Exact Algorithms for Integrated Facility Location and Production Planning Problems

Thomas C. Sharkey,[1] Joseph Geunes,[2] H. Edwin Romeijn,[3] Zuo-Jun Max Shen[4]

[1] *Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, Troy, New York 12180*

[2] *Department of Industrial and Systems Engineering, University of Florida, Gainesville, Florida 32611-6595*

[3] *Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109-2117*

[4] *Department of Industrial Engineering and Operations Research, University of California, Berkeley, California 94720-1777*

Received 13 June 2009; accepted 8 February 2011
DOI 10.1002/nav.20458
Published online 31 May 2011 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: We consider a class of facility location problems with a time dimension, which requires assigning every customer to a supply facility in each of a finite number of periods. Each facility must meet all assigned customer demand in every period at a minimum cost via its production and inventory decisions. We provide exact branch-and-price algorithms for this class of problems and several important variants. The corresponding pricing problem takes the form of an interesting class of production planning and order selection problems. This problem class requires selecting a set of orders that maximizes profit, defined as the revenue from selected orders minus production-planning-related costs incurred in fulfilling the selected orders. We provide polynomial-time dynamic programming algorithms for this class of pricing problems, as well as for generalizations thereof. Computational testing indicates the advantage of our branch-and-price algorithm over various approaches that use commercial software packages. These tests also highlight the significant cost savings possible from integrating location with production and inventory decisions and demonstrate that the problem is rather insensitive to forecast errors associated with the demand streams. © 2011 Wiley Periodicals, Inc. Naval Research Logistics 58: 419–436, 2011

Keywords: facility location; production planning; branch and price; order selection

1. INTRODUCTION

The design of a supply chain network structure has significant impacts on operating costs. The locations of supply facilities, the assignment of customer demands to these facilities, and the management of production and inventory at the facilities are all important contributors to overall supply chain performance. When these location, assignment, and production planning decisions are undertaken separately, the supply chain may end up incurring significantly (and unnecessarily) high operating costs. A potential for significant cost savings can often arise as a result of considering these decisions simultaneously in the planning phase. This article considers such an integrated location and production planning problem, where individual customers have dynamic demand streams over a finite horizon. We must determine a set of operational facilities and assign the demand of each customer in every time period to an open facility. We must then manage the production and inventory levels within each open facility in order to meet assigned demands in every time period. The goal is to minimize total facility opening costs, assignment costs, and production planning costs incurred while meeting all customer demands.

Typical facility location problems (see, for example, Daskin [7]) seek an assignment of customers to open facilities that minimizes the sum of the assignment costs (of customers to open facilities) and the facility opening costs. These facility location problems are often static in nature, meaning that they cannot, except at a very coarse level, model problems where customer demands vary over time. Dynamic facility location problems (see Wesolowsky [28], Van Roy and Erlenkotter [20], and Chardaire et al. [4]) are discrete models in which customer demands vary over time where we wish to assign the demand of each customer in each time period to an open facility. These dynamic models are not only interested in determining where to locate facilities but also

Correspondence to: T.C. Sharkey (sharkt@rpi.edu)

when to open the facilities. In these dynamic models, each customer must be assigned to an open facility in every time period and the facility must meet the demand of its assigned customers. However, these dynamic location models do not capture the typical economies of scale observed in production at the facilities. Further, and more importantly, these dynamic models assume that all demand assigned to a facility in a time period is produced during that time period, i.e., they do not permit facilities to carry inventory between time periods.

Typical modeling approaches attempt to capture the production and inventory costs at a facility as a component of the linear assignment or connection cost incurred when serving a customer from the facility. However, this can typically only serve as an approximation because of the economies of scale often seen in production (for example, fixed-charge production set up costs), which require knowledge of the entire set of demands assigned to the facility before computing total production cost. Effective model precision and accuracy requires location models that explicitly account for the production and inventory decisions at each facility. Recent research has incorporated production and/or inventory decisions and associated costs within facility location problems to varying degrees. Daskin et al. [6], Shen et al. [23], and Shu et al. [24] consider joint inventory-location models where each customer has uncertain demand (although the parameters of the associated probability distributions are static over time) and each facility must determine when to produce (or reorder) and the amount of safety stock to hold to meet a prescribed system-wide service level. Huang et al. [12] consider a continuous-time single-sourcing problem where each customer has a constant, static demand rate and each facility must determine its optimal order policy based on the total demand rate assigned to it in the face of classic economic order quantity costs. These models incorporate production and/or inventory decisions at facilities, but customer demands (or the distributions of these demands) are assumed to be stationary over time.

The previous research on assigning customers to facilities when customers have dynamic demand has focused on the multi-period single-sourcing problem (MPSSP); however, this problem does not account for strategic facility opening decisions (see Romeijn and Romero Morales [16–18] and Freling et al. [9]). Moreover, previous work on the MPSSP has made the simplifying assumption that production and inventory costs are linear (although production and inventory quantities may be capacitated), meaning that this work does not account for economies of scale in production. Freling et al. [9] developed an exact branch-and-price algorithm for a special case of the MPSSP where the demand patterns follow the same seasonal pattern and the goal is to assign a customer to the same facility for the entire horizon. This special case of the MPSSP can be reduced to a location problem with static demand, for which the branch-and-price algorithm was developed.

The integrated facility location and production planning models developed in this article overcome several of the shortcomings of the aforementioned work by including facility location decisions, dynamic customer demands, and economies of scale in production within the model. These models will make the practical assumption, typically seen in distribution systems, that we operate under a single-sourcing strategy. This means that a customer's demand in a period is not split among multiple supply facilities. This strategy offers several practical advantages, such as reduced coordination complexity between facilities, and has been assumed in the joint inventory-location model, the continuous-time single-sourcing problem, and the MPSSP. In the context of our problem where customers have dynamic demand streams, we will interpret the single-sourcing strategy to mean that a customer's demand in any given period is not split across multiple facilities. We will, however, allow assignment decisions in our location model to vary dynamically. The branch-and-price approaches previously developed for the joint inventory-location model, the continuous-time single-sourcing problem, and the MPSSP did not permit dynamic assignments.

As noted previously, we develop exact branch-and-price algorithms to solve our integrated facility location and production planning models. Romeijn et al. [19] considered a similar class of integrated location and production planning problems, but they were concerned with developing approximation algorithms for this class of problems. In particular, Romeijn et al. [19] require that a customer is assigned to the same facility over the problem's horizon and prove that it is unlikely that there exists a constant factor approximation algorithm for this problem. For the special case where the demand patterns of the customers follow the same seasonal trend, Romeijn et al. [19] develop a 1.52-approximation algorithm. The problems we consider in this article are strategic planning problems and fast solution times are, therefore, not of paramount importance. This means that finding a good solution quickly may not be as important as finding an optimal solution in a time frame that can aid in supply chain planning activities. Further, we are able to explore certain characteristics of our problem class by determining an optimal solution using a branch-and-price algorithm. This permits characterizing potential cost savings by comparing the integrated model solution to that obtained using a sequential approach that first solves the classical facility location problem and then determines the optimal production and inventory levels for the corresponding facility set. We also explore the sensitivity of our problem class to forecast errors in the customer demand streams. These comparisons would be less meaningful if we were to base them on heuristic solutions.

We use a branch-and-price algorithm to solve the set-covering formulation of the integrated facility location and production planning problem. Savelsbergh [21] first developed a branch-and-price algorithm for a set-partitioning formulation of the Generalized Assignment Problem. There has been much success (see, for example, Shen et al. [23], Freling et al. [9], Huang et al. [12], and Shu et al. [24]) in applying branch-and-price algorithms to set-partitioning or set-covering formulation(s) of customer assignment models. It turns out that the pricing problems that arise in these algorithms are often interesting supply chain planning problems in their own right. In particular, the pricing problem is typically a supply chain planning problem with customer (or order) selection. In these problems, the supplier selects a subset of customers that maximizes profit, i.e., revenues from customers minus the costs necessary to serve the customers. In the context of our integrated facility location and production planning model, this pricing problem takes the form of a production planning problem with order selection. This problem is a generalization of the problem studied by Geunes et al. [10] and can also be interpreted as a joint pricing and production planning problem. The pricing problem in this article considers general concave production cost functions, as opposed to the fixed-charge plus linear production cost functions in Geunes et al. [10]. Therefore, the contribution of this article lies not only in the integrated location and production planning problem but also in developing algorithms to solve this production planning problem with order selection and various extensions thereof.

The remainder of this article is organized as follows. We formally describe and formulate the integrated facility location and production planning problem and discuss its applications in Section 2. The set-covering formulation and the branch-and-price algorithm for this problem are then discussed in Section 3. We discuss the pricing problem, i.e., the production planning and order selection problem, in Section 4. We provide the results of our computational testing in Section 5. These computational tests include comparing our algorithm to various approaches using commercial solvers. They also analytically characterize the value of integrating location with production and inventory decisions and examine the sensitivity of our model to forecast errors in customer demand streams. We discuss how to modify the branch-and-price algorithm to account for several important variants of our integrated location and production planning problem in Section 6. These extensions lead to several variants of the production planning and order selection problem. We conclude the article in Section 7.

2. PROBLEM STATEMENT

The integrated location and production planning problem can be mathematically formulated as follows. We are given a set of m facilities, n customers, and T time periods. The demand of customer j in period t is given by d_{jt} $(j = 1, \ldots, n, t = 1, \ldots, T)$. A connection cost, c_{ijt}, is associated with facility i and customer j in time period t, which is expressed as a cost per unit of demand. Each facility has an associated opening cost of f_i, which is incurred if we assign any customers to the facility. Each facility i has an associated concave cost function that corresponds to the cost of producing p units in time period t, $P_{it}(p)$, and a per-unit inventory cost, h_{it}, for holding a unit of inventory in time period t. We allow customer assignments to vary dynamically over the horizon, i.e., in each time period t we must determine which facility will supply the demand of customer j. We thus define a binary variable x_{ijt} for $i = 1, \ldots, m$, $j = 1, \ldots, n$, and $t = 1, \ldots, T$, which equals 1 if we assign customer j to facility i in time period t, and equals 0 otherwise. The binary variable y_i for $i = 1, \ldots, m$ is equal to 1 if facility i is open. We further define the variables p_{it} and I_{it} for $i = 1, \ldots, m$ and $t = 1, \ldots, T$ as the production and inventory levels, respectively, at facility i in time period t. Our integrated facility location and production planning problem can be formulated as follows:

$$\text{minimize} \sum_{i=1}^{m} \left(f_i y_i + \sum_{t=1}^{T} (P_{it}(p_{it}) + h_{it} I_{it}) \right) + \sum_{i=1}^{m} \sum_{j=1}^{n} \sum_{t=1}^{T} d_{jt} c_{ijt} x_{ijt}$$

subject to (P)

$$\sum_{i=1}^{m} x_{ijt} \geq 1 \qquad \text{for } j = 1, \ldots, n,\ t = 1, \ldots, T \tag{1}$$

$$x_{ijt} \leq y_i \qquad \text{for } i = 1, \ldots, m,\ j = 1, \ldots, n,\ t = 1, \ldots, T \tag{2}$$

$$x_{ijt} \in \{0, 1\} \qquad \text{for } i = 1, \ldots, m,\ j = 1, \ldots, n,\ t = 1, \ldots, T \tag{3}$$

$$y_i \in \{0, 1\} \qquad \text{for } i = 1, \ldots, m \tag{4}$$

$$I_{i,t-1} + p_{it} = \sum_{j=1}^{n} d_{jt} x_{ijt} + I_{it} \qquad \text{for } i = 1, \ldots, m,\ t = 1, \ldots, T \tag{5}$$

$$I_{i0} = 0 \qquad \text{for } i = 1, \ldots, m \tag{6}$$

$$p_{it}, I_{it} \geq 0 \qquad \text{for } i = 1, \ldots, m,\ t = 1, \ldots, T. \tag{7}$$

