

SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NY 14853-7501

TECHNICAL REPORT NO. 916

July 1990

HEURISTICS FOR A ONE WAREHOUSE
MULTI-RETAILER DISTRIBUTION PROBLEM
WITH PERFORMANCE BOUNDS¹

by

Yale Herer and Robin Roundy

¹This work was supported in part by NSF grants ECS-8404641 and DMC-8451984, and also by AT&T Information Systems and DuPont Corporation.

Abstract

We investigate the one warehouse multi-retailer distribution problem with traveling salesman tour vehicle routing costs. We model the system in the framework of the more general production/distribution network with arbitrary non-negative monotone joint order costs. Our goal is to develop heuristics whose costs are provably close to the cost of an optimal policy for these systems.

We show that if we are given a submodular function which is close to the true order cost in a general production/distribution system, then we can find a power of two policy whose cost is close to the cost of an optimal policy.

In addition we propose a dynamic program that computes optimal power of two policies for the one warehouse multiretailer system, assuming only that the order costs are non-negative monotone.

Finally, we perform computational tests that show that our power of two policies for one warehouse multi-retailer distribution systems have less error than our theoretic worst case analysis would lead us to believe.

1 Introduction

The problem studied herein is that of when to deliver products from a central warehouse to geographically dispersed retailers, and when to replenish the inventories at the central warehouse, in order to minimize long run average cost. We use the term facility to refer to either the warehouse or one of the retailers. The goods inventoried and delivered can be a single product or a group of products, and the set of products delivered to one retailer need not be the same as those delivered to another retailer. The costs involved can be thought of as an extension of the costs in the classical EOQ model (Harris 1915).

The central warehouse receives deliveries from an outside supplier(s). Inventories can be held at the central warehouse or at the retailers. When an order is placed by one or more retailers, a truck departs from the central warehouse and visits all of the retailers that are currently ordering. This delivery is assumed to be made instantaneously. Demand occurs at the retailers, and is constant, continuous, and known. This type of demand assumption is common in the literature.

Two types of costs are involved, ordering costs and inventory costs. The two costs pull the system in two different directions. Ordering costs induce the retailers and central warehouse to order less often, thus driving up inventory levels. However inventory holding costs induce the retailers and central warehouse to order more often, thus driving up the number of orders placed.

The ordering costs are assumed to be the sum of two cost functions. The first

is a monotone non-negative submodular cost function. In this context submodular cost functions can be thought of as concave, or as representing economies of scale. This is a very general cost function, and can represent many costs commonly encountered. Examples include the cost of hiring a truck, loading a retailer's goods at the central warehouse, unloading the goods at the retailer, processing a retailer's order, a fixed cost for ordering at the warehouse, and many others. A mathematical definition of submodular functions will be given below.

The second cost function is a per mile charge times the length of the travelling salesman tour through the central warehouse and the retailers that are visited. This cost function is intended to include all costs that are proportional to the distance the delivery truck must travel. These costs may include, but are not limited to, gasoline, labour hours, and truck depreciation. This type of cost is present in most delivery systems, and is a major component in many of them.

The inventory costs are the same as in the standard EOQ model. They are linear, and depend only on the stock on hand. We allow the inventory holding cost rate for each retailer's goods to differ both at the retailer and at the warehouse. We make only one assumption regarding the inventory holding cost rates. We assume that the inventory holding cost rate at the warehouse, for a particular retailer's goods, is no more than the inventory holding cost rate at the retailer. We feel that this assumption is justifiable because the warehouse is usually created to hold the inventory where it is easier and cheaper to handle. Furthermore, inventory holding

costs rates are often set to reflect value added, and certainly no value is taken away by transferring items from the warehouse to a retailer.

This problem arises in many real world situations where the delivery costs are significant, and have an important component which is proportional to the distance travelled. One example is the delivery of items in a grocery store chain. The chain has a central warehouse and retailers. Items are cheaper to store at the warehouse, and deliveries must be made to the retailers via a truck. Other situations where this model may apply include delivery of gasoline to service stations and delivery of raw materials to manufacturing plants.

We will solve this problem by calculating power of two reorder intervals. A reorder interval is a length of time that a facility waits between placing successive orders and thus replenishing its inventory. By a power of two reorder interval, we mean that for each facility the reorder interval is a positive or negative integer power of two times some common base planning period. For example, if the base planning period were one, then the allowable reorder intervals are $\dots, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1, 2, 4, 8, 16, \dots$. This may at first seem like a very restrictive assumption. But, as we show below, the cost of the policy we calculate with this restriction is not far from the cost of the optimal policy.

We will call a collection of power of two reorder intervals a power of two policy. A power of two policy is implemented by having each facility place an order at time zero. Each facility follows its own reorder interval thereafter. The order quantities

are such that a facility orders only when its physical inventory is zero. At each point in time when one or more retailers places an order, it is assumed that a delivery vehicle follows a travelling salesman tour through the warehouse and the retailers that are currently placing orders.

Power of two policies have many advantages. They are easier to deal with than general policies. A manager of a facility can more easily understand and plan other activities (such as production) if he knows he will be placing an order and receiving deliveries once every four days. If general policies were allowed or optimal policies were used, deliveries might arrive at a facility once every π days in one month, and once every $\sqrt{7}$ days in the next month. Power of two policies are also cyclic, and thus they have a regenerative nature. A power of two policy's cycle length is equal to $\tau_{max} = \max_{\{0 \leq i \leq N\}} \tau_i$. Power of two policies are also easier to implement. The truck driver needs to learn only a few different routes. In fact if we let $\tau_{min} = \min_{\{0 \leq i \leq N\}} \tau_i$ then the number of routes is at most $\log_2(\frac{\tau_{max}}{\tau_{min}})$. This number will typically be small. In fact as pointed out by Federgruen and Zheng (1988), it is hard to imagine a situation where it would be greater than ten. If it is ten, and $\tau_{min} =$ one day then there would be at least one facility that would receive deliveries only every 1024 days, or once every 2.8 years.

One assumption that is made implicitly in this model is that the delivery truck has infinite capacity. This assumption is not as constraining as it might appear at first. There are many situations where this assumption holds. These situations

include urban distribution systems where delivery trucks are often not filled to capacity. This is generally because the short distances involved and the high cost of floor space tend to insure short reorder intervals, and thus small order quantities. We note that the bounds we develop on the cost of our heuristic policies are based on travelling salesman tour heuristics that typically return to the central warehouse several times in the middle of a route (see Herer 1990). At these times the delivery truck can pick up more goods to deliver. Another option is to subdivide the route into subroutes at these points, and use multiple trucks. In addition, it is our belief that this work will prove to be useful in studying more complicated inventory distribution systems that model limited capacity in the delivery vehicles. Possible extensions of this work will be discussed in Section 6

Another assumption we have made is that of instantaneous delivery. This assumption is justified if the time to deliver is small as compared to the reorder intervals. This is typically the case.

Some of the results derived herein actually apply to a wider class of problems than those described in the introduction. But we feel that the situation described here is the most interesting and useful. Furthermore the computational results section deals with the model described here in the introduction.

The one warehouse multi-retailer distribution system with travelling salesman tour vehicle routing costs has been studied with a slightly different model. Anily and Federgruen (1988) use the same model except that no inventory is allowed

at the warehouse, trucks have a limited capacity, the holding costs at all of the retailers are the same, and the order cost consist of a cost per mile and a fixed cost for hiring a truck. They restrict their search of policies to policies where the demand is split into regions, and whenever any retailer in a region orders all retailers in that region order. They propose a heuristic which they show to be asymptotically optimal within the given class of policies. Gallego and Simchi-Levi (1988) use the same model as Anily and Federgruen (1988) except retailers are allowed to have different holding costs, there is no cost of hiring a truck, and the order costs can also include a fixed cost for each retailer ordering. They propose a direct shipping heuristic whose cost is no more than 1.061 times the cost of an optimal policy if the economic order quantity for each retailer is at least 71% of a truck load.

This problem is closely related to the inventory routing problem. In that problem trucks are generally considered to have finite capacities. We are given the demands over a finite number of discrete time periods, and must construct routes to service our customers while trying to keep the distance travelled to a minimum. The problem is usually solved on a rolling horizon basis. Dror, et. al. (1985), Dror and Levy (1986), and Dror and Ball (1987) consider problems in the fuel and gas industries Their objective is to minimize cost subject to not allowing stock-outs. Beltrami and Bodin (1974) examine an inventory routing problem that arises in garbage collection. Federgruen and Zipkin (1984) study an inventory routing

problem with a scarce resource.

Burns, et. al. (1985) is the first paper, to our knowledge, that attempted to minimize long run average cost, combining travelling salesman tour length costs with inventory costs. They used information on the spatial density of the retailers rather than their exact locations.

A related problem, with general production/distribution networks and with less general cost functions, is well solved in a practical sense. Roundy (1985, 1986) demonstrated a heuristic to obtain a power of two policy that is within two percent of optimal (six percent if the base planning period is fixed). His results hold for the case of ordering costs that fit into a non-negative family model. This model of costs is described below. Federgruen, et. al. (1989), Queyranne (1985), and Zheng (1987) have extended Roundy's results to the case where the ordering costs are non-negative monotone submodular (submodularity will be defined below). Roundy's algorithms are polynomial if the number of families is polynomial. The algorithms for submodular order costs are polynomial if the order cost function can be evaluated in polynomial time. The main contribution of this work is to extend these bounds to order cost functions which are close to being submodular. One such function which has been shown to be close to submodular in Herer (1990), is the function whose values are the travelling salesman tour lengths through the central warehouse and a set of retailers.

The rest of this paper is organized as follows. In Section 2 we present a mathe-

mathematical formulation of our problem, and show that the cost of our solution is close to the cost of the optimal solution. In Section 3 we relate these results to the one warehouse multi-retailer problem, and propose heuristics. In Section 4 we present a dynamic programming formulation to find the optimal power of two policy. In Section 5 we present computational results. Finally, in Section 6 we present some directions of future research.

2 Cost Bounds for General Production/Distribution Systems

In this section we deal with general production/distribution systems. The one warehouse multi-retailer system, a special case of general production/distribution systems, is used as an illustration. We will show that there is a power of two policy whose cost is close to the cost of an optimal policy, if the ordering costs are close to being submodular. In the next section we will relate these results to the one warehouse multi-retailer distribution system. We will also use these results to develop a heuristic for the one warehouse multi-retailer distribution system.

