Inventory and Facility Location Models with Market Selection

Retsef Levi [⋆1], Joseph Geunes [⋆⋆2], H. Edwin Romeijn [⋆⋆⋆3], and David B. Shmoys [†4]

[1] School of ORIE, Cornell University, Ithaca, NY 14853. rl227@cornell.edu
[2] Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL 32611-6595. geunes@ise.ufl.edu
[3] Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL 32611-6595. romeijn@ise.ufl.edu
[4] School of ORIE and Department of Computer Science, Cornell University Ithaca, NY 14853. shmoys@cs.cornell.edu

Abstract. We consider important generalizations of a wide class of traditional deterministic inventory and facility location models that we call *inventory/facility location models with market selection*. Instead of the traditional setting, we are given a set of *markets*, each specified by a sequence of demands and associated with a revenue. Decisions are made in two stages. We first make a decision of what markets to select, where all other markets are rejected. Next we have to construct a minimum-cost production plan (facility layout) to satisfy all of the demands of all the selected markets. The goal is to minimize the overall lost revenues of rejected markets and the production (facility openings and connection) costs. We show how to leverage existing approximation results for the traditional models to corresponding results for the counterpart models with market selection. More specifically, any LP based $\alpha-$approximation for the traditional model can be leveraged to a $\frac{1}{1-e^{-\frac{1}{\alpha}}}-$approximation algorithm for the counterpart model with market selection. Our techniques are also applicable to an important class of covering problems.

1 Introduction

Traditional deterministic inventory theory provides streamlined optimization models where the goal is to satisfy a sequence of specified demands over a given planning horizon with minimum overall cost, optimally balancing different costs in the supply chain. Recent trends in supply chain research and practice have led to a broadened perspective that considers not only decisions on the supply side but on the demand side as well. More specifically, the demands to which the supply chain must respond and commit to are not completely exogenous parameters, but may be influenced by endogenous decisions such as pricing, promotions, and other strategic marketing-based

⋆ Research supported partially by a grant from Motorola and NSF grants CCR-9912422& CCR-0430682.
⋆⋆ Research supported partially by NSF grants DMI-0322715 and DMI-0355533.
⋆⋆⋆ Research supported partially by NSF grant DMI-0355533.
† Research supported partially by NSF grants CCR-9912422& CCR-0430682.

factors. In particular, a supplier should strive to optimally fit the *shape* of the demand to the supply chain's production capabilities. A fundamental aspect of these issues is the choice of markets (or customers) to which the supplier commits to sell commodities. That is, given products and a set of potential markets, which markets should the supplier target for sales? Clearly, this decision is impacted by the potential profitability of each market. While marketing models can be used to predict the potential *revenue* available in a market, economies of scale in production make it difficult to isolate the *cost* of serving a market, in particular, because this cost depends on the collective set of markets being served. Aiming to address such *market selection* decisions in the planning phase, we propose the following integrated market selection and production planning problem.

We consider generalizations to a wide class of traditional deterministic inventory models that capture the above mentioned issues. In these more general models, the input consists of a set of *markets* (or customers), where each is specified by a sequence of demands over the given planning horizon and is associated with a potential revenue. The decisions are made in two stages. First we decide which markets we are going to *select* and which ones will be *rejected*. By selecting a market we commit to satisfy *all* of its demands over the horizon. Rejecting a market implies that we can ignore all of its demands, but we lose its potential revenue. Once a subset of markets is selected, we need to construct a minimum cost production plan to satisfy all of the demands of the selected markets on time. As in the traditional inventory models the production cost usually consists of a *fixed ordering cost* that is incurred in each period an order is placed, and of *holding cost* for carrying excess inventory from period to period. Our goal is to minimize the overall cost, namely the *lost revenue* incurred by market rejection and the *production cost* of satisfying the demands of the selected markets. We call this new class of extended models *inventory models with market selection*.

All of the inventory models we consider in this paper can be formulated as facility location type integer programs that admit strong LP relaxations. We show how to modify these LP's to capture the more general models with market selection. We then describe general LP-rounding and primal-dual frameworks that leverage any LP-based approximation algorithm for the traditional inventory models to a corresponding algorithm for its counterpart with market selection. This results in a stream of approximation algorithms for these more general market selection models that have constant worst-case performance guarantees. That is, for any instance of the specific market selection problem, the algorithm provides a solution with cost at most constant multiple of the optimal cost.

The models. Following are the details of some of the models to which our LP-rounding and primal-dual frameworks are applicable. In the *joint replenishment problem with market selection*, we consider a set of m markets, denoted by $\mathcal{M}$. Each market $j \in \mathcal{M}$ is specified by a sequence of demands for N commodities over a finite planning horizon of T periods and is associated with a potential revenue, r_j. The demands required by market j for item i are denoted by $d^j_{i1}, \dots, d^j_{iT}$ (for $i = 1, \dots, N$). The production cost structure is identical to the traditional joint replenishment problem (JRP). Each order incurs a *fixed joint ordering cost*, K_0, regardless of the combination of items being ordered. In addition, for each item included in the order, we incur a *fixed item ordering cost*, K_i, regardless of the number of units being ordered. The fixed ordering cost is balanced by the *holding cost* for carrying excess inventory over periods. The holding cost structure follows [13], i.e., for each $i = 1, \ldots, N$, $t = 1, \ldots, T$ and $s \leq t$, we have a parameter h^i_{st} that denotes the per unit cost of providing item i in period t from period $s \leq t$. The parameters h_{ist} are assumed to be non-negative and non-increasing in s. Per unit ordering cost can be incorporated into the h parameters. This generalizes the traditional holding cost structure, where we have a parameter $h_{it} \geq 0$ for each i and t, which denotes the per unit cost to hold inventory from period t to $t + 1$.