Constraints (1)–(4) are traditional (dynamic) facility location constraints and constraints (5)–(7) are production planning constraints at each facility to ensure that, in each time period, we meet the demand of all customers assigned to the facility. If we know the set of customers assigned to a given facility in all time periods, then we simply need to solve an uncapacitated production planning problem at the facility. This production planning problem is a generalization of the classic economic lot-sizing problem (see Wagner and Whitin [27]), where concave production cost functions replace fixed-charge plus linear production costs. This subproblem can be solved via a dynamic programming algorithm in $O(T^2)$ time (see Veinott [25] or Denardo [8]). In terms of the facility location model, we can essentially view each customer j as T separate customers, one for each time period, since we can assign the demand of customer j to different facilities in different time periods.

The application of problem (P) is appropriate in the strategic planning phase when evaluating supply chain design options. The purpose of such models is typically to evaluate supply chain design alternatives and to aid in long-range budget planning. As noted by Shapiro [22], accounting for inventory costs in such models serves to approximate the way in which inventory plans influence supply chain costs. Thus, our model provides a mechanism to more closely approximate these production- and inventory-related costs when economies of scale in production strongly influence total costs. Application of the model therefore requires judicious balancing between long-term (location) costs and shorter-range (production- and inventory-related) costs. While a deterministic demand assumption is often employed in strategic planning models in the literature (and in location models in particular), the deterministic demands used in our model are clearly used to represent projected customer demands over a long-range (or strategic) planning horizon, and how these contribute to operational costs. Although a deterministic demand assumption is often an abstraction from reality, as we later show in Section 5.3, the longer-term strategic location decisions our model provides are reasonably insensitive to errors in forecasted demand, while the resulting model costs provide good estimates of overall costs for use in budget planning. We also recognize the limitations introduced by ignoring facility capacity limits. This limitation may necessitate use of a capacity requirements planning or finite-capacity scheduling function for application to existing facilities with hard capacity limits (see Zipkin [30] for a discussion on how similar limitations are addressed in existing material requirements planning systems). Alternatively, for new facility locations, we might account for a linear approximation of the cost of capacity as part of the facility connection cost parameters, and plan facility capacity levels based on the demands assigned to a facility in the optimal solution.

Problem (P) can be used to solve problems where the costs and customer demand patterns repeat over time. In such situations, the problem horizon is then interpreted as the duration of the pattern before it repeats. We note that if the cost or demand patterns at different customers have different cycle lengths, then the overall cycle length (or horizon of (P)) will be the least common multiple of the different cycle lengths of the cost and demand patterns. In these settings, the facility opening costs, f_i, might contain (i) an annualized opening cost and (ii) the operating cost of the facility over the horizon. It may also be appropriate to model the problem so that the inventory entering the first period equals the inventory leaving the last period (i.e., $I_{i0} = I_{iT}$ for all $i = 1, \dots, m$). It is not difficult to show that in an optimal solution to (P), we have $I_{i0} = I_{iT} = 0$ for all $i = 1, \dots, m$, but we may want to allow these inventory levels to take non-negative values in practice. We show that our branch-and-price algorithm can solve this extension of (P), which we refer to as cyclic, in Section 6.2.

Problem (P) can also be applied to multi-item production planning problems with component substitutions (see, for example, Balakrishnan and Geunes [1] and Wu and Golbasi [29]). In this setting, each facility corresponds to a particular "component" that may be used in producing different end products (the customers). Each end product has an associated demand stream over the finite horizon. The connection cost of end product j to component i is then interpreted as the "substitution" or "conversion" cost for using component i in end product j. The facility opening cost for component i can then be interpreted as the cost of acquiring the capabilities of producing the component (e.g., a design cost). We then account for economies of scale in producing component i in each of the time periods over the horizon of our production planning problem, and the ability to hold components in inventory. Balakrishnan and Geunes [1] considered this multi-item production planning problem with substitutions in the absence of component acquisition costs and when the production cost of a component consists of a fixed charge setup cost plus a per unit variable production cost. Our problem (P) can thus be used to model more general situations than considered by Balakrishan and Geunes [1], where we have component acquisition costs and a more general form of economies of scale in component production.

3. A BRANCH-AND-PRICE ALGORITHM

In this section, we develop a branch-and-price algorithm for (P) that will solve the problem to optimality. The current formulation of (P) is a large-scale mixed integer nonlinear programming problem. We develop an equivalent formulation of (P) that removes the nonlinearity of the objective function of (P) by posing the problem as a set-covering problem. However, this set-covering formulation contains an exponential number of variables. To address this problem, we use the column generation approach of Gilmore and Gomory [11] to solve the relaxation of the set-covering formulation. This procedure will be used to calculate the linear programming (LP) bounds at each node of a branch-and-bound tree, leading to a branch-and-price algorithm to solve (P).

It is easy to see that any feasible solution to (P) can be viewed as a partition of the customer/time period pairs into m subsets, each of which is assigned to a facility. Denote the number of distinct subsets that can be feasibly assigned to a facility by L. In particular, $L = 2^{nT}$ since each customer/time period pair can be included or not included in the subset. Furthermore, we let the binary matrix $\alpha^{\ell}_{i\cdot\cdot}$ represent the ℓth subset associated with facility i, where $\alpha^{\ell}_{ijt} = 1$ if customer/time period pair (j, t) belongs to the subset, and 0 otherwise. For consistency with past approaches (for example, Savelsbergh [21]), we will also refer to $\alpha^{\ell}_{i\cdot\cdot}$ as the ℓth column associated with facility i. Furthermore, let $\zeta_i(\alpha^{\ell}_{i\cdot\cdot})$ denote the cost associated with serving the set of customers/time periods given by $\alpha^{\ell}_{i\cdot\cdot}$ at facility i. If we then define the decision variable y^{ℓ}_i to be equal to 1 if facility i serves the ℓth associated subset and 0 otherwise, the set-covering formulation (SC) of (P) reads:

$$\text{minimize} \sum_{i=1}^{m}\sum_{\ell=1}^{L} \zeta_i\left(\alpha^{\ell}_{i\cdot\cdot}\right) y^{\ell}_i$$

subject to (SC)

$$\sum_{i=1}^{m}\sum_{\ell=1}^{L} \alpha^{\ell}_{ijt} y^{\ell}_i \geq 1 \quad \text{for } j = 1,\ldots,n,\ t = 1,\ldots,T \tag{8}$$

$$y^{\ell}_i \in \{0,1\} \quad \text{for } i = 1,\ldots,m,\ \ell = 1,\ldots,L.$$

Without loss of generality, we assume that $\alpha^{1}_{i\cdot\cdot}$ represents the empty set, so clearly $\zeta_i(\alpha^{1}_{i\cdot\cdot}) = 0$. Otherwise, we have

$$\zeta_i\left(\alpha^{\ell}_{i\cdot\cdot}\right) = f_i + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}\alpha^{\ell}_{ijt} + g_i\left(\sum_{j=1}^{n} d_{j1}\alpha^{\ell}_{ij1}, \ldots, \sum_{j=1}^{n} d_{jT}\alpha^{\ell}_{ijT}\right), \tag{9}$$

where $g_i(D_1, \ldots, D_T)$ is equal to the optimal solution value of an associated production planning problem:

$$\text{minimize} \sum_{t=1}^{T} (P_{it}(p_{it}) + h_{it}I_{it})$$

subject to

$$I_{i,t-1} + p_{it} = D_t + I_{it} \quad \text{for } t = 1,\ldots,T$$
$$I_{i0} = 0$$
$$p_{it}, I_{it} \geq 0 \quad \text{for } t = 1,\ldots,T.$$

For convenience, throughout the remainder of this paper, we will refer to v(*Problem*) as the value of the optimal solution to *Problem*. We will also refer to *ProblemR* as the problem that is obtained if the integrality constraints in *Problem* are relaxed. It is easy to see that v(SC) is equal to v(P). Therefore, v(SCR) is a lower bound on v(P). We will prove that v(PR) provides a tighter bound on v(P) than v(SCR) in Section 3.2. However, we first describe the branch-and-price algorithm that we will use to solve (SC).

3.1. Components of the Algorithm

We next discuss a branch-and-price algorithm for solving (SC). We solve the relaxation of (SC) through column generation (see, for example, Gilmore and Gomory [11]) at each node of a branch and bound tree. There are many important factors involved in designing an efficient and effective branch-and-price algorithm (e.g., Barnhart et al. [3]). In this section, we focus on two of the main theoretical factors that arise in the design of a branch-and-price algorithm for our problem: (i) the pricing problem in the algorithm and (ii) the branching rules in the algorithm.

3.1.1. The Pricing Problem

The most vital factor in the success of a branch-and-price algorithm is the ability to solve (either exactly or heuristically) the pricing problem that arises as a result of the column generation approach. In this section, we derive the optimization problem that arises as the pricing problem for solving problem (P). Suppose that N is the current set of columns in the reduced LP relaxation of (SC). We let (SCR(N)) denote the corresponding relaxation and let $\mu^*(N)$ be the optimal dual multipliers associated with constraints (8). The columns will be evaluated for each facility $i = 1, \ldots, m$ in order to determine if there is an eligible nonbasic variable (or column) to enter the basis. In particular, we must determine if a column exists such that the optimal solution value to the problem

$$\min_{x_{i\cdot\cdot}\in\{0,1\}^{n\times T}} \left\{ \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}x_{ijt} + g_i\left(\sum_{j=1}^{n} d_{j1}x_{ij1}, \ldots, \sum_{j=1}^{n} d_{jT}x_{ijT}\right) - \sum_{j=1}^{n}\sum_{t=1}^{T} \mu^*_{jt}(N)x_{ijt} + f_i \right\} \tag{10}$$

is negative. We may disregard the constant f_i in solving this problem and define $r_{ijt} = \mu^*_{jt}(N) - d_{jt}c_{ijt}$, which can be interpreted as the revenue associated with selecting customer j. This means that the pricing problem can be formulated as

$$\max_{x_{i\cdot\cdot}\in\{0,1\}^{n\times T}} \left\{ \sum_{j=1}^{n}\sum_{t=1}^{T} r_{ijt}x_{ijt} - g_i\left(\sum_{j=1}^{n} d_{j1}x_{ij1},\ldots,\sum_{j=1}^{n} d_{jT}x_{ijT}\right)\right\},$$

where $g_i(\sum_{j=1}^{n} d_{j1}x_{ij1},\ldots,\sum_{j=1}^{n} d_{jT}x_{ijT})$ represents the optimal cost of serving the demands of the selected customers through production and inventory decisions at facility i. This formulation of the pricing problem can be interpreted as selecting the set of customers that maximize our profits, i.e., the revenues received from them minus the costs associated with serving the customers. We refer to this problem as the production planning and order selection problem (PPOSP), which we will study in Section 4.