Some notation that will be used throughout this paper is as follows:

- N The number of retailers.
- r_i Retailer i .
- r_0 The central warehouse.

S	A subset of V .
$K(S)$	A submodular estimate of $T(S)$.
τ_v	The power of two reorder interval for node v .
τ^e	The power of two reorder interval for route e .
b	The base planning period.

We now give a mathematical description of our model. We assume we are given a bill of material network. We let $V = \{v_1, v_2, \dots, v_{|V|}\}$ be the set of nodes in the bill of material network. The nodes of the bill of material network represent operations. This may include machining operations (such as grinding), assembly operations, disassembly operations, and/or the movement of material. The inventory of the end product of an operation is associated with the corresponding node. The arcs of the bill of material network represent material requirements. An arc from node v_i to node v_j means that $\lambda_{v_i v_j}$ units of product v_i are required to make one unit of product v_j . The $\lambda_{v_i v_j}$ s are called the gozinto parameters. The only assumption we make about the network is that there are no directed cycles, i. e. there is no part that is eventually an input to itself. This is not always the case, as when a product is a catalyst in a reaction to make itself. But in most manufacturing settings (and certainly in the one warehouse multi-retailer distribution system) there are no directed cycles. A more complete description of the network representation of an arbitrary production/distribution system can be found in Roundy (1986). Our model differs only in the types of setup cost functions allowed.

A route is any directed path of nodes in the bill of material network that ends in a product for which there is external demand. We let $E = \{e_1, e_2, \dots, e_{|E|}\}$ be the set of all of the routes in the bill of material network. We call the node at which the route starts the tail of the route, while we call the node at the end of the route the head of the route. If the tail of route e is v_i and the head is v_j then λ^e is the number of v_i 's required, on this route, to make one v_j . We define the demand of a route to be $d^e = \lambda^e d_{v_j}$ where d_{v_j} is the external demand for node v_j . The inventories at the nodes are assigned to the routes they will eventually follow. If a node is the tail of four different routes, its inventory is conceptually spilt into four different inventories, one for each route. Each item in inventory is assigned to the route that it will follow. Consequently if an item is associated with route e , then it will visit each of the nodes in e in succession and will eventually be sold as a finished product at the head of route e . Holding costs will be computed using route-based inventories. A more complete discussion of routes can be found in Roundy (1986).

As an illustration the bill of material network for a one warehouse multi-retailer production/distribution system is contained in Figure 1. The node labeled r_i , $0 \leq i \leq N$, represents facility i . The arc from r_0 to r_i , $1 \leq i \leq N$ represents the flow of material from the central warehouse to retailer i . The routes are $(r_0, r_1), (r_0, r_2), \dots, (r_0, r_N)$ and $(r_1), (r_2), \dots, (r_N)$.

We are interested in minimizing the long run average cost. There are two costs




Figure 1: The bill of material network for a one warehouse multi-retailer distribution system.

involved in our model, inventory costs and ordering costs.

Holding costs are computed using routes and echelon inventories, as was done in (Roundy 1986). For each route there is a linear echelon holding cost rate. This rate is assumed to be greater than zero. Let H_v be the traditional holding cost rate at node v and let \mathcal{D}_v be is the set of immediate predecessors of node v . Then the echelon holding cost rate is $h_v = H_v - \sum_{k \in \mathcal{D}_v} \lambda_{kv} H_k$. If v is the tail of route e then we let $h^e = h_v$. For the one warehouse multi-retailer system the echelon holding cost rate at the warehouse is equal to the traditional holding cost rate, while the echelon holding cost rate at a retailer is the traditional holding cost rate at the retailer minus the holding cost rate at the warehouse.

The ordering costs can be represented as a set function $L(S)$ that maps subsets S of the set V to \mathfrak{R}^+ . $L(S)$ is a general non-negative monotone real valued set function which is the sum of two set functions. The first, $G(S)$, is any non-negative monotone submodular set function. A submodular set function is a set function which maps subsets of the set V to \mathfrak{R} , and satisfies the following inequality:

$$G(S \cup L \cup M) - G(S \cup L) \leq G(S \cup M) - G(S) \quad \forall \text{ disjoint } S, L, M \subset V. \quad (1)$$

By non-negative monotone we mean $G(S \cup M) \geq G(S) \geq 0$ for all disjoint $S, M \subset V$. The second component of the order cost function $L(S)$ is $mT(S)$, which is any non-negative monotone set function. Ordering costs are incurred each time a set of nodes places an order. If at a particular point in time the set S of nodes places an order then a cost $L(S) = G(S) + mT(S)$ is incurred. For the one warehouse multi-retailer distribution system $mT(S)$ is the cost m per mile travelled times the travelling salesman tour length $T(S)$ through the retailers in S and the central warehouse. Recall from the introduction that $G(S)$ can include the cost of hiring a truck, fixed warehouse order costs, fixed retailer order costs, etc.

As far as we know, the only structural property of policies for this problem that are optimal is that of zero ordering. That is, if there exists an optimal policy, then there exists an optimal policy where nodes order only when their inventory is zero (their conventional inventory, not their echelon inventory). This was shown

by Federgruen et. al. (1989) for monotone non-negative submodular order cost functions. Their proof remains valid for any non-negative order cost function.

We will show that the cost of an optimal power of two policy is close to the cost of an optimal policy. We will also show that if one has a good non-negative monotone submodular approximation of the true order costs then one can use this approximation to compute reorder intervals. If one uses these reorder intervals in the original inventory system then one incurs a cost which is close to the cost of an optimal policy. A power of two policy is a policy in which each node $v \in V$ orders once every $\tau_v = 2^{z_v} b$ time units, starting at time zero. Here z_v is a positive or negative integer. Recall that b is a fixed base planning period. We also require that in a power of two policy facility i receive orders only when its conventional inventory is zero.

Routes have power of two reorder intervals that are closely related to the power of two reorder intervals for the nodes in the bill of materials network. If route $e = (v_1, \dots, v_m)$ where v_1, \dots, v_m are nodes in the traditional bill of material network then $\tau^e = \max_{v \in e} \tau_v$. The maximum is taken because the tail of the route orders goods that will travel along the route, only once every $\max_{v \in e} \tau_v$ time units (see Roundy 1986). If v_i is the tail of route e and $\tau_{v_i} < \tau^e$ then node v_i does not order for route e every time it orders. For optimal power of two policies $\tau_v = \min_{\{e|v \text{ is the tail of } e\}} \tau^e$.

The echelon inventory of route e is the total amount of inventory in the system

that is or once was associated with route e . Since power of two policies are zero ordering, the echelon inventory (for route e) follows a saw toothed pattern (see Roundy 1986).

The long run average cost of an arbitrary policy can be very hard to calculate. However the cost of a power of two policy has a simple functional form. Let σ be a permutation of the indices $\{1, \dots, |V|\}$ such that $\tau_{\sigma(1)} \leq \tau_{\sigma(2)} \leq \dots \leq \tau_{\sigma(|V|)}$. Let $\frac{1}{\tau_{\sigma(|V|+1)}} = 0$. The long run average cost is then equal to

$$\sum_{e \in E} \frac{1}{2} d^e h^e \tau^e + \sum_{i=1}^{|V|} \left(\frac{1}{\tau_{\sigma(i)}} - \frac{1}{\tau_{\sigma(i+1)}} \right) L \left(\bigcup_{j=1}^i \{v_{\sigma(j)}\} \right). \quad (2)$$

The first term represents the holding cost charged at the tail of the route for the echelon inventory of the route.

The second term represents the order cost. If $\tau_{\sigma(i)} = \tau_{\sigma(i+1)}$ then the cost $L \left(\bigcup_{j=1}^i \{v_{\sigma(j)}\} \right)$ is never incurred because $\frac{1}{\tau_{\sigma(i)}} - \frac{1}{\tau_{\sigma(i+1)}} = 0$. Thus the second term is non-zero only when $\tau_{\sigma(i)} < \tau_{\sigma(i+1)}$. Once every $\tau_{\sigma(i)}$ time units an order is placed by all of the nodes in the set $\bigcup_{j=1}^i \{v_{\sigma(j)}\}$. Once every $\tau_{\sigma(i+1)}$ time units this order coincides with orders placed by other nodes, and an order cost larger than $L \left(\bigcup_{j=1}^i \{v_{\sigma(j)}\} \right)$ is incurred.

We can use (2) as the objective function in an optimization problem to find the optimal power of two policy. The optimization problem is:

$$\min \left[\sum_{e \in E} \frac{1}{2} d^e h^e \tau^e + \sum_{i=1}^{|V|} \left(\frac{1}{\tau_{\sigma(i)}} - \frac{1}{\tau_{\sigma(i+1)}} \right) L \left(\bigcup_{j=1}^i \{v_{\sigma(j)}\} \right) \right]$$

subject to:

$$\tau_v = 2^{z_v} \quad z_v \in \{\dots, -2, -1, 0, 1, 2, 3, \dots\} \text{ for all } v \in V \quad (3)$$

$$\tau^e = \max_{v \in e} \tau_v \quad \text{for all } e \in E. \quad (4)$$

If $T(S) \equiv 0$ then $L(S)$ is a monotone non-negative submodular order cost function. The problem would essentially be solved if we had an oracle that could evaluate $G(S)$ in polynomial time (see Federgruen, et. al. 1989, Queyranne 1985, and Zheng 1987). Federgruen et. al. (1989) show that when $L(S)$ is monotone non-negative submodular, if one replaces (3) with

$$\tau_v \geq 0 \quad \text{for all } v \in V \quad (5)$$

and solves the resulting problem one obtains a lower bound on the cost of any policy, not just power of two policies. This new optimization problem is the continuous relaxation of the previous optimization problem. Queyranne (1985) and Zheng (1987) give a polynomial time algorithm (given an oracle that can evaluate $G(S)$ for any set S in polynomial time) to compute the optimal power of two policy. Federgruen et. al. (1989) then show that the cost of this optimal power of two policy is no more than 1.021 times the value of the optimal solution to the continuous relaxation. They also show that $C^*(G)$ as defined below is equal to the value of the optimal solution to the continuous relaxation.

$$C^*(G) = \max \left(\sum_{v \in V} 2\sqrt{q_v(G)h'_v(G)} \right) \quad (6)$$

subject to:

$$\begin{aligned} q(G) &\in \mathcal{P}(G) \\ \frac{1}{2}h^e d^e &= \sum_{v \in e} x_{ev} \quad e \in E \end{aligned}$$

$$\sum_{\{e|v \in e\}} x_{ev} = h'_v(G) \quad v \in V$$

$$q(G) \geq 0, x \geq 0.$$

Here $\mathcal{P}(G)$ is the polymatroid $\{q | \sum_{v \in S} q_v \leq G(S) \text{ for all } S \subseteq V\}$. The variables are the q 's, x 's, and h 's. We let $q^*(G)$ and $h^*(G)$ be the optimal values that solve the above maximization problem with order cost function $G(S)$.