The *single item lot-sizing problem with market selection* is a special case of the JRP problem above, where we have only one commodity, i.e., $N = 1$. As a result there is no point to consider separate joint and item ordering cost, and instead we have just time dependent fixed ordering cost K_s (for each $s = 1, \ldots, T$).

The *one-warehouse multi-retailer (OWMR) problem with market selection* is a generalization of the JRP problem, where we assume that each item corresponds to a different retailer, and the retailers are supplied by a central warehouse that can hold inventory. We follow the general holding cost structure described in [18, 19].

In the *assembly problem with market selection*, each market is again specified by a sequence of T demands for a single commodity (end-product). However, the end-product is assembled from N components (see [13, 14] for details).

We will also consider some of these models with the additional constraints that orders are placed in *batches*, each with a given capacity. For each additional batch ordered, we incur an additional fixed ordering cost. These constraints are usually called *soft capacity constraints*.

In addition, our techniques are applicable to several market selection variants of facility location problems. This includes the classical uncapacitated metric facility location problem and its variants with soft capacities, and with service installation costs (see [22] for details).

The goal in all of these models is to first select a subset of the markets, and then construct a minimum cost production plan (or facility layout) for all of the demands of the selected markets, such that the overall cost incurred by market rejection and the production (facility layout and connection) plan is minimized.

Related literature. The traditional inventory models mentioned above, have attracted the attention of many researchers throughout the years. The single item lot-sizing problem can be solved to optimality using dynamic programming [23], including the case with soft capacity constraints [17]. The JRP and the OWMR problems are known to be NP-hard [1]. The assembly problem with general holding cost structure as in [13, 14] is also known to be NP-hard, and it is a long standing open problem whether or not it is NP-hard in the case of traditional holding cost. In several recent papers [13, 14, 18, 19], Levi, Roundy and Shmoys have developed LP-based approximation algorithms for these inventory models. In particular, they have provided a primal-dual 2-approximation algorithm for the JRP and the assembly problem, and an LP-rounding 2.398-approximation algorithm for the OWMR problem.

The facility location problem is known to be NP-hard, and there is a huge body of literature on approximation algorithms for this problem. Mahdian, Ye and Zhang have the currently best known result, an LP-based 1.52-approximation algorithm [15] (we refer the reader to [20] for a survey on approximation results on this problem).

There are several other models, such as the *prize-collecting travelling salesman problem* [2, 6], *prize-collecting Steiner tree* [10, 9] and *multi-processor scheduling with rejection* [4], where classical combinatorial optimization problems were extended to consider *rejection/selection* decisions. However, in all of these models the rejection /selection decisions are made independently for each element (node or a job). Our market selection models are different in that the rejection/selection is made with respect to a collection of elements as a set. We note that in their full generality our techniques and results are applicable to the market selection variants of the prize-collecting problems mentioned above, as well as to market selection variants of additional covering problems.

In [8], Geunes, Romeijn and Taaffe have considered a single item, single location *order selection model* where each market is specified by a single demand in a single period; they have provided an optimization algorithm based on dynamic programming. In this paper, we consider multi-item, multi-stage and capacitated inventory models with more complex selection decisions that are made with respect to markets specified over the time horizon rather than in a single period.

Our approach, when applied to the metric facility location problem, generalizes the model of *facility location with penalties* discussed by Charikar, Khuller, Mount and Narasimhan in [7]. They have also considered a model where the rejection/selection decision is made per demand point independently rather than with respect to subsets of demands. They have provided a 3-approximation algorithm for this model, extending the novel primal-dual algorithm of Jain and Vazirani [11] for the facility location problem. For the more general model considered in this paper, we provide a significantly better approximation factor of 2.075 using LP-rounding. We also provide a 2.542-approximation algorithm for the market selection variant of the facility location problem with soft capacities. However, our LP-rounding framework is not applicable to some of the facility location models they have considered.

