

Coordinating Multi-Attribute Procurement Auctions Subject to Finite Capacity Considerations

Jiong Sun and Norman M. Sadeh

April 2004
CMU-ISRI-03-105

e-Supply Chain Management Laboratory
Institute for Software Research International
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891, USA

Email Addresses: jiongs@andrew.cmu.edu, sadeh@cs.cmu.edu

Also appears as Computer Science Technical Report
CMU-CS-03-184

The research reported in this paper has been funded by the National Science Foundation under ITR Grant 0205435.

Keywords: Supply Chain Management, Procurement, Reverse Auction, Finite Capacity Scheduling

Abstract

Reverse auctions offer the prospect of more efficiently matching suppliers and producers in the face of changing market conditions. Prior research has ignored the temporal and finite capacity constraints under which reverse auctioneers typically operate. In this paper, we consider the problem faced by a manufacturer that can procure multiple key components from a number of possible suppliers through multi-attribute reverse auctions. Bids submitted by prospective suppliers include a price and a delivery date. The manufacturer has to select a combination of supplier bids that will maximize its overall profit, taking into account its own finite capacity and the prices and delivery dates offered by different suppliers for the same components. The manufacturer's profit is determined by the revenue generated by the products it sells, the costs of the components it purchases as well as late delivery penalties it incurs if it fails to deliver products in time to its own customers. We provide a formal model of this important class of problems, discuss its complexity and introduce rules that can be used to efficiently prune the resulting search space. We also introduce a branch-and-bound algorithm and an efficient heuristic search procedure. Computational results show that our heuristic procedure typically yields solutions that are only a few percent from the optimum. They also indicate that taking into account finite capacity considerations can significantly improve the manufacturer's bottom line.

1. Introduction

Today's global economy is characterized by fast changing market demands, short product lifecycles and increasing pressures to offer high degrees of customization, while keeping costs and lead times to a minimum. In this context, the competitiveness of both manufacturing and service companies will increasingly be tied to their ability to identify promising supply chain partners in response to changing market conditions. With the emergence of e-business standards, such as ebXML, SOAP, UDDI and WSDL, the Internet will over time facilitate the development of more flexible supply chain management practices.

Today, however such practices are confined to relatively simple scenarios such as those found in the context of MRO (Maintenance, Repair and Operations) procurement. The slow adoption of dynamic supply chain practices and the failure of many early electronic marketplaces can in part be attributed to the one-dimensional nature of early solutions that forced suppliers to compete solely on the basis of price. Research in the area has also generally ignored key temporal and capacity constraints under which reverse auctioneers typically operate. For instance, a PC manufacturer can only assemble so many PCs at once and not all PCs are due at the same time. Such considerations can be used to help the PC manufacturer select among bids from competing suppliers.

In this paper, we present techniques aimed at exploiting such temporal and capacity constraints to help a reverse auctioneer select among competing multi-attribute procurement bids that differ in prices and delivery dates. We refer to this problem as the *Finite Capacity Multi-Attribute Procurement* (FCMAP) problem. It is representative of a broad range of practical reverse auctions, whether in the manufacturing or service industry. This article provides a formal definition of the FCMAP problem, discusses its complexity and introduces several rules that can be used to prune its search space. It also presents a branch-and-bound algorithm, a simulated annealing procedure and an efficient pseudo-early/tardy heuristic search procedure that all take advantage of these pruning rules. Computational results show that accounting for the reverse auctioneer's finite capacity can significantly improve

its bottom line, confirming the important role played by finite capacity considerations in procurement problems. Results are also presented that compare the performance of our heuristics search procedures both in terms of solution quality and computational requirements under different bid profile assumptions. These results suggest that our pseudo-early/tardy procedure is generally capable of generating solutions that are just within a few percent of the optimum and that it scales nicely as problem size increases.

The balance of this paper is organized as follows. Section 2 provides a brief review of the literature. In section 3, we introduce a formal model of the FCMAP problem. Section 4 identifies three rules that can help the reverse auctioneer (or manufacturer) eliminate non-competitive bids or bid combinations. Section 5 introduces a branch-and-bound algorithm that takes advantage of our pruning rules. This is followed by the presentation of two heuristic search procedures that also take advantage of our pruning rules. In particular, Section 6 details a randomized pseudo-early/tardy heuristic that exploits a property of the FCMAP problem introduced in Section 4. In Section 7, a second heuristic search procedure is presented that combines Simulated Annealing (SA) search with a cost estimator based on the well-known “Apparent Tardy Cost” rule first introduced by Vepsalainen and Morton [25]. Section 8 also introduces a post-processing procedure that can further improve the quality of a solution. An extensive set of computational results are presented and discussed in Section 9. Section 10 provides some concluding remarks and discusses future extensions of this research.

2. Literature Overview

Surprisingly little research has been reported on coordinating procurement and finite capacity production planning. A notable exception is the work of Bassok and Akella [3] who explore a single-period, single-machine model that integrates production and raw material ordering decisions in a manufacturing facility with a single type of raw material and one or more finished products with stochastic demand. Raw material delivery is assumed to be stochastic: the manufacturer typically receives just a

fraction of what it orders. The authors focus on determining the quantities in which products are released into the system, taking into account the system's capacity and expectations about the fraction of ordered raw materials likely to arrive. The objective is to minimize the sum of backlog costs, production costs, ordering costs, as well as raw material and finished goods holding costs. Gurnani et al. [12] study a similar problem, where a manufacturer faces stochastic demand for a single finished product that requires two critical components. A computational study is used to demonstrate the benefits of coordinating procurement and production decisions over a more traditional approach that relies on non-coordinated procurement and production policies. In contrast to these earlier studies, the research we present in this paper integrates procurement and finite capacity production planning problem in environments where orders are placed for different types of finished products, each possibly requiring a different combination of components. An important aspect of solving this more general problem involves coordinating the procurement of the multiple components required by a given order. For instance, there is no point in paying a premium for having one component delivered early, if the other components can only be acquired much later. Another distinguishing feature of our work is the granularity at which we model demand, with each order having its own delivery date and its own marginal penalty for not meeting that date. By differentiating between different orders, their individual due dates, tardiness costs and component requirements, it becomes possible to develop solutions that capture the finer tradeoffs associated with their procurement requirements.

Another relevant body of literature revolves around research that has looked at procurement decisions subject to uncertain demand or supply conditions. For instance, Song et al. [23] have studied an assembly environment with a single finished product type requiring components, each procured from a pre-identified supplier and delivered subject to a random lead time. There is a one-time demand of random quantity with a known delivery date for the single finished product. The assembler has to decide how much to order of each component and at what time without knowing the demand quantity for the finished product. The objective is to minimize the sum of procurement costs, holding costs and backlog costs. In this research,

assembly time is assumed to be zero, thereby eliminating the need for finite capacity production planning. Gurnani et al. [11] consider a product with a single random demand and two components; there is an independent supplier for each component and a joint supplier that can supply the components in pairs. Components arrive in the current period with some probability or in the subsequent period otherwise. Once again, the assembly stage takes no time. The key decisions are how much to purchase from each supplier to minimize the sum of purchasing costs, component holding costs and backlog costs. Other related research includes the work of Gallien and Wein [10] who study the production of a product whose single component has a stochastic procurement lead time. Production is assumed to take no time (which amounts to having infinite capacity). Yano [26] considers the assembly of a product requiring two components subject to stochastic procurement and production lead times. Kumar [16], Hopp and Spearman [13] and Shore [22] consider an environment with multiple components, random procurement lead times and instantaneous production. Chu et al. [4] consider multiple components with random procurement lead times and a deterministic production lead time. In contrast to the above, the work we present in this paper models the finite capacity of the manufacturer/assembler and allows for environments with multiple finished products, each requiring a possibly different set of components. In addition, for each component, the manufacturer has to select from a number of bids from different sets of prospective suppliers with each bid possibly differing in price and delivery date. The present paper also assumes deterministic conditions.

A third line of relevant research has been concerned with multi-period finite capacity production planning models. For instance, Ciarallo et al. [5] study a multi-period aggregated production planning problem with a single-product single-stage manufacturer. The manufacturer faces a stationary random demand and has a stochastic capacity that varies from one period to the next. The manufacturer has to select the quantity of the product to release into the system to minimize the sum of holding and backlog costs. The authors show that the optimal finite horizon policy is of the “order-up-to-level” form. Jain and Silver [14] consider a single-period variant of the problem studied by Ciarallo et al. where the manufacturer can pay a premium

to purchase dedicated capacity from a supplier. Karmarkar and Lin [15] consider a multi-period production planning problem in which demand and yield is random. Zipkin [27] also considers a combined inventory and production problem for a facility facing stochastic demand. Production is assumed to be organized in large, discrete batches and is modeled as a queueing system, where production times can be product-specific and batch-size-dependent. Standard inventory and queueing submodels are combined into a classical optimization over batch sizes and safety stocks with respect to holding cost of finished goods plus penalty cost of backlogged orders under different possible control policies (e.g., First-Come First-Served). More recent work has also looked at models, where additional capacity can be purchased (e.g. outsourcing). For instance, Angelus and Porteus [1] consider a combined capacity and production planning problem for a make-to-stock product under stochastic demand. Rajagopalan and Swaminathan [19] report results for a combined capacity and production planning problem involving multiple finished products and demand growth. Van Mieghem [24] studies the coordination between subcontracting and production decisions in a two-stage, two-player stochastic game with uncertain demand. The above research results all rely on different combinations of aggregate demand, capacity and procurement models. In contrast, the work presented in this paper models demand, capacity and procurement requirements at a more detailed level, enabling for the development of solutions that exploit finer tradeoffs between different procurement and production options. This paper also focuses on deterministic scenarios.

3. The Finite Capacity Multi-Attribute Procurement Problem

The Finite Capacity Multi-Attribute Procurement (FCMAP) problem revolves around a reverse auctioneer – referred below as the “manufacturer”, though it could also be a service provider. The manufacturer has to satisfy a set of customer commitments or orders $O_i, i \in M = \{1, \dots, m\}$ (see Figure 1). Each order i needs to be completed by a

due date dd_i , and requires one or more components (or services), which the manufacturer can obtain from a number of possible suppliers. The manufacturer has to wait for all the components before it can start processing the order (e.g., waiting for all the components required to assemble a given PC). For the sake of simplicity, we assume that the processing required by the manufacturer to complete work on customer order O_i has a fixed duration du_i , and that the manufacturer can only process one order at a time (“capacity constraint”).