We wish to investigate what happens if $T(S)$ is non-negative, monotone, and approximately submodular. One might at first be worried that a slight deviation in the estimate of the order cost might cause a large error in the cost of the policy computed when compared to the optimal cost. However this does not happen because of the robustness of the EOQ model. This robustness has been known at least since Brown (1978 p. 184). He notes that if the quantity ordered is off from the optimal by no more than a factor of $\sqrt{2}$ then the impact on cost is no more than six percent. In fact for our problem, if the true ordering cost is non-negative monotone submodular and we under or over estimate it by a non-negative monotone submodular cost function that is within a factor of $\sqrt{2}$ of the true cost, then the deviation in the cost of the optimal power of two policy is no more than six percent. Even though our true order cost function is not non-negative monotone submodular we will see in Theorem 1 that having a non-negative monotone submodular function close to $T(S)$ will mean that we can find a power of two policy for our production/distribution system that has a cost which is close to the cost of the optimal policy.

We wish to quantify what we mean when we say that a cost function is close to being non-negative monotone submodular. We use the following optimization problem to quantify close:

$$\min \alpha$$

subject to:

$$K(S \cup L \cup M) - K(S \cup L) \leq K(S \cup M) - K(S) \quad \forall \text{ disjoint } S, L, M \subset V \quad (7)$$

$$K(S) \leq K(S \cup M) \quad \forall \text{ disjoint } S, M \subset V \quad (8)$$

$$T(S) \leq K(S) \leq \alpha T(S) \quad \forall S \subseteq V. \quad (9)$$

(7) forces $K(S)$ to be submodular, (8) forces $K(S)$ to be monotone, while (9) along with the objective function forces $K(S)$ to be close to $T(S)$. Note that (9) also guarantees that $K(S)$ is greater than or equal to zero. We say that $K_1(S)$ is closer to $T(S)$ than $K_2(S)$ is, if the α associated with $K_1(S)$ is less than the α associated with $K_2(S)$.

We will develop an approach for computing power of two policies, and we will show that the cost of a power of two policy computed using our approach is close to the cost of an optimal policy. We assume the existence of a monotone non-negative submodular set function $K(S)$ such that $T(S) \leq K(S) \leq \alpha T(S)$ for all $S \subseteq V$. We define $J(S) = \frac{K(S)}{\eta}$, $1 \leq \eta \leq \alpha$ and we define θ by $\theta = \frac{\alpha}{\eta}$. Therefore $\frac{J(S)}{\theta} \leq T(S) \leq \eta J(S)$. We define γ to be the smallest value that satisfies

$$G(S) + mK(S) \leq \gamma(\alpha G(S) + mK(S)) \quad \text{for all } S \subseteq V.$$

Note that this equation always holds for $\gamma = 1$ and that $\gamma \geq \frac{1}{\alpha}$. We will see that a smaller value of $\alpha\gamma$ will give rise to tighter bounds. We also note that as $G(S)$ becomes a larger part of the total cost, the factor $\alpha\gamma$ decreases. We define three new monotone non-negative submodular set functions $L^l(S) = G(S) + \frac{mK(S)}{\alpha}$ and $L^u(S) = G(S) + mK(S)$, which are respectively lower and upper bounds on $L(S)$, and $L^j(S) = G(S) + mJ(S)$.

We will refer to a few different systems in the following theorem and corollary. They are all identical except for the order costs incurred. The original system we will refer to as the L system. We will also refer to the original system as the $G+mT$ system because $L(S) = G(S) + mT(S)$. The other systems will be referred to by their ordering cost function. We will call the system that is identical to the original system but incurs an order cost $L^u(S)$ instead of $L(S)$ the L^u system. The L^l system and the L^j system are similarly defined.

Theorem 1 *The cost of an optimal policy for the L system is bounded from below by $C^*(L^l)$, the optimal value of (6) for the L^l system, and from above by the cost of an optimal power of two policy for the L^u system. The ratio of the upper and lower bounds is less than $1.021\sqrt{\alpha\gamma}$ if the base planning period b is variable, and $1.061\sqrt{\alpha\gamma}$ if b is fixed. If $J(S) = K(S)$ (or equivalently, $\eta = 1$) then these bounds also apply to the cost of computing an optimal power of two policy for the L^j system and using that policy in the L system.*

Proof:

Claim 1.1 *Let $G(S) + mK(S) \leq \gamma(\alpha G(S) + mK(S))$ for all $S \subseteq V$. If $q \in \mathcal{P}(G + mK)$ then $\frac{q}{\alpha\gamma} \in \mathcal{P}\left(G + \frac{mK}{\alpha}\right)$.*

Proof: Let S be an arbitrary subset of V . We are given that $\sum_{v \in S} q_v \leq G(S) + mK(S)$. Therefore

$$\begin{aligned} \sum_{v \in S} \frac{q_v}{\alpha\gamma} &\leq \frac{G(S) + mK(S)}{\alpha\gamma} \\ &\leq \frac{\gamma(\alpha G(S) + mK(S))}{\alpha\gamma} \\ &= G(S) + \frac{mK(S)}{\alpha}. \end{aligned}$$

Since S was chosen arbitrarily we have that $\frac{q}{\alpha\gamma} \in \mathcal{P}\left(G + \frac{mK}{\alpha}\right)$.

This completes then proof of Claim 1.1.

Now the proof of Theorem 1 is completed as follows. Recall that q^* and h'^* are the optimal q and h' in optimization problem (6).

$$C^*(L^I) = \sum_{v \in V} 2\sqrt{q_v^*(L^I) h'_v{}^*(L^I)} \quad (10)$$

$$\leq \text{optimal policy cost for the } L^I = G + \frac{mK}{\alpha} \text{ system} \quad (11)$$

$$\leq \text{optimal policy cost for the } L^J = G + mJ \text{ system} \quad (12)$$

$$\leq \text{optimal power of two policy cost for the } L^J \text{ system} \quad (13)$$

$$\leq \text{optimal power of two policy cost for the } L^u = G + mK \text{ system} \quad (14)$$

$$< 1.021 \left(\sum_{v \in V} 2\sqrt{q_v^*(L^u) h_v'^*(L^u)} \right) \quad (15)$$

$$= 1.021 \sqrt{\alpha\gamma} \left(\sum_{v \in V} 2\sqrt{\frac{q_v^*(L^u)}{\alpha\gamma} h_v'^*(L^u)} \right) \quad (16)$$

$$\leq 1.021 \sqrt{\alpha\gamma} \left(\sum_{v \in V} 2\sqrt{q_v^*(L^l) h_v'^*(L^l)} \right). \quad (17)$$

(11) follows from (10) by a result from Federgruen et. al. (1989) mentioned above. (12) follows from (11) because $\frac{K(S)}{\alpha} \leq J(S)$. (13) follows from (12) because the set of allowable policies in (13) is strictly smaller. (14) follows from (13) because $J(S) \leq K(S)$. (15) follows from (14) by a result from Federgruen et. al. (1989) mentioned above. (16) follows from (15) through simple algebra.

From claim 1.1 we know that if $q^*(L^u) \in \mathcal{P}(G + mK)$ then $\frac{q^*(L^u)}{\alpha\gamma} \in \mathcal{P}\left(G + \frac{mK}{\alpha}\right)$. Clearly $q^*(L^u) \in \mathcal{P}(G + mK)$; in fact it is the optimal q for optimization problem (6) for the L^u system. Thus $\frac{q^*(L^l)}{\alpha\gamma}$ is feasible for optimization problem (6) for the L^l system. We also observe that if a h' is feasible for optimization problem (6) with one order cost function then it is feasible for optimization problem (6) with any order cost function. (17) follows from (16) because $\frac{q^*(L^u)}{\alpha\gamma}$ and $h'^*(L^u)$ are feasible for optimization problem (6) for the L^l system and thus the term in parentheses in equation (16) is less than or equal to the optimal value for optimization problem (6) for the L^l system.

The cost of the optimal policy for the L system is greater than (11) and less than (14). Thus it is clearly bounded from below by $C^*(L^l)$, the optimal value of (6) for the L^l system, and bounded above by the cost of an optimal power of two

policy for the L^u system. Since (10) and (17) differ by at most $1.021\sqrt{\alpha\gamma}$, we know that (10) and (14) differ by at most $1.021\sqrt{\alpha\gamma}$. Suppose that $J(S) = K(S)$ and that we compute an optimal power of two policy for the L^J system and use it in the L system. The resulting cost is less than (14) because $L(S) \leq L^J(S)$ for all $S \subseteq V$. The same proof holds when we fix b , except that 1.061 is substituted for 1.021.

□

Theorem 1 holds independent of our knowledge of α , θ , and η . This is the reason two types of tightness were investigated in Herer (1990). If one had a $K(S)$ satisfying (9), and felt that $\beta = \min_{S \subseteq V} \frac{K(S)}{T(S)}$ was greater than one, one could use $J(S) = \eta K(S)$ for any η , $\frac{1}{\alpha} \leq \eta \leq 1$ to compute the power of two reorder intervals, and thus obtain a better policy.

Corollary 1.1 *The cost of the optimal power of two policy for the L system is less than $1.021\sqrt{\alpha\gamma}$ (or $1.061\sqrt{\alpha\gamma}$ for fixed b) times the cost of any feasible policy for the L system.*

This is because the cost of the optimal power of two policy for the L system is greater than (11) and less than (14).

As an approach to solving general production/distribution problems with order costs that are close to submodular we suggest determining the optimal power of two reorder intervals in the L^J system. We then suggest using these reorder intervals in the L system.

3 One Warehouse Multi-Retailer Heuristics

In this section we study the one warehouse multi-retailer distribution system. We develop heuristics to solve this problem. We note that this system is a special case of the general production/distribution system discussed in the previous section. We therefore will relate the bounds developed there to the one warehouse multi-retailer distribution problem.