Our techniques and results. Our LP-rounding framework builds on known facility location type LP relaxations for the traditional models. We modify the LP's to capture the corresponding market selection problems by adding a variable for each market that indicates whether it is rejected or not. If we think about this variable as representing a special 'order' (or facility), then we either serve all of the demands of the corresponding market by this special 'order' and lose the revenue of the market, or we have to assign each one of its demands to a regular order (facility). We solve the *market selection LP* relaxation and consider its optimal solution. The rounding algorithm is conceptually very simple. For some threshold $\frac{1}{\beta}$ (for some $\beta \geq 1$), we reject all the markets that the LP fractional solution rejects by more than $\frac{1}{\beta}$. Clearly, this gives a bound on the rejection cost that our solution incurs. The main observation is that once we decide on the selected markets, then the market selection LP is reduced to the *traditional LP* for the inventory (facility location) model, for an instance defined by the selected markets aggregated. We now construct the production plan (facility openings) for these markets using existing LP-based α-approximation algorithm for the traditional inventory (facility location) model. It is clear that the cost incurred is at most α times the optimal value of the (reduced) traditional LP. Finally, observe that if a certain market is fractionally rejected to an amount less than $\frac{1}{\beta}$, then the fractional production plan satisfies at least

$1 - \frac{1}{\beta}$ of each of its demands. Therefore, we can scale the fractional solution that corresponds to the selected markets by no more than $1-\frac{1}{\beta}$ to get a feasible fractional solution to the reduced traditional LP. This bounds our production cost to be at most $\alpha/(1 - \frac{1}{\beta})$ times the optimal fractional production cost of the market selection LP. Optimizing deterministically over β we can now leverage any existing LP-based α-approximation algorithm for the traditional inventory problem to an $(\alpha + 1)$-approximation algorithm for the corresponding market selection counterpart. However, drawing $\frac{1}{\beta}$ at random from the appropriate distribution, as described by Goemans in [9], provides a randomized algorithm with improved expected guarantees. Using derandomization we can now leverage any α-approximation into a deterministic $\frac{1}{1-e^{-\frac{1}{\alpha}}}$ guarantee. As a by product we also show that if the traditional model admits a strong LP relaxation so does its market selection counterpart. We note that somewhat similar ideas were used independently in [21] in the context of 2-stage stochastic models.

In addition, we show how to extend the primal-dual framework described in [13, 14] to provide combinatorial approximation algorithms for the market selection variants of the single item lot-sizing problem, the JRP problem and the assembly problem. We provide worst-case guarantees of 2, 3 and 3, respectively.

The rest of the paper is organized as follows. In Section 2, we start with the illustrative simple example of the single item lot-sizing problem. Then in Section 3, we describe the general LP-rounding framework and the approximation results it implies. Finally, in Section 4, we briefly discuss how to extend the primal-dual framework of [13, 14].

2 The Single Item Lot-Sizing Problem - An Illustrative Example

In this section, we will use the illustrative simple example of the single item lot-sizing problem with market selection to demonstrate the underlying techniques of our LP-rounding framework.

Next we describe a facility location type LP-relaxation of this selection problem that is based on the LP relaxation of its traditional counterpart. We then show how to round the optimal solution of this LP into a feasible solution to the market selection variant with cost at most 1.582 times the optimal cost.

The LP. The LP is as follows:

$$\begin{aligned}
&\text{minimize} && \sum_{j=1}^{m} z_j r_j \;+\; \sum_{s=1}^{T} y_s K_s \;+ \sum_{j=1}^{m}\sum_{t=1}^{T}\sum_{s=1}^{t} x_{st}^{j} H_{st}^{j} && && \text{(P)}\\
&\text{subject to} && \sum_{s=1}^{t} x_{st}^{j} \;+\; z_j = 1, && j = 1,\dots,m,\; t = 1,\dots,T, && (1)\\
& && x_{st}^{j} \le y_s, && j = 1,\dots,m,\; t = 1,\dots,T, && (2)\\
& && && s = 1,\dots,t, &&\\
& && x_{st}^{j},\; y_s \ge 0, && j = 1,\dots,m,\; s = 1,\dots,T, && (3)\\
& && && t = s,\dots,T. &&
\end{aligned}$$

The *selection variable* z_j is equal to 1 if we decided to reject market j and 0 otherwise. The *assignment variable* x_{st}^{j} indicates whether the demand of market j in period t, d_t^j, was provided from period s. Correspondingly, we let H_{st}^{j} be the overall cost of providing the demand d_t^j from period s, i.e., $H_{st}^{j} = h_{st} d_t^j$. The *order variable* y_s indicates whether an order was placed in period s. The *assignment constraint* (1) ensures that each demand in period t of a selected market is satisfied from some time period $s \le t$. Note that if we have selected market j and set $z_j = 0$, then for each period t, constraint (1) implies that $\sum_{s\le t} x_{st}^{j} = 1$. Hence, all of the demands of market j must be satisfied on time. Conversely, if $z_j = 1$, then we do not have to satisfy any of the demands of market j, and in turn we incur a cost of r_j. The *order constraint* (2) ensures that no demand can be provided from period s without placing an order in period s. It is straightforward to verify that the induced integer programm provides a correct formulation for the single item lot-sizing problem with market selection, and therefore, the optimal solution of the LP-relaxation provides a lower bound on the cost of any feasible solution to the problem.