3.1.2. The Branching Rule

If the optimal solution to (SCR) is not integral, one can add cutting planes to (SCR) or apply branch-and-bound in order to obtain an integral optimal solution to solve (SC). A natural choice for a branching rule would be to branch on the variables y_i^ℓ, but as mentioned in Savelsbergh [21], this branching rule destroys the structure of the pricing problem to be solved at successive nodes of the branch-and-bound tree. In particular, this branching rule may require us to determine a kth best solution to the pricing problem, where k is equal to the depth of the current node in the branch and bound tree. If the solution to (SCR) is fractional, there exists at least some y_i^ℓ with $\ell \geq 2$ that is fractional. This implies that at least one underlying location variable, y_i, and one underlying customer assignment variable, x_{ijt}, in (P) are also fractional. This leads to two potential branching rules: (i) branch on the underlying location variable or (ii) branch on the underlying customer assignment variable. The fact that we branch on these underlying variables preserves the structure of our pricing problems at the nodes of the branch and bound tree (see, for example, Savelsbergh [21]). In most situations, it will be advantageous to branch on the underlying location variable first, since we can often expect that a set of customer assignment variables will be fractional in the solution to (SCR).

3.2. Tightness of the Set Covering Formulation

In this section, we will focus on comparing the tightness of (SC), i.e., the value v(SCR), and the tightness of (P), i.e., the value v(PR). We will prove that v(PR) provides a better lower bound on v(P) than v(SCR). The main idea behind our proof is to show that: (i) the problem (PR) is equivalent to another problem, which we denote by (P′R), and (ii) any solution to (P′R) (and hence (PR)) can be converted to a feasible solution to (SCR) with an objective function that is less than or equal to the objective function in (P′R). We summarize this result in the following theorem.

THEOREM 3.1: $v(\text{PR}) \geq v(\text{SCR})$.

PROOF: The full proof appears in the appendix. □

Despite this negative result about the tightness of (SCR) compared with the tightness of (PR), there is still an inherent advantage to solving (SCR) rather than (PR). This is because (SCR) is linear (although the pricing problems are nonlinear) and (PR) is nonlinear. Therefore, if we can solve the pricing problem for (SCR) effectively, we generally expect that the branch-and-price algorithm will perform well. Further, our computational results indicate that (SCR) (and, therefore, (PR)) often provides the optimal objective function value to (P).

4. THE PRODUCTION PLANNING AND ORDER SELECTION PROBLEM

In this section, we consider the pricing problem associated with (SCR). Recall that we determined previously that the pricing problem associated with (SCR) can be formulated as

$$\max_{x_{i\cdot\cdot}\in\{0,1\}^{n\times T}} \left\{ \sum_{j=1}^{n}\sum_{t=1}^{T} r_{ijt}x_{ijt} - g_i\left(\sum_{j=1}^{n} d_{j1}x_{ij1},\ldots,\sum_{j=1}^{n} d_{jT}x_{ijT}\right)\right\},$$

where g_i represents the optimal cost of serving the demands of the selected customers through production and inventory decisions at facility i. We can suppress the index i in the decision variables and include the production and inventory decisions explicitly in our formulation of the pricing problem, i.e., we can formulate it as

$$\text{maximize} \sum_{j=1}^{n}\sum_{t=1}^{T} r_{jt}x_{jt} - \left(\sum_{t=1}^{T}(P_t(p_t) + h_t I_t)\right)$$

subject to (PPOSP)

$$\begin{aligned} I_{t-1} + p_t &= \sum_{j=1}^{n} d_{jt}x_{jt} + I_t \quad \text{for } t = 1,\ldots,T \\ I_0 &= 0 \\ p_t, I_t &\geq 0 \quad \text{for } t = 1,\ldots,T \\ x &\in \{0,1\}^{n\times T}. \end{aligned}$$

The (PPOSP) can be intuitively described as follows: a supplier is given a set of orders, each with an associated demand level, revenue, and time period. The supplier wishes to determine a selection of orders that maximizes profit, i.e., the revenues received from selected orders minus the production and inventory costs associated with meeting those orders. The (PPOSP) is closely related to the problem considered by Geunes et al. [10]. In their problem, they consider an integrated pricing and lot-sizing model where the demand to be satisfied in each time period is determined by the price level. They assume that the revenue function in each time period is piecewise linear and concave and show that the problem can be interpreted as an order selection problem. This means that the problem of Geunes et al. [10] is a special case of the (PPOSP) where the production cost functions are those of the traditional economic lot-sizing problem, i.e., a fixed-charge setup cost plus a per-unit production cost. They propose a dynamic programming algorithm for this special class of the (PPOSP) that runs in $O(nT^2)$ time where n is the largest number of orders in a single time period (i.e., the number of customers in the problem).

The dynamic programming (DP) algorithm we use to solve (PPOSP) is quite similar to the DP algorithm for the classic production planning problem (see Denardo [8]). It relies on the fact that we can decompose the optimal solution to the (PPOSP) into a sequence of subplans defined by the periods in which we hold zero inventory. A subplan is defined by two time periods t and t' such that $t < t'$ with the property that we have no inventory entering and leaving the subplan (i.e., $I_{t-1} = I_{t'-1} = 0$), and we produce all the selected demand between periods t and $t' - 1$ using production in time period t. It is well known that if production and inventory costs are concave in the production planning problem, then an optimal solution exists that is composed of a series of subplans (see, for example, Wagner [26], Veinott [25], and Love [14]). This holds true for the (PPOSP) as well because if we know the optimal selection of customers, the problem reduces to a production planning problem with concave cost functions.

For our problem, we define the function $p(t, t')$ (with $t < t'$) to be equal to the maximum profit obtained in the subplan (t, t'). The value of $p(t, t')$ is equal to the optimal solution value of the problem

$$\text{maximize} \sum_{s=t}^{t'-1}\sum_{j=1}^{n} r_{js}x_{js} - P_t\left(\sum_{s=t}^{t'-1}\sum_{j=1}^{n} d_{js}x_{js}\right) - \sum_{s=t}^{t'-1}\sum_{s'=t}^{s-1} h_{s'}\sum_{j=1}^{n} d_{js}x_{js}$$

subject to (SP-P)

$$x_{js} \in \{0, 1\} \text{ for } j = 1, \ldots, n, \text{ for } s = t, \ldots, t' - 1.$$

We can simplify the objective function of this problem by grouping together the revenues and inventory holding costs of each customer in the problem. In other words, we define $R_{js} = r_{js} - \sum_{s'=t}^{s-1} h_{s'}d_{js}$ and we can formulate (SP-P) as

$$\text{maximize} \sum_{s=t}^{t'-1}\sum_{j=1}^{n} R_{js}x_{js} - P_t\left(\sum_{s=t}^{t'-1}\sum_{j=1}^{n} d_{js}x_{js}\right)$$

subject to (SP-P′)

$$x_{js} \in \{0, 1\} \text{ for } j = 1, \ldots, n, \text{ for } s = t, \ldots, t' - 1.$$

It has been shown (see Shen et al. [23] and Huang et al. [12]) that a problem of the form (SP-P′) with a concave function P_t can be solved by (i) sorting the variables in nonincreasing order of the ratio R_{js}/d_{js} and (ii) evaluating each solution that selects the first ℓ customers in this ordering. In the worst case, step (i) requires $O(n(t' - t)\log n(t' - t)) = O(nT \log nT)$ to sort the variables/customers and step (ii) requires $O(n(t' - t)) = O(nT)$ time. This means that to calculate $p(t, t')$ for all $t = 1, \ldots, T$ and $t' = t + 1, \ldots, T + 1$ requires $O(nT^3 \log nT)$ time.

By using the value of a subplan, $p(t, t')$ for $t = 1, \ldots, T$ and $t' = t, \ldots, T + 1$, we can provide both a forward and a backward dynamic programming algorithm. We define the function $C(t)$ to be the maximum profit obtained in the problem (PPOSP) by only considering time periods $1, \ldots, t$. We can set $C(0) = 0$ and define our recurrence relationship as

$$C(t) = \max_{s=0,\ldots,t-1}\{C(s) + p(s + 1, t + 1)\} \qquad (11)$$

for $t = 1, \ldots, T$. We can also define the function $C'(t)$ to be the maximum profit obtained in the problem (PPOSP) by only considering time periods $t, t+1, \ldots, T$. We set $C(T+1) = 0$ and define the recurrence relationship as

$$C'(t) = \max_{s=t+1,\ldots,T}\{p(t, s) + C(s)\}.$$

for $t = 1, \ldots, T$. In both these cases, the total time required to solve the (PPOSP) by dynamic programming is $O(nT^3 \log nT)$.

5. COMPUTATIONAL TESTING

In this section, we discuss our computational testing on several issues surrounding the integrated location and production planning problem. We note that all tests were performed on a Dell Optiplex 755 with an Intel Core 2 Duo CPU E8400 3.0 GHz processor and 4 gigabytes of memory. All linear programs and integer programs that arise in the testing were solved using the default options of CPLEX 11.0.

Section 5.1 focuses on the performance of our branch-and-price algorithm for the general problem (P) with concave production cost functions. We compare its performance to the global optimization software BARON and also to an algorithm that solves successive linear approximations of (P) using CPLEX 11.0 (see, for example, Magnanti et al. [15]). For problems with fixed-charge plus linear production production cost functions, preliminary testing demonstrated that CPLEX 11.0 is extremely effective in solving a formulation of (P) based on the plant-location formulation (see Krarup and Bilde [13]) of the economic lot-sizing problem. We, therefore, have focused our testing on the suitability of our approach and the approaches using commercial software packages for problems with general concave production cost functions. Section 5.2 focuses on the value of integration of the location and production planning decisions. Section 5.3 is concerned with examining the impact of errors in the forecasted demand vectors of the customers on the costs and the solutions provided by our model.

There are many characteristics of the branch-and-price algorithm that need to be set in the computational testing. We have used the dual simplex method in CPLEX 11.0 to solve the relaxation(s) of the set-covering problem in the branch-and-price approach. The initial pool of columns in the reduced representation of (SCR) included, for each facility, the column representing the assignment of all the customers and time periods to that facility. The column generation approach used a multiple pricing scheme where we solved the pricing problem for each facility and included any column that had a negative reduced cost (i.e., we include up to m columns, one for each facility, at each iteration of the column generation approach). After the column generation approach solved (SCR), if the current optimal solution was non-integral, we then solved an integer programming problem over the set of columns currently in the reduced formulation to determine if there was an alternative integral optimal solution to (SCR).