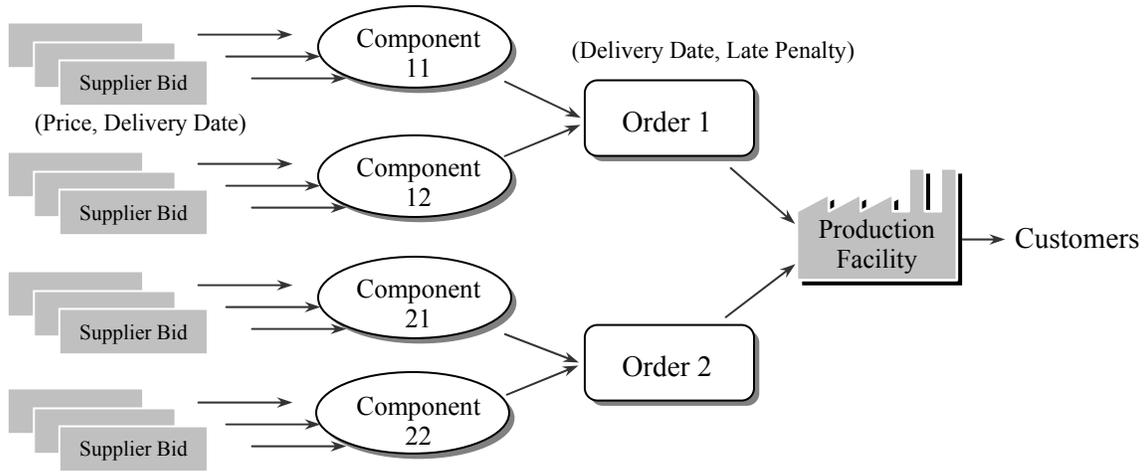


Figure 1. Finite capacity multi-attribute procurement problem

Formally, for each order O_i and each component $comp_{ij}$, $j \in N_i = \{1, \dots, n_i\}$, the manufacturer organizes a reverse auction for which it receives a set of multi-attribute bids $\beta_{i,j} = \{B_1^{ij}, \dots, B_{n_{ij}}^{ij}\}$ from prospective suppliers. Each bid B_k^{ij} includes a bid price bp_k^{ij} and a proposed delivery date dl_k^{ij} . Below we use the notation $B_k^{ij} = (dl_k^{ij}, bp_k^{ij})$.

Failure by the manufacturer to meet an order O_i 's due date results in a penalty $tard_i \times T_i$, where T_i is the time by which delivery of the product or service is late, and $tard_i$ is the marginal penalty for missing the delivery date. Such penalties, which are commonly used to model manufacturing scheduling problems, reflect actual

contractual terms, loss of customer goodwill, interests on lost profits or a combination of the above [18].

A solution to the FCMAP problem consists of:

- a selection of bids: $Bid_Comb = \{Bid_Comb_1, \dots, Bid_Comb_m\}$, where Bid_Comb_i ($i \in M$) is a combination of n_i bids - one for each of the components required by order O_i , and
- a collection of start times: $ST = \{st_1, \dots, st_m\}$, where st_i is the time when the manufacturer is scheduled to start processing order O_i , and $st_i \geq dl^j, \forall j=1, \dots, n_i$, since orders cannot be processed before all the components they require have been delivered by suppliers.

Given a solution (Bid_Comb, ST) , the profit of the manufacturer is the difference between the revenue generated by its customer orders (once they have been completed) and the sum of its procurement costs and tardiness penalties. This is denoted:

$$prof(Bid_Comb, ST) = \sum_{i \in M} rev_i - \sum_{i \in M} \sum_{j \in N_i} bp^{ij} - \sum_i tard_i \times T_i \quad (1)$$

where,

- rev_i is the revenue generated by the completion of order O_i (i.e., the amount paid by the customer),
- bp^{ij} is the price of component $comp_{ij}$ in Bid_Comb , and
- $T_i = Max(0, st_i + du_i - dd_i)$ with st_i being the start time of order O_i in ST .

Note that because we assume a given set of orders, the term $\sum_{i \in M} rev_i$ is the same across

all solutions. Accordingly, maximizing profit in Equation (1) is equivalent to minimizing the sum of procurement and tardiness costs: $cost(Bid_Comb, ST)$

$$= \sum_{i \in M} \sum_{j \in N_i} bp^{ij} + \sum_{i \in M} tard_i \times T_i .$$

It is worth noting that the above model contrasts with earlier research in dynamic supply chain formation, which has generally assumed manufacturers with infinite

capacity or fixed lead times and ignored delivery dates and tardiness penalties ([6], [7], [9], [20]).

From a complexity standpoint, it can easily be seen that the FCMAP problem is strongly NP-hard, since the special situation where all components are free and available at time zero reduces to the single machine total weighted tardiness problem, itself a well known NP-hard problem [8].

An example of an exact procedure to solve FCMAP problems involves looking at all possible procurement bid combinations and, for each such combination, solving to optimality a single machine weighted tardiness problem with release dates (e.g., using a branch-and-bound algorithm). A release date is a date before which a given order is not allowed to be processed. Given a combination of procurement bids Bid_Comb_i , an order O_i has a release date:

$$r_i = \text{Max}_{j \in N_i} [dl^{ij}] \quad (2)$$

where dl^{ij} denotes the delivery date of component $comp_{ij}$ in Bid_Comb_i . In other words, the component that arrives the latest determines the order's release date.

Clearly, with the exception of fairly small problems, the requirements of the above procedure are computationally prohibitive. Below, we identify a number of rules that can be used to efficiently prune the search space associated with FCMAP problems.

4. Pruning the Search Space

Pruning Rule 1: Eliminating Expensive Bids with Late Delivery Dates

Consider an FCMAP problem P with an order O_i requiring a component $comp_{ij}$ for which the manufacturer has received a set of bids $\beta_{i_j} = \{B_1^{ij}, \dots, B_{n_y}^{ij}\}$ from prospective suppliers. Let $B_k^{ij} = (dl_k^{ij}, bp_k^{ij})$ and $B_l^{ij} = (dl_l^{ij}, bp_l^{ij})$ be two bids in β_{i_j} such that:

$$dl_l^{ij} \geq dl_k^{ij} \text{ and } bp_l^{ij} \geq bp_k^{ij}.$$

Then problem P' with $\beta'_{i_j} = \beta_{i_j} \setminus \{B_l^{ij}\}$ admits the optimal solutions with the exact same profit as problem P .

The correctness of this rule should be obvious. Its application is illustrated in Figure 2, where an order requires two components: component 1 and component 2. The manufacturer has received bids for each component. Using Rule 1, it can be determined, for instance, that bid_{14} is not competitive given that it is more expensive than bid_{13} and arrives late. Similarly, bid_{22} and bid_{24} can also be pruned.

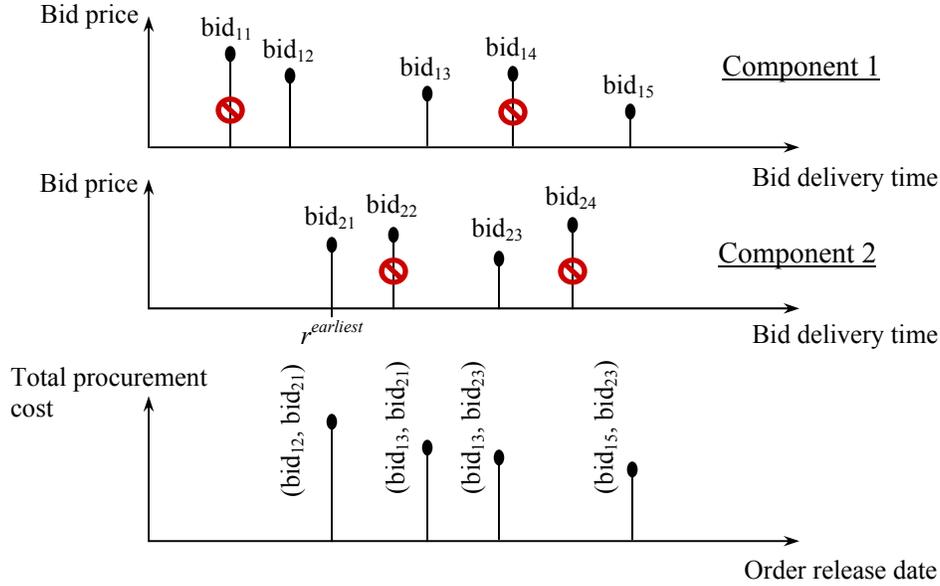


Figure 2. From 20 bid combinations to 4 non-dominated ones

Pruning Rule 2: Eliminating Expensive Bids with Unnecessarily Early Delivery Dates

Consider an FCMAP problem P with an order O_i requiring a set of components $comp_{ij}$, $j \in N_i = \{1, \dots, n_i\}$. Let $\beta_{ij} = \{B_1^{ij}, \dots, B_{n_{ij}}^{ij}\}$ be the set of bids received by the manufacturer for each component $comp_{ij}$ with $B_k^{ij} = (dl_k^{ij}, bp_k^{ij})$. We define $r_i^{earliest}$ as the earliest possible release date for order O_i . It can be computed as:

$$r_i^{earliest} = \text{Max}_{1 \leq j \leq n_i} \text{Min}_{1 \leq k \leq n_{ij}} dl_k^{ij}$$

Let B_k^{ij} and B_l^{ij} be two bids for component $comp_{ij}$ such that:

$$bp_l^{ij} \geq bp_k^{ij} \text{ and } dl_l^{ij} \leq dl_k^{ij} \leq r_i^{earliest}.$$

Then problem P' with $\beta'_{ij} = \beta_{ij} \setminus \{B_l^{ij}\}$ admits the exact same set of optimal solutions as problem P .

An intuitive explanation should suffice to convince the reader. While bid B_l^{ij} has an earlier delivery date than bid B_k^{ij} , this earlier date is not worth paying more for: it does not add any scheduling flexibility to the manufacturer since the start of order O_i remains constrained by $r_i^{earliest} \geq dl_l^{ij}$. A formal proof can easily be built based on this observation.

Note that, in general, it is not possible to prune bid B_k^{ij} . This is because other bids for component $comp_{ij}$ may have delivery dates that are after $r_i^{earliest}$, which would reduce the number of available scheduling options possibly leading to lower quality solutions. Application of this rule is also illustrated in Figure 2, where it results in the pruning of bid_{11} . This is because both bid_{11} and bid_{12} arrive before the order's earliest release date, $r^{earliest}$, and bid_{11} is more expensive than bid_{12} .

Pruning Rule 3: Eliminating Expensive Bid Combinations with Unnecessarily Early Delivery Dates

Consider an FCMAP problem P whose search space has already been pruned using Rule 1. In other words, given two bids $B_k^{ij} = (dl_k^{ij}, bp_k^{ij})$ and $B_l^{ij} = (dl_l^{ij}, bp_l^{ij})$, $k \neq l$, for the same component $comp_{ij}$, if $dl_l^{ij} > dl_k^{ij}$, then $bp_l^{ij} < bp_k^{ij}$.

Let $Bid_Comb_i^a = \{B_a^{i1}, \dots, B_a^{in_i}\}$ be a combination of bids for the n_i components required by order O_i . Suppose also that there exist two bids $B_a^{ik} = (dl_a^{ik}, bp_a^{ik}) \in Bid_Comb_i^a$ and $B_b^{ik} = (dl_b^{ik}, bp_b^{ik}) \in Bid_Comb_i^b$, and a component $comp_{il}$, $l \neq k$ such that $dl_a^{ik} < dl_b^{ik} \leq dl_a^{il}$, then $Bid_Comb_i^a$ is dominated by $Bid_Comb_i^b$, where $Bid_Comb_i^b = (Bid_Comb_i^a \setminus \{B_a^{ik}\}) \cup \{B_b^{ik}\}$. By "dominated" we mean that, for every solution to problem P involving $Bid_Sel_i^a$, there is a better solution where $Bid_Comb_i^a$ is replaced by $Bid_Comb_i^b$.