Observe that once we decide on the selected markets (i.e., assign zero/one value to each variable z_j), then the above LP is reduced to an LP of a traditional single item lot-sizing problem defined by the selected markets. The demand in period t is equal to the aggregated demands d_t^j over all markets $j \in \mathcal{M}$ that were selected (i.e., the demand in period t is equal to $\sum_{j\in\mathcal{M}} d_t^j$). Moreover, it is well known that this LP always provides an integer solution i.e., it has an integrality property (for different proofs see [12, 3, 5, 13]). However, we note that the market selection LP, (P) above, does not have an integrality property (i.e., for some instances the optimal solution is fractional).

A Rounding Algorithm Let $(\hat{z}, \hat{x}, \hat{y})$ be the optimal solution of (P) above. We next describe an extremely simple procedure to round this optimal fractional solution into a feasible solution for the single item lot-sizing problem with market selection, denoted by $(\bar{z}, \bar{x}, \bar{y})$, and then show that its cost is always at most twice the optimal cost.

Let $\mathcal{Z}_{\frac{1}{2}} := \{j \in \mathcal{M} : \hat{z}_j \ge frac12\}$, be the set of all markets that are rejected by 'amount' of at least $\frac{1}{2}$ in the fractional optimal solution. Let $\mathcal{Z}_S := \mathcal{M} \setminus \mathcal{Z}_{\frac{1}{2}}$. We now reject all the markets $j \in \mathcal{Z}_{\frac{1}{2}}$ and select all other markets, i.e., we set $\bar{z}_j = 1$ if and only if $\hat{z}_j \ge frac12$. Next consider the single item lot-sizing problem defined by the selected markets, i.e., by the markets $j \in \mathcal{Z}_S$, where again the demand in period t is equal to $\sum_{j\in\mathcal{Z}_S} d_t^j$. This problem can be efficiently solved to optimality using dynamic programming (see [23] for details). We take this solution as our production plan. This concludes the description of the algorithm.

Analysis. We start with a lemma that bounds the rejection cost incurred in $(\bar{z}, \bar{x}, \bar{y})$. The proof is rather straightforward and follows from the construction of the algorithm.

Lemma 1. *The rejection cost of the solution $(\bar{z}, \bar{x}, \bar{y})$ is at most twice the rejection cost of the optimal fractional solution $(\hat{z}, \hat{x}, \hat{y})$, i.e., $\sum_{j=1}^m \bar{z}_j \le 2\sum_{j=1}^m \hat{z}_j$.*

To complete the analysis, it is enough to show that the production cost of the solution $(\bar{z}, \bar{x}, \bar{y})$ is at most twice the production cost of the solution $(\hat{z}, \hat{x}, \hat{y})$.

Lemma 2. *The production cost of the solution $(\bar{z}, \bar{x}, \bar{y})$ is at most twice the production cost of $(\hat{z}, \hat{x}, \hat{y})$, i.e., $\sum_{s=1}^T \bar{y}_s K_s + \sum_{j=1}^m \sum_{t=1}^T \sum_{s=1}^t \bar{x}_{st}^j H_{st} \le 2(\sum_{s=1}^T \hat{y}_s K_s + \sum_{j=1}^m \sum_{t=1}^T \sum_{s=1}^t \hat{x}_{st}^j H_{st})$.*

Proof. Recall that once we decide on the selected markets, then the problem and the LP are reduced to a traditional lot-sizing problem defined by the markets in $\mathcal{Z}_S$. Moreover, since the facility location LP of the lot-sizing problem has an integrality property, we know that the optimal value of that LP is equal to the value of the optimal solution achieved by applying dynamic programming to this problem. It is then enough to observe a feasible solution to this reduced LP that has cost at most twice the production cost in $(\hat{z}, \hat{x}, \hat{y})$. We next describe how to construct such a solution, denoted by (x, y).

For each $j \in \mathcal{Z}_{\frac{1}{2}}$ we set $x_{st}^j = 0$ for each $s \le t$. Now for each $j \in \mathcal{Z}_S$ we set $x_{st}^j := \frac{\hat{x}_{st}^j}{1-\hat{z}_j}$. Finally, for each $s = 1, \ldots, T$ we set $y_s := \max\{\min_{j\in\mathcal{Z}_S}(\frac{\hat{y}_s}{1-\hat{z}_j}), 1\}$. It is readily verified that (x, y) constructed above is a feasible fractional solution for the LP of the lot-sizing problem defined by the markets in $\mathcal{Z}_S$.

Observe that for every $j \in \mathcal{Z}_S$ we have $1 - \hat{z}_j > frac12$. Hence, for each $j \in \mathcal{Z}_S$ and $s \le t$ we have $x_{st}^j \le 2\hat{x}_{st}^j$ and for each $s = 1, \ldots, T$ we have $y_s \le 2\hat{y}_s$. Now by construction of the algorithm $(\bar{x}, \bar{y})$ provides an optimal solution to the LP of the lot-sizing problem defined by the markets in $\mathcal{Z}_S$. We get that,

$$\sum_{s=1}^T \bar{y}_s K_s + \sum_{j=1}^m \sum_{t=1}^T \sum_{s=1}^t \bar{x}_{st}^j H_{st} \le \sum_{s=1}^T y_s K_s + \sum_{j=1}^m \sum_{t=1}^T \sum_{s=1}^t x_{st}^j H_{st} \le$$
$$2(\sum_{s=1}^T \hat{y}_s K_s + \sum_{j=1}^m \sum_{t=1}^T \sum_{s=1}^t \hat{x}_{st}^j H_{st}).$$

This concludes the proof of the lemma.