In the computational testing, we examine three intervals of facility opening costs: (i) f_i is set equal to zero, i.e., (P) becomes a generalization of the uncapacitated multi-period single-sourcing problem (see Romeijn and Romero Morales [16] for the definition of this problem) with nonlinear (concave) production cost functions, (ii) f_i is uniformly generated on the interval $[100, 250]$, and (iii) f_i is uniformly generated on the interval $[200, 500]$. We have chosen to use square root functions to model economies of scale in production, a typical choice for concave production cost functions. In particular, the production cost function at facility i in time period t is defined by $P_{it}(p_{it}) = A_{it}\sqrt{p_{it}}$, where the term A_{it} is uniformly generated on the interval $[5, 15]$. The holding costs, h_{it}, were uniformly generated on the interval $[1, 4]$. The demands of the customers, d_{jt}, were generated on the interval $[5, 15]$. In order to determine the connection costs, we generated the location of each facility and customer uniformly on the box $[0, 10]^2$. The connection costs per unit of demand, c_{ijt}, were then set equal to the Euclidian distance from the location of facility i to customer j.

5.1. Performance of the Branch-and-Price Algorithm on the General Problem

We first analyze the performance of the branch-and-price algorithm on randomly generated instances of (P) by comparing the performance of the branch-and-price algorithm with the commercial global optimization software BARON. To solve (P), it is necessary to use the mixed-integer nonlinear programming option of BARON. We note that BARON was called using its GAMS interface.

We next compare the performance of the branch-and-price algorithm with the solution of (P) through a series of linear approximations. This type of approach was shown to be successful for other supply chain planning problems with nonlinear economies of scale in Magnanti et al. [15]. At each iteration of this approach (which we will refer to as the successive linear approximation (SLA) approach), we approximate each of the concave production cost functions, P_{it}, by a piecewise linear concave function with δ_{it} breakpoints (we call this function $\hat{P}_{it}$). At each of these δ_{it} breakpoints, the value of the piecewise linear function is equal to the value of the production cost function. This means that the piecewise linear concave function is a lower approximation of P_{it}, and an increased number of breakpoints results in a better approximation. We then replace the production cost function P_{it} with the approximation $\hat{P}_{it}$ in (P). Further, we can represent the piecewise linear concave function by introducing a set of continuous and binary variables. In particular, we formulated each of the piecewise linear functions, $\hat{P}_{it}$, using the so-called multiple choice model (see, for example, Balakrishnan and Graves [2] or Croxton et al. [5]) for converting piecewise linear functions into equivalent mixed-integer linear programming formulations. We have included the details of this conversion in the Appendix. We then solve the resulting mixed-integer linear programming formulation using the integer programming solver within CPLEX 11.0. After determining the optimal solution to this problem, we then either update the collection of breakpoints for one or several of the production cost functions or determine that the solution is also optimal to (P). In particular, if the optimal solution to the approximation has a production level of $\hat{p}_{it}$, we include the value of $\hat{p}_{it}$ as a breakpoint in the approximation of P_{it} if it has not already been included. If all production levels are already breakpoints in the approximations of their respective production cost functions, then the solution is optimal to (P).

We present the results for problems with a modest number of customers ($n = 10$ and $n = 15$) first and then examine problems with a larger number of customers. We examine

Table 1. Computational results for $n = 10$ and $m = 5$ for the branch-and-price algorithm, the SLA approach, and the software package BARON.

		B&P algorithm			SLA				BARON			
f_i	T	Time (s)	S.D.	Wins	Time (s)	S.D.	Wins	% Error	Time (s)	S.D.	Wins	% Error
0	5	0.9	0.6	9	6.3	6.4	1	0	3600^{10}	0	0	0.7
$U[100, 250]$	5	2.8	1.0	7	7.4	4.7	3	0	3600^{10}	0	0	0.2
$U[200, 500]$	5	3.3	1.8	3	3.6	4.3	7	0	2943.5^{8}	1313.5	0	0.7
0	10	68.3	98.4	6	457.0^{1}	1064.3	4	0	3600^{10}	0	0	5.4
$U[100, 250]$	10	111.2	48.1	5	1130.4^{2}	1381.2	5	0.1	3600^{10}	0	0	2.5
$U[200, 500]$	10	166.8	118.8	5	275.9	369.1	5	0	3600^{10}	0	0	1.5

problems with $T = 5$ and $T = 10$ time periods and consider each class of the facility opening costs. We observed that for all instances considered, the branch-and-price algorithm solves (SC) at the root node of the branch and bound tree. This phenomenon has been observed before in solving uncapacitated customer assignment models with concave cost functions (see, for example, Shen et al. [23] and Huang et al. [12]).

5.1.1. Ten Customer Problems

In Tables 1 and 2, we report the average time (over 10 randomly generated instances) required to solve (P) with $n = 10$ customers by our algorithm, the SLA approach, and BARON for $m = 5$ and $m = 10$ facilities, respectively. We imposed a time limit of 3600 seconds (1 hour) in our testing. In the column that reports the time required to solve (P), the superscript indicates the number of instances in which the SLA approach or BARON did not obtain the optimal solution (or verify the optimality) to (P) within an hour. In calculating the average time for classes of problems in which the SLA approach or BARON did not solve an instance within an hour, we included these instances and set their time to 3600 s. This means that the entries are actually lower bounds on the average time required by BARON or the SLA approach. The percentage error column reports the average error of the best solutions obtained within the time limit by the SLA approach or BARON on instances that were not solved. Further, we have presented the standard deviations associated with the running times of each of the three methods and the number of "wins" (out of the 10 instances) for each method. In other words, the number of wins for a method represents the number of instances in which the method found an optimal solution in the shortest amount of time.

The results clearly indicate that both our branch-and-price algorithm and the SLA approach outperform BARON for these tests. For $m = 5$, the branch-and-price algorithm has a better average running time than the SLA approach for all classes and is more robust in the sense that it has a smaller standard deviation (as a percentage of the average running time) for all classes. For $m = 10$, the branch-and-price algorithm outperforms the SLA approach: it has a much better average running time for all classes and outperforms the SLA approach on almost all instances. It is also interesting to note that the SLA approach performs notably worse for $m = 10$. This can be partially attributed to the fact that the ratio of the number of customers to the number of facilities plays an important role in the quality of the linear relaxations of standard customer assignment models. In particular, as the ratio increases the quality of the relaxations become extremely good and, therefore, the commercial software packages benefit from these improved bounds (see, e.g., Savelsbergh [21]). These results indicate that this relationship seems to apply to our class of problems as well.

Table 2. Computational results for $n = 10$ and $m = 10$ for the branch-and-price algorithm, the SLA approach, and the software package BARON.

		B&P algorithm			SLA				BARON			
f_i	T	Time (s)	S.D.	Wins	Time (s)	S.D.	Wins	% Error	Time (s)	S.D.	Wins	% Error
0	5	1.7	1.0	10	13.1	9.9	0	0	3600^{10}	0	0	9.3
$U[100, 250]$	5	2.2	1.0	10	17.4	11.9	0	0	3600^{10}	0	0	2.2
$U[200, 500]$	5	3.2	0.8	6	10.1	8.9	4	0	3600^{10}	0	0	1.1
0	10	13.3	4.6	10	2660.7^{7}	1507.6	0	1.3	3600^{10}	0	0	11.5
$U[100, 250]$	10	63.5	34.1	10	2686.6^{7}	1458.4	0	1.1	3600^{10}	0	0	4.0
$U[200, 500]$	10	243.3	264.0	9	1647.0^{4}	1617.0	1	0.5	3600^{10}	0	0	2.1

Table 3. Computational results for $n = 15$ and $m = 5$ for the branch-and-price algorithm, the SLA approach, and the software package BARON.

		B&P algorithm			SLA				BARON			
f_i	T	Time (s)	S.D.	Wins	Time (s)	S.D.	Wins	% Error	Time (s)	S.D.	Wins	% Error
0	5	3.4	1.4	8	23.6	42.3	2	0	3600^{10}	0	0	1.0
$U[100, 250]$	5	19.7	14.2	4	23.8	39.1	6	0	3600^{10}	0	0	0.7
$U[200, 500]$	5	27.8	16.2	2	22.7	33.7	8	0	3600^{10}	0	0	0.9
0	10	528.2	817.1	5	1626.0^{3}	1643.2	5	0.1	3600^{10}	0	0	4.2
$U[100, 250]$	10	1683.5	1062.9	8	3083.6^{8}	1126.9	2	0.3	3600^{10}	0	0	3.6
$U[200, 500]$	10	1882.1^{1}	1178.1	5	2059.4^{5}	1618.2	5	0.6	3600^{10}	0	0	2.9

5.1.2. *Fifteen Customer Problems*

Tables 3 and 4 report the results of our computational testing for $n = 15$ customers with $m = 5$ and $m = 10$ facilities. Once again, both the branch-and-price algorithm and the SLA approach outperform BARON for these test problems. For $m = 5$, although the SLA approach does outperform the branch-and-price algorithm for a number of instances, the results indicate that the branch-and-price algorithm has a better average running time and is more robust than the SLA approach. For $m = 10$, the branch-and-price algorithm significantly outperforms the SLA approach in both average running time and the number of instances in which it has the fastest running time. The results indicate the branch-and-price algorithm actually solves problems with $m = 10$ more effectively than problems with $m = 5$. We believe that this can be attributed to the fact that there is more opportunity for facilities with low operational costs to be located near customers and, therefore, the customer assignment decisions become clearer.

5.1.3. *Larger Scale Problems*

We next present the results of applying the branch-and-price algorithm and the SLA approach to problems of increased size. We will focus on testing the boundaries of each of these approaches, meaning that we set a high time limit to see if these approaches can provide an optimal solution (or a solution of high-quality) within this limit. The main reason that we do not examine the global optimization software package BARON is that we have already reached the limits of this package for problems with $n = 10$ and $n = 15$. We note that since (P) is a strategic problem used in supply chain planning, it is not unreasonable to run either of these approaches for a few hours in order to determine an optimal solution. Therefore, we set the time limit for these tests to 8 hours (480 minutes). For these results, we present the number of minutes required to solve the problem, as opposed to the number of seconds presented in Sections 5.1.1 and 5.1.2. Table 5 presents the computational results for an instance of each problem from classes with $n = 20$ and $T = 10$. Table 6 reports the computational results for instances with $n = 30$ and $T = 10$. The SLA approach either failed to find the optimal solution to (P) or CPLEX 11.0 ran out of memory while solving one of the linear approximations of the problem for every instance considered with $n = 20$ and $n = 30$. We note that the instances where the SLA approach ran out of memory are denoted by a "*" and we report the current best known bounds at the time of termination of the approach. Further, in calculating the error for the branch-and-price algorithm for the instances it did not solve to optimality, we compared it with the best known lower bound given by the SLA approach. The branch-and-price algorithm outperforms (in computational time and solution quality) the SLA

Table 4. Computational results for $n = 15$ and $m = 10$ for the branch-and-price algorithm, the SLA approach, and the software package BARON.