Given that $Bid_Comb_i^a$ includes a bid for a second component $comp_{il}$ that gets delivered at time $dl_a^{il} \geq dl_b^{ik} > dl_a^{ik}$, replacing bid B_a^{ik} with bid B_b^{ik} will not delay the start of order O_i and can only help reduce the cost of its components since $bp_a^{ik} > bp_b^{ik}$ (as indicated earlier, we assume that Rule 1 has already been applied to prune bids). It is straightforward to build a formal proof based on the above observation. Note also that Rule 3 actually subsumes Rule 2 – Rule 2 is easier to visualize and also introduces the notion of earliest possible release date, which we use later in this article.

The three pruning rules we just identified can be used to prune the set of bids to be considered. This is illustrated in Figure 2, where the combination of the three rules brings the number of bid combinations to be considered from 20 to just 4 non-dominated combinations. In particular, the application of Rule 3 helps us prune bid combination (bid_{12}, bid_{23}) . This is because this combination is dominated by (bid_{13}, bid_{23}) , which results in the same release date but is cheaper. Another bid combination pruned using Rule 3 is (bid_{15}, bid_{21}) .

It should be clear that, for each order, Rules 1 and 2 can be applied in $O(c \cdot b \cdot \log b)$ time, where b is an upper-bound on the number of bids received for a given component and c an upper-bound on the number of components required by a given order. It can also be shown that, for a given order O_i , Rule 3 can be applied in $O(tb \cdot \log tb)$ time, where tb is the total number of bids received for order O_i across all the components it requires. This is done as follows:

1. For each component $comp_{ij}$, create a sorted list $\lambda_{i,j} = \langle B_1^{ij}, \dots, B_{n_{ij}}^{ij} \rangle$ such that $dl_k^{ij} < dl_{k+1}^{ij}$. Create an overall list of delivery dates for all the bids received for O_i (i.e., for all the components required by the order) and sort the delivery dates in increasing order. Let Λ_i be this sorted list
2. For each date r_i in Λ_i , keep only those non-dominated combinations of bids that are compatible with having r_i as order O_i 's release date. Note that such bid combinations are of the form $\{B_a^{i1}, \dots, B_a^{in_i}\}$ where $dl_a^{ij} \leq r_i, \forall j$, and there is

no other bid B_b^{ij} such that $dl_a^{ij} < dl_b^{ij} \leq r_i$. In other words, for each component $comp_{ij}$, B_a^{ij} is the latest bid compatible with release date r_i (and hence also the cheapest such bid). Finding such bids requires very little time, given the sorted bid lists $\lambda_{i,j}$ created in step 1.

As a parenthesis, it is worth noting that the three pruning rules we just introduced apply to scenarios with more complex constraints, as they only take advantage of the release constraint that requires each order to have all its components before it can be processed. For instance, this includes problems where the manufacturer is modeled as a job shop, the capacity of some machines is greater than one and there are sequence dependent setup times. It can also be shown that the pruning rules can be extended to accommodate problems with inventory holding costs, as long as orders are not allowed to be shipped before their due dates – this assumption corresponds to having finished goods inventory and is representative of many supply chain situations.

Consider the non-dominated bid combinations resulting from the application of our three pruning rules to an FCMAP problem. Let the non-dominated bid combinations of order O_i be denoted:

$$Bid_Comb_i^* = \{Bid_Comb_i^1 = (r_{i1}, pc_{i1}), \dots, Bid_Comb_i^{m_i} = (r_{im_i}, pc_{im_i})\},$$

where r_{ik} is the release date of bid combination $Bid_Comb_i^k$, as defined in Equation (2), and pc_{ik} is its total procurement cost, defined as the sum of its component bid prices. It follows that:

Property 1: *For each order O_i , $i \in M$, it must hold that, if $r_{ia} < r_{ib}$, then $pc_{ia} > pc_{ib}$, $\forall a, b \in \{1, \dots, m_i\}, a \neq b$. In other words, the total procurement costs of non-dominated bid combinations strictly decrease as their release dates increase.*

Proof:

We have already shown that, following the application of Rule 1, the bids that remain for a given component have prices that strictly decrease as their delivery dates increase.

Let $Bid_Comb_i^a$ be a non-dominated bid combination for order O_i – following the application of Rules 1 through 3. Let its release date r_{ia} be determined by the delivery

date of component j , namely $r_{ia} = dl_a^{ij}$. Note that, by definition, the release date of a bid combination is always determined by one or more of its components. Given that Rule 3 has already been applied, the delivery date dl_a^{ik} of any component k must be the latest delivery date among those bids for component k that satisfy $dl_a^{ik} \leq dl_a^{ij}$.

Consider another non-dominated bid combination $Bid_Comb_i^b$ for order O_i such that $r_{ib} > r_{ia}$. Let l be the index of one of the components determining the release date of bid combination $Bid_Comb_i^b$, namely $r_{ib} = dl_b^{il} > r_{ia} = dl_a^{ij}$. Just as for bid combination $Bid_Comb_i^a$, the fact that Rule 3 has been applied implies that the delivery date dl_b^{ik} of any component k in $Bid_Comb_i^b$ must be the latest delivery date among those bids for component k that satisfy $dl_b^{ik} \leq dl_b^{il}$. Given that $r_{ib} = dl_b^{il} > r_{ia} = dl_a^{ij}$, it also follows that, for any component k , we have $dl_b^{ik} \geq dl_a^{ik}$ with a strict inequality for at least one component, namely component l . Given that Rule 1 has been applied, it also follows that, for any component k , $bp_b^{ik} \leq bp_a^{ik}$ with a strict inequality for at least one component (component l). Hence, $pc_{ia} = \sum_{1 \leq k \leq n_i} bp_a^{ik} > pc_{ib} = \sum_{1 \leq k \leq n_i} bp_b^{ik}$. \square

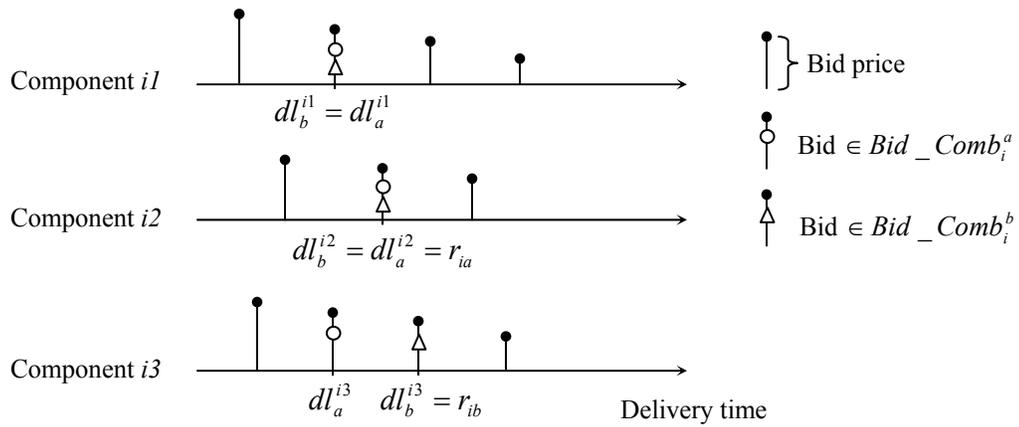


Figure 3. Illustration of Property 1

Property 1 is illustrated in Figure 3, where we have two bid combinations $Bid_Comb_i^a$ and $Bid_Comb_i^b$ for an order O_i that requires three components. In this particular example, r_{ia} is determined by the delivery date of component 2, while r_{ib} is determined by that of component 3. The two bid combinations share the same delivery dates for two out of three of the components required by order O_i : components 1 and 2. The difference in procurement cost comes from the lower price associated with the later delivery of component 3 in bid combination $Bid_Comb_i^b$ (namely, $dl_b^{i3} = r_{ib} > dl_a^{i3}$).

Note also that, if after application of the pruning rules there exist two non-dominated bid combinations, $Bid_Comb_i^a$ and $Bid_Comb_i^b$, such that $r_{ia} = r_{ib}$, it must hold that $pc_{ia} = pc_{ib}$.

In the following sections, we introduce a branch-and-bound algorithm to solve the FCMAP problem along with two (significantly faster) heuristic search procedures. All three procedures take advantage of the pruning rules we just introduced. One of the two heuristic procedures also takes advantage of Property 1.

5. A Branch-and-Bound Algorithm

Following the application of the pruning rules introduced in the previous section, optimal solutions to the FCMAP problem can be obtained using a branch-and-bound procedure. Branching is done over the sequence in which orders are processed by the manufacturer and over the release dates of non-dominated bid combinations of each order. Specifically, the algorithm first picks an order to be processed by the manufacturer then tries all the release dates (of non-dominated bid combinations) available for this order. Note that, as orders are sequenced in this fashion, some of their available release dates become dominated, given prior sequencing decisions. For instance, consider two orders O_1 and O_2 , with O_2 having two release dates r_{21} and r_{22} with $r_{21} < r_{22}$ - following the application of Pruning Rules 1 through 3. Suppose that, at the current node, O_1 is sequenced before O_2 and that O_1 's earliest completion

date is greater than r_{22} . It follows that release date r_{21} is strictly dominated by release date r_{22} at this particular node. Release dates that become dominated as a result of prior assignments can be pruned on the fly, thereby further speeding up the search procedure. Given a node n in the search tree, namely a partial sequence of orders and a selection of release dates for each of the orders already sequenced, it is possible to compute an lower-bound for the profit of all complete solutions (i.e., leaf nodes) compatible with this node:

$$LB_n = \sum_{i \in OS_n} (pc_i + tard_i \times T_i) + \sum_{i \notin OS_n} [tard_i \times \max(0, cd_{OS_n} + du_i - dd_i) + mpc_i],$$

where:

- OS_n is the set of orders sequenced at node n ;
- pc_i is the total procurement cost associated with the non-dominated release date (or bid combination) assigned to order $O_i \in OS_n$ and T_i is its tardiness. Note that each order is scheduled to start as early as possible, given prior sequencing decisions and the release date assigned to it: there are no benefits to starting later;
- cd_{OS_n} is the completion date of the last order in OS_n ;
- mpc_i is the minimum possible procurement cost of order O_i – this cost is node-independent.

If the lower bound of a node n is greater than the best feasible solution found so far, the node n and all its descendants are pruned.