As a corollary of Lemmas 1 and 2 we get the following theorem.

Theorem 1. *The algorithm provides a 2-approximation algorithm for the lot-sizing problem with market selection, i.e., the cost $(\bar{z}, \bar{x}, \bar{y})$ is at most twice the optimal cost.*

2.1 Randomized Algorithm

Next we briefly discuss how randomized rounding yields an improved approximation algorithm. Instead of using the threshold $\frac{1}{2}$, we follow the ideas in [9], and choose a threshold $\frac{1}{\beta}$ uniformly at random from $(0, \delta]$, where $0 < \delta \leq 1$ (the value of δ will be specified later). Once $\frac{1}{\beta}$ is chosen we proceed in the same way described above, i.e., we reject all markets $j \in \mathcal{M}$ with $\hat{z}_j \geq \frac{1}{\beta}$, and construct a minimum-cost production plan for the selected markets (using dynamic programming).

Following are several lemmas that establish the expected performance guarantee of the randomized algorithm described above.

Lemma 3. *For each market $j \in \mathcal{M}$, the probability that market j is rejected is at most $\frac{\hat{z}_j}{\delta}$, i.e., the expected rejection cost of the solution is at most $\frac{\sum_j \hat{z}_j r_j}{\delta}$.*

Proof. For each market j with $\hat{z}_j \leq \delta$ the claim follows trivially. For each market j with $\hat{z}_j > \delta$, we reject the market with probability 1, however, $\frac{\hat{z}_j}{\delta} > 1$. This concludes the proof.

Lemma 4. *The expected production cost of the solution is at most $\frac{\ln(\frac{1}{1-\delta})}{\delta}(\sum_{s=1}^{T} \hat{y}_s K_s + \sum_{j=1}^{m} \sum_{t=1}^{T} \sum_{s=1}^{t} \hat{x}_{st}^{j} H_{st})$.*

Proof. By similar arguments to Lemma 2 above, it is readily verified that, for each value of $\frac{1}{\beta}$, the production cost is at most $\frac{1}{1-\frac{1}{\beta}}$ times the production cost of the optimal LP solution, $(\hat{z}, \hat{z}, \hat{y})$. Taking the integral $\frac{1}{\delta}\int_0^{\delta} \frac{1}{1-\frac{1}{\beta}} d(\frac{1}{\beta})$ gives the required result.

Setting δ equal to $1 - e^{-1}$ provides and a solution with expected cost at most 1.582 times the optimal cost. Observe that the solution constructed by the algorithm is uniquely determined by the set of rejected markets. Hence, there are at most m relevant values of $\frac{1}{\beta}$ that will yield distinct solutions. We can then derandomize the algorithm by enumerating only the m relevant values of $\frac{1}{\beta}$ and get the following theorem.

Theorem 2. *There exists a 1.582-approximation algorithm for the market selection variant of the lot-sizing problem.*

3 General LP-Rounding Framework

In this section, we will generalize the example discussed in Section 2 above, and propose a general LP-rounding framework that can be applied to market selection variants of a wide class of classical deterministic inventory and facility location models. The ingredients for our LP-rounding framework are a facility location type LP relaxation and an LP-based α-approximation algorithm for the traditional model without market selection. We show how, given these ingredients, one can leverage the α-approximation algorithm into an $\frac{1}{1-e^{-\frac{1}{\alpha}}}$-approximation algorithm for the corresponding counterpart model with market selection.

Our algorithms can use any facility location type LP with *assignment variables*, x, and *order variables*, y, and with *assignment constraints* (e.g., constraint (1) above), *ordering constraints* (e.g., constraint (2) above) and possibly *soft capacity constraints*. Here the order variables y indicate how many batches were ordered in each order, where each batch incurs an additional fixed ordering cost. The number of batches installed dictates the capacity of the order, or in other words, how many units of demand can be satisfied by this order. For the lot-sizing example discussed in Section 2, but with soft capacity constraints, we would have the constraint, $\sum_{t \geq s} \sum_{j=1}^{m} x_{st}^{j} d_t^j \leq y_s U_s$, for each period $s = 1, \ldots, T$, where U_s is the batch capacity associated with the order in period s. By an LP-based α-approximation algorithm we mean an algorithm that constructs a solution with cost guaranteed to be at most α times the optimal LP value.

We now state and prove a general theorem that relates the approximation results for traditional deterministic inventory (and facility location models) to approximation results for their corresponding counterparts with market selection.