Recall that holding costs are calculated using the echelon inventory method. The echelon holding cost rate at the warehouse (h_i^w) for retailer i 's goods is identical to the traditional holding cost rate at the warehouse. The echelon holding cost rate at retailer i (h_i) is equal to the traditional holding cost rate at retailer i minus h_i^w . We can allow h_i^w to differ for different retailers by adjusting the $\lambda_{r_0 r_i}$ s.

Recall that the order cost function $L(S)$ has two components. The first is $G(S)$, a general monotone non-negative submodular set function, while the second is $mT(S)$ the travelling salesman tour length through the retailers in S and the central warehouse, multiplied by a cost per mile travelled. Recall that $G(S)$ might not be equal to $G(S \cup \{r_0\})$, but $T(S) = T(S \cup \{r_0\})$.

One very broad class of functions that are non-negative monotone submodular are the ones that fit into the non-negative family model (see Roundy 1986). A set function \mathcal{F} is said to fit into a non-negative family model if $\mathcal{F}(S)$ can be written as

$$\mathcal{F}(S) = \sum_{M \subseteq V} 1_{(S \cap M \neq \emptyset)} F(M),$$

$$\text{where } 1_R = \begin{cases} 1 & \text{if relation } R \text{ is true} \\ 0 & \text{if relation } R \text{ is false} \end{cases}$$

and F is any set function that maps subsets of V to \mathfrak{R}^+ . Some examples of submodular costs that fit into the non-negative family model and their associated non-zero F values are:

- Cost of hiring a truck if any retailer is ordering. $F(V \setminus \{r_0\}) =$ the cost of hiring the truck.
- Cost of making a stop at retailer i if retailer i orders, and/or the cost of processing retailer i 's order. $F(\{r_i\}) =$ the cost of retailer i ordering.
- Cost of the warehouse ordering. $F(\{r_0\}) =$ the warehouse order cost.

Non-negative linear combinations of monotone non-negative submodular functions are monotone non-negative submodular functions. The same applies to set functions that fit into the family model. Thus any non-negative linear combination of the above functions is a monotone non-negative submodular function that fits into the family model. There are monotone non-negative submodular functions that do not fit into the non-negative family cost model. A simple example is given by Queyranne (1985).

We will develop heuristics to compute power of two policies. By a power of two policy we mean that facility i , $0 \leq i \leq N$ orders once every $\tau_i = 2^z b$ time units,

starting at time zero. Here z_i is a positive or negative integer. Recall that b is some fixed base planning period.

The long run average cost of an arbitrary policy for the one warehouse multi-retailer distribution problem can be very hard to calculate. However the cost of a power of two policy has a simple functional form. Let σ be a permutation of the indices $\{0, 1, \dots, N\}$ such that $\tau_{\sigma(0)} \leq \tau_{\sigma(1)} \leq \dots \leq \tau_{\sigma(N)}$. Let d_i be the demand rate for retailer i . Let $\frac{1}{\tau_{\sigma(N+1)}} \equiv 0$. For one warehouse multi-retailer systems (2) becomes

$$\sum_{i=1}^N \frac{1}{2} d_i h_i^w \max(\tau_i, \tau_0) + \sum_{i=1}^N \frac{1}{2} d_i h_i \tau_i + \sum_{i=0}^N \left(\frac{1}{\tau_{\sigma(i)}} - \frac{1}{\tau_{\sigma(i+1)}} \right) L \left(\bigcup_{j=0}^i \{r_{\sigma(j)}\} \right). \quad (18)$$

The echelon method is used to compute the holding costs. The first term represents the holding cost charged at the warehouse rate for all items in the system (both at the warehouse and at the retailer) that will eventually be sold by retailer i . The maximum is taken because the warehouse orders goods for retailer i once every $\max(\tau_i, \tau_0)$ time units. The second term represents the holding cost charged at the retailer rate for all items at the retailer. The third term represents the order cost.

The cost of an optimal power of two policy is obtained by minimizing (18) subject to

$$\tau_i = 2^{z_i} \quad z_i \in \{\dots, -2, -1, 0, 1, 2, 3 \dots\} \text{ for all } 0 \leq i \leq N.$$

If $T(S)$ were non-negative monotone submodular and could be evaluated in polynomial time then the one warehouse multi-retailer distribution problem with

travelling salesman tour vehicle routing costs would, in all practicality, be solved (see Federgruen, et. al. 1989, Queyranne 1985, and Zheng 1987). However $T(S)$ can not be evaluated in polynomial time unless $\mathcal{P} = \mathcal{NP}$. Even if $\mathcal{P} = \mathcal{NP}$ the set function $T(S)$ is not submodular, even in the euclidean plane (see Anily and Federgruen 1988). A simple example is found in Figure 2, where all horizontal and vertical distances between neighboring facilities are one. $T(S \cup L \cup M) = 3\sqrt{2} + 2$, $T(S \cup L) = 2\sqrt{2} + 2$, $T(S \cup M) = 2\sqrt{2} + 2$, $T(S) = 4$. Thus the left hand side of (1) is equal to 1.41, while the right hand side of (1) is equal to .83.




Figure 2: An example showing $T(S)$ is not submodular.

But $T(S)$ is close to submodular. This has been shown in Herer (1990). Four heuristics for estimating $T(S)$ by a submodular function were proposed. For these heuristics the worst case values of α were investigated theoretically and the average values of α were investigated computationally. The submodular estimates of α were each non-negative monotone. These results combined with Theorem 1 of Section 2

suggest a heuristic for the one warehouse multi-retailer distribution problem with travelling salesman tour vehicle routing costs.

We note here one very useful fact. The four heuristics for approximating $T(S)$ by a submodular function can be modeled by the family model described above. In Appendix A we give the families and their costs.

Our heuristic for the one warehouse multi-retailer problem is as follows. One selects one of the non-negative monotone submodular functions found in Herer (1990) and uses it as $J(S)$. Since $J(S)$ is the length of a feasible tour, $J(S) = K(S)$ and $\eta = 1$. The function $L^J(S) = G(S) + mJ(S)$ is used to compute optimal power of two reorder intervals. To compute the power of two reorder intervals we exploit the fact that the approximate order costs $L^J(S)$ fit into the family model. Thus we are able to use the divide and conquer algorithm found in Maxwell and Muckstadt (1985) to compute reorder intervals. These reorder intervals will then be used in the L system. If desired the upper and lower bounds of Theorem 1 can also be computed using the divide and conquer algorithm.

If one wishes to further reduce their costs, then instead of using the reorder intervals calculated assuming that the true costs were $L^J(S)$ one could use more of the available information. We note that the divide and conquer algorithm minimizes (18) subject only to $\tau_i \geq 0$. This gives the “natural” reorder intervals. The natural reorder intervals are then rounded to powers of two to get a power of two policy (see Roundy 1985). One could partition the facilities into clusters according

to their natural reorder intervals. One could then take the clusters created by the divide and conquer algorithm and use the true cost $L(S)$ to compute new natural reorder intervals, which would then be rounded as before. We let the set S be the set of all facilities in a particular cluster, and we let the set M be the set of all of the facilities that have strictly smaller reorder intervals than the reorder intervals of the facilities in S . Clearly $S \cap M = \emptyset$. Then the facilities in the set S have new natural reorder intervals equal to $\sqrt{\frac{L(S \cup M) - L(M)}{\mathcal{H}(S)}}$. Here $\mathcal{H}(S) = \sum_{\{i|r_i \in S\}} \frac{1}{2} h_i d_i$ if $r_0 \notin S \cup M$; $\mathcal{H}(S) = \sum_{\{i|r_i \in S\}} \frac{1}{2} (h_i + h_i^w) d_i + \sum_{\{i|r_i \in M\}} \frac{1}{2} h_i^w d_i$ if $r_0 \in S$; and $\mathcal{H}(S) = \sum_{\{i|r_i \in S\}} \frac{1}{2} (h_i + h_i^w) d_i$ if $r_0 \in M$. However one must be careful to preserve order. That is if the natural reorder interval for facility i is less than the natural reorder interval for facility j in the divide and conquer solution, the same relationship must hold for the new natural reorder intervals. This is due to the way costs are modeled. The cluster based formula for the costs incurred becomes invalid if order is not preserved (see Federgruen, et. al. 1989).

We use this method of further reducing costs in all of our computational results found in Section 5. We recompute the reorder intervals starting at the smallest reorder interval and work our way up to the largest reorder interval. If we run into a cluster of facilities whose new natural reorder interval is less than a new natural reorder interval already calculated, then we set the reorder interval for the facilities in the cluster equal to the largest new natural reorder interval already calculated.

One can use theoretical worst case bounds on α to come up with worst case

bounds on the relative cost of our heuristic and the optimal policy for the one warehouse multi-retailer distribution problem. In Herer (1990) worst case expressions were derived for α for each of the three estimates of $T(S)$, the Loop estimate, the Box estimate, and the Circle estimate. Using Theorem 1 we obtain the following worst case bounds on the deviation from the optimal policy cost that each heuristic would yield.

- **Loop:** $102.1\sqrt{\left\lceil\frac{N}{3}\right\rceil}$ percent.
- **Box:** $102.1\sqrt{\sqrt{2}\left\lceil\log_{\frac{4}{3}}(N)\right\rceil(2\lceil\log_2(N+1)\rceil+1)} = O(\log N)$ percent.
- **Circle:** $102.1\sqrt{\sqrt{\frac{N\pi}{2} + \frac{3\pi}{2}}} = O(\sqrt[4]{N})$ percent.

For N in the interval $[1, 36]$ the Loop heuristic has the best worst case bound. The Circle heuristic is dominant for N ranging from 37 to somewhere in excess of 8,000,000, and the Box heuristic is dominant for $N > 8,000,000$. So in a worst case sense Box does not seem to be of practical interest. Some sample values of N and the associated worst case bounds for the three heuristics is contained in Table 1. In Herer (1990) one can see that the average values of α are much smaller than the worst case values used in Table 1, and hence the bound on the cost of the policy developed is on average much smaller than the worst case. We will see in Section 5 that the actual percent deviation from optimality is less than both the worst case analysis and the computational results of Herer (1990) would lead us to believe.

Table 1: Sample worst case bounds on the percent cost deviation from optimal.

N	Loop	Box	Circle
10	204	1030	301
20	270	1336	328
50	421	1638	376
100	595	1881	424

4 Dynamic Programming

In this section we propose a dynamic programming algorithm for the one warehouse multi-retailer distribution problem. The dynamic program will find optimal power of two policies given a fixed base planning period assuming only that the order costs are non-negative and monotone. In light of the results in Section 2 we feel that optimal power of two schedules could be very useful.