6. An Early/Tardy Heuristic

Property 1 tells us that, following the application of the pruning rules, the procurement costs of non-dominated bid combinations strictly decrease as release dates increase. Figure 4 plots the total procurement cost and tardiness cost of an order for different possible start times. While tardiness costs increase linearly with start times that miss the order's due date, procurement costs vary according to a decreasing

step-wise function. Specifically, the circles in Figure 4 represent the order's non-dominated bid combinations. For instance, if the order starts at time st , its procurement cost is pc_i , namely the procurement cost of the latest non-dominated bid combination compatible with this start time (Bid_Comb^i). Its tardiness cost is equal to $tard \times \max(0, st + du - dd)$, where $tard$ is its marginal tardiness penalty, dd its due date and du its duration (or processing time). The resulting problem can be viewed as a pseudo early/tardy scheduling problem. It bears a lot of similarity to a traditional early/tardy scheduling problem (e.g., [17]) but is also slightly different because procurement costs associated with different bid combinations lead to:

- 1) Step-wise earliness costs – in contrast to linear earliness costs found in a traditional early/tardy problem; and
- 2) Potential savings for completing the order past its due date, to the extent that there are late bid combinations that are so cheap that it is worth finishing the order late. Again this is different from a traditional early/tardy problem, where finishing an order on time always leads to the lowest cost for that order.

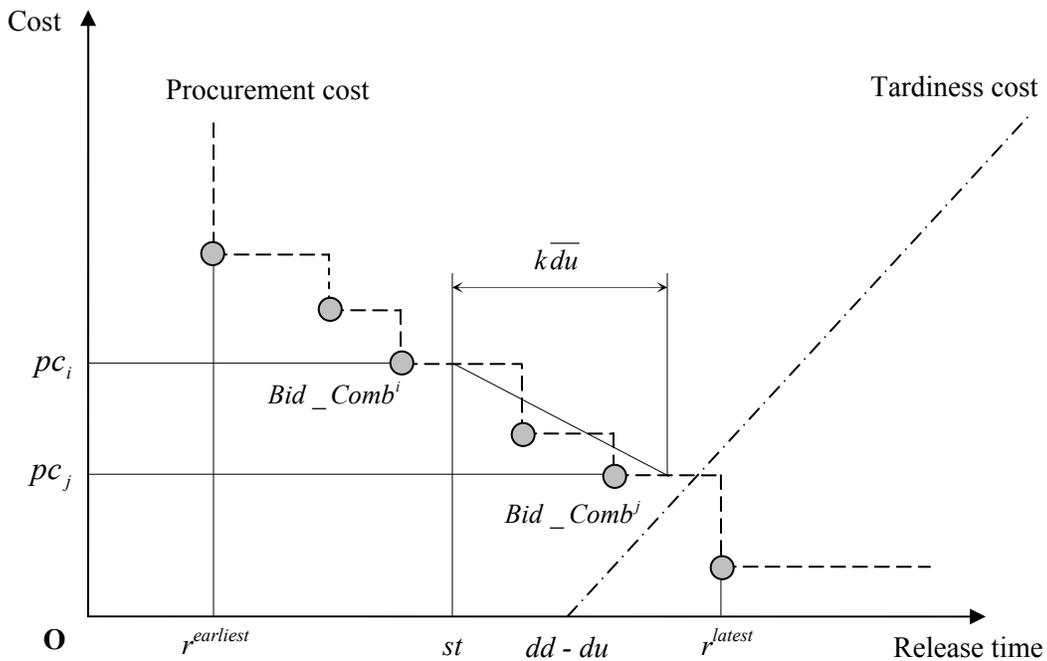


Figure 4. An order's tardiness and procurement costs

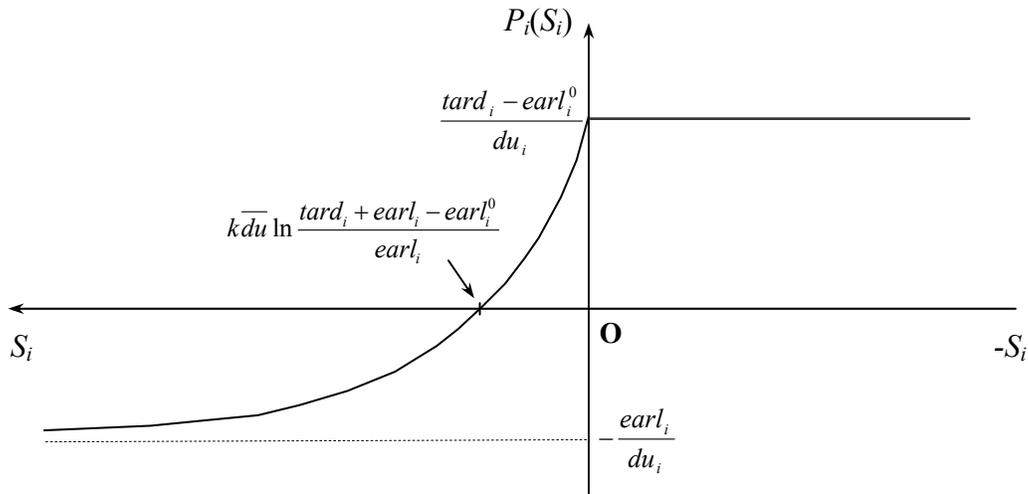


Figure 5. Priority function for Pseudo-Early/Tardy heuristic

Ow and Morton [17] have introduced an early/tardy dispatch rule for one-machine scheduling problems subject to linear earliness and tardiness costs. Because our earliness costs are not linear, this heuristic can not readily be applied. Below, we briefly review some of its key elements and discuss how we have adapted it to produce a family of heuristic search procedures for the FCMAP problem.

Ow and Morton's dispatch rule interpolates between two extreme cases. The first situation is one where all orders are assumed to have plenty of time and where only earliness costs need to be minimized. The second case is one where all orders are assumed to be late and where only tardiness needs to be minimized. In the former case, it can be shown that an optimal solution can be built by sequencing orders according to a *Weighted Longest Processing Time* dispatch rule, where each order receives a priority:

$$P_i = -earl_i / du_i ,$$

where P_i is the priority of order O_i , du_i is its processing time and $earl_i$ is its marginal earliness cost – namely the penalty incurred for every unit of time the order finishes before its due date. Conversely, in the latter case, when all jobs are assumed to be tardy, it can be shown that an optimal solution can be built by sequencing orders according to a *Weighted Shortest Processing Time* dispatch rule of the form:

$$P_i = tard_i / du_i ,$$

where $tard_i$ is its marginal tardiness penalty.

Like Ow and Morton's, our early/tardy heuristic interpolates between two extreme cases: one where all orders have plenty of time and one where all orders are late. The priority associated with this latter situation is different however from the one in Ow and Morton's rule. This is because later start times for orders that are late may still result in reductions in procurement costs. Accordingly, the priority associated with this latter situation is:

$$P_i = (tard_i - earl_i^0) / du_i ,$$

where $earl_i^0$ is the earliness weight at $st = dd_i - du_i$.

The resulting early/tardy heuristic assigns each order a priority that varies with its slack S_i :

$$P_i(S_i) = -\frac{earl_i}{du_i} + \frac{tard_i + earl_i - earl_i^0}{du_i} \times \exp\left[-\frac{(S_i)^+}{k \cdot \bar{du}}\right] \quad (3)$$

where \bar{du} is the average processing time of an order, k is a look-ahead parameter. $(X)^+$ denotes $\text{Max}(X, 0)$, and slack S_i at time t is defined as:

$$S_i = dd_i - du_i - t .$$

The above formula can easily be seen to reduce to the *Weighted Shortest Processing Time* dispatch rule with a marginal tardiness cost of $tard_i - earl_i^0$ when slack $S_i \leq 0$ and to the *Weighted Longest Processing Time* dispatch rule when $S_i \rightarrow \infty$. The look-ahead parameter k can intuitively be thought of as the average number of orders that would be tardy if order O_i is selected to be scheduled next. The value of k basically controls the transition between the two extreme scenarios between which this rule interpolates. Higher values of k make the transition start earlier. This can be interpreted as being more sensitive to tardiness when a larger number of orders stand to be late.

In the FCMAP problem, an order O_i cannot start before its earliest possible release date r_i^{earliest} (see Pruning Rule 2 – it should be clear that this release date is never pruned by Rule 2). In addition, earliness costs vary according to a step function.

A marginal earliness cost can however be obtained through regression, whether locally or globally. Specifically, we distinguish between the following two approaches to computing marginal earliness costs for an order in the FCMAP problem:

1) *Local Earliness Weight*: At time t , the local marginal earliness cost associated with an order O (see Figure 4) can be approximated as the difference in procurement costs associated with the latest non-dominated bid combinations compatible with processing the order at respectively time t (namely Bid_Comb^i) and time $t + k \cdot \overline{du}$ (namely Bid_Comb^j):

$$earl^L = \frac{pc_i - pc_j}{k du},$$

2) *Global Earliness Weight*: An alternative involves computing a single global marginal earliness cost for each order. This can be done using a Least Square Regression:

$$earl^G = \frac{\sum pc \cdot rd - n \cdot \overline{pc} \cdot \overline{rd}}{\sum rd^2 - n \cdot \overline{rd}^2},$$

where \overline{pc} is the average procurement cost of non-dominated bid combinations for the order, and rd is their average release date.

The simplest possible release policy for the FCMAP problem involves releasing each order O_i at its earliest possible release date, namely $r_i^{earliest}$. We refer to this policy as an *Immediate Release Policy*. It might sometime result in releasing some orders too early and hence yield unnecessarily high procurement costs. An alternative is to use an *Intrinsic Release Policy*, which releases orders when their early/tardy priority $P_i(S_i)$ becomes positive. $P_i(S_i)$ can be viewed as the marginal cost incurred for delaying the start of order O_i at time t . As long as this cost is negative, there is no benefit to releasing the order. The tipping point, where $P_i(S_i) = 0$, is the order's intrinsic release date:

$$\hat{r}_i = dd_i - du_i + k \overline{du} \cdot \ln \frac{earl_i}{tard_i + earl_i - earl_i^0}. \quad (4)$$

Here again, one can use either the local or global earliness weights associated with an order. Intuitively, one would expect the global earliness weight to be more appropriate for the computation of an order's release date and its local earliness weight to be better suited for the computation of its priority at a particular point in time. This has generally been confirmed in our experiments. In Section 9, we only present results where priorities are computed using local earliness weights. We do however report results, where release dates are computed with both local and global earliness weights, and we found our heuristic performs better with global earliness weight.

Rather than limiting ourselves to deterministic adaptations of Ow and Morton's dispatch rule, we have also experimented with randomized versions, where order release dates and priorities are modified by small stochastic perturbations. This enables our procedure to make up for the way in which it approximates procurement costs, sampling the search space in the vicinity of its deterministic solution. The resulting pseudo-early/tardy search heuristic operates by looping through the following procedure for a pre-specified amount of time. As it iterates, the procedure alternates between the immediate and intrinsic release policies discussed earlier and successively tries a number of different values for the heuristic's look-ahead parameter k . The following outlines one iteration – i.e. with one particular release policy and one particular value of the look-ahead parameter.