Theorem 3. *Consider a deterministic inventory or a facility location model that can be formulated as a facility location type integer program with assignment and order variables, and with assignment constraints, order constraints and possibly soft capacity constraints. Also assume that the model has an LP-based α-approximation. Then there exists an $(\alpha+1)$-approximation algorithm for the counterpart of this model with market selection.*

Proof. We first adjust the LP relaxation for the traditional model to capture the counterpart model with market selection. Let $\mathcal{M}$ be again the set of m markets. As before we let the selection variable z_j be equal 1 if we reject market j and 0 otherwise. For each demand of market j, we have a separate set of assignment variables (for each facilty/order that can serve it). Similarly there is an assignment constraint for each of its demands that includes the z_j variable. The order constraints and possibly the soft capacity constraints are adapted accordingly. Finally, we modify the objective function and add the market rejection part $\sum_{j=1}^{m} z_j r_j$. We now solve this LP and let $(\hat{z}, \hat{x}, \hat{y})$ be the respective optimal solution.

Now for some parameter $\beta > 1$ that will be determined later, let $\mathcal{Z}_{\frac{1}{\beta}} := \{j : \hat{z}_j \geq \frac{1}{\beta}\}$, i.e., the set of markets that are rejected in the market selection LP optimal solution by 'amount' at least $\frac{1}{\beta}$. Let $\mathcal{Z}_S$ be the set of all other markets (i.e., $\mathcal{Z}_S = \mathcal{M} \setminus \mathcal{Z}_{\frac{1}{\beta}}$). We now reject all the markets in $\mathcal{Z}_{\frac{1}{\beta}}$ and select all markets $j \in \mathcal{Z}_S$. Next consider the instance for the traditional variant of the model defined by the markets in $\mathcal{Z}_S$, where again, for each period (location) and each commodity, we aggregate the corresponding demands of all the selected markets. We then apply the α-approximation algorithm to the instance defined by $\mathcal{Z}_S$, and get a production plan (facility location solution).

It is straightforward to see that the rejection cost incurred by the solution is at most $\beta \sum_{j=1}^{m} \hat{z}_j r_j$. Also observe that by a similar scaling to the one described in Lemma 2, we can observe a feasible solution to the traditional LP defined by the markets in $\mathcal{Z}_S$ of cost at most $\frac{\beta}{\beta-1}$ times the cost of the (fractional) production plan in $(\hat{z}, \hat{x}, \hat{y})$. However, by our assumption, the cost of the production plan constructed by the algorithm, is at most α times the cost of the optimal solution to the above traditional LP defined by the markets in $\mathcal{Z}_S$. Therefore, it is at most $\alpha(\frac{\beta}{\beta-1})$ times the cost of the production plan in

$(\hat{z}, \hat{x}, \hat{y})$. Setting $\beta = \alpha + 1$, we get that the cost of the solution is at most $\alpha + 1$ times the cost of $(\hat{z}, \hat{x}, \hat{y})$. This concludes the proof.

Similar to the lot-sizing case, if we choose $\frac{1}{\beta}$ at random uniformly from $(0, \delta]$ (where again $0 < \delta \leq 1$), then the expected rejection cost is at most $\frac{1}{\delta}$ the rejection cost in the optimal LP solution, $(\hat{z}, \hat{x}, \hat{y})$ (i.e., at most $\frac{\sum_j \hat{z}_j r_j}{\delta}$), and the expected production cost is at most $\frac{\alpha \ln(\frac{1}{1-\delta})}{\delta}$ times the production cost in $(\hat{z}, \hat{x}, \hat{y})$ (i.e., at most $\frac{\alpha \ln(\frac{1}{1-\delta})}{\delta}(\sum_{s=1}^{T} \hat{y}_s K_s + \sum_{j=1}^{m} \sum_{t=1}^{T} \sum_{s=1}^{t} \hat{x}_{st}^{j} H_{st})$). Setting δ equal to $1 - e^{-\frac{1}{\alpha}}$ guarantees that expected cost of the algorithm is at most $\frac{1}{1-e^{-\frac{1}{\alpha}}}$ times the optimal cost. By the same derandomization techniques, we achieve the same guarantee deterministically.

Approximation results. The *single item lot-sizing problem with soft capacity constraints* can be solved optimally [17]. In [18], Levi, Roundy and Shmoys have showed that the integrality gap of the facility location type LP of this model is at most 2. We get the following theorem.

Theorem 4. *There exists a 2.542-approximation algorithm for the market selection single item lot-sizing problem with soft capacity constraints.*

We consider a sequence of LP-based approximation algorithms described in [13, 14, 18, 19], namely, a 2-approximation algorithms for the JRP and assembly problems, a 4-approximation algorithm for the JRP with soft capacities, a 2.398-approximation algorithm for the OWMR and a 4.769-approximation algorithm for the OWMR with soft capacities. We then conclude the following theorems.

Theorem 5. *There exist a 2.54-approximation algorithms for the market selection variants of the joint replenishment and the assembly problems, and a 4.521-approximation algorithm for market selection variant of the JRP with soft capacities.*

Theorem 6. *There exist a 2.993-approximation algorithm and a 5.287-approximation algorithm for the market selection variants of the OWMR and the OWMR with soft capacities, respectively.*

Finally, in [15, 16], Mahdian, Ye and Zhang have provided a 1.52-approximation algorithm and 2-approximation algorithm for the facility location problem and its variant with soft capacities, respectively. Both algorithms are LP-based. In [22], Shmoys, Swamy and Levi have provided a 6-approximation algorithm for the facility location problem with service installation costs. We again conclude a corresponding theorem.