We note that the joint replenishment problem is a special case of the one warehouse multi-retailer distribution problem. See Federgruen and Zheng (1988) for a discussion of the joint replenishment problem with submodular order costs. The difference between the two problems is that in the joint replenishment problem the warehouse can not hold inventory. Queyranne (1987) has proposed a dynamic programming algorithm to find optimal equal reorder interval schedules for the joint replenishment problem with general order costs. His solution assigns each

product to a group and each group is given its own reorder interval. He assumes that costs are additive over groups. If one thinks of a schedule merely as a listing of products and their reorder intervals, then this method does not find an optimal equal reorder interval schedule (Goyal 1987). The basic reason for this is that if S and M are different groups, and at a given point in time orders are placed by the retailers in $S \cup M$, the algorithm would charge a cost of $L(S) + L(M)$ rather than $L(S \cup M)$.

In this section we propose a dynamic programming algorithm to compute optimal power of two schedules, given a fixed base period b , for the one warehouse multi-retailer distribution problem. Our dynamic programming algorithm has the same time and space complexity as Queyranne (1987). The algorithm requires only that the cost of ordering $L(S)$ be monotone and non-negative. We do not claim that these schedules are optimal in any global sense. We do not even claim that they produce optimal cyclic schedules. However, in light of the possible cost savings in combining orders and the results of Section 2, we feel that power of two policies are usually very good. In addition, since this algorithm produces, after backtracking, optimal power of two schedules, Corollary 1.1 applies even if a specific submodular estimate of $L(S)$ is not known.

With highly non-submodular order costs, power of two schedules can be arbitrarily bad as the following example shows. Consider the joint replenishment problem with $N = 2$, $L(\{r_1\}) = L(\{r_2\}) = 0$, and $L(\{r_1, r_2\}) = \infty$. Then clearly,

as long as the holding costs are non-negative and finite, an optimal schedule is to continuously replenish r_1 and r_2 one after the other, but never at the same time. The inventory levels are virtually zero, and essentially no costs are incurred. With a power of two schedule r_1 and r_2 will be replenished together at some point in time, yielding an infinite cost. However we feel that cost functions such as these are not common in practice, particularly in regard to routing problems, and that we can expect power of two policies to work very well for real world problems.

Recall that h_i is the echelon holding cost rate at retailer i , H_i is the traditional holding cost rate at retailer i , and that d_i is the demand rate for retailer i . We define $h(S) = \sum_{i \in S} \frac{1}{2} h_i d_i$ and $H(S) = \sum_{i \in S} \frac{1}{2} H_i d_i$. We define $r(x) = 2^{\lceil \log_2(\frac{x}{b}) + .5 \rceil} b$. The function r is used to round a reorder interval to the (geometrically) closest power of two times the base planning period.

Let $C_l(S)$ be the long run average cost for the one warehouse multi-retailer distribution problem, ignoring costs incurred at the warehouse and all the retailers not in S , and assuming that the warehouse has a predetermined reorder interval that is greater than that of any of the reorder intervals of the retailers in S . Hence we only look at the echelon holding costs at the retailers in S and the ordering costs for the retailers in S . We also define $\tau_l(S, M) = r \left(\sqrt{\frac{L(S) - L(M)}{h(S \setminus M)}} \right)$, where M is a subset of S . This represents the optimal reorder interval for the set $S \setminus M$ if all retailers in the set $S \setminus M$ have the same reorder interval and all of the retailers in the set M have a strictly smaller reorder interval. $C_l(S)$ is computed for all $S \subseteq V$

as follows.

$$C_l(S) = \min_{\{M \subset S \mid \tau_l(S, M) > \tau_l^*(M)\}} \left(C_l(M) + \frac{L(S) - L(M)}{\tau_l(S, M)} + \tau_l(S, M)h(S \setminus M) \right).$$

If M obtains the minimum then we let $\tau_l^*(S) = \tau_l(S, M)$.

Boundary condition: $C_l(\emptyset) = 0$ and $\tau_l^*(\emptyset) = 0$.

This recursion can be computed in any order as long as $C_l(M)$ is computed before $C_l(S)$ for all $M \subset S$. For example one could compute $C_l(S)$ for all sets of a given cardinality, from one up to N . Note that $S \subseteq V$ and that the warehouse is not in the set V . For the retailers in S , $C_l(S)$ includes the costs in the second and third terms of (18).

The reason this recursion works is that for each set S we consider the possibility of the set always ordering together ($M = \emptyset$). We also consider the possibility that some subset ($S \setminus M$) of the retailers in S has a larger reorder interval than another subset (M) of the retailers. This occurs when the optimal set M satisfies $\emptyset \neq M \neq S$.

We now define two more recursions to obtain the optimal power of two reorder intervals, assuming a fixed base planning period. The first computes $C_e(S)$, which represents the optimal cost considering only the warehouse and the retailers in S , given that the warehouse has a reorder interval which is at least as large as the largest reorder interval of any retailer. The second computes $C_g(S)$, which is the cost of an optimal power of two schedule for the one warehouse multi-retailer distribution problem considering only the warehouse and the retailers in S . The

function r is as defined above. We define $\tau_e(S, M) = r \left(\sqrt{\frac{L(S \cup \{r_0\}) - L(M)}{H(S \setminus M) + H(M) - h(M)}} \right)$ and $\tau_g(S, M) = r \left(\sqrt{\frac{L(S \cup \{r_0\}) - L(M \cup \{r_0\})}{H(S \setminus M)}} \right)$.

$$C_e(S) = \min_{\{M \subseteq S \mid \tau_e(S, M) > \tau_l^*(M)\}} \left(C_l(M) + \frac{L(S \cup \{r_0\}) - L(M)}{\tau_e(S, M)} + \tau_e(S, M)(H(S \setminus M) + H(M) - h(M)) \right).$$

$$C_g(S) = \min \left(C_e(S), \min_{\{\emptyset \neq M \subset S \mid \tau_g(S, M) > \tau_g^*(M)\}} \left(C_g(M) + \frac{L(S \cup \{r_0\}) - L(M \cup \{r_0\})}{\tau_g(S, M)} + \tau_g(S, M)H(S \setminus M) \right) \right).$$

The variables $\tau_e^*(S)$ and $\tau_g^*(S)$ have definitions parallel to that of $\tau_l^*(S)$, i. e., if M obtains the minimum in the recursion for $C_e(S)$ then we let $\tau_e^*(S) = \tau_e(S, M)$. The same holds for $\tau_g^*(S)$. Again as boundary conditions we let $\tau_e^*(\emptyset) = 0$ and $\tau_g^*(\emptyset) = 0$. The recursions can be computed in any order as long as $C_l(M)$ is computed before $C_e(S)$ for all $M \subseteq S$ and $C_g(M)$ and $C_e(S)$ are computed before $C_g(S)$ for all $M \subset S$. For example one could first compute $C_e(S)$ and then compute $C_g(S)$ for all sets of a given cardinality, from one up to N .

In the recursion for $C_e(S)$, setting $M = \emptyset$ represents the central warehouse and all of the retailers in S having the same reorder interval. Setting $M = S$ represents the central warehouse having a reorder interval strictly greater than that of all of the retailers in S . In the recursion for $C_g(S)$, setting $M = \emptyset$ is disallowed because setting $\tau_0 < \tau_i$ for all $i \geq 1$ is strictly sub-optimal.

The quantity $C_g(V)$ represents the cost of the optimal power of two policy for the one warehouse multi-retailer distribution problem. The actual reorder intervals

can be computed via a straight forward backtracking procedure.

To use our recursions for the joint replenishment problem we set $C_i(S) = \infty$ for all $S \neq \emptyset$. Thus $C_e(S) = \frac{L(S \cup \{r_0\})}{\tau_e(S, \emptyset)} + \tau_e(S, \emptyset)H(S)$ for all $S \subseteq V$. This is the appropriate modification of our recursions because $\tau_i < \tau_0$ is not allowed in the joint replenishment problem.

5 Computational Study

In this section we investigate the usefulness of the techniques we have developed. We do this by generating random one warehouse multi-retailer distribution problems and solving them. For small problems ($N \leq 16$) we use the dynamic programming recursion in Section 4 to compute the optimal power of two policies and their associated costs. To compute $T(S)$ we use a dynamic programming recursion that is almost identical to that of Held and Karp (1962). We compare these optimal power of two policies to the power of two policies produced by the heuristic of Section 3. The heuristic solutions were obtained using an algorithm developed by Maxwell and Muckstadt (1985). We were able to use this algorithm rather than the general algorithm for production/distribution systems with submodular costs, because all of the submodular estimates of the order costs fit into the family model (see Appendix A). If the order costs had not fit into the family model then we would have had to use one of the more general algorithms found in Queyranne (1985) and Zheng (1987). These are polynomial but require either mul-

multiple submodular function minimizations or multiple polymatroid maximum flow computations. Hence the computations and the programming time were greatly reduced by using the structure of the costs.

We look at systems in which the location of the retailers are generated uniformly in the unit circle and in the unit square. The central warehouse is located in the center of the circle or square. We looked at both circles and squares because the expected value of the number of retailers on the convex hull of the points differ dramatically for circles and squares (Edelsbrunner 1987 p. 174). The order cost has two components. The first is a warehouse order cost. This cost is simply a fixed cost that is incurred each time the warehouse orders. The second component is $mT(S)$. Since the holding costs grow linearly in the number of retailers N and the travelling salesman tour costs grow as \sqrt{N} (Beardwood, et. al. 1959) we set $m = \sqrt{N}$ so that the order costs remain at the same magnitude relative to the holding costs. We did not put any other submodular component into the order costs so as to accentuate the possible optimality loss caused by the non-submodularity of $T(S)$.

The warehouse order cost was generated from a uniform distribution on the interval $[.5, 1.5]$. Recall that d_i is the demand rate at retailer i , h_i is the echelon holding cost rate at retailer i , and h_i^w is the holding cost rate at the warehouse for retailer i 's goods. The factors $\frac{1}{2}h_i d_i$ and $\frac{1}{2}h_i^w d_i$ were generated from a uniform distribution on the interval $[.1, 1]$ and then cubed. The uniform distribution did

not include zero as to avoid infinitely large reorder intervals. The distribution was cubed in order to create a more uniform spread of optimal reorder intervals. Without cubing most facilities tended to have the same reorder interval. We also set b , the base planning period to one.