1. For each order O_k , $k \in M = \{1,2,\dots,m\}$, compute the order's release date. When using the immediate release policy, this simply amounts to setting the order's release date $RD_k = r_k^{earliest}$. When using the intrinsic release policy, the order's release date is computed as $RD_k = \text{Max}\{r_k^{earliest}, (1 + \alpha) \times \hat{r}_k\}$, where α is randomly drawn from the uniform distribution $[-dev_1, +dev_1]$ (dev_1 is a parameter that controls how widely the procedure samples the search space);
2. Dispatch the orders, namely let $t_0 = \text{Min}_{k \in M} RD_k$
 - 1) For all those orders O_k that have not yet been scheduled and whose release dates are before t_0 , compute the order's priority at time t_0 as:

$$PR_k(t_0) = (1 + \beta) \cdot P_k(dd_k - du_k - t_0),$$

where P_k is the pseudo-early/tardy priority defined in (3) and β is randomly drawn from the uniform distribution $[-dev_2, +dev_2]$ (dev_2 is a parameter that controls how widely the procedure samples the search space);

- 2) Let order O_i be the order with the highest priority. Schedule O_i to start at time t_0 ;
- 3) If all orders have been scheduled, then Stop. Else, let $t_1 = t_0 + du_i$ and t_2 be the earliest release date among those orders that have not yet been scheduled. Set $t_0 = \text{Max}\{t_1, t_2\}$ and repeat Steps 1-3.
- 4) Compute the profit of the resulting solution. If it is higher than the best solution obtained so far, make this the new best solution.

A deterministic version of this procedure simply amounts to setting dev_1 and dev_2 to zero.

7. A Simulated Annealing Search Procedure

A second heuristic search procedure for the FCMAP problem involves using Simulated Annealing (SA) to explore different combinations of bids. Given a selection of non-dominated bid combinations $Bid_Comb = \{Bid_Comb_1, \dots, Bid_Comb_m\}$ – one combination per order, the procedure computes the release date r_i of each order O_i and sequences the orders, using the Apparent Tardiness Cost (ATC) dispatch rule first introduced in [25]. ATC is known to generally yield high quality schedules for the one-machine total weighted tardiness problem and has a $O(m \cdot \log m)$ complexity. As such it is an excellent estimator for the best solution compatible with a given selection of bid combinations. The following further details the SA procedure:

Step 1 – Initialization:

Set an initial temperature $Temp = Temp_0$, and an initial bid selection $Bid_Comb^1 = \{Bid_Comb_1^1, \dots, Bid_Comb_m^1\}$;

Use the ATC dispatch rule to build a schedule. Let $cost^1 = cost(Bid_Comb^1, ST^1)$, where $ST^1 = \{st_1^1, \dots, st_m^1\}$ is the set of start times assigned by ATC to orders O_1 through O_m . Set $Bid_Comb^{opt} = Bid_Comb^1$ and $cost^{opt} = cost^1$.

Step 2 – Search:

Perform the following step N times:

Select $Bid_Comb = neighbor(Bid_Comb^1)$ (randomly or through some heuristic), and compute $cost = cost(Bid_Comb, ST)$, where ST is the set of order start times assigned by the ATC dispatch rule;

If $cost^1 \geq cost \geq cost^{opt}$, set $Bid_Comb^1 = Bid_Comb$;

Else if $cost > cost^1$ and $rand() \leq exp((cost^1 - cost)/Temp)$, set $Bid_Comb^1 = Bid_Comb$;

Else if $cost < cost^{opt}$, set $Bid_Comb^{opt} = Bid_Comb^1 = Bid_Comb$.

If Bid_Comb^{opt} was not modified in the last N iterations, decrease the temperature $Temp = Temp \cdot \alpha$. Go to Step 3.

Step 3 – Termination Condition:

If Bid_Comb^{opt} has not been improved over the past K steps, then STOP and return $(Bid_Comb^{opt}, ST^{opt})$ as the best solution found by the procedure, otherwise go to Step 2.

The initial bid combination Bid_Comb^1 is randomly generated. Note also that the ATC dispatch rule is itself a parametric dispatch rule with a look-ahead parameter [25]. In our experiments, we systematically run ATC with values of the look-ahead parameter equal to 0.5, 1.0, 1.5, ..., 6.0 and pick the best of the 12 solutions we have generated.

We have studied variations of this procedure that rely on different movesets. In particular, we have considered a one-bid moveset variation, where we modify the selection of a single bid (for a given component), and a two-bid moveset variation, where two bid selections (for two different components) are modified at once. For both types of movesets, we have also experimented with two ways of selecting moves:

- a random mechanism, where a move in the moveset is randomly selected, and

- an organized mechanism that replaces the bid(s) that reduce most the profit of the current solution (Bid_Comb^1 , ST^1), namely, those bids for which $bp^{ij} + tard_i \cdot T_i$ is the greatest.

The experiments presented in Section 9 use the organized mechanism, as we found it to generally yield higher quality solutions than the random one. We have not found a great difference between variations of our procedure using one-bid movesets and two bid movesets.

8. Right-Shifting Solution Improvement Procedure

Since the total cost function of each order is the sum of a linearly increasing tardiness cost function and a step-wise non-increasing earliness cost function, the total cost function has multiple local minimum points, as shown in Figure 6. Meanwhile, the above pseudo-early/tardy heuristic dispatches orders as soon as the machine becomes available. Hence, the solution produced by the pseudo-early/tardy heuristic can sometimes be further improved by right-shifting orders to lower local minimum points without changing the order processing sequence. Let $\langle 1, 2, \dots, m \rangle$ denote an order processing sequence produced by the above algorithm, st_i denote the start time of order O_i , $i \in \{1, 2, \dots, m\}$. As shown in Figure 6, the local minimum points have a one-to-one correspondence with the non-dominated bid combinations. The improvement method processes each order i from m to 1 as follows:

1. Right-shift early orders: If $st_i < dd_i - du_i$ and $st_i < st_{i+1} - du_i$, right-shift order O_i until $Min\{dd_i - du_i, st_{i+1} - du_i\}$, i.e., let $st_i = Min\{dd_i - du_i, st_{i+1} - du_i\}$. Even if its cost remains the same, right-shifting O_i creates more room to right-shift earlier orders, and therefore may allow to decrease the costs of the other orders. Go to Step 2.
2. Under any of the following situations, stop right-shifting:
 - $st_i = st_{i+1} - du_i$;
 - There is no local minimum point between st_i and $st_{i+1} - du_i$; or

- All the local minimum points between st_i and $st_{i+1} - du_i$ have higher costs than the current cost, tc_i , of order O_i .

Otherwise, let j^* be a local minimum point, the cost of which is the lowest among all the points between st_i and $st_{i+1} - du_i$. Right-shift order O_i to the local minimum point j^* , i.e., let $st_i = st_{j^*} - du_i$.

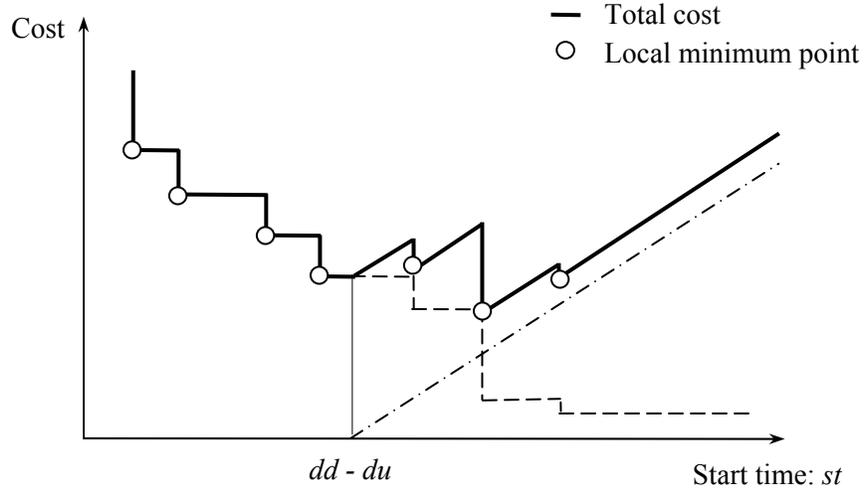


Figure 6. An order's total cost

9. Computational Evaluation

A number of experiments have been run to evaluate the impact of our pruning rules, the performance of our heuristic search procedures, and the benefits of our FCMAP model over traditional reverse auction models that ignore the manufacturer's finite capacity. These experiments are further detailed below.

Empirical Setup

Problems were randomly generated to cover a broad range of conditions by varying the distribution of bid prices and bid delivery dates as well as the overall load faced by the manufacturer. Results are reported for 2 groups of problems:

1. **Problems with 10 orders, 5 required components per order and 20 supplier bids per component:** These problems were kept small enough so

that they could be solved to optimality with our branch-and-bound algorithm. Two sets of problems were designed. Key parameter values of the first set were drawn from the following uniform distributions:

- Order processing time: $U[5,25]$
- Order marginal tardiness cost: $U[1,10]$
- Order due dates: 2 distributions:
 - i. Medium Load (*ml*) problems: $U[100,300]$
 - ii. Heavy Load (*hl*) problems: $U[100,200]$
- Component bid deliveries: 2 distributions:
 - i. Narrow bid delivery distribution (*nd*): $U[0,50]$
 - ii. Wide bid delivery distribution (*wd*): $U[0,100]$
- Component bid prices: 2 distributions:
 - i. Narrow bid price distribution (*np*): $U[5,35]$
 - ii. Wide bid price distribution (*wp*): $U[5,65]$

The other set of problems has all the same distributions except for the component bid deliveries distributions:

- i. Narrow bid delivery distribution (*nd*): $U[0,150]$
- ii. Wide bid delivery distribution (*wd*): $U[0,200]$

The first set of problems represents relatively easy problem instances, and in almost all cases, it holds that $dd - du > r^{latest}$, where r^{latest} is the latest release date determined by the cheapest non-dominant bid combination. Namely, the cost function is quasiconvex and has no local minimum point after $dd - du$. We call this set of problems *Early-Bid* problems. The other set of problems, which we call *Mixed-Bid* problems, represents relatively hard problem instances, where, in many situations, $dd - du < r^{latest}$, i.e., the cost function has multiple local minimum points after $dd - du$ (see Figure 6). A total of 20 problems were generated in each category (ml/hl, nd/wd, np/wp), yielding a total of 320 problems.

2. **Problems with 500 orders, 5 required components per order and 20 supplier bids per component:** While these problems were too large to be solved with branch-and-bound (even with our pruning rules), they were used to

validate results obtained on the smaller sets of problems. This includes, determining how our heuristic search procedures scale up and evaluating the benefits of our FCMAP model over traditional reverse auction models and policies that ignore the manufacturer’s finite capacity – these latter being simply referred to below as “infinite capacity” policies. Key parameter values were drawn from the following uniform distributions:

- Order processing time: $U[1,5]$
- Order marginal tardiness cost: $U[1,10]$
- Order due dates:
 - i. Medium Load (*ml*) problems: $U[500,1500]$
 - ii. Heavy Load (*hl*) problems: $U[500,1000]$
- Component bid deliveries: 2 distributions:
 - i. Narrow bid delivery distribution (*nd*): $U[0,800]$
 - ii. Wide bid delivery distribution (*wd*): $U[0,1000]$
- Component bid prices: same 2 distributions as 10-order problems (*np/wp*)

A total of 20 problems were generated in each category for a total of 160 problems.