		B&P algorithm			SLA				BARON			
f_i	T	Time (s)	S.D.	Wins	Time (s)	S.D.	Wins	% Error	Time (s)	S.D.	Wins	% Error
0	5	3.6	1.7	9	118.7	184.6	1	0	3600^{10}	0	0	7.8
$U[100, 250]$	5	11.5	4.4	8	66.1	94.9	2	0	3600^{10}	0	0	1.7
$U[200, 500]$	5	26.7	17.2	4	23.4	24.0	6	0	3600^{10}	0	0	1.9
0	10	42.0	25.7	10	3002.8^{8}	1261.0	0	3.2	3600^{10}	0	0	10.0
$U[100, 250]$	10	613.1	320.6	10	3600^{10}	0	0	3.1	3600^{10}	0	0	4.2
$U[200, 500]$	10	1489.3	1683.2	10	3600^{10}	0	0	2.5	3600^{10}	0	0	2.6

Table 5. Computational results for $n = 20$ and $T = 10$ for the branch-and-price algorithm and the SLA approach.

		B&P algorithm		SLA	
f_i	m	Time (min)	% Error	Time (min)	% Error
0	5	42.5	0	480[1]	0.003
$U[100, 250]$	5	333.9	0	480[1]	2.2
$U[200, 500]$	5	367.4	0	480[1]	1.1
0	10	1.2	0	391.4*	1.8
$U[100, 250]$	10	33.1	0	480[1]	2.2
$U[200, 500]$	10	55.1	0	400*	1.2

approach for all the instances considered. We again see that the branch-and-price algorithm performs better on instances with a lower customer-to-facility ratio: it solves all problems within the time limit for $n = 30$ and $m = 15$, but only one instance within the time limit for $n = 30$ and $m = 10$, and no instances for $n = 30$ and $m = 5$. Based on these observations, we have tested our algorithm on instances with $n = 40$, $m = 20$, and $T = 10$ in Table 7. The results in Tables 5–7 indicate that our algorithm is especially suitable for large-scale problems with a low customer-to-facility ratio that focus simply on the assignment decisions since it solves the instances of the problem with $(f_i = 0, n = 20, m = 10)$, $(f_i = 0, n = 30, m = 15)$, and $(f_i = 0, n = 40, m = 20)$ in under an hour.

5.2. The Value of Integration

In this section, we characterize the "value" of integrating the facility location and production planning decisions for each of the problem classes considered in Section 5.1. In other words, we determine the percentage of additional costs incurred by undertaking the location and assignment decisions separately from the production planning decisions, i.e., we consider them sequentially. This means that we determine the location and assignment decisions based on solving a facility location problem with m facilities and nT customers and then determine the production and inventory levels at each facility by solving a production planning problem based on the optimal assignments from the facility location problem. If we let v^* denote the total cost (location, assignment, and production planning costs) incurred in an optimal solution for (P) and let $v(S.A.)$ denote the total cost obtained by the sequential approach to the problem, then the value of the integration is defined as

$$\text{Value of Integration} = 100 \times \frac{v(S.A.) - v^*}{v^*}.$$

Table 6. Computational results for $n = 30$ and $T = 10$ for the branch-and-price algorithm and the SLA approach.

		B&P algorithm		SLA	
f_i	m	Time (min)	% Error	Time (min)	% Error
0	5	480[1]	.8	480[1]	1.9
$U[100, 250]$	5	480[1]	0.7	480[1]	1.5
$U[200, 500]$	5	480[1]	0.8	480[1]	1.4
0	10	14.1	0	307.1*	4.7
$U[100, 250]$	10	480[1]	4.0	396.6*	5.8
$U[200, 500]$	10	480[1]	3.5	480[1]	13.6
0	15	10.2	0	480[1]	8.6
$U[100, 250]$	15	291.2	0	386.6*	8.9
$U[200, 500]$	15	448.5	0	480[1]	12.0

Table 7. Computational results for $n = 40$ and $T = 10$ for the branch-and-price algorithm and the SLA approach.

		B&P algorithm		SLA	
f_i	m	Time (min)	% Error	Time (min)	% Error
0	20	54.2	0	248.2*	26.7
$U[100, 250]$	20	480[1]	21.2	380*	28.2
$U[200, 500]$	20	480[1]	22.5	480[1]	32.6

To "normalize" the costs of a particular instance of (P), we have removed the base connection costs from the problem (which does not change the optimal solution to (P) or the facility location problem). In particular, for each customer and time period, we determined the minimum cost associated with transporting the demand of the customer/time period over all facilities and subtracted this minimum cost from the model.

Note that there is flexibility in determining the assignment costs of customer j in time period t to facility i in the traditional facility location problem during the sequential planning process. We set this assignment cost equal to $\widehat{ap}_{it} + c_{ijt}$ where $\widehat{ap}_{it}$ is an approximation of the production/inventory costs of a unit of demand assigned to facility i in time period t and c_{ijt} is the connection cost per unit demand of customer j to facility i in time period t. We examine three different classes of $\widehat{ap}_{it}$: (i) the FL class: $\widehat{ap}_{it} = 0$ meaning that the facility location problem disregards the production/inventory costs, (ii) the 20% class: $\widehat{ap}_{it} = P_{it}(0.2d_t)/0.2d_t$, where d_t is the cumulative demand of the customers in time period t, which represents the average cost per unit demand of producing 20% of d_t in period t at facility i, and (iii) the 40% class: $\widehat{ap}_{it} = P_{it}(0.4d_t)/0.4d_t$ which represents the average cost per unit demand of producing 40% of d_t in period t at facility i. Note that (ii) and (iii) essentially assume that we produce all demand assigned to facility i in time period t with production in time period t. Note that (ii) would be better suited than (iii) for problems where we expect to have a larger number of facilities operational.

Table 8 reports the average value of integration for each of the different classes of approximation production costs

Table 8. The value of integrating location and production planning problems with $n = 10$ and $n = 15$ customers.

		$n = 10, m = 5$			$n = 10, m = 10$			$n = 15, m = 5$			$n = 15, m = 10$		
f_i	T	FL	20%	40%	FL	20%	40%	FL	20%	40%	FL	20%	40%
0	5	18.8	11.2	14.3	22.0	12.2	14.5	16.8	10.5	12.2	21.1	13.9	15.8
$U[100, 250]$	5	3.8	4.2	3.9	5.0	4.3	4.3	5.3	6.3	6.8	5.5	6.7	6.8
$U[200, 500]$	5	1.7	2.0	1.8	3.6	2.2	2.1	2.9	2.3	3.2	6.2	2.8	3.4
0	10	12.5	8.0	8.1	21.7	11.3	13.5	9.4	6.8	7.3	19.0	11.1	12.1
$U[100, 250]$	10	5.2	3.4	3.9	9.2	5.6	5.8	6.6	5.3	5.4	7.6	9.3	7.9
$U[200, 500]$	10	3.6	3.6	3.4	4.5	6.3	5.2	4.2	7.2	6.8	8.0	6.1	6.1

over the 10 instances considered in each problem class from Sections 5.1.1 and 5.1.2. Table 9 reports the value of integration for the instances considered in Section 5.1.3, where the value of integration is calculated based on the best known current solution to the problem since we did not solve the problem to optimality for certain instances with $n = 30$ and $n = 40$. This means the reported values are lower bounds on the actual value of integration. It is clear from these results that there are benefits (significant, in some cases) from integrating facility location and production planning decisions for all problem classes, even when the location and assignment decisions are made by incorporating an approximation of the production/inventory costs into the facility location portion of the sequential decision-making process. Further, the results indicate that we can obtain a better understanding of the costs involved with serving our customer base by integrating the facility location and production planning decisions.

5.3. Sensitivity of the Model to Forecasting Errors

This section examines the impact of customer demand forecast errors on the key strategic decisions, i.e., the location decisions and optimal total costs in our model. Understanding these impacts is important since we may not know the customer demand vectors with certainty at the time that we must make strategic facility location decisions. Therefore, we are interested in examining whether the optimal locations obtained when using demands containing forecast errors are the same as (or similar to) the optimal locations with the actual demand vector. Our model (P) will be quite useful in supply chain planning activities if it can provide optimal strategic decisions that are robust to changes in demands in addition to an accurate estimate of the long-range costs.

We examine the same 10 problems considered in each class of problems from Sections 5.1.1 and 5.1.2 that used positive facility opening costs. For each of these instances, we examined 10 “forecasted instances”: five forecasted instances with “low” errors and five forecasted instances with “high” errors. The procedure to create a forecasted instance with a low error involved generating T errors uniformly on the interval $[-2, 2]$ for customer j. We then sorted these errors from the smallest to largest according to their absolute values. This provides us with our error vector for customer j. The reason that we sorted the errors was that we can expect that the forecasting errors will increase as we go further out in the horizon. We then added the error vector of customer j to the actual demand vector of customer j to get the forecasted demand vector for customer j. The procedure was identical for generating a forecasted instance with a high error, except that the errors were generated uniformly on the interval $[-4, 4]$. We note that the magnitude of the absolute expected error is 6.6% for the low errors and 13.3% for the high errors. For each actual instance and error class, we are interested in: (i) the average absolute error of the cost of the forecasted instances (taking into account the expected “base cost” as we did in Section 5.2) and (ii) the number of forecasted instances that provide a different location vector than optimal location vector for the problem with the actual demand.

Tables 10 and 11 report the results on the sensitivity of the model to forecasting errors. We present the absolute error of the cost of the forecasted instances across all actual instances (the “Abs. % Error” column), the number of actual instances where all forecasted instances provide the same

Table 9. The value of integrating location and production planning problems with $n = 20$, $n = 30$, and $n = 40$ customers.

		$n = 20, m = 5$			$n = 20, m = 10$			$n = 30, m = 5$			$n = 30, m = 10$			$n = 30, m = 15$			$n = 40, m = 20$		
f_i	T	FL	20%	40%	FL	20%	40%	FL	20%	40%	FL	20%	40%	FL	20%	40%	FL	20%	40%
0	10	11.7	9.8	9.7	19.4	11.9	15.5	21.1	18.7	19.5	19.9	10.9	11.7	34.7	16.2	19.6	27.7	17.6	22.0
$U[100, 250]$	10	4.4	1.5	1.2	12.8	1.1	8.4	7.9	7.4	7.4	10.5	10.5	11.7	17.5	14.0	14.3	13.0	9.6	12.3
$U[200, 500]$	10	8.2	6.2	6.3	0.0	0.1	0.1	2.5	1.5	1.4	6.4	5.9	5.7	7.8	10.1	5.8	7.8	7.2	7.5

Table 10. Computational results for $n = 10$ on the sensitivity of the model to forecasting errors.