The results are organized by the number of retailers in the problem instance (referred to as problem size) and by the submodular estimate of $T(S)$ used. For each problem size we generated twenty problem instances, and solved each instance using each of the four submodular estimates. We generate ten instances in which the retailers' locations were generated uniformly in the unit circle. We also generated ten instances in which the retailers' locations were generated uniformly in the unit square.

The results are summarized in Tables 2 through 5. In each of these tables we report $100 \left(\frac{\text{heuristic cost}}{\text{optimal power of two cost}} - 1 \right)$. The column marked Naive represents the optimal cost, assuming all of the facilities order together. This is the policy that (Anily and Federgruen 1988) would yield for uncapacitated trucks. We note however that our models are not identical. Among other differences Anily and Federgruen (1988) attempt to solve the joint replenishment problem, whereas we attempt to solve the one warehouse multi-retailer distribution problem.

The run times were dominated by the dynamic programming recursion used to calculate the optimal power of two reorder intervals. The runs were performed on a Sun Sparc station. The run times ranged from almost instantaneous for $N = 1$

Table 2: Mean percent error for retailers distributed in the unit circle.

N	Circle	Loop	Box	MST	Naive
16	24.5	3.0	4.4	1.0	3.7
15	25.4	1.5	6.7	2.3	3.2
14	19.9	2.0	5.5	1.7	2.4
13	29.8	1.6	4.8	1.6	3.3
12	26.0	2.2	6.9	1.3	3.0
11	22.4	2.1	5.0	1.5	4.4
10	23.9	1.6	5.1	0.8	3.0
9	18.8	1.4	4.9	0.8	4.8
8	18.9	0.2	5.0	1.3	3.9
7	23.1	1.1	2.5	0.5	7.4
6	20.2	0.4	2.6	1.8	3.4
5	21.2	1.1	2.1	0.2	6.1
4	23.6	0.6	2.1	0.1	7.4
3	25.4	0.0	0.8	0.1	9.0
2	24.4	0.0	0.5	0.0	2.2
1	26.1	0.0	0.0	0.0	7.9

Table 3: Maximum percent error for retailers distributed in the unit circle.

N	Circle	Loop	Box	MST	Naive
16	42.2	8.9	9.6	2.0	7.0
15	40.9	4.2	11.5	9.7	6.9
14	33.2	6.3	14.7	3.4	5.2
13	43.3	6.7	10.9	4.1	7.8
12	35.6	5.8	20.6	4.6	4.6
12	35.7	5.1	10.6	3.9	7.6
10	36.6	6.4	12.4	3.7	15.4
9	28.2	6.3	12.1	2.6	12.4
8	30.8	1.9	9.7	4.1	7.0
7	37.3	4.9	7.3	1.5	17.6
6	38.8	2.6	4.5	4.5	9.8
5	39.9	4.6	9.1	1.3	28.8
4	45.2	4.1	4.1	0.5	22.5
3	54.1	0.0	7.3	0.8	25.7
2	42.8	0.0	2.9	0.0	14.6
1	70.8	0.0	0.0	0.0	46.9

Table 4: Mean percent error for retailers distributed in the unit square.

N	Circle	Loop	Box	MST	Naive
16	26.9	3.2	9.0	1.7	2.4
15	29.3	2.0	6.0	1.9	1.3
14	27.3	3.0	7.7	1.6	1.8
13	23.9	2.6	9.2	2.3	2.7
12	24.9	3.1	6.4	1.8	4.3
11	18.0	3.5	5.1	2.0	3.2
10	17.6	3.8	4.4	1.8	3.3
9	20.3	1.2	5.1	3.0	3.9
8	13.4	1.7	4.7	2.0	2.6
7	18.4	2.4	3.1	1.4	6.6
6	20.0	1.5	5.0	1.5	5.3
5	14.9	0.2	1.7	1.4	3.3
4	10.3	1.1	2.1	0.9	4.4
3	21.5	0.0	1.3	0.4	15.4
2	20.8	0.0	0.4	0.2	6.9
1	16.5	0.0	0.0	0.0	4.8

Table 5: Maximum percent error for retailers distributed in the unit square.

N	Circle	Loop	Box	MST	Naive
16	32.8	10.6	21.7	2.7	5.0
15	37.0	9.0	13.5	3.4	2.8
14	36.9	9.4	15.5	4.9	4.4
13	29.6	6.7	18.4	7.8	12.8
12	34.9	9.4	12.6	4.7	16.6
11	25.5	6.9	11.7	6.2	10.5
10	25.9	9.9	9.5	5.7	18.4
9	34.8	3.0	12.6	4.8	10.6
8	25.7	5.1	9.7	6.1	12.0
7	39.7	11.4	9.3	4.5	17.2
6	41.5	5.4	11.7	4.0	14.0
5	37.1	1.1	8.3	7.6	7.6
4	25.3	5.7	7.5	5.2	11.9
3	43.1	0.0	4.3	2.7	46.8
2	42.6	0.0	4.0	1.6	27.0
1	29.9	0.0	0.0	0.0	22.5

to three hours for $N = 15$ to twenty hours for $N = 16$. One reason for the large jump in run time when going from $N = 15$ to $N = 16$ is the paging required when virtual memory is used.

As can be seen from the tables, the submodular estimates from the Loop and MST heuristics seem to give the best results. We feel that this is because these heuristics use the most detailed information on the actual locations of the retailers. The Box heuristic uses this same information. We feel its poor performance is due to the fact that this heuristic creates artificial divisions without regard to where the retailers are actually located.

We also performed experiments for points distributed uniformly in the unit circle using the same distributions for the parameters, for $N = 20, 40, 60, 80, 100$. Because of the size of N , neither optimal power of two policies nor optimal tours were calculated. We calculated the order costs using the minimal spanning tree heuristic with short cuts, and running the tour thus obtained through a 2-opt procedure. These results are presented in Table 6. For Table 7 we evaluated the order costs using the nearest neighbor heuristic run through a 2-opt procedure. We used the minimum of the tour lengths produced by these two methods for the remaining tables. See Lawler et. al. (1985) for a more complete discussion of these heuristics. Note that the minimal spanning tree method just described is different from the submodular estimate of $T(S)$ which we call MST. For the latter we find the minimal spanning tree through all of the facilities and use twice the

size of the subtree induced by the set $S \cup \{r_0\}$ as our estimate of $T(S)$. For the former, we compute the minimal spanning tree for each set S for which we wish to obtain a tour length and improve it by taking short cuts and running it through a 2-opt procedure. Below we report the mean over ten problems of the factor $100 \left(\frac{\text{heuristic cost}}{\text{best heuristic cost}} - 1 \right)$.

The run times for all the heuristics were small, especially when one considers that this is a planning algorithm. The longest run times were for $N = 100$ for the Loop heuristic. For this heuristic it took approximately three and a quarter minutes to compute the family structure of the order costs, five minutes to compute the reorder intervals and three and three quarter minutes to construct the routes. The next longest run times for $N = 100$ were required by the MST heuristic. It required less than five seconds to construct the family structure, approximately two minutes to compute the reorder intervals and three and three quarter minutes to construct the routes. All calculations were done on a Sun Sparc station, except the reorder intervals were calculated on a personal computer. We wish to express our appreciation to Peter Jackson for lending us his code to determine reorder intervals.

For each of the different scenarios in Tables 6 through 13 we used identical random number seeds. The retailer order costs, for the scenarios that have retailer order costs, were generated from a uniform distribution on the interval $[.1,1]$ and then cubed. The retailers that were generated with exponential distances from the

warehouse were generated so that the average distance from the central warehouse was $\frac{1}{2}$. The scenarios that had large order costs (both retailer and warehouse) merely had their order costs multiplied by a factor of ten. The tables clearly indicate that MST is the most robust of all the heuristics. The other heuristics do a fairly good job of determining reorder intervals for one situation or another, but no heuristic other than MST does well for all of the different scenarios.

Table 6: Mean percent above minimum for retailers distributed in the unit circle using minimal spanning tree.

N	Circle	Loop	Box	MST	Naive
100	24.2	9.8	12.1	0.2	1.4
80	20.2	8.0	11.0	0.2	1.4
60	14.1	6.6	8.8	0.7	1.3
40	8.8	1.7	4.3	0.2	2.5
20	4.2	1.2	4.4	0.3	2.9

Table 14 gives the approximate percentages that the different cost components contributed to the total cost for the different scenarios. The values are only approximate because they changed with different heuristics and with different problem sizes. A change in heuristic and problem size usually had little effect on the percentages, but for some scenarios the change in heuristic had a large effect. For example for the scenario represented by Table 13 most heuristics had retailer order

Table 7: Mean percent above minimum for retailers distributed in the unit circle using nearest neighbor.

N	Circle	Loop	Box	MST	Naive
100	24.6	9.6	12.1	0.5	1.7
80	20.9	8.6	11.7	0.5	1.3
60	14.3	6.6	8.5	0.8	1.1
40	8.8	1.9	4.5	0.3	1.8
20	4.5	1.3	4.2	0.4	3.5

Table 8: Mean percent above minimum for retailers distributed in the unit circle using the best tour.

N	Circle	Loop	Box	MST	Naive
100	24.7	9.7	12.2	0.4	1.6
80	21.2	8.6	11.7	0.5	1.4
60	14.5	6.7	8.8	0.8	1.2
40	9.1	2.0	4.7	0.5	1.9
20	4.3	1.2	4.4	0.3	3.1

Table 9: Mean percent above minimum for retailers distributed in the unit circle with retailer order costs.

N	Circle	Loop	Box	MST	Naive
100	3.6	15.3	2.8	0.0	7.4
80	2.2	11.8	2.0	0.0	5.6
60	2.1	10.7	1.8	0.0	5.7
40	3.0	8.9	2.5	0.1	5.8
20	8.4	3.3	2.9	0.2	3.9

Table 10: Mean percent above minimum for retailers distributed in the unit circle with large retailer order costs.

N	Circle	Loop	Box	MST	Naive
100	1.6	133.7	0.4	0.0	20.5
80	2.0	120.2	0.2	0.0	20.7
60	3.8	108.7	0.3	0.1	21.7
40	4.0	96.7	0.9	0.1	21.0
20	4.4	57.5	1.0	0.1	15.1

Table 11: Mean percent above minimum for retailers distributed with exponential distances.