Note that order revenues are irrelevant, since the orders to be produced are fixed. In other words, all solutions admit the same overall revenue and overall profit is solely determined by the sum of tardiness and procurement costs associated with a given solution – see equation (1). Accordingly, we report overall costs rather than overall profits.

Distance from the Optimum

Tables 1 and 2 summarize results obtained on 10-order/*Early-Bid* problems. Tables 3 and 4 summarize 10-order/*Mixed-Bid* problems. These tables provide the average distance from the optimum of solutions obtained with different variations of our search heuristics across 16 problem sets (four medium load, *Early-Bid* problem sets in Table 1, four heavy load, *Early-Bid* problem sets in Table 2, four medium load, *Mixed-Bid* problem sets in Table 3, and four heavy load, *Mixed-Bid* problem sets in

Table 4) with each problem set including a total of 20 problems. This distance from the optimum was computed as:

$$[\text{cost}(\text{solution}) - \text{cost}(\text{optimal_solution})] / \text{cost}(\text{optimal_solution}).$$

Standard deviations are provided between parentheses. Optimal solutions were obtained using the branch-and-bound procedure introduced in Section 5. Results are reported for the following techniques:

- **Infinite capacity:** This is a technique that reflects traditional reverse auction practices, where the manufacturer's capacity is ignored. Specifically, for each order and each component, the manufacturer selects the cheapest bid compatible with the order's due date. Orders are then scheduled according to the ATC dispatch rule, namely the same rule used in our Simulated Annealing procedure (See Section 7).
- **Finite capacity:** Results are reported for a number of variations of our search heuristics:
 - Simulated Annealing (SA) procedure: This is the procedure introduced in Section 7 with $Temp_0=300$, $\alpha=0.95$, $N=60$ and $K=40$. The procedure was run five times on each problem and we report both average performance and best performance over 5 runs.
 - Pseudo-Early/Tardy (PET) Procedure: this is the pseudo-early/tardy heuristic introduced in Section 6. Here again we report results for several variations of this heuristic:
 - *G-L* uses global earliness weights in its release policy and local earliness weights in its priority computations,
 - *L-L* uses local earliness weights for both release date and priority computations,
 - *Det* is a deterministic variation of the pseudo-early/tardy heuristic described in Section 6, namely $dev_1 = dev_2 = 0$,
 - *Rand* is a stochastic variation of the same heuristic with $dev_1 = 0.3$ and $dev_2 = 0.3$. For comparison sake, the CPU time given to this heuristic was the same CPU time required by an average SA run in the same problem category,

- *Rand-RS* improves the solutions produced by *Rand* using the post-processing procedure described in Section 8,
- *Hybrid* is a heuristic that runs SA once, PET/L-L/Rand-RS once, PET/G-L/Rand-RS once and takes the best of the resulting solutions.

Table 1. Percentage deviation from the optimum – 10-order/early-bid/medium-load problems (Standard deviations are provided between parentheses)

	Inf. Cap.	Finite Capacity								Hyb.
		SA		PET						
		Avg. of 5 Runs	Best of 5 Runs	L-L			G-L			
				Det	Rand	Rand-RS	Det	Rand.	Rand-RS	
np/nd	35.04 (36.84)	26.68 (24.04)	22.39 (21.29)	1.82 (333)	0.29 (0.64)	0.19 (0.57)	1.89 (4.69)	0.03 (0.12)	<u>0.00</u> (0.00)	<u>0.00</u> (0.00)
wp/nd	37.57 (26.42)	30.87 (20.14)	23.04 (14.21)	2.06 (4.36)	0.03 (0.12)	<u>0.03</u> (0.12)	2.71 (5.28)	0.04 (0.18)	<u>0.03</u> (0.18)	<u>0.00</u> (0.00)
np/wd	92.25 (84.21)	42.43 (38.62)	25.76 (24.71)	7.31 (5.25)	3.45 (4.12)	1.88 (2.42)	3.01 (2.86)	0.48 (0.81)	<u>0.46</u> (0.79)	<u>0.37</u> (0.61)
wp/wd	76.58 (72.95)	39.01 (28.34)	24.79 (18.88)	9.74 (9.20)	4.08 (5.70)	3.19 (4.97)	4.75 (5.52)	1.14 (2.52)	<u>1.03</u> (2.51)	<u>0.80</u> (2.45)

Table 2. Percentage deviation from the optimum – 10-order/early-bid/heavy-load problems (Standard deviations are provided between parentheses)

	Inf. Cap.	Finite Capacity								Hyb.
		SA		PET						
		Avg. of 5 Runs	Best of 5 Runs	L-L			G-L			
				Det	Rand	Rand-RS	Det	Rand	Rand-RS	
np/nd	37.16 (34.79)	25.40 (22.56)	20.12 (18.84)	7.85 (5.76)	2.31 (2.71)	<u>2.15</u> (2.55)	8.21 (6.62)	2.65 (3.91)	2.39 (3.67)	<u>1.24</u> (1.64)
wp/nd	47.83 (41.49)	37.97 (31.14)	33.00 (26.83)	14.64 (12.11)	3.59 (6.32)	3.51 (6.28)	13.75 (12.34)	3.13 (4.46)	<u>3.07</u> (4.31)	<u>1.93</u> (2.92)
np/wd	131.11 (80.67)	56.17 (25.44)	36.53 (19.36)	14.51 (7.84)	6.50 (4.03)	4.92 (3.11)	10.23 (8.74)	3.09 (4.04)	<u>3.02</u> (4.09)	<u>2.28</u> (2.16)
wp/wd	165.93 (100.54)	77.20 (36.86)	52.08 (23.56)	20.65 (8.90)	10.34 (6.33)	9.47 (6.04)	18.74 (8.33)	6.13 (5.12)	<u>5.75</u> (4.36)	<u>5.23</u> (3.63)

Table 3. Percentage deviation from the optimum – 10-order/mixed-bid/medium-load problems (Standard deviations are provided between parentheses)

	Inf. Cap.	Finite Capacity								Hyb.
		SA		PET						
		Avg. of 5 Runs	Best of 5 Runs	L-L			G-L			
		Det	Rand	Rand-RS	Det	Rand	Rand-RS			
np/nd	89.73 (78.07)	52.86 (40.77)	30.24 (22.09)	16.92 (7.87)	8.75 (6.81)	6.30 (5.02)	3.90 (4.25)	0.85 (1.25)	<u>0.62</u> (1.13)	<u>0.57</u> (1.14)
wp/nd	59.71 (52.40)	36.41 (26.04)	24.23 (16.98)	17.78 (9.38)	11.06 (8.85)	7.30 (6.15)	4.72 (5.03)	2.04 (3.10)	<u>1.94</u> (3.08)	<u>1.41</u> (2.01)
np/wd	86.09 (84.32)	55.73 (58.56)	35.44 (42.25)	17.80 (10.17)	10.40 (7.29)	7.92 (6.82)	5.66 (3.36)	2.29 (2.41)	<u>1.65</u> (1.80)	<u>1.62</u> (1.82)
wp/wd	106.61 (79.84)	81.53 (65.82)	61.93 (59.13)	22.60 (10.25)	15.24 (10.17)	11.29 (7.82)	8.46 (5.58)	4.30 (3.61)	<u>3.44</u> (2.80)	<u>3.22</u> (2.80)

Table 4. Percentage deviation from the optimum – 10-order/mixed-bid/heavy-load problems (Standard deviations are provided between parentheses)

	Inf. Cap.	Finite Capacity								Hyb.
		SA		PET						
		Avg. of 5 Runs	Best of 5 Runs	L-L			G-L			
		Det	Rand	Rand-RS	Det	Rand	Rand-RS			
np/nd	187.43 (112.91)	79.69 (43.21)	41.71 (17.80)	19.04 (6.10)	13.82 (5.94)	13.29 (6.17)	14.13 (8.44)	4.17 (4.02)	<u>4.10</u> (3.87)	<u>3.98</u> (3.62)
wp/nd	126.71 (74.62)	71.04 (53.18)	40.29 (30.19)	21.37 (9.47)	12.45 (5.93)	10.74 (5.36)	15.18 (7.68)	6.32 (5.01)	<u>6.06</u> (5.01)	<u>5.68</u> (4.46)
np/wd	253.74 (170.31)	146.23 (127.82)	76.65 (82.99)	17.91 (5.77)	12.68 (5.29)	12.37 (5.31)	11.65 (6.52)	5.08 (4.69)	<u>4.93</u> (4.55)	<u>4.41</u> (3.60)
wp/wd	126.73 (104.02)	93.21 (77.45)	72.92 (71.58)	28.07 (10.77)	16.67 (9.13)	15.84 (9.24)	15.74 (8.71)	6.47 (5.11)	<u>6.41</u> (5.09)	<u>6.41</u> (5.09)

Tables 1-4 yield a number of observations:

- **Importance of the FCMAP Model:** All finite capacity heuristics yield solutions with significantly lower costs than the infinite capacity one, thereby confirming the importance of the FCMAP model. Taking into account finite capacity considerations and tightly coordinating the procurement of the multiple components required by each order subject to these finite capacity considerations significantly improve the manufacturer’s bottom line. The results are generally most impressive on problem categories np/wd and wp/wd, where our hybrid

heuristic respectively reduces total costs by at least 70% on *Early-Bid* problems and 80% on *Mixed-Bid* problems.

- **Distance from the Optimum:** Our hybrid heuristics yields solutions that are respectively less than 3.3% from the optimum on medium-load/*Mixed-Bid* problems and 6.5% from the optimum on heavy-load/*Mixed-Bid* problems. Also, our solutions are respectively less than 0.8% from the optimum on medium-load/*Early-Bid* problems and 5.3% from the optimum on heavy-load/*Early-Bid* problems. In particular, our PET-GL-Rand-RS gets the optimal solutions on all 40 randomly generated *Early-Bid* ml/np/nd and ml/wp/nd problems.
- **Effectiveness of Property 1:** Even deterministic versions of the PET heuristic using G-L yields solutions that are respectively within 4.8% from the optimum on medium-load/*Early-Bid* problems and 18.8% from the optimum on heavy-load/*Early-Bid* problems. Even without right-shifting improvement, our stochastic version of the PET heuristic using G-L yields solutions that are respectively within 1.2% from the optimum on medium-load/*Early-Bid* problems and 6.2% from the optimum on heavy-load/*Early-Bid* problems. This strongly suggests that Property 1 and the way in which our PET heuristic approximates earliness costs are rather effective. Note also that this deterministic version of our heuristic takes only a tiny fraction of a second on these problems.
- **PET heuristic versus SA heuristic:** Given the same amount of CPU time, the PET heuristic performs significantly better than the SA search procedure.
- **Effectiveness of Right-Shifting to Improve PET:** Even for the *Early-Bid* problems, **PET-Rand-RS** respectively improves the solutions given by **PET-Rand** by 0.47% in average using L-L and 0.12% in average using G-L. A look at the results on *Mixed-Bid* problems confirms the effectiveness of **PET-Rand-RS** in the cases where $dd - du > r^{latest}$, and the improvements are respectively 2.00% in average using L-L and 0.29% in average using G-L.
- **Global versus Local Earliness Weights:** The G-L variation of the PET heuristic generally performs much better than L-L on both medium and heavy load problems, particularly on category wp/wd: respectively as much as 3.7% better on *Early-Bid* problems and 9.4% better on *Mixed-Bid* problems (L-L performs

slightly better than G-L only on category *Early-Bid/hl/np/nd*). Additional results not reported here show however that local earliness weights yield significantly better results than global earliness weights when it comes to priority computations. These results confirm our intuition that local earliness weight computations better capture the changing profiles of non-dominated bid combinations, and is better suited for the computation of its priority at a particular point in time, while the global earliness weight is more appropriate for the computation of an order’s release date.

