pp. 453–467.
- [17] Jin, Y. 2006. On The Integration of Production, Transportation, and Inventory Control. PhD thesis, University of Massachusetts Amherst.
- [18] Joneja D. 1990. The Joint Replenishment Problem: New Heuristics and Worst Case Performance Bounds. *Operations Research* **38**(4) pp. 711–723.
- [19] Levi, R., R. Roundy and D. Shmoys. 2005. First Constant Approximation Algorithm for the One-Warehouse-Multi-Retailer Problem. To appear in *Management Science*.

- [20] Li, C. L., V. N. Hsu and W. Q. Xiao. 2004. Dynamic Lot Sizing with Batch Ordering and Truckload Discounts. *Operations Research* **52**(4), pp. 639–654.
- [21] Lippman, S. 1969. Optimal Inventory Policy with Multiple Set-up Costs. *Management Science* **16**(1), pp. 118–138.
- [22] Lu, L. and M. Posner. 1994. Approximation Procedures for the One-Warehouse Multi-Retailer System. *Management Science* **40**(10), pp. 1305–1316.
- [23] Lee, C.-Y., S. Çentikaya and W. Jaruphongsa. 2003. A Dynamic Model for Inventory Lot-Sizing and Outbound Shipment Consolidation at a Third Party Warehouse. *Operations Research* **51**, pp. 735–747.
- [24] Muriel, A. and D. Simchi-Levi. 2004. Supply Chain Design and Planning - Applications of Optimization Techniques for Strategic and Tactical Models. In A. G. de Kok and S. C. Graves (eds) “Supply Chain Management: Design Coordination and Operation,” Handbooks in Operations Research and Management Science, Volume 11.
- [25] Pochet, Y. and L. A. Wolsey. 1993. Lot-Sizing with Constant Batches: Formulation and Valid Inequalities. *Mathematics of Operations Research* **18**(4), pp. 767 - 785.
- [26] Polyak, B. T. 1967. A General Method for Solving Extremum Problems (in Russian). *Doklady Adademii Nauk SSSR* **174**, pp. 33 - 36.
- [27] Roundy, R. 1985. 98%-Effective Integer-Ratio Lot-Sizing For One-Warehouse Multi-Retailer Systems. *Management Science* **31**(11), pp. 1416 - 1430.
- [28] Schwarz, L. B. 1973. A Simple Continuous Review Deterministic One-Warehouse N-Retailer Inventory Problem. *Management Science* **19**(5), pp. 555 - 566.