Theorem 7. *There exist a 2.075, 2.542 and 6.514 approximation algorithms for the market selection variants of the uncapacitated facility location problem, the facility location problem with soft capacities and facility location with service installation costs.*

4 Extended Primal-Dual Framework

In [13, 14], Levi, Roundy and Shmoys have described a general primal-dual framework that solves the lot-sizing problem to optimality, and provides a 2-approximation for the JRP and assembly problem. In this section, we will briefly discuss how to modify their algorithms to provide combinatorial approximation algorithms (i.e., algorithms that do not require solving the LP) for the counterpart models with market selection. Using these modified algorithms, we provide performance guarantees of 2, 3 and 3 for these models, respectively. For simplicity and due lack of space, we next focus on the single item lot-sizing problem. Similar extensions exist for the JRP and assembly problem.

The following is the dual of (P) from Section 2.

$$\text{maximize} \quad \sum_{j=1}^{m}\sum_{t=1}^{T} b_t^j \tag{D}$$

$$\text{subject to} \quad b_t^j \le H_{st}^j + l_{st}^j, \qquad j=1,\ldots,m,\ t=1,\ldots,T. \quad s=1,\ldots,t. \tag{4}$$

$$\sum_{j=1}^{m}\sum_{t=s}^{T} l_{st}^j \le K_s, \qquad s=1,\ldots,T, \tag{5}$$

$$\sum_{t=1}^{T} b_t^j \le r_j, \qquad j=1,\ldots,m, \tag{6}$$

$$l_{st} \ge 0, \qquad t=1,\ldots,T,\ s=1,\ldots,t. \tag{7}$$

The above dual program is very similar to the dual used in [13, 14], and the main difference is the additional constraint (6). As in [13, 14], the primal-dual algorithm works in two stages. We first construct a feasible dual solution, and then use this solution to construct a cheap feasible integer solution for the primal LP.

Constructing a dual solution. We associate a budget b_t^j with each demand point (j,t) (i.e., the required demand of market j in period t, d_t^j). We now start to increase the budget using a mechanism that is called *wave form*. Think of a wave the moves backward in time, and let τ be the indicator of the *wavefront* location. We initialize the wavefront variable τ to T. The algorithm consists of a series of iterations as the value of τ is (continuously) decreased through the interval $[T,1]$. This parameter controls the values of the budgets b_t^j of each unfrozen demand point (j,t). The budget b_t^j is kept equal zero as long as $\tau > t$. Once $\tau \le t$ and until it is frozen, b_{it}^j is always equal $H_{\tau t}^j$. As the wave moves backward in time, we will *temporarily open orders* and *freeze* budgets (i.e., stop increasing then) of demand points; as the budgets are increased we identify the following events:

Event 1 When $\tau = s$ (for $s = T, T-1, \ldots, 1$), we consider all unfrozen demand points (j,t) with $t \ge s$, and start increasing the variable l_{st}^j at the same rate as b_t^j, i.e., we keep $b_t^j = H_{\tau t}^j = H_{st}^j + l_{st}^j$; (Note that as the wavefront reaches s and the budget b_t^j increases to H_{st}^j the constraint (4) becomes tight).

Event 2 Suppose that for some s, we have that $\sum_{j=1}^{m} \sum_{t \geq s} l_{st}^{j} = K_s$. (Note that this means that we can no longer increase any variables l_{st}^{j} without violating the constraint (5)). We then temporarily open an order in period s, and freeze all unfrozen budgets of demand points (j, t) with $t \geq s$.

Event 3 Suppose that for some market j we have $\sum_{t=1}^{T} b_t^j = r_j$ (i.e., constraint (6) becomes tight). We then *mark* market j and freeze *all* the budgets of its demand points. This includes also budgets of demand points with $t < \tau$ that will be kept at zero even after τ will advance and cross time t.

We continue this procedure until all the budgets are frozen. Let $(\bar{b}, \bar{l})$ be the dual solution at the end of the first phase of the algorithm. It is straightforward to verify that this solution is feasible with respect to (D) above. Next we show how to use this solution to construct a feasible solution for the market selection problem.

Constructing an integer primal solution. We reject all markets that were marked by event 3 above. We then construct a production plan for the rest of the markets following a similar procedure to the one described in [13, 14]. Let $R = \{s_1 = 1 < s_2 < \cdots < s_p\}$ be the set of the time periods of all temporarily opened orders. For each $s \in R$, let $open(s)$ be the location of the wavefront when the order at s was temporarily opened. We say that the interval $[open(s), s]$ is the *shadow interval* of s. Furthermore, r and s in R are said to be *dependent* if and only if their shadow intervals intersect. We consider the periods s_i, $i = 1, \ldots, p$, in increasing order of s_i, and permanently open an order s_q whenever its associated shadow interval does not intersect the shadow interval of any earlier s_i, $i = 1, \ldots, q-1$, that has already been permanently opened. The set of permanently opened orders provides a feasible production plan, where each demand point of a selected market is satisfied from the latest possible opened order.