Impact of Ignoring the Manufacturer’s Capacity on Larger Problems

Figure 7 and Table 5 summarize results evaluating the impact of ignoring the manufacturer’s finite capacity on larger problems with 500 orders in medium load situations. Similar results for heavy load situations are provided in Figure 8 and Table 6. It can be seen that the PET-Rand-RS systematically yields much better results than the infinite capacity policy. A look at the cost breakdowns provided in Table 5 and 6 indicates that PET is capable of selectively sacrificing procurement costs to yield significant reductions in tardiness costs. On some problems, PET reduces overall costs by more than 90%. These results further validate the benefits of the FCMAP model advocated in this paper and the way in which our PET heuristic leverages Property 1.

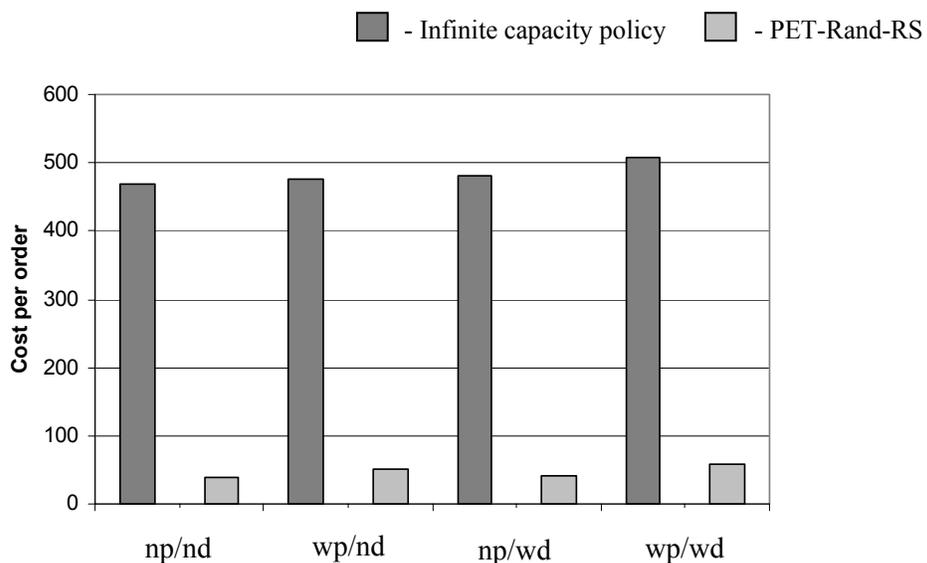


Figure 7. Infinite capacity policy vs. PET-Rand-RS heuristic: average overall cost per order (500-order/medium-load)

Table 5. Cost breakdown (500-order/medium-load)

	ML	Total Cost per Order	Procurement Cost	Tardy Cost
np/nd	Inf. Cap.	469.2 (42.9)	30.6 (0.1)	438.6 (42.8)
	PET-Rand-RS	39.6 (3.0)	36.6 (0.4)	2.9 (3.0)
wp/nd	Inf. Cap.	475.9 (43.8)	38.2 (0.3)	437.7 (43.7)
	PET-Rand-RS	52.1 (2.6)	49.6 (1.0)	2.5 (2.3)
np/wd	Inf. Cap.	482.1 (40.9)	31.4 (0.2)	450.7 (41.0)
	PET-Rand-RS	41.2 (2.9)	38.5 (0.5)	2.8 (2.3)
wp/wd	Inf. Cap.	508.2 (43.0)	40.0 (0.5)	468.2 (43.1)
	PET-Rand-RS	57.7 (2.9)	54.0 (1.0)	3.6 (2.6)

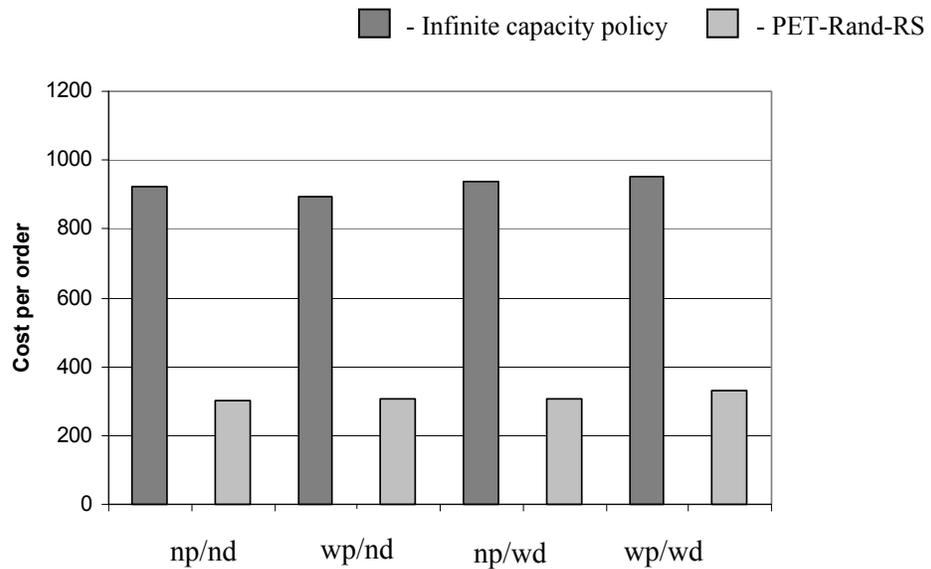


Figure 8. Infinite capacity policy vs. PET-Rand-RS heuristic: average overall cost per order (500-order/heavy-load)

Table 6. Cost breakdown (500-order/heavy-load)

	ML	Total Cost per Order	Procurement Cost	Tardy Cost
np/nd	Inf. Cap.	922.8 (81.6)	31.1 (0.1)	891.6 (81.6)
	PET-Rand-RS	301.3 (36.8)	37.0 (0.7)	264.3 (36.6)
wp/nd	Inf. Cap.	895.8 (62.2)	39.2 (0.3)	856.6 (62.2)
	PET-Rand-RS	307.5 (31.1)	49.1 (0.8)	258.5 (31.2)
np/wd	Inf. Cap.	936.4 (60.5)	32.7 (0.2)	903.7 (60.5)
	PET-Rand-RS	304.8 (30.0)	39.2 (0.7)	265.6 (30.1)
wp/wd	Inf. Cap.	953.2 (78.4)	42.3 (0.4)	910.9 (78.4)
	PET-Rand-RS	328.0 (33.4)	53.5 (0.8)	274.4 (33.4)

Effectiveness of Pruning Rules

To test the effectiveness of pruning rules, the CPU times used by our branch-and-bound algorithm to solve to the optimality are reported in Figure 9 respectively with and without applying the three pruning rules. As can be seen in Figure 9, without pruning the search space, the CPU times increase super-exponentially respectively with the number of orders, the number of components per order and the number of bids per component. If the three pruning rules are applied, the CPU times increase much more slowly with the problem size, and are only around 0.1 second in all tested problems. This confirms the effectiveness of our three pruning rules in reducing the search space. CPU times here were obtained using a 1GHz Pentium-III computer.

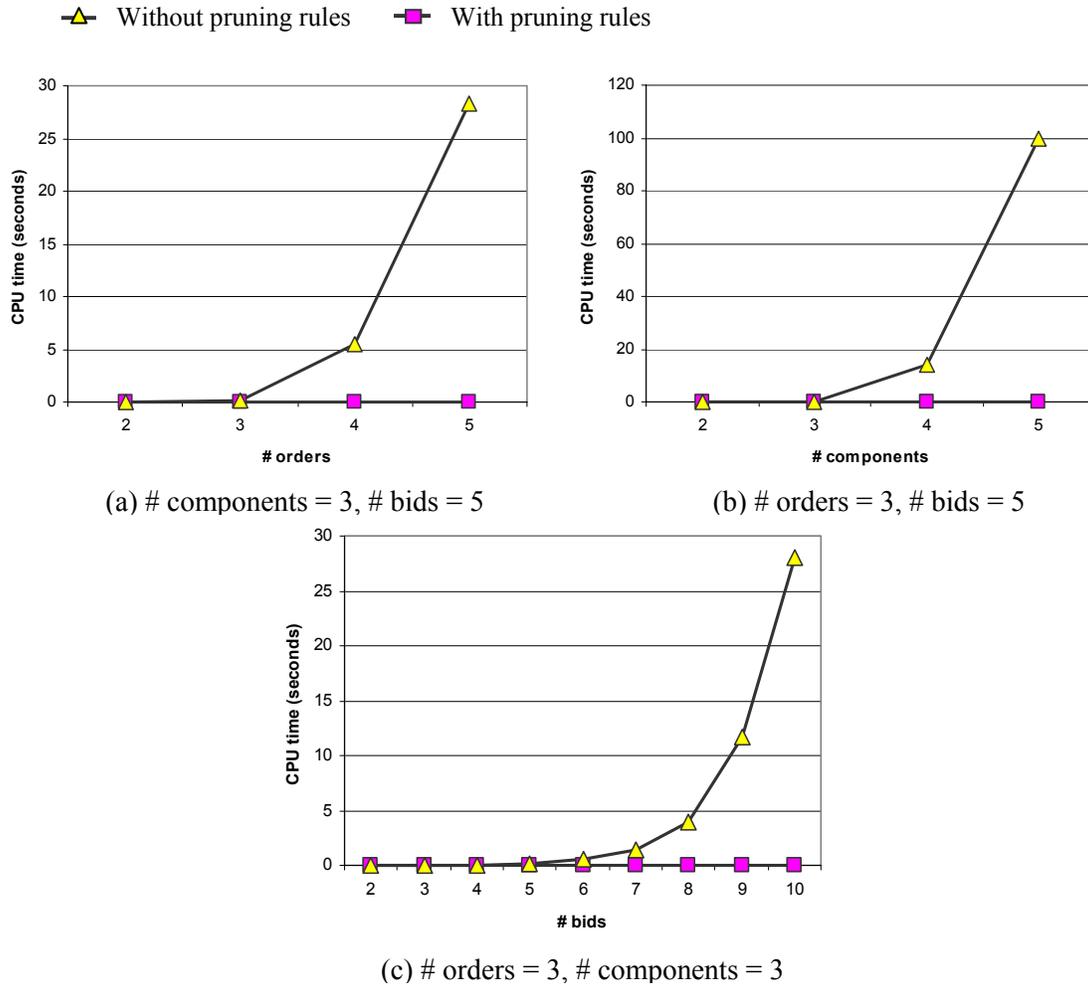


Figure 9. Effectiveness of pruning rules

Computational Requirements

Our computational results suggest that the CPU time required by the SA procedure increases almost linearly with problem size, while that of branch-and-bound grows exponentially (see Figure 10). By design, the CPU time allocated to the PET-Rand heuristic was set to be equal to that of the SA procedure. CPU times of Infinite Capacity Policy, PET-Rand and PET-Rand-RS on large problems of 50 to 500 orders are reported in Figure 11. As expected, the infinite capacity heuristics is the fastest one, though PET does not require more than 12 seconds on these large problems and almost no time in improving the solutions.