			Low variability				High variability			
f_i	T	m	Abs. % error	# All cor.	Avg. # diff.	Max. # diff.	Abs. % error	# All cor.	Avg. # diff.	Max. # diff.
$U[100, 250]$	5	5	1.3	7	0.7	3	3.2	7	0.7	3
$U[200, 500]$	5	5	2.0	7	0.3	1	4.7	7	0.5	3
$U[100, 250]$	10	5	1.0	9	0.1	1	4.0	9	0.2	2
$U[200, 500]$	10	5	2.2	10	0	0	4.3	9	0.1	1
$U[100, 250]$	5	10	1.3	9	0.3	3	4.3	7	0.3	1
$U[200, 500]$	5	10	1.4	8	0.3	2	3.1	5	0.7	2
$U[100, 250]$	10	10	1.5	8	0.2	1	4.1	4	0.9	3
$U[200, 500]$	10	10	1.7	10	0	0	2.3	8	0.2	1

location vector as the optimal location vector (the "# All Cor." column), the average number of forecasted instances that provide a different location vector (the "Avg. # Diff." column), and the maximum number of forecasted instances (across all actual instances) that provide a different location vector (the "Max. # Diff." column). These results indicate that for low forecast errors, applying (P) to the forecasted demand streams of the customers provides a very accurate estimate of the costs (between 0.5% and 2.7% in absolute error) associated with the problem with the actual demand stream. Further, for most instances in each problem class, the strategic location decisions were the same in all of the forecasted problems with low variability. For high forecast errors, applying (P) to the forecasted demand streams does provide a less accurate estimate of the costs (between 2.2% and 5.5% in absolute error). However, for most instances in each problem class, the strategic location decisions were the same in all of the forecasted problems with high variability. Therefore, (P) can be quite useful in planning activities, even if there are errors in the forecasts associated with the customer demand streams.

6. EXTENSIONS OF THE GENERAL PROBLEM

In this section, we will discuss several extensions of our integrated facility location and production planning problem, (P), and how to modify the branch-and-price algorithm to solve them. The key factor in this modification lies in effectively solving the pricing problem, which will be a generalization of the (PPOSP). We will actually show that for several of these generalizations, we can solve the problem by using a similar dynamic program as the one for (PPOSP) without increasing the worst-case complexity of $O(nT^3 \log nT)$ time. In Section 6.1, we first consider the situation in which the facility opening/closing decisions are dynamic, i.e., we can open a new facility at any point over the horizon and close it throughout the horizon. We then discuss a model in which the demand patterns of the customers are cyclic (see Section 6.2), i.e., the patterns repeat over time.

6.1. The Problem with Dynamic Openings and Closings

Our formulation of the integrated facility location and production planning problem, (P), assumes that it is necessary for all open facilities to be acquired (or built) prior to the first time period and that all facilities will be operational for the entire problem horizon. In certain situations, it may not be necessary (or possible) for all facilities to be acquired prior to the first time period, or it may be possible to sell off (or close) a facility prior to the end of the horizon. In other words, we may need to model the possibility that facility i can be

Table 11. Computational results for $n = 15$ on the sensitivity of the model to forecasting errors.

			Low variability				High variability			
f_i	T	m	Abs. % error	# All cor.	Avg. # diff.	Max. # diff.	Abs. % error	# All cor.	Avg. # diff.	Max. # diff.
$U[100, 250]$	5	5	1.3	9	0.2	2	5.4	7	0.5	3
$U[200, 500]$	5	5	2.7	10	0	0	5.5	10	0	0
$U[100, 250]$	10	5	2.2	10	0	0	4.4	8	0.3	2
$U[200, 500]$	10	5	2.2	10	0	0	3.6	7	0.6	2
$U[100, 250]$	5	10	1.7	9	0.2	2	4.6	7	0.5	3
$U[200, 500]$	5	10	1.9	9	0.1	1	2.2	7	0.3	1
$U[100, 250]$	10	10	1.3	8	0.3	2	2.7	6	0.5	2
$U[200, 500]$	10	10	0.5	10	0	0	2.5	8	0.3	2

acquired/opened in any time period $t = 1, \ldots, T$ and the possibility that facility i can be closed in any time period. A salvage value might be obtained from selling the facility, and we can no longer operate the facility (i.e., we cannot sell off a facility and reacquire it in a later time period). To modify our integrated facility location and production planning problem (P) to model the situation with dynamic openings and closings (P-DOC), we introduce binary variables y_{it} and w_{it} for $i = 1, \ldots, m$ and $t = 1, \ldots, T$ that represent the decision of opening facility i in time period t and closing facility i in time period t, respectively. We denote the facility opening cost and salvage value of facility i in time period t as f_{it} and s_{it}. To formulate the (P-DOC) of the problem with dynamic openings and closings, we replace the facility opening costs in the objective function of (P) with

$$\sum_{i=1}^{m} \sum_{t=1}^{T} (f_{it} y_{it} - s_{it} w_{it}).$$

In the constraints of (P-DOC) we need to ensure that (i) we only assign the demand of a customer in a time period to a facility which is currently operational (i.e., it has been opened but not yet closed) in that time period, (ii) we open each facility at most once, (iii) we only close a facility in a time period after which it is opened, and (iv) we close a facility at most once. To ensure that (i) occurs, we replace the constraints (2) of (P) with the constraint

$$x_{ijt} \leq \sum_{s=1}^{t} y_{is} - \sum_{s=1}^{t} w_{is} \text{ for } i = 1, \ldots, m, \quad j = 1, \ldots, n, t = 1, \ldots, T. \tag{12}$$

If we open a facility i in time period t' and close it on or before time period t, then the right side of this equation will be equal to 0, so that we cannot assign a customer to facility i in time period t, so that (i) is satisfied. The binary constraints on x_{ijt} and constraints (12) ensure that (iii) holds. It is not difficult to ensure that (ii) and (iv) hold by including the constraints

$$\sum_{t=1}^{T} y_{it} \leq 1 \text{ for } i = 1, \ldots, m \tag{13}$$

$$\sum_{t=1}^{T} w_{it} \leq 1 \text{ for } i = 1, \ldots, m. \tag{14}$$

To effectively solve the (P-DOC), we will need to modify the branch-and-price algorithm to solve (P) from Section 3. Note that, again, the definition of the cost of assigning the ℓth column to facility i will change. In particular, we let t_ℓ and t'_ℓ be the first and last time period in which the demand of any customer is selected in the ℓth column associated with facility i. We then define the cost of assigning the ℓth column to facility i as

$$\zeta_i\left(\alpha^{\ell}_{i\cdot\cdot}\right) = \min_{t=1,\ldots,t_\ell} f_{it} + \min_{t=t'_\ell,\ldots,T} (-s_{it}) + \sum_{j=1}^{n} \sum_{t=1}^{T} d_{jt} c_{ij} \alpha^{\ell}_{ijt} + g_i\left(\alpha^{\ell}_{i\cdot\cdot}\right).$$

This definition of the cost ensures that our facility opening cost is as cost-effective as possible given the set of customers/time periods associated with the ℓth column, and that we recover the largest possible salvage value for closing the facility when it no longer needs to be operational. For the column generation approach to solve (SCR), we need to account for the flexibility of opening/closing the facility in the pricing problem. This means that we end up with an extension of the (PPOSP) in which we must determine the time period to open a facility and (potentially) determine the time period to close the facility. We refer to this extension of the (PPOSP) as the (PPOSP-DOC). We can extend the dynamic programming algorithm for (PPOSP) to solve this variant of the (PPOSP-DOC). In particular, we simply disregard (for now) the facility opening/closing decisions in the DP algorithm. We then solve for the maximum profit for each of the subplans, $p(t, t')$ for all $t = 1, \ldots, T$, $t' = t, \ldots, T + 1$. In the (PPOSP-DOC), it is possible that we open the facility in any time period τ and close the facility in any time period τ' (where $\tau \leq \tau'$). Therefore, for any two time periods τ and τ', we will need to determine the maximum amount of profit obtained by only selecting customers with demand between τ and τ'. Therefore, we define the function $C(\tau, \tau')$ to be the maximum amount of profit obtained by only selecting demands in time periods $\tau, \ldots, \tau'$. It is not difficult to see that we can determine the values $C(\tau, \tau')$ for $\tau = 1, \ldots, T$, $\tau' = \tau, \ldots, T$ by applying a forward recurrence relationship similar to (11) from the starting point τ. This would require $O(nT^3 \log nT)$ time to determine $p(t, t')$ and $O(T^2)$ time to apply each of the recurrence relations from the starting points τ for $\tau = 1, \ldots, T$. The optimal solution value to the (PPOSP-DOC) would then be equal to

$$\max_{\tau=1,\ldots,T} \max_{\tau'=\tau,\ldots,T} \left\{ C(\tau, \tau') - \min_{s=1,\ldots,\tau} f_s + \max_{s=\tau',\ldots,T} w_s \right\}.$$

The time required to determine this value would include determining $p(t, t')$ for $t = 1, \ldots, T$, $t' = t, \ldots, T$ ($O(nT^3 \log nt)$ time), determining $C(\tau, \tau')$ for each starting point τ ($O(T^3)$ time), and determining the maximum value given by the equation ($O(T^2)$ time). This means that the bottleneck operation of this algorithm to solve the (PPOSP-DOC) would be calculating $p(t, t')$ and, therefore, we can solve the (PPOSP-DOC) in $O(nT^3 \log nT)$ time.

We have currently assumed in the (P-DOC) (and the (PPOSP-DOC)) that the decisions related to the opening and closing of a facility have no effect on the production and inventory costs associated with the facility. It may be the case that the facility requires a "ramp-up time" and a "ramp-down time" to open or close. If the facility cannot operate during these times, then the only aspect of the ramp-up/ramp-down time that affects our decision making is the first/last time period in which the facility can operate and thus the previously discussed model and algorithm for the pricing problem can solve this problem. However, it may be the case that we can operate the facility during the ramp-up and ramp-down times, but the production and inventory costs are affected. To solve the pricing problem that arises in this variant of the (P-DOC), we can solve $O(T^2)$ problems of the form (PPOSP). In particular, we examine the problem (PPOSP) that corresponds to starting the ramp-up time (i.e., opening the facility) in time period t and ending the ramp-down time (i.e., closing the facility) in time period t', where $t < t'$. Since we know the opening time period and closing time period of the facility, we will know the production and inventory costs associated with each time period, so we can solve the (PPOSP) for these opening/closing time periods. For each pair of values (t, t'), we then incorporate the opening cost and the salvage value of the facility and then take the maximum over all pairs to determine the optimal solution to this variant of the (PPOSP-DOC).