Analysis. The analysis follows along the lines of [13, 14]. By similar arguments to the ones used there, we can show that it is possible to 'pay' for the cost of the production plan of the solution, using every budget $\bar{b}_t^j$ at most once. In addition, it is clear that the rejection cost can be 'paid' with the budgets associated with the demands of the rejected markets. Therefore, the total cost of the solution can be 'paid' using each budget $\bar{b}_t^j$ at most twice. For the JRP and assembly problem the modification in the algorithm is similar. Here it can be shown that the production cost can be 'paid' using each budget at most twice, and the rejection cost is again 'paid' by the budgets of the rejected markets. We then conclude the following theorem.

Theorem 8. *The modified primal-dual framework provides approximation algorithms with worst-case performance guarantees of 2, 3 and 3, respectively, for the market selection variants of the lot-sizing, the JRP and the assembly problem.*

References

1. E. Arkin, D. Joneja, and R. Roundy. Computational complexity of uncapacitated multi-echelon production planning problems. *Operations Research Letters*, 8:61–66, 1989.
2. E. Balas. The prize collection travelling salesman problem. *Networks*, 19:621–636, 1989.

3. I. Bárány, T. J. Van Roy, and L. A. Wolsey. Uncapacitated lot-sizing: the convex hull of solutions. *Mathematical Programming Study*, 22:32–43, 1984.
4. Y. Bartal, S. Leonardi, A. M. Spaccamela, J. Sgall, and L. Stougie. Multiprocessor scheduling with rejection. *SIAM Journal on Discrete Mathematics*, 13:64–78, 2000.
5. D. Bertsimas, C. Teo, and R. Vohra. On dependent randomized rounding algorithms. *Operations Research Letters*, 25:105–114, 1999.
6. D. Bienstock, M. X. Goemans, D. Simchi-Levi, and D. Williamson. A note on the prize collecting traveling salesman problem. *Mathematical Programming*, 59:413–420, 1993.
7. M. Charikar, S. Khuller, D. M. Mount, and G. Narasimhan. Algorithms for facility location problems with outliers. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 642–651, 2001.
8. J. Geunes, H. E., Romeijn, and K. Taaffe. Requirements planning with pricing and order selection flexibility. To appear in Operations Research, 2004.
9. M. X. Goemans. Approximate solutions of hard combinatorial problems, 1996. Lecture notes in the schoolof ASHCOMP.
10. M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24:296–317, 1995.
11. K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM 48*, pages 274–296, 2001.
12. J. Krarup and O. Bilde. Plant location, set covering and economic lot sizing: an O(mn) algorithm for structural problems. In *Numerische Methoden Bei Optimierungsaufgaben*, volume 3, pages 155–180, 1977.
13. R. Levi, R. O. Roundy, and D. B. Shmoys. Primal-dual algorithms for deterministic inventory problems. Technical Report TR1042, ORIE Department, Cornell University, 2004. Submitted.
14. R. Levi, R. O. Roundy, and D. B. Shmoys. Primal-dual algorithms for deterministic inventory problems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 353–362, 2004.
15. M. Mahdian, Y. Ye, and J. Zhang. Improved approximation algorithms for metric facility location. In *Proceedings of 6th APPROX*, pages 229–242, 2002.
16. M. Mahdian, Y. Ye, and J. Zhang. A 2-approximation algorithm for the soft-capacitated facility location problem. In *Proceedings of 7th APPROX*, pages 129–140, 2003.
17. Y. Pochet and L. A. Wolsey. Lot-sizing with constant batches: *Formulation* and valid inequalities. *Mathematics of Operations Research*, 18:767–785, 1993.
18. R.Levi, R. Roundy, and D. Shmoys. A constant approximation algorithm for the one-warehouse multi-retailer problem. Technical Report TR1408, ORIE Department, Cornell University, 2004. Submitted, extended abstract will appear in SODA 2005.
19. R.Levi, R. Roundy, and D. Shmoys. A constant approximation algorithm for the one-warehouse multi-retailer problem (Extended Abstract). In *Proceedings of the 16th Annual SIAM-ACM Symposium on Discrete Algorithms*, 2005.
20. D. B. Shmoys. The design and analysis of approximation algorithms: facility location as a case study. In S. Hosten, J. Lee, and R. Thomas, editors, *Trends in Optimization. AMS Proceedings of Symposia in Applied Mathematics*, 2004. To appear.
21. D. B. Shmoys and C. Swamy. Stochastic optimization is (almost) as easy as deterministic optimization. In *Proceedings of the 45th Annual IEEE Symposium on the Foundations of Computer Science, 2004.*, pages 228–237, 2004.
22. D. B. Shmoys, C. Swamy, and R. Levi. Facility location with service installation costs. In *Proceedings of the 15th Annual SIAM-ACM Symposium on Discrete Algorithms*, pages 1081–1090, 2004.

23. H. M. Wagner and T. M. Whitin. Dynamic version of the economic lot sizing model. *Management Science*, 5:89–96, 1958.