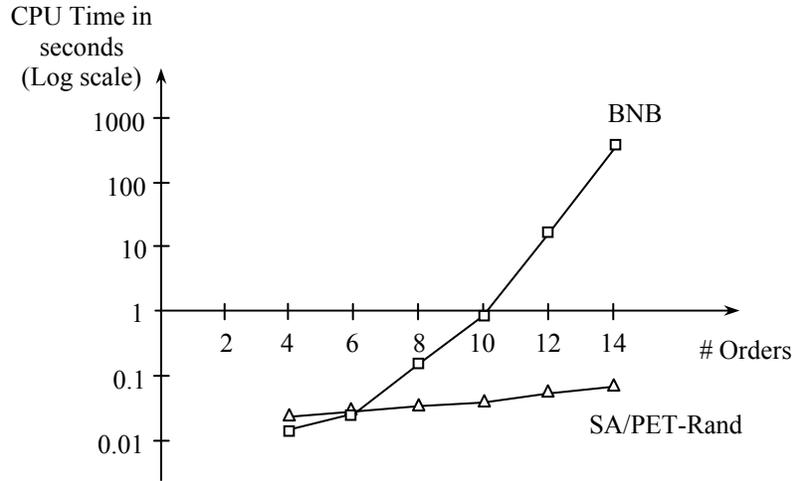


Figure 10. CPU time on small problems

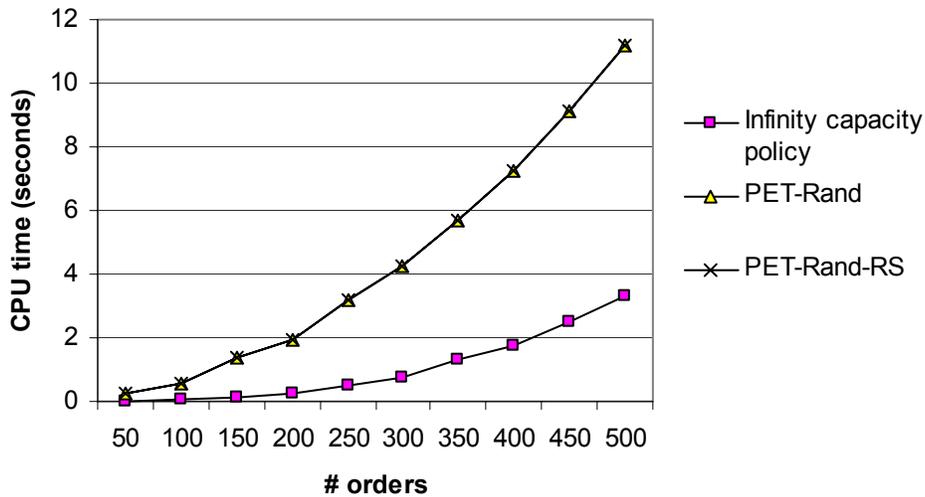


Figure 11. CPU time on large problems

10. Concluding Remarks

Prior work on dynamic supply chain formation has generally ignored capacity and delivery date considerations. In this paper, we have introduced a model for finite capacity multi-attribute procurement problems faced by manufacturers who have to select among supplier bids that differ in terms of prices and delivery dates. We have

identified several dominance criteria that enable the manufacturer to quickly eliminate uncompetitive combinations of bids and have shown that the resulting problem can be modeled as a pseudo-early/tardy problem with stepwise earliness costs. A branch-and-bound algorithm, a randomized pseudo-early/tardy search heuristic and a Simulated Annealing procedure have been introduced to help the manufacturer select a combination of bids that maximizes its overall profit, taking into account its finite capacity as well as the prices and delivery dates associated with different supplier bids. We have shown that these procedures greatly improve over simpler infinite capacity bid selection models. Comparison with optimum solutions obtained using branch-and-bound, suggest that a hybrid heuristic that combines our PET and SA procedures generally yields solutions that are within 6-7% of the optimum.

It should also be noted that the model and techniques presented in this paper can easily be generalized to accommodate situations where the manufacturer can process multiple orders at the same time (non-unary capacity) or where the manufacturer incurs setup times for switching production between different product families. This is true for the pruning rules we introduced as well as the branch-and-bound procedure and two heuristic search procedures. At the same time, we have not attempted to evaluate our techniques on these problems and hence do not know, for instance, how far our heuristic search procedures would be from the optimum. It is also worth noting that our pruning rules also apply to situations where the manufacturer is modeled as a more complex job shop environment, where each order has to flow through a (possibly different) succession of machining (or service) centers. It can also be shown that our model and techniques can be extended to accommodate those problems with inventory holding costs, as long as orders are not allowed to be shipped before their due dates – this assumption corresponds to having finished goods inventory and is representative of many supply chain situations.

Future work will aim to refine our model in support of dynamic profitable-to-promise functionality, where the manufacturer needs to determine how to respond to requests for bids from prospective customers while also selecting among procurement bids from prospective suppliers [2, 20].

Acknowledgements

The research reported in this paper has been conducted as part of the MASCHINE project, a joint research effort between Carnegie Mellon University and the University of Michigan funded by the National Science Foundation under ITR Grant 0205435.

References

- [1] A. Angelus and E.L. Porteus, Simultaneous capacity and production management of short-life-cycle, produce-to-stock goods under stochastic demand, *Management Sci* 48(3) (2002), 399-413.
- [2] R. Arunachalam and N.M. Sadeh. Design of the supply chain trading competition. Proceedings of Workshop on Trading Agent Design and Analysis, IJCAI-03, Acapulco, Mexico, August 2003.
- [3] Y. Bassok and R. Akella, Ordering and production decisions with supply quality and demand uncertainty, *Management Sci* 37(12) (1991), 1556-1574.
- [4] C. Chu, J.M. Proth, Y. Wardi and X. Xie, Supply management in assembly systems: the case of random lead times. *Lecture Notes in Control and Information Sciences* 199: 11th International Conference on Analysis and Optimization of Systems: Discrete Event Systems (1994), 443-448.
- [5] F.W. Ciarallo, R. Akella and T.E. Morton, A periodic review production planning model with uncertain capacity and uncertain demand optimality of extended myopic policies, *Management Sci* 40(3) (1994), 320-332.
- [6] J. Collins, C. Bilot, M.L. Gini and B. Mobasher, Decision processes in agent-based automated contracting, *IEEE Internet Computing* 5 (2001), 61-72.
- [7] R. Davis and R.G. Smith, Negotiation as a metaphor for distributed problem solving, *Artificial Intelligence* 20 (1983), 63-109.
- [8] J. Du and J.Y.T. Leung, Minimizing total tardiness on one processor is NP-hard, *Mathematics of Oper Res* 15 (1990), 483-495.

- [9] P. Faratin and M. Klein, Automated contract negotiation as a system of constraints, Proceedings of the Workshop on Distributed Constraint Reasoning, IJCAI-01, Seattle, WA, August 2001, 33-45.
- [10] J. Gallien and L.M. Wein, A simple effective component procurement policy for stochastic assembly systems, Working Paper, Sloan School of Management, Massachusetts Institute of Technology (1998).
- [11] H. Gurnani, R. Akella and J. Lehoczky, Optimal order policies in assembly systems with random demand and random supplier delivery, IIE Trans 28 (1996), 865-878.
- [12] H. Gurnani, R. Akella and J. Lehoczky, Supply management in assembly systems with random yield and random demand, IIE Trans 32 (2000), 701-714.
- [13] W. Hopp and M. Spearman, Setting safety lead times for purchased components in assembly systems, IIE Trans 25 (1993), 2-11.
- [14] K. Jain and E.A. Silver, The single period procurement problem where dedicated supplier capacity can be reserved, Naval Res Logist, 42 (1995), 915-934.
- [15] U. Karmarkar and S. Lin, Production planning with uncertain yield and demand, Working Paper, William E. Simon Graduate School of Business Administration, University of Rochester (1986).
- [16] A. Kumar, Component inventory costs in an assembly problem with uncertain supplier lead-times, IIE Trans 21 (1989), 112-121.
- [17] P.S. Ow and T.E. Morton, The single machine early/tardy problem, Management Sci, 35 (1989), 177-191.
- [18] M. Pinedo, Scheduling: Theory, Algorithms, and Systems, Prentice Hall, Upper Saddle River NJ, 1995.
- [19] S. Rajagopalan and J.M. Swaminathan, A coordinated production planning model with capacity expansion and inventory management, Management Sci 47(11) (2001), 1562-1580.
- [20] T.W. Sandholm, An implementation of the contract net protocol based on marginal cost calculations, The Eleventh National Conference on Artificial Intelligence (AAAI-93), Washington DC, July 1993, 256-262.
- [21] N.M. Sadeh, R. Arunachalam, R. Aurell, J. Eriksson, N. Finne and S. Janson, TAC'03: a supply chain trading competition, AI Magazine 24 (2003), 92-94.

- [22] H. Shore, Setting safety lead-times for purchased components in assembly systems: A general solution procedure, *IIE Trans* 27 (1995), 634-637.
- [23] J. Song, C.A. Yano and P. Lerssrisuriya, Contract assembly: dealing with combined supply lead time and demand quantity uncertainty, *Manufacturing and Service Oper Management* 2(3) (2000), 287-296.
- [24] J.A. Van Mieghem, Coordinating Investment, Production, and Subcontracting, *Management Sci* 45(7) (1999), 954-971.
- [25] A. Vepsalainen and T.E. Morton, Priority rules and lead time estimation for job shop scheduling with weighted tardiness costs. *Management Sci* 33 (1987), 1036-1047.
- [26] C. Yano, Stochastic lead times in two-level assembly systems. *IIE Trans* 19 (1987), 371-378.
- [27] P. Zipkin, Models for design and control of stochastic, multi-item batch production systems, *Oper Res* 34(1) (1986), 91-104.