N	Circle	Loop	Box	MST	Naive
100	22.1	5.5	8.7	0.0	14.3
80	20.3	4.0	7.7	0.0	16.2
60	24.3	2.6	7.5	0.0	12.9
40	26.3	3.2	7.0	0.0	12.5
20	20.5	1.3	4.9	0.6	7.2

Table 12: Mean percent above minimum for retailers distributed in the unit circle with large warehouse order costs.

N	Circle	Loop	Box	MST	Naive
100	8.7	5.8	6.5	0.5	1.3
80	4.2	2.4	2.3	0.1	2.7
60	2.3	1.5	1.8	0.0	2.6
40	2.7	0.7	2.0	0.1	2.1
20	15.6	0.3	3.6	0.2	3.9

Table 13: Mean percent above minimum for retailers distributed with exponential distances, large retailer order costs, and large warehouse order costs.

N	Circle	Loop	Box	MST	Naive
100	2.8	115.4	0.4	0.0	17.6
80	4.2	93.4	0.6	0.0	16.7
60	4.3	79.6	0.8	0.0	17.0
40	3.4	89.1	0.5	0.0	17.6
20	3.1	53.5	0.4	0.0	11.1

costs accounting for thirty five percent of the total cost, but the Loop heuristic had retailer order costs accounting for eighty five percent of the total cost.

6 Future Research

We have investigated the one warehouse multi-retailer distribution system with travelling salesman tour vehicle routing costs. We have proposed heuristics for finding power of two reorder intervals for these systems. We have shown that the ratio of the cost of the power of two policy calculated by these heuristics and the cost of an optimal policy is bounded by a factor that depends on the closeness of the travelling salesman tour length $T(S)$ to submodularity (see Herer 1990). The theoretical bounds apply to any production/distribution system that

Table 14: Approximate percentage of the costs' contribution to the total for the different scenarios.

Scenario	Holding	TSP tour	retailer order	warehouse order
8	65	30	0	5
9	55	20	20	5
10	50	5	40	5
11	65	25	0	5
12	65	20	0	15
13	55	5	35	5

has order costs which are close to being submodular. Heuristics that are similar to our one warehouse multi-retailer heuristic could easily be developed for any production/distribution system that has order costs which are close to submodular, and bounds on the deviation from optimality could be given.

We have also presented a dynamic programming algorithm to compute optimal power of two reorder intervals for one warehouse multi-retailer distribution systems with arbitrary monotone non-negative costs.

More work still has to be done in this line of research. In particular we feel more work can be done to develop heuristics for other inventory systems that have non-submodular order costs. One such inventory system that immediately comes to mind is where the order costs are based on the the length a truck must travel to

visit all of the retailers, given that the truck has a finite capacity. i. e., the truck must return to the central warehouse to pick up more goods once it has delivered all that it can carry. This is a characteristic of the vehicle routing problem. Other inventory systems of interest also emerge from the vehicle routing literature. For example, if one retailer is only open in the evenings, and another is only open in the mornings, one obtains a vehicle routing problem with time windows.

We do not feel that our bound on the cost of our heuristic for the production/distribution problem (see Theorem 1) is tight. We feel that through a more careful analysis one may be able to develop a tighter bound.

We are developing a heuristic to solve the multi-warehouse multi-retailer distribution problem with travelling salesman tour vehicle routing costs. In this system we require that a retailer choose a single warehouse from which to receive a particular good. However a retailer can receive two different goods from two different warehouses. The warehouses order their supplies from geographically dispersed suppliers, and thus incur a travelling salesman tour length based order cost as well.

A Family Structures of the Submodular Approximations of $T(S)$

In this appendix we give the family structure of the submodular approximations to $T(S)$ that were presented in Herer (1990). We give the family structure by defining the function $F(S)$ discussed in Section 3. We handle the four heuristic approximations, Loop, Box, Circle, and MST in turn.

A.1 Loop Family Structure

The Loop heuristic partitions the retailers into loops. The family structure of a loop is bit complicated. Since the family structure of multiple loops is simply the union of the family structures of the individual loops we will focus on a single loop.

There are two types of loops. The first type is a ring. We let the retailers around the ring be labeled $(r_0, r_1, \dots, r_N, r_{N+1} = r_0)$. We define $A(r_i, r_j) = \{r_i, r_{i+1}, \dots, r_j\}$ for all $1 \leq i \leq j \leq N$. We define $F(S)$ as follows,

$$\begin{aligned}
 F(A(r_i, r_j)) &= D(r_i, r_{j+1}) + D(r_{i-1}, r_j) - D(r_i, r_j) - D(r_{i-1}, r_{j+1}) && \text{for all } 1 \leq i \leq j \leq N \\
 F(S) &= 0 && \text{otherwise.}
 \end{aligned}$$

Note that $F(S) \geq 0$ for all $S \subseteq V \cup \{r_0\}$ because of the triangle inequality.

Theorem 2 *The family costs defined above correctly model the distance around the ring.*

Proof: We assume $r_j \notin S$. We let r_i be the first retailer counter clockwise from r_j that is both on the ring and in the set S . We let r_k be the first retailer clockwise from r_j that is both on the ring and in the set S . All we need to show to prove the theorem correct is that adding retailer j to the set S incurs a family cost of $D(r_i, r_j, r_k) - D(r_i, r_k)$. Recall that $D(r_i, r_j, r_k) = D(r_i, r_j) + D(r_j, r_k)$. The family cost incurred by adding r_j to S is

$$\sum_{l=i+1}^j \sum_{m=j}^{k-1} F(A(r_l, r_m)). \quad (19)$$

We will prove that (19) is equal to $D(r_i, r_j, r_k) - D(r_i, r_k)$ through two inductions.

Claim 2.1

$\sum_{l=i+1}^j F(A(r_l, r_{k-1})) = D(r_j, r_k) + D(r_i, r_{k-1}) - D(r_j, r_{k-1}) - D(r_i, r_k)$ for all $0 \leq i < j < k \leq N + 1$.

Proof: The proof is by induction on $j - i$.

Base case ($j - i = 1$):

$$\begin{aligned} \sum_{l=i+1}^j F(A(r_l, r_{k-1})) &= F(A(r_{i+1}, r_{k-1})) \\ &= D(r_{i+1}, r_k) + D(r_i, r_{k-1}) - D(r_{i+1}, r_{k-1}) - D(r_i, r_k) \\ &= D(r_j, r_k) + D(r_i, r_{k-1}) - D(r_j, r_{k-1}) - D(r_i, r_k). \end{aligned}$$

Now we assume that the claim is true for $j - i = n$ and show that it is true for $j - i = n + 1$.

$$\sum_{l=i+1}^j F(A(r_l, r_{k-1})) = \sum_{l=i+1}^{j-1} F(A(r_l, r_{k-1})) + F(A(r_j, r_{k-1}))$$

$$\begin{aligned}
&= D(r_{j-1}, r_k) + D(r_i, r_{k-1}) - D(r_{j-1}, r_{k-1}) - D(r_i, r_k) + \\
&\quad D(r_j, r_k) + D(r_{j-1}, r_{k-1}) - D(r_j, r_{k-1}) - D(r_{j-1}, r_k) \\
&= D(r_j, r_k) + D(r_i, r_{k-1}) - D(r_j, r_{k-1}) - D(r_i, r_k).
\end{aligned}$$

Claim 2.2 $\sum_{l=i+1}^j \sum_{m=j}^{k-1} F(A(r_l, r_m)) = D(r_i, r_j, r_k) - D(r_i, r_k)$ for all $0 \leq i < j < k \leq N + 1$.

Proof: The proof is by induction on $k - i$.

Base case ($k - i = 2$):

$$\begin{aligned}
\sum_{l=i+1}^j \sum_{m=j}^{k-1} F(A(r_l, r_m)) &= F(A(r_j, r_j)) \\
&= D(r_j, r_{j+1}) + D(r_{j-1}, r_j) - D(r_j, r_j) - D(r_{j-1}, r_{j+1}) \\
&= D(r_i, r_j, r_k) - D(r_i, r_k).
\end{aligned}$$

Now we assume that the claim is true for $k - i = n$ and show that it is true for $k - i = n + 1$. We assume without loss of generality that $j \neq k - 1$. If $j = k - 1$ then we reverse the orientation of the ring.

$$\begin{aligned}
\sum_{l=i+1}^j \sum_{m=j}^{k-1} F(A(r_l, r_m)) &= \sum_{l=i+1}^j \sum_{m=j}^{k-2} F(A(r_l, r_m)) + \sum_{l=i+1}^j F(A(r_l, r_{k-1})) \\
&= D(r_i, r_j, r_{k-1}) - D(r_i, r_k) + \\
&\quad D(r_j, r_k) + D(r_i, r_{k-1}) - D(r_j, r_{k-1}) - D(r_i, r_k) \\
&= D(r_i, r_j, r_k) - D(r_i, r_k).
\end{aligned}$$

□

The second type of loop is a group of any one, two, or three retailers. Note however that any group of one or two retailers is also a ring. For the case where

we have three retailers we assume without loss of generality that $(r_0, r_1, r_2, r_3, r_0)$ is an optimal tour through the central warehouse and the retailers in S . We let the function $F(S)$ be defined as follows,

$$\begin{aligned}
F(\{r_1\}) &= D(r_0, r_1, r_2) - D(r_0, r_2) \\
F(\{r_2\}) &= D(r_1, r_2, r_3) - D(r_1, r_3) \\
F(\{r_3\}) &= D(r_2, r_3, r_0) - D(r_2, r_0) \\
F(\{r_1, r_2\}) &= D(r_1, r_3) + D(r_0, r_2) - D(r_1, r_2) - D(r_0, r_3) \\
F(\{r_2, r_3\}) &= D(r_2, r_0) + D(r_1, r_3) - D(r_2, r_3) - D(r_1, r_0) \\
F(\{r_1, r_2, r_3\}) &= D(r_1, r_0, r_3) - D(r_1, r_3) \\
F(S) &= 0 \qquad \qquad \qquad \text{otherwise.}
\end{aligned}$$

We note that this definition of $F(S)$ is what one would get if $(r_0, r_1, r_2, r_3, r_0)$ were a ring. $F(\{r_1, r_2\}) \geq 0$ and $F(\{r_2, r_3\}) \geq 0$ because $(r_0, r_1, r_2, r_3, r_0)$ is an optimal tour. In fact $F(S) \geq 0$ for all $S \subseteq V \cup \{r_0\}$. It is tedious but straight forward to verify that these family costs correctly model $T(S)$ when a loop consists of only three retailers.