6.2. The Cyclic Problem

It may be possible that the demand for each customer and the production and inventory costs at each facility are in an equilibrium situation, meaning that the demand and the costs are stationary with a cycle length T. In other words, we have that $d_{j,T+1} = d_{j1}, d_{j,T+2} = d_{j2}, \dots$, for each customer j, $P_{i,T+1}(p) = P_{i1}(p), P_{i,T+2}(p) = P_{i2}(p), \dots$, for each facility i and production level p, and $h_{i,T+1} = h_{i1}, h_{i,T+2} = h_{i2}, \dots$, for each facility i. In other words, the production planning problems faced by each facility are cyclic in nature. The MPSSP has been considered when the demand patterns are cyclic, see, for example, Romeijn and Romero Morales [16, 17].

This cost and demand structure imply that the inventory levels at each of the facilities will be cyclic as well. In other words, rather than require that the initial inventory level I_{i0} is equal to zero for each facility, we assume that the starting inventory level of a facility during a cycle must be equal to the ending inventory level of the facility, i.e., constraints (6) of (P) are replaced by

$$I_{i0} = I_{iT} \text{ for } i = 1, \dots, m. \tag{15}$$

This is the only modification to the model (P) that is necessary to solve the integrated facility location and production planning problem with cyclic demands (P-CD). We note, however, that the interpretation of the facility opening costs, f_i, needs to be slightly modified in (P-CD) since we will not be "opening" the facilities in every cycle of the problem. However, we may think of the facility opening costs in (P-CD) as a composition of (i) the fixed operating costs of that facility during a cycle (such as the salaries of workers and/or maintenance costs of the equipment in the facility) and (ii) an amortized portion of the initial long-term investment required to obtain the facility. If we use this interpretation of the facility opening costs f_i, we see that we are minimizing all relevant costs incurred during a cycle to meet the demand of the customers in our problem. To solve (P-CD) with the branch-and-price algorithm from Section 3, we (again) need to redefine the cost function associated with the ℓth column of facility i. In particular, we interpret the production planning cost function $g_i(\alpha^{\ell}_{i\cdot\cdot})$ as the optimal production and inventory costs in the cyclic production planning problem to meet the demand of the customer/time periods in the ℓth subset. This leads to another variant of the (PPOSP), which we call the production planning and order selection problem with cyclic demand (PPOSP-CD), where the initial inventory must equal the inventory in the last time period.

In the traditional setting (i.e., demand in each time period is fixed and given) with cyclic demand, it is not difficult to show that there exists a time period in which we carry no inventory in the optimal production plan; however, this period is not necessarily the Tth time period as is the case with normal production planning problems. In other words, there exists a period which we can refer to as the "starting period" in our problem; however, this starting period does not necessarily have to be period 1. As in the traditional problem, it is not difficult to show that there exists an optimal solution to (PPOSP-CD) where a time period exists in which we carry no inventory. Using this property, we can modify the dynamic programming algorithm for (PPOSP) to solve (PPOSP-CD). In particular, we will calculate the maximum profit obtained in a subplan, $p(t, t')$, for all $t = 1, \dots, T$ and $t' = 1, \dots, T$, since we may produce a unit of demand in time period t and hold it in inventory in time periods $t, t + 1, \dots, T, 1, \dots, t'$. This, again, requires $O(nT^3 \log nT)$ time. We then apply the dynamic programming recurrence relationship from each possible "starting" time period $t = 1, \dots, T$. In other words, define the function $V(t)$ to be the maximum amount of profit in (PPOSP-CD) when we consider time period t as the starting period for the problem. We can calculate $V(t)$ through either the forward or backward recurrence relationship by considering time period t as time period 1. We can solve for $V(t)$ in $O(T^2)$ time for each $t = 1, \dots, T$, so we require $O(T^3)$ time to calculate $V(t)$ for all $t = 1, \dots, T$. The optimal solution value for (PPOSP-CD) is then equal to $\max_{t=1,\dots,T} V(t)$. The bottleneck operation is again the calculation of all $p(t, t')$ and, therefore, (PPOSP-CD) can be solved in $O(nT^3 \log nT)$ time.

7. CONCLUSIONS

In this article, we considered a problem that integrates three important supply chain management decisions: the location of supply facilities, the assignment of customers to these facilities, and the management of the production and inventory levels within the facilities. We formulated this problem and developed a branch-and-price algorithm for its set-covering formulation. It turns out that the pricing problem that arises in this algorithm is an important supply chain planning problem in its own right. In particular, this problem is a production planning and order selection problem in which the supplier is given a set of potential orders and must select a subset of these orders to maximize its profit, i.e., the revenues received from selected orders minus the production planning costs incurred in meeting the demand of the selected customers. Computational testing of our proposed algorithm indicates its advantages over a commercial global optimization software package. These tests also characterized the value of integrating facility location and production planning decisions and indicated that there are significant potential benefits from simultaneously considering these decisions within an integrated model. Further, our computational testing indicates that our model is rather insensitive to errors in the forecasted demand streams and can, therefore, be valuable in supply chain planning activities. We also considered several extensions of the integrated location and production planning problem, as well as the resulting extensions of the production planning and order selection pricing problem that arise in applying the branch-and-price algorithm. We provided polynomial-time dynamic programming algorithms for each of the extensions of the pricing problem.

ACKNOWLEDGMENTS

We thank Chase Rainwater for his help in coding a branch-and-price algorithm for general customer assignment models, which was then specialized to our integrated problem, and for his discussions about several issues that arise in the computational testing of branch-and-price algorithms for customer assignment models. We thank the anonymous referees and associate editor whose suggestions have greatly strengthened the material presented in the article. This work was supported in part by a National Science Foundation Graduate Research Fellowship, the National Science Foundation under grant nos. CMMI-0927930, DMI-0355533, CMMI-0926508, CMMI-0621433, and CMMI-0727640 and by the NSFC under grant no. 70731003 and 90924012.

APPENDIX

PROOF OF THEOREM 3.1: To prove this result, we will focus on comparing v(PR) and v(SCR) with the optimal solution value to the following problem:

$$\text{minimize } \sum_{i=1}^{m}\left(f_i y_i + g_i\left(\sum_{j=1}^{n} d_{j1}x_{ij1},\ldots,\sum_{j=1}^{n} d_{jT}x_{ijT}\right)\right) + \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}x_{ijt}$$

subject to (P′R)

$$\sum_{i=1}^{m} x_{ijt} \geq 1 \quad \text{for } j = 1,\ldots,n,\ t = 1,\ldots,T \tag{16}$$

$$x_{ijt} \leq y_i \quad \text{for } i = 1,\ldots,m,\ j = 1,\ldots,n,\ t = 1,\ldots,T \tag{17}$$

$$x_{ijt} \geq 0 \quad \text{for } i = 1,\ldots,m,\ j = 1,\ldots,n,\ t = 1,\ldots,T \tag{18}$$

$$y_i \geq 0 \quad \text{for } i = 1,\ldots,m \tag{19}$$

Our first result shows that this problem is equivalent to (PR).

LEMMA A.1: $v(\text{PR}) = v(\text{P}'\text{R})$.

PROOF: Given an optimal solution, (x^*, y^*, p^*, I^*), to (PR), we will examine the cost of the solution (x^*, y^*) to (P′R). The cost of this solution in (P′R) must equal the cost of the solution (x^*, y^*, p^*, I^*) to (PR) or we contradict (i) the optimality of (x^*, y^*, p^*, I^*) to (PR) or (ii) the definition of the functions $g_i(\sum_{j=1}^{n} d_{j1}x_{ij1},\ldots,\sum_{j=1}^{n} d_{jT}x_{ijT})$. Given an optimal solution $(\bar{x}, \bar{y})$ to (P′R), we examine the solution $(\bar{x}, \bar{y}, \bar{P}, \bar{I})$ where $\bar{P}_{it}$ and $\bar{I}_{it}$ are the optimal production and inventory levels associated with the optimal solution to $g_i(\sum_{j=1}^{n} d_{j1}\bar{x}_{ij1},\ldots,\sum_{j=1}^{n} d_{jT}\bar{x}_{ijT})$. By definition, $(\bar{x}, \bar{y}, \bar{P}, \bar{I})$ is feasible to (PR) and has the same objective function value as the optimal solution to (P′R). Therefore, the optimal solution to (PR) has a solution in (P′R) with the same objective function value and the optimal solution to (P′R) has a solution in (PR) with the same objective function value. This implies our desired result. □

We will prove that any solution to (P′R) has an equivalent solution in (SCR) with an objective function that is less than or equal to the objective function in (P′R). The fact that the functions $g_i(D_1,\ldots,D_T)$ are concave will help relate the objective functions of these solutions.

LEMMA A.2: The functions $g_i(D_1,\ldots,D_T)$ for $i = 1,\ldots,m$ are concave.

PROOF: This can be proven in a similar manner as Lemma 3.1 in Romeijn et al. [19]. □

We now present the main result of this section.

THEOREM 3.1: $v(\text{PR}) \geq v(\text{SCR})$.

PROOF: Consider any solution, (x, y), to (P′R). We construct a solution to (SCR) with an objective function that is less than or equal to the objective function value of the solution in (P′R). We initially set $y_i^\ell = 0$ for all $i = 1,\ldots,m$ and $\ell = 1,\ldots,L$. For facility i, we set the values of y_i^ℓ that will be positive in the solution to (SCR) recursively. At any point in this recursive procedure, we let the value $\bar{x}_{ijt}$ denote the remaining amount of the assignment decision x_{ijt} in (P′R) that has yet to be assigned in the solution y to (SCR). Initially, we set $\bar{x}_{ijt} = x_{ijt}$ for all $j = 1,\ldots,n$ and $t = 1,\ldots,t$. An iteration of the recursive procedure for facility i proceeds as follows: we have a set of customers and time periods, denoted by ℓ, with strictly positive values of $\bar{x}_{ijt}$. We then set $y_i^\ell = \min_{(j,t)\in\ell}\bar{x}_{ijt}$ and $\bar{x}_{ijt} = \bar{x}_{ijt} - y_i^\ell$. At least one $(j,t) \in \ell$ will now have $\bar{x}_{ijt} = 0$. We then continue this procedure until all $\bar{x}_{ijt} = 0$. It is not difficult to show (through an inductive argument) that at any point in the recursive procedure we have $\sum_{\ell=1}^{L} \alpha_{ijt}^{\ell} y_i^{\ell} = x_{ijt} - \bar{x}_{ijt}$. This means that at the end of the recursive procedure that $x_{ijt} = \sum_{\ell=1}^{L} \alpha_{ijt}^{\ell} y_i^{\ell}$. Therefore,

$$\sum_{i=1}^{m}\sum_{\ell=1}^{L} \alpha_{ijt}^{\ell} y_i^{\ell} = \sum_{i=1}^{m} x_{ijt} \geq 1$$

so that Eq. (8) is satisfied in (SCR). We now set $y_i^1 = 1 - \sum_{\ell=2}^{L} y_i^{\ell}$ (where $\ell = 1$ represents the empty set) so that $\sum_{\ell=1}^{L} y_i^{\ell} = 1$. We further note that $\sum_{\ell=2}^{L} y_i^{\ell} = \max_j \max_t x_{ijt} \leq y_i$. We now show that objective function of the solution in (SCR) is less than or equal to the objective function of the solution in (P′R). Now observe that the solution value of y with respect to facility i in (SCR) is equal to