For each ring the Loop heuristic creates $O(N^2)$ non-zero families. Thus the Loop heuristic creates at most $O(N^2)$ non-zero families. See Figure 3 for an example of the family structure for the Loop heuristic. The labels of the families contain the names of the retailers in the family.

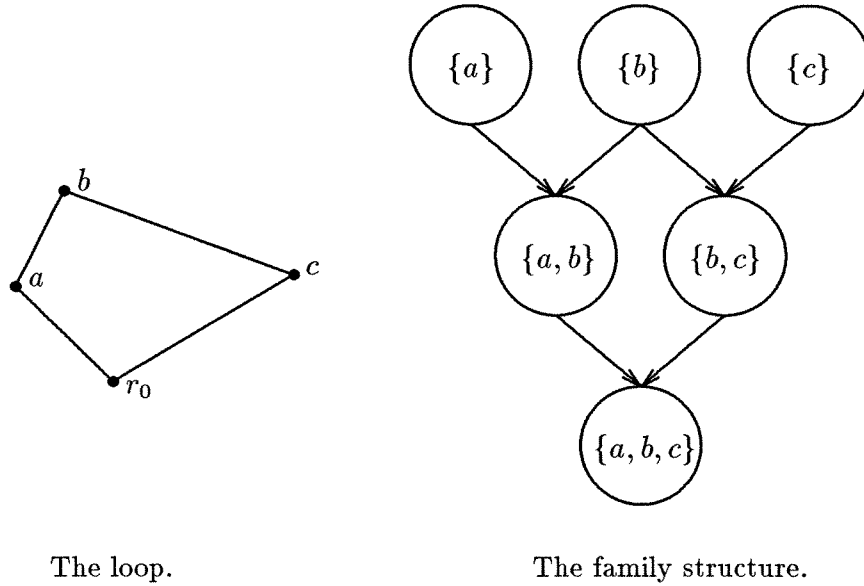


Figure 3: The family structure for Loop.

A.2 Box Family Structure

The Box heuristic builds tree structured road systems. All tree structured road systems have a simple family structure. Let (i, j) be an arc in the tree, i. e., one of the roads. We define the set $R_{(i,j)}$ to be all the retailers whose unique path to the central warehouse includes the arc (i, j) . We let $D(i, j)$ be the length of arc (i, j) . We define $F(S)$ as follows,

$$\begin{aligned}
 F(R_{(i,j)}) &= 2D(i, j) && \text{for all arcs } (i, j) \\
 F(S) &= 0 && \text{otherwise.}
 \end{aligned}$$

The reason this family structure yields the desired costs is that arc (i, j) must be transversed in both directions if any retailer in $R_{(i,j)}$ places an order. In addition if

none of the retailers in $R_{(i,j)}$ places an order, then arc (i,j) need not be transversed.

Since the Box heuristic can be easily modified to create a new depot only when two retailers are split between branches, and any tree has one less arc than nodes, we know that there are at most $O(N)$ non-zero families. In addition it is not hard to show that these families have a tree structure. See Figure 4 for an example of a tree structured road system and its associated family structure. The labels of the families contain the names of the retailers in the family.

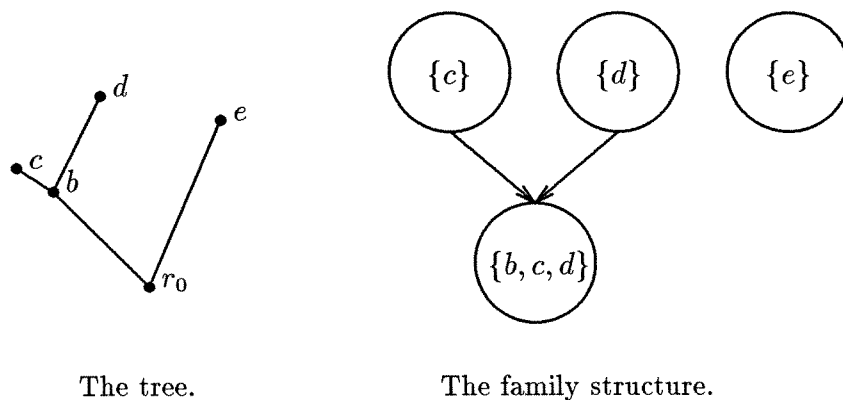


Figure 4: The family structure for Box and MST.

A.3 Circle Family Structure

The Circle heuristic has the simplest family structure of the four heuristics. We order the retailers so that $D(r_0, r_{i-1}) < D(r_0, r_i)$ for all $1 \leq i \leq N$. We define $F(S)$ as follows,

$$F\left(\bigcup_{j=i}^N \{r_j\}\right) = \left\lceil \sqrt{\frac{N\pi}{2}} + \frac{3\pi}{2} \right\rceil 2(D(r_0, r_i) - D(r_0, r_{i-1}))$$

$$F(S) = \begin{cases} 0 & \text{for all } 1 \leq i \leq N \\ & \text{otherwise.} \end{cases}$$

The reason this family structure yields the desired costs is that the families that have a member retailer ordering also contain $r_{max}(S)$ as a member. Recall that $r_{max}(S)$ is the retailer in S that is furthest from the central warehouse.

The Circle heuristic clearly has at most N non-zero families. See Figure 5 for an example of the family structure for the Circle heuristic. The labels of the families contain the names of the retailers in the family.




Figure 5: The family structure for Circle.

A.4 MST Family Structure

The MST heuristic builds tree structured road systems. Section A.2 shows how to create the family structure for any tree structured road system. We note here

that since the MST heuristic yields a tree that has one node per retailer, the MST heuristic yields a family structure with $O(N)$ non-zero families.

B Bibliography

Anily, S., and A. Federgruen. 1988. One Warehouse Multiple Retailer Systems with Vehicle Routing Costs. *Working Paper, Faculty of Commerce and Business Administration, University of British Columbia.*

Beardwood, J., J. H. Halton and J. M. Hammersley. 1959. The Shortest Path Through Many Points. *Proc. Cambridge Philos. Soc.* **55**, 299–327.

Beltrami, L., and L. Bodin. 1974. Networks and Vehicle Routing for Municipal Waste Collection. *Networks* **4**, 65–94.

Brown, R. G. 1978. Inventory Control. In *Handbook of Operations Research; Vol. 2*. Moder, J. J. and S. E. Elmaghraby (eds.). Van Nostrand Reinhold, New York.

Burns, L. D., R. W. Hall, D. E. Blumenfeld and C. F. Daganzo. 1985. Distribution Strategies that Minimize Transportation and Inventory Costs. *Operations Research* **33**, 469–490.

Dror, M., and M. Ball. 1987. Inventory/Routing: Reduction from an Annual to a Short-Period Problem. *Naval Research Logistics* **34**, 891–905.

Dror, M., M. Ball and B. Golden. 1985. A Computational Comparison of Algorithms for the Inventory Routing Problem. *Annals of Operations Research* **4**, 3–23.

Dror, M., and L. Levy. 1986. A Vehicle Routing Improvement Algorithm Comparison of a “Greedy” and a Matching Implementation for Inventory Routing. *Comput. and Operations Research* **13**, 33–45.

Edelsbrunner, H. 1987. *Algorithms in Combinatorial Geometry*. Brauer, W., G. Rozenberg and A. Salomaa (eds.). Springer-Verlag.

Federgruen, A., M. Queyranne and Y. S. Zheng. 1989. Simple Power of Two Policies are Close to Optimal in a General Class of Production/Distribution Networks with General Joint Setup Costs. *Working Paper, Wharton School, University of Pennsylvania.*

Federgruen, A., and Y. S. Zheng. 1988. The Joint Replenishment Problem with General Joint Cost Structures: General Solution Methods and Performance Bounds. *Working Paper, Wharton School, University of Pennsylvania.*

Federgruen, A., and P. Zipkin. 1984. A Combined Vehicle Routing and Inventory Allocation Problem. *Operations Research* **32**, 1019–1037.

Gallego, G., and D. Simchi-Levi. 1988. On the Effectiveness of Direct Shipping Strategy for the One Warehouse Multi-Retailer R-Systems. *Working Paper, Department of Industrial Engineering and Operations Research, Columbia University*.

Goyal, S. K. 1987. Comment on “A Dynamic Programming Approach for Joint Replenishment Under General Order Cost Functions”. *Management Science* **33**, 131–135.

Harris, F. 1915. *Operations and Cost*, 48–53. (Factory Management Series). A. W. Shaw Co., Chicago.

Held, M., and R. M. Karp. 1962. A Dynamic Programming approach to sequencing problems. *SIAM J. Appl. Math.* **10**, 196–210.

Herer, Y. 1990. Submodularity and The Travelling Salesman Problem. *Tech. Report No. 915, School of Operations Research and Industrial Engineering, Cornell University*.

Lawler, E. L., J. K. Lenstra, A. H. G. Rinnooy Kan and D. B. Shmoys (eds.). 1985. *The Travelling Salesman Problem, A Guided Tour of Combinatorial Optimization*. John Wiley & Sons Ltd., Chichester.

Maxwell, W. L., and J. A. Muckstadt. 1985. Establishing Consistent and Realistic Reorder Intervals in Production-Distribution Systems. *Operations Research* **33**, 1316–1340.

Nemhauser, G. L., L. A. Wolsey and M. L. Fisher. 1978. An Analysis of Approximations for Maximizing Submodular Set Functions-I. *Mathematical Programming* **14**, 265–294.

Queyranne, M. 1985. A Polynomial-Time, Submodular Extension to Roundy’s 98%-Effective Heuristic for Production/Inventory Systems. *Working Paper No. 1136, Faculty of Commerce and Business Administration, University of British Columbia*.

Queyranne, M. 1987. Comment on “A Dynamic Programming Approach for Joint Replenishment Under General Order Cost Functions”. *Management Science* **33**, 131–133.

Roundy, R. 1985. 98%-Effective Integer-Ratio Lot-Sizing for One-Warehouse Multi-Retailer Systems. *Management Science* **31**, 1416–1430.

Roundy, R. 1986. A 98%-Effective Lot-Sizing Rule for a Multi-Product, Multi-Stage Production/Inventory System. *Mathematics of Operations Research* **11**, 699–727.

Zheng, Y. S. 1987. Replenishment Strategies for Production/Distribution Networks with General Joint Setup Costs. *Ph.D. thesis, Graduate School of Arts and Sciences, Columbia University.*