$$\begin{aligned}
\sum_{\ell=1}^{L} \zeta_i(\alpha_{i\cdot\cdot}^{\ell}) y_i^{\ell} &= \sum_{\ell=1}^{L} g_i\left(\sum_{j=1}^{n} d_{j1}\alpha_{ij1}^{\ell}, \dots, \sum_{j=1}^{n} d_{jT}\alpha_{ijT}^{\ell}\right) y_i^{\ell} \\
&\quad + \sum_{\ell=2}^{L}\left(f_i + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}\alpha_{ijt}^{\ell}\right) y_i^{\ell} \\
&= \sum_{\ell=1}^{L} g_i\left(\sum_{j=1}^{n} d_{j1}\alpha_{ij1}^{\ell}, \dots, \sum_{j=1}^{n} d_{jT}\alpha_{ijT}^{\ell}\right) y_i^{\ell} \\
&\quad + \sum_{\ell=2}^{L}\left(f_i y_i^{\ell} + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}\alpha_{ijt}^{\ell} y_i^{\ell}\right) \\
&= \sum_{\ell=1}^{L} g_i\left(\sum_{j=1}^{n} d_{j1}\alpha_{ij1}^{\ell}, \dots, \sum_{j=1}^{n} d_{jT}\alpha_{ijT}^{\ell}\right) y_i^{\ell} + f_i\left(\sum_{\ell=2}^{L} y_i^{\ell}\right) \\
&\quad + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}\left(\sum_{\ell=2}^{L} \alpha_{ijt}^{\ell} y_i^{\ell}\right) \\
&\leq \sum_{\ell=1}^{L} g_i\left(\sum_{j=1}^{n} d_{j1}\alpha_{ij1}^{\ell}, \dots, \sum_{j=1}^{n} d_{jT}\alpha_{ijT}^{\ell}\right) y_i^{\ell} \\
&\quad + f_i y_i + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}x_{ijt} \\
&\leq g_i\left(\sum_{\ell=1}^{L}\sum_{j=1}^{n} d_{j1}\alpha_{ij1}^{\ell} y_i^{\ell}, \dots, \sum_{\ell=1}^{L}\sum_{j=1}^{n} d_{jT}\alpha_{ijT}^{\ell} y_i^{\ell}\right) \\
&\quad + f_i y_i + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}x_{ijt} \\
&= g_i\left(\sum_{j=1}^{n} d_{j1}\left(\sum_{\ell=1}^{L} \alpha_{ij1}^{\ell} y_i^{\ell}\right), \dots, \sum_{j=1}^{n} d_{jT}\left(\sum_{\ell=1}^{L} \alpha_{ijT}^{\ell} y_i^{\ell}\right)\right) \\
&\quad + f_i y_i + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}x_{ijt} \\
&= g_i\left(\sum_{j=1}^{n} d_{j1}x_{ij1}, \dots, \sum_{j=1}^{n} d_{jT}x_{ijT}\right) \\
&\quad + f_i y_i + \sum_{j=1}^{n}\sum_{t=1}^{T} d_{jt}c_{ijt}x_{ijt},
\end{aligned}$$

where the last inequality holds since the functions $g_i(D_1, \dots, D_T)$ are concave and $\sum_{\ell=1}^{L} y_i^{\ell} = 1$. This implies that the optimal solution value to (P′R) is greater than or equal to the optimal solution value to (SCR). Therefore, by Lemma A.1, $v(\text{PR}) \geq v(\text{SCR})$. □

Mixed-Integer Formulation of Piecewise Linear Approximation of Production Cost Functions

We now describe the procedure to obtain a mixed-integer programming formulation of the problem where we approximate the production cost functions with a piecewise linear function with δ breakpoints. Let $\alpha_{1it}, \alpha_{2it}, \dots, \alpha_{\delta it}$ denote the breakpoints of our piecewise linear function approximation (call this function $\hat{P}_{it}$) of P_{it}. Further, let β_{kit} and c_{kit} for $k = 1, \dots, \delta$ denote y-intercept and the slope of the kth segment of $\hat{P}_{it}$. We introduce continuous variables $\hat{z}_{kit}$ and binary variables $\hat{y}_{kit}$ for each segment of $\hat{P}_{it}$. The following constraints are placed on these variables

$$p_{it} = \sum_{k=1}^{\delta} \hat{z}_{kit} \tag{20}$$

$$\alpha_{kit}\hat{y}_{kit} \leq \hat{z}_{kit} \leq \alpha_{(k+1)it}\hat{y}_{kit} \tag{21}$$

$$\sum_{k=1}^{\delta} \hat{y}_{kit} = 1. \tag{22}$$

Note that if $\hat{y}_{kit} = 1$, then the value p_{it} appears in the interval $[\alpha_{kit}, \alpha_{(k+1)it}]$. Constraints (20)–(22) ensure that p_{it} belongs to exactly one of these intervals. We also have the following relationship and therefore may replace $\hat{P}_{it}$ with a linear function of the variables in the objective function:

$$\hat{P}_{it}(p_{it}) = \sum_{k=1}^{\delta}(c_{kit}\hat{z}_{kit} + \beta_{kit}\hat{y}_{kit}).$$

We note that in obtaining this approximate formulation of (P), we always define the following two points as the first and last breakpoints of the function $\hat{P}_{it}$: the first breakpoint was always 0 and the last breakpoint was always $\sum_{j=1}^{n}\sum_{s=t}^{T} d_{jt}$, which is the tightest upper bound on production during a particular time period. Further, due to the concavity of the piecewise linear (approximate) production cost functions (observed by Balakrishnan and Graves [2]), it is not necessary to include the lower bound constraint in (21) or constraint (22). However, we have included these constraints into the mixed-integer programming formulation of this problem because preliminary computational testing suggested that these redundant constraints improve the computational performance of CPLEX 11.0 on the problem.

REFERENCES

[1] A. Balakrishnan and J. Geunes, Requirements planning with substitutions: Exploiting bill-of-materials flexibility in production planning, Manuf Service Oper Manage 2 (2000), 166–185.

[2] A. Balakrishnan and S. Graves, A composite algorithm for a concave-cost network flow problem, Networks 19 (1989), 175–202.

[3] C. Barnhart, E.L. Johnson, G.L. Nemhauser, M.W.P. Savelsbergh, and P.H. Vance, Branch-and-price: Column generation for solving huge integer programs, Oper Res 46 (1998), 316–329.

[4] P. Chardaire, A. Sutter, and M.C. Costa, Solving the dynamic facility location problem, Networks 28 (1996), 117–124.

[5] K.L. Croxton, B. Gendron, and T.L. Magnanti, A comparison of mixed-integer programming models for nonconvex piecewise linear cost minimization problems, Manage Sci 49 (2003), 1268–1273.

[6] M. Daskin, C. Coullard, and Z.-J. Shen, An inventory-location model: Formulation, solution algorithm, and computational results, Ann of Oper Res 10 (2001), 83–106.

[7] M.S. Daskin, Network and discrete location: Models, algorithms, and applications, Wiley, New York, 1995.

[8] E.V. Denardo, Dynamic programming: Models and applications, Prentice-Hall, Englewood Cliffs, New Jersey, 1982.

[9] R. Freling, H.E. Romeijn, D. Romero Morales, and A.P.M. Wagelmans, A branch-and-price algorithm for the multi-period single-sourcing problem, Oper Res 51 (2003), 922–939.

[10] J. Geunes, H.E. Romeijn, and K. Taaffe, Requirements planning with dynamic pricing and order selection flexibility, Oper Res 54 (2006), 394–401.

[11] P.C. Gilmore and R.E. Gomory, A linear programming approach to the cutting-stock problem, Oper Res 9 (1961), 849–859.

[12] W. Huang, H.E. Romeijn, and J. Geunes, The continuous-time single-sourcing problem with capacity expansion opportunities, Naval Res Logist 52 (2005), 193–211.

[13] J. Krarup and O. Bilde, "Plant location, set covering, and economic lot size: An $O(mn)$-algorithm for structured problems," L. Collatz and W. Wetterling (Editors), Numerische Methoden Bei Optimierungsaufgaben, Band 3: Optimierung Bei Graph-Theoretischen Und Ganzahligen Problemen, Birkhauser, Basel, 1977, pp. 155–180.

[14] S.F. Love, Bounded production and inventory models with piecewise concave costs, Manage Sci 20 (1973), 313–318.

[15] T. Magnanti, Z.-J. Shen, J. Shu, D. Simchi-Levi, and C.-P. Teo, Inventory placement in acyclic supply chain networks, Oper Res Lett 34 (2006), 228–238.

[16] H.E. Romeijn and D. Romero Morales, A probabilistic analysis of the multi-period single-sourcing problem, Discrete Appl Math 112 (2001), 301–328.

[17] H.E. Romeijn and D. Romero Morales, An asymptotically optimal greedy heuristic for the multi-period single-sourcing problem: The cyclic case, Naval Res Logist 50 (2003), 412–437.

[18] H.E. Romeijn and D. Romero Morales, Asymptotic analysis of a greedy heuristic for the multi-period single-sourcing problem: The acyclic case, J Heurist 10 (2004), 5–35.

[19] H.E. Romeijn, T.C. Sharkey, Z.J. Shen, and J. Zhang, Integrating facility location and production planning decisions, Networks 55 (2010), 78–89.

[20] T.J. Van Roy and D. Erlenkotter, A dual-based procedure for dynamic facility location, Manage Sci 28 (1982), 1091–1105.

[21] M.P.W. Savelsbergh, A branch–and–price algorithm for the generalized assignment problem, Oper Res 45 (1997), 831–841.

[22] J.F. Shapiro, Modeling the supply chain, Duxbury, Pacific Grove, CA, 2001.

[23] Z.J. Shen, C. Coullard, and M.S. Daskin, A joint location-inventory model, Transport Sci 37 (2003), 40–55.

[24] J. Shu, C.P. Teo, and Z.J. Shen, Stochastic transportation-inventory network design problem, Oper Res 53 (2005), 48–60.

[25] A.F. Veinott, Minimum concave cost solutions of Leontief substitution models of multi-facility inventory systems, Oper Res 17 (1969), 262–291.

[26] H.M. Wagner, A postscript to dynamic problems of the theory of the firm, Naval Res Logist Quart 7 (1960), 7–12.

[27] H.M. Wagner and T.M. Whitin, Dynamic version of the economic lot size model, Manage Sci 5 (1958), 89–96.

[28] G.O. Wesolowsky, Dynamic facility location, Manage Sci 19 (1973), 1241–1248.

[29] S.D. Wu and H. Golbasi, Multi-item, multi-facility supply chain planning: Models, complexities, and algorithms, Comput Optim Appl 28 (2004), 325–356.

[30] P.H. Zipkin, Foundations of inventory management